17

Object Reconstruction Using Close-Range All-Round Digital Photogrammetry for Applications in Industry

Justin Davey

University of Cape Town

1999

Thesis submitted in partial fulfilment of the requirements for the degree of Master of Science in Engineering.

Declaration

I declare that this is my original work and that it has not been submitted for a degree, in this or similar form, to any University.

Justin John Davey

Abstract

Photogrammetry has many inherent advantages in engineering and industrial applications, which include the ability to obtain accurate, non-contact measurements from data rapidly acquired with the object *in situ*. Along with these advantages, digital photogrammetry offers the potential for the automation or semi-automation of many of the conventional photogrammetric procedures, leading to real-time or near real-time measurement capabilities. However, all-round surface measurement of an object usually benefits less from the above advantages of photogrammetry. To obtain the necessary imagery from all sides of the measurement object, real-time processing is nearly impossible, and it becomes difficult to avoid moving the object, thus precluding *in situ* measurement. However, all-round digital photogrammetry and, in particular, the procedure presented here, still offer advantages over other methods of full surface measurement, including rapid, non-contact data acquisition along with the ability to store and reprocess data at a later date.

Conventional or topographic photogrammetry is well-established as a tool for mapping simple terrain surfaces and for acquiring accurate 3-D point data. The complexities of all-round photogrammetry make many of the standard photogrammetric methods all but redundant. The work presented in this thesis was aimed at the development of a reliable method of obtaining complete surface data of an object with non-topographic, all-round, close-range digital photogrammetry. A method was developed to improve the integrity of the data, and possibilities for the presentation and visualisation of the data were explored. The potential for automation was considered important, as was the need to keep the overall time required to a minimum. A measurement system was developed to take as input an object, and produce as output an accurate, representative point cloud, allowing for the reconstruction of the surface. This system included the following procedures:

- a novel technique of achieving high-accuracy system pre-calibration using a cubic control frame and fixed camera stations,
- separate image capture for the control frame and the object,

- surface sub-division and all-round step-wise image matching to produce a comprehensive 3-D data set,
- point cloud refinement, and
- surface reconstruction by separate surface generation.

The development and reliability of these new approaches is discussed and investigated, and the results of various test procedures are presented. The technique of system pre-calibration involved the use of a mechanical device - a rotary table - to impart precisely repeatable rotations to the control frame and, separately, the object. The actual repeatability precision was tested and excellent results achieved, with standard deviations for the resected camera station coordinates of between 0.05 and 0.5 mm. In a detailed test case, actual rotations differed from the desired rotations by an average of 0.7″ with a standard deviation of less than 2′. The image matching for the test case, from a set of forty-eight images, achieved a satsifactory final accuracy, comparable to that achieved in other similar work.

The meaningful reconstruction of surfaces presented problems, although an acceptable rendering was achieved, and a thorough survey of current commercially available software failed to produce a package capable of all-round modelling from random 3-D data.

The final analysis of the results indicated that digital photogrammetry, and this method in particular, are highly suited to accurate all-round surface measurement. The potential for automation - and, therefore, for near real-time results - of the method in the stages of image acquisition and processing, calibration, image matching and data visualisation is great. The method thus lends itself to industrial applications. However, the need for a robust and rapid method of surface reconstruction needs to be fulfilled.

Acknowledgements

During the time spent working on this dissertation, I was assisted greatly by many people:

My supervisor, Prof. Dr. Heinz Rüther, Head of the Department of Geomatics at UCT, deserves special mention for all the constructive criticism and advice he gave to me and for the expert knowledge he shared with me.

Dr. Julian Smit was also particularly helpful in familiarising me with digital photogrammetric software and methods.

Mrs. Val Atkinson, Ms. Sue Binedell, Ms. Andrea Court and Mr. Kari Laatikainen, along with Mr. Sidney Smith and Mr. Michael Haywood, have my sincere appreciation for all their assistance.

Fellow postgraduate students in the Department over the last two years – Glynnis Barodien, Ulrike Brüßler, Simon Crone, Malcolm Dingle, Nathan Geffen, Justin Hill, Simon Hull, Bonita Kleyn, Siddique Motala, Mbinji Mufalo, Yasson Ouma, Terry Richards, Lani Roux, Ross Rozendaal, Simon Taylor and Robert Zimba - were also tremendously helpful and encouraging, and my thanks go to them.

This research would not have been possible without the generous financial assistance of the Foundation for Research Development (FRD). Their support is much appreciated.

Lastly, to my family and Tessa Davidson, thanks for your support and encouragement.

Table of Contents

List of Figures

List of Tables

List of Symbols

c	principal distance
g	grey-scale value
h	number of columns in a digital image
i	variable integer
j	variable integer
N	normal vector
p, P	point
r	terms of the rotation matrix
R	rotation matrix
s	scale factor
s_x	pixel dimension in x-direction
s_y	pixel dimension in y-direction
v	number of rows in a digital image
V	vector
x	x-coordinate in image space
x_p	x-coordinate of principal point
x_{pix}	x-coordinate in pixel coordinate system
X	x-coordinate in object space
X_c	x-coordinate of perspective centre
y	y-coordinate in image space
y_p	y-coordinate of principal point
y_{pix}	y-coordinate in pixel coordinate system
Y	y-coordinate in object space
Y_c	y-coordinate of perspective centre
Z	z-coordinate in object space
Z_c	z-coordinate of perspective centre
α	angle
θ	incidence angle
κ	rotation about the Y-axis
ϕ	rotation about the Z-axis
ω	rotation about the X-axis

1. Introduction

1.1 RESEARCH OBJECTIVE AND SCOPE

The research undertaken had as its objective the investigation and development of a system for 3-D object reconstruction by means of all-round close-range digital photogrammetry. The term 'close-range' has no precise definition, but may be understood as referring to applications where the camera-to-object distance is of the order of a few metres or less. Close-range photogrammetry differs from conventional topographic photogrammetry in a number of important areas. These include the typical object-to-camera range, as mentioned above, and the use of strongly convergent (as opposed to parallel) photography, therefore precluding the use of stereoplotters and stereocomparators. Close-range photogrammetry is also referred to as non-topographic photogrammetry and has numerous possible applications in industry, including quality control and reverse engineering.

This research was of an investigative nature, in the form of an indepth feasibility study of the all-round measurement methodology which was developed. The entire digital photogrammetric procedure, from network design, through image capture and image processing, to final data presentation, was explored and the varying requirements and demands of the objective were taken into account.

The vastness of the topic precluded a definitive solution and the research was limited to a general exploration of possibilities, although several developments of some specific aspects of the procedure are presented. The system developed here by no means attempts to obtain real-time processing. Rather, the foundations for a complete digital photogrammetric procedure are laid out, with the emphasis placed on accuracy, time-efficiency and the potential for automation.

1.2 THESIS OVERVIEW

This thesis comprises seven chapters:

Chapter 1 provides a general introduction to the research undertaken and to photogrammetry itself.

Chapter 2 presents basic photogrammetric theory, upon which the principles of the subsequent work are founded.

Chapter 3 discusses digital image processing techniques as used in digital photogrammetry.

Chapter 4 details the photogrammetric procedure followed in this research, from the design of the network, through the stages of image capture and processing, to the presentation of the acquired data.

Chapter 5 reports on the testing and results of the detailed test case which was explored.

Chapter 6 presents a discussion and conclusions based on the major findings of the research.

Chapter 7 concerns various recommendations which have been made regarding this research and future work in this field.

1.3 INTRODUCTION TO PHOTOGRAMMETRY

1.3.1 Definition of Photogrammetry

Two similar definitions of photogrammetry are as follows:

"Photogrammetry is the art, science, and technology of obtaining reliable quantitative information about physical objects and the environment through the process of recording, measuring and interpreting photographic images and patterns of radiant imagery derived from sensor systems." -- Karara (1989).

"Photogrammetry is the science, and art, of determining the size and shape of objects as a consequence of analysing images recorded on film or electronic media." -- Atkinson (1996).

These definitions, while by no means complete, summarise well the objectives of the photogrammetrist, namely the determination of a spatial quantity using imagery.

1.3.2 Digital Photogrammetry

Digital photogrammetry may be defined as the utilisation of digital image capture and digital image processing techniques to obtain spatial data. While conventional images are analysed with traditional operator-controlled equipment, digital images lend themselves to automated processing and analysis. Automation, in turn, is suited to applications where the time from image capture to final data delivery is very short. Indeed, many such applications require near real-time processing, and digital photogrammetry offers this possibility. The requirement for high-speed processing is most often the case in industrial applications - such as a production line - where immediate responses are required and thus very limited time is available for the generation of results.

1.3.3 Industrial Photogrammetry

Generally the engineer in industry is unaware of the potential of photogrammetry to perform useful measurement and quality control functions in an industrial environment. Several inherent characteristics of photogrammetry offer great advantages over other methods of industrial measurement. Some of these advantages are listed below.

- Photogrammetry offers remote, non-contact measurement.
- Objects of virtually any size can be measured photogrammetrically.
- Data can be gathered rapidly, reducing downtime.
- Dynamic events can be measured in real-time.
- Measurements can be made in hazardous environments.
- Photographic data can be stored for future analysis.
- High accuracies are achievable.

Nowadays, photogrammetry is being used in increasing numbers of industrial and engineering applications, with great potential for many further applications in these sectors. Photogrammetric techniques may be employed in an engineering or industrial environment in the areas of research, planning, production engineering, manufacture, testing, monitoring, repair and reconstruction. (Meyer, 1973; Adams, 1989). Typical

areas of application include the aerospace, automobile, and manufacturing industries. Photogrammetric techniques may be used to control robotic machinery, measure deformations in engineering structures, retrieve spatial data in reverse engineering applications, and so on.

Photogrammetric data can be obtained by numerous methods, including the use of metric and non-metric cameras, X-ray equipment, scanning electron microscopes and digital photography. To quote Karara (1989), "anything that can be photographed can be measured." Digital cameras were used exclusively in this research, as their potential in automated applications and real-time data production was considered of paramount importance.

2. Basic Photogrammetric Theory

2.1 OBJECT AND IMAGE SPACE

The perspective projection is the projection of a point in three-dimensional space (object space), through the perspective centre of the camera onto a plane (image space). The ray from an object point (X_i, Y_i, Z_i), passes through the optical centre of the lens, represented by the perspective centre (X_c, Y_c, Z_c), and forms a corresponding image point (x_i, y_i). The perpendicular distance between the image plane and the perspective centre is the principal distance, c. The principal point (x_p, y_p) is the projection of the perspective centre onto the image plane. The perspective projection onto the positive image plane is shown in Figure 1 below.

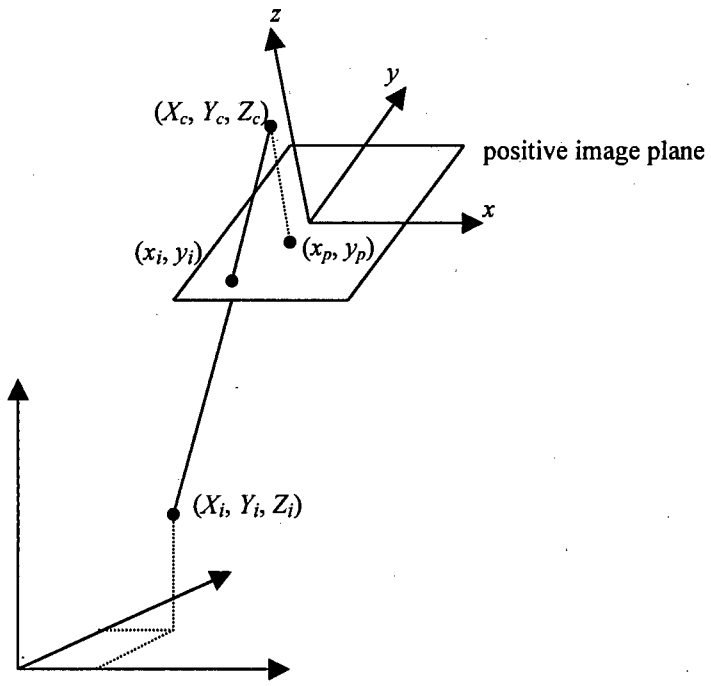

Figure 1. The perspective projection of a point onto the image plane.

Photogrammetric theory assumes that, firstly, all image points lie in the image plane and secondly, any point in object space, the perspective centre and the corresponding image point are collinear. To transform a point (x_i, y_i) in image space into object space coordinates (X_i, Y_i, Z_i), it is necessary to perform three rotations, three translations and a scaling. Mathematically,

$$\begin{pmatrix} x_i - x_p \\ y_i - y_p \\ c \end{pmatrix} = sR \begin{pmatrix} X_i - X_c \\ Y_i - Y_c \\ Z_i - Z_c \end{pmatrix} \qquad \textbf{Equation 1}$$

where R is an orthogonal rotation matrix containing nine elements, which are functions of the three rotations ω about X, κ about Y and ϕ about the Z axis. R may be calculated as the product of the three individual matrices R_X, R_Y and R_Z and is dependant on the order of multiplication of these matrices. In a right-hand fixed axis system, for the rotation sequence ω, κ, ϕ, the elements of R are:

$$\begin{aligned} r_{11} &= \cos\kappa \cos\phi \\ r_{12} &= -\cos\kappa \sin\phi \\ r_{13} &= \sin\kappa \\ r_{21} &= \cos\omega \sin\phi + \sin\omega \sin\kappa \cos\phi \\ r_{22} &= \cos\omega \cos\phi - \sin\omega \sin\kappa \sin\phi \\ r_{23} &= -\sin\omega \cos\kappa \\ r_{31} &= \sin\omega \sin\phi - \cos\omega \sin\kappa \cos\phi \\ r_{32} &= \sin\omega \cos\phi + \cos\omega \sin\kappa \sin\phi \\ r_{33} &= \cos\omega \cos\kappa \end{aligned} \qquad \textbf{Equation 2}$$

and

$$R_{(\omega,\kappa,\phi)} = \begin{pmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{pmatrix} \qquad \textbf{Equation 3}$$

2.2 CAMERA CALIBRATION AND ORIENTATION

Camera calibration and orientation are critical to the photogrammetric process. Before unknown object space coordinates can be measured, knowledge of the interior and exterior orientations of at least two cameras is necessary. A camera network may be said to be calibrated when the nine parameters - x_p, y_p, c, X_c, Y_c, Z_c, ω, κ and ϕ - of the interior and exterior orientations for each camera station are known.

Interior orientation parameters include c, the principal distance of the camera, x_p and y_p, the location of the principal point in image space, and additional parameters which account for effects such as lens distortion, or deviation from collinearity.

Exterior orientation refers to the position and attitude of a camera in object space. The six parameters of exterior orientation are X_c, Y_c and Z_c, the location of the perspectitve centre, and the three rotations ω, κ and ϕ about the X, Y and Z axes.

Relative orientation describes the spatial relationship between two images in object space. The six relative orientation parameters for two images are simply the differences between the six exterior orientation parameters of the first image and the six corresponding exterior orientations of the second image. Thus the relative orientation of two images is known if their respective exterior orientations are known. Knowledge of any five of the relative orientation parameters defines the relative orientation.

2.3 THE THEORY OF COLLINEARITY

The theory of collinearity is one of the fundamental theories of photogrammetry, and states that any object point, the perspective centre of the camera and the corresponding image point are collinear. The collinearity equations, Equations 4 and 5, are derived from the transformation equation, Equation 1, as follows:

$$x_i - x_p = c\frac{r_{11}(X_i - X_c) + r_{12}(Y_i - Y_c) + r_{13}(Z_i - Z_c)}{r_{31}(X_i - X_c) + r_{32}(Y_i - Y_c) + r_{33}(Z_i - Z_c)} \qquad \textbf{Equation 4}$$

$$y_i - y_p = c\frac{r_{21}(X_i - X_c) + r_{22}(Y_i - Y_c) + r_{23}(Z_i - Z_c)}{r_{31}(X_i - X_c) + r_{32}(Y_i - Y_c) + r_{33}(Z_i - Z_c)} \qquad \textbf{Equation 5}$$

Here, the r_{ij} terms are the nine elements of R as given in Equation 2 above. Equations 4 and 5 could include additional terms, Δx and Δy, to account for image plane unflatness, lens distortion, and any deviations from collinearity. The solution of the collinearity equations requires that at least four and a half image space coordinates, corresponding to five known non-coplanar control points, are known. This allows the nine linearly independent elements, i.e. the r_{ij} terms, of the equations to be determined. The collinearity equations may be used to solve for either:

(i) the camera orientation parameters, when good approximations to these unknowns are available and provided sufficient control point coordinates are known, or

(ii) unknown object space coordinates, when good approximations to these unknowns are available and provided the interior and exterior orientation parameters are known.

