

MAPPING OF GDOP ESTIMATES THROUGH THE USE OF LIDAR DATA

KRISTA AMOLINS

June 2008



**TECHNICAL REPORT
NO. 259**

MAPPING OF GDOP ESTIMATES THROUGH THE USE OF LIDAR DATA

Krista Amolins

Department of Geodesy and Geomatics Engineering
University of New Brunswick
P.O. Box 4400
Fredericton, N.B.
Canada
E3B 5A3

June 2008

© Krista Amolins 2008

PREFACE

This technical report is a reproduction of a thesis submitted in partial fulfillment of the requirements for the degree of Master of Science in Engineering in the Department of Geodesy and Geomatics Engineering, June 2008. The research was supervised by Drs. Peter Dare and Yun Zhang, and support was provided by the Natural Sciences and Engineering Research Council of Canada and by the New Brunswick Innovation Foundation.

As with any copyrighted material, permission to reprint or quote extensively from this report must be received from the author. The citation to this work should appear as follows:

Amolins, Krista (2008). *Mapping of GDOP Estimates Through the Use of Lidar Data*. M.Sc.E. thesis, Department of Geodesy and Geomatics Engineering Technical Report No. 259, University of New Brunswick, Fredericton, New Brunswick, Canada, 125 pp.

ABSTRACT

The positioning accuracy of the Global Positioning System (GPS) and other Global Navigation Satellite Systems is affected by the configuration of visible satellites. Dilution of Precision (DOP) values are a measure of the strength of the satellite configuration but the software tools currently available for calculating DOP values have a limited ability to take into account obstructions. Determining when the best satellite configuration will be observable at a particular location requires identifying obstructions in the area and ascertaining whether they are blocking satellite signals.

In this research, Light Detection and Ranging (LiDAR) data were used to locate all the obstructions around each terrain point by extracting and comparing two surfaces, one that represented obstructions and one that represented the terrain. Once all the obstructions in a selected area had been identified, GPS satellite location data were used to determine satellite visibility at different epochs and to calculate GDOP (Geometrical DOP) at locations where at least four satellites were visible. Maps were then generated for each epoch showing the GDOP values over the selected area. Some small differences were noted between the clear sky GDOP values calculated by the proposed method and those output by an available software planning tool and in a few cases there was a discrepancy in the number of visible satellites identified due to slight differences in the calculated satellite elevations. Nevertheless, the maps produced by the proposed method give a better representation of the GDOP values in the field than do traditional methods or other software tools.

TABLE OF CONTENTS

	Page
ABSTRACT	ii
TABLE OF CONTENTS	iii
LIST OF FIGURES	v
LIST OF TABLES	vii
1. INTRODUCTION	1
1.1. Global Positioning System.....	3
1.2. Airborne Laser Scanning	6
1.3. Existing Research	9
1.3.1. <i>Digital Terrain Models</i>	9
1.3.2. <i>Use of LiDAR Data in DOP Prediction</i>	12
1.4. Research Overview	14
1.4.1. <i>Objective</i>	14
1.4.2. <i>Proposed method</i>	15
1.4.3. <i>Outline</i>	17
2. LIDAR DATA MANIPULATION.....	20
2.1. Description of Sample Data.....	21
2.2. Filtering.....	25
2.3. Potential Obstructions.....	27
2.3.1. <i>Point Cloud Representation</i>	28
2.3.2. <i>Interpolated Surface</i>	31
2.3.3. <i>Removing Small Obstructions</i>	34
2.4. Terrain Test Area	35
3. DSM GENERATION	39
3.1. Surface Considerations	40
3.2. Interpolation to Grid	42
3.3. Removing Non-surface Objects.....	46
4. IDENTIFYING OBSTRUCTIONS.....	53
4.1. Elevation Angle	55
4.2. Direction Angle.....	57
4.3. Generating Obstruction Images	60
5. CALCULATING GDOP VALUES.....	65

5.1. Satellite Data.....	65
5.2. Converting Coordinates	66
5.3. Identifying Visible Satellites	68
5.4. Dilution of Precision	71
6. DISPLAYING RESULTS	76
6.1. Creating Images	76
6.2. Adding Features	78
6.3. Static Images.....	80
6.3.1. <i>Difference Images</i>	83
6.3.2. <i>Optimal Time</i>	85
6.3.3. <i>Maximum and Minimum</i>	86
6.4. Dynamic Images	87
7. ASSESSMENT	90
7.1. Assessment Data	91
7.2. Generated DSM	93
7.3. Visible Satellites and GDOP Calculations.....	95
7.4. Effect of Obstructions.....	98
8. CONCLUSIONS AND RECOMMENDATIONS	104
8.1. Conclusions.....	105
8.1.1. <i>Quantifying Obstructions</i>	106
8.1.2. <i>Calculating and Mapping GDOP Values</i>	107
8.1.3. <i>Research Accomplishments</i>	109
8.2. Recommendations.....	110
8.2.1. <i>Performance Tests</i>	110
8.2.2. <i>Obstruction Classification</i>	111
8.2.3. <i>Interpolating between Epochs</i>	112
8.2.4. <i>Extending Planning Capability</i>	113
8.2.5. <i>Extracting Surfaces from Other DTM Sources</i>	113
8.2.6. <i>User Interface</i>	114
REFERENCES.....	115
APPENDIX I: COORDINATE CONVERSION.....	118
APPENDIX II: ADDITIONAL ASSESSMENT.....	121
CURRICULUM VITAE	

LIST OF FIGURES

	Page
1.1 Schematic of proposed method.....	16
2.1. Flight pattern for airborne laser survey over Fredericton, NB.	22
2.2. Representation of campus elevation.	23
2.3. Along-track point spacing. a) Single flight strip. b) Overlapping flight strips.....	24
2.4. Noise at centre of flight lines.	25
2.5. Three-dimensional campus representation with noise.....	26
2.6. First returns, point representation. a) Full campus. b) Power lines.	30
2.7. Effect of grid resolution. a) Original LiDAR data. b) Interpolated surface with 0.35 m resolution. c) Interpolated surface with 0.5 m resolution. d) Interpolated surface with 0.74 m resolution.	33
2.8. Full campus and selected subarea. a) Elevation of full campus. b) Elevation of selected area. c) Intensity of selected area.....	37
3.1. a) Last returns in selected test area. b) Terrain visualization of last returns averaged over a 1 x 1 m grid.	43
3.2. Empty cells when a) cell boundaries used, b) circle radius $1/\sqrt{2}$ used in terrain... 44	44
3.3. a) Surface elevations after each of the first five iterations of erosion. b) Elevation differences between consecutive iterations.	49
3.4. Enlarged section of the test area: a) original interpolated surface; b) surface after third erosion iteration. Highlighted areas: (i) cars; (ii) building; (iii) vegetation.. 50	50
3.5. Difference between original interpolated surface and surface after first erosion iteration.....	50
4.1. a) Three-dimensional representation of obstruction point cloud. b) Three- dimensional representation of DSM.	54
4.2. Effect of building height and proximity on sky view.	55
4.3. a) Elevation angle θ and b) direction angle ψ of a potential obstruction relative to a surface point (x_0, y_0, z_0)	56
4.4. Effect of building location on sky view.....	58
4.5. Relative elevation angle of obstructions by quadrant: a) six levels b) eighteen levels.....	61
4.6. a) Obstructions from interpolated surface, 6 levels and 4 sectors. b) Obstructions from point cloud, 18 levels and 8 sectors.	63
5.1. a) Direction vectors in UTM coordinate system. b) Representation of test area in ECEF coordinate system.	68

5.2.	Satellite a) elevation and b) direction angles.....	69
5.3.	Number of satellites visible at epoch 01:15:00. a) Five degree elevation mask. b) Ten degree elevation mask. c) Fifteen degree elevation mask. d) Twenty degree elevation mask.....	72
5.4.	GDOP values at epoch 00:00:00, scaled to seven classifications.....	75
6.1.	GDOP values at epoch 00:00:00, mapped using 240 colours representing seven classifications a) without object outlines b) with object outlines.....	79
6.2.	a) GDOP values for five epochs, from 00:00:00 to 00:01:00. b) Change in GDOP values between epochs.....	82
6.3.	Optimal two-hour interval between 06:00:00 and 22:00:00.....	86
6.4.	GDOP value extremes between 08:00:00 and 10:00:00. a) Minimum. b) Maximum	87
6.5.	Possible web browser-based interface for viewing GDOP maps.....	89
7.1.	Three-dimensional views of assessment target locations: a) commercial parking lot; b) residential rooftop; c) Head Hall parking lot; d) Provincial Archives parking lot; e) sidewalk at Sir Howard Douglas Hall. f) Aerial view of target point locations.....	92
7.2.	Standard deviation of elevations within grid cells: a) interpolated terrain, scale range 0 – 7 m; b) DSM, scale range 0 – 1 m.....	95
7.3.	Visible satellites and GDOP values from <i>Trimble Planning</i> and calculated by proposed algorithm for epochs from 22:30:00, 20 January 2008 to 13:15:00, 21 January 2008.....	96
7.4.	Satellite visibility from 22:30:00, 20 January 2008 to 13:15:00, 21 January 2008...	96
7.5.	Sky view plot, 22:30, 20 January 2008.....	99
7.6.	Sky view plot, 03:00, 21 January 2008.....	101
7.7.	Sky view plot, 07:00, 21 January 2008.....	102
II.1.	Sky view plot, 00:00, 21 January 2008.....	122
II.2.	Sky view plot, 01:30, 21 January 2008.....	123
II.3.	Sky view plot, 04:30, 21 January 2008.....	123
II.4.	Sky view plot, 06:00, 21 January 2008.....	124
II.5.	Sky view plot, 09:00, 21 January 2008.....	124
II.6.	Sky view plot, 10:30, 21 January 2008.....	125
II.7.	Sky view plot, 12:00, 21 January 2008.....	125

LIST OF TABLES

	Page
2.1. Points by return.....	27
2.2. Filtered first returns.	35
7.1. Target point location descriptions and sector obstruction levels.....	92
7.2. Algorithm-determined satellite visibility at 22:30, 20 January 2008.	99
7.3. Algorithm-determined satellite visibility at 03:00, 21 January 2008.	101
7.4. Algorithm-determined satellite visibility at 07:00, 21 January 2008.	102
I.1. Conversion constants.	118
I.2. Eccentricity coefficients.	118
II.1. Satellites visible at target locations.	121

CHAPTER 1

INTRODUCTION

Knowledge advances by steps, and not by leaps.

~ Lord Macaulay

The development of Global Navigation Satellite Systems (GNSS), in particular the Global Positioning System (GPS) allows positions on the Earth to be determined with increasing accuracy and precision. The systems have applications ranging from site surveying and ground monitoring, to navigation aids and geo-caching. For the latter, which are applications that the general public most associate with GPS, positioning accuracy within a few metres is sufficient; for the former, much greater accuracy is needed [Baltsavias, 1999].

Several factors affect the overall accuracy of GPS, including errors and uncertainties in satellite orbits and clocks, atmospheric effects, and configuration geometry of the visible satellites. Dilution of precision (DOP) is a scalar value that represents the geometrical relationship between positioning accuracy and satellite configuration. There are different variants of DOP to reflect the accuracy of different position components, such as vertical or horizontal accuracy. The most readily available data are for Position DOP (PDOP): the GPS Operations Center publishes maps showing the maximum predicted PDOP (Position DOP) values globally for a 24-hour period [US Air Force Space Command, 2007]. Geometric DOP (GDOP), however, reflects overall accuracy including time, rather than just position accuracy, and so is of greater interest in this

study. It is calculated using the positions of four visible satellites and the position of the receiver. The ideal DOP value is one. Values of two to three are considered excellent, while values greater than nine are considered high and should be discarded.

Generally, more than four satellites are above the receivers' horizon at any time and the GDOP calculation can incorporate as many satellites as the receiver is able to track. Satellite signals, however, are not able to penetrate solid objects. Some signal detection is possible below trees or inside houses, particularly if the receiver is located near a window, but the position accuracy is likely to be greatly reduced. While it may still be adequate for use in the increasingly popular hand-held GPS navigation systems, the reduced accuracy is not adequate for most other applications. Furthermore, larger buildings can block signal reception entirely, both inside the building and in a shadow area around the building.

Ideally, receivers should only be placed at locations where there is a clear view of the sky but, since this is not always possible, the presence of sky obstructions needs to be considered when a GPS survey is being planned. Existing methods to incorporate the presence of obstructions in DOP estimates can be inaccurate, time-consuming, and/or require knowledge about the height and location of potential obstructions. Software tools exist to assess the ideal satellite configuration or DOP value at a given time; however, they require the user to know the intended location of the receiver and the ability to include obstructions is limited. In many cases, the receiver location is definite, but where there is flexibility in selecting the receiver location, it would be useful to be able to incorporate obstructions in the calculation of DOP values and compare these values over

a wide area. It could also be useful after data are collected to verify that the DOP values were sufficiently low over the entire collection period to ensure the quality of the results. In non-surveying applications, such as in-car navigation, by incorporating obstructions locations where positioning and navigation are unreliable could be identified, overall or at specific times.

Airborne Light Detection and Ranging (LiDAR) systems, also known as airborne laser scanning systems, provide detailed information about terrain and object elevations in the scanned area. The raw LiDAR data form a point cloud from which different types of surfaces can be extracted. For instance, the terrain surface can be extracted and, with appropriate processing, serve as a map of potential receiver locations. An obstruction surface could also be extracted that would include both objects, such as trees and buildings, and the terrain itself, which could obstruct part of the sky from view. By comparing these surfaces, the potential obstructions can be assessed, and the results can then be used in conjunction with GPS satellite data to create maps showing DOP values that reflect the situation in the field.

1.1. Global Positioning System

GPS is a satellite-based positioning system designed and maintained by the US military. Intended primarily for military purposes, the signals transmitted by the satellites are nonetheless freely available to anyone with the proper receiving equipment. The number of non-military applications in which GPS data are used, and the number of devices that provide at least rudimentary positioning, now greatly outnumber military

ones; the decision by the US Government in 2000 to discontinue Selective Availability, which was an intentional degradation of the signal accuracy for users without access to the Precise Positioning Service, further increased the utility of the Standard Positioning Service and the ubiquity of GPS devices. The primary applications are mapping and surveying, as well as land, sea, and air positioning and navigation [US Coast Guard, n.d.].

The basis of GPS is a constellation of a minimum of 24 satellites in six orbital planes inclined 55° to the Earth's equator. The placement of the satellites within the constellation ensures that at least four satellites are visible from any location on Earth at any time. With 31 satellites currently in the constellation, more satellites may be visible, and usually are, potentially improving the positioning accuracy. Some of the extra satellites are in place in case of satellite failures; others are new satellites launched as part of the GPS modernization process. The satellites have a nominal altitude of 20 200 km, corresponding to a period of 12 sidereal hours [US Coast Guard, n.d.].

Individual satellites are identified by a pseudo-random noise code, which repeats every millisecond. The range to a satellite is determined by aligning a receiver-generated replica of the code with the arriving signal and determining the time offset. Each satellite also transmits a message containing specific information about the satellite itself, including its clock and ephemeris data, and about the other satellites in the constellation. Once a receiver has successfully received the message from one satellite, it can be used to quickly locate and start tracking other satellites. To obtain a position fix, a minimum of three satellite ranges must be observed. Each range can be likened to the radius of a

sphere, centred about the associated satellite, on the surface of which the receiver is located; three spheres intersect at two points, one of which can generally be easily discarded and the other of which is the receiver position. Satellites have very accurate atomic clocks, but the receiver clock is much less accurate and so a fourth satellite is needed to resolve clock bias. The position is calculated in an Earth-Centred, Earth-Fixed (ECEF) coordinate system, in x, y, and z coordinates, and transformed by the receiver into latitude, longitude, and height [Dana, 2000].

A number of factors affect positioning accuracy, such as ephemeris uncertainties, propagation errors, clock and timing errors, and receiver noise. These combined contribute to the measurement accuracy, or the user equivalent range error. Generally, the longer a satellite is tracked, the more accurate the resulting position. The relationship of satellites, however, is a factor that can have a very large effect on accuracy. DOP measures the strength of satellite configuration, or the size of the confidence region for position solutions. The DOP value multiplied by the measurement accuracy gives the positioning accuracy [Wells et al., 1986].

The variants of DOP are calculated from the trace of the covariance matrix of the direction vectors between the receiver and visible satellites. If the satellites are widely spread out, the diagonal entries of the covariance matrix will be low and so the DOP value will be low; if the satellites are clustered together, the opposite is true. Some of the DOP variants are: TDOP, time; VDOP, vertical; HDOP, horizontal; PDOP, position; and GDOP, geometric. GDOP is the variant used in this study because it takes into account all four coordinates (x, y, z, and time) and is an overall indication of the effect of the

satellite geometry. The lowest GDOP value is achieved when one satellite is directly overhead and the other three are on the horizon, 120° apart. This cannot be achieved in practice, however, because to reduce biases from incorrect modelling of tropospheric delays, a mask angle between 8° and 15° is set, cutting out satellites close to the horizon [Wells et al., 1986].

The International GNSS Service (IGS) is an organization that provides data and products related to GNSS. Primarily, the data are from GPS as currently the Russian GNSS (GLONASS) is incomplete, with only thirteen satellites in operation, and the European system Galileo is not expected to be operational until 2013. In addition to Earth-based products such as Earth rotation parameters and coordinates of IGS tracking centres, IGS makes available GPS satellite ephemerides. These are provided in a standard file format, denoted SP3, as Ultra-Rapid, Rapid, or Final values. The files contain satellite positions, clock corrections, and other information for all active satellites over 24 hours at fifteen-minute intervals. The Ultra-Rapid orbits consist of predicted satellite positions, with approximately 10 cm accuracy, and observed satellite positions, with better than 5 cm accuracy, updated four times daily. The Rapid values have a 17-hour latency while the Final values have a 13-day latency. Both have a better than 5 cm accuracy [Beutler et al., 2005].

1.2. Airborne Laser Scanning

A LiDAR system collects data through the use of a laser-ranging unit, which sends out high-energy pulses in short intervals and detects the energy reflected back from

objects within the field of view. LiDAR systems are used in bathymetry and atmospheric studies, as well as terrestrial surveys, but for simplicity, further discussion and use of the term will relate only to the latter. The transmitting and receiving apertures in the system are mounted so as to have the same optical path. LiDAR calculates the distance between the sensor and any objects in the illuminated spot that return energy to the sensor from the time delay between transmission and reception. Theoretically, the maximum range is limited by the maximum value of the time counter, but in practice, the maximum range is limited by laser energy loss and other factors [Wehr and Lohr, 1999].

To convert the range into a position in a particular coordinate system, the position and orientation of the scanner with respect to that coordinate system is needed. This is obtained using an onboard Positioning and Orientation System (POS). Since travel times are recorded to nearly one nanosecond, giving a range accuracy on the order of 10 cm or better, the POS must have an equivalent or better accuracy; this can be achieved by including both an Inertial Measurement Unit (IMU) and Differential GPS (DGPS). The reference coordinate system is WGS84. At the completion of a scan, the POS data, calibration data and mounting parameters for the POS and for the laser scanner with respect to the POS, and the ranges with instantaneous scanning angles are combined to determine x, y, and z coordinates for each return [Wehr and Lohr, 1999]. Accuracy specifications for LiDAR data vary between systems and manufacturers, with typical values being 15 cm vertical and less than 1 m horizontal [National Oceanic and Atmospheric Administration, 2008.].

The instantaneous field of view is defined by the narrow divergence of the laser. Typically, it is between 0.3 mrad and 2 mrad. The motion of the laser across the flight direction and the forward motion of the plane define the scanned area. On flat and bare terrain, the point spacing in the across-flight direction would mainly be determined by the pulse repetition frequency, while in the along-flight direction it would be determined by the plane's ground speed and the period of the scan line. Different scanners produce different scan patterns, such as zigzag, bidirectional meander, parallel lines, and elliptical. The actual scan pattern on the ground depends on the flying direction, ground speed, and terrain topography [Wehr and Lohr, 1999]. The resulting data are therefore sub-randomly distributed point clouds, with density affected by flying height in addition to the other factors [Axelsson, 1999]. Point densities as high as 2.5 or 3.0 points/m² are achievable.

The LiDAR points represent the terrain or any objects, for example, trees and other vegetation, buildings, cars, and power lines, that lie above the terrain. The laser cannot penetrate solid surfaces and can generally only partially penetrate vegetation, typically to between 20% and 40% below the upper surface. The beams that succeed in penetrating the vegetation canopy may be reflected from within the tree crown, from objects located below the tree crown, from the ground, or from any or all of these, resulting in multiple returns. Other locations where multiple returns may be generated are the edges of rooftops and power lines. Up to five returns can be recorded and processed for each pulse [Ackermann, 1999].

