CALCULATION OF ON- AND OFF-ROAD SHORTEST TIME PATHS USING PARALLEL ALGORITHMS

Timothy M. Telcik
Department of Geographic Information Systems
School of Computing
Curtin University of Technology
GPO Box U1987
Perth, Western Australia 6001
Email: tim@marsh.cs.curtin.edu.au

ABSTRACT Most previous work on the determination of shortest time paths has dealt with travel over pre-defined networks or with off-road mobility. Many real-life problems involve combinations to reach the desired destination. This problem has been termed the Cross-Country Problem (CCP); commercial systems that allow such computations are lacking. A solution is discussed which uses a Digital Elevation Model (DEM) called a Triangulated Irregular Network (TIN). Also considered is a comparison of the TIN and regular grid DEM for the CCP, and the role of a massively parallel computer for path-finding.

INTRODUCTION

There has been an increased interest in efficient shortest path algorithms because of their application to robot and vehicle navigation, a desire to understand how human beings solve navigation problems in complex environments, and the increasing availability of powerful computational techniques and devices for solving such problems [Smith and Parker 1987]. While many Geographic Information Systems (GIS) software packages provide typical linear network routing algorithms, they lack the ability to find a minimal cost path over a general terrain.

This paper examines the problem of finding an optimum path for a land-based vehicle between a source point (***s***) and destination point (***d***) over intervening terrain using a combination of on-road/off-road access. The problem has been termed the Cross-Country Problem (CCP). One solution to the CCP is presented using a Digital Elevation Model (DEM) called a Triangulated Irregular Network (TIN). Also discussed is how the TIN and regular grid DEM fare for the CCP, and the applicability of a massively parallel computer to path-finding research at Curtin University.

THE TRIANGULATED IRREGULAR NETWORK APPROACH

A Triangulated Irregular Network (TIN) is a surface modelling tool which can be used to represent terrain. It is constructed from a set of irregularly spaced (x,y,z) data points selected from a surface which are triangulated to produce a continuous mesh of non-overlapping triangles (see Figure 1).

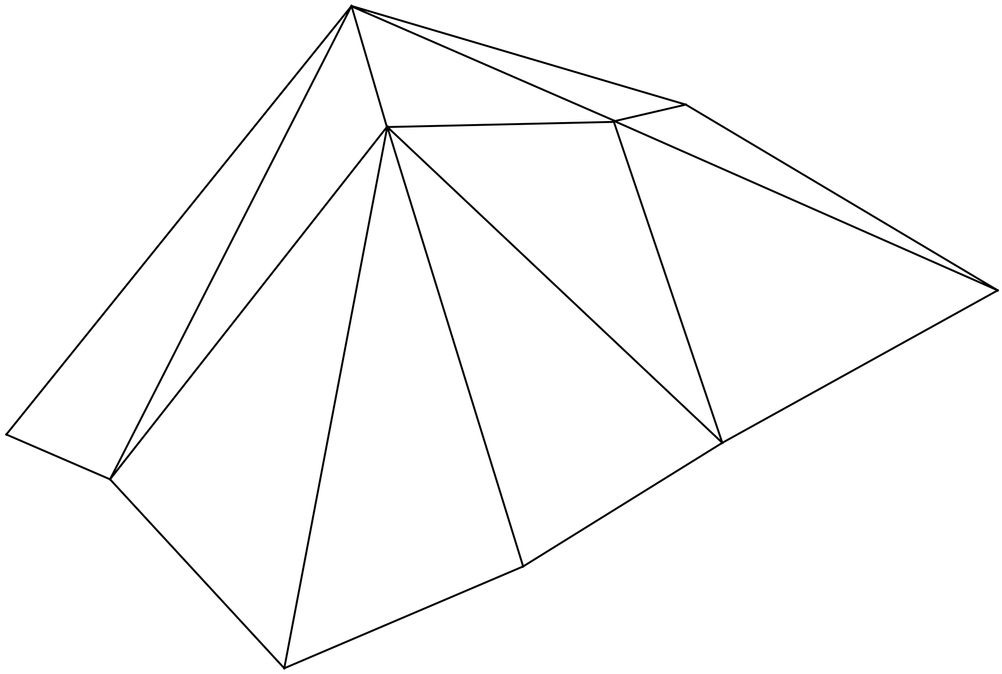

Figure 1. Perspective view of a TIN containing 10 points.