2.4 THE BUNDLE ADJUSTMENT

The bundle adjustment is an elegant method of combining (i) and (ii) in 2.3 above to simultaneously solve for the unknown orientation parameters, object space coordinates and for additional parameters such as lens distortion. Van der Vlugt (1991), author of the UCT Geomatics Department's bundle adjustment program *Photonet*, states that, for simultaneous solution of the bundle adjustment, the following conditions must be met:

- each unknown point must be visible in at least two images,
- for camera orientation, a minimum of five well-distributed known points must be visible on each image, and
- full coordinates for at least two control points, and an ordinate from a third point, must be known to define the network position and orientation.

Increasing the redundancy, i.e. increasing the number of images and/or control points, helps to smooth the effect of inevitable random observation errors and improves the accuracy of the bundle adjustment.

2.5 THE DIRECT LINEAR TRANSFORMATION

The direct linear transformation (DLT) may be used as an alternative to the bundle adjustment to determine the transformation parameters between image and object space coordinates. The main advantage of the DLT is that no prior knowledge of either exterior or interior orientation, or of initial estimates for these unknowns, is required. Indeed, the DLT has a rapid convergence, even given inaccurate initial values. The DLT equations, Equations 6 and 7, relate measured image coordinates x_i and y_i to measured object space coordinates X_i, Y_i and Z_i by means of the eleven DLT transformation parameters.

$$x_i + \Delta x_i = \frac{b_{11}X_i + b_{12}Y_i + b_{13}Z_i + b_{14}}{b_{31}X_i + b_{32}Y_i + b_{33}Z_i + 1} \qquad \textbf{Equation 6}$$

$$y_i + \Delta y_i = \frac{b_{21}X_i + b_{22}Y_i + b_{23}Z_i + b_{24}}{b_{31}X_i + b_{32}Y_i + b_{33}Z_i + 1} \qquad \textbf{Equation 7}$$

For the purposes of camera calibration, five and a half image coordinates, relating to six known control points, are required to perform the necessary space resection and solve for the unknown transformation parameters in the DLT. A least squares adjustment can be performed when redundant information is available. Using the DLT to perform a resection requires a good spatial distribution of control points for high accuracy. Knowing the DLT parameters also enables the coordinates of unknown points in object space to be determined by performing a space intersection using two or more images.

The DLT is often used to obtain approximate interior and exterior orientation parameters as provisional values for the bundle adjustment.

3. Digital Image Processing

Digital photogrammetric procedures differ from those of traditional photogrammetry in that the images analysed and measured are in digital rather than analog or hardcopy format. Digital images therefore lend themselves to analysis by computer.

Digital images consist of discrete picture elements - pixels - of a uniform intensity. The digital images used in this research are all grey-scale images in which brightness, but not colour, is represented. In a grey-scale image, by assigning eight bits to each pixel, 256 (i.e. 2^8) levels of brightness are possible, ranging from 0 (black) to 255 (white). A digital camera sensor consists of an array of pixels. A typical CCD array might comprise 512 rows and 512 columns, producing a 256 kB raw image. The camera used for some of this research, a Kodak DCS 420, has a sensor array of 1012 rows and 1524 columns, resulting in a raw image size of 1.5 MB. It is not hard to see that digital imaging often involves striking a balance between the low cost, high processing rate and poor resolution of a small CCD and the high cost, low processing rate and good resolution of a large CCD.

In a digital image, pixel centre locations are given according to their row and column numbers in the image. Pixel (0,0) is the pixel in the top left corner of the image, with row numbers increasing downwards and column numbers increasing to the right. As shown in Figure 2 below, this differs from the metric image coordinate system, as described in Chapter 2.1.

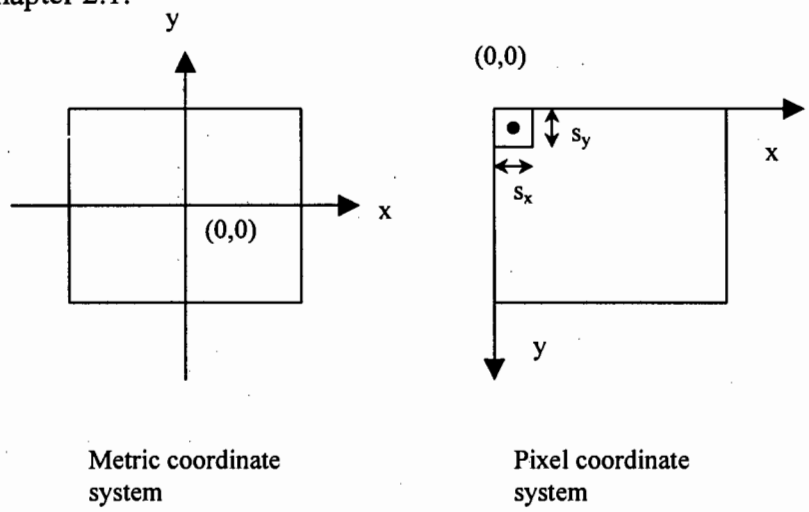

Figure 2. Metric and pixel coordinate systems.

Transformations are required to coordinate pixels in image space, and the equations for these transformations are as follows,

$$x = s_x \left(x_{pix} - \frac{h - 1}{2} \right) \qquad \textbf{Equation 8}$$

$$y = s_y \left(\frac{v - 1}{2} - y_{pix} \right) \qquad \textbf{Equation 9}$$

where x and y are the pixel centre's coordinates in image space, x_{pix} and y_{pix} are in the pixel coordinate system, h and v are the number of columns and rows in the image, and s_x and s_y are the pixel dimensions in the x- and y-directions.

3.1 TARGET DETECTION AND CENTERING

Camera calibration - the determination of the interior and exterior orientation parameters - in digital photogrammetry requires that previously coordinated control points be measured in image space. The coordination of these targets is often required to sub-pixel accuracy in close-range applications. In order that the targets are readily identifiable by human operators and by processing algorithms they are most often made from circular discs of retroreflective material. Due to varying image orientations, the circular targets invariably appear as elliptical shapes in the image.

The measurement of a control point in the image plane is achieved by the following:

1. target detection, usually by means of a thresholding procedure or by edge detection
2. determination of the boundary of the target, usually by means of an edge-following algorithm
3. determination of the centre of the target to sub-pixel accuracy, usually by means of the weighted centre of mass algorithm

These steps are discussed in more detail below.

3.1.1 Thresholding

In order to isolate possible targets from the surrounding pixels, a thresholding procedure may be performed. Provided care has been taken with illumination during the image acquisition, retroreflective targets appearing in the image will generally have the highest grey-scale values of all pixels, i.e. close to 255. The thresholding procedure sets all pixels with a grey-scale value below a specified threshold to zero. The result is that only the bright patches of the image remain, including the targets. A further step may be performed, which involves setting the remaining pixels to a grey-scale value of 255. This has the effect of creating a binary image in which each pixel is either black or white. Otherwise, the grey-scale information of the isolated targets is retained in a weighted centre-of-mass approach.

The threshold value itself can either be set interactively by a computer operator, viewing the thresholded image on a computer screen, or determined by means of empirically derived equations, two of which are given below.

$$g_t = g_{min} + \frac{g_{max} - g_{min}}{4} \qquad \textbf{Equation 10}$$

$$g_t = \frac{g_{min} + g_{mean}}{2} \qquad \textbf{Equation 11}$$

Equations 10 and 11 are most effective when applied to a small window surrounding the target. Here g_t is the threshold value, while g_{max}, g_{min} and g_{mean} are the maximum, minimum and mean greyscale values for the selected portion of the image.

3.1.2 Edge Following

The extent of each target must be determined so that a window can be applied for the subsequent centering algorithm. An edge-following routine is employed for this purpose. This routine tracks the edge pixels in the binary image, thus defining the horizontal and vertical extent of the target in the image.

A popular method of edge following uses the so-called chain-coding algorithm. This algorithm sequentially searches for edge pixels in a clockwise direction. Figure 3

below indicates the starting direction of the next search based upon the direction of the last successful search.

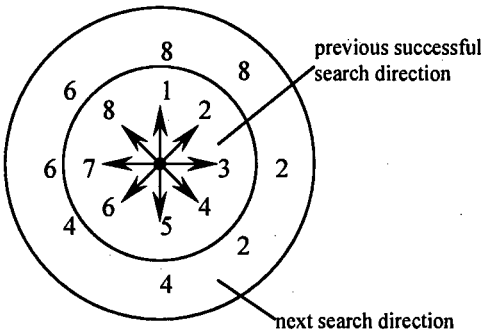

Figure 3. The search sequence used by the chain-coding algorithm.

The starting pixel is found by scanning through each row starting with the first row. Therefore, the direction in which the first pixel is found is always 3. According to the algorithm, the direction in which to begin the next search is 2. If this search direction is unsuccessful, meaning that the pixel scanned is not a target pixel, the search direction proceeds clockwise, until the next edge pixel is located. The sequence continues in this fashion, tracing the edge of the target, and finally returning to the starting pixel. This is shown in Figure 4, which depicts a portion of an imaginary image, including a target, which has been converted into a binary image with a threshold procedure.

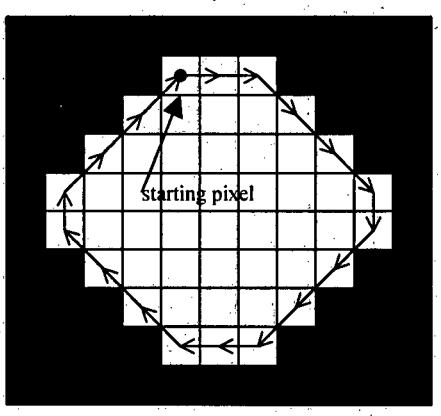

Figure 4. Imaginary target showing sequence of chain-coding.

Beginning at the starting pixel, Table 1 shows the search procedure.

Pixel Number	Last Search Dir.	Starting Search Dir.	Dirs Attempted
1	3	2	2, 3
2	3	2	2, 3
3	3	2	2, 3, 4
4	4	2	2, 3, 4
5	4	2	2, 3, 4
6	4	2	2, 3, 4, 5
7	5	4	4, 5, 6
8	6	4	4, 5, 6
9	6	4	4, 5, 6
10	6	4	4, 5, 6, 7
11	7	6	6, 7
12	7	6	6, 7, 8
13	8	6	6, 7, 8
14	8	6	6, 7, 8
15	8	6	6, 7, 8, 1
16	1	8	8, 1, 2
17	2	8	8, 1, 2
18	2	8	8, 1, 2
1 – finish	2	-	-

Table 1. Search sequence of chain coding of imaginary binary image.

The chain coding algorithm is robust and is capable of successfully tracking complex target perimeters.

3.1.3 Target Centre Determination

The best target centering is achieved using the weighted centre-of-mass algorithm, which is also the method most commonly used. This is called an area-based algorithm. Other methods of determining the centre of a target involve edge-based algorithms which use edge information to approximate the centre, or an area-based method of least squares template matching.

The weighted centre-of-mass algorithm has been shown to produce excellent estimates of the target centre, with an accuracy of up to 0.001 pixels (Hoffa, 1996), although accuracies of around 0.01 pixels are more usual. The algorithm uses the grey-scale values of target pixels above the threshold level.

The equations for the centre location in image coordinates are:

$$x = \frac{\sum_i \sum_j g_{ij} . i}{\sum_i \sum_j g_{ij}} \qquad \textbf{Equation 12}$$

and

$$y = \frac{\sum_i \sum_j g_{ij} \cdot j}{\sum_i \sum_j g_{ij}} \qquad \textbf{Equation 13}$$

Here, g_{ij} is the grey-scale value of the pixel at location (i, j). In cases where the background pixels have not been set to zero by the threshold routine, pixels below the threshold level should be taken as having a grey-scale value of zero for this calculation.

3.2 IMAGE ENHANCEMENT, EDGE DETECTION AND INTEREST OPERATORS

The detection of edges and points of interest in an image is an important component of the photogrammetric procedure, as it reduces the volume of data to be processed while retaining information relating to the structure of the object. Edges define the boundaries of surfaces which constitute an object and interest points identify important features on the surface. Both edges and interest points are marked by steep grey-scale gradients in an image. Recovering the object space coordinates of such edges and surface points on an object allows the object's shape to be recovered. This is the objective of the image matching routine, as described in 3.3.

3.2.1 Image Enhancement

Image enhancement techniques increase the usefulness of digital images by:

- increasing contrast,
- removing random noise, and
- reducing image degradation.

These enhancements to an image may greatly assist in the detection of edges and interest points and, consequently, in the image matching routine. The objective of image enhancement varies according to the application and so enhancement

algorithms are numerous and often *ad hoc*. Ekstrom (1984) covers image enhancement in detail. Some common image enhancement techniques include:

- **Grey-scale modification.** This involves applying a transformation to each image pixel to improve the contrast or dynamic range.
- **Low-pass filtering.** In a typical image the low-frequency component is dominant. Using Fourier analysis, the high-frequency noise is reduced, the low-frequency data is retained and the signal-to-noise ratio is improved. The disadvantage of low-pass filtering is that edge definition is lost and the image may appear blurred.
- **High-pass filtering.** High-pass filtering reduces the low-frequency components, while emphasising the higher frequency components. Since the high frequency components correspond to edges or interest points, the high-pass filter has the effect of sharpening the image but also tends to increase the amount of background noise.
- **Median filtering.** This method is useful for reducing random or so-called salt-and-pepper noise, while retaining edge definition. A window is passed over the image and the current pixel's grey-scale value is replaced by the median value of the pixels in the window.

3.2.2 Edge Detection

An edge in an image appears as a sudden change in pixel brightness. Various operators perform the edge detection in digital images. The Canny edge operator, which uses the first derivative of the Gaussian function to find maxima in the first directional derivative of grey-scale, the Sobel edge operator, which convolves a 3×3 pixel mask with the image to detect maximum grey-scale gradients, and the maximum gradient method, which finds the direction of maximum gradient at every pixel, all locate edges to single pixel accuracy as well as providing edge directions. The moment preserving method (Tabatabai and Mitchell, 1984), used subsequently to any of the above methods, yields edge locations to sub-pixel accuracy. This method of moment preserving calculates the sub-pixel location of an instantaneous edge – the boundary at which the grey-scale changes instantaneously from one steady value to another - using the actual edge information. This is shown in Figure 5.

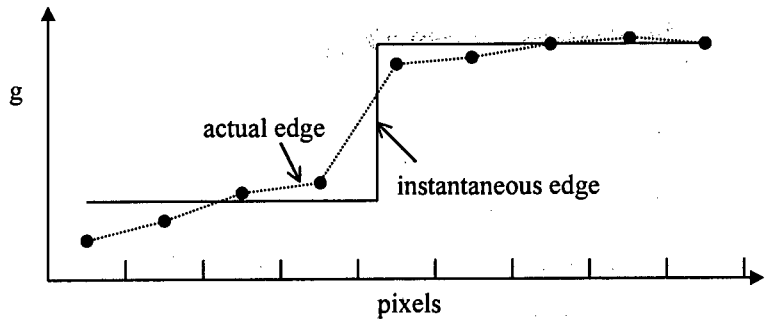

Figure 5. Edge located to sub-pixel accuracy by the moment preserving method.