A ground surface model can be derived through mathematical modelling. The accuracy may vary within the model because the LiDAR system cannot be controlled to obtain measurements from specific targets, resulting in denser or sparser point clusters [Ackermann, 1999]. The terrain points will clearly be sparser under vegetation, though multiple returns can be useful in filtering out the vegetation without losing the terrain points. Since some information can be lost if the points are interpolated to a grid, the best strategy is therefore to use the original data throughout the processing stages [Axelsson, 1999].

1.3. Existing Research

1.3.1. Digital Terrain Models

Airborne laser scanning has been under development for decades. As the technology improved, the potential applications became more widely recognized and studied. One of the main focuses has been deriving a Digital Terrain Model (DTM) from the LiDAR point cloud that is at least as accurate as what could be achieved with photogrammetry. Kraus and Pfeifer [1998] recognized that although no more than 25% of the LiDAR rays were able to penetrate the tree canopy, this was a better result than with photogrammetry. However, they also detected significant systematic errors that indicated to them deficiencies in the LiDAR system that would have to be resolved before the full potential of LiDAR could be realized.

The approach proposed by Kraus and Pfeifer [1998] uses linear prediction to derive a DTM. They determined an initial surface from the data by assigning equal weight to all

points. Once the surface was generated, they calculated residuals by subtracting the height of the surface from the height of the data points. These residuals were then used to determine new weights: terrain points were likely to have negative residuals, while vegetation points were likely to have small negative or positive residuals. By applying a shift value, negative residuals could be given full weight and large positive residuals could be given zero weight. The surface was then recomputed using the new weights and the process repeated. Roggero [2001] also used linear regression, using local slope to determine weights and a threshold to classify points as ground or non-ground. He determined that the results were very sensitive to the parameters used in the linear regression.

A different approach was proposed by Vosselman [2000]. He criticized the methods of other researchers for their assumption of a locally horizontal terrain, since this necessitated different parameters for different terrain types, but accepted the premise that a large height difference between two nearby points was unlikely to be due to the terrain slope. A challenge that he identified is the lack of a logical order within the point cloud, and the loss of information implicit in interpolating the data to a regular grid. His proposed method therefore involved a Delaunay triangulation to help locate all the points within a specified distance of a target point. Ground points were distinguished from non-ground points on the basis that where there was a height difference between two points, the higher point was unlikely to be terrain.

Although some information is lost by interpolating the data to a regular grid, the step allows standard image analysis techniques to be applied. Zhang et al. [2003] used

mathematical morphology, theorizing that using a varying filter size could help to better detect non-ground objects of different sizes. They constructed a regularly spaced surface from the data by selecting the point with the minimum elevation value in each cell and assigning the nearest value to empty cells, suggesting this was an improvement on other interpolation methods because the cell values were directly from the source data. To separate ground from non-ground, repeated openings were applied, with the window size increasing at each iteration, and points were classified as non-ground if they exceeded an elevation difference threshold.

Brovelli and Cannata [2004] used edge detection, edge thinning, and region growing. Once the data were interpolated to a grid, their method detected, classified, and thinned edges based on the sharp rise of the surface and positive residuals inside the objects compared to negative residuals outside. Region growing was then applied to detect and classify the inside edges of features. Miliaris and Kokkas [2007] applied median filtering to locate seed pixels for region growing as median filtering removes very high and very low values from the surface while preserving edges. The seed cells were building edges, identified by mid-range positive differences between the original surface and the median filtered surface. Region growing segmentation used slope criteria.

Most algorithms developed for deriving DTMs perform sufficiently well with general terrain types but less well with particular terrain types, such as complex urban areas and dense forest cover areas. Shan and Sampath [2005] proposed a one-dimensional, bidirectional labeling method for urban areas in which the raw LiDAR data were examined sequentially, on the assumption that consecutive points follow the scan

pattern. Ground points are identified by the slope and elevation difference: the slope between two points is expected to be larger at the boundary between building and terrain than it is within a typical terrain area, and ground points have lower elevations than those in a building neighbourhood. The labeling was performed first in the forward direction and then in the reverse direction, with ground points only being identified as such if they were labeled as ground in both directions. In contrast, Kobler et al. [2007] focused on steep relief under dense forest cover. They named their method REIN, for REpetitive INterpolation. Multiple elevation estimates were made for ground points, and the DTM was generated using the distribution of these estimates. This method relies on removing all negative outliers and most off-ground returns in the initial filtering stage.

A common consideration is that the range of scales over which features vary is smaller than the range of scales over which terrain can vary. Silvan-Cardenas and Wang [2006] proposed using the multiscale Hermite transform to adaptively erode gridded LiDAR data. The signal was analyzed by convolution with a bank of Gaussian derivative filters. After each convolution, the signal had been separated into approximation and detail coefficients. Non-ground features were detected using a gradient threshold and removed by subtracting the gradient from the approximation coefficient and setting the detail coefficients to zero.

1.3.2. Use of LiDAR Data in DOP Prediction

Applications of LiDAR data beyond DTM generation have also been studied and explored. A few recent studies use LiDAR data to predict GPS error. While these

recognize the potential for using LiDAR data to identify sky view obstructions, they do not attempt to explicitly calculate DOP values but rather rely on existing software to do so or only determine the number of satellites visible at different locations. Making use of existing software can be an efficient way to accomplish a task; however, there are limitations. As previously mentioned, software tools that calculate DOP values do not offer much flexibility in specifying location, allowing only a single input accurate to the nearest minute. More importantly, they are not designed in such a way that varying obstructions can be easily included. This means that ultimately, unless only a very small number of locations were under consideration, relying on existing software would be unnecessarily time consuming.

Beesley [2003] tried to model PDOP values as a function of percentage of sky visible. She used building footprints derived from aerial photographs to identify building points in the LiDAR data and a similar approach to identify vegetation points. The sky was divided into 36 wedges, each 10° wide. The maximum upward angle was determined based on the height of objects within a wedge and their distance from the reference point. The percentage of sky visible was then calculated from the maximum upward angle in all the wedges. Beesley attempted to draw conclusions on the effect of percentage of sky visible on PDOP values by comparing software-predicted values with field-observed values. Although there is some validity to the claim that when a lower percentage of the sky is visible the PDOP values will be higher, a direct relationship cannot be drawn between the two values because the position of the satellites relative to the sky obstructions is a determining factor.

Lohani and Kumar [2007] focused on the effect of trees, specifically trying to model the probability of satellite signals being obstructed by trees and therefore the probability that the GDOP value is increased. They used field data extensively as well as the features of Leica *SkiPro*TM software to develop and test their model at four specific terrain points with known obstructions. Where satellites were determined to be behind trees, they calculated the GDOP value for all possible combinations of visible satellites and used the probability of a satellite being obstructed by the tree to determine the probability of that GDOP value. Their results confirm that trees affect GDOP values, but are location-specific and rely on detailed knowledge of the receiver location, including the type of trees present.

In Taylor et al. [2005], the objective was to predict the number of satellites visible at any location at any time, in real-time. This was accomplished by applying line-of-sight analysis to a digital surface model generated from LiDAR data. Although this approach was deemed successful, as the authors note a minimum number of visible satellites does not ensure a good position quality, again because the position of the visible satellites must be considered. As with the other two studies, further work is needed to make this method useful for a broad range of applications.

1.4. Research Overview

1.4.1. Objective

The objective of this work is to develop a method through which maps can be generated that show GDOP values over a selected area that reflect the presence of

obstructions. These maps could be used for survey planning, survey verification, or overall navigation and positioning accuracy assessment in any area where LiDAR data are available over any time period for which GPS satellite position data are available. The software tools currently available that calculate DOP values for planning or verification purposes are limited in terms of incorporating obstructions, even though obstructions that block satellite signals can have a negative effect on DOP values and therefore on positioning accuracy. Other researchers have attempted to model the effect of obstructions on DOP values using LiDAR data; however, these attempts have relied on the existing software or have not gone beyond identifying visible satellites. As a result, their methods, in effect, must be applied on a point-by-point basis and have limited applicability. By creating maps of the values over a selected area, the proposed method not only shows the user how GDOP values may be affected at target locations but also allows them to quickly compare values at any number of alternate locations, which is particularly valuable in planning.

1.4.2. Proposed method

There are two major tasks to accomplish: quantifying obstructions and identifying visible satellites. Because of the large volume of LiDAR data and the large number of GDOP values calculated, additional steps were needed to facilitate the identification of obstructions and to organize the results in a coherent manner. The proposed method is depicted in Figure 1.1. The number of the chapter where each process or data set is described is given in brackets.

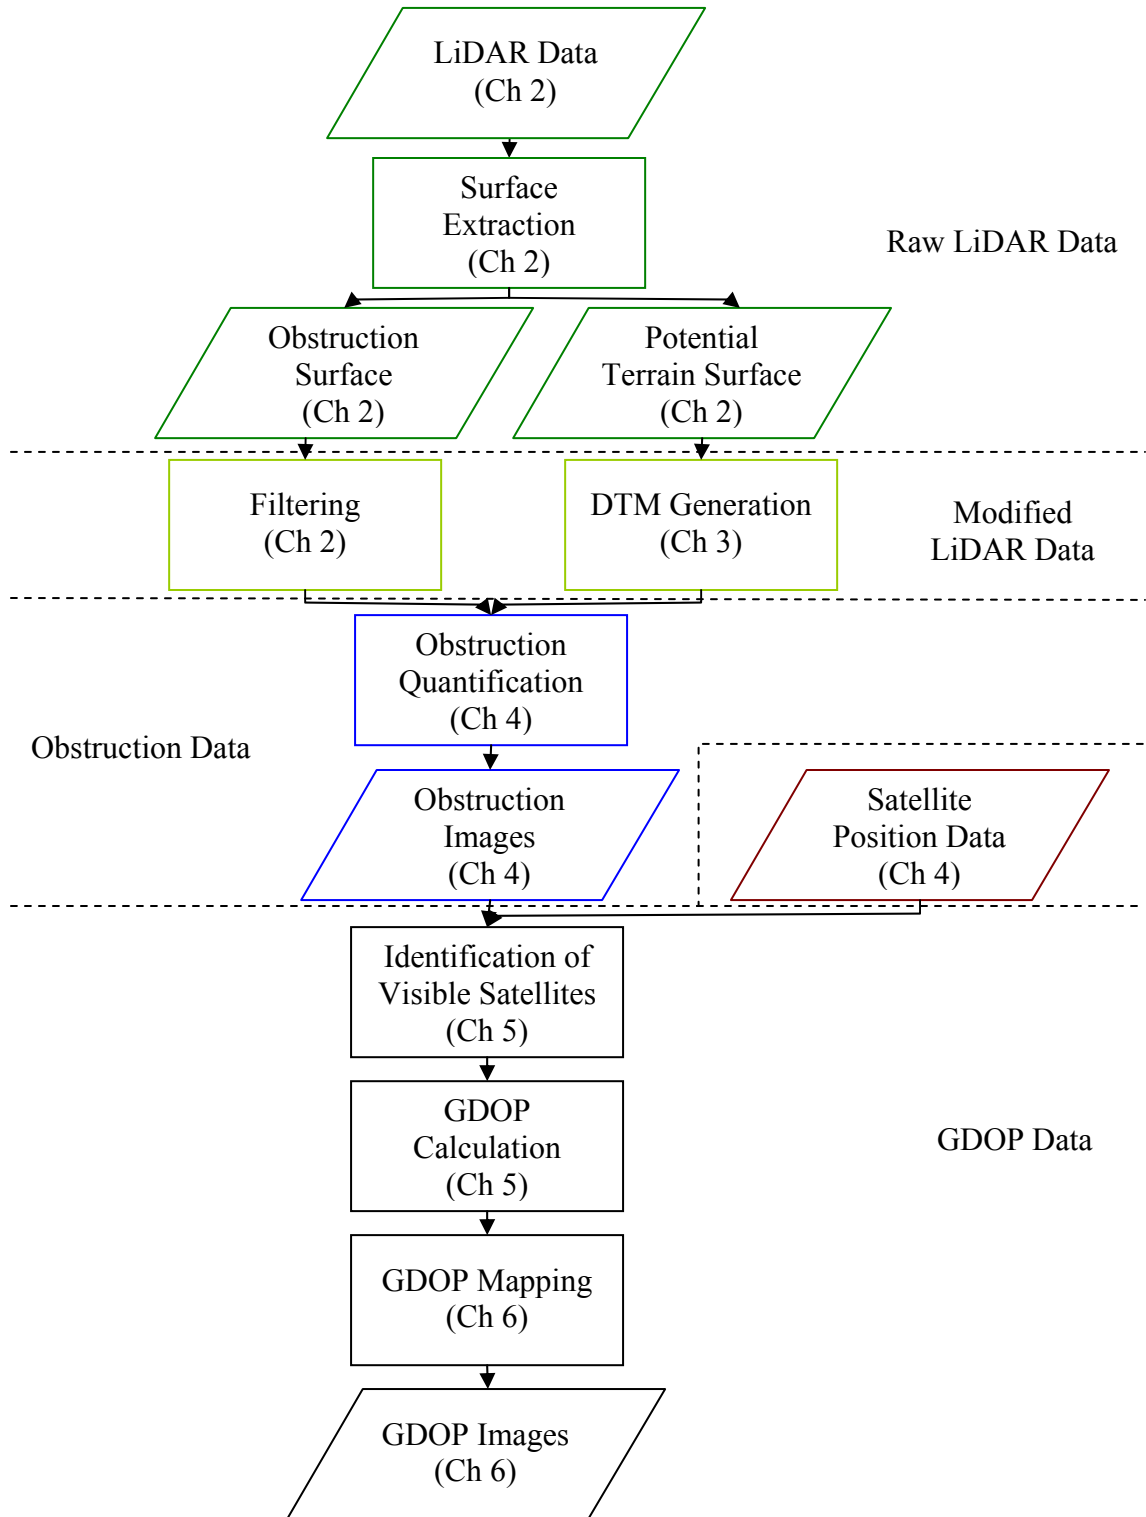


Figure 1.1
Schematic of proposed method.

Although there are various software applications available for working with LiDAR data, including ones that supposedly extract bare earth, or bald earth, points to generate DTMs and that offer various methods of interpolating the data, they are generally proprietary software applications that have a cost associated with their use. Where free versions are available the functionality is reduced, for instance, the ‘Save’ function is disabled. Furthermore, the software tested did not offer sufficient flexibility or control over the output in terms of specifying parameters, selecting data, or creating new output types. Therefore, all of the stages of the proposed method were implemented using code written in C by the author. Two applications were used for display purposes only: a free demonstration version of *LASEditTM* from Cloud Peak Software, which displays LiDAR data as individual points; and *3DEMTM* from Visualization Software LLC, which displays LiDAR data as a smooth surface and can also display data stored as a surface matrix.

1.4.3. Outline

Chapter 2 describes in depth the LiDAR data used in testing and the preliminary manipulation of the data, which includes filtering and surface extraction. Some initial filtering was required due to what was perceived as high noise in the LiDAR point cloud. To identify obstructions using the LiDAR data, two surfaces needed to be created and compared, the first an obstruction point cloud, containing all the highest object points in addition to terrain points, and the second a potential terrain cloud, containing all the points most likely to be part of the terrain.

In Chapter 3 the process by which a Digital Surface Model (DSM) was derived from this second point cloud is detailed. A DSM differs from a DTM in that it may contain non-ground points whereas a DTM contains only ground points. While ideally a DTM should be used, certain assumptions can be made that make a DSM acceptable for the purposes of this work. Any of the approaches described above could have been applied, but for the specific aims of this application, more complex methods were not warranted. Instead, a simple erosion method was developed.

Chapter 4 describes the process through which the first task was accomplished: identifying and quantifying obstructions through comparison of the obstruction point cloud and the DSM. This stage required significant processing time due to the volume of LiDAR data. Upon completion of the processing, significant obstructions for each terrain point are represented by two angles, which allows for easy comparison with satellite data.

The satellite data are introduced in Chapter 5, which describes the process through which the second task, identifying visible satellites, was accomplished, as well as the process of calculating GDOP values. Since LiDAR data and GPS data in SP3 format use different coordinate systems, some coordinate transformations were required. While using SP3 files restricts the planning capability of the proposed method and fixes the epoch interval at 15 minutes, it does not require an added step of calculating satellite positions. Certain assumptions were made to keep the number of coordinate transformations to a minimum. The satellite positions were then represented using the same two angles as for obstructions so that they could be compared to the obstruction

values to determine which satellites were not obstructed, and also which ones were above the elevation mask. Where at least four satellites were visible, the GDOP value could be calculated.

Even using the standard 15-minute interval, a large number of GDOP values are calculated for the test area for a single day. Chapter 6 presents various options developed for displaying those values, including options that organize the information in ways that might be helpful to a user trying to determine a time and location for GPS data collection.

In Chapter 7, the results of different stages of the proposed method, from the generated DSM to the effect of obstructions on satellite visibility, are assessed. Finally, in Chapter 8, conclusions and recommendations are made, including a suggestion for extending the planning capability of the proposed method beyond what is possible using SP3 files.

CHAPTER 2

LIDAR DATA MANIPULATION

A line is length without breadth.

~ Euclid

The size of a LiDAR data set depends on several factors, including flying height, pulse repetition rate, number of returns recorded, and survey area. These factors affect the overall spacing of single points, which would be expected to be regular under ideal conditions on a flat terrain, and the overall expected number of points. Regardless, there is generally a large volume of data produced during any survey. The LiDAR LAS format is a standard, public, binary file format that allows the data to be easily accessed; however, the point cloud contained in a LAS file is not necessarily the most useful format for interpretation or manipulation. Furthermore, while a large number of points provides greater terrain and object detail, not all of that detail is needed to identify obstructions.

Several pre-processing steps are required in order to extract only the most useful data. The sample data contain a small number of points with anomalous elevations. These are filtered out of the point cloud. The remaining data are then manipulated into two overlapping subsets that are needed to determine the relative elevation angles of obstructions. One subset contains all the points that are potential obstructions. This includes the highest points of non-terrain objects as well as terrain points. The other subset contains the points that are most likely to be terrain points. Since it is not possible

to determine conclusively whether a point is terrain or non-terrain solely from its coordinates and intensity value, this second subset requires more extensive manipulation, which will be discussed in Chapter 3.

2.1. Description of Sample Data

LiDAR data were collected on Julian Day 143, 2006 (May 23), over Fredericton, New Brunswick, by the Applied Geomatics Research Group (AGRG) at the Centre of Geographic Sciences (COGS) for the purpose of flood modelling and provided to the University of New Brunswick for use in other research, including this work. During the 69-minute flight, clear flying conditions were experienced and fourteen flight lines were completed; Figure 2.1 shows the flight pattern, with the University of New Brunswick, Fredericton Campus (UNBF) outlined in blue. The flying height was 1000 m. The laser was operating for approximately 30 minutes with a half scan angle of 18° and at a pulse repetition frequency of 50 kHz and a scan frequency of 33 Hz (Goulden and Hopkinson, 2006). This means that fifty thousand pulses were emitted per second and thirty-three flight lines whose end points were no more than 18° from nadir were completed per second. Up to four returns were recorded per pulse.

