Terrain Characterization from Ground-Based LADAR[1]

by

Christoph Witzgall, Geraldine S. Cheok, David E. Gilsinn
National Institute of Standards and Technology

ABSTRACT

Test runs of Army Research Laboratory's (ARL) autonomous vehicle (Experimental Unmanned Vehicle, XUV) were followed by the acquisition of high resolution scans of selected regions of the test course. These scans were used to *i*) determine terrain features (e.g., heavy vegetation, ditches, etc.) which may hamper the autonomous navigation of the XUV and *ii*) develop the ability to quantify terrain features such as vegetation or roughness. Those tasks require determination of "ground" or "bare earth", which is a major issue of ongoing research into terrain characterization. Point clouds collected by ground-based LADAR (laser distance and ranging) pose a particular challenge because they are extremely dense in close proximity to the instrument and progressively sparse at larger distances. This work focuses on the National Institute of Standards and Technology's (NIST) procedures for ground determination and the development of gauges for vegetation coverage and slope variability.

Keywords: *autonomous vehicle, ground, LADAR, slope analysis, terrain characterization, triangulated irregular networks, vegetation coverage, vehicle mobility.*

.

1. INTRODUCTION

The use of LADAR scanners for the purpose of terrain representation and analysis has been increasing steadily over the past two decades as they can rapidly capture large amounts of 3D information – several million points per scan. A typical LADAR locates a point via range and angular information, providing coordinates centered at the instrument. The "angle-angle-range" data, or polar coordinates, are usually converted to Cartesian coordinates, or *x, y, z* data, and processed as such by most LADAR data processing software.

In terrain applications, ground cover, trees, bushes, rocks, tall grass, and various kinds of artifacts are parts of the scanned scene. A pervasive problem is to automatically extract "ground truth" or "ground" or "bare earth" from a point cloud.

What constitutes ground or ground truth? Movable artifacts are clearly not part of ground, neither are trees and bushes. Whether grass cover is considered part of the ground depends on the application. Short grass on a lawn may be acceptable whereas tall grass may not. Large rocks and some man made structures may be considered either as part of the ground surface or as artifacts superimposed on that surface. The definition of ground is thus flexible and application dependent. In this paper, ground is considered from the point of view of vehicle mobility. That is, trees and bushes need to be *i*) identified for vehicle path planning/crash avoidance and *ii*) removed to the extent that meaningful determination of terrain slopes is possible.

In this paper, procedures for ground determination (Section 3) and subsequent terrain characterization are described, addressing vegetation coverage (Section 4), slope and path analysis (Section 5). In Section 6, experimental applications are reported based on high resolution LADAR scans of test courses at Tooele Army Depot, UT (Fig. 1) and Ft. Indiantown Gap, PA.

The work described in this paper was funded in part by the ARL as part of its Technology Readiness Level 6 (TRL-6) program. The experiments were conducted to verify the readiness level of an XUV.

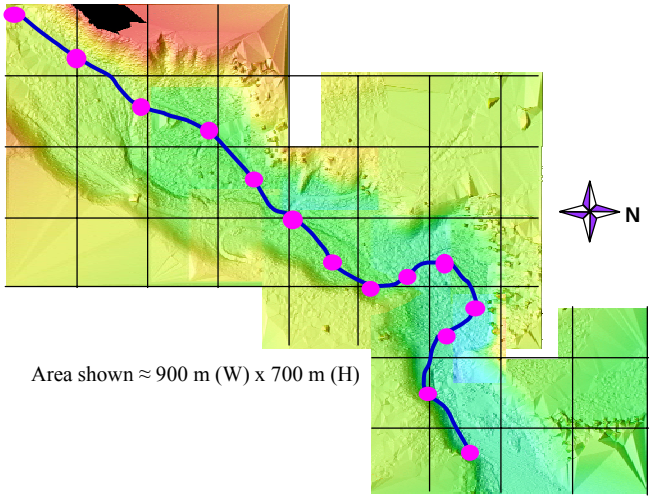

Figure 1. 3D surface of "wash" region , Tooele, UT with waypoints and trace of a XUV run. The difference in elevation between top and bottom of valley is ≈110 m.

[1] U. S. Government work and not subject to copyright in the U.S.