For an indepth discussion of edge detection methods consult Hoffa (1996) and Smit (1997).

3.2.3 Interest Operators

An interest point is any point on the surface of an object which is readily identifiable, and which contributes to the reconstruction of that object. Various interest operators exist, among them the Förstner interest operator (which is designed to detect points, line intersections, corners and centres of arcs), the Moravec interest operator (which detects regions containing large fluctuations in grey-scale values) and the angle-based interest operator (which is particularly suited to the extraction of corner points). Hoffa (1996) covers interest operators in more detail.

3.3 IMAGE MATCHING

Recovering the object space coordinates of an imaged point requires that the point appears and is measured in at least two images. Provided the interior and exterior orientations are known, this allows a space intersection to be performed which yields the coordinates of the point in three dimensional space. In a typical digital photogrammetric application, many thousands of points are required to be coordinated. Therefore, an automated procedure to identify conjugate points in pairs or sets of images is virtually essential to the digital photogrammetric process. Such automation is achieved with image matching routines.

According to Schenk (1996) the problem of image matching can be summarised in the following four steps:

1. Identify an interest point or feature in the template image.

2. Find its corresponding (conjugate) entity in the search image(s).
3. Calculate the 3-D coordinates of the matched entity in object space.
4. Assess the quality of the match.

The second step is the hardest to solve. A number of approaches to solving this problem have been taken, including

Area-based Matching. A commonly used technique is that of area-based least squares matching which compares the grey-scale values of a small (usually square, e.g. 9 × 9 pixels) patch of pixels with a reference patch. The reference patch is matched to the search patch by a transformation, including two scale factors, a rotation, a shear angle and radiometric shift corrections. The similarity of the patches is measured by correlation or least-squares techniques. The match is successful when the grey-scale differences between these two patches are minimised. A recently developed area-based approach is that of multi-photo geometrically constrained (MPGC) matching. This method utilises the interior and exterior orientation parameters to geometrically constrain the search for matching patches in sets of multiple search images. Baltsavias (1991) gives a thorough treatment of MPGC matching.

Feature-based Matching. This approach uses edges or features extracted from original imagery are compared to determine conjugate features. The similarity of the features is measured with a cost function. Feature-based matching is used primarily in computer vision applications.

Symbolic Matching. This is a method of comparing symbolic descriptions of features, rather than features themselves. The similarity is measured by means of a cost function.

4. The Photogrammetric Procedure

A photogrammetric project involves various stages, all of which are vital to the successful outcome of the project. These stages include planning (which involves network and control field design), photography and processing (which involves the system calibration, bundle adjustment, image matching, etc.). This chapter covers the following seven stages of the photogrammetric procedure, detailing the approach taken in each in this research:

1. Network design
2. Control
3. Image capture and system calibration
4. Bundle adjustment
5. Image matching and surface mensuration
6. Point cloud refinement
7. Surface modelling, object reconstruction and visualisation

The difference between the all-round close-range approach taken in this work and the conventional topographic approach is highlighted. For the purposes of this thesis, conventional photogrammetry is taken to mean photogrammetry in which the photography occurs to one side of the object, usually from camera locations which are located approximately in a plane parallel to the object.

4.1 NETWORK DESIGN

The position and orientation of a camera relative to the object and relative to the other camera locations in a network is critical to the success of a photogrammetric survey. A well designed camera network, with strong geometry, in conjunction with high quality imagery, makes accurate and reliable data much easier to obtain.

Mason (1994) details ten constraints and objectives in camera network design. These are summarised as follows:

- **Image Scale Constraint.** Measurement precision is approximately proportional to image scale. For a desired precision, a constraint is placed upon the camera-to-object distance.

- **Resolution Constraint.** Precise measurement requires that targets and features in the image have a sufficiently high resolution. This imposes a constraint either upon the size of the artificial targets/control points to be used, upon the camera-to-object distance, or both.

- **Workspace Constraint.** The geometry and size of the working area in which the photography is to be performed, along with limitations on the possible camera positions, constrains the network design.

- **Depth-of-field Constraint.** It is important, particularly for matching but to a lesser extent also for target identification and centering, that sharp focus, throughout the depth of the image, is achieved.

- **Incidence Angle Constraint.** The incidence angle, θ, of the ray, r, from a feature on the measured object to the camera must not be too small. The minimum acceptable incidence angle depends on the geometry, texture and radiometric properties of the feature, as well as on the lighting conditions. (Figure 6).

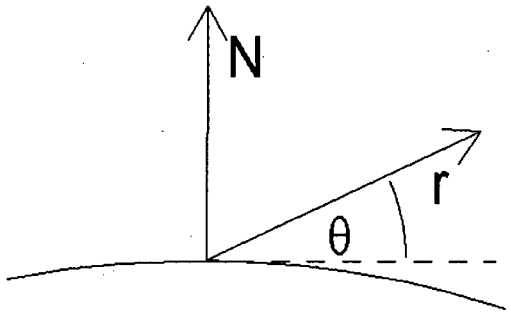

Figure 6. Incidence angle.

- **Number and Distribution of Image Points Constraint.** For the purposes of restitution, a suitable distribution and density of control points should be visible from each station.

- **Illumination Constraint.** Illumination has a marked effect on the accuracy of a photogrammetric survey. This should be taken into account when designing a network.

- **Contribution to Intersection Angles.** Each camera station should have strong geometry relative to the other stations in the network to ensure good intersections of rays. This usually entails a reasonable separation between camera stations.

- **Field of View.** Ideally, the entire target field should appear in each station's field of view. This simplifies the network, and produces stronger geometry.

- **Visibility.** Occlusions should be kept to a minimum.

The above criteria give the photogrammetrist the information required to design a strong network and define a possible camera placement area, as in Figure 7 (after Mason).

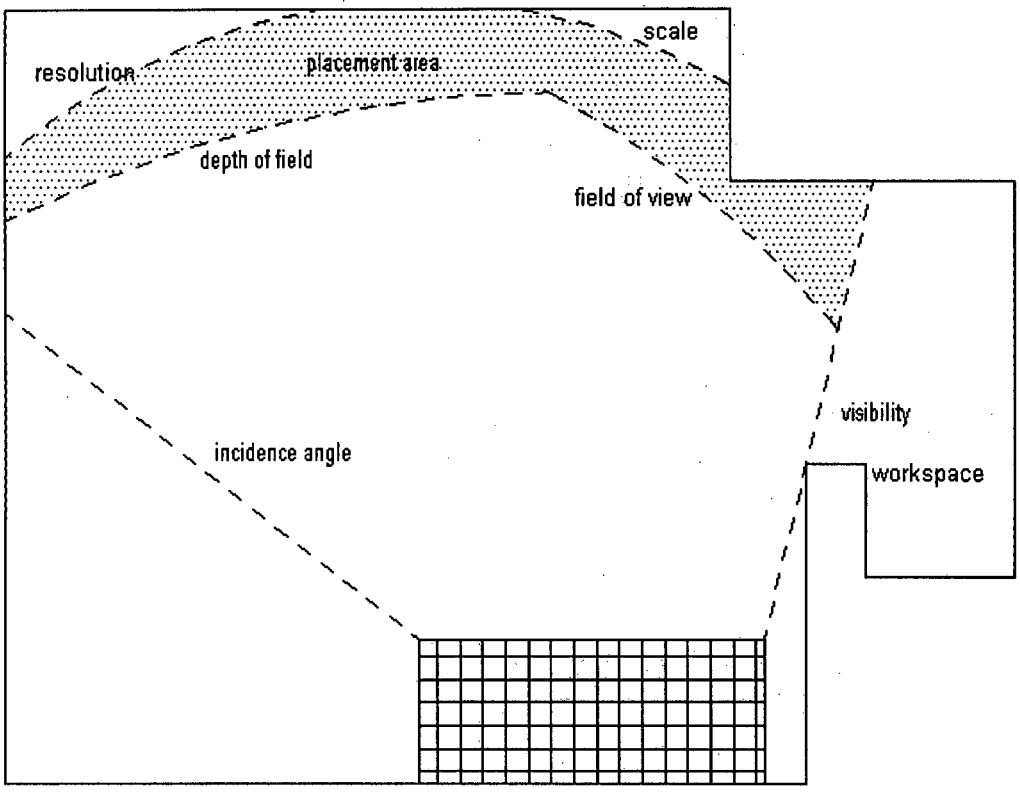

Figure 7. Constraints imposed upon a network, and resultant placement area.

In the case of all-round photogrammetry, the network is considerably more complicated than in the case of conventional topographic photogrammetry. This complication is brought about because, whereas in the conventional case the object to be measured is usually a simple topographic surface requiring photography from one side only, in the all-round case the network has to cover the object from all aspects.

This is certainly one of the chief difficulties involved in all-round close-range photogrammetry.

A typical conventional network is shown in Figure 8. Notice how the cameras are positioned only on one side of the object. Compare this with the all-round network, shown in Figure 9, where cameras are positioned all around the object. (For the sake of clarity, only nine of the required stations are shown in Figure 9.)

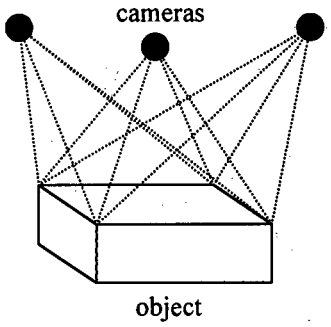

Figure 8. Conventional network.

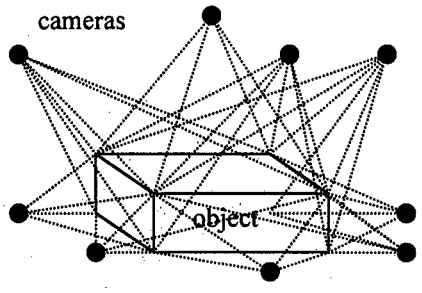

Figure 9. All-round network.

4.2 CONTROL

In addition to their effect on network design, the differences between conventional and all-round photogrammetry influence the design of control fields. In the conventional case, cameras are usually located exclusively to one side of the object. This means that when a control field is designed for a conventional photogrammetric project, targets need only face in one direction.

The all-round case is more complex. Images need to be captured from all sides of the object, and these images must contain sufficient control to allow the interior and exterior orientations of the camera to be determined with acceptable accuracy. To allow for imaging from all sides, various possibilities of control frame design were considered, including using spherical control points connected by thin rigid rods, and using a cubic or spherical control frame. However, the occlusion of portions of the object by the control frame is more of a problem in all-round than in conventional photogrammetry, and it is nearly impossible to design a control frame which produces no occlusions. A method of network calibration was developed by the author, using a cubic control frame, which solved the problem of occlusions.

4.2.1 The Cubic All-Round Control Frame

In general, the control field should have points that are well distributed in three dimensions. This improves the accuracy of the interior orientation parameters which are recovered from measurement of the control field in the image. For the purposes of this research, a control frame was built in the form of a cubic framework. This frame is shown in Figures 10 and 11.

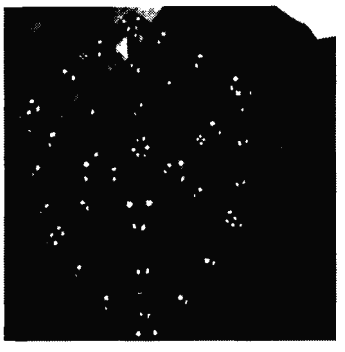

Figure 10. Faces 1, 3 and 4.

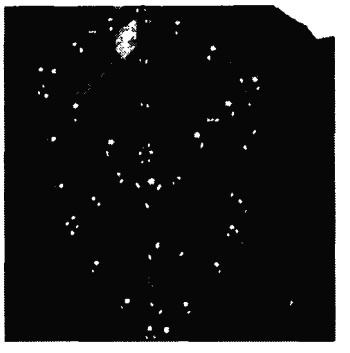

Figure 11. Faces 3, 4 and 6.

180 retro-reflective targets were attached to the frame, with an average of thirty points per side. (See Appendix 2). The faces were numbered from 1 - the top face - through 6 - the bottom. Faces 2 through 5 were denoted by two patches of tiny retro-refelective targets on each of these faces, corresponding to the face number. The numbering patches for Faces 3 and 4 are visible in Figures 10 and 11. This face numbering was devised to assist the photogrammetrist in keeping track of the orientation of the control frame in each image. A mounting rod of circular cross-section - visible in the images above - was attached to Face 6 to allow the frame to be held firmly in the chuck of a rotary table. The reasons for this are discussed in 4.3 below.

The unknown target coordinates were determined by a free network adjustment. Firstly, an ITC CCD camera, which was to be used in the latter stages of the project, was fixed at a set aperture and focal length. The interior orientation parameters for this camera were obtained by performing a bundle adjustment using a known control frame. Secondly, the uncalibrated frame targets were provisionally coordinated with a tape measure, and vernier calipers were used to measure distances between various points to scale the network. The control frame was coordinated in a right-hand

cartesian system, with the point (1000, 1000, 1000) provisionally placed at the corner of faces 2, 5, and 6.

Fifty images were taken of the frame with the calibrated CCD camera. The control points were extracted with a threshold routine and identified in each image. Provisional exterior and interior orientation parameters were obtained by performing a DLT. The provisional control point coordinates were improved by performing a free network adjustment. The bundle adjustment was performed with 5389 degrees of freedom. The control point errors in object space from this bundle adjustment were as follows:

Quantity	Value (mm)
A posteriori standard deviation of unit weight, σ_0	0.0165
RMS error, σX	0.774
RMS error, σY	0.807
RMS error, σZ	0.753

Table 2. First bundle adjustment results.

The results presented in Table 2 show that, using the provisional control point coordinates, the control point errors in object space were sub-millimetre and that the observation standard deviation in image space was 16.5 microns. The provisional control points were updated with the values from the bundle adjustment. Poor observations were removed from the adjustment, decreasing the degrees of freedom to 4838. The interior orientation was held fixed and the network adjustment was performed again with the updated target object space coordinates, yielding errors as follows:

Quantity	Value (mm)
A posteriori standard deviation of unit weight, σ_0	0.0045
RMS error, σX	0.230
RMS error, σY	0.256
RMS error, σZ	0.225

Table 3. Second bundle adjustment results.

The second bundle adjustment improved the control point errors to approximately a quarter of a millimetre, with an *a posterioiri* standard deviation of 4.5 μm. The two images shown are a plot in *AutoCAD* of the provisional control point coordinates

(Figure 12), and a plot in *AutoCAD* of the same provisional coordinates overlaid with a plot of the final adjusted coordinates (Figure 13).

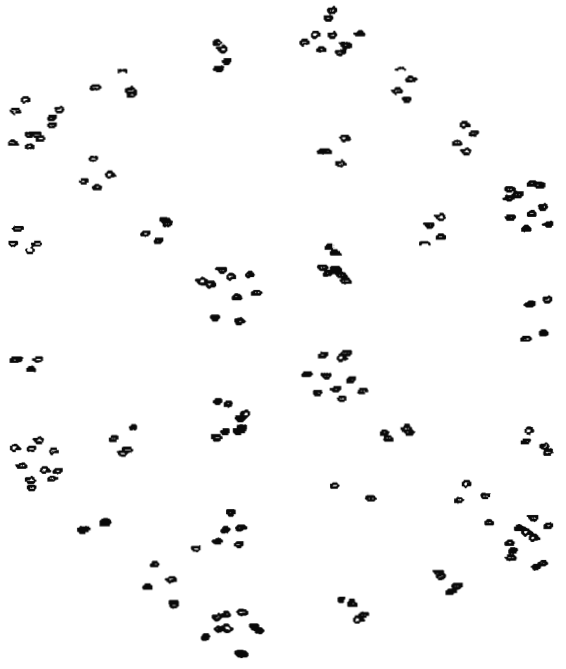

Figure 12. Provisional control point coordinate plot.