The data for 13 flight lines were provided in LAS file format. In total, these files contain nearly sixty-three million data points, of which 75.7% are first returns, 20.6% are second returns, 3.6% are third returns, and the remaining 0.1% are fourth returns. (62 880 499 total: 47 600 228 first; 12 928 041 second; 2 286 301 third; 65 929 fourth). A file containing only the data points for UNBF and a portion of the residential area to

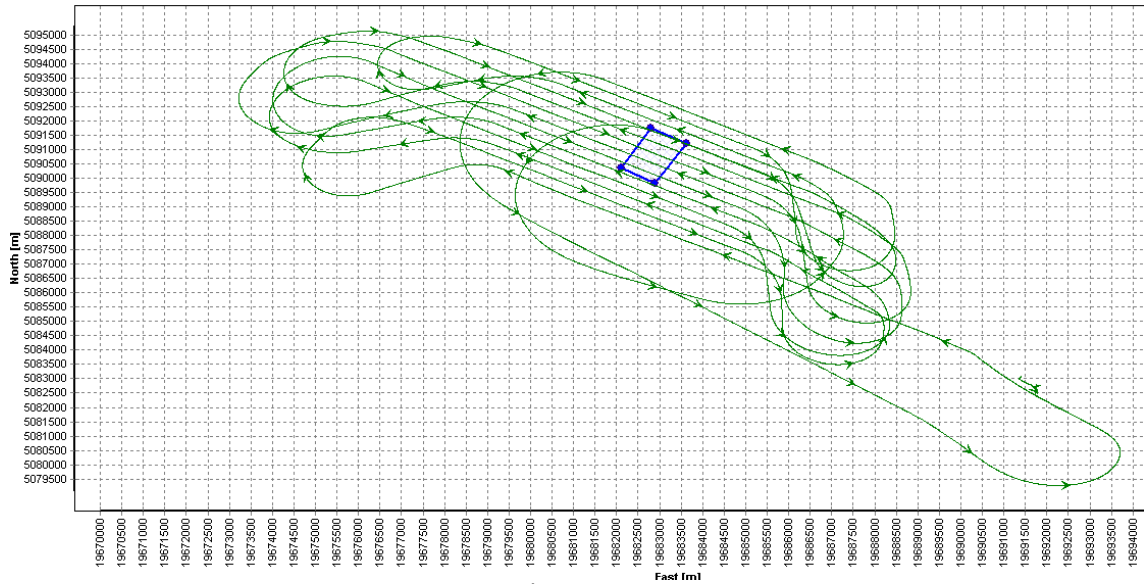


Figure 2.1
 Flight pattern for airborne laser survey over Fredericton, NB.
 (from Goulden and Hopkinson, 2006)

the north of campus, east of College Field, was also provided. The UNBF point cloud consists of over four million points, approximately 6.8% of the total data. The proportion of each return type for this subset is slightly different than for the complete data set, with more first returns and fewer second returns: 77.5% and 18.6%, respectively. The x and y coordinates are specified in the Universal Transverse Mercator (UTM) coordinate system. UNBF is in Zone 19 T, extending from 682100.00 m to 683624.03 m East and from 5089836.86 m to 5091750.32 m North. In geographic coordinates, this corresponds to 66.63° - 66.65° W and 45.94° - 45.96° N. The rectangular area specified by these values is approximately 2.92 km^2 ; however, as can be seen in both Figure 2.1 and Figure 2.2, a surface representation of the campus data set, UNBF is oriented towards the northeast and only occupies half of this rectangular area, roughly 1.51 km^2 .




Figure 2.2
Representation of campus elevation.

The point density for UNBF is approximately 2.84 points/m^2 , yet this value does not give a true indication of how accurately the point cloud represents the surface because the data points are not evenly distributed over the area. In the along-track direction of a flight strip, the point spacing is affected by the zigzag pattern that results when the laser scanner reverses direction at the edge of the strip (see Figure 2.3a). Since there is significant overlap between strips, the effect is reduced to a great extent but not eliminated (see Figure 2.3b). The spacing in both along-track and across-track directions is affected by the elevation of terrain objects, particularly at larger scan angles. Furthermore, recording up to four returns per pulse increases the number of data points and therefore the point density, despite multiple returns being clustered. While the point density can be used to help determine the optimum resolution to create a regular grid surface from the point data, which will be discussed in Chapter 3, there is not a direct relationship between the two values. For instance, if the grid resolution of the sample data were the inverse of the point density, that is $0.35 \text{ m}^2/\text{point}$, more than a third of the cells would be empty while almost a quarter of the cells would contain more than one point. Although 99% of these cells would likely contain no more than four points, some cells might contain more than a dozen points.




Figure 2.3
Along-track point spacing. a) Single flight strip. b) Overlapping flight strips.

2.2. Filtering

There is a substantial amount of overlap in the data collected in adjacent flight lines: the perpendicular distance between adjacent flight lines is approximately equal to the swath width. There is also a significant amount of “noise” in each of the flight line data files along the centre of the flight lines, which coincides with the edges of adjacent flight lines. This is illustrated in Figure 2.4: the noise in Line 12, along its centre, coincides with the edges of Lines 11 and 13 and the noise in Lines 11 and 13, along their centres, coincides with the edges of Line 12. Although the noise points most visible over the river in the illustration, they are equally present over dry land. These points have very high elevation values, around the same elevation as the flying height, and might indicate gross systematic problems with the laser scanner at nadir; however, it is beyond the scope of this research to speculate about the cause. Figure 2.5 shows a three-dimensional representation of the area in Figure 2.2 before the noise was removed from the data. The four bands of noise that can be distinguished clearly do not represent terrain or surface object elevations. They interfere with interpretation and processing of the surface and need to be filtered out of the data set.




Figure 2.4
Noise at centre of flight lines.




Figure 2.5
Three-dimensional campus representation with noise.

The exact range of elevations for the noise is unknown. An arbitrary value of 800 m was selected for the filter threshold, based on the knowledge that the actual surface points have much lower elevation values while the sampled noise points have much higher elevation values. A little more than 5000 points, only 0.12% of the total number of points for UNBF, were removed from the data set using this threshold, causing a decrease in point density of less than 0.005 points/m^2 . Table 2.1 gives the distribution of filtered points by return number. The majority of the filtered points were identified as second returns, which was unexpected since the points with the highest elevations should be first returns when multiple returns are recorded for a pulse. The data were therefore more closely examined using the following assumptions: data points are stored

in the order they are received; all returns for a single pulse are received and recorded before those of the next pulse; all returns for a single pulse have the same value in the user bit field. It was determined that the noise points with return numbers greater than one do not have corresponding points with lower return numbers; that is, a noise point identified as a second, third, or fourth return for a pulse is in fact the first return in the data set recorded for that pulse. This is a further indication of systematic problems, but again those are beyond the scope of this research.

Table 2.1
Points by return.

Return number		1	2	3	4	Total
Number of points	All	3,326,697	796,454	163,894	5,290	4,292,335
	Remaining	3,326,361	791,942	163,726	5,290	4,287,319
	Filtered	336	4,512	168	0	5,016

2.3. Potential Obstructions

An obstruction is defined as an object that blocks a portion of the sky from view at a given terrain point. The obstruction can be thought of as casting a satellite shadow: just as when one stands in a shadow cast by the sun, one cannot see the sun, a receiver placed in a satellite shadow cannot ‘see’ the satellite. Sun shadows can be observed to move and change as the sun travels through the sky; similarly, satellite shadows move and change as the satellites move through their orbits. There are generally many GPS satellites above the horizon at any given time and the effect of the changing satellite configuration will be discussed more in depth in a later chapter. What is needed in the

first processing stage is a representation of the objects that will potentially cast satellite shadows.

2.3.1. Point Cloud Representation

One option for this representation is simply a subset of the LiDAR point cloud. The points that are selected for this subset represent the highest points of non-terrain objects, such as building rooftops and treetops. As mentioned, if multiple returns are recorded for a pulse, the first return should correspond to the highest point of the object encountered by the laser, or at least it is presumably close to the highest point. For example, if the laser hits the perimeter of a building, there might be one return from the edge of the roof and one return from the ground. The roof point represents part of a higher-elevation object but will not necessarily have the same elevation as the highest point of the building. Multiple returns are also often generated when the laser hits a tree; however, although the first return is very likely to be from the upper branches, it is not likely to be from the absolute highest point of the tree.

It is more difficult to determine whether a single return represents an object or whether it represents a terrain point. A single return might be generated by a rooftop, by a road or path, by low vegetation, or occasionally even by high vegetation. Yet, since the terrain itself, which includes roads, paths, low vegetation, and bare earth, can also obstruct a portion of the sky, it is not necessary to distinguish between terrain and non-terrain for the obstruction surface. All points that are the only returns for a pulse are included in the highest point subset, along with the points that are first returns of

multiples. Figure 2.6 is a point representation of all first returns, including single returns. Elevation change is indicated by hue; magenta points have the lowest elevation, green points have mid-value elevations, and red and white points have the highest elevation. Black dots indicate spots where no data were collected; the larger black areas correspond to a pond and part of the Saint John River, where water reflected the laser energy away from the sensor.

The main advantage of using a subset of the LiDAR point cloud is that this preserves the exact coordinate values, including elevation, for each potential obstruction point. One of the disadvantages is that it is difficult to efficiently process the large number of randomly ordered points. Focusing on a smaller test area within UNBF helps to reduce processing time, but only to a certain degree since where possible the boundaries of the potential obstructions subset must extend beyond the boundaries of the test area to include high points that lie outside the test area but may cast shadows into it. Some degree of familiarity with the test area is needed to select the optimum distance by which the potential obstructions subset should be extended: in areas where there are many tall objects, such as multi-storey buildings, the distance should be greater, whereas in areas where the objects are generally low, the distance may be shorter. Another consideration in determining the extended distance is the satellite elevation mask to be used. In general, satellites that are less than 10° to 15° above the horizon are excluded from position determination because they increase the positioning error. Consequently, potential obstruction points that are below the elevation mask can be discarded. For example, if a 15° elevation mask is applied, an object at a distance of 150 m from a




Figure 2.6
First returns, point representation. a) Full campus. b) Power lines.

given ground point must be at least 40 m higher than that point for it to be considered an obstruction; if the difference in elevation is not expected to be that great then a shorter extended distance can be used. Regardless of the extent of the potential obstruction point cloud, because the obstruction points are randomly ordered but the surface points are in a regular order, it is faster to calculate obstruction angles from a potential obstruction point with respect to the surface points within range rather than from a surface point with respect to the potential obstructions. This, however, means that for every new area of interest, the entire obstruction point cloud must be processed.

2.3.2. Interpolated Surface

An alternative to using the point cloud directly is to create a surface by transferring the data into a regular grid. This allows the data to be accessed more quickly and easily; however, there are many factors that must be considered to determine the optimal resolution for the surface, including memory requirement, preservation of object details, and actual point distribution. A finer resolution grid will preserve more of the original data since fewer cells will contain more than one data point, but it will occupy more memory and will have a larger number of empty cells than a coarser resolution grid. Regardless of the resolution, some cells will contain multiple data points. The simplest way to handle this is to select the point with the highest elevation and discard the others. This approach results in a loss of information and accuracy, creating a worst-case potential obstruction surface; a different approach may be more suitable if a more moderate end result is desired.

The effect of different resolutions on the percentage of points that are discarded, the amount of empty space left in the grid, and the surface features, is illustrated using a rectangular area selected from the UNBF data within the boundaries of campus. When a 0.35 m resolution is used, 10.0% of the data points are discarded while 79.8% of the grid cells are empty. When the resolution is increased to 0.5 m, the percentage of discarded points is 19.0% and the percentage of empty cells is 63.0%. The corresponding files are approximately two-thirds and one-third the size of the original LAS file, respectively. This decrease is largely due to not having to store the x and y coordinates for each point; however, the intensity value and additional point information are also discarded.

The change in surface features can be seen in Figure 2.7. The original LiDAR point distribution is depicted in Figure 2.7a. As the grid resolution increases, building edges, such as the one circled in Figure 2.7, become less clearly defined. Yet, the changes are relatively minor: each cell represents only 0.12 m² with a 0.35 m resolution (Figure 2.7b), and only 0.25 m² with a 0.5 m resolution (Figure 2.7c). The selected points are no more than a quarter metre from the centre of the cell in the first case and only slightly more than a third of a metre in the second case, which is sufficiently accurate for locating potential obstructions. Figure 2.7d represents a 0.74 m resolution, which corresponds to the drop in point density to 1.83 points/m² that results from using only first returns to generate the obstruction surface. At this resolution, 35.5% of the data points are discarded; 35.6% of the cells are empty, 38.6% contain only one point, 18.2% contain two points, 6.2% contain three points, and 1.4% contain four or more points.

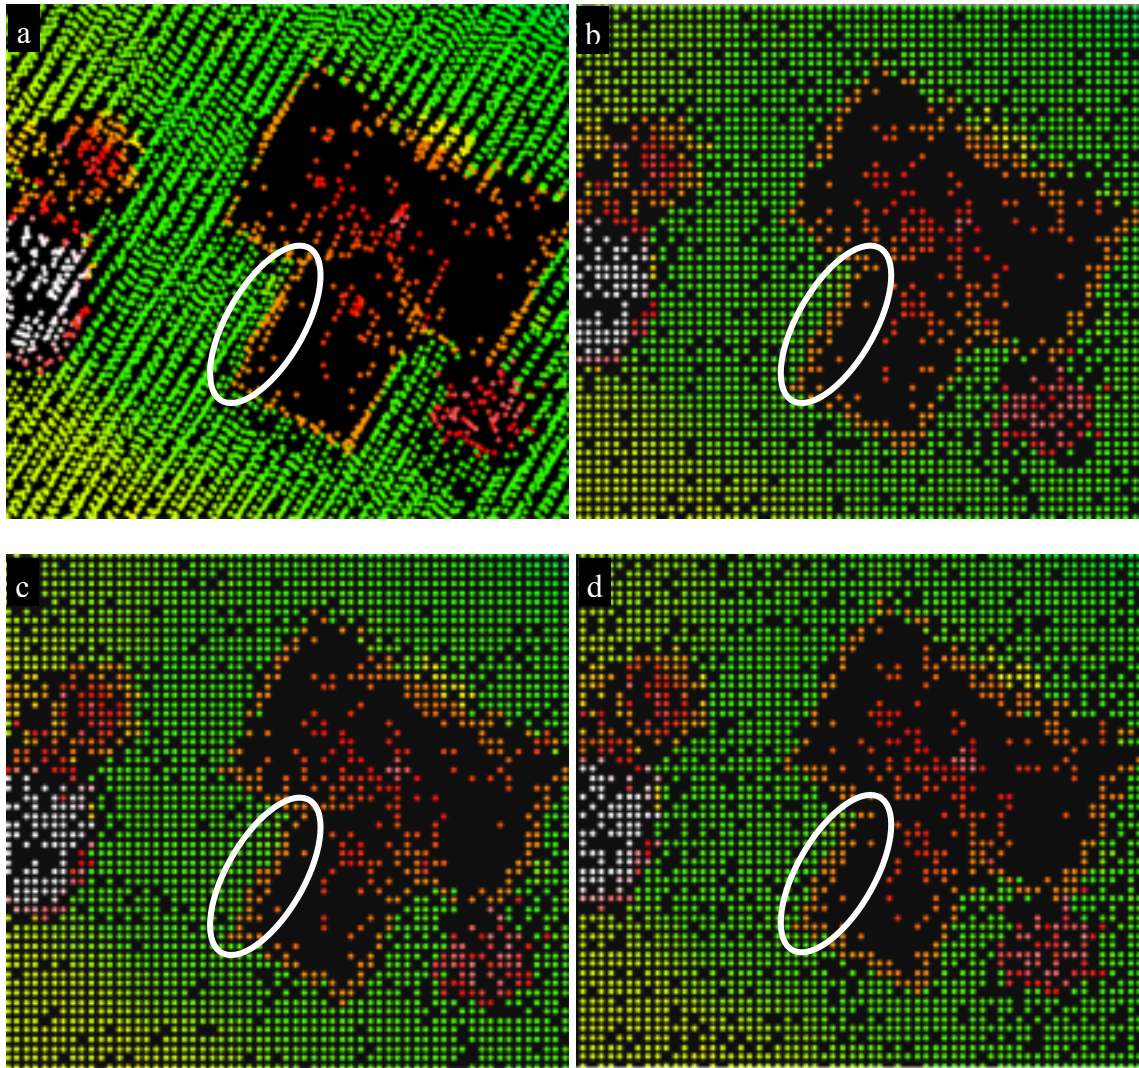


Figure 2.7
Effect of grid resolution. a) Original LiDAR data. b) Interpolated surface with 0.35 m resolution. c) Interpolated surface with 0.5 m resolution. d) Interpolated surface with 0.74 m resolution.

2.3.3. Removing Small Obstructions

There is one additional aspect to consider, which is highlighted in Figure 2.5b. The majority of first returns are either from large objects that cast significant satellite shadows or from the terrain; however, a small number are from power lines, telephone poles, and other small objects that are not likely to block satellite signals. Ideally, these points should be removed to avoid false obstruction detections, particularly if they are a significant height above the terrain, as is the case with power lines. Many of the points have two characteristics that can be used to differentiate them from other points: they are the first return of two from a pulse and they have low intensity values. Any points that satisfy either of those conditions are filtered out of the subset of highest points but since many points other than those from small objects also satisfy one condition or the other, some of the points that are filtered out should in fact remain in the subset.

Table 2.2 gives the results of filtering with a low intensity threshold set at eight; the threshold value was determined through testing. Approximately 16.6% of the total number of first returns were removed. Most of the points were removed because they were identified as the first return of two for a pulse. This condition is satisfied not only by power lines, but also at the edge of buildings and by vegetation. It is unlikely that removing the vegetation points will significantly affect the representation of potential obstructions since vegetation will generate more than two returns for some pulses and also single returns. In general, it is also unlikely that removing points from building edges will significantly affect the representation since points will remain from the rooftop; however, it is possible that buildings with roof materials that generate low

intensity returns might disappear. Nevertheless, it is better to remove all points that could cause false obstruction detections, and also ensure that when the final results are displayed, it is clear where buildings are located.

Table 2.2
Filtered first returns.

	All first returns	First of Two	Low Intensity	Both	Final number
Number of points	3,326,361	510,652	102,982	61,074	2,773,801
Percentage	100	15.4	3.1	1.8	83.4

2.4. Terrain Test Area

The second subset that is extracted from the LiDAR point cloud contains all potential terrain points. It is not possible to determine whether a point lies on the terrain or is part of a non-terrain object solely from its co-ordinates, nor even when intensity is included. When more than one return is generated from a pulse, the last return is likely to be from the terrain but may not be. When a single return is generated, the likelihood that it is a terrain point depends to a certain extent on the relative proportion of terrain areas and flat, non-terrain objects such as rooftops. Both last returns and single returns are included in the potential terrain point cloud. The method applied to transform this point cloud into a regular grid surface and remove a significant amount of the non-terrain points is discussed in the next chapter.

A large amount of data remains even after it is transformed into a regular surface. A smaller area was therefore selected as a test area. The test area is 450 m x 450 m, centred on the northernmost section of Head Hall; the corresponding UTM values are 682550 m

to 683000 m East and 5091000 m to 5091450 m North. Part of this square does not contain any data since it is beyond the edge of campus (see Figure 2.8). The size and position of the test area were determined arbitrarily, but nevertheless it contains a variety of terrain slopes and surface objects. In the upper (northern) portion, the terrain is relatively flat, as it includes a couple of parking lots and a portion of College field as well as roads and paths. The remainder of the terrain is sloped to varying degrees, with some flat areas around campus buildings.

Figures 2.8b and c are visualizations of the elevation and intensity values for all the points in the test area. The most evident non-terrain objects, as can be seen in Figure 2.8b, are trees and large buildings. There are also smaller buildings on the western side; these are somewhat obscured by surrounding vegetation but many of those vegetation points are not included in the potential terrain cloud point. Other objects include power lines (linear segments), cars (isolated bumps), and the flagpoles in front of Sir Howard Douglas Hall. Different objects are visible in Figure 2.8c due to the different reflecting and refracting properties of natural and manufactured materials. In particular, roads, pathways and grass are all considered part of the terrain since they have relatively low elevation and are indeed largely indistinguishable in terms of elevation. Asphalt, however, tends to return little of the laser energy while grass returns significantly more energy, allowing the two types of ground cover to be easily distinguished in the intensity image. While the intensity values are not used in identifying obstructions, they do provide complementary information that might be helpful in interpreting the end results.




Figure 2.8
Full campus and selected subarea. a) Elevation of full campus. b) Elevation of selected area. c) Intensity of selected area.

Once the initial LiDAR manipulation steps are completed, the large volume of LiDAR data has been reduced to a more manageable amount, from over four million points to less than six hundred thousand points. The obstruction point cloud extracted from the data does not need to be further manipulated after noise and small obstructions, such as power lines, have been removed through filtering. The data in the potential terrain point cloud, however, need to be organized in a manner that facilitates obstruction detection. This is accomplished with a DSM. The process used to create the DSM is discussed in the next chapter.

CHAPTER 3

DSM GENERATION

*When it is not necessary to change, it is
necessary to not change.*

~ Viscount Falkland

The terms digital terrain model (DTM) and digital surface model (DSM) are similar; however, a DTM can be considered to be a model of only the surface of the earth, whereas a DSM generally includes features that lie above the earth's surface, such as buildings and trees. In this paper, a DTM is defined as a representation of the bare earth with all non-terrain objects removed and a DSM is defined as a representation of the bare earth that may also contain non-terrain objects. Although DTMs of the test area can be obtained from other sources, for instance Service New Brunswick, generating a DTM from the LiDAR data might reduce errors and improve the accuracy of the final results. This is because a significant difference between the ground elevation in the DTM, regardless of the source, and the corresponding ground elevation from the obstruction point cloud will affect the calculation of obstructions: if the DTM value is much higher than the LiDAR value, the results of the obstruction calculations will be artificially low; if the DTM value is sufficiently lower than the LiDAR value, the terrain itself will be calculated as an obstruction.

