

Path Planning on a Compressed Terrain

Daniel M. Tracy
Rensselaer Polytechnic
Institute
tracyd@cs.rpi.edu

W. Randolph Franklin
Rensselaer Polytechnic
Institute
frankwr@rpi.edu

Barbara Cutler
Rensselaer Polytechnic
Institute
cutler@cs.rpi.edu

Franklin T. Luk
Rensselaer Polytechnic
Institute
luk@cs.rpi.edu

Marcus Andrade
Federal University of Viçosa,
Brazil
marcus.ufv@gmail.com

ABSTRACT

We present a better algorithm for path planning on complex terrain in the presence of observers and define several metrics related to path planning to evaluate the quality of various terrain compression strategies.

The path-planning algorithm simulates a smuggler and border guards scenario. First, we place observers on a terrain so as to optimize their visible coverage area. Next, we compute a path that a smuggler would take to minimize detection by an observer, path length, and uphill movement. The smuggler is allowed the full range of Euclidean motion on the 2-dimensional plane, unlike alternate path planning schemes that strictly avoid obstacles. We use two runs of the A* algorithm to efficiently compute this path.

We introduce new application-specific error metrics for evaluating lossy terrain compression. The target terrain applications are the optimal placement of observers on a landscape and the navigation through the terrain by smugglers. The error metrics compare the observer visibility and the cost of the optimal smuggler's route on the reconstructed terrain to the original terrain.

Categories and Subject Descriptors

I.3.5 [Computing Methodologies]: Computer Graphics—*Computational Geometry and Object Modeling*; E.4 [Data]: Coding and Information Theory—*Data compaction and compression*

Keywords

Path planning, observer viewshed, terrain compression.

1. INTRODUCTION

Recent advances in LIDAR and related scanning technologies have allowed for the acquisition of increasingly high resolution terrain data, in the form of rasters of elevation data.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM GIS '08, November 5-7, 2008, Irvine, CA, USA
Copyright 2008 ACM 978-1-60558-323-5/08/11 ...\$5.00.

A sample urban dataset we are using for testing covers approximately 25 city blocks, containing 12 million 3D points and requires about 300 MB for storage. Publicly-available terrain datasets covering the continental United States, sampled horizontal approximately every 30 meters require 40 GB of storage. Higher resolution data sampled every 10 meters or less is available and becoming commonplace.

These types of data sets have many commercial and military applications. For example, a soldier may wish to download the elevation data for a war zone onto a PDA, in order to plan a safe and efficient route through the terrain. Due to limited bandwidth and storage, it is typically not practical to download and process the original data. Thus, we need to transmit and store this information in a compressed format.

Depending on the application, a certain level of information loss can be acceptable. In the field of image processing, the usual way to evaluate a lossy compression scheme is to calculate the per-pixel average and the maximum errors between the original and reconstructed geometries. The standard image quality metrics may not be appropriate for the common tasks performed on terrain data, and it may be beneficial to consider more domain-specific applications. For example, correct evaluation of inter-elevation point terrain visibility is essential for surveying, cell phone tower placement, military surveillance, etc.

We consider the problem of multiple observer siting on a compressed and reconstructed terrain. An observer is placed at a point on the terrain and can see other points on the terrain only if no other part of the landscape obstructs a direct line of sight between the points, and the point is within the radius of visibility. One application of observer siting is to maximize the amount of visible terrain by specifying the optimal placement for a fixed number of watchtowers. In many real world scenarios, this problem must be solved using only the reconstructed terrain because the original terrain is prohibitive for transmission and storage.

Another interesting terrain data application is that of path planning. Given a complex mountainous or urban dataset, determine the shortest path through the terrain minimizing various relevant quantities such as total path length, distance traveled through terrain visible by one or more observers, and distance traveled uphill. We present a method for efficiently computing the best path through a terrain and show how this target application can be used as a metric for evaluating terrain compression techniques.

The contributions of our paper are:

- A new application-specific method for evaluating the

quality of the reconstructed geometry provided by a particular compression scheme.