The TIN concept and some of its applications were put forward by Peucker *et al.* [1976, 1978]. Poiker [1990] (formerly Peucker) notes two methods for topologically storing TINs:

(i) **Triangle by triangle** (see Figure 2 and Table 1).

(ii) **Points and their neighbours** (see Figure 2 and Table 2).

A typical TIN (as used here) is a 2-and-a-half-dimensional (2½D) structure (*i.e.* only one z value for any x, y). By assigning each triangle and edge in the TIN with a weight called a *traversability index*, the TIN can be used as the base model for solving the Cross-Country Problem. The traversability index for a triangle is derived from known information about terrain, such as slope, soil type and vegetation type. TIN edges are assigned a traversability index that is less than, or equal to, that of their adjacent triangles.

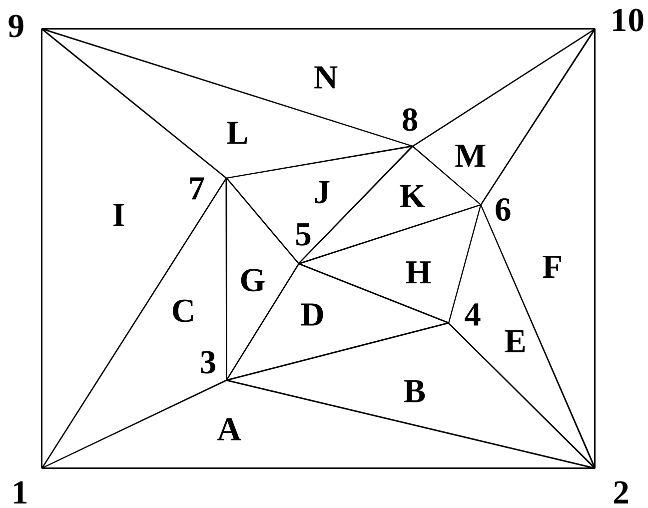

Figure 2. A TIN with labelled points and triangles (from Telcik 1992).

ID	V1(x,y,z)	V2(x,y,z)	V3(x,y,z)	TRI NBRS
A	3(5, 3, 3)	2(15, 0, 2)	1(0, 0, 0)	B -1 C
B	4(11, 5, 4)	2(15, 0, 2)	3(5, 3, 3)	E A D
C	7(5, 10, 9)	3(5, 3, 3)	1(0, 0, 0)	A -1 I
		. . .		
N	10(15,15, 2)	8(10,11, 7)	9(0, 5, 6)	-1 M L

Table 1. Triangle topology for Figure 2 (from Telcik 1992).

ID	(x,y,z)	POINT NEIGHBOURS
1	0, 0, 0	9 7 3 2 -1
2	15, 0, 2	10 -1 1 3 4 6
3	5, 3, 3	7 5 4 2 1
		. . .
10	15, 15, 2	2 6 8 9 -1

Table 2. Point topology for Figure 2 (from Telcik 1992).

Note: Neighbours are ordered clockwise starting from north. A value of '-1' represents a neighbour outside the TIN. Thus, a point or triangle with a neighbour value of '-1' is on the TIN's boundary.

THE TIN CROSS-COUNTRY PROBLEM

The Weighted Region Problem (WRP) was introduced by Mitchell [1986] and Mitchell and Papadimitriou [1986, 1987, 1991]. It forms a basis for the TIN Cross-Country Problem (TIN CCP) which is defined as:

> *Find the minimal cost path located on the surface of a weighted TIN between a source point (**s**) and destination point (**d**)* (see Figure 3).

A weighted path is the sum of all sub-path costs between ***s*** and ***d*** across the TIN (*i.e.* a path does not pass above or below the surface). Each sub-path cost is calculated from its Euclidean (or straight line) distance across the interior of a surface triangle (or along an edge) multiplied by the triangle's (or edge's) weighting.