2. TIN PROCEDURES

Along with regular rectangular grids and quadtrees [1], Triangulated Irregular Networks (TINs) are used for the purpose of representing – interpolating or approximating – point clouds by surfaces. Since the TIN technique is still less known, a short description will be provided. For details on TIN procedures used at NIST see Witzgall et al. [2].

TINs provide a meshing of points (x_k, y_k, z_k) in the form of a piecewise triangular surface in 3D with these points as vertices, constructed above a 2D triangulation of the locations (x_k, y_k) of the vertices in a base plane (Fig. 2). The term "triangulation" is understood to indicate a partition of the base plane into non-overlapping triangles. In the context of TINs, triangulations are usually constructed using the "Delaunay" principle, which stipulates that cirmcumcircles of the triangles do not contain the locations of vertices in their interior.

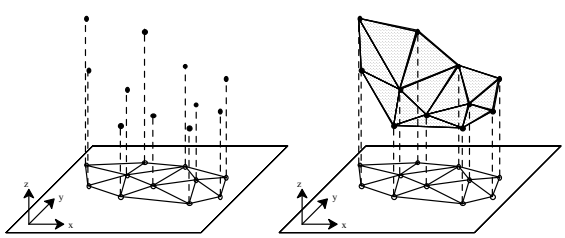

Figure 2. Triangulation of data locations and the corresponding TIN surface.

The NIST-developed TIN algorithms used in this work are based on the insertion method for Delaunay triangulation. Here, the meshing is constructed incrementally, one point at a time.

The insertion method may be terminated after a "partial triangulation" of the base plane for a specified number of vertices has been reached, rather than continuing the process until it terminates with a "full triangulation." A surface based on full triangulation interpolates the data set, while one based on partial triangulation approximates the data set. Partial triangulations depend critically on the sequence in which data points are inserted. The NIST routines offer the following options for selecting the next insertion point:

- Select the point of largest deviation from the current meshed surface
- Select the point for which the product of its deviation and the area of its 2D triangle is largest.

The first option offers aggressive insertion at areas of intensive elevation variation, creating small triangles in those areas. The second option achieves more homogeneous triangle sizes while still maintaining the adaptive advantage of the insertion method.

3. GROUND DETERMINATION PROCEDURES

Several procedures for ground determination are described in this paper. They fall into four categories, according to the following tasks:

- Selection/Rejection of ground points
- Meshing of the selected ground points to create a ground surface
- Smoothing to remove vertices that are outliers
- Mowing to remove points whose elevations above the ground surface are in excess of a specified tolerance

These procedures are selectively employed and typically repeated. One purpose of mowing, for instance, is to edit data in preparation for a repeat of the selection/rejection procedure.

The above procedures also yield information that can be used to identify tall vegetation, trees and bushes. In addition, they pave the way for terrain slope determination. Slope variability is one of the indicators for terrain roughness.

Grid based binning is used for the selection of ground points. TIN based meshing is used to generate the surfaces with respect to which the point clouds may be mowed. TINs were also used to "mask" areas where high elevation variability indicates the presence of tall vegetation, supplementing the bin-based selection/rejection process.

3.1 Selecting Ground Points

Once a grid has been specified, the points (x_k, y_k, z_k) in the point cloud are arranged by bins (i,j) according to their locations (x_k, y_k) in the base plane. This enables local comparisons and statistics, which can be used to select potential ground points – one bin at a time – for the purpose of interpolating these points by a ground surface.

One approach is to simply select in each non-empty bin a point with the minimum elevation z_k, the reason being that this lowest point will most likely be a ground point.

However, if the elevations of the points in a particular bin vary extensively, then tall vegetation or artifacts may be found in that bin, and there is little confidence in its lowest point

being a ground point. Therefore, a rejection mechanism has been put in place, where a bin (i, j) is rejected if:

- there are fewer points in bin (i,j) than a specified minimum
- the RMS (root-mean-square) of the elevations z_k in the bin (i, j) is above a specified tolerance z_{tol}

$$\sqrt{\frac{1}{n(i,j)}\sum\left[z_k - z_{mean}(i,j)^2\right]} > z_{tol} \ .$$

The summation extends over the points in bin (i,j) where

$n(i, j)$ = number of points in bin (i,j)
z_{mean} = mean elevations of z_k in bin (i,j).