- A novel scheme for computing a minimum length path between two terrain points while avoiding detection.
- New compression error metrics based on path length and visibility errors

2. BACKGROUND

2.1 Path Planning Techniques

Path planning around obstacles in E^2 and E^3 is a classic research topic with diverse important applications.

The well-known Lee maze-running algorithm [9] and High-tower line-search algorithm [5] consider movement in only four directions on a two-dimensional terrain.

In robotics, the goal is often to find the safest path, keeping as far away from obstacles as possible. To accomplish this, a Voronoi diagram of the obstacle boundaries may be constructed. Performing a graph search on the Voronoi diagram can then yield the optimal path [1]. Another approach is to construct the visibility graph for a two-dimensional terrain from the vertices of polygonal obstacles. Performing a graph search on the visibility graph will yield the shortest path. It is not necessary to restrict the obstacles or the path to a grid [10]. This will work well if simply the Euclidean distance is to be minimized. However, for more complex cost functions, a different approach is needed.

Three-dimensional path-planning problems are much more difficult. The general case of computing the shortest Euclidean path between two points, while avoiding polyhedral obstacles, is NP-hard. A common problem in robotics is to compute a collision-free path of a 3-dimensional object, such as a 6-DOF human arm, moving through a 3-dimensional workspace [8]. Kavraki [7] computes collision-free paths for robots moving among stationary obstacles. This method is particularly interesting for robots with many degrees of freedom (five or more).

The A* algorithm is a general-purpose graph search algorithm which can be adapted to solve any path-planning problem whose search space can be expressed as a tree [4]. This serves as the base for our path planning algorithm, described later in the paper.

2.2 Multiple Observer Siting

The goal in the placement of observers in the “smugglers and border guards scenario” is to maximize their *joint viewshed*, which is the area of the terrain that is visible by at least one observer. Due to the inherent complexity of computing the visibility between every pair of points on the terrain, it is impractical to compute the exact optimal solution to the multiple observer siting problem, especially when considering applications on small portable devices used by soldiers out in the field. Instead the randomized multiple observer siting algorithm by Franklin and Vogt [3] is used to approximate the optimal siting of observers.

3. OUR CONTRIBUTIONS

In this section we describe our efficient algorithm for planning smuggler’s paths through complex terrain in the presence of observers. We then describe a set of metrics for evaluating how well the terrain details are preserved with different lossy compression schemes.

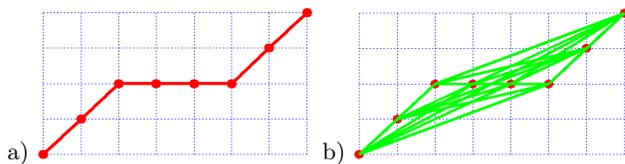


Figure 1: A sample result of the first pass of our path planning algorithm. For the second pass, we consider movement along the green lines.

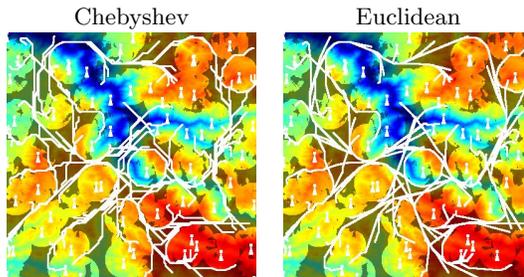


Figure 2: A comparison of the Chebyshev and Euclidean metrics. 50 pairs of random start and end points are used for more robust testing. The white towers are the observers.

3.1 Path Planning on Complex Terrain

After adapting multiple observer siting for border guard placement, we developed a path planning algorithm to determine the optimal smuggler’s route. We assume that the smuggler has complete knowledge of the observers’ positions, and that the smuggler would like to minimize the risk of detection while taking the quickest route possible.