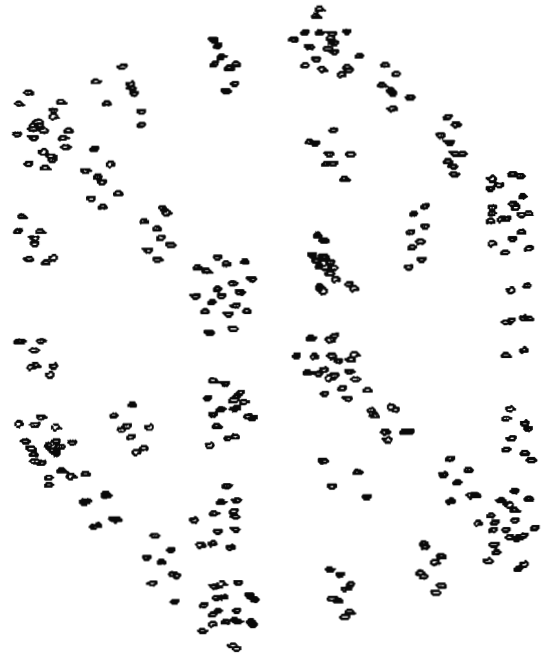

Figure 13. Provisional control point coordinates, overlaid with final adjusted control points.

4.3 IMAGE CAPTURE AND SYSTEM CALIBRATION

With the eventual aim of producing an all-round surface model of an object, it is necessary to capture images that fully cover the surface of the measurement object. The acquired images, after matching, should yield a sufficient number of points to produce a representative point cloud containing sufficient data to enable the reconstruction of the object. This objective is difficult to achieve in the all-round case as occlusions inevitably occur, making impractical numbers of images necessary for adequate coverage. However, a method of image capture and system calibration was developed, involving the separate acquisition of "control" and "object" images, which solved the problem of occlusion of portions of the measurement object by the control frame.

4.3.1 System Pre-Calibration

In a usual photogrammetric survey, the object to be measured is placed within a control field, or a control field is placed around the object. Images are acquired from various camera stations, with the camera being moved between stations. The presence of the control field in the images allows for restitution. The method developed here for an all-round photogrammetric survey, utilises stationary cameras and a rotating object, rather than the usual stationary object and variable camera positions, and involves acquiring two images from every station. One image depicts the control frame, and the second image, taken from an identical position, depicts the object. Thus the "control" image provides control for the "object" image. The procedure involves initially capturing a complete sequence of control images, with one image for every required camera station. The control frame is removed and replaced by the object. A set of object images is now captured, each object image having a corresponding control image taken from an identical location. This is shown in Figure 14.

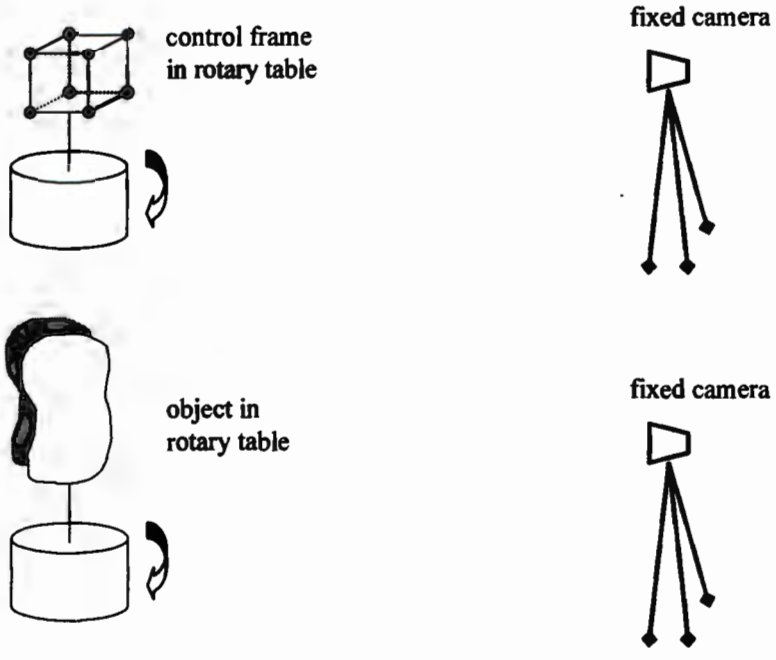

Figure 14. Separate "control" and "object" image acquisition.

The control frame and object are moved separately relative to a set of stationary cameras, but the effect is of a set of cameras being moved first around a stationary frame and then, through identical stations, around a stationary object. This method clearly relies on being able to accurately position the object in the same location as the control frame, for each station. To achieve the accurate repositioning required for pairs of corresponding images, as described above, to be captured, a rotary table from the Department of Materials Engineering at UCT was used. A rotary table is a device used for accurate machining operations.

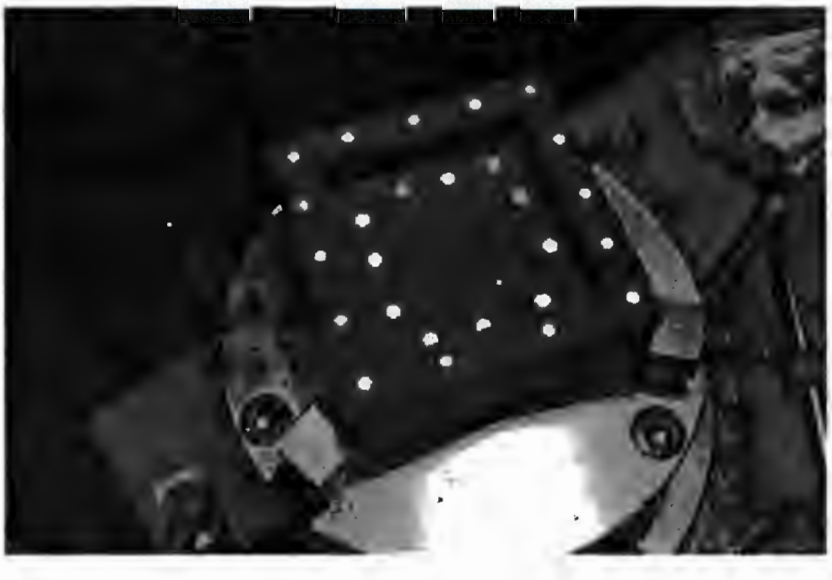

Figure 15. Pratt & Burnerd rotary table with control frame in position for testing.

The Pratt & Burnerd table used for these tests, and shown in Figure 15, allows an object to be rotated smoothly and accurately through 360° with positioning precision, marked on the positioning dial, of better than 0.1°. One full rotation of the crank handle turns the rotary table 4°. A mechanism for translating the control frame and object was not considered, as translations would be more difficult to reproduce precisely.

4.3.2 Testing of Repositioning Precision

In order to test the effectiveness of the rotary table in reproducing a given position precisely, a test was performed using a small control frame, as shown in Figure 15. The frame was held firmly in the rotary table chuck and the crank handle adjusted to zero degrees. A Kodak DCS 420 digital camera was used for image capture. The camera was positioned on a tripod about 0.5 m away from the object and, using a remote trigger to eliminate camera movement, an image was captured of the control frame at a crank rotation of zero degrees. The control frame was rotated away from the zero mark in a clockwise direction, then returned to zero and another image was captured. This procedure was repeated for five images. A further five images were taken after the control frame had been rotated anti-clockwise, and then back to zero. Five more images were captured after arbitrary rotations away from and then back to zero, some of them up to 180°. This variation in rotation direction was performed to account for the inevitable hysteretic inaccuracies of the rotary table.

Having obtained fifteen images of the control frame, all taken at the nominal zero mark, each image was processed and a bundle adjustment was performed, as described in 4.4. This procedure recovered the interior and exterior orientation parameters of the camera for each of the fifteen images. These parameters, in the ideal case, should have been identical for each image. By analysing the resultant orientation discrepancies, the repositioning precision of the rotary table could be gauged. The standard deviations, mean values and ranges for each of the interior and exterior orientation parameters are presented in Table 4 below.

Parameter	Unit	Std. Deviation, σ	Mean	Range
x_p	mm	0.004304	-0.07746	0.01477
y_p	mm	0.026785	0.242362	0.068056
c	mm	0.034428	15.41219	0.103045
ω	°	0.116922	2.055413	0.288262
κ	°	0.027318	0.632778	0.074436
ϕ	°	0.043808	-53.2438	0.094505
X_c	mm	0.175671	1029.727	0.628729
Y_c	mm	0.054039	1044.655	0.175713
Z_c	mm	0.505952	1203.352	1.454655

Table 4. Results of the repositioning precision test for the rotary table.

A viable method of image capture, with simultaneous control, has been demonstrated. This method precludes a common problem associated with all-round photogrammetry - that of shadows and occlusions - while also allowing each object image to contain the maximum amount of useful information and permitting optimum lighting conditions for both the control frame and the object. Because the camera remains at a fixed radius from the rotary table, the method also offers the advantages of maintaining a constant camera-to-object distance, and therefore scale and resolution, and of strongly convergent geometry, which are all considerations in network design (4.1). The repeatability of the method, summarised in Table 4, has been demonstrated. Accuracies of better than 0.5 mm for the perspective centre location, better than 50 μm for the principal point location and between 1′ and 7′ for the rotations in object space were achieved. The rotary nature of the procedure does mean that camera stations are restricted to an arc of constant radius, which necessitates careful planning of the network.

4.4 PERFORMING THE BUNDLE ADJUSTMENT

The next step in the photogrammetric procedure is that of the bundle adjustment. This allows for the high redundancy, simultaneous solution of interior and exterior orientation parameters, 'free' object points, lens distortion parameters etc. using multiple images. The DPG's in-house bundle adjustment software, *Photonet* (Van der Vlugt, 1994), was used in this project. *Photonet* performs a constrained or free network adjustment, either using known fixed control point coordinates or using provisional free coordinates. *Photonet* requires image coordinates for each control point observation, provisional exterior orientations for each camera station, and

provisional interior orientations and digital parameters for each camera used. The formats of the files used may be found in Appendix 1.

The bundle adjustment produces a result file, which includes all orientation parameters, object points, image observations, standard deviations and a list of outliers – observations which have residuals greater than three times the *a posteriori* standard deviation for all observations. In order to improve the accuracy of the adjustment, these observations may be excluded from the adjustment, which is then repeated without the weak observations. In this way, *a posteriori* standard deviations of approximately 2μm for the control point observations in image space may be obtained.

4.5 IMAGE MATCHING AND SURFACE MENSURATION

4.5.1 Preliminary Procedures

Image matching is performed using the DPG's in-house matching program (Smit). The preliminary steps involve edge detection and/or interest point determination, and selection of the area to be matched in the template image. The method of edge detection used was the maximum gradient method with moment preserving which detects edges to sub-pixel accuracy. The area to be matched is defined manually by an operator who views the template image on a monitor. The operator can exclude or include portions of the template image for attempted image matching by marking polygons on the image.

4.5.2 Step-wise Image Matching

A method of step-wise image matching was used to obtain an all-round point cloud which would faithfully represent the surface of an object. This method involves performing the matching procedure repeatedly, each time with a different selection of the acquired images. For each match, one image is the template and adjacent images are the search images. In this way, the object points are recovered by successive separate matches each covering a small portion of the surface. The sets of data, when combined, constitute a complete surface point cloud. In a typical all-round case, it is not uncommon to have twenty or more images covering the object. Thus, the step-

wise matching routine is time-consuming but effective at recovering a representative point cloud.

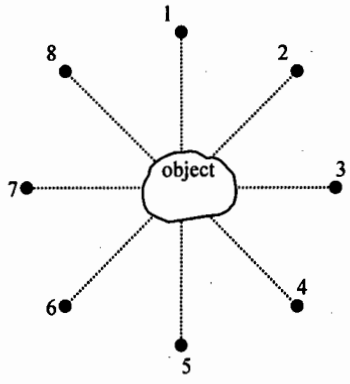

Figure 16. Top view of possible tier of camera stations for all-round image acquisition.

Figure 16 shows the top view of a possible configuration of eight camera stations, located about the equator of an object, and produced by sequential rotations of the rotary table, as described in 4.3. In this case eight rotations, each of 45°, would produce the network configuration shown in Figure 16. Step-wise image matching might occur as follows, in this case:

- Using image 1 as the template, matching would be performed with adjacent images as search images.
- Next, image 2 would become the template, with its neighbouring images as the search images.
- This procedure would continue, using each image in turn as the template image, and producing eight point clouds covering the equator of the object.

In an all-round survey, it might be necessary to employ two more cameras positioned above and below the central camera in a common vertical plane and with convergent axes to provide complete coverage. Designating these upper stations as 1a to 8a, and the lower stations as 1b to 8b, it would be possible to perform sixteen further matches following the procedure outlined above. Figure 17 shows how the cameras might be arranged, with current stations 1, 1a, 1b and stations 2, 2a, 2b shown after a 45° rotation.

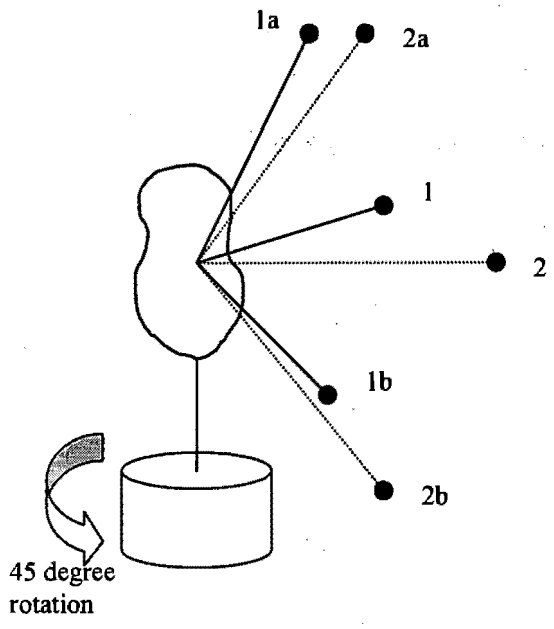

Figure 17. Two sets of convergent camera stations, separated by a 45° rotation.

Note that the cameras are never moved. Rather, their position relative to the control frame and object appears to change due to the rotation of the rotary table. Table 5 details the image sets which might be used for each match, for one tier of images.

Match No.	Search Right	Template	Search Left
1	8	1	2
2	1	2	3
3	2	3	4
4	3	4	5
5	4	5	6
6	5	6	7
7	6	7	8
8	7	8	1

Table 5. Image grouping for matching in an imaginary all-round photogrammetric project.

The procedure outlined in Table 5 would be repeated for images 1a to 8a, and for images 1b to 8b. The result of the step-wise matching in this case would be twenty-four point clouds, all in a common coordinate system, which, when combined, would cover the surface of the object.

This is an example of how image matching in an all-round network might be performed. Due to the geometric diversity of possible measurement objects and network requirements, the possible configurations are numerous and the solutions empirical and *ad hoc*. For instance, in the example above, images above and/or below the template image might have been used as search images, in addition to the left and right images.

4.6 POINT CLOUD REFINEMENT

4.6.1 Matching Reliability and Causes of Mismatches

The matching procedure produces a large quantity of "matched" points. Some of these points are accurate representations of the actual surface, while others lie a considerable distance away from the true surface. Such points have usually been mismatched. This can be as a result of several factors (Baltsavias, 1991) including:

- indistinct surface features or weak surface texture,
- occlusions caused by convoluted surfaces, surface discontinuities or poor network design,
- perspective differences caused by abrupt changes in surface normal
- temporal changes in the image capture conditions,
- unsuitable lighting conditions during image capture,
- noise in the form of radiometric distortions caused by reflections etc., and
- oblique surface orientation

While most of these causes of mismatches - weak texture, occlusions, poor lighting and so on - can be overcome with careful planning in the network design and image capture stages of a photogrammetric project, the problem of oblique surface orientation is inevitable, except when dealing with the most simple of objects. Since it is nearly impossible to prevent oblique surface mismatches in the planning or photography stages, if a method of identifying and eradicating these mismatches can be developed, it will be possible to improve the quality of a cloud of matched points. Consequently, a better representation of the mapped surface can be achieved.