- bin (i, j) has been blocked by containing a point from a "masking" file, determined in the separate TIN-based masking procedure described in Section 3.1.1.

The reason for specifying minimum occupancy for bins is so as to have sufficient comparative information to support the selection of a ground point. However, some of the results reported in this paper accept single occupancy.

Bins, whose minimum elevation points were rejected because their RMS was too large, or which have been blocked by masking, are candidates for containing vegetation. The "reject" information has thus been used to derive statistics on vegetation coverage as described later in this paper.

The key to this method is the selection of an appropriate grid size – large or small. A small grid size is preferred as this enables better capture of such terrain features as ridgelines, ditches, valleys, etc. On the other hand, a large grid size may be required when *a priori* knowledge indicates that a majority of the points may be above ground.

As mentioned earlier ground determination is an iterative process. Grid size selection is also dependent on where the current stage of the process. For example, a larger grid size may be suitable in the early stages of the process where a rough ground surface may be sufficient for use in further screening and mowing the data points. In the later stages of the processing, a smaller grid size may be required for better resolution and capture of terrain features.

Grid size may also depend on the type of terrain. For example, for an open or semi-open terrain, a smaller grid size may be chosen because the chances of selecting a ground point is increased as compared to wooded terrain where a larger grid size may be required to increase the chances that at least one point in the bin is a ground point.

The problem of choosing an appropriate grid size is particularly challenging in the case of ground based LADARs, where there is a large discrepancy in the density of the locations (x_k, y_k) of the scan points (x_k, y_k, z_k). Indeed, for some scans, it was observed that about 70 % of the data points were within 10 m of the instrument.

3.1.1 Masking

In addition to their other roles in ground determination, TIN techniques can assist in ground point selection by identifying areas from which ground points should not be selected such as areas occupied by trees, bushes, and other artifacts. This technique is, therefore, referred to as "masking".

In this technique, the TIN algorithm is applied to the point cloud, but terminated after only 10 % to 15 % of the data points have been inserted. Of the two insertion selection options described in Section 2, the first one is selected. The result is a partial triangulation with small triangles concentrated at "hot spots" of elevation variations such as caused by trees and bushes. After deleting all triangles with edges of, say, 10 cm or greater, the remaining triangle vertices are the masking points referred to earlier. As an example of the masking process, an overhead view of a point cloud is provided in Fig. 3a, where the locations of vegetation are clearly seen. Fig. 3b shows the triangulation with no triangle deletion and Fig. 3c shows only the masked triangles after all triangles with edge length greater than 10 cm are deleted. The identification of tall vegetation in Fig. 3c clearly matches that shown in Fig. 3a.

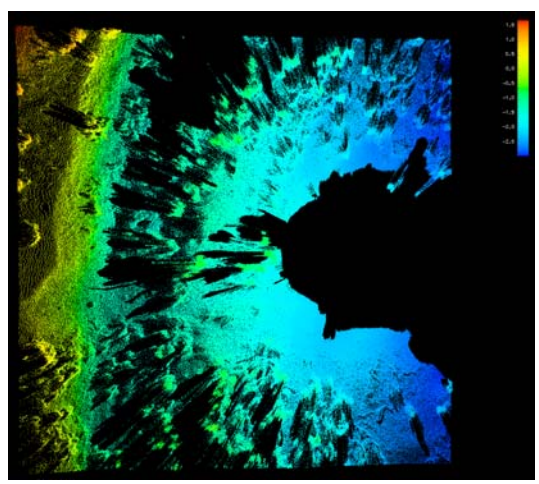

(a) Point cloud – Overhead view.
(b)
Figure 3. Vegetation ID Using Masking Technique (cont.)

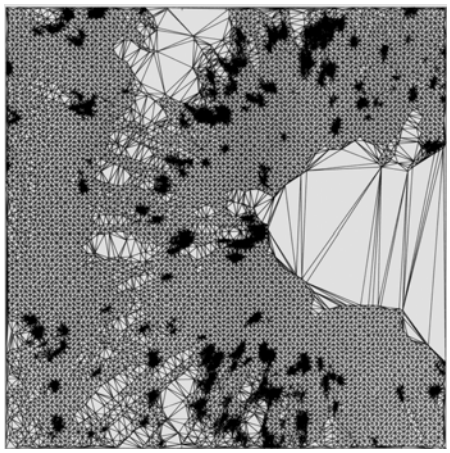

(b) Triangulation of mask file without deletion.