The path finding routine is an adaptation of the A* algorithm. The first cost metric we evaluated was simply the number of grid points visited, ignoring the elevations and forbidding the visible regions (similar to [9]). For the A* algorithm, the terrain is represented as an $n \times n$ 2D matrix of height values, each point on the terrain is a separate node, and each node has up to eight children in the search tree, corresponding to eight neighboring points. This method has the limitation of minimizing only the Chebyshev distance between the end points [12].

3.2 Chebyshev vs. Euclidean Distance

A naive method to allow for a full range of Euclidean motion would be to include edges between all grid point pairs in the search space. However, this increases the size of the search space from $O(n^2)$ to $O(n^4)$, as there are $O(n^4)$ possible edges to consider. Also, to compute the true cost of each edge requires $O(n)$ time, because the edge must be segmented as described below. This is clearly too expensive for larger terrains.

To speed up the algorithm, we designed a two-pass system. On the first pass, all points on the terrain are included in the search space, and the Chebyshev path is computed as described in the previous section (Figure 1a). On the second pass, the only nodes in the search space are the points that are retained in the first path, and an edge is added to every other node in the search tree (see Figure 1b). Thus, for any pair of points in the search space, the smuggler may traverse a straight line connecting them. In practice, computing this

second pass is more efficient than the first pass. Examples of both passes are illustrated in Figure 2.

3.3 Alternate Cost Functions

We also investigated additional extensions to the cost metric. First we added a penalty for traveling uphill. Next, we lifted the restriction on traveling through regions that were visible from one or more of the observer positions to allow this movement but carry a stiff penalty. The cost of moving from one point to an adjacent point is given by $Cost = \sqrt{h^2 + v^2} * S * V$, where h is the horizontal distance between the points, and v is the elevation difference. S is the slope penalty, which is $(1 + v/h)$ when going uphill and 1 otherwise. V is the visibility penalty, which is 100 if the new cell is visible and 1 otherwise. A penalty of 100 was found to be appropriate to discourage movement through a viewshed when it is avoidable.

The cost metric parameters may be adapted to different types of terrain scenarios. For example, for urban LIDAR data, a visibility penalty of 10 and a slope penalty of $1 + v^2/h$ works better in discouraging the smuggler to climb up the side of a building.

Calculating the cost to move between adjacent points is trivial. However, for the two-pass system that which allows a full range of Euclidean motion, the cost to traverse a straight line that connects two distant points must be computed. This line is not likely to pass through grid points exactly. Here the elevation and visibility at several points along the line must be interpolated.

Our path planning procedure takes a cost function defined on a uniform grid, with non-uniform edge weights, and computes the path that minimizes the cost function while allowing a full range of Euclidean motion.

3.4 Novel Evaluation Metrics

We have innovated a new application-specific protocol for evaluating the quality of terrain compression. First, the original terrain representation is compressed then reconstructed to obtain the alternate representation. This is the representation available to field agents who will be making tactical decisions. Second, the multiple observer siting is performed on the alternate representation to produce a set of observers, and their joint viewshed is computed. Third, this same set of observer positions is transferred back to the original representation, where their true joint viewshed is computed. This simulates the errors and non-optimality that would result if observer placement was computed in the field using a compressed representation. In the fourth step, our path planning algorithm is applied to find the optimal paths on both the original and the reconstructed representations.

Three new error metrics were derived for the smugglers and border guards scenario. They target typical applications and tasks that use terrain data. If artifacts from a proposed compression scheme lead to significant errors in any of these metrics, the compression scheme should not be recommended for critical terrain applications, even if the reconstructed terrain is without *visual* artifacts.

The new error metrics are as follows:

1. *Viewshed Error*, or area of the symmetric difference of the cumulative viewsheds, implying suboptimal observer siting.
2. *Path Cost Error*, or difference in the path costs computed on the original and alternate, implying that the smuggler’s path is suboptimal.

3. *Path Visibility Error*, or percent of the alternate path that is visible when it is applied back to the original terrain. The smuggler plans the path using the alternate representation, believing that path to be safe, but which might inadvertently traverse a guard’s field of vision.