The subset of potential terrain points extracted from the LiDAR data set contains non-terrain points, even though it consists of only last returns, as the laser cannot penetrate solids objects and can generally only partially penetrate vegetation. These non-

terrain points need to be removed for the subset to be used as the basis of a DTM for the test area. The random order of the LiDAR points, however, makes it difficult to differentiate terrain points from non-terrain points since it is not always trivial to identify all the nearby points that are needed to provide a reference frame. Various software tools and applications have been designed to handle LiDAR data and to extract or manipulate DTMs and DSMs, yet at this time many do not offer a 'bare earth' extraction feature that can generate a true DTM with all non-terrain objects removed. Over the past decade, researchers have developed various approaches for generating DTMs from LiDAR data, including morphology, slope-based filtering, and building footprint detections. The objective in all of these approaches is to remove all non-terrain points and only non-terrain points. Although this objective is important in many applications, certain assumptions can be made here that makes it less important for this work, and allow the use of a DSM rather than a DTM. To facilitate obstruction detection, the LiDAR data were interpolated to a grid through simple averaging. Non-terrain objects were partially or fully removed through erosion, by comparing the elevations of adjacent cells and biasing the new centre values toward the lower elevations.

3.1. Surface Considerations

It is important to have an accurate DTM so that results calculated at each stage reflect, as much as possible, what will actually be encountered in the field, whether it is the presence of obstructions or the expected GDOP values. Nevertheless, there are certain assumptions that can be made about the effect of non-terrain objects on results

that allow some flexibility, and in fact allow a DSM to be used instead of a DTM. One assumption is that for a ground surface point below a building, the sky is completely obstructed; that is, a receiver placed inside a building will not receive satellite signals. Conversely, there is a clear view of the sky from rooftops. Although this is not always true, such as in dense urban areas where taller buildings may block part of the sky view from the rooftops of lower buildings or in areas where there are mature trees that may overhang rooftops, for development purposes it was assumed to be a generally valid statement. These two assumptions indicate the expected results when buildings are removed from the DSM and when they remain in the DSM; the question of whether or not to remove them is therefore largely a question of convenience. The decision does affect how the results should be interpreted and there are other considerations, such as whether any buildings were fully removed from the obstruction point cloud due to low intensity returns from the roof, whether the presence or absence of buildings will impact processing time. Regardless, unless building rooftops are considered accessible during a survey and the final results must include accurate values for rooftops, the two options are essentially equivalent as otherwise the values generated for rooftop points irrelevant.

Another assumption is that a GPS receiver would not be placed under dense foliage. It is possible that satellite signals could penetrate the tree canopy, just as rays from the sun are sometimes able to penetrate the tree canopy, but it is not certain; intuitively, the ideal placement for a receiver is in a location where one can see the sky. This assumption suggests that vegetation could be considered similar to buildings, in that it is not necessary to remove it from the DSM; however, since a receiver would not be placed

in the treetops, it makes more sense to try to remove vegetation as much as possible. An additional consideration with vegetation is the seasonal variation in the degree of obstruction it creates. Conceivably, locations that are almost completely obstructed during the summer could be feasible receiver locations during the winter, when all the leaves have fallen. Determining whether a sufficient number of satellites are visible in such instances is a far more complex problem and would likely need to be dealt with on a case-by-case basis.

A final assumption is that certain objects are small enough that they do not adversely affect results if they are not removed from the ground surface. For instance, a GDOP value calculated for a location in a parking lot will not be appreciably different if the ground point is actually on top of a car rather than on asphalt. This assumption greatly simplifies the process of generating a DSM since it could be difficult to try to remove these small objects without also removing small terrain variations.

3.2. Interpolation to Grid

Figure 3.1a is a representation of the potential terrain point cloud, which consists of all last returns, including single returns, extracted from the LiDAR data. The point density is nominally 2.8 pts/m² for the entire data set; this drops to 2.21 pts/m² for potential terrain points, or one point for every 0.45 m². While this might seem to be sufficiently dense so that interpolation is not needed, because the surface is randomly sampled there are likely gaps in the coverage that must be filled. In addition,

interpolation organizes the data into a regular grid, which makes access and manipulation faster and easier.




Figure 3.1

a) Last returns in selected test area. b) Terrain visualization of last returns averaged over a 1 x 1 m grid.

The grid resolution should be sufficiently fine that little detail is lost due to multiple points falling within the same grid cell, yet sufficiently coarse that there are not a significant number of empty cells. Using a grid with cell area equal to the inverse of the point density might seem like a reasonable choice but a 0.67 m resolution was found to result in more than a third of the cells in the test area being empty. Doubling the cell area gives a resolution of just under one metre; since one metre is a more convenient quantity than 0.95 m, it was used instead. With a 1 m by 1 m grid, slightly more than two points are expected to fall within each grid cell. There will therefore be some loss in accuracy, but there will also be fewer empty cells. Furthermore, the larger cell size is more realistic for the purpose of determining the optimum location at which to place a receiver.

There are different ways to determine whether a point falls within a cell and to calculate the elevation of cell. The simplest approach is to have the cells be distinct and non-overlapping, delineated by the cell bounds, and to take the average of the elevations of all the extracted LiDAR points that fall within each grid cell. With this approach, almost 10% of cells in the test area are empty. Approximately half of these empty cells are in the no-data area, seen as a black triangle in the upper-left corner of Figure 3.2; the other half are gaps in the DSM that will have to be filled, seen as black points in Figure 3.2.




Figure 3.2
Empty cells when a) cell boundaries used b) circle radius $1/\sqrt{2}$ m used in terrain interpolation.

It is possible to reduce the number of empty cells by expanding the area of consideration for each cell, for instance by including all points that fall within a certain radius of the centre of the cell. Figure 3.2b shows the location of empty cells when a

radius of $1/\sqrt{2}$ m is used, which is the smallest radius that contains the entire cell, as compared to Figure 3.2a for which the cell boundaries were used. The number of empty cells is greatly reduced using this expansion, to 4.7%, but there is also a significant reduction in the number of cells containing only one point, from 24% to 6.75%. While this might seem to be an undesirable effect, a single point is not likely to lie at the centre of a cell and so does not necessarily better represent the cell elevation than does an average of multiple points. In addition, a circular area of consideration is less dependent on the grid alignment than a square area of consideration, where points that are between 0.5 m and 0.71 m may be included or excluded depending on where they lie in relation to the grid.

Elevations need to be calculated for any remaining empty grid cells. This was accomplished by averaging the elevations of the adjacent eight grid cells, on the condition that at least half of those cells contained data points so that cells that lay in the no-data area were not assigned false elevations. If less than four adjacent cells contained elevation data, a 'no-data' value was assigned to the empty grid cell.

Regardless of which approach is used to delineate cell data, a point that is included in the elevation calculation for an empty cell could lie over two metres from the centre of that cell. The important difference between the two options is that there are far fewer gaps to fill when the area of consideration is slightly expanded; presumably, the overall effect of using a larger area of consideration is less significant than the overall effect of interpolating elevations for a larger number of empty cells. The 'no-data' cells that remained after filling are in the upper-left, or north-west, corner of the test area, as can

be seen in the terrain visualization depicted in Figure 3.1b, which was generated using a one metre resolution and a circular area of consideration with a $1/\sqrt{2}$ m radius. This ‘no-data’ region lies outside UNBF boundaries. Large or deep bodies of water would also create ‘no-data’ regions due to the low scattering coefficient of water but are not present in the test area.

In the gridded surface generated as described above, the z_0 value of a triplet (x_0 , y_0 , z_0) represents the elevation for a 1 m^2 surface area delimited by the lines x_0 to the west, y_0 to the south, $x_0 + 1$ to the east, and $y_0 + 1$ to the north. The elevation is approximate and, due to the use of averaging, may or may not be a good representation of the terrain elevation for that surface area. The values calculated for cells that contained multiple data points with similar elevations are likely to most closely resemble the actual terrain elevation, as are the values for empty cells whose adjacent cells all had similar elevations. Where the data points or adjacent cells had significantly different elevations - as would occur around building edges, near tall vegetation, and on steep terrain - the calculated value is not likely to represent an actual terrain elevation. It is these values, in particular, that should be removed from the DSM to ensure that it best represents the actual terrain.

3.3. Removing Non-surface Objects

Extracting a DTM from LiDAR data is not a trivial exercise. Due to the considerations outlined above, however, it is not necessary in this case to remove all objects completely – a reduced DSM is sufficient. The approach developed uses the

values in a window around each grid cell to slowly erode large surface objects while leaving the terrain unchanged. The new elevation for each grid cell was determined through three steps. First, the slope was calculated in the horizontal, vertical, and diagonal directions using the eight adjacent cell values and the average was taken of the lower values of any directional pairs whose slope exceeded a set threshold. For example, if the horizontal slope and one of the diagonal slopes were greater than the threshold, with the value in the cell on the right less than the value on the left and the value in the upper right cell less than the value in the lower left cell, then the new value for the centre cell would be the average of the right and upper right cells. Second, the average was taken of all eight adjacent cells for cells for which none of the slopes of the directional pairs exceed the threshold. If the difference between this average and the value of the centre cell exceeded a set threshold, the average became the new elevation. Third, the original elevation was restored if the new elevation was more than a set distance above the original.

A small number of iterations are sufficient to remove, or at least significantly reduce the height of trees; the number of iterations required to remove buildings depends on the size of the structures. To avoid false erosions at the edge of no-data regions, the calculations are only performed with cells that contain elevation values. The second step smoothes anomalies in the terrain that only occupy one grid cell and are therefore missed by the directional pairs, while the third step prevents grid cell elevations from being substantially increased. There might be cases when the original elevation value for the grid cell is below the terrain, if, for instance, there happened to be an open hole at the

time of data collection that would normally be filled. Nevertheless, this limitation on increasing elevations is necessary because it is more likely that the irregular shape of vegetation caused a value higher than the terrain to be calculated.

All of the thresholds are arbitrarily set. The first threshold is the square root of the iteration number, up to a maximum value of three: during the first iteration it has a value of one, during the second iteration it has a value of $\sqrt{2}$, and during the ninth iteration and any subsequent iterations it has a value of three. This varying threshold seems to remove buildings and vegetation without causing small dips in the terrain to grow, as would occur with a constant threshold due to the bias towards lower elevations. The second threshold is two metres, and the third threshold is one metre. Figure 3.3 shows (a) the resulting surface after each pass and (b) the difference between the resulting surface and the previous surface for five iterations (the first iteration is the bottom image). Green areas in Figure 3.3b indicate areas where there was a small decrease in elevation (~ 1 m), yellow areas indicate mid-range decreases in elevation, and red areas indicate large decreases in elevation (~ 15 m). There were no increases in elevation, or only very small increases.

This process does not fully remove large surface objects such as buildings and trees and does not remove small objects such as cars. However, the results are superior to those that were obtained when two software applications with terrain manipulation tools were tested. As can be seen by comparing Figures 3.4a and 3.4b, the sections of vegetation that remain after three iterations of the process described above have been reduced in height, placing them below the obstruction surface. Building footprints are




Figure 3.3
 a) Surface elevations after each of the first five iterations of erosion. b) Elevation differences between consecutive iterations.




Figure 3.4
 Enlarged section of the test area: a) original interpolated surface; b) surface after third erosion iteration. Highlighted areas: (i) cars; (ii) building; (iii) vegetation.



Figure 3.5
 Difference between original interpolated surface and surface after first erosion iteration.

smaller, and five iterations are sufficient to reduce the Lady Beaverbrook residence to little more than a skeleton, but rooftop elevations at the centre of the buildings have not changed. This is not a concern, again because of the considerations described above, in particular because the rooftops are not considered accessible and can therefore be ignored. Nevertheless, it is important to be aware of how the buildings have been altered when interpreting later results. Figure 3.5 is a monochromatic version of the difference image after the first iteration, the bottom image in Figure 3.3b, with black indicating simply where an elevation value changed. It shows the outlines of buildings and the areas covered by vegetation. Overlaying this image on later results could help with visual interpretation since these object features may not be easily distinguishable in later stages; obtaining this image is one of the key advantages of this proposed method.

A final step that was considered was applying a smoothing function to the results of the final iteration. One can see in Figure 3.4b that the terrain has a rough texture both under vegetation and around buildings where the differences between direction pairs were not large enough to cause any further erosion. While smoothing would visually improve the DSM, it likely would not appreciably improve later results. Furthermore, it is possible that smoothing would increase the differences between the DSM and the obstruction surface around small objects such as cars, thereby exaggerating the importance of those small objects.

The DSM used as the basis for locating obstructions was generated after several assumptions were made about the importance of different types of objects. The LiDAR data were interpolated to a regular grid and a simple erosion operator was developed that

determines new terrain elevations from the lowest elevations of the adjacent grid cells. After a number of iterations, vegetation and small buildings have been almost entirely removed and large building footprints have been reduced. The next step is to compare the DSM to the obstruction point cloud to locate and quantify obstructions.

CHAPTER 4

IDENTIFYING OBSTRUCTIONS

*Character is like a tree and reputation like
its shadow. The shadow is what we think of
it; the tree is the real thing.*

~ Abraham Lincoln

In the preceding chapters, two data sets were extracted from the LiDAR data, the first an obstruction point cloud and the second a DSM. Both are representations of surfaces in a three-dimensional space, as can be seen in Figure 4.1. The models were rotated to better show the increase in terrain elevation from north to south of the test area; the origin is at the south-west corner, which is the upper right corner in Figure 4.1 and the lower left corner in all preceding and subsequent figures. These surfaces must be compared to determine at which terrain points an object blocks part, or all, of the sky from view and is therefore an obstruction. Quantification of the obstruction should be relative, not dependent on either the distance to the obstruction or its height. To achieve this, each object was reduced to an elevation and a direction angle with respect to a terrain point. The plane was divided into a specified number of direction sectors and only the greatest elevation angle was recorded for each sector. The resulting outcome is a set of images, each one corresponding to a direction sector, with the physical elevation of each grid cell replaced by the elevation angle of the most significant obstruction for the cell in that direction sector.

Although the number of points in the obstruction point cloud was reduced through filtering, as described in Sections 2.2 and 2.3, the number of obstruction points that need to be compared to each terrain point can be further reduced by only considering those points within a specified radius of that point. A smaller radius will include fewer obstruction points, thereby necessitating fewer calculations, but it may exclude large obstructions that cast shadows over a terrain point despite their distance. In the selected area, the terrain has an elevation range of approximately 37 m; that is, the lowest terrain point is 37 m lower than the highest terrain point. Non-terrain objects rise an additional 15 m above the highest terrain elevation. The obstruction radius used was 150 m. At this distance, an object must have a relative elevation of at least 40 m to be considered an obstruction when a 15° GPS mask angle is applied. If a smaller GPS mask angle were used, a larger radius might be more appropriate: Head Hall, in particular, casts a long shadow over lower-lying terrain to the north.




Figure 4.1

a) Three-dimensional representation of obstruction point cloud. b) Three-dimensional representation of DSM.

4.1. Elevation Angle

An object is an obstruction if its elevation is sufficiently greater than the surface elevation of a terrain point that it casts a shadow over that point, but merely comparing elevations is not enough to identify the presence of an obstruction: the height, width, and distance of an object all affect the degree to which that object is an obstruction. The terrain itself can be an obstruction but a mid-range slope inclination may only block a small portion of the sky for downhill terrain points while a large building may block fully half the sky from view. Since buildings are likely to be the primary obstructions, where a building is defined as a structure large enough for a person to enter, they can be used to illustrate different degrees of sky view obstruction. Around the perimeter of a building, regardless of its size, a significant portion of the sky is blocked. As one moves away from the building, as in Figure 4.2b as compared to Figure 4.2a, more of the sky becomes visible. With taller buildings, it is necessary to move further away to get out of their shadows, but there may be shadows from other objects, as in Figure 4.2c. The elevation angle incorporates both height and distance of an obstruction into one value that indicates how much of the sky is visible in a general sense, thereby allowing comparison of any obstructions around a terrain point.




Figure 4.2
Effect of building height and proximity on sky view.

The symbol θ is used to represent the elevation angle of an object point with respect to a terrain point. It is measured up from a plane parallel to the xy plane at height z_0 , which is the height of the terrain point (Figure 4.3a). Most values of θ will be positive, though negative values are possible, such as when the object point lies down-slope from the terrain point. For a point (x, y, z) in the obstruction point cloud, the elevation angle relative to a DSM point (x_0, y_0, z_0) can be calculated using the following equation:

$$\theta = \tan^{-1}\left(\frac{z - z_0}{\sqrt{(x - x_0)^2 + (y - y_0)^2}}\right). \quad (4.1)$$

If the elevation angles are calculated for all points in the obstruction cloud that lie within a selected radius of a typical surface point, there could be values of θ differing by as little as a fraction of a degree, or by more than 90 degrees. Regardless, determining overall how much, or how little, of the sky is visible would be a simple matter of selecting the maximum value of θ . This would give an accurate measure of the highest obstructions but might require a significant amount of processing, particularly for larger terrain areas or higher density obstruction point clouds.



Figure 4.3

a) Elevation angle θ and b) direction angle ψ of a potential obstruction relative to a surface point (x_0, y_0, z_0)

An alternate approach is to use obstruction levels, each perhaps five or ten degrees wide. This could require fewer calculations, but some accuracy would be lost. Rearranging Eq. 4.1 yields the equation

$$z_m = z_0 + \tan \theta \sqrt{(x - x_0)^2 + (y - y_0)^2}, \quad (4.2)$$

where z_m is the elevation in metres of a point with same x and y coordinates as the obstruction point (x, y, z) and an elevation angle θ with respect to the terrain point (x_0, y_0, z_0) . Geometrically, for a fixed θ and (x_0, y_0, z_0) , the set of points that satisfy Eq. (4.2) form a cone; this cone can be thought of as a boundary between the obstructed and unobstructed sky for that value of θ . The value z_m is the minimum absolute elevation at which a LiDAR point at (x, y) is considered to be an additional obstruction for the specified terrain point. It is therefore sufficient to compare the actual elevation of the LiDAR point to this minimum elevation. Using obstruction levels rather than exact elevation angles could decrease processing time, since once it is determined that there is an obstruction point at the highest level for a particular terrain point, no further calculations are needed for that terrain point. However, if the obstruction level divisions are too fine, this approach could increase processing time.

4.2. Direction Angle

It is not enough to identify obstructions solely by elevation angle. In Figures 4.2 (a) and (c), the elevation angle for the building to the left of the terrain point is the same. If this angle is the only value used to determine how much of the sky is visible in Figure 4.2a, since there are no obstructions in any other direction, a large portion of visible sky

would be considered obstructed; in Figure 4.2c, there is a larger obstruction to the right of the terrain point, which blocks an additional portion of the sky in that direction. The inadequacy of one angle is further illustrated in Figure 4.4. If the terrain point is located directly beside a building, as in Figures 4.4 (a) and (b), as much as three-quarters of the sky might be visible, or as little as half. As one moves away from the building (Figure 4.4c), or other obstruction, more of the sky becomes visible but there are also connected obstruction zones to consider (Figure 4.4d). Consequently, a direction angle needs to be calculated to relate the terrain point and the obstruction point in the horizontal plane.




Figure 4.4
Effect of building location on sky view.

As shown in Figure 4.3b, the direction angle, denoted ψ , is measured from the positive x axis (east) in the xy plane, towards the obstruction point. It is calculated from the x and y coordinates of the terrain point (x_0, y_0) and the obstruction point (x, y) :

$$\psi = \tan^{-1}((y - y_0)/(x - x_0)). \quad (4.3)$$

It is unlikely that two obstruction points will have exactly the same direction angle, which means that direction angles, in themselves, are not useful for comparing the

location of obstructions. Comparisons can be made, however, by dividing the sky into a number of sectors of equal width, starting at 0° . For example, if there are four sectors, each one is 90° wide; the first sector corresponds to the northeast quadrant, 0° to 90° , the second sector corresponds to the north-west quadrant, 90° to 180° , the third sector corresponds to the south-west quadrant, 180° to 270° , and the fourth sector corresponds to the southeast quadrant, 270° to 360° . The highest elevation angle of all the obstruction points that lie within a sector determines the obstruction level for that sector.

For every terrain point, calculations were performed in turn for each obstruction point within a specified radius. Initially, the obstruction level values were set to zero for each sector. As higher obstructions were identified in different sectors, the obstruction level values were incremented accordingly. This was accomplished by first calculating the direction angle and determining the direction sector for an obstruction point and then comparing its elevation to z_m for that sector. The value of z_m was calculated using Eq. (4.2) with the value of θ corresponding to one obstruction level higher than that already determined for the appropriate direction sector. For instance, if an obstruction point lay in Sector Three and previous calculations for Sector Three determined that the obstruction level was five, this new point was compared to the z_m value for obstruction level six.