Two generalisations are made about the TIN CCP:

> **(i)** The vehicle traversing the terrain is of negligible size in comparison to the scale of the terrain model.
>
> **(ii)** Direction of travel does not affect path cost (*eg.* travelling up, down or across a hill does not matter).

The TIN CCP algorithm incorporates planar unfolding, Snell's Law of Refraction, and heuristics and pruning techniques. These concepts and their relevance to the TIN CCP will be described in the following sections.

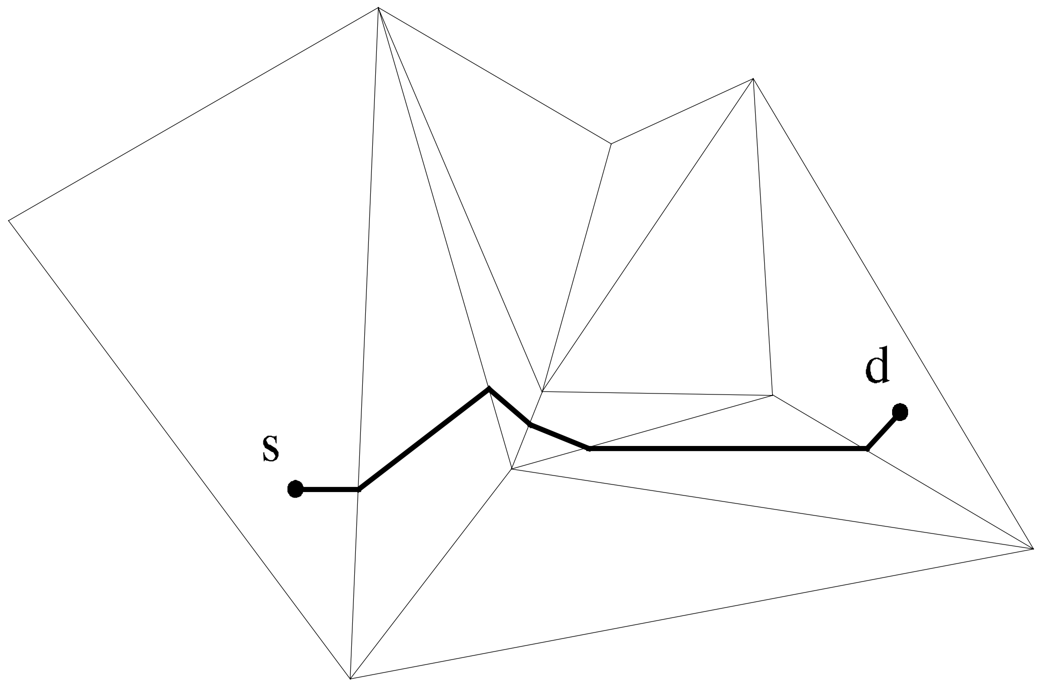

Figure 3. Minimal cost path across weighted TIN in Fig.2.

Planar Unfolding

For the CCP, travel occurs on the surface of the TIN. Path computation is simpler when the triangles crossed by a path are unfolded onto a common *x*-*y* plane using a concept called *planar unfolding* [Bajaj and Moh 1988; Barrera and Vázquez-Gómez 1989; Lyusternik 1964; Mount 1985a, 1985b; Mitchell *et al.* 1987; Sharir and Schorr 1986]. Planar unfolding is an *affine* transformation which retains parallel lines. Simply projecting triangles onto the *x*-*y* plane is inappropriate.

The unique shortest path between two points in a plane is a straight line. This principle can be extended to points on the surface of convex and non-convex polyhedra, such as a TIN. For a convex polyhedron, the shortest path after planar unfolding is a straight line (see Figure 4). For a non-convex polyhedron, the shortest path after planar unfolding is a piecewise straight line which could pass through the vertex of neighbouring facets (see Figure 5).