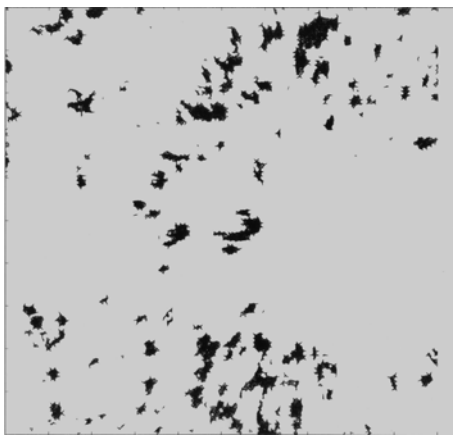

(c) Masked triangles only.

Figure 3 (cont.). Vegetation ID Using Masking Technique

3.2 Ground Surface Determination

The stage is now set for constructing an initial ground surface. Such a surface will be used for "mowing" the original data set, that is, removing points above the ground surface. The final ground surface will provide the ground estimate for the vegetation identification process (Section 4). Gridded data are also needed for visualization.

TIN meshing is used for interpolating the ground points, requiring a full triangulation of their locations. The reason for choosing interpolation rather than approximation based on partial triangulation is that the selected points are already a sample of the full data set.

Experience has shown that among the selected points there are still some points whose high elevations make them unlikely ground points. In the TIN surface, the elevation of such a point exceeds the median elevation of its neighbors by more than a specified tolerance, say, 25 cm. The current procedure is to delete such points.

3.3 Mowing

The purpose of mowing is to remove as many as possible of those points in a point cloud which are above ground such as canopy points. Repeating the ground determination procedures for the mowed data will reduce the instances of rejection, say, on the basis of RMS. It will thus lead to a richer collection of ground points and therefore to improved ground surface.

As indicated earlier, mowing determines for each point in a data set how high it is above a current ground surface. If that height is larger than a specified mowing tolerance, say 10 cm, then the point is removed from the data set.

4. VEGETATION IDENTIFICATION

An initial approach at quantifying coverage by tall vegetation is described. It relies on a predetermined ground surface, and it utilizes the "reject" information gathered during the select/reject process. Two attributes "tree" and "bush" were assigned to bins of the underlying grid.

The first application of the selection/rejection procedure provides key information in the form of elevation statistics for those bins (i,j) which had been rejected on the basis of RMS or masking blocks:

- $n(i,j)$ = number of points in bin (i,j)
- $(x_{min}(i,j), y_{min}(i,j), z_{min}(i,j))$ = minimum elevation point
- $z_{mean}(i,j)$ = mean of elevations
- $z_{max}(i,j)$ = maximum elevation

At location $(x_{min}(i,j), y_{min}(i,j))$, the ground surface assumes the elevation $z_{surface}(i,j)$, which will be considered ground elevation unless the elevation z_{min} is even lower. Thus

- $z_{ground}(i,j) = \min\{z_{urface}(i,j), z_{min}(i,j)\}$.

The criterion for trees is checked first:

- $n(i,j) > 20$
- $z_{max}(i,j) - z_{ground}(i,j) > 3.75$ m
- $z_{mean}(i,j) - z_{ground}(i,j) \geq 0.75$ m

followed by the criterion for bushes:

- $n_{bin}(i,j) > 20$
- $z_{mean}(i,j) - z_{ground}(i,j) \geq 0.25$ m

If neither set of criteria is met, the bin is not classified. The numerical parameters were selected experimentally.

Vegetation coverage is now determined as the ratios:

- *tree_coverage = #tree bins / #non-empty bins*
- *bush_coverage = #bush bins / #non-empty bins*

It is clear, that this characterization process does not necessarily identify individual trees and bushes, as any of them may impact several bins. Also, the canopy spread of a mature tree may cover several bins, which could well be traveled by the XUV with only the trunk as an obstacle. For a mobility analysis, therefore, one might consider mowing at, say, 3 m, prior to the process of selection/rejection. The elevation statistics of the reject bins can then be used for screening no-go bins.