As with the number of obstruction levels, the number of sectors has an effect on processing time but to a lesser degree. Having fewer sectors results in a faster processing time but is less accurate; having more sectors only slightly increases the processing time and gives a more accurate picture of where obstructions are located with respect to the

terrain points. Processing time is more significantly affected by the number of obstruction points under consideration. One of the limitations of working with a point cloud is that the points are not associated with objects and therefore each must be treated separately, whereas if a group of points were known to belong to the same obstruction, the outermost points would be sufficient to determine the effect of that obstruction on nearby terrain points and the processing time would be reduced.

4.3. Generating Obstruction Images

When all the calculations are complete, the resulting output is a set of images that represent the obstructions in each sector. The shadow cast by an obstruction falls in the opposite direction from its direction angle; for example, if the direction angle for an obstruction is 65° , its shadow falls along 245° . Figure 4.5a displays the obstructions in the test area, as determined using four direction sectors and six obstruction levels; Figure 4.5b displays the same, as determined using eighteen levels. The top image in both sets shows the effect of obstructions lying to the southeast: shadows are present on the northwest side of objects and decrease in significance along that direction. The dark green areas represent terrain points where there are no obstructions, or only obstructions below 15° in (a) and below 5° in (b). The dark red areas represent terrain points that are completely obstructed, having obstructions above 75° in (a) and above 85° in (b). There are more level distinctions in Figure 4.5b, but otherwise the image sets are the same. An additional two-thirds of an hour was required to complete the calculations using eighteen




Figure 4.5
 Relative elevation angle of obstructions by quadrant: a) six levels b) eighteen levels.

levels, which is a 12.5% increase in processing time. The consequences of not fully removing buildings can be seen in all of the images: there is a red outline around all of the large buildings, in addition to their shadows, where the elevations were successfully reduced and green areas within the outlines that correspond to the rooftop elevations that remained in the DSM. In reality, satellite visibility will be similar for all areas of a rooftop, and this is what is represented by the green areas; the red areas represent the terrain below the rooftop, which is completely obstructed by the building. It should also be noted that the no-data area in a sense casts its own shadow over terrain points within 150 m, since there may be large, unseen objects in that area that cannot be identified as obstructions. To avoid this problem in general, the boundaries of the obstruction point cloud were extended 150 m beyond the test area.

For comparison, Figure 4.6 shows the results achieved using different parameters or obstruction data. The image set in Figure 4.6a is the output generated using four sectors and six obstruction levels, and also the interpolated obstruction surface rather than the obstruction point cloud. There are mostly only slight differences between this image set and those in Figure 4.5. The most significant differences result from the fact that power lines were not removed through filtering before the obstruction surface was generated.

The image set in Figure 4.6b was generated using eight sectors and eighteen levels; the processing time was only slightly longer for (b) than for (a), which indicates that there is no advantage, time-wise, in using an interpolated obstruction surface. The difference in processing time, for eighteen obstruction levels, between eight sectors and four sectors was approximately ten minutes, or a 2.8% increase. The decision on the




Figure 4.6
a) Obstructions from interpolated surface, 6 levels and 4 sectors. b) Obstructions from point cloud, 18 levels and 8 sectors.

number of direction sectors, therefore, mostly depends on how much accuracy is desired in the horizontal plane, while the decision on the number of obstruction levels must factor in time delay as well as vertical accuracy.

The proposed method for describing obstructions uses two angles, elevation and direction, relative to each terrain point. The range of elevation angles were divided into levels and the range of direction angles were divided into sectors to allow obstructions to be more quickly and easily compared; increasing the number of levels or sectors increases the amount of detail stored regarding the most significant obstructions but also increases processing time. Once the obstructions have been quantified, the LiDAR data are no longer needed: GPS satellite data can be compared to the obstruction images to determine which satellites are visible. The approach used to identify satellites is described in the next chapter, as well as the process for calculating GDOP values.

CHAPTER 5

CALCULATING GDOP VALUES

*To one who has been long in city pent,
'Tis very sweet to look into the fair
And open face of heaven.*

~ John Keats

Once obstructions have been identified, and their effect on sky visibility for the surrounding terrain has been mapped, the LiDAR data are no longer needed. The obstruction images generated after comparing the obstruction point cloud with the DSM were used along with satellite position data to determine which satellites are visible from a given terrain point at a particular epoch. Since the satellite position data, available from IGS, use a different coordinate system from that used by the LiDAR data, and therefore in the obstruction images, some coordinate conversions first needed to be applied. After the visible satellites were identified for each terrain point, the GDOP values were calculated to give an overall idea of satellite configuration strength.

5.1. Satellite Data

Satellite position data can be obtained from IGS [Beutler et al., 2005]. The IGS archives contain files dating back to GPS week 723, which was mid-November, 1993. Final orbits are released approximately two weeks after the observation day and have an accuracy better than 5 cm. Each file contains the observed satellite positions and clock times for one twenty-four hour period at fifteen-minute intervals. Rapid and Ultra-Rapid

orbits are also available for applications where more immediate position data are required. Rapid orbits are updated daily and have a latency of 17 hours. The Ultra-Rapid files contain satellite orbit data for two days, the first day observed and the second day predicted. They are updated four times daily, giving a three-hour latency for the observed half and real-time data for the predicted half. Predicted orbits have an accuracy of approximately 10 cm.

All of these files follow a standard format. The file header includes the start date and time in both international date format and GPS week. It also lists the number of epochs in the file, the length of an epoch in seconds, the number of satellites observed, and other relevant details. Satellite positions are given in x, y, z coordinates in the IGS05 coordinate system, the most recent IGS implementation of the International Terrestrial Reference Frame (ITRF), which is an Earth-Centred, Earth-Fixed (ECEF) reference frame. The coordinates in the file are divided by epoch and sorted by satellite identification number.

The data selected for testing are the final calculated orbits from 1 October 2006 (GPS week 1395, day 0). There were twenty-nine observable satellites on that day.

5.2. Converting Coordinates

The two data sets are in different coordinate systems, the LiDAR data in UTM coordinates and the satellite data in IGS05; therefore, a coordinate conversion needed to be performed before proceeding with the next stage. Transforming either all of the satellite coordinates to UTM or all of the surface coordinates to IGS05, however, would

have been time-consuming, and to a certain degree unnecessary. GPS satellites have a nominal orbit of 20 000 km; in comparison, the Easting, Northing and elevation differences in the test area are negligible. For larger terrain areas, or areas where there is a much greater difference in terrain elevations, that might not be the case. Nevertheless, in general it can be assumed that the differences are not large enough to affect the final GDOP values, particularly if these are only calculated to one or two decimal places. Consequently, it was decided that only five terrain points needed to be converted: the centre point of the test area and the midpoints of each of the four sides. The centre point was used as a representative point from which to determine satellite ranges, while the midpoints were used to construct a plane in the IGS05 coordinate space from which to measure satellite elevation.

In each case, the conversion was applied in two steps, first from UTM to geodetic coordinates (latitude and longitude) and then to IGS05. The formulae for the transformations are included in Appendix I. Both the UTM and the geodetic coordinate systems have only two principal coordinates, whereas IGS05 has three. To complete the conversion to IGS05, an elevation value was needed. While the actual elevations of the points could have been used, doing so would have skewed the plane constructed from the four midpoints according to the slope of the terrain. In the case of the centre point, a test of the effect of changing its elevation by 30 m resulted in a difference in GDOP of less than 10^{-4} points, which is not significant enough to warrant finding the actual elevation for the terrain cell closest to the centre of the area. Furthermore, by using 0 m for the elevation, the points are placed on the surface of the ellipsoid that represents the

Earth's shape, which is a reasonable simplification for the test area and for GDOP calculations.

From the midpoint coordinates two vectors were constructed. These were intended as direction vectors only and were therefore normalized to have unit length; one vector corresponds to the positive x axis in the original UTM coordinates, E , and the other corresponding to the positive y axis, N . Since only the centre point was used for checking satellite elevations, the direction vectors are not in themselves useful: a third vector, n , corresponding to the positive z axis, was obtained by taking their cross product, $\vec{n} = \vec{E} \times \vec{N}$ (see Figure 5.1). This third vector is the normal to the plane EN and therefore perpendicular to the horizon, pointing away from the centre of the Earth. It was used to determine the elevation angle of the satellites, as described in the next section.




Figure 5.1

a) Direction vectors in UTM coordinate system. b) Representation of test area in ECEF coordinate system.

5.3. Identifying Visible Satellites

Intuitively, if the terrain is simplified to a plane in the UTM coordinate system, as it is with the vectors E and N , a satellite can be assumed to be visible if its elevation is greater than zero. When the plane is translated into the IGS05 system, it is more

complicated to determine whether or not a satellite is visible, and the introduction of potential obstructions adds the necessity of finding the satellite direction angle as well as its elevation angle. For that purpose, a vector d_{sat} was constructed from the centre point of the selected area to the satellite under consideration (see Figure 5.2a). The elevation angle of the satellite, θ , was then calculated using the equation

$$\sin \theta = \cos(90 - \theta) = (\vec{n} \cdot \vec{d}_{sat}) / (\|\vec{n}\| \|\vec{d}_{sat}\|), \quad (5.1)$$

where $\|x\|$ represents the length of vector x ; since the vector n has unit length, $\|n\|$ could be removed from the equation. Eq. (5.1) is from the definition of the scalar product, but since the desired angle is the complement of the angle between the two vectors, it is modified with the identity $\cos(90 - \theta) = \sin(\theta)$.




Figure 5.2
Satellite a) elevation and b) direction angles.

Determining the direction of the satellite required several calculations because it needed to be measured with respect to the positive x axis in the DSM, that is with respect to E . To reduce the overall number of calculations, only satellites above the horizon, or rather above the elevation mask, were considered. Typically, when the elevation mask is 15° , there are seven or eight satellites visible but there may be as many as a dozen. In order to calculate the GDOP value, there need to be at least four satellites

visible. For each satellite that was at a sufficiently high elevation, its direction was calculated as follows. To start, d_{sat} was projected onto the EN plane by applying cross products. The result of a cross product of two vectors is a third vector that is perpendicular to both; since n is the normal vector for the plane EN , the result of a cross product involving n must lie on EN . A second cross product was needed so that $d_{sat,proj}$ was in the same direction as d_{sat} rather than perpendicular to it. The resulting equation is

$$\vec{d}_{sat,proj} = proj(\vec{d}_{sat}, EN) = \vec{n} \times (\vec{d}_{sat} \times \vec{n}). \quad (5.2)$$

Next, the scalar projections of $d_{sat,proj}$ along E and N were determined. Since E and N are unit vectors, the scalar projections are equal to the scalar products:

$$d_{sat,E} = proj_{\vec{E}} \vec{d}_{sat,proj} = \vec{E} \cdot \vec{d}_{sat,proj} \quad \text{and} \quad d_{sat,N} = proj_{\vec{N}} \vec{d}_{sat,proj} = \vec{N} \cdot \vec{d}_{sat,proj}. \quad (5.3)$$

Finally, the direction angle ψ was calculated as:

$$\psi = \tan^{-1}(d_{sat,N} / d_{sat,E}). \quad (5.4)$$

Although this angle could be calculated directly from the projection using the same manipulation of the scalar product formula as in Eq. 5.1, that would not give sufficient direction information because of the nature of the cosine and arccosine functions. A similar issue arises with tangent and arctangent; however, the atan2 function uses the signs of the two components to determine the correct quadrant and return a value in the range $[-\pi, \pi]$.

Once both the elevation and direction angles were known for all the satellites that were above the elevation mask, they could be used along with the obstruction images to determine which satellites were visible at each terrain point. A satellite's direction sector

was determined from its direction angle. If its elevation angle was greater than the upper bound of the obstruction level of that direction sector for a given terrain point, the satellite was deemed to be visible at that point. For example, if for a particular terrain point the obstruction level for Sector Two was four out of six levels, there was an obstruction between 60° and 75° and therefore any satellites in Sector Two had to have an elevation angle greater than 75° to be visible.

The comparison between obstruction levels and satellite elevations was repeated for each terrain point for every satellite above the elevation mask at each epoch. The output from this stage could therefore require a significant amount of memory, even if all that is recorded for every terrain point at every epoch is the identification numbers of the visible satellites. This memory requirement can be reduced by immediately performing the GDOP calculations once all the satellites visible at a terrain point have been identified, rather than separating the two tasks. For visualization purposes, the process was run for different elevation mask values and the number of satellites visible at one epoch was recorded. Figure 5.3 shows the results for a small portion of the test area.

5.4. Dilution of Precision

Dilution of precision values are a measure of the geometric strength of the visible satellite configuration. A minimum of four satellites are needed to calculate GDOP; if there are more than four visible, the calculations determine the best configuration of four. The first step is to construct the range matrix A . This matrix has four columns and as many rows as there are visible satellites. The matrix entries are determined using the

following equations:

$$R_i = \sqrt{(s_{i,x} - c_x)^2 + (s_{i,y} - c_y)^2 + (s_{i,z} - c_z)^2} \quad (5.5)$$

$$A[i][0] = (s_{i,x} - c_x)/R_i \quad A[i][1] = (s_{i,y} - c_y)/R_i \quad A[i][2] = (s_{i,z} - c_z)/R_i \quad A[i][3] = 1, \quad (5.6)$$

where i is the satellite index, starting from 0, R_i is the range to the i^{th} satellite s_i with coordinates $(s_{i,x}, s_{i,y}, s_{i,z})$, (c_x, c_y, c_z) are the coordinates of the receiver, and $A[i][j]$ is the entry in the i^{th} row and j^{th} column of matrix A . The GDOP value is determined from the diagonal entries of the covariance matrix P :

$$P = (A^T A)^{-1} \quad (5.7)$$

$$GDOP = \sqrt{P[0][0] + P[1][1] + P[2][2] + P[3][3]}. \quad (5.8)$$




Figure 5.3
Number of satellites visible at epoch 01:15:00. a) Five degree elevation mask. b) Ten degree elevation mask. c) Fifteen degree elevation mask. d) Twenty degree elevation mask.

The calculations were performed at each epoch for every terrain point at which there were at least four satellites visible, using the centre point of the test area as the location of the receiver. While it is possible that using the actual terrain point coordinates would not increase processing time significantly, it would also not significantly affect the GDOP values and therefore there is no benefit to adding the extra conversion step. When the calculations have been completed, the output is a set of up to 96 images per SP3 file, each corresponding to the GDOP values for the test area at a specific 15-minute epoch.

These images, in and of themselves, are not necessarily of much use. Figure 5.4 helps to illustrate this point. It is an image representation of the GDOP values in the test area for the first epoch on 1 October 2006. Since GDOP values range from one to infinity, a scale proposed by Jon Person [n.d.] was applied for display purposes. There are seven classifications in this scale: ideal (1), excellent (2-3), good (4-6), moderate (7-8), fair (9-20), poor (21-50), and unacceptable (greater than 50). Any terrain points where fewer than four satellites were visible were classified as having an unacceptable value, although it is possible that there are enough satellites visible to give an acceptable position reading. Restricting the number of colours makes it easier to see where the boundaries are of the areas that are classified as having good or better GDOP values, but it is difficult to associate these areas with the physical terrain.

The satellite data and LiDAR data are in different coordinate systems, requiring some coordinate transformations before they could be compared. Since the dimensions of any target area are much smaller than the nominal satellite orbital radius, a single point was used to represent the terrain. Vectors were constructed to represent Easting,

Northing and elevation so that the satellite positions could be expressed using the same two angles, elevation and direction, as the obstructions. This allowed visible satellites to be identified and, where at least four satellites were visible, GDOP values to be calculated. A variety of methods for displaying the GDOP maps is discussed in the next chapter.




Figure 5.4
GDOP values at epoch 00:00:00, scaled to seven classifications.

CHAPTER 6

DISPLAYING RESULTS

Evolution... is – a change from an indefinite, incoherent homogeneity, to a definite coherent heterogeneity.

~ Herbert Spencer

It is not difficult to display the results for a single epoch, using either exact values or scaled values. Applications such as *3DEM* and *MacDEM*TM, a similar application available for the classic Macintosh operating system, can read binary files and allow the colour palette to be changed to reflect the values in the file. It is also possible to generate a bitmap image, and to explicitly specify as many colours as needed. The goal, however, is to jointly display results for multiple epochs in a manner that allows the user to quickly and easily compare them. This goal can be accomplished in various ways. For the most part, the display options proposed here combine the important information into a single image, or a series of images that the user must scroll through, but it is also possible to create dynamic displays.

6.1. Creating Images

The first consideration was the file format to be used to display the images. It is relatively simple to output a binary file containing the calculated GDOP values, yet without a program that can read the file and translate it into an image display, this format is not useful. In contrast, while generating a file in a standard image format is more

complicated, the image can be viewed in a number of widely available applications, including web browsers. The most common image formats are bitmap, Graphics Interchange Format (GIF), and Joint Photographic Experts Group (JPEG) format. Whereas the GIF and JPEG image formats are more advanced and generally produce smaller image files, the bitmap image format is more straightforward to generate and also has different options for number of bits per pixel, allowing flexibility in balancing file size with image detail.

Bitmap images can have four, eight, or twenty-four bits per pixel. In all cases, colours are specified using one byte for each of the red, blue, and green colour channels, giving a theoretical palette of over 16 million colours; in practice, the human eye can only discriminate between a fraction of this number. With four-bit and eight-bit images, the colour palette is specified in the file header and the pixel values are the index values of the palette, which further limits the number of colours available. Four-bit images can have a maximum of sixteen colours, which is sufficient to display the scaled GDOP values with two gradations per classification. Eight-bit images can have a maximum of 256 colours, which is sufficient for up to twenty gradations per classification. Twenty-four bit images can use the full range of colours but there is likely little advantage in doing so because of the difficulty of differentiating similar hues. The images used to illustrate different display options in the following sections are eight-bit images.

The colours used in an image could be selected explicitly, with the user having full control over which GDOP value ranges should be most easily distinguishable. However, as the number of colours in the palette increases, this task would become tedious.

Generating a palette automatically, based on the desired number of colours, greatly reduces user control but ensures the selected colours are spread across the visible spectrum.

6.2. Adding Features

By this stage, the images no longer have easily identifiable features; terrain objects have been replaced by irregular shapes of varying colour. The colour palette used in Figure 5.4 does not distinguish between areas where the GDOP value is unacceptable (greater than 50) and areas where the GDOP value cannot be calculated. This results in a significant portion of the image being red, making it difficult to locate specific physical features. The difficulty can be partially resolved by changing the colour palette, as in Figure 6.1a where the lowest GDOP values are displayed in blue, the highest GDOP values are displayed in magenta, no data areas are in black, and areas where the GDOP value otherwise cannot be calculated are in white. However, there are still large, featureless areas in the image. In the DSM generating process, an image was extracted (Figure 3.5) that shows the approximate outline of buildings and trees. The terrain points directly below the outlines are completely obstructed and therefore the corresponding points in the GDOP map should be white. Consequently, the outlines can be overlaid on the GDOP map without losing any information, as shown in Figure 6.1b.

With the features restored, a question is raised about whether there is any benefit in removing them from the DSM at all. As can be seen in Fig. 6.1b, the GDOP values on building rooftops are good or excellent, as expected, while no GDOP values can be




Figure 6.1
 GDOP values at epoch 00:00:00, mapped using 240 colours representing seven classifications a) without object outlines b) with object outlines.

calculated for terrain points lying directly below buildings, as can be seen by the white areas within the black building outlines where the rooftop elevations were eroded. Yet, if trees were not removed, the GDOP map might contain misleading information that corresponds to the sky view from treetops. Therefore, at least one erosion iteration does need to be performed, unless the tree points in the original LiDAR data can be properly classified and removed accordingly. One advantage in not removing buildings completely from the DSM is that the best GDOP value for a particular epoch can be determined by looking at the areas corresponding to rooftops, since these areas can be expected to have a clear view of all visible satellites.

6.3. Static Images

The simplest, most straightforward method of displaying images for multiple epochs is as a sequence of static images. With this approach, the user must scroll through the images and determine which locations are most suitable based on simple visual comparisons and interpretation. While it is possible to discern trends and patterns in the changes in GDOP values, since the values are not directly related to the satellites' movement, these trends and patterns could change suddenly. For example, there could be a decreasing trend at a location, meaning the location seems to become more suitable over time, followed by a sudden increase. This could occur as a satellite moving towards the horizon will often improve the configuration despite the movements of other visible satellites. Once it falls below the elevation mask or behind an obstruction, however, the resulting configuration of the remaining satellites could be poor.