4. RESULTS

The new protocol and error metrics were tested on two compression schemes: JPEG 2000 [6] [11] and ODETLAP (Overdetermined Laplacian Solver) [2]. 400×400 terrains sampled from DTED-2 data were used because both JPEG and ODETLAP can process them efficiently, though the path planning itself has been observed to be efficient on 3200×3200 terrains. Three hilly (hill1, hill2, hill3) and three mountainous datasets (mtn1, mtn2, mtn3) were chosen to standardize all testing on. Sample paths on each are shown in Figure 3. Some initial experiments have also been performed on some urban LIDAR data from Ottawa, as shown in Figure 4.

The scheme is robust, working equally well on JPEG 2000 and ODETLAP. The numbers agree with what is seen visually. The greater the compression, the more the terrain features are blurred, and the larger the computed errors. The visibility is usually greater on the alternate representation than on the original. This is because the compression removes detail and smoothes out the terrain, eliminating visibility obstructions. The increased visibility will sometimes block off important passages, forcing the smuggler to take a long detour. The path visibility errors tend to be very small because a portion of the terrain that is biased towards the nonvisible areas is being sampled. This is a good indication that we are computing correct paths.

Also, our heuristic path planning algorithm was compared against the brute-force method. The average difference in the lengths of the computed paths was less than 0.1%, while the average speedup was greater than 100.

5. CONCLUSIONS

Observers were optimally sited on a terrain and good smuggler’s paths were computed through the terrain. An algorithm was developed that minimizes path cost, as measured by path length plus penalties for uphill movement and observer detection, while also permitting the range of Euclidean motion, rather than simply minimizing Chebyshev distance. New methods were also developed for evaluating terrain compression. These application-specific error metrics test how well the reconstructed terrain performs on the smugglers and border guards problem, with regards to the observers’ visibility and the smuggler’s path cost. When a user wants to plan a smugglers and border guards scenario on a reduced representation of a terrain, these new error metrics will help select the appropriate representation.

6. FUTURE WORK

Alternative observer placement policies may also be considered. For example, rather than seeking to maximize the total coverage area, the observers could form a perimeter and seek to minimize the “gaps” in that perimeter.

One useful extension of the smugglers and border guards scenario is to consider mobile observers. Each guard patrols a specified path. The smuggler may have to pause at certain

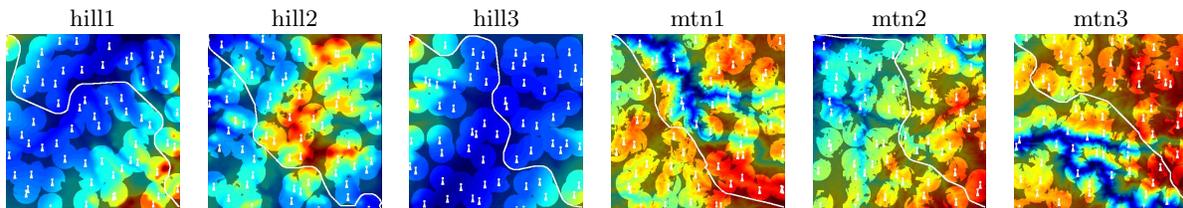


Figure 3: A path computed on a sited terrain. The white towers represent observers. The red areas correspond to high elevations, and the blue areas correspond to low elevations. The bright areas are visible by at least one observer, while the dark areas are hidden.