5. SLOPE DETERMINATION

A major application requiring good ground surface is the determination of terrain slopes. Areas of steep slopes need to be identified for path planning as well as for post-travel assessments of areas where the XUV had problems. In addition, slope statistics may be used as an indicator for terrain roughness.

5.1 Surface Contours and Path Analysis

Given the TIN-meshed ground surface, gridded surface points (x_{ij}, y_{ij}, z_{ij}) can be extracted to facilitate surface visualization, and contouring. Given a vehicle path in the base plane by 2D points (x_p, y_p), elevations, z_p, can be inferred, and pitch and roll of the vehicle along the path can be calculated. Such calculations provide a means for ground verification.

The gridded data approximate a continuous ground surface $z = z(x, y)$. Similarly, the discrete path points approximate a smooth curve with defined travel directions. At any path point (x_p, y_p), let

- $\mathbf{t}$ = travel direction in the base plane; $\|\mathbf{t}\| = 1$
- $\mathbf{s}$ = right handed perpendicular to $\mathbf{t}$: $\mathbf{s} = \mathbf{t} \times \mathbf{z}$,

where $\mathbf{z}$ denotes the unit vector in the direction of the z-axis. The vector $\mathbf{s}$, therefore, is also a unit vector and lies in the base plane (see Figure 4). In the continuous model, the surface slopes in those directions are their inner products with the surface gradient $\nabla z = (z_x, z_y)$. They are also the tangents of the pitch and roll angles, respectively. Thus

- pitch = arctan(< $\mathbf{t}$, ∇z >)
- roll = arctan(< $\mathbf{s}$, ∇z >) .

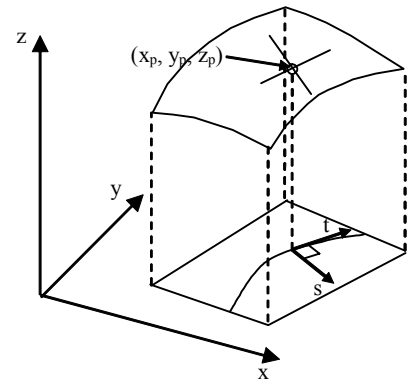

Figure 4. At path point (x_p, y_p, z_p) on the ground surface $z(x, y)$, tangential slopes in directions ***t*** and ***s*** define pitch and roll of the vehicle

The partial derivatives z_x , z_y and the travel direction $\mathbf{t}$ can be estimated locally using the – smoothed --gridded surface data and the 2D points describing the path traveled. Median filters are used for smoothing in this work.

5.2 Terrain roughness

There are several aspects of terrain roughness. The surface may be rough because it is covered with gravel, rocks, or ruts. Here, large slope differences occur within short distances of, say, 10 cm. Or, from the point of view of mobility, one would be interested in whether the terrain exhibits a large degree of undulation. Here, slope differences at distances of, say, 1 m are at issue. An effort to quantify this latter kind of slope variation is based on fitting least-squares planes to the points in individual bins, excluding those bins which for any reason were rejected during the selection/rejection process. This process yields for each accepted bin (i,j) , the quantities:

- $xslope(i,j)$ = slope of plane in x-direction
- $yslope(i,j)$ = slope of plane in y-direction

Statistics are derived for those quantities:

- $xslope_{RMS}$ = RMS of the slopes $xslope(i,j)$
- $yslope_{RMS}$ = RMS of the slopes $yslope(i,j)$

The following gauge for terrain roughness has been implemented:

- $roughness = \sqrt{(xslope_{RMS})^2 + (yslope_{RMS})^2}$

It was chosen, because it could be shown to satisfy the following two properties:

- it vanishes if and only if the surface is a plane
- if it assumes the same value in two areas, then it assumes the same value in the union of the areas.

6. EXAMPLES

6.1 Ground Surface

As the point density is extremely high around the scanner and decreases with distance away from the scanner, the region of interest was selected as an area 20 m x 20 m around the instrument to reduce the amount of no data areas.