4.6.2 Mismatch Eradication Routine

This theory was put into practice by coding a routine in *C++* (Appendix 3) to eliminate these erroneous points. The routine functions as follows:

1. The exterior orientations of the camera stations are read from file.
2. The camera network centre is determined for the portion of surface in question.
3. The file containing the matched points for the area being processed is read.

4. A best-fitting plane is calculated through every point, using the n^* closest neighbours.
5. The normal of this plane is determined at every point.
6. The angle between the normal vector and the vector from the point to the network centre is obtained, allowing the incidence angle to be calculated and compared with a threshold value.
7. Points with an incidence angle below the threshold - i.e. on surfaces unnacceptably oblique - are discarded.
8. The remaining points are written to a new file.

The file formats used for this routine are the standard formats used in the Digital Photogrammetry Group at UCT, and may be found in Appendix 1.

Figure 18 below depicts a photogrammetric system configuration. Here a network of three cameras is used, with the camera network centre (the arithmetic mean location of all camera station's perspective centre) located at $(X,Y,Z)_n$. Two points obtained from image matching are indicated at p_1 and p_2. The vectors V_1 and V_2 connect each point with the network centre, while the vectors N_1 and N_2 are normal to the surface at each point. In the diagram, the incidence angle (see Figure 6) of V_1 at p_1 is considerably larger than that of V_2 at p_2.

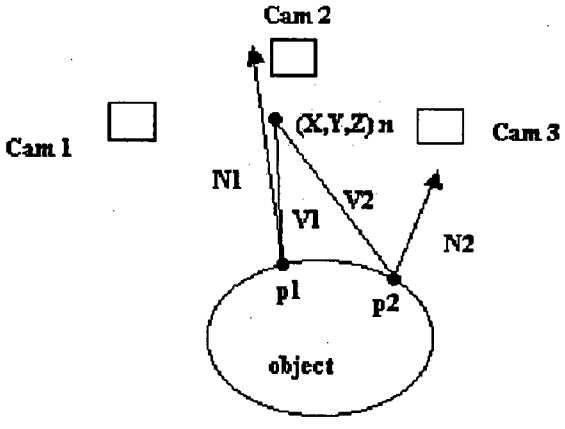

Figure 18. Surface obliqueness relative to a camera network.

* A minimum of two neighbours is required to allow solution of the matrix equations which describe the plane. Two neighbours allow a rapid zero degrees-of-freedom solution, but more neighbours provide redundancy and therefore better accuracy. A good compromise is to use about five neighbours giving three degrees of freedom.

Given that other factors such as lighting and surface texture - as discussed above - are equal, we can expect the accuracy of the matched point p_1 to be superior to that of p_2. Therefore, the incidence angle formed by the vector V_i and a surface at a particular point p_i may be treated as a direct measure of the probable accuracy of the point. In this way, points on a surface with a weak orientation relative to the camera network, may be removed from the point cloud.

The mathematics of the routine, which the author developed from the matrix method for best-fitting straight lines in 2-D by Allison (1993), are described below, beginning with the least squares best-fitting plane equations. In three-dimensional space, a plane has the equation

$$Z = AX + BY + C \qquad \textbf{Equation 14}$$

The normal vector to this plane is given by

$$\vec{N} = \langle A, B, -1 \rangle \qquad \textbf{Equation 15}$$

The vertical distance between any point (X_i, Y_i, Z_i) and the best-fitting plane, i.e. the error, is

$$\Delta Z = AX_i + BY_i + C - Z_i \qquad \textbf{Equation 16}$$

For N points, the sum of the squares of these errors is:

$$f(A,B,C) = \sum_N (AX_i + BY_i + C - Z_i)^2 \qquad \textbf{Equation 17}$$

The values A, B and C which minimise this function can be obtained by setting each of the partial derivatives of f with respect to A, B and C equal to zero. Mathematically,

$$\frac{\partial f}{\partial A} = \frac{\partial f}{\partial B} = \frac{\partial f}{\partial C} = 0 \qquad \textbf{Equation 18}$$

Therefore, the following three equations result:

$$A\sum_N X_i^2 + B\sum_N X_iY_i + C\sum_N X_i = \sum_N X_iZ_i \qquad \textbf{Equation 19}$$

$$A\sum_N X_iY_i + B\sum_N Y_i^2 + C\sum_N Y_i = \sum_N Y_iZ_i \qquad \textbf{Equation 20}$$

$$A\sum_{N} X_i + B\sum_{N} Y_i + CN = \sum_{N} Z_i \qquad \textbf{Equation 21}$$

If the following eight abbreviations are employed,

$$\overline{X} = \sum_{N} X_i \qquad \textbf{Equation 22}$$

$$\overline{Y} = \sum_{N} Y_i$$

$$\overline{Z} = \sum_{N} Z_i$$

$$\overline{X^2} = \sum_{N} X_i^2$$

$$\overline{Y^2} = \sum_{N} Y_i^2$$

$$Q = \sum_{N} X_i Y_i$$

$$R = \sum_{N} X_i Z_i$$

$$S = \sum_{N} Y_i Z_i$$

Then, using Equations 22, Equations 19, 20 and 21 can be represented in matrix form as:

$$\begin{bmatrix} \overline{X}^2 & Q & \overline{X} \\ Q & \overline{Y}^2 & \overline{Y} \\ \overline{X} & \overline{Y} & N \end{bmatrix} \begin{Bmatrix} A \\ B \\ C \end{Bmatrix} = \begin{Bmatrix} R \\ S \\ \overline{Z} \end{Bmatrix} \qquad \textbf{Equation 23}$$

If a matrix M is produced such that:

$$M = \begin{bmatrix} X_1 & Y_1 & 1 \\ X_2 & Y_2 & 1 \\ : & : & : \\ X_N & Y_N & 1 \end{bmatrix} \qquad \textbf{Equation 24}$$

and the following vectors are produced:

$$\vec{A} = \begin{Bmatrix} A \\ B \\ C \end{Bmatrix}$$ **Equation 25**

$$\vec{Z} = \begin{Bmatrix} Z_1 \\ Z_2 \\ \vdots \\ Z_N \end{Bmatrix}$$ **Equation 26**

then Equation 23 is reduced to:

$$M^T M\vec{A} = M^T \vec{Z}$$ **Equation 27**

and so:

$$\vec{A} = \left(M^T M\right)^{-1} M^T \vec{Z}$$ **Equation 28**

Recall from Equation 25 that vector $\boldsymbol{A}$ in Equation 28 contains the coefficients A, B and C of the best-fitting plane, and so the normal vector $\boldsymbol{N}$ may be obtained from Equation 15. Since the camera network centre and current point coordinates are also known, the vector $\boldsymbol{V}$ can be constructed as follows:

$$\vec{V} = \left\langle X_n - X_i, Y_n - Y_i, Z_n - Z_i \right\rangle$$ **Equation 29**

Using dot products and the cosine rule:

$$\cos\alpha = \frac{\vec{V} \bullet \vec{N}}{\left\|\vec{V}\right\| \left\|\vec{N}\right\|}$$ **Equation 30**

the angle between the normal vector and the object point vector can be determined. Thus, the incidence angle θ of the vector V is 90°-α, and a measure of surface obliqueness has been obtained. The routine uses this information, which it calculates for every point, to determine if a point has an incidence angle below a user-defined threshold. This allows points to be sorted based on the orientation of the surrounding surface.

4.6.3 Testing of the Routine

In order to establish the effectiveness of the routine described above, three point clouds were obtained photogrammetrically, i.e. by image matching. All the mapped surfaces had portions of surface obliquely oriented relative to the camera network, and so the resultant data sets contained a significant number of poor quality points to test the routine.

Threshold Angle, θ_t	Data Set 1 Retained	Data Set 1 %Rejected	Data Set 2 Retained	Data Set 2 %Rejected	Data Set 3 Retained	Data Set 3 %Rejected
0° (no processing)	8752	-	4189	-	1856	-
20°	7009	19.9	3908	6.7	1786	3.8
45°	4049	53.7	2730	34.8	1278	31.1
75°	679	92.2	366	91.3	248	86.6

Table 6. Effect of threshold angle on number of points rejected for three test data sets.

Points with an incidence angle below the threshold are discarded. Table 6 shows how, with increasing incidence angle threshold, more points are rejected from each data set. A suitable threshold angle will vary from project to project, depending on the quality of imagery, the accuracy of the matching results, the presence or absence of texture or reflections on the surface, and occlusions of the object. Therefore, a threshold should be selected empirically in each individual case, keeping in mind that a compromise needs to be reached between improving the quality of the data, and indiscriminately rejecting valid data.

As shown in Table 6, threshold angles of 20°, 45° and 75° degrees were used in the processing of the three test data sets, resulting in three reduced point clouds for each surface. A grid for each of these surface point clouds, including the original data sets, was then created by Kriging using the *Surfer* package. This enabled the effectiveness of the routine at removing mismatches and 'spikes' to be gauged. The grids for each surface are presented in Figures 19, 20 and 21.

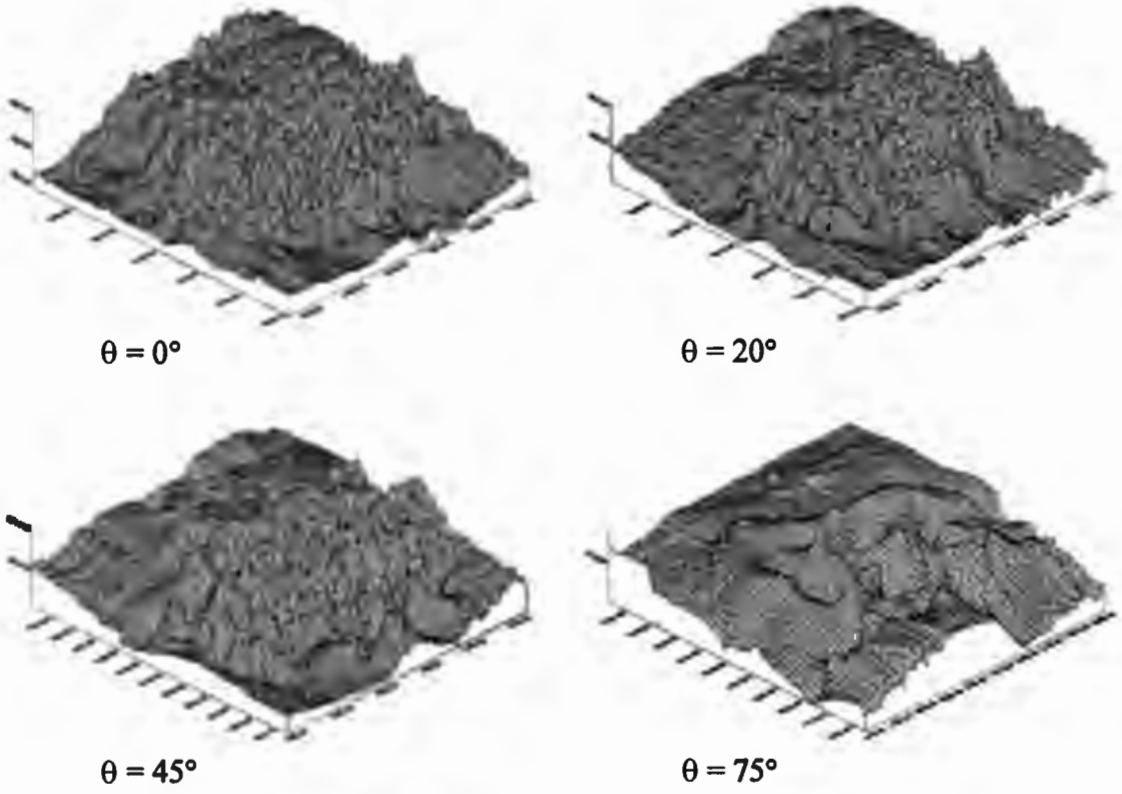

Figure 19. Four DSMs generated for Data Set 1.

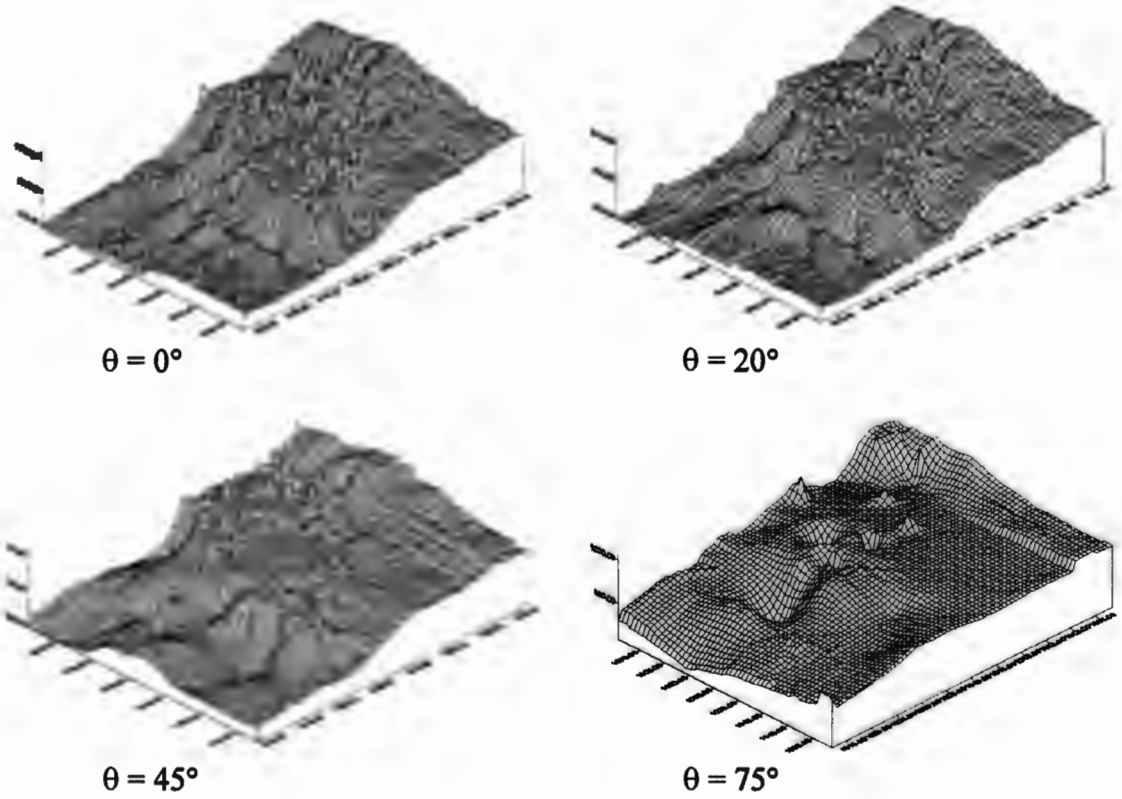

Figure 20. Four DSMs generated for Data Set 2.