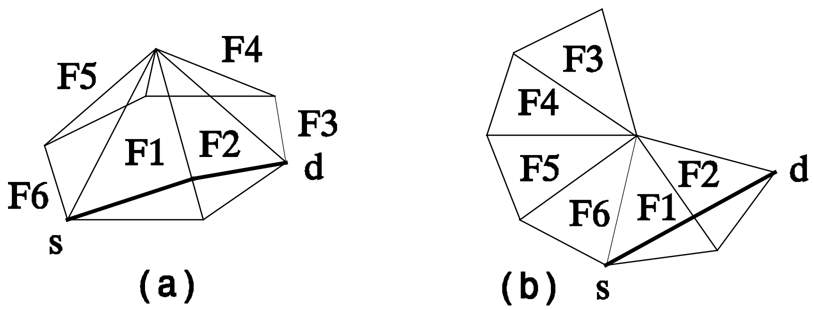

Figure 4. (a) The shortest path over a convex polyhedra before unfolding, and (b) after unfolding (from Barrera and Vázquez-Gómez 1989).

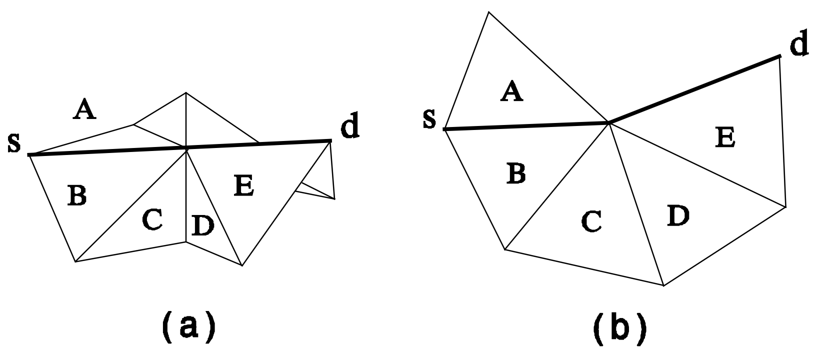

Figure 5. (a) The shortest path over a non-convex polyhedra before unfolding, and (b) after unfolding (from Barrera and Vázquez-Gómez 1989).

To unfold two triangular facets, f_1 and f_2, which are *edge-adjacent* at their common edge e_1 (see Figure 6), unfold f_1 onto the same plane as f_2 such that points in f_1 fall on the opposite side of e_1 to points in f_2. The transformation can be achieved using matrices or vector algebra.

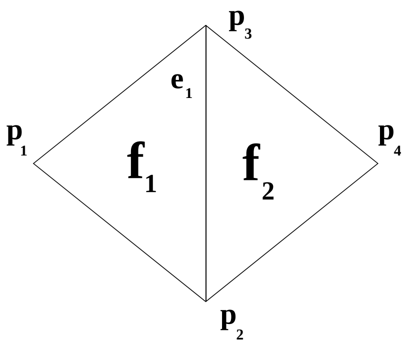

Figure 6. Two adjacent triangles.

To unfold a series of triangles, f_1, f_2, ..., f_n, with the *edge-sequence* of common edges e_1, e_2, ..., e_n, unfold f_1 about e_1 until it lies in the same plane as f_2, then unfold f_1 and f_2 about e_2 onto the same plane as f_3 etc. until all triangular facets f_1, f_2, ..., f_{n-1} lie in the same plane as f_n [Mitchell *et al.* 1987].

Planar unfolding enables weighted triangles within the TIN to be imagined as optical prisms through which light rays (or paths) pass. When a path extends from one triangle to its neighbour, the current triangle is unfolded onto the plane of its neighbouring triangle and Snell's Law of Refraction applied in the common coordinate system.

Snell's Law of Refraction

A path's cost is proportional to its Euclidean (or straight line) distance through a weighted region or along a weighted edge (*i.e.* distance **x** weight). Path planning across the boundaries between homogeneous weighted regions is analogous to Fermat's principle of optics which states that light rays always follow a path between two locations that is a minimum, maximum, or stationary point with respect to time (although it is usually a minimum) [Rowe and Richbourg 1990]. Thus, optimal paths across weighted regions will be piecewise-linear, turning only at region boundaries [Lyusternik 1964; Mitchell 1988; Rowe and Alexander 1992; Mitchell and Papadimitriou 1986, 1987, 1991; Richbourg *et al.* 1986, 1987; Rowe and Richbourg 1987]. For example, a stick appears shorter in water because light rays *bend* when passing from water to air (which is less dense).