The sequence of procedures to determine the ground is highly subjective. The general procedure followed to determine the ground in this work was :

- create a mask file – file of points to block bins as part of the rejection process.
- perform gridded binning to select initial ground points (1 m grid).
- mesh and screen the ground points
- mow original data file (1m)
- screen the mowed file
- create a mask file using the mowed file
- perform gridded binning to select second set of ground points
- mesh and screen the ground points
- mow (10 cm) the original data file using the surface generated by the second set of ground points
- perform gridded binning to select third set of ground points
- mesh/screen the ground points

The results of two ground surfaces are shown in Figs. 5 and 6. Figure 6 shows an overhead view of the ground surface in a wooded region in Ft. Indiantown Gap (FTIG), PA. The monotone gray regions in the Fig. 5b are no-data regions. Figure 6 was obtained at Tooele, UT, representative of an arid environment. The surface shown in Fig. 6 is of the dam region – the same region as shown in Fig. 8.

(a) Panoramic photo

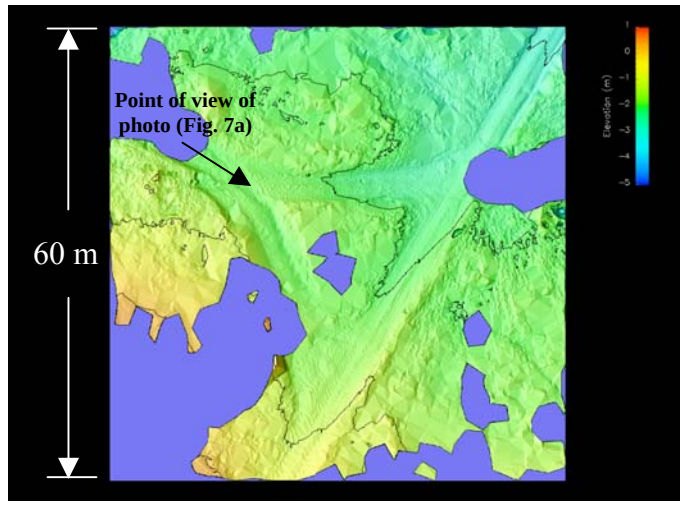

(b) Ground surface

Figure 5. 3D surface of region of interest (E-Stop 2) in a wooded area. Most of the trees were successfully removed.

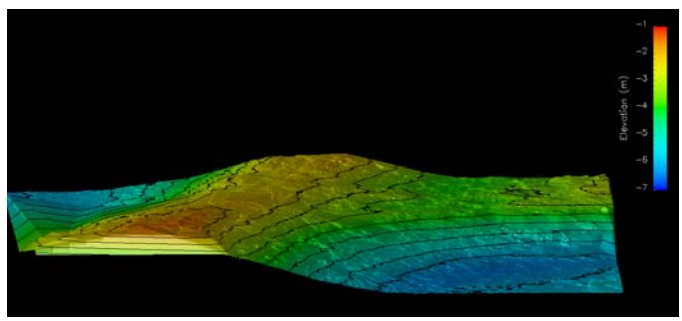

Figure 6. Ground surface of dam region.

6.2 Vegetation ID

The sequence of procedures to identify the vegetation was slightly less involved than that for ground determination. The general procedure was:

- create a mask file – file of points to block bins as part of the rejection process
- perform gridded binning to select initial ground points (1 m grid).
- mesh and screen the ground points
- tree-bush ID

In the wooded environment, the mowing technique had to be applied as the tree canopy resulted in “false” trees being identified.

Some results of the vegetation ID are shown in Fig. 7.

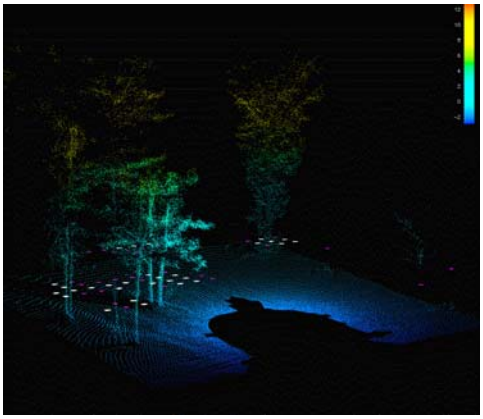

(a) "Stop 20"