Figure 6.2a shows an example of a sequence of images corresponding to the GDOP values at five epochs, with the earliest epoch, 00:00:00, at the bottom and the latest epoch, 00:01:00, at the top. At the start of the sequence, there are many areas that have GDOP values that are moderate or better. The strength of the configuration is overall slightly better in the second epoch, as evidenced by the slightly darker blue on the rooftops, yet many of the areas that had moderate GDOP values now have fair or poor values; this likely indicates that a satellite is close to the horizon, or elevation mask, and is obstructed from view at most terrain points. Initially, there are several good options for receiver locations but at the third epoch reception is at best rated fair (the lowest GDOP value at this epoch is 9.28).

The map for the fourth epoch is curious: an insufficient number of satellites are visible at most terrain points to calculate GDOP, yet rooftops seem to have excellent values (3.15). An analysis of satellite positions reveals that there are eleven satellites above the horizon, five of which are below the 15° elevation mask. Of the remaining six satellites, two are between 15° and 30° and one is between 30° and 45°. Clearly, the obstructions and the satellites are lined up at this particular epoch in such a way that most of the satellites are blocked throughout the area. By the next epoch, sufficient satellites are again visible and the GDOP values improve. This example shows that relying on GDOP values at individual epochs to identify trends can be misleading. It also illustrates the importance of considering the presence of obstructions when determining optimal locations for satellite receivers.

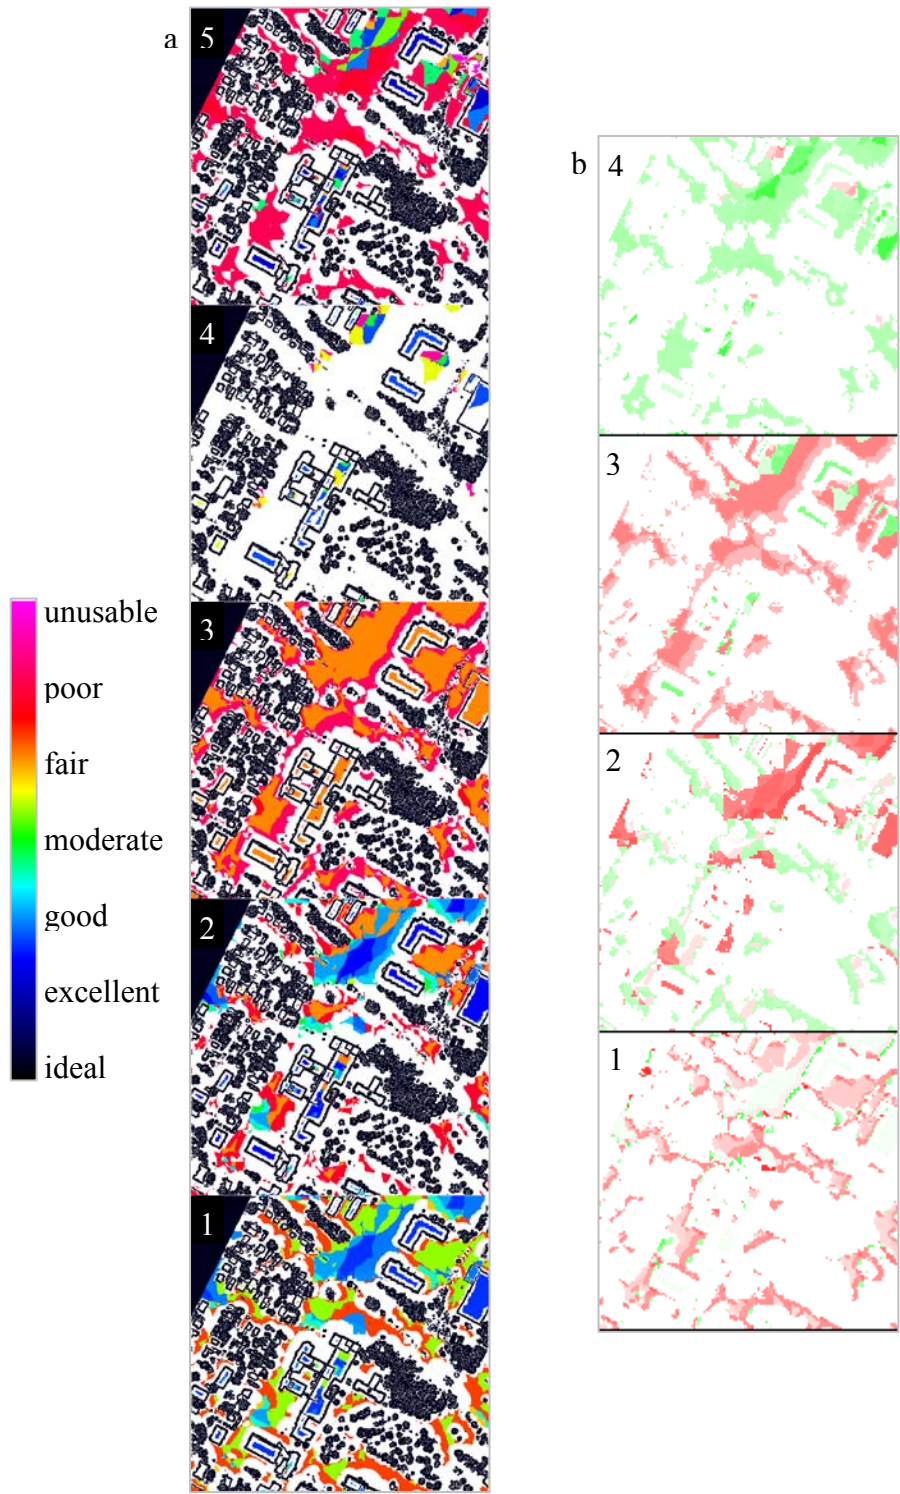


Figure 6.2
 a) GDOP values for five epochs, from 00:00:00 to 00:01:00. b) Change in GDOP values between epochs.

6.3.1. Difference Images

One of the problems with relying solely on a series of GDOP maps is that it can be difficult for a user to ascertain whether overall the values are improving or deteriorating between epochs. Images providing additional information could be helpful in the interpretation. One example is images that show how much the GDOP values have changed between epochs. This idea is illustrated in Figure 6.2b. White indicates no change, while increasing colour saturation indicates increasing change with green hues corresponding to changes towards ideal GDOP values and red hues corresponding to changes towards unusable GDOP values. These changes give a better sense of how the satellites are moving relative to each other and the terrain than the GDOP maps do. In the first difference image, the bottom image in Figure 6.2b, there is a slight to moderate deterioration at almost all locations, with almost imperceptible improvements on the rooftops and slight improvements at a few locations. This means that there is a satellite that has improved the satellite configuration only at locations with a clear sky view; elsewhere, the satellite configuration has become poorer.

In the next image, there is a strong deterioration (a difference of more than two points) at many locations but also a moderate improvement (a difference of close to one point) at many locations. Since the rooftops are among the locations where the GDOP values deteriorated, it is clear that an important satellite has been lost from view; that particular satellite was likely already out of sight at the other set of locations and so its disappearance did not adversely affect them. The deterioration continues in the next image, with GDOP value increases of one to two points, but again there are a few areas

where the GDOP value has increased significantly, decreasing by more than two points. These few areas correspond to the areas in Figure 4.5a where there are no obstructions in the northwest sector (top image), which indicates that a satellite has risen above the elevation mask in the northwest sector but is still below most obstructions. Either it, or another satellite, has become visible throughout the area by the next epoch, causing general improvement in GDOP values.

By themselves, the difference or change images are not useful since they do not show the associated GDOP values and therefore do not give context to the change. A deterioration at one terrain point could be a deterioration from good to moderate and at another it could be a deterioration within the good range yet if the relative deterioration is the same, the colour of the two terrain points will be the same in the difference image. Similarly, an improvement at one terrain point could be an improvement within the poor range and at another it could be an improvement from moderate to fair, but both could be represented by the same shade of green in the difference image. Regardless, the difference image does provide a quick overview of how the values are changing over time. However, since the epoch interval is fifteen minutes and satellites can change their positions significantly in that time, one must not assume that the GDOP values at intervening times can be accurately calculated from the values represented in these images – that requires interpolating satellite position data.

6.3.2. Optimal Time

While the difference images can illustrate trends between epochs, examining a series of difference images does not make the task of selecting an optimal time and location for GPS data collection any easier than does a series of GDOP maps when there is a large number of epochs to consider. In this case, an image showing the time at which the GDOP value is the best at each location could be beneficial. The optimal time could correspond to a single epoch or the average over an interval of any length, from half an hour to half a day, although longer intervals are more likely to have a larger range between the best GDOP value and the worst GDOP value. Figure 6.3 is an example of an optimal time image. The GDOP values for sixty-four epochs, from 06:00:00 to 21:45:00, were reduced to average values over two-hour intervals. The intervals start on the hour and end eight epochs later. Each interval is represented by a different colour. White areas are those where no GDOP values were calculated or the average value does not improve after the first interval; black areas are object outlines.

As with the difference images, optimal time images do not indicate what the actual GDOP values are, nor even the average GDOP value for the optimal interval. If a user has a preferred location, he or she could quickly determine the optimal time for data collection at that location but would then still need to consult the GDOP maps for that optimal time to determine if it is suitable. Another limitation is that two, or more, time intervals might have the best GDOP value. For simplicity, the earliest interval is designated as the optimal time. Since the best GDOP value is determined by averaging, however, an additional condition could be imposed to reduce the likelihood that an

optimal time is in fact unsuitable: the optimal time interval cannot contain a GDOP value below moderate.




Figure 6.3
Optimal two-hour interval between 06:00:00 and 22:00:00.

6.3.3. Maximum and Minimum

Using averaging to determine the optimal time for data collection can be deceptive since several good GDOP values could compensate for one poor GDOP value in any set time interval. Examining the GDOP maps corresponding to a selected optimal time interval is one way to verify whether the interval contains any poor values; another way is to generate images showing the minimum and maximum GDOP values for the interval, or equivalently best and worst GDOP values. In Figure 6.3, there is a long area

to the west of two residential buildings where the optimal data collection time is between 8 and 10 am (shown in dark blue); Figure 6.4 shows the minimum and maximum GDOP values for this time interval. In Figure 6.4a, it can be seen that the minimum GDOP value for the area of interest is in the good range, whereas in Figure 6.4b, it can be seen that the maximum GDOP values for the area are fair or poor. These images do not give any indication of whether there is another time interval when the average might be lower but the highest GDOP values might also be lower; what they do show is that 8:00 to 10:00 am is not a suitable time for data collection at this location.




Figure 6.4
GDOP value extremes between 08:00:00 and 10:00:00. a) Minimum. b) Maximum.

6.4. Dynamic Images

The images described in the above sections organize the information from several epochs in different ways to emphasize different aspects, such as how the GDOP values have changed or the optimal time for data collection, yet they can only supplement the information in the GDOP maps, not replace it. It is therefore desirable to have an

interface through which the user can view GDOP maps in a controlled order, and which can be run on any computer. One possible variation of the interface is shown in Figure 6.5. It utilizes HTML (HyperText Markup Language) and JavaScript, which can be viewed in any web browser, to create a slideshow. Since pixel dimensions of the GDOP maps are likely to be large, the epoch value can be embedded in the first pixel with no loss of information, which enabled the slideshow generator that was developed to include the time of each epoch in the set in a drop-down menu, allowing the user to jump to a specific epoch. Buttons are also included to allow the user to jump to the start or the end of the set, to advance or back-up one epoch at a time, and to start or stop automatically advancing through the images. These functions might not be needed for small sets of GDOP maps, but they become more important as the size of the set increases. In particular, displaying the epoch time rather than its number allows the user to more easily identify suitable or unsuitable data collection times. With HTML, the colour scale can also be included easily alongside the image and potentially other data could be included, such as exact GDOP values and coordinates.

Since the number of GDOP values calculated in even a small area for a small number of epochs is so large, the information needs to be organized in a way that can be easily interpreted. Some of the options for displaying the information are difference images, optimal time images, and maximum and minimum images. Before any of these can be implemented, however, the entire process from generating the DSM to identifying visible satellites and calculating GDOP values should be assessed to ensure that the results are valid. This is the topic of the next chapter.

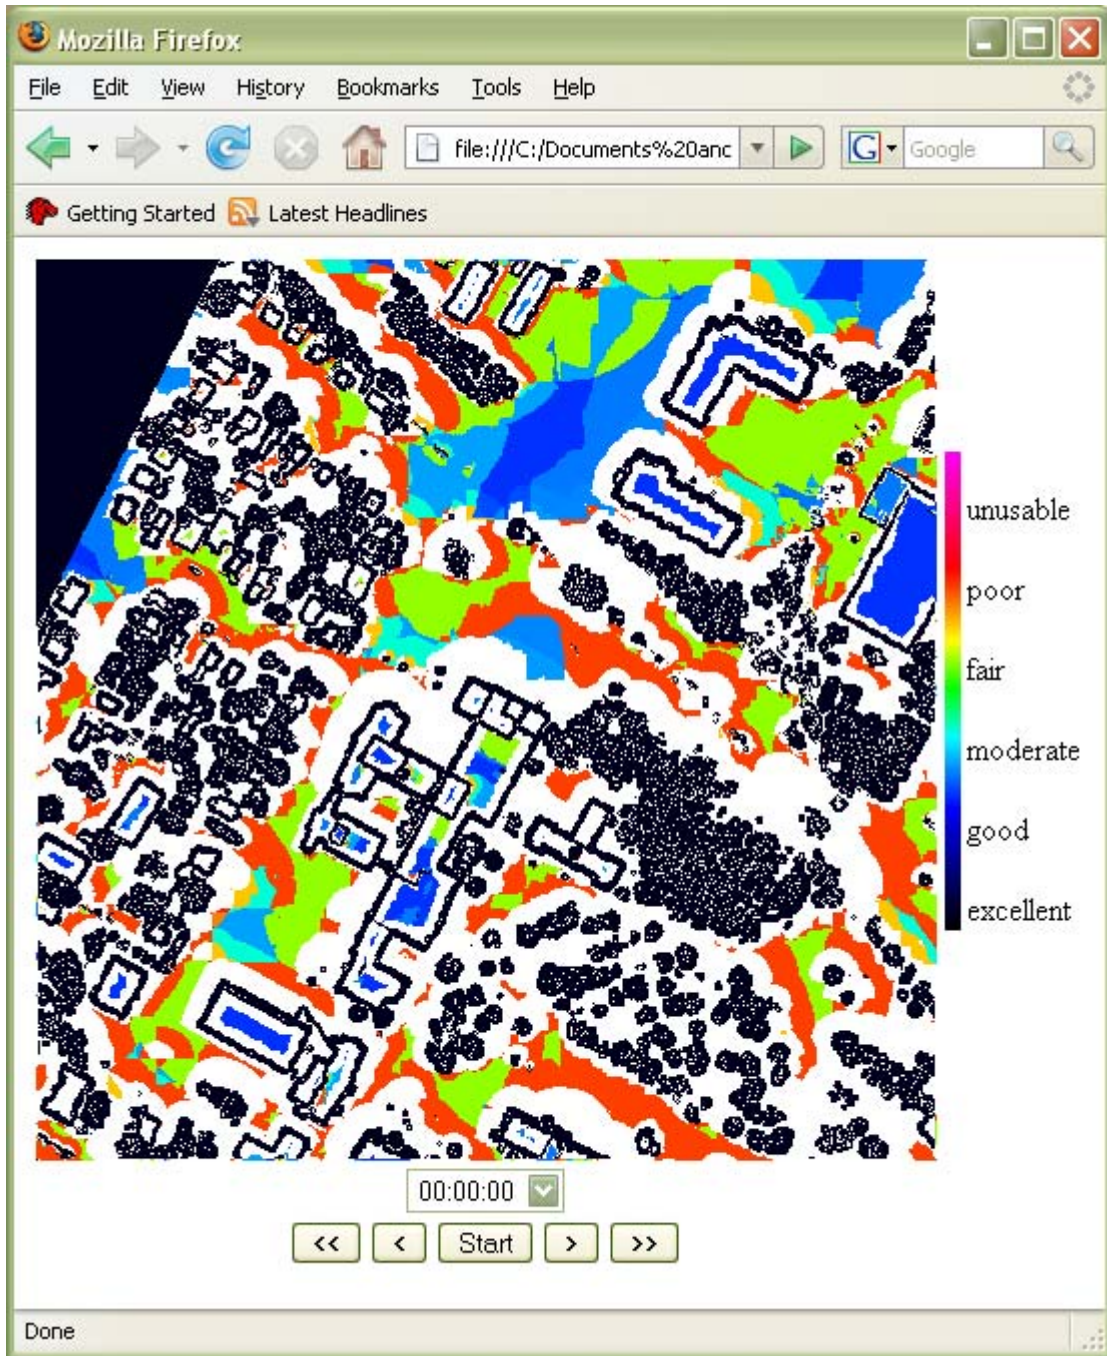


Figure 6.5
Possible web browser-based interface for viewing GDOP maps.

CHAPTER 7

ASSESSMENT

*Conscience is the inner voice that warns
that somebody may be looking.*

~ H. L. Mencken

The final stage in the development of a method to generate GDOP maps that reflect the presence of local obstructions is an assessment of the results. Since the results are primarily visual in nature, the assessment is also primarily visual. However, two software tools were used to support the evaluation: *GPSViewTM* and *Trimble PlanningTM*. *GPSView* is an application that runs on the Macintosh operating system. It displays sky plots of visible satellites at a specific location and time. *Trimble Planning* runs on the Windows operating system. It has several options for displaying satellite visibility and also DOP values.

The most important aspect to assess is whether the GDOP maps produced are accurate representations of the situation in the field. Arguably, if the final results are sufficiently accurate then the results at intermediate stages are also sufficiently accurate, yet any improvements to the process are likely to be caused by modifications to the earlier stages. Therefore, three aspects were assessed: DSM generation, to determine whether it produces a suitable basis for potential receiver locations; identifying visible satellites and calculating best GDOP values, to determine whether the algorithm correctly locates satellites above the horizon and calculates GDOP values to a

satisfactory degree of accuracy; and effect of obstructions, to determine whether obstructed satellites are correctly identified.

7.1. Assessment Data

While the first two aspects can be assessed over the whole test area, specific terrain points must be used to properly assess the effect of obstructions. Five points were selected for this purpose. Four are ground points and the fifth is located on a rooftop. The images in Figure 7.1 show the locations of the target points within the test area and 3D views of the immediate surroundings. The orientation of the 3D scenes was rotated so as to allow the clearest view of the target points, which are located approximately in the centre of the scenes. Arrows above each image indicate the direction of North for that particular scene.

The obstructions are identified using six levels in four sectors, as described in Chapter 4 and illustrated in Figure 4.5a. The obstruction values for the target points are listed in Table 7.1. These values should be interpreted as in the following example: The selected point in the Head Hall A-level parking lot has an obstruction level of 2 in Sector 3. Since the levels are numbered 0 to 5, the second level corresponds to the interval 30° to 45° . Sector 3 is the south-west quadrant. Therefore, satellites in the south-west quadrant of the sky must have an elevation above 45° to be deemed visible. A slight exception in interpretation occurs with Level 0 obstructions. The rooftop point has Level 0 obstructions in all directions and while it is very likely that there are, in fact, no obstructions for that point, satellites must nevertheless have an elevation above 15° to be




Figure 7.1

Three-dimensional views of assessment target locations: a) commercial parking lot; b) residential rooftop; c) Head Hall parking lot; d) Provincial Archives parking lot; e) sidewalk at Sir Howard Douglas Hall. f) Aerial view of target point locations.

Table 7.1

Target point location descriptions and sector obstruction levels.

Description	Northing	Easting	Sector obstruction levels			
			0	1	2	3
A Parking lot beside commercial buildings	5091406	682803	1	1	1	0
B Roof of residential building	5091315	682888	0	0	0	0
C Head Hall A-level parking lot	5091265	682765	1	1	2	2
D Sidewalk in front of Sir Howard Douglas Hall	5091018	622818	1	2	4	3
E Parking lot beside Provincial Archives	5091104	682650	0	2	2	2

deemed visible because that might not be the case for Level 0 obstructions at ground points. To avoid any negative consequences on satellite visibility, the value of the elevation mask was selected to coincide with the minimum elevation needed to clear Level 0 obstructions.

Different GPS data were used for the assessment than those that were used during development. GPS almanacs provide satellite orbital parameters for an epoch typically two days after their time of release and are used by planning software, including *Trimble Planning* and *GPSView*, to calculate approximate satellite positions. At the time the assessment was performed, the current GPS almanac had a time of applicability 503808.0 s from the start of GPS week 1463, that is, approximately 20:00 Friday 25 January 2008, while the most recent satellite position data available from IGS were from the beginning of that week. A difference between the time of applicability and the time of position data introduces a greater probability of a difference between predicted and actual satellite positions, but it should not be significant enough to adversely affect the assessment. The assessment data were from Ultra-Rapid file *igu14631_18.sp3*, which contains data for GPS week 1463, starting at 18:00:00 on day 0 (20 January 2008), and spanning 196 epochs. The selected time interval was from 22:30:00, 20 January 2008, to 13:15:00, 21 January 2008, or from epoch 18 to 77.