Snell's Law of Refraction is a corollary to Fermat's principle which Mitchell [1988] relates as:

> *The path of a light ray passing through a boundary e between regions f and f ' with respective indices of refraction* α *and* β *obeys the relationship that* $\alpha \sin \theta = \beta \sin \phi$, *where* θ *and* ϕ *are the respective angles of incidence and refraction* (see Figure 7).

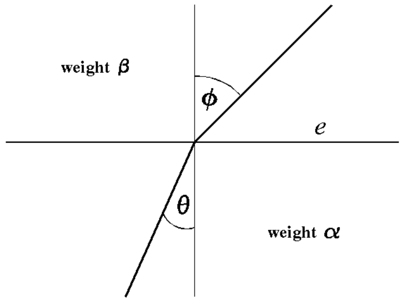

Figure 7. Light ray crossing the boundary *e* between weighted regions α and β (adapted from Mitchell 1988).

A road network can be represented by assigning TIN edges with weights less than their adjacent triangles. Paths obeying Snell's Law will follow a road rather than travel through the interior of a region because the cost per unit distance is less.

Mitchell [1988] notes one important modification to Snell's Law with respect to travel: a path, unlike a light ray, cannot be *critically reflected* (*i.e.* a path cannot rebound from an edge if its incident angle is too high). Instead, the path will *travel* along the edge and exit at a later point.

Three important ideas arise from Snell's Law and its modification:

(i) Paths only bend, or change direction, at neighbouring regions with different weights. Thus, the path between two identical homogeneous weighted regions is a straight line.

(ii) If the weight of a source region is less than its neighbour and common edge then a path will travel parallel to the common edge just inside the source region because it is *cheaper*.

(iii) Paths will travel along edges which have weights less than their surrounding faces. This allows for optimisation of on-road/off-road access across a TIN.

The Snell's Law analogy is useful but does not carry too far. Finding the exact path of refraction from ***s*** to ***d*** using Snell's Law of Refraction cannot be solved in closed form because a ray of light striking a vertex will *diffract* akin to a ray of light passed through a thin slit, because of *critical reflections* (mentioned above), and because of the equations which must be solved [Joseph Mitchell, pers. comm. 1992; Mitchell and Papadimitriou 1991]. Instead, a simple bisection iteration (or coordinate descent) method is used to calculate where paths optimally intersect triangle edges [Mitchell 1986, 1988; Mitchell and Papadimitriou 1986, 1987, 1991; Rowe and Alexander 1992; Rowe and Richbourg 1990]. This method does not provide an exact solution, but approximates the optimal weighted path.

Snell's Law was chosen as the optimisation technique for weighted path planning because a *brute-force* approach which iteratively adjusts path intersection points on edges to achieve a global minimum was computationally inefficient.

Heuristics and Pruning Techniques

A *heuristic* is an assumption that a particular condition *usually* exists. It can be used to reduce the amount of search computation. Heuristics offer an improvement in efficiency in most cases, but do not always provide the best solution. Any method which reduces the search space of a graph problem is called *pruning*.

The following heuristics and pruning techniques are used for the TIN CCP:

(i) **Maintaining a global upper bound (*U*).** This is the least cost path found so far. It is initially a weighted piecewise line path between ***s*** and ***d***. Whenever a weighted path reaches the destination, *U* is replaced if the new weighted path is less than it. Paths with a *cost* > *U* cannot be optimal and can therefore be removed from further consideration. The global upper bound is a basis for A^* search [Hart *et al.* 1968].

(ii) **Using an adjustable bounding ellipse.** An ellipse is formed using ***s*** and ***d*** as loci and using *U* to define the minor and major axis. It is adjusted when *U* is changed during the search. Paths which stray outside the bounding ellipse are discarded from the search as they are probably too costly. A fast approximation is to use a bounding box which surrounds the ellipse.

(iii) **Removing spiralling paths.** Any path which spirals back onto itself is discarded because it is degenerate (*i.e.* it will not reach the destination).