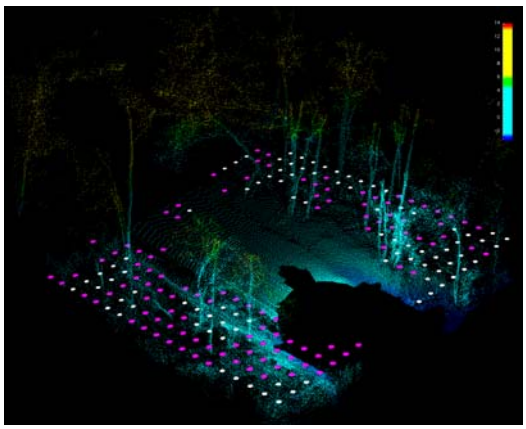

(b) "Stop 7"

Figure 7. FTIG, Black Course

6.3 Path Analysis

Figure 8 provides an overhead view of a 3D surface with steep slopes, with contours drawn every 0.2 m over a 200 m x 200 m region ("dam" region). It also shows the path of the XUV over the dam beginning at the right of the figure and ending at the top. There are marks at 50m, 100 m, 150 m, 200 m, and 250 m along the path. These represent distances along the path from the starting point. Figure 9 shows a profile of the pitch and roll of the XUV along the path in Figure 8.

At approximately 75 m along the path there are large negative (downward) and positive (upward) pitches indicated. There is indication that the upward pitch is also accompanied by the large positive (counterclockwise) roll.

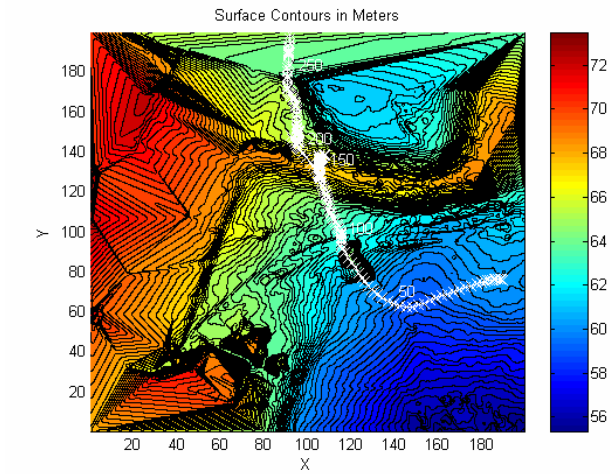

Figure 8. 200 m x 200 m, Overhead view

The path profile in Figure 10 seems to indicate a forward and backward motion at the top of the dam. This appears as the sawtooth portion of the plot in Figure 10.

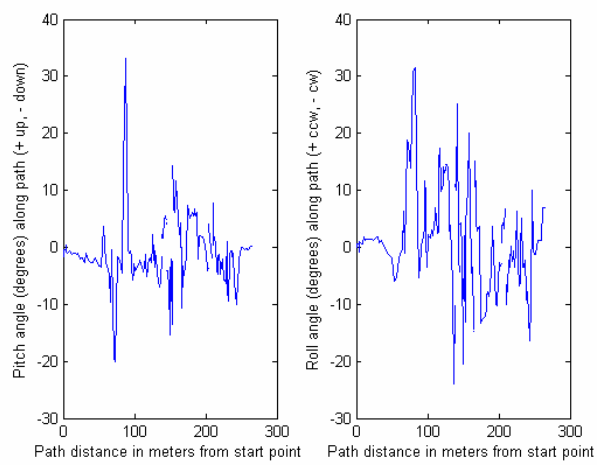

Figure 9. Pitch and Roll profiles along the path in Fig. 8.

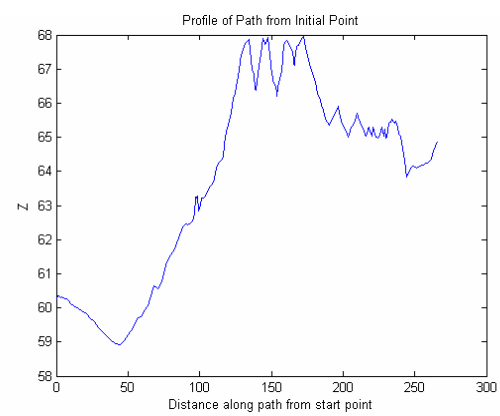

Figure 10. Path profile in Fig. 8.