7.2. Generated DSM

One of the reasons for using the LiDAR data to generate a DSM rather than obtaining one from another source was to minimize the possibility of height differences

in the terrain when detecting obstructions. Given that terrain points are included in the potential obstruction point cloud, if the DSM terrain heights were consistently below the LiDAR terrain heights, the entire area would appear completely obstructed, whereas if the DSM terrain heights were consistently above the LiDAR ones, obstructions would be regularly underestimated. Either scenario could also occur, however, if the approach used to generate the DSM is not suitable. A one-metre height offset was applied when detecting obstructions to better reflect the position of the receiver relative to the ground. This had the effect of moderating the impact of minor, local variations in the terrain, but the variations between the DSM and the LiDAR elevations should be much less than a metre for the DSM to be considered a suitable basis for potential receiver locations.

Since the DSM was generated by averaging points' elevations, one method of assessing the result is to examine the standard deviation of elevations within each grid cell. In Figure 7.2a, cells in the average terrain where the standard deviation is greater than 7 m are red and those with values below 4 m are in blue, fading to white as the value approaches 0 m. The highest standard deviations are in areas of high vegetation and at building perimeters; terrain and rooftop points, which are the potential receiver locations, have small standard deviations. Figure 7.2b shows the standard deviations of the DSM, which is the average terrain after five erosion iterations. Cells with a standard deviation greater than 1 m are in red, values decreasing from 1 m to 0.5 m have increasing blue tones, and values below 0.5 m are blue, fading to white. The red areas are those areas where erosion occurred and are therefore not of interest. The main source of small terrain differences is cars in parking lots, visible as small blue shapes. There are

also areas of the terrain where the slope introduces greater variation, but overall the terrain variation within each cell is very small, which indicates that the DSM can be used for identifying obstructions and is suitable as a basis for potential receiver locations.




Figure 7.2

Standard deviation of elevations within grid cells: a) interpolated terrain, scale range 0 – 7 m; b) DSM, scale range 0 – 1 m.

7.3. Visible Satellites and GDOP Calculations

To correctly identify satellites that are blocked by obstructions, or conversely satellites that are visible from a particular ground location, the algorithm must first correctly identify satellites above the horizon or above the elevation mask. This was assessed by comparing the list generated using *Trimble Planning* of the number of satellites visible at each epoch with the values output during processing. In Figure 7.3, there are two sets of points. The first set is the number of visible satellites; pink squares mark the number of visible satellites according to *Trimble Planning*, while purple

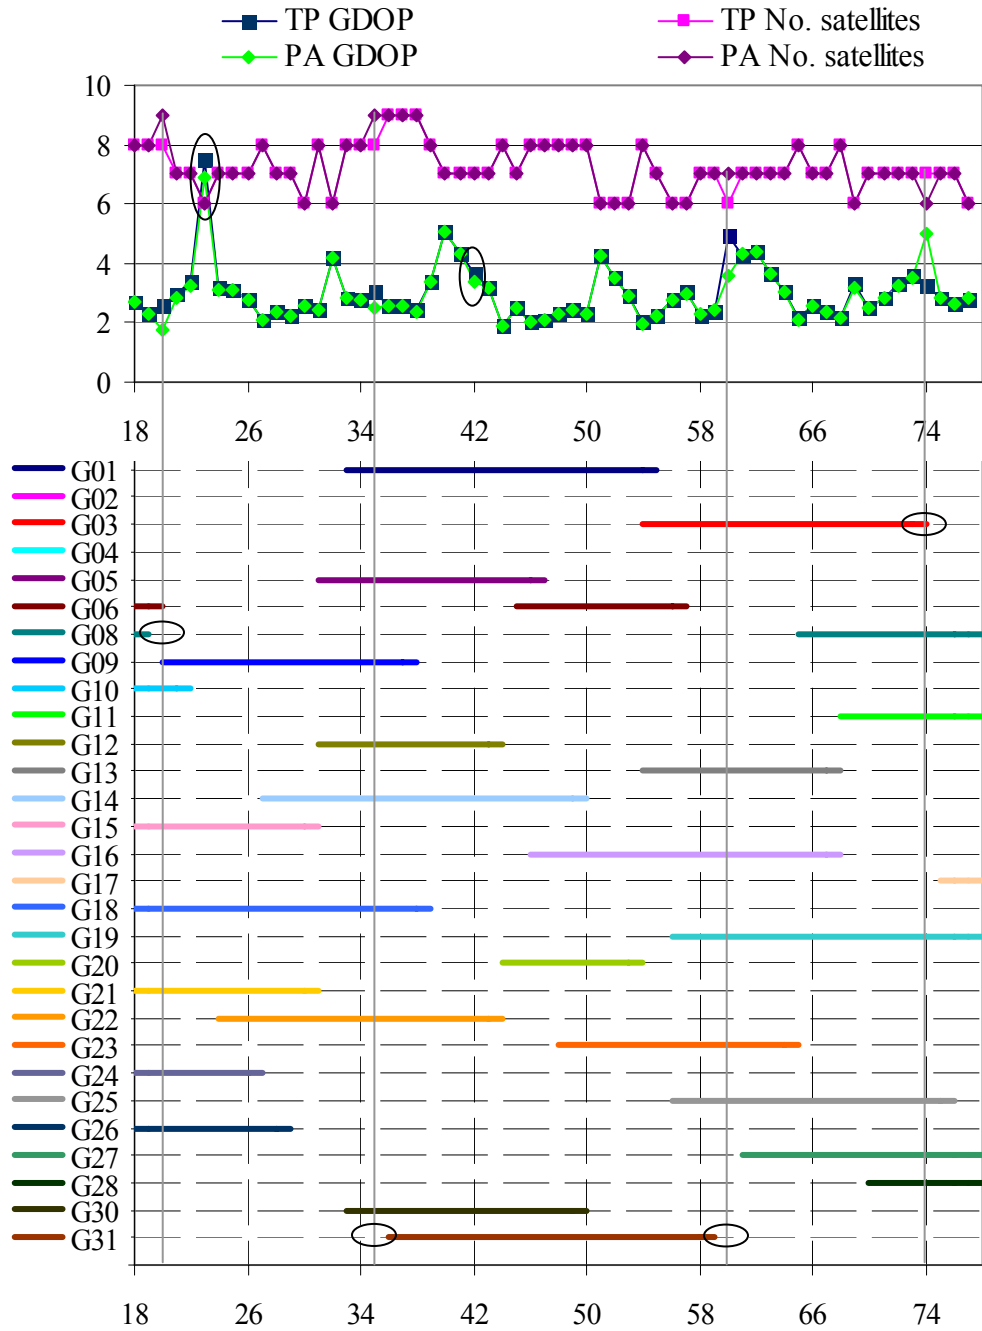


Figure 7.3

Visible satellites and GDOP values from *Trimble Planning* and calculated by proposed algorithm for epochs from 22:30:00, 20 January 2008 to 13:15:00, 21 January 2008.

Figure 7.4

Satellite visibility from 22:30:00, 20 January 2008 to 13:15:00, 21 January 2008.

diamonds mark the number of visible satellites detected by the proposed algorithm at each epoch. At all but four epochs, the two sets are coincident and the difference is only one satellite at each of those four epochs. Figure 7.4 is a plot of the epochs at which each satellite is visible, again according to *Trimble Planning*. At each of the epochs where there is a discrepancy, there are satellites that have just stopped being visible or are about to become visible; differences in the way a satellite's elevation is calculated, though otherwise irrelevant, may have placed the satellite on different sides of the elevation mask at critical epochs.

The second set of points shows the GDOP values for clear sky view locations, with the blue squares marking the values calculated by *Trimble Planning* and the green diamonds marking the values calculated by the proposed algorithm at each epoch. The largest differences are at the epochs where a different number of visible satellites were detected: 0.81 at epoch 20 (23:00), 0.51 at epoch 35 (02:45), 1.35 at epoch 60 (09:00), and 1.77 at epoch 74 (12:30). There are two other epochs at which the GDOP values are noticeably different: epoch 23 (23:45), with a difference of 0.59, and epoch 42 (04:30), with a difference of 0.3. The reason for these differences is unknown. At all other epochs, the difference is less than 0.1, with the majority having a difference of less than 0.05. While further investigation may be beneficial to identify aspects of the algorithm that need refinement, given that in the final results GDOP values are converted to colours in a continuous scale, these differences are sufficiently small that the proposed algorithm can be said to be identifying visible satellites and calculating GDOP values correctly.

7.4. Effect of Obstructions

The advantage of *GPSView* over *Trimble Planning* is in the way sky view plots are displayed: each plot shows the position of satellites above the horizon, including those below the elevation mask, at a single epoch. The elevation mask is shaded and the concentric rings can be set to 15° distances, making comparison between the algorithm-determined list of visible satellites and the sky view plot straightforward. The following figures, 7.5 – 7.7, are sky view plots for three selected epochs. The obstruction profiles for the target points, excluding point B, which does not have any significant obstructions, have been overlaid to further aid in comparing satellite visibility (A in blue, C in green, D in red, and E in yellow). The outer-most ring, shaded grey, corresponds to Obstruction Level 0, 0° - 15° ; satellites in this ring are below the elevation mask. The innermost circle corresponds to Obstruction Level 5, 75° - 90° . The obstruction sectors start at northeast and are counted counter-clockwise (northeast, northwest, southwest, southeast). Satellites located in the ring corresponding to the identified obstruction level and below in a given sector for a target location are not visible at that target location.

Figure 7.5 is the *GPSView* plot for 22:30, 20 January 2008 (epoch 18). It shows eleven satellites available, eight of which are visible. Table 7.1 summarizes the algorithm-determined visibility for this epoch, limited to the eight satellites above the elevation mask. Although all the points have different obstruction profiles, the same satellites were determined to be visible at points A, C, and D. The reason for this can be seen in Figure 7.5: satellites PG06, PG08 and PG18 are in level 1, which is obstructed in

Table 7.2
Algorithm-determined satellite visibility at 22:30, 20 January 2008.

ID	Visible at Target Point			
	A	C	D	E
6	no	no	no	no
8	no	no	no	yes
10	yes	yes	yes	yes
15	yes	yes	yes	yes
18	no	no	no	no
21	yes	yes	yes	yes
24	yes	yes	yes	yes
26	yes	yes	yes	yes
Total	5	5	5	6

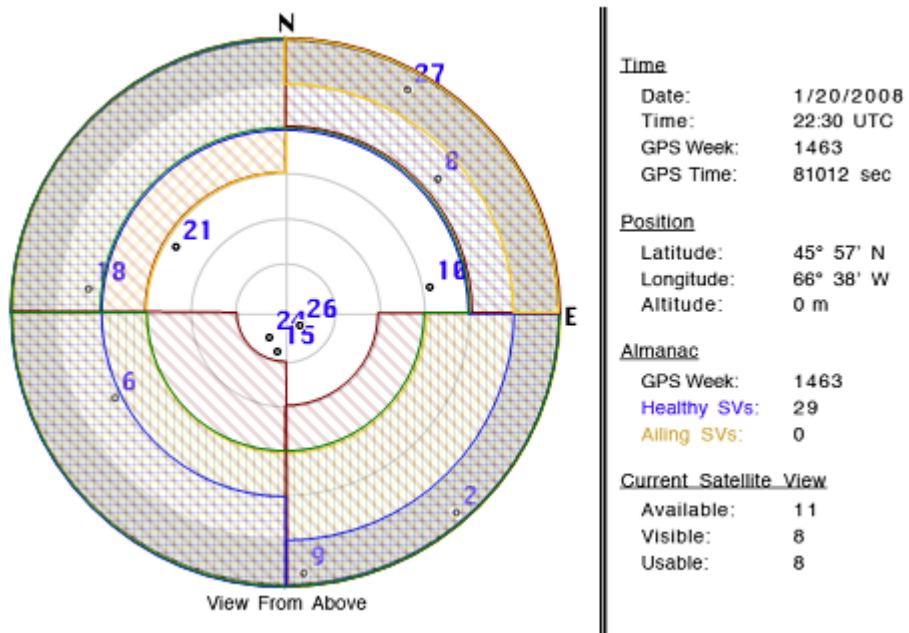


Figure 7.5
Sky view plot, 22:30, 20 January 2008.

Sectors 2 and 3 for all the target points, and only clear for point E in Sector 1 and for point A in Sector 4. At all four target ground points, the GDOP value is only fair, due to three of the satellites being clustered together; at the rooftop point (target point B), the GDOP value is excellent. In all cases, the satellites identified by the algorithm as being visible correspond to the configuration displayed in *GPSView*.

The overall satellite configuration is better at 03:00, 21 January 2008, (epoch 36), in terms of satellites being present in all four sectors, as can be seen in Figure 7.6. Yet, since fewer satellites are at a high elevation, at two of the target points there are not enough visible satellites to calculate GDOP; algorithm-determined satellite visibility is summarized in Table 7.3. At target point D, only satellites PG09 and PG14 are visible: PG22 is a few degrees too low to be visible above the Level 4 obstruction in Sector 2. It is visible at target point E, while satellites PG05 and PG12 are almost visible. In the sky view plot, PG05 appears to be above 45° , although the algorithm determined its elevation to be slightly below that threshold. This is a further indication that slight differences in calculation methods can have a negative impact on satellite detection, and also suggests that 15° obstruction levels are too coarse. The remaining target points have good or excellent GDOP values, even though at target point C, the satellites in Sector 3 are all deemed to be obstructed, as is satellite PG31.

In the above two cases, twice the number of satellites needed for a GDOP calculation were above the elevation mask. At several other epochs, only six satellites were potentially visible from the target points. Their configuration at 07:00 21 January 2008 (epoch 52) can be seen in Figure 7.7 and the algorithm-determined visibility is

Table 7.3
Algorithm-determined satellite visibility at 03:00, 21 January 2008.

ID	Visible at Target Point							
	A	C	D	E				
1	yes	yes	no	no				
5	yes	no	no	no				
9	yes	yes	yes	yes				
12	yes	no	no	no				
14	yes	yes	yes	yes				
18	yes	no	no	no				
22	yes	yes	no	yes				
30	yes	no	no	no				
31	no	no	no </tr <tr> <td>Total</td> <td>8</td> <td>4</td> <td>2</td> <td>3</td> </tr>	Total	8	4	2	3
Total	8	4	2	3				




Figure 7.6
Sky view plot, 03:00, 21 January 2008

Table 7.4
Algorithm-determined satellite visibility at 07:00, 21 January 2008.

ID	Visible at Target Point			
	A	C	D	E
1	yes	no	no	no
6	yes	yes	yes	yes
16	yes	yes	no	yes
20	no	no	no	no
23	yes	yes	no	no
31	yes	yes	yes	yes
Total	5	4	2	3




Figure 7.7
Sky view plot, 07:00, 21 January 2008.

summarized in Table 7.4. The result is comparable to the result of the configuration at epoch 36: at target point D, only two satellites are visible, PG06 and PG31 in Sector 0; at target point E a third satellite, PG16 in Sector 2, is visible. The remaining target points, however, have poorer GDOP values at epoch 52 than at epoch 36. A similar assessment could be performed on the remaining epochs in the selected time interval; additional sky view plots and corresponding algorithm output are included in Appendix II. The three selected epochs, though, are an adequate sample to demonstrate that the algorithm is able to identify and quantify obstructions, correctly identify most satellites that are not blocked by obstructions, and calculate GDOP values sufficiently accurately when at least four satellites are visible. Further work may be needed to resolve the possible misidentification of satellites near the elevation thresholds.

CHAPTER 8

CONCLUSIONS AND RECOMMENDATIONS

*A man of destiny knows that beyond this hill
lies another and another. The journey is
never complete.*

~ F. W. de Klerk

The objective of this research was to develop a method for mapping GDOP values that takes into account local obstructions. There were two major tasks to accomplish: quantifying obstructions and identifying visible satellites. Since available software tools for generating DSMs and calculating GDOP values based on satellite visibility do not offer the flexibility or control needed to accomplish these tasks, all stages of the proposed method were implemented with C code written by the author. LiDAR data were used to identify and quantify obstructions in a target area, while GPS satellite position data in SP3 format were used to locate visible satellites and calculate GDOP values. Various options were proposed for displaying the mapped GDOP values. Given that the maps use a continuous colour scale, assessment of the results was primarily visual. More rigorous testing is recommended, particularly in conjunction with the development of display options that include showing numerical values. Other recommendations are to create a user interface to simplify processing and display, to interpolate GPS satellite position data to one-minute epochs, to extend the planning capability of the proposed method, and to classify the obstructions to determine the likelihood of signal obstruction.

8.1. Conclusions

LiDAR data represent terrain and non-terrain objects as points in a cloud. Often, objects or parts of objects generate more than one point return per laser pulse, in particular vegetation and building perimeters. This makes it possible to extract two distinct surfaces, albeit with some commonalities where single returns are generated. One surface, consisting of the first returns, contains the highest local elevations – rooftops, treetops, and also ground points. The highest local elevations, including the ground points, are potential obstructions that may block satellite signals from receivers in their shadow. The second surface, consisting of the last returns, contains the lowest local elevations – rooftops and ground points around building perimeters, tree crowns and ground points below trees if the laser achieved full penetration, and terrain points. These are all potential receiver locations, although terrain points not located below trees or in too close proximity to buildings are the most likely candidates. The proposed method includes a process through which these two surfaces can be compared and obstructions can be quantified for every potential receiver location. It also includes a process through which GPS satellite position data can be compared to the obstructions to determine satellite visibility and then can be used to calculate and map GDOP values that could be expected in the field. By means of these processes, the research objective was successfully achieved.

8.1.1. Quantifying Obstructions

Identifying and assessing obstructions using the proposed method requires a significant amount of processing time, on the order of several hours even for small target areas. This is due to the large volume of LiDAR data, which is not in any easily predictable order, particularly when data from multiple flight lines are merged into one file. It was determined that the potential obstruction surface should be left as a point cloud because this would provide the most accurate representation. However, to decrease processing time and facilitate obstruction detection, the potential terrain points were interpolated to a grid. Non-terrain objects were partially removed using a simple erosion operator that was developed so that the resulting surface contained primarily terrain values only. Once the comparison of the two surfaces is complete and the obstructions are quantified for each terrain point, the LiDAR data are no longer needed and the remaining processing steps of the proposed method can be completed in a matter of minutes. Furthermore, once the obstructions in a target area have been quantified, the resulting obstruction images can be used any number of times with different satellite data to produce GDOP maps for the desired time intervals.

For each terrain point, obstructions were reduced to two angles: elevation and direction. The elevation angle indicates how high a satellite must rise in the sky to be visible above the obstruction while the direction angle indicates which portion of the sky is obstructed. The elevation range is divided into levels. Six levels were used during development and testing but any number of levels can be used. Using more levels gives a more accurate picture of the relative elevation of obstructions; it also causes a

significant increase in processing time. The direction range is divided into sectors. As with elevation levels, more sectors correspond to a more detailed picture of the location of obstructions, and a modest increase in processing time. Four sectors were used during development and testing.

8.1.2. Calculating and Mapping GDOP Values

The LiDAR data use the UTM coordinate system while GPS satellite position data, as available from IGS in SP3 format, use the IGS05 coordinate system. Coordinate transformation was therefore necessary to relate the two sets of data. Since GPS satellites have an orbital altitude much greater than the dimensions of any anticipated target areas, it was determined that representing the target area as a single point would be sufficiently accurate, thereby necessitating only a single coordinate transformation for terrain location. Additional transformations were required, however, to relate satellite positions to obstruction angles. The midpoints of the target area were transformed from UTM to IGS05 and used to construct direction vectors corresponding to Northing and Easting. These direction vectors along with the normal vector, which was constructed by taking the cross product of the two direction vectors, were then used at each epoch to determine the elevation angle of each satellite relative to the plane of the target area and the direction angle of each satellite above the elevation mask relative to Easting.

By developing a process that describes satellite positions using the same pair of angles as was used in the process developed to quantify obstructions, the satellite positions could be directly compared to the obstructions at every terrain point in the

target area to identify which satellites were not blocked by local obstructions. GDOP values were then calculated at terrain points where at least four satellites were visible. However, the satellite elevation angles calculated with this approach are not identical to those calculated using available GPS planning software such as *GPSView* and *Trimble Planning*. This could result in a different number of satellites being identified as visible, depending on how close a satellite is to the visibility threshold, and consequently GDOP values different from those of other software tools being calculated. At present, since the GDOP values are translated into colours in a continuous scale for display, these potential differences have been deemed not to have a significant negative impact on the final results. Further testing is needed to determine what effect they may have if numerical GDOP values were retained in the final results.