TIN CCP Algorithm

The TIN Cross-Country Problem Algorithm is based on work by Barrera and Vázquez-Gómez [1989], Gewali *et al.* [1988], Mitchell [1986], Mitchell *et al.* [1987], Mitchell and Papadimitriou [1986, 1987, 1991], Mount [1985a, 1985b], Rowe [1990], Rowe and Richbourg [1990], and Sharir and Schorr [1986].

Three simple data structures are used:

(i) *agenda* - contains partial paths generated during the search.
(ii) *edge lists* - each edge in the TIN has a list containing path intersection points of various partial paths under consideration.
(iii) *plist* - list containing completed paths.

A simplified version of the TIN CCP algorithm follows:

initialise *edge lists*, *agenda* and *plist* as empty
s ← source point
d ← destination point
src-tri ← triangle containing ***s***
dest-tri ← triangle containing ***d***
U ← weighted length of unfolded piecewise line between ***s*** and ***d***
box ← bounding box approximating ellipse based on *s, d,* and *U*
push *src-tri* path neighbours onto *agenda*

```
While the agenda is not empty Do Begin
    pop a partial_path from agenda
    If partial-path reached dest-tri Then Begin
        add partial-path to plist
        cost ← weighted length of partial-path between s and d
        If cost < U Then
            adjust U and box
    End
    Else Begin
        prune partial-path
        For each neighbouring-triangle
            If partial-path Not blocked Then Begin
                propagate partial-path to neighbouring triangles using planar
                  - unfolding and Snell's Law
                add partial-path to agenda
            EndIf
        EndFor
    EndIfElse
EndWhile
```

When the *agenda* is empty the search is complete because all triangles in the TIN have been examined. The algorithm outputs to *plist* the edges crossed by each non-discarded weighted path (i.e. *edge-sequence*). An optimal path between ***s*** and ***d*** may be found by examining the edge-sequences.

TIN VERSUS REGULAR GRID FOR THE CROSS-COUNTRY PROBLEM

TINs and regular grid (RG) DEMs form the basis for most terrain modelling and spatial data processing currently undertaken within Geographic Information Systems (GIS) and computer cartography. A solution to the CCP using these structures will be affected by their respective strengths and weaknesses.

A TIN can adapt to a change in the roughness of terrain, use minimal data points to represent a surface (*i.e.* it is *surface specific*), and can be topologically encoded so that adjacency analysis is easily computed [Aronoff 1989].

Specifying a suitable resolution for surface representation with an RG is difficult because a terrain's features often change over a region causing redundant grids where the terrain is uniform. However, storing a grid as a 2D array implies adjacency information (in the row and column), and allows information to be quickly altered by assigning a new value to the grid element.

Grid vertices are generally *interpolated* (or inferred) from a given data set which means that spatial accuracy is usually lost during its creation. If TIN vertices are *selected* from an RG then any further values interpolated from it are no more accurate than values derived from the grid.

For the Cross-Country Problem a regular grid (RG) has an inherent problem of rapidly accumulating arithmetic *round-off error* when calculating paths across cell boundaries [Rowe 1990; Rowe and Richbourg 1990], and is affected by *digitisation-bias* (*i.e.* a path can only exit a cell in four or eight directions) [Mitchell 1988; Rowe and Richbourg 1990].

THE ROLE OF PARALLEL PROCESSING IN PATH-FINDING

To solve the CCP a large number of paths must be explored before the best path can be found. Curtin University recently acquired a massively parallel machine, a DEC MasPar MP-1, configured with 2048 individual RISC processors called Processing Elements (PEs). The use of such a machine is one solution to reduce the amount of time required for path-planning.

The MP-1 is based on the SIMD (Single Instruction Multiple Data) processing paradigm in which PEs perform the same instructions on different data items at the same time, in what is known as *lockstep*. Algorithms designed for a SIMD machine decompose the data they operate on into elements, or groups of elements, that can be assigned to different PEs.