7. FUTURE WORK

The selection/rejection procedure described in this paper has two problems:

- it fails at steep slopes because the elevation distribution mimics that of tall vegetation, resulting in "false negatives"
- it fails to take into account the large discrepancies of point location densities which typically result from ground-based LADAR scanning

The solution to the first problem is to fit planes to the points in each individual bin, and to interpret elevations relative to those planes, respectively. In this framework, it would also make sense to consider a sliding bin as a device for finding more ground points.

As to the second problem, it is proposed to modify the grid so that its bins are arranged concentrically around the instrument. The bins in that grid should be proportional to each other and have shapes that are approximately square. To this end, it is stipulated that the bins should share with squares the property that

- their two diagonals are perpendicular to each other.

The shape and size of a bin is then uniquely determined by the angle of its two straight sides and the shorter of its circular arcs, and the entire grid is determined by an inner radius, an outer radius, and the angle increment between bins. Figure 11 illustrates the generation of the grid from that geometric definition. Figure 12 presents an example of a concentric grid. In both figures, the bins in the outermost ring have been extended to accommodate residual bins due to the specification of a particular outer radius.

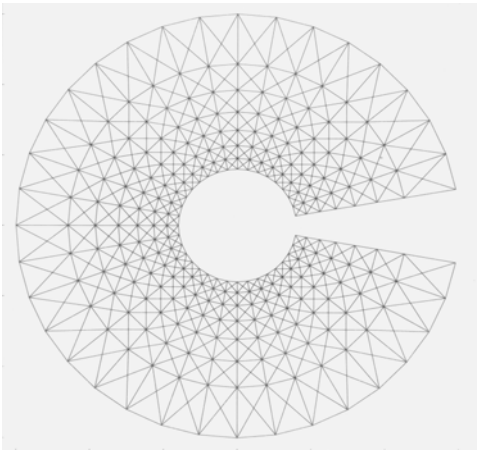

Figure 11. Generation of Grid Pattern by Perpendicular Diagonals.

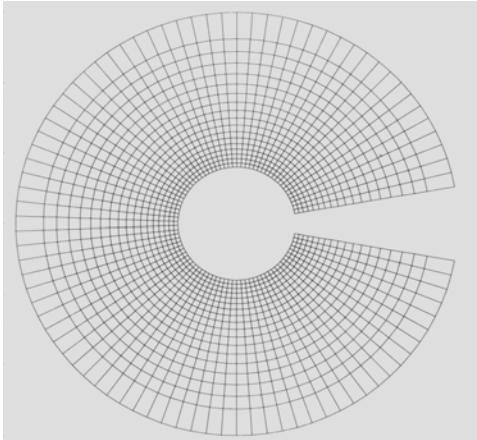

Figure 12. Instance of a Concentric Grid.

8. SUMMARY

Ongoing work on issues of terrain analysis and characterization has been reported. Those issues center around the determination of ground from point clouds. In most instances, this task is difficult because a large portion of the data represent points above ground, such as foliage. The task is particularly difficult for ground-based LADAR, because of the large discrepancy of data density; the bulk of the data is found in the immediate vicinity of the instrument. At larger distances, the points are sparsely distributed and mostly above ground.

Several procedures for ground determination have been described. These procedures employed grid based binning and TIN based methods. Applied in various combinations and repetitions, these procedures have yielded promising results for data collected at Tooele, UT, and Ft. Indiantown Gap, PA. Slopes were determined based on ground surface. Other applications developed in this work are preliminary numerical indicators or gauges for *i*) coverage by tall vegetation and *ii*) terrain roughness. A further goal is to gather experience with a variety of data sets in order to analyze, evaluate, and improve ground-based procedures.

REFERENCES

[1] Samet, H. [1990], "*The Design and Analysis of Spatial Data Structure*," Addison Wesley, Reading, PA.

[2] Witzgall, C., Bernal, J., and Cheok, G. S. [2003], "TIN Techniques in Data Analysis and Surface Construction," NISTIR in preparation.