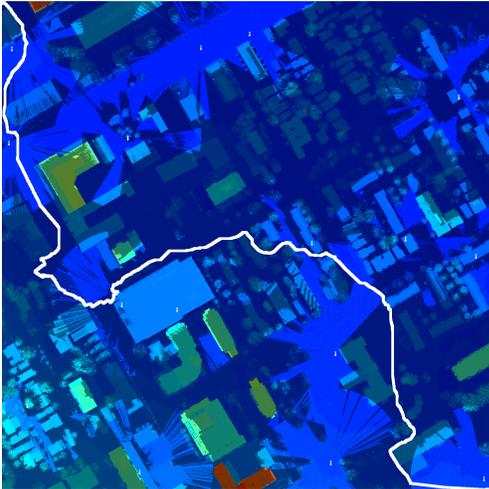


Figure 4: Path planning on a 2000×2000 terrain derived from Ottawa LIDAR data.

intervals to wait for the guards to reposition themselves. An element of unpredictability may also be added to the guards' movements. The smuggler will not know the observers' future positions, though it would retain the ability to track the observers' current positions.

7. ACKNOWLEDGEMENTS

This research was supported by NSF grants CCR-0306502 and DMS-0327634, by DARPA/DSO/GeoStar, and by CNPq - the Brazilian Council of Technological and Scientific Development.

8. ADDITIONAL AUTHORS

Additional authors: Metin Inanc (Rensselaer Polytechnic Institute, email: inancm@gmail.com), Zhongyi Xie (Rensselaer Polytechnic Institute, email: xie@cs.rpi.edu), and Jared Stookey (Rensselaer Polytechnic Institute, email: jstookey@jstookey.com).

9. REFERENCES

- [1] H. Choset. Incremental construction of the generalized Voronoi diagram, the generalized Voronoi graph, and the hierarchical generalized Voronoi graph. In *Proceedings of the First CGC Workshop on Computational Geometry*, October 1997.
- [2] W. R. Franklin, M. Inanc, Z. Xie, D. M. Tracy, B. Cutler, M. V. A. Andrade, and F. Luk. Smugglers and border guards - The GeoStar Project at RPI. In *15th ACM International Symposium on Advances in Geographic Information Systems (ACM GIS 2007)*, Seattle, WA, USA, Nov 2007.
- [3] W. R. Franklin and C. Vogt. Tradeoffs when multiple observer siting on large terrain cells. In A. Riedl, W. Kainz, and G. Elmes, editors, *Progress in Spatial Data Handling: 12th International Symposium on Spatial Data Handling*, pages 845–861. Springer, Vienna, 2006. ISBN 978-3-540-35588-5.
- [4] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths in graphs. *IEEE Trans. Syst. Sci. and Cybernetics*, SSC-4(2):100–107, 1968.
- [5] D. W. Hightower. A solution to line-routing problems on the continuous plane. In *DAC '69: Proceedings of the 6th Annual Conference on Design Automation*, pages 1–24, New York, NY, USA, 1969. ACM.
- [6] JPEG Committee. JPEG 2000. <http://www.jpeg.org/jpeg2000/> (accessed 14 Dec 2007), 2004.
- [7] L. Kavraki, P. Svestka, J.-C. Latombe, and M. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, August 1996.
- [8] J. Kuffner and S. LaValle. RRT-connect: An efficient approach to single-query path planning. In *Proc. IEEE Int'l Conf. on Robotics and Automation (ICRA'2000)*, San Francisco, CA, Apr. 2000.
- [9] C. Lee. An algorithm for path connections and its applications. *IRE Transactions on Electronic Computers*, EC-10, pages 346–365, Sept. 1961.
- [10] T. Lozano-Pérez and M. A. Wesley. An algorithm for planning collision-free paths among polyhedral obstacles. *Commun. ACM*, 22(10):560–570, 1979.
- [11] M. J. Owen and M. W. Grigg. The compression of digital terrain elevation data (DTED) using JPEG 2000. Technical Report DSTO-TR-1548, Defense Science and Technology Organization Salisbury, Australia, February 2004.
- [12] D. M. Tracy, W. R. Franklin, B. Cutler, M. A. Andrade, F. T. Luk, M. Inanc, and Z. Xie. Multiple observer siting and path planning on lossily compressed terrain. In *Proceedings of SPIE Vol. 6697 Advanced Signal Processing Algorithms, Architectures, and Implementations XVII*, San Diego CA, 27 August 2007. International Society for Optical Engineering. paper 6697-16.