Finding the shortest path between two nodes in a graph is called the single source one-to-one shortest path problem [Deo and Pang 1984]. Most algorithms solving this problem construct a graph containing a set of vertices (nodes), and edges (arcs) linking those vertices. Each edge is assigned a weight which describes the cost of travelling along it. If a path of travel exists between source and destination vertices in the graph, the algorithm can determine a set of edges which minimise the cost of travel.

A parallel approach to the single-source problem on a SIMD machine requires that each PE must have a copy of the original graph [Harlan 1991]. To store a large complex graph in the local memory of each PE is not feasible because of physical memory limits. Similarly, a parallel approach to the TIN CCP using a SIMD machine requires that each PE must contain a copy of the data set and associated structures. Thus, for a large, or real-world, data set initial investigation shows that a SIMD machine is not well adapted to solving the TIN CCP.

CONCLUSION

The TIN Cross-Country Problem is an important extension to the TIN concept developed by Peucker *et al.* [1976, 1978] which can be applied to routing problems, such as vehicle navigation, and resource allocation, such as dispatching a fire crew to the scene of a fire.

Initial investigation into parallelizing the TIN CCP using a SIMD machine has shown the problem to be unfeasible because of the problem type and machine constraints. However, the SIMD approach does warrant further investigation, and a regular grid solution to the CCP using the MasPar will form a basis for further research on parallel path-finding at Curtin University.

While the TIN CCP algorithm presented may not be optimal it paves the way for future algorithms which could be incorporated into commercial GIS software supporting the TIN.

ACKNOWLEDGMENTS

I must firstly thank Steve Kessell for suggesting this problem, keeping the show *on the road* and proofing this paper. Thank you to Doug Hudson for discussing this problem. David Vitali's discussion on parallel techniques is greatly appreciated. Thanks to the following people for providing/discussing their work: Joseph Mitchell, Neil Rowe, Renato Barrera, David Mount and Charles Weisbin. Thanks (yet again) to Jill Kessell for proofing my work. Finally, thanks must go to my parents and Trish Evans for their on-going support.

BIBLIOGRAPHY

Aronoff, S. (1989) **Geographic Information Systems: A Management Perspective**, Ottawa, Canada: WDL Publications.

Bajaj, C. and T.T. Moh (1988), "Generalized Unfoldings for Shortest Paths", **International Journal of Robotics Resources** 7, 1, 71 - 76.

Barrera, R. and J. Vázquez-Gómez (1989), "A Shortest Path Method for Hierarchical Terrain Models", **Proceedings of Auto-Carto 9**, pp. 156 - 163, Baltimore, Maryland: American Society for Photogrammetry and Remote Sensing, and American Congress on Surveying and Mapping.

Deo, D. and C. Pang (1984), "Shortest Path Algorithms: Taxonomy and Annotation", **Networks** 14, 275 - 323.

Gewali, L., A. Meng, J.S.B. Mitchell and S. Ntafos (1988), "Path Planning in 0/1/∞ Weighted Regions with Applications (extended abstract)", **Proceedings of the Fourth Annual Symposium on Computational Geometry**, pp. 266 - 278, Urbana-Champaign, Illinois: ACM Press.

Harlan, R.M. (1991), "Searching in Parallel: A Case Study with the Single-Source Shortest Path Problem", **SIGCSE Bulletin** 23, 1, 254 - 259.

Hart, P.E., N.J. Nilsson and R. Raphael (1968), "A Formal Basis for the Heuristic Determination of Minimum Cost Paths", **IEEE Transactions on Systems Science and Cybernetics** SSC-4, 2, 100 - 107.

Lyusternik, L.A. (1964), **Shortest Paths: Variational Problems**, New York: Macmillan.

Mitchell, J.S.B. (1986), "Planning Shortest Paths", PhD Dissertation, Department of Operations Research, Stanford University, Stanford, California.

Mitchell, J.S.B. (1988), "An Algorithmic Approach to Some Problems in Terrain Navigation", **Artificial Intelligence** 37, 171 - 201.

Mitchell, J.S.B., D.M. Mount and C.H. Papadimitriou (1987), "The Discrete Geodesic Problem", **SIAM Journal of Computing** 16, 4, 647 - 668.

Mitchell, J.S.B. and C.H. Papadimitriou (1986), "The Weighted Region Problem", Technical Report, Department of Operations Research, Stanford University, Stanford, California.