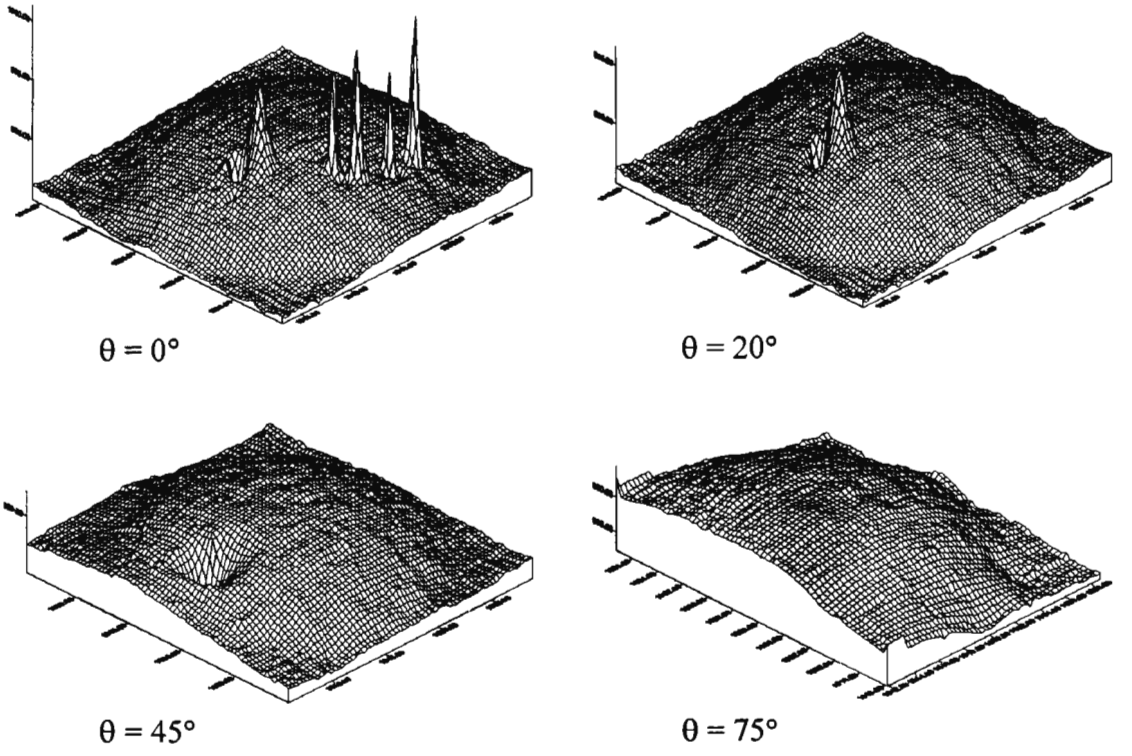

Figure 21. Four DSMs generated for Data Set 3.

The twelve DSMs shown above indicate mixed results for this method of point cloud refinement. Figures 19 and 20 show a marked smoothing effect, while some obvious mismatches still remain. Figure 21 shows the most convincing evidence for the routine. Notice how in the plot for a threshold angle of 20° four big spikes have been removed, with the loss of less than 4% of the data (see Table 6). Note also how the subsequent plots - at 45° and 75° - show how the boundary of the point cloud has receded due to the increasing threshold angle. This is demonstrated in Figure 22 below.

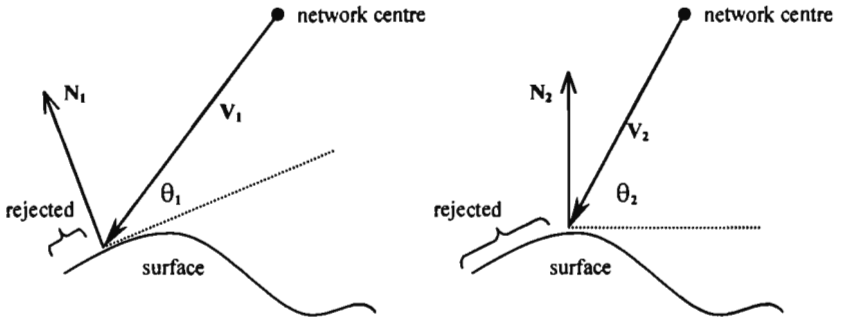

Figure 22. Surface point rejection based on surface obliqueness.

Figure 22 shows a surface and camera network centre. In each case, a vector (V_1, V_2) is constructed from the network centre to a surface point. The surface normal (N_1, N_2) is obtained by finding a best fitting plane at the current point. Thus the angle (θ_1, θ_2) is obtained. Note how the larger incidence angle threshold, θ_2, causes a greater number of surface points to be rejected.

4.6.4 Discussion of Refinement Routine

The results presented in 4.6.3 show that the method of refining a 3-D point data set on the basis of surface obliqueness (i.e. incidence angle) is at least partially successful in improving point quality. Data Set 3 was significantly improved after less than 4% of the data had been removed by the routine.

The selection of camera stations to be included in the network centre calculation should be done with care. Recall from the discussion on step-wise image matching in 4.5 that points on an object's surface are recovered from a number of successive image matching routines. These separate data sets should be processed with the refinement routine individually before being combined into an all-round point cloud. It is important to remember that only camera stations that were used in the image matching for the data being processed, should be used to determine the network centre for use in the refinement procedure. Figure 23 shows why this is important.

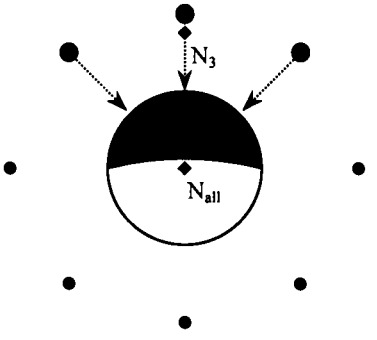

Figure 23. The importance of correct camera network centre determination.

Figure 23 shows a possible convergent network of eight camera stations, as might be produced by the system pre-calibration procedure described in 4.3. Each station is positioned at a constant radius from a spherical object, and is separated by 45°. Let us

assume that the points on the shaded portion of the surface were recovered by matching the images obtained from the top three camera stations. The point N_3 is the network centre for these three stations, while N_{all} is the network centre for all eight stations. If an incidence angle threshold of 60° had been selected for point refinement, but N_{all} had been obtained by incorrectly using all cameras stations in the network centre calculation, no points would be rejected because of the erroneous network centre. Correct calculation of N_3 as the network centre would allow oblique surface points to be rejected, as desired.

Indiscriminate rejection of points, particularly when the accuracies obtained in the image matching are good, is not recommended. Rather, the conservative use of this routine - in conjunction with mismatch removal by inspection and the use of other criteria, such as known geometric constraints, and approximate point accuracies provided by the image matching routine - is advocated.

4.7 SURFACE MODELLING, OBJECT RECONSTRUCTION AND VISUALISATION

The successfully completed photogrammetric procedure produces a cloud of points, closely representing the object. However, this may not be enough to actually visualise the object. Some sense has to be made of these numerous points, and relationships between neighbouring points must be established, in order that the surface of the object may be reconstructed digitally.

4.7.1 Digital Surface Models (DSMs)

The mapping of a terrain or topography, one of the best established and most traditional applications of photogrammetry, is relatively easy to accomplish. Having obtained a surface data representing the terrain, a means is required whereby the surface may be visualised using the discrete point data. Many methods to achieve this end have been developed, all of which resort to a system of gridding, meshing or triangulating resulting in what will be referred to in this project as a DSM.

Maune (1996) defines a DEM (digital elevation model) as "the digital cartographic representation of the elevation of [a surface] (Z) at regularly spaced intervals in X and

Y directions". Maune differentiates between a DEM and a DTM (digital terrain model) by noting that, while both models consist of surface heights above a lattice or grid of regularly-spaced nodes the DTM also incorporates significant surface features and breaklines. However, these definitions are by no means universally accepted and, in fact, the terms have become synonymous (Ackermann, 1996) and are often used to denote virtually any surface model consisting of heights above a datum grid, and is even used to refer to triangulated models. Both the DEM and DTM use methods of interpolation (cubic splines, polynomials etc.) to compute height values at each grid point from the randomly scattered data.

A triangulated irregular network (TIN) is a set of adjacent non-overlapping triangles computed from an irregular set of 3-D points. A TIN data structure contains information detailing which points in the data set form each triangle, and which triangles are neighbours. This method of surface modelling allows for analysis and display of terrains and other similar types of surface. Maune regards TINs as generally superior to DEMs and DTMs in their effectiveness at accurately modelling a surface.

One thing that is true of any DSM is that it is a single-valued function. This means that any DSM describes a surface where $Z = f(X,Y)$. In other words, for any given X, Y position, one Z ordinate will be determined for the DSM. Therefore, the visualisation of a more complex surface presents unique problems. An object's surface may contain many discontinuities, convolutions and other features not encountered in a typical terrain. Indeed, in the all-round case, it is inevitable that double or even multiple heights will be encountered at any node. Consequently, the software required to render a surface model of a complex object must be much more powerful than currently available terrain modelling software. In fact, no suitable software was found which was capable of performing the required surface modelling.

4.7.2 3-D Delaunay Triangulation

A DSM is most often created using one of a number of well-established interpolation or triangulation techniques. As discussed above, this is not sufficient to allow the visualisation of a complex object. However, a method of 3-D Delaunay triangulation (Cignoni *et al*, 1992) has been developed which has the potential to create a solid

body from 3-D point data. This is an adaptation of the 2-D Delaunay triangulation algorithm to the requirements of the 3-D case. (Here, 2-D refers not to the data, which still consists of 3-D coordinates, but rather to the *XY*-plane in which the triangulation occurs.) In the 2-D case, point triples are connected to form triangles. If the unique circle passing through each triple of points contains any other points, then that included point is used to create new triangles.

In the 3-D case, the basic units are tetrahedra rather than triangles. This requires that point quadruples are joined to form a tetrahedron. In this way, a solid body is formed from a point cloud. A 3-D Delaunay triangulator, *DeWall* (Cignoni *et al*, 1992), was tested on the point clouds generated in this research. The input required is simply a list of the *X*-, *Y*- and *Z*-coordinates of all the points produced by image matching. The output is a file containing the indices (relating to the point numbers) of the four vertices of each tetrahedron created.

This method of triangulation provided mixed results with some point clouds being successfully triangulated and others causing the program to fail or return an error message. According to Cignoni, this is due to the fact that large datasets (typically of 20000 points or more) may cause the algorithm to fail. Additionally, the algorithm does not cope with datasets containing four or more cospherical points. While this is generally unlikely to occur in a random dataset, it was the cause of failure in a number of attempted triangulations.

The tetrahedra resulting from this method of 3-D Delaunay triangulation were drawn in *AutoCAD*, using a dedicated scripting routine programmed by the author. The resultant wire-frame models indicated that the 3-D Delaunay method could not be considered entirely successful. See Chapter 5 for an example of the 3-D Delaunay method's output.

4.7.3 Reconstruction and Visualisation

The reconstruction of an object in three dimensions from photogrammetric data presents problems for the photogrammetrist. Obtaining reliable point data from an image matching procedure is possible, and the methods discussed in 4.6 assist in ensuring the integrity of these points, but presenting this information in a sensible and

realistic manner is a complex and difficult requirement. As Hoppe (1994) states: "To adequately scan [a complex object], multiple view points must be used. Merging the data generated from multiple view points to reconstruct a polyhedral surface representation is a non-trivial task."

Haala and Brenner (1999) present their work on urban area reconstruction. Aerial laser scanning was used to obtain 3-D point information while 2-D GIS provided the ground plans. Buildings were segmented into primitive solids which were then fitted to the data. The authors note that manual refinement, via a graphical user interface, is necessary where a high degree of detail is required. Visualisation is achieved by overlaying the DSM with orthophotos. While this work suffices for the relatively simple shapes encountered in an urban scene, it is anadequate for the requirements of the all-round, complex body case. Consider, for example, that a simple building can be completely represented by about ten or twenty points, while many thousands of points might be needed to constitute a thorough representation of a complex object such as a sculpture.

Chikatsu *et al* (1997) have done work on real-time 2.5 dimensional object reconstruction. Chikatsu's work involved the use of line lasers and one CCD camera, allowing the recovery of one side of the measurement object, hence his use of the term 2.5-D. The system proved useful for the measurement of simple surfaces, but problems presented by holes and occlusions in objects were not overcome.

Van Der Vlugt's (1995) research on an automated 3-D surface measurement system included work done towards all-round object measurement. The all-round measurement of a wooden shoe mould, for subsequent CAD modelling and CNC milling, was performed. Van Der Vlugt reports that "intermediate software to convert random surface data into evenly distributed data is needed between the photogrammetric and *AutoCAD* packages."

Hoppe (1994) has developed software for the reconstruction of surfaces from random 3-D point clouds. His work looks extremely promising. He describes and demonstrates an algorithm that takes as input an unorganized set of points in 3-D representing an unknown surface *S*, and produces as output a surface that approximates *S*. Previous surface reconstruction methods have typically required

additional knowledge, such as structure in the data, known surface genus, or orientation information. In contrast, the method outlined in Hoppe's thesis requires only the 3-D coordinates of the data points. From the data, the method is able to automatically infer the topological type of the surface, its geometry, and the presence and location of features such as boundaries, creases, and corners.

The method used in this work to allow the visualisation of the object surface involved the use of *Surfer* to obtain a regular grid, and *CosmoWorlds* virtual reality modelling software for the rendering of the surface.

5. Test Case

The photogrammetric routine detailed in Chapter 4 was tested on a suitable measurement object. The object selected was a five-litre plastic chemical bottle. This bottle was of a suitable size to be mounted in the rotary table, and provided an all-round surface for measurement. A steel mounting rod was fitted in the bottle's screw-on cap. This enabled the bottle to be mounted, in an inverted position, in the rotary table. In order to have sufficient texture on the otherwise shiny surface, the bottle was painted with a matt white finish, and a speckle pattern in matt black was added by hand. The bottle is shown in Figure 24.

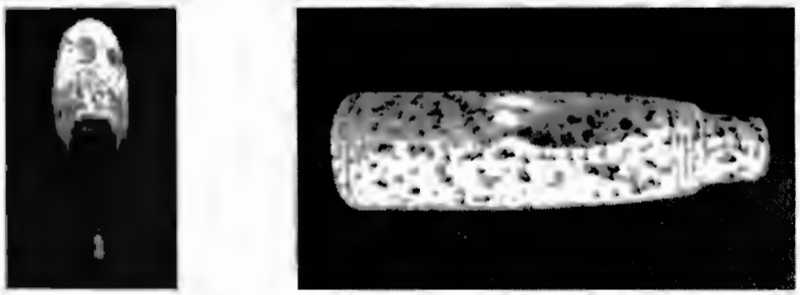

Figure 24. The plastic bottle, prepared for measurement.

Note the matt finish, added texture and matt black steel mounting arrangement.

5.1 PHOTOGRAPHY

The photography for the plastic bottle required adequate coverage to enable subsequent image matching to recover a complete surface point cloud. The network also required strong geometry with an even spread of camera stations.

5.1.1 Network Design

The network chosen consisted of three cameras mounted in roughly the same vertical plane. The upper camera was directed downwards at about 60°, the lower camera upwards at about 60°, and the middle camera was approximately level. The separation between the stations in the horizontal plane was produced by incremental rotations of the rotary table of 22.5°. The resultant network was a three-tiered belt of 48 camera stations. The ray from the perspective centre location of each station, which was

obtained from the bundle adjustment discussed in 5.2, to the centre of the control frame was plotted in *AutoCAD* along with the control point coordinates, as shown in Figure 25. Note the strongly convergent geometry, and spatial regularity of the camera placement.

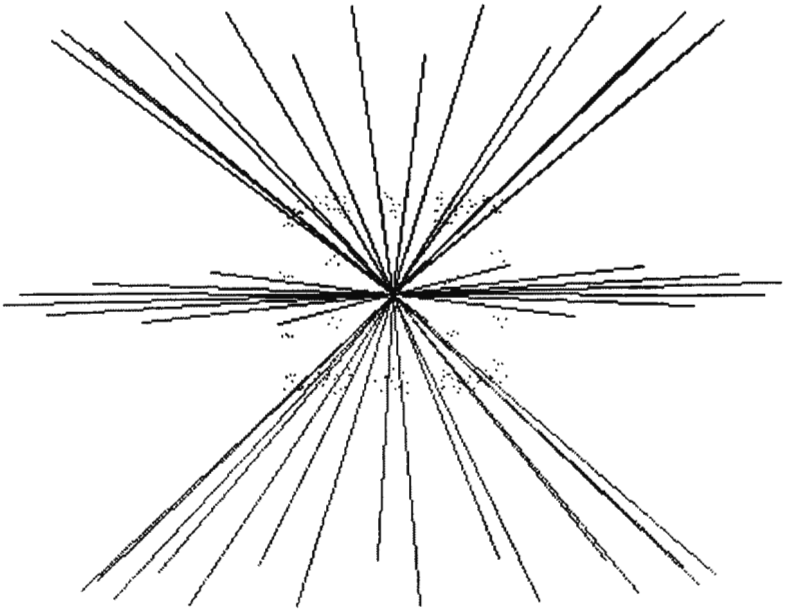

Figure 25. Control points with camera stations, as plotted in *AutoCAD*.