Each epoch in the GPS satellite data is fifteen minutes long. If a GPS data collection survey lasts two hours, a minimum of eight epochs must be mapped and examined; if the user wishes to determine the optimal collection time interval, many more epochs may need to be examined. Several different display options were developed to assist in this determination. One option is a difference image, which shows the change in GDOP values between two epochs and is an indication of possible trends in the overall GDOP values in the target area. Another option is an optimal time image, which indicates during which time interval the average GDOP value is the best for each terrain point in the target area. The final option is minimum and maximum images, which display the best and worst GDOP values during a specified time interval. These options can be used independently or in conjunction with one another to assist the user in interpreting the

GDOP maps. The images can also be organized in a slideshow that can be viewed in a web browser to facilitate viewing.

8.1.3. Research Accomplishments

DOP values are a measure of the strength of the satellite configuration and give an indication of positioning accuracy. The satellite configuration visible during a GPS survey can be affected by the presence of local obstructions, yet existing methods for determining DOP values do not allow obstructions to be easily incorporated into the calculations. Other researchers have made attempts in recent years to use LiDAR data to model the effect of obstructions on DOP values but these attempts invariably relied on existing software and were therefore limited in applicability.

The method proposed in this work uses only code developed by the author to manipulate the LiDAR data, quantify obstructions, identify visible satellites, and calculate and map GDOP values, thereby extending its applicability and also avoiding any issues that may arise with proprietary software. Because it uses SP3 files for satellite position data, which are available at most a day in advance, currently the proposed method is suitable only for verification purposes after a survey has been completed. In addition, it is limited to locations where LiDAR data is available. Nevertheless, the GDOP maps that are produced accurately reflect the conditions experienced in the field over a wide area, as compared to the idealized conditions at specific locations assumed by existing software tools, and therefore the research objective has been successfully achieved.

8.2. Recommendations

The method proposed satisfies the research objective of creating GDOP maps that reflect the presence of local obstructions. Nevertheless, there are a number of possible ways to further improve and refine the results, including performance tests, classifying obstructions, interpolating satellite data, extending planning capability, investigating other data sources for the obstruction surface and DSM, and creating a user interface.

8.2.1. Performance Tests

The assessments performed were primarily visual in nature. The number of satellites that the algorithm identified as visible at each epoch and the best GDOP value calculated at each epoch were compared with the number and value determined by GPS planning software, but the validity of the GDOP values calculated at obstructed locations was assessed only according to the positions of the satellites in sky view plots. Some differences between the sky view plots and the algorithm output were noted due to the way satellite elevation was determined; field tests would help to resolve these differences. However, because of the coarseness of the obstruction elevation levels and direction sectors, field-determined values might not correspond to the worst-case values determined by the algorithm and would therefore need to be carefully analyzed. Assessments should also be carried out on the results produced when the size of the elevation levels and obstruction sectors is decreased.

Another factor that could affect field tests is leaf-on/leaf-off conditions. When deciduous trees are bare of leaves, from late fall to early spring, they are likely only

minor obstructions for satellite signals; even full foliage in mid-summer will not necessarily fully block satellite signals. Because the algorithm treats all obstructions as fully opaque, it may determine that a satellite is not visible while field tests determine otherwise.

8.2.2. Obstruction Classification

One method of minimizing the probability that a satellite obstructed by vegetation is incorrectly identified as not being visible is to classify the obstructions using an approach similar to that proposed by Lohani and Kumar [2007]. By distinguishing between solid obstructions, namely buildings, and potentially transparent obstructions, namely vegetation, the algorithm could produce results that not only reflect the presence of local obstructions but also the nature of local obstructions. The two classes could be easily distinguished by the number of returns generated by each: solid obstructions generate only one return, whereas potentially transparent obstruction can generate multiple returns, though some proximity-based corrections would be needed. Vegetation could be further differentiated into deciduous and evergreen, potentially making additional terrain available as receiver locations during the winter months. Every additional obstruction classification creates additional GDOP calculations and maps, particularly as the probability of satellite signals being obstructed varies. While in some cases it might be desirable to carry out all possible calculations, the simplest approach would be to generate one map showing the best-case scenario, where only solid

obstructions are considered, and one map showing the worst-case scenario, where all obstructions are considered.

8.2.3. Interpolating between Epochs

The satellite data available from IGS give the positions of satellites at fifteen-minute intervals. However, since the satellites have a twelve-hour orbital period, there could be a significant change in satellite positions between epochs and consequently a significant change in GDOP values. At times or locations where GPS signal reception is generally poor, it could be desirable to know within a minute or two when a satellite has become visible or has stopped being visible, or correspondingly when the GDOP value has noticeably improved or deteriorated. Simply interpolating GDOP values between fifteen-minute epochs to generate maps for one- or two-minute epochs would not give an accurate picture of the changes because the movements of the satellites and the changes in their configuration are not reflected in the GDOP values. Instead, the satellite positions must be interpolated. Different methods for doing so are discussed in Shenewerk [2003]. While this may help to determine to a much narrower margin of error when GDOP values are sufficiently good for data collection purposes, it will create fifteen times as many GDOP maps, if one-minute intervals are used. Additional display methods may need to be devised that focus on specific times or locations to manage the large volume of information.

8.2.4. Extending Planning Capability

Using satellite position data from IGS restricts the proposed method to use in verification, either after a survey has been completed or at most one day before the survey after standard planning methods have been used. For it to be considered as a planning tool, this method must be able to produce sufficiently accurate GDOP maps several weeks in advance. This could be accomplished by using GPS almanacs instead of IGS data, which would require the added step of calculating satellite positions from the orbital parameters in the almanacs. A further advantage to using GPS almanacs is that position data would not need to be interpolated to obtain shorter epoch intervals.

8.2.5. Extracting Surfaces from Other DTM Sources

Currently, LiDAR data is only available for a small percentage of areas where the proposed method might be of use and the cost associated, in both time and money, with a LiDAR survey means that it is not viable to collect data solely for mapping GDOP values. However, the approach could potentially be modified to accept other data sources. To generate obstruction images, all that is needed is an obstruction surface and a terrain surface. These are most easily extracted from multiple-return LiDAR data but could, in theory, be extracted from any elevation model provided the heights of objects in the model are sufficiently accurate. By investigating this possibility, as well as the extent to which the process of generating obstruction images would need to be modified and the accuracy of the resulting images, the number of areas where the proposed method could be applied could be significantly increased.

8.2.6. User Interface

Each stage in the proposed method has specific input parameters, from the number of erosion iterations in the DSM generation to the number of elevation levels and direction sectors in the obstruction analysis, and from the start epoch and number of epochs for the GDOP maps to the time intervals for the displays. Keeping track of each of these parameters can be difficult. A user interface that allowed each of the parameters to be set would be helpful in managing the input and output for each stage. Ideally, once set, the interface would be able to run through the entire process, starting from the raw LiDAR data to the final display output, or start and end at any selected point. For instance, it would need to accept LiDAR data input as well as have an option not to process obstructions, but could also have an option to only process obstructions. The more complete the interface, the greater its utility, and the more possibilities for expanding, developing, and refining the process.

REFERENCES

- Ackermann, F. (1999). "Airborne laser scanning - present status and future expectations." *ISPRS Journal of Photogrammetry & Remote Sensing*, 54, 64-67.
- Axelsson, P. (1999). "Processing of laser scanner data - algorithms and applications." *ISPRS Journal of Photogrammetry & Remote Sensing*, 54, 138-147.
- Baltsavias, E. P. (1999). "Airborne laser scanning: basic relations and formulas." *ISPRS Journal of Photogrammetry & Remote Sensing*, 54, 199-214.
- Beesley, B. J. (2003). "Sky viewshed modelling for GPS use in the urban environment." *Proceedings of the Twenty-Third Annual ESRI User Conference*. San Diego, 7-11 July.
- Beutler, G., M. Rothacher, S. Schaer, T. A. Springer, J. Kouba, and R. E. Neilan (2005). "The International GPS Service (IGS): Celebrating the 10th anniversary and looking to the next decade." *Advances in Space Research*, 36(3), 320-326.
- Brovelli, M. A., and M. Cannata (2004). "Digital terrain model reconstruction in urban areas from airborne laser scanning data: the method and an example for Pavia (northern Italy)." *Computers & Geosciences*, 30, 325-331.
- Dana, P. H. (2000). *Global Positioning System Overview*. The Geographer's Craft Project, Department of Geography, The University of Colorado at Boulder. [On-line] 18 January 2008. <http://www.colorado.edu/geography/gcraft/notes/gps/gps.html>
- Goulden, T. and C. Hopkinson (2006). "Fredericton LiDAR Survey." Presentation to CARIS.
- Kobler, A., N. Pfeifer, P. Ogrinc, L. Todorovski, K. Ostir, and S. Dzeroski (2007). "Repetitive interpolation: A robust algorithm for DTM generation from Aerial Laser Scanner Data in forested terrain." *Remote Sensing of Environment*, 108, 9-23.
- Kraus, K., and N. Pfeifer (1998). "Determination of terrain models in wooded areas with airborne laser scanner data." *ISPRS Journal of Photogrammetry & Remote Sensing*, 53, 193-203.
- Lohani, B., and R. Kumar (2007). "A model for GPS-GDOP prediction in urban environment using LiDAR data." 8th Conference on Optical 3-D Measurement Techniques. Zurich, Switzerland, 9-12 July.

Miliaresis, G., and N. Kokkas (2007). "Segmentation and object-based classification for the extraction of the building class from LIDAR DEMs." *Computers & Geosciences*, 33, 1076-1087.

National Oceanic and Atmospheric Administration (2008). "Remote Sensing for Coastal Management." *NOAA Coastal Services Center*. [On-line] 6 June 2008. http://www.csc.noaa.gov/crs/rs_apps/sensors/lidar.htm

Person, J. (n.d.), "Mastering GPS Programming: Part Two, Geometric Dilution of Precision." *GIS and GPS Components for Visual Studio.NET – GeoFrameworks*. [On-line] 18 January 2008. http://www.geoframeworks.com/Articles/WritingApps2_3.aspx

Roggero, M. (2001), "Airborne laser scanning: Clustering in raw data." *International Archives of Photogrammetry and Remote Sensing*, 34(3/4), 227-232.

Shan, J., and A. Sampath (2005). "Urban DEM generation from raw lidar data: A labeling algorithm and its performance." *Photogrammetric Engineering & Remote Sensing*, 71(2), 217-226.

Shenewerk, M. (2003). "A brief review of basic GPS orbit interpolation strategies." *GPS Solutions*, 6, 265-267.

Silvan-Cardenas, J. L., and L. Wang (2006). "A multi-resolution approach for filtering LiDAR altimetry data." *ISPRS Journal of Photogrammetry & Remote Sensing*, 61, 11-22.

Taylor, G., J. Li, D. Kidner, and M. Ware (2005). "Surface modelling for GPS satellite visibility." *Proceedings of the 5th International Workshop on Web and Wireless Geographical Information Systems, Lecture Notes in Computer Science 3833*, Eds. K.-J. Li and C. Vangenot. Lausanne, Switzerland, 15-16 December, 281-295.

US Air Force Space Command (2007). "Performance Reports." *GPS Operations Center*. [On-line] 18 January 2008. http://gps.afspc.af.mil/gpsoc/performance_reports.aspx

US Coast Guard GPS (n.d.). "General Information." *USGC Navigation Center*. [On-line] 18 January 2008 <http://www.navcen.uscg.gov/gps/default.htm>

Vosselman, G. (2000). "Slope based filtering of laser altimetry data." *International Archives of Photogrammetry and Remote Sensing*, 33(3/2).

Wehr, A., and U. Lohr (1999). "Airborne laser scanning - an introduction and overview." *ISPRS Journal of Photogrammetry & Remote Sensing*, 54, 68-82.

Wells, D., et al. (1986). *Guide to GPS Positioning*, Department of Geodesy and Geomatics Engineering Lecture Note No. 58 ed., 291 pp., University of New Brunswick, Fredericton, New Brunswick.

Zhang, K., S. Chen, D. Whitman, M. Shyu, J. Yan, and C. Zhang (2003). "A progressive morphological filter for removing nonground measurements from airborne LIDAR data." *IEEE Transactions on Geoscience and Remote Sensing*, 41(4), 872-882.

APPENDIX I: COORDINATE CONVERSION

A.1. Conversion Constants

The constants and formula coefficients used in the coordinate transformations from UTM to geodetic and from geodetic to ECEF are listed in Tables I.1 and I.2.

Table I.1
Conversion constants.

Constant	Symbol	Value
Semi-major axis	a	6378137 m
Semi-minor axis	b	6356752.314 m
Scale along central meridian	k_0	0.9996
Eccentricity of Earth's elliptical cross-section	e	$\sqrt{1 - b^2/a^2}$

Table I.2
Eccentricity coefficients.

Symbol	Formula
e^2	$e^2/(1-e^2)$
e_1	$(1 - \sqrt{1 - e^2}) / (1 + \sqrt{1 - e^2})$
x_1	$3 e_1/2 - 27 e_1^3/32$
x_2	$21 e_1^2/16 - 55 e_1^4/32$
x_3	$151 e_1^3/96$
x_4	$1097 e_1^4/512$

A.2. UTM to Geodetic Coordinate Conversion

Variables: cN, Northing value of point; cE, Easting value of point; UTM zone.

$$K = cN/k_0$$

$$\mu = K / a / (1 - e^2/4 - 3e^4/64 - 5e^6/256)$$

$$L = \mu + x_1 \sin(2\mu) + x_2 \sin(4\mu) + x_3 \sin(6\mu) + x_4 \sin(8\mu)$$

$$M = e'^2 \cos^2(L)$$

$$N = \tan^2(L)$$

$$O = a / \sqrt{1 - e^2 \sin^2(L)}$$

$$P = a(1 - e^2) / \sqrt[3]{1 - e^2 \sin^2(L)}$$

$$Q = (500000 - cE) / O / k_0$$

$$R = O \tan(L) / P$$

$$S = Q^2/2$$

$$T = (5 + 3N + 10M - 4M^2 - 9e'^2) Q^4 / 24$$

$$U = (61 + 90N + 298M + 45N^2 - 3M^2 - 252e'^2) Q^6 / 720$$

$$\phi = 180(L - R(S+T+U)) / \pi$$

$$W = (1 + 2N + M) Q^3 / 6$$

$$X = (5 + 28N - 2M - 3M^2 + 24N^2 + 8e'^2) Q^5 / 120$$

$$Y = (Q - W + X) / \cos(L)$$

$$Z = 6 \text{ zone} - 183$$

$$\lambda = Z - 180Y / \pi$$

A.3. Geodetic to ECEF Coordinate Conversion

Variables: φ , point latitude; λ , point longitude; h , point altitude.

$$\chi = \sqrt{1 - e^2 \sin^2 \varphi}$$

$$x = (a / \chi + h) \cos \varphi \cos \lambda$$

$$y = (a / \chi + h) \cos \varphi \sin \lambda$$

$$z = (a(1 - e^2) / \chi + h) \sin \varphi$$

APPENDIX II: ADDITIONAL ASSESSMENT

Table II.1 lists the satellites visible at selected epochs at the five assessment target locations described in Section 7.1 and illustrated in Figure 7.1 and Table 7.1. The points are: A, parking lot beside commercial buildings; B, roof of residential building; C, Head Hall A-level parking lot; D, sidewalk in front of Sir Howard Douglas Hall; E, parking lot beside Provincial Archives. Obstruction levels for the four direction sectors are given in brackets, with each level corresponding to a 15° elevation range.

Table II.1
Satellites visible at target locations.

Time	Epoch	Satellites visible at target locations				
		A (1,1,1,0)	B (0,0,0,0)	C (1,1,2,2)	D (1,2,4,3)	E (0,2,2,2)
00:00	24	09 15 18 21 24 26	09 15 18 21 22 24 26	15 18 21 26	15 18 26	15 18 21 26
01:30	30	09 14 08 22	09 14 15 18 21 22	09 18 22	09 18 22	09 15 18 22
04:30	42	01 05 12 14 22 30 31	01 05 12 14 22 30 31	01 05 12 14 30 31	01 05 12 14 30	01 05 12 14 30 31
06:00	48	01 06 14 16 20 31	01 06 14 16 20 23 30 31	01 20 31	01 31	01 30 31
09:00	60	03 13 16 19 23 25	03 13 16 19 23 25 27	03 13 16 23 25	03 13 06	03 13 16 23
10:30	66	03 13 19 25 27	03 08 13 16 19 25 27	03 19 25 27	03 16 19 25 27	03 19 25 27
12:00	72	08 11 19 25 27	03 08 11 19 25 27 28	08 19 25 27	08 19	03 08 19 25 27

The following Figures II.1 – II.7 are the *GPSView* sky view plots for each of the epochs listed in Table II.1 with the obstruction profiles for target points A, C, D, and E overlaid as in Figures 7.5 – 7.7 (in blue, green, red, and yellow, respectively). Each ring represents a 15° elevation interval, and the corresponding obstruction level. The outermost ring, shaded grey, is Level 0, 0°-15°; satellites in this ring are below the elevation mask. The innermost circle is Level 5, 75°-90°. The obstruction sectors start at northeast and are counted counter-clockwise (northeast, northwest, southwest, southeast). As noted in Section 7.4, there may be slight differences in satellite elevation as determined by the algorithm and as determined by *GPSView* that could result in misidentification of visible satellites. However, the satellites visible above the obstruction profile for each target point should otherwise correspond to the algorithm-determined visible satellites listed in Table II.1.




Figure II.1
Sky view plot, 00:00, 21 January 2008.




Figure II.2
Sky view plot, 01:30, 21 January 2008.




Figure II.3
Sky view plot, 04:30, 21 January 2008.




Figure II.4
Sky view plot, 06:00, 21 January 2008.




Figure II.5
Sky view plot, 09:00, 21 January 2008.




Figure II.6
Sky view plot, 10:30, 21 January 2008.

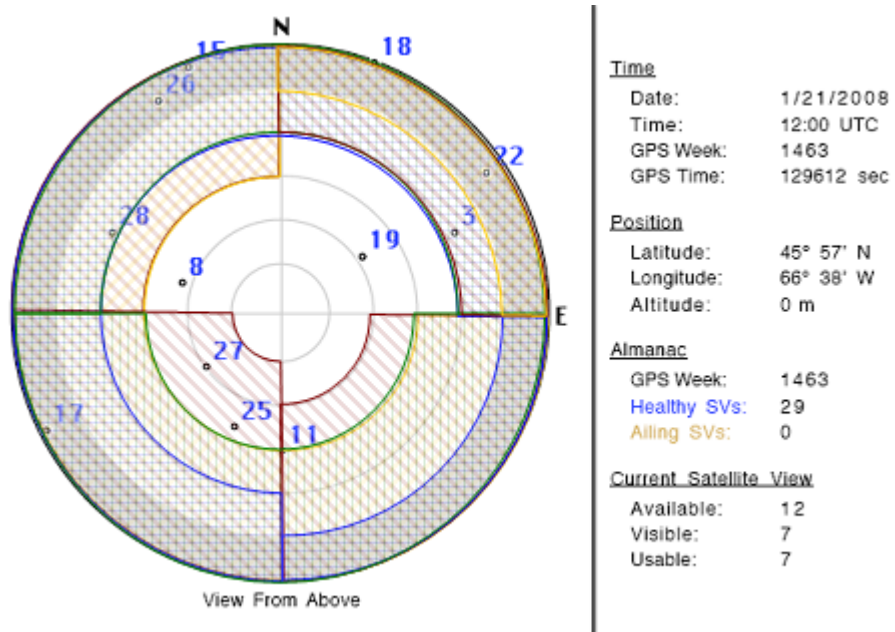


Figure II.7
Sky view plot, 12:00, 21 January 2008.

CURRICULUM VITAE

Candidate's full name: Krista Marija Amolins

Universities attended:

Queen's University, B.Sc.E., 2004

Queen's University, B.A., 2005

Publications:

Amolins, K., Y. Zhang, and P. Dare (2007). "Wavelet base image fusion techniques – An introduction, review and comparison." *ISPRS Journal of Photogrammetry & Remote Sensing*, 62, 249-263.

Conference Presentations:

Amolins, K., Y. Zhang, and P. Dare (2008). "Classification of LiDAR data using standard deviation of elevation and characteristic point features." Proceedings of the 2008 IEEE International Geoscience & Remote Sensing Symposium. Boston, Massachusetts, U.S.A., 6-11 July.

Amolins, K., Y. Zhang, and P. Dare (2007). "Applications of wavelet transforms in image fusion." Proceedings of the Urban Remote Sensing Joint Event 2007. Paris, France, 11-13 April.