Mitchell, J.S.B. and C.H. Papadimitriou (1987), "The Weighted Region Problem (extended abstract)", **Proceedings of the Third Annual Symposium on Computational Geometry**, pp. 30 - 38, Waterloo, Ontario, Canada: ACM Press.

Mitchell, J.S.B. and C.H. Papadimitriou (1991), "The Weighted Region Problem: Finding Shortest Paths through a Weighted Planar Subdivision", **Journal of the ACM** 38, 1, 18 - 73.

Mount, D.M. (1985a), "On Finding Shortest Paths on Convex Polyhedra", Computer Science Technical Report 1495 (CS-TR-1495), University of Maryland, College Park, Maryland.

Mount, D.M. (1985b), "Voronoi Diagrams on the Surface of a Polyhedron", Computer Science Technical Report 1496 (CS-TR-1496), University of Maryland, College Park, Maryland.

Peucker, T.K., R.J. Fowler, J.J. Little and D.M. Mark (1976), "Digital Representation of Three-Dimensional Surfaces by Triangulated Irregular Networks", Technical Report # 10, Office of Naval Research (ONR) Geography Programs 1-63, prepared for the ONR under contract # N00014-75-C-0886.

Peucker, T.K., R.J. Fowler, J.J. Little and D.M. Mark (1978), "The Triangulated Irregular Network", **Proceedings of the Digital Terrain Model (DTM) Symposium**, pp. 516 - 540, St. Louis, Missouri: American Society of Photogrammetry.

Poiker, T.K. (1990), "The TIN Model", **Technical Issues in GIS - NCGIA Core Curriculum** (M.F. Goodchild & K.K. Kemp eds.), pp. 39/3 - 39/15, National Centre for Geographic Information and Analysis, University of California: Santa Barbara, California.

Richbourg, R., N. Rowe, M. Zyda and R. McGhee (1986), "Exploiting Capability Constraints to Solve Global Two-Dimensional Path Planning Problems", Technical Report NPS-86-006, Department of Computer Science, U.S. Naval Postgraduate School, Monterey, California.

Richbourg, R., N. Rowe, M. Zyda and R. McGhee (1987), "Solving Global Two-Dimensional Routing Problems Using Snell's Law and A^* Search", **Proceedings IEEE Conference on Robotics and Automation**, pp. 1631 - 1636, Raleigh, North Carolina: IEEE Computer Society Press.

Rowe, N.C. (1990), "Roads, Rivers, and Obstacles: Optimal Two-Dimensional Path Planning Around Linear Features for a Mobile Agent", **The International Journal of Robotics Research** 9, 6, 67 - 74.

Rowe, N.C. and R.S. Alexander (1992), "Constructing Optimal-Path Maps for Path Planning Across Weighted Regions", Department of Computer Science, U.S. Naval Postgraduate School, Monterey, California.

Rowe, N.C. and R.F. Richbourg (1987), "A New Method for Optimal Path Planning through Non-Homogeneous Free Space", Technical Report NPS52-87-003, Department of Computer Science, U.S. Naval Postgraduate School, Monterey, California.

Rowe, N.C. and R.F. Richbourg (1990), "An Efficient Snell's Law Method for Optimal-Path Planning across Multiple Two-Dimensional, Irregular, Homogeneous-Cost Regions", **The International Journal of Robotics Research** 9, 6, 48 - 66.

Sharir, M. and A. Schorr (1986), "On Shortest Paths in Polyhedral Space", **SIAM Journal of Computing** 15, 1, 193 - 215.

Smith, T.R. and R.E. Parker (1987), "An Analysis of the Efficacy and Efficiency of Hierarchical Procedures for Computing Trajectories over Complex Surfaces", **European Journal of Operational Research** 30, 327 - 338.

Telcik, T.M. (1992), "Contouring a Triangulated Irregular Network", **AURISA News** 43. Also appears in: Telcik, T.M. (1991), "Contouring a Triangulated Irregular Network in Theory and Practice", Computing Project 394, School of Computing, Curtin University of Technology, Perth, Western Australia.