The upper camera stations (numbers 1 to 16) were numbered beginning at the zero degree rotation of the rotary table. The middle stations received numbers 17 to 32, and the lower tier 33 to 48.

5.1.2 Image Acquisition

The cameras selected for this project were ITC CCD video cameras. The specifications include a chip size of 7.95 × 6.45 mm (795 × 596 pixels) and a signal-to-noise ratio of better than 48 dB. The advantages of using these, rather than higher resolution digital cameras, included:

- image acquisition was possible via a framegrabber, allowing rapid image download,
- the smaller image size (approximately $1/6^{th}$ the size of the Kodak DCS 420) meant quicker processing and smaller total data storage requirements, and

- the cameras were able to be fitted in the required workspace with less difficulty than would have been experienced with a bulkier camera.

A PIP512B framegrabber was used to capture the analogue video signal and convert into a digital image. This produces an image of 512 × 512 pixels. The framegrabber card accepts three coaxial inputs, the analog video signals from each camera, and sends one signal to a video monitor. Image acquisition was done singly, as the PIP512B does not permit simultaneous image capture. Nevertheless, the system of video cameras and framegrabber allowed the ninety-six images necessary for the project to be captured in less than forty-five minutes, including the time it took for successive rotations of the rotary table, and to switch the control frame and object after the 48 control images had been taken.

The method of system precalibration, as discussed in 4.3, was used for the image acquisition. In this case its advantages included:

- the drastic reduction of occlusions due to an obstructing control frame,
- greater ease of image capture, as it was unnecessary to move the cameras between stations
- the ability to maintain a backdrop of matt black sheeting in one position, thus improving the quality of imagery, and
- the facility to adjust lighting conditions for the separate acquisition of control and object images, thereby obtaining the best possible lighting for each set of images.

The all-round control frame, detailed in 4.2, was fixed in the rotary table. The table was rotated to the zero degree mark, and images captured from each of the three fixed cameras. A rotation of 22.5° was made, and the next three images were taken. Subsequent rotations of 22.5° were made, with images captured after each rotation, resulting in forty-eight 'control' images. See 5.2 for repositioning precision data.

With the cameras fixed in their separate stations, the control frame was removed from the rotary table and the measurement object – the plastic bottle – was fixed in place.

The rotary table was returned to the 0° mark, and forty-eight images, from stations corresponding exactly to the 'control' image stations, were taken of the object.

5.2 BUNDLE ADJUSTMENT

5.2.1 Preliminary Procedures

The forty-eight control images were processed by the standard photogrammetric methods. The targets, including control points, were isolated by running each image through a threshold routine. The control points were then identified and labelled, following the control point numbering system described in 4.2, and detailed in Appendix 2. Using the identified control points, the DLT was calculated for each image, allowing a space resection to be performed which yielded provisional interior and exterior orientation parameters.

The next step was to combine all the data in a bundle adjustment, to recover the precise orientation parameters for each camera station. The bundle adjustment was run as a constrained network, with only the exterior and interior orientation parameters and three terms for radial and decentering lens distortion kept free.

5.2.2 Bundle Adjustment Results

The bundle adjustment converged after five iterations, after running for almost two hours on a 486 DX4-100 PC. Outliers (control point observations with a standard deviation greater than three times the *a posteriori* standard deviation value) were then removed from the adjustment, which was performed again, this time converging after four iterations with an *a posteriori* standard deviation of 0.0043 mm. The errors for the control point coordinates in object space were all in the sub-millimetre region. The full bundle adjustment results are presented in Table 7.

	First Bundle Adjustment	Second Bundle Adjustment
Degrees of freedom	4006	3918
Iterations to converge	5	4
σ_X	1.32 mm	0.34 mm
σ_Y	0.88 mm	0.69 mm
σ_Z	0.81 mm	0.26 mm
σ_0	0.0099 mm	0.0043 mm

Table 7. Bundle adjustment results.

5.2.3 Repositioning Precision

The bundle adjustment yielded accurate exterior orientation parameters for each camera station. This information allowed the actual rotations, experienced by each camera due to the rotation of the rotary table, to be reconstructed. Comparing these actual rotations with the desired value of 22.5° determined the repositioning precision for this project. Recall that the repositioning precision was determined in a test in 4.3. The results are as tabulated below, for each camera tier, and overall.

	TopTier	Middle Tier	Bottom Tier	Overall
Mean	22.50001°	22.50001°	22.50056°	**22.5002° = 22° 30′ 0.7″**
Std. Dev.	0.025867°	0.032905°	0.030485°	**0.028567° = 1′ 43″**
Max	-	-	-	**22.5548° = 22° 33′ 17″**
Min	-	-	-	**22.4374° = 22° 26′ 14″**

Table 8. Repositioning precision for the test case.

The overall standard deviation of 1′ 43″ above translates into a displacement of approximately 0.4 mm over the average camera-to-object distance, which was about 0.8 m. The mean rotation was a mere 0.0002° or 0.7″ from the desired rotation of 22.5°. The error in the actual experienced rotation had two main sources:

1. The actual rotation imparted to the control frame by the rotary table was determined manually, and errors would have occurred due to rotations not being of precisely the required amount.
2. The control point coordinates determined for the control frame had associated errors (see 4.2) which would have been transferred to the resected camera perspective centre locations.

5.3 IMAGE MATCHING

5.3.1 Edge Detection

Before image matching could commence, edges were first detected in each image. This was done using the maximum gradient method, followed with moment preserving, to detect edge locations to sub-pixel accuracy. The area required to be matched was then selected in each template image. Figure 26 shows a binary image of edge locations, detected for one template image.

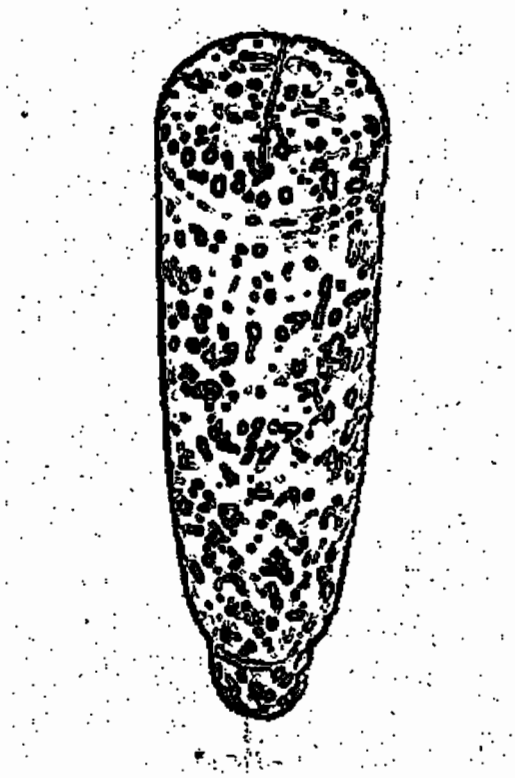

Figure 26. Binary image showing edge locations.

Note how the dark, featureless background used during the photography and the matt black mounting rod have produced a negligible quantity of edges.

5.3.2 Surface Subdivision

The bottle was subdivided along the natural breaklines of its surface to allow the matching to be run. This was necessary because the matching routine requires a search axis approximately perpendicular to the surface. The images of the top and bottom portions were matched using the *Z*-axis, which ran along the length of the

bottle, as the depth axis. In this way points on the flat portion of the lid and base of the bottle could be matched. The middle portion of the bottle was matched using as the depth axis whichever of the *X*- and *Y*-axes was orientated most towards the camera station. These subdivisions, and the orientation of the coordinate system, are shown in Figure 27.

Figure 27. The surface subdivision for matching.

5.3.3 Step-Wise Image Matching Configuration

The method of step-wise image matching discussed in 4.3 was used in this project. Using the subdivisions outlined above, image sets were selected which, when matched, would produce points belonging only to one of the subdivisions. The image strips for each camera are shown in Figure 28.

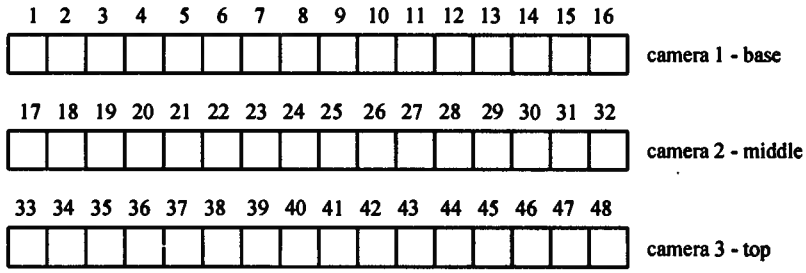

Figure 28. Strips of images captured from each station.

The upper row of images was taken by Camera 1 and these were used to measure points on the base of the bottle, which was mounted in an inverted position. Cameras 2 and 3 provided the images for the matching of the sides and top of the bottle. Figure

29 depicts how the image sets were configured for matching. For the sake of clarity only twenty-one of the forty-eight images, shown in Figure 28, are included. The arrows in the diagram, emanating from each respective template image, indicate which adjacent images were used as search images in the image matching. For example, when image No. 7 was used as a template image, the arrows indicate the search images as Nos. 6 and 8. When image No.25 was used as the template, the search for matching points occurred in images Nos. 9, 24, 26 and 41.

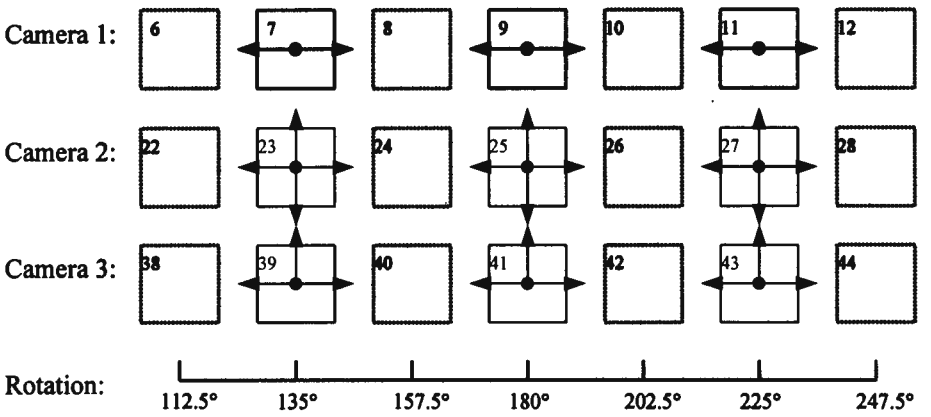

Figure 29. Diagrammatic layout of images, showing sets for image matching.

Using the forty-eight images and sub-pixel edge locations for each of the twenty-four template images, twenty-four separate matching procedures were performed. Only images at odd-numbered stations (those at 0°, 45°, 90°, .., 315°) were used as template images, while the intermediate images (those at 22.5°, 67.5°, 112.5°, .., 337.5°) were used only as search images. These 'search-only' images are shown with dotted lines in Figure 29.

5.3.4 Image Matching Results

The image matching was successfully performed in about twenty-four hours. Strict quality controls were used in an attempt to reduce the number of mismatches and blunders. The result of these stringent measures was that less object points were produced by the matching than would have been the case with less strict controls. Even so, the twenty-four separate point clouds which were produced contained a total of over three thousand points. A second image matching procedure was performed, with less stringent parameters, for portions of the bottle which had less than sufficient numbers of points. This increased the total number of points to almost seven thousand, with a consequent decrease in matching accuracy.

The matching accuracy compared very favourably with results obtained by Van der Vlugt (1995), who obtained values of σ_0 of 3.8 for a propellor, 4.5 for a plastic lid - both measured conventionally from one side - and 6.0 for a wooden shoe mould, which was measured from all sides. In this project, the overall results obtained from the two matching procedures for the plastic bottle were as tabulated below in Table 9.

	First Match	Second Match
Mean σ_0	3.98	4.75
Maximum σ_0	4.21	6.50
Minimum σ_0	3.57	3.80

Table 9. Matching results.

For each set of matching results, single point clouds were obtained for the top and the base of the bottle by amalgamating the eight separate point clouds for each respective surface. The sides of the bottle, represented by a further eight point clouds from each matching procedure, were divided into quadrants. Each quadrant contained two of the eight point clouds, and had as its depth axis either *X* or *Y*. The point clouds obtained for the bottle from both matching procedures were plotted in *AutoCAD*. Figure 30 depicts the inverted bottle, as plotted in *AutoCAD*.

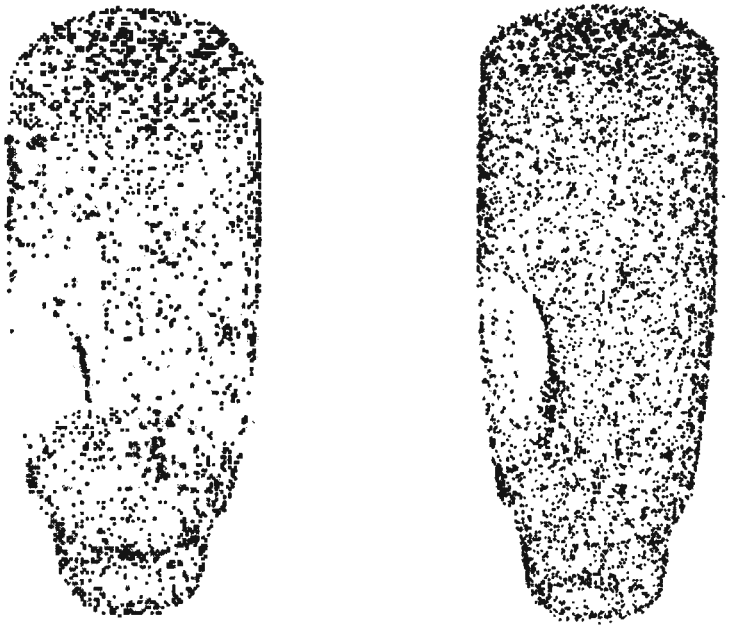

Figure 30. *AutoCAD* point cloud models of 3145 points (left) and 6717 points (right) of the bottle for the two separate matches.

Obvious mismatches were removed for each point cloud, and these totalled twenty-two for the left hand cloud and 364 for the right hand cloud. These mismatches were either detected by eye, and using a simple routine which eliminated any points not lying between the central axis of the bottle and the camera. (These points could not possibly have been correctly matched.) Point elimination based on local surface orientation – detailed in 4.6 – was also performed on each point cloud, using a conservative threshold angle of 30°, although the overall quality of points was good enough to warrant not using this procedure.

5.4 SURFACE RECONSTRUCTION

5.4.1 Digital Surface Models

The surface presentation package, *Surfer*, while inadequate for all-round modelling, as discussed in 4.7, allowed each point cloud to be represented as a wire frame terrain. Contour maps were also plotted for each portion of the surface. In each case, the axis used as the depth axis in the image matching was used as the height axis for the plotting of the DSM. Surfaces plotted facing in the negative axis direction appeared as positive shapes, while surfaces facing in a positive axis direction – such as the bottle's top which faced along the negative *Z*-axis - appeared as negative shapes. Three examples of *Surfer's* output are shown in Figures 31, 32 and 33 below.

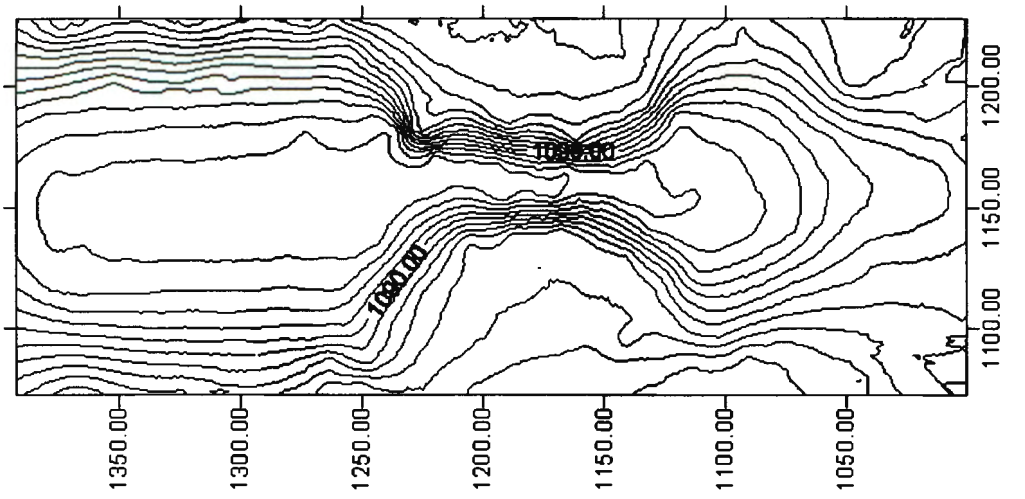

Figure 31. *Surfer* contour plot of bottle, including handle, from station 17.

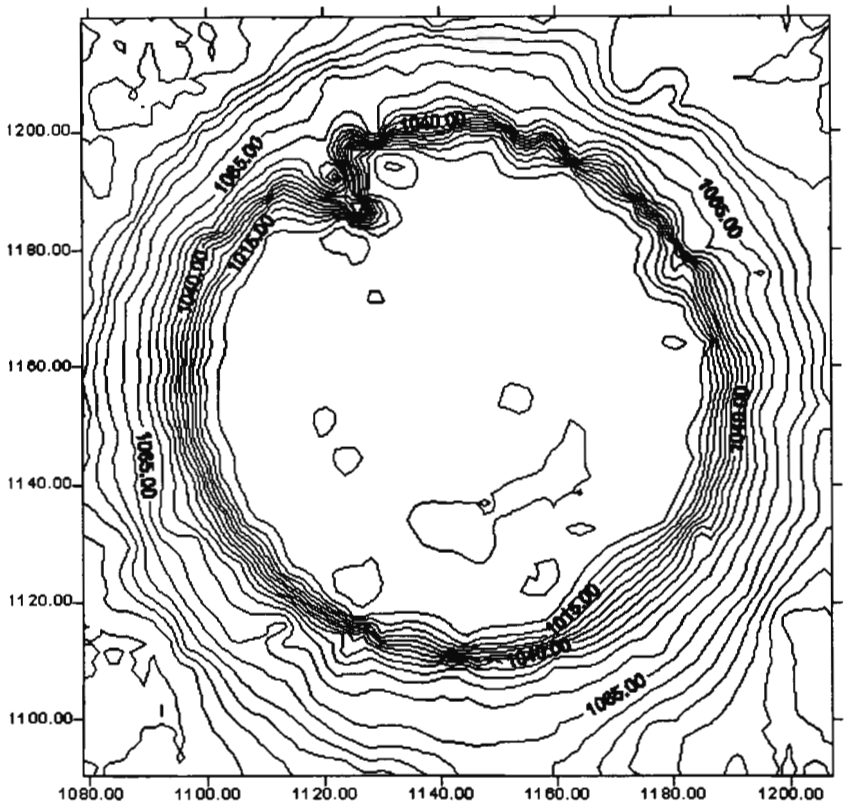

Figure 32. ***Surfer*** **contour plot of the top of the bottle.**

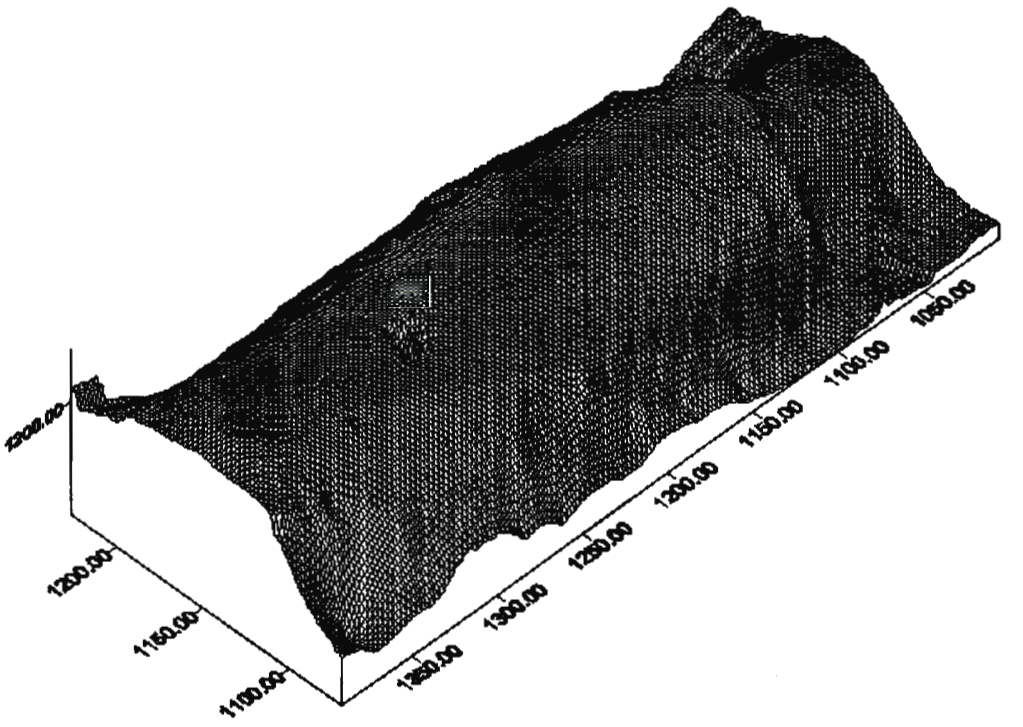

Figure 33. ***Surfer*** **DSM plot from station 25.**

The method of 3-D Delaunay triangulation (see 4.7) was tested on the data obtained for the bottle. The result of this method of triangulation for the base of the bottle is presented, as plotted in *AutoCAD*, in Figure 34.

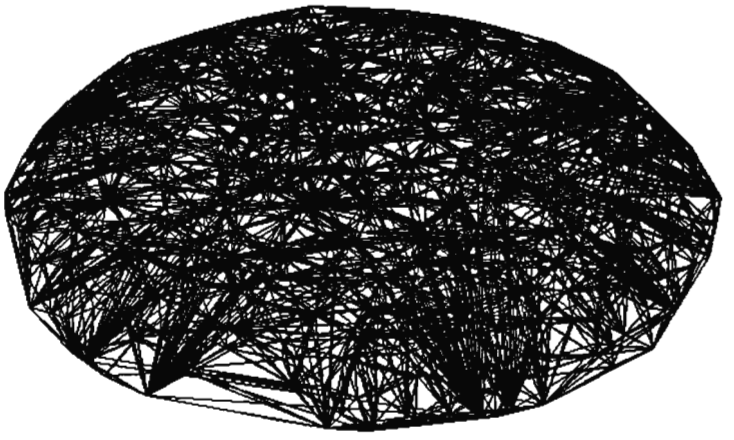

Figure 34. Tetrahedra generated by the 3-D Delaunay method of Cignoni.

5.4.2 All-Round Model

CosmoWorlds virtual reality modelling software was used to visualise the complete surface of the object. This software, while not designed with engineering applications in mind, is quite flexible, allowing texture mapping, image draping, as well as interactive zooming, rotations, walk-throughs etc. The software does not actually include the ability to model random point clouds as solid bodies, but a method was devised which allowed the complete object surface to be rendered. This method involved piecing together six surface models, two models - one in the positive direction and one in the negative direction - for each of the *X*-, *Y*- and *Z*- axes. Figure 35 shows how this was done.

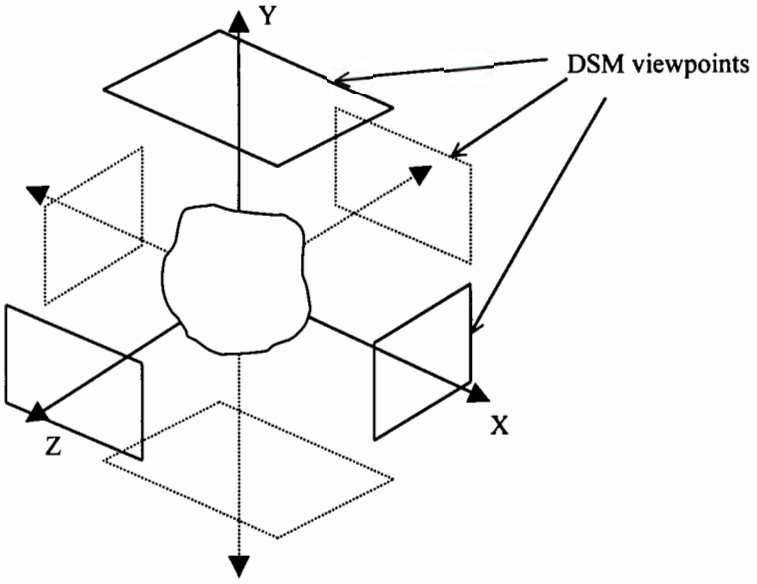

Figure 35. *CosmoWorlds* axis system, and configuration of six DSMs.

Note in Figure 35 that the height axis of the DSM in *CosmoWorlds* is the *Y*-axis. All the grid data was rearranged so as to have *Y*-values as the heights. The grid data used to produce the DSMs was obtained from *Surfer*, and converted into the required format by a dedicated routine, programmed in C++. All the DSMs were initially plotted in the *XZ*-plane, as required by *CosmoWorlds*, but were subsequently shifted by a rotation and translation, to the required orientation. *CosmoWorlds* also includes a feature which allows either the positive or negative side of a DSM to be rendered. Thus the problem of negative shapes, obtained in DSMs plotted looking in the positive axis direction (as encountered in *Surfer*), was overcome. Figures 36 and 37 show separate DSMs modelled in *CosmoWorlds*.

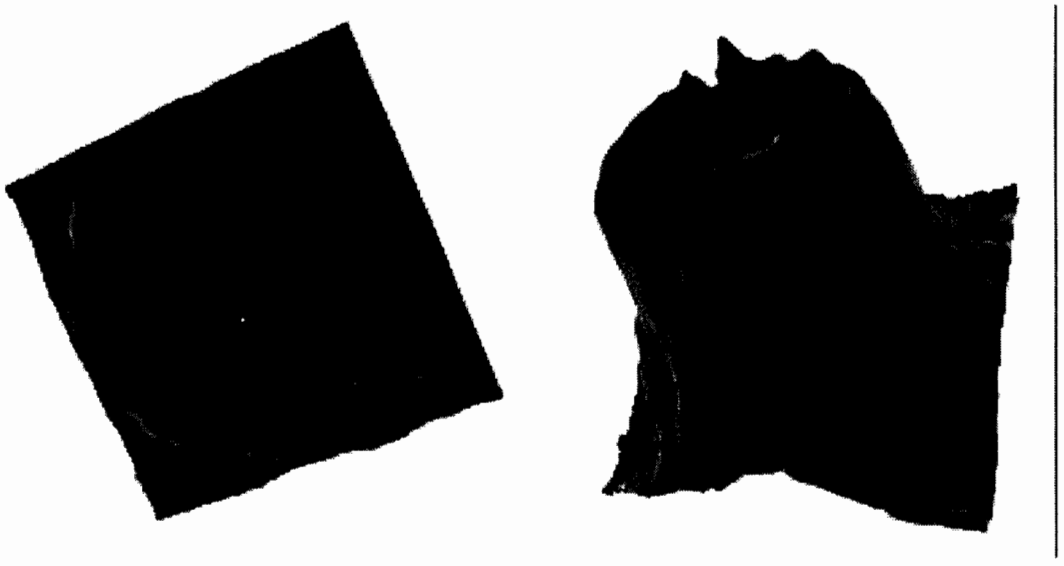

Figure 36. Base and top of the bottle.

Figure 37. Three sides of the bottle including the handle.

As described above, the separate DSMs were rotated and translated into their correct positions in CosmoWorlds, to produce the all-round model shown in Figure 38.

Figure 38. The bottle viewed obliquely from below showing the handle

6. Discussion and Conclusions

The objective of this research, as stated in Chapter 1, was the development of a system of all-round surface measurement using close-range digital photogrammetry. Conventional photogrammetric techniques were questioned and re-evaluated, and the photogrammetric method, detailed in Chapter 4, was developed based on these techniques, and on the novel approaches developed in this work. The developed method was tested at various stages of the research, and a detailed test case was undertaken to establish the potential of the method.

The results achieved in the testing of the photogrammetric procedure, presented in Chapter 4, and in the detailed test case, presented in Chapter 5, were sufficiently satisfactory to suggest that the approach taken in this research has the potential to be successfully applied in situations requiring all-round surface measurement. The approach offers several benefits over conventional methods of photogrammetry.

The novel method of system pre-calibration, and separate control frame and object image acquisition, offers the following advantages:

- minimal occlusions and shadows occur in the imagery and, therefore, maximum use is made of the available pixels
- once-off calibration and bundle adjustment for several measurement objects
- optimum lighting is possible for both the control field (during the once-off calibration) and the subsequent measurement surfaces
- regularly-spaced and strongly convergent network design
- static camera stations eliminate tedious camera relocation between images and make for rapid image capture of less than 30 seconds per image
- excellent potential for automation, by means of stepper- or servo-motors for the rotary table

The pre-calibration approach was tested for rotation repeatability, and the results were good. In a test performed to establish the repeatability of the rotary table placement, standard deviations for the perspective centres of resected camera stations were between 0.05 and 0.5 mm, for the rotations in object space, between 10″ and 7′, and

for the interior orientation parameters, between 4 and 34 μm. A pre-calibrated camera network could allow several objects to be measured without any need for calibration, but the system should be periodically recalibrated for optimum accuracy.

The step-wise image matching routine also has good potential for automation and proved to be a secure method of obtaining comprehensive surface data. Some degree of user intervention will always be required, but the repetitive nature of the process means that, once the photogrammetrist has determined the configuration of template and search images to be used, automated matching would be possible.

Having obtained separate sets of 3-D point data, the method of point cloud refinement has the ability to eliminate mismatches and oblique portions of a surface from a data set. In the case of data of known reliable accuracy, the routine also has the capacity to filter object points into separate point clouds, for subsequent modelling. In this case, the routine would be used to transfer surface points from one point cloud, in which they appear obliquely oriented relative to the camera network, to a point cloud in which they appear less oblique. This routine should be used conservatively, taking into account other factors such as known matching accuracy and known geometric properties of the measurement object. Mismatch removal by inspection is a fool-proof, but tedious, method of removing obvious erroneously matched points. This method could be used in conjunction with the oblique surface threshold routine.

The rendering of the surfaces, in both conventional wire-frame format and in *CosmoWorlds*, was successfully performed but the final accuracy achieved would not have been acceptable in an industrial or engineering application. However, the technique has been developed here, and higher resolution CCD cameras and more advanced software would improve this shortfall of the system.

The reconstruction of an object from a random 3-D point cloud is anything but a trivial task. A survey of commercially available software revealed that, to the best of this author's knowledge, no software is capable of successfully reconstructing a model of a measured object automatically. The approach taken here of rendering separate DSMs for portions of the surface, and then translating and rotating these

'faces' into a single coordinate system has potential. However, the work of Hoppe (1994), which is commercially unavailable, appears one of the most promising, and would perfectly complement the bulk of the work presented here. Hoppe has developed software capable of regenerating a 3-D surface from an irregular data set, without requiring any *a priori* knowledge of the geometric properties of the surface. Use of Hoppe's routine, in conjunction with the photogrammetric method presented in this thesis, would mean that, even without automation, an object's surface could be imaged, measured photogrammetrically with image matching, and reliably reconstructed in a reasonably short time. Automation of parts of the process would reduce this time considerably.

The shortcomings of the method presented in this thesis include:

- *In situ* and real-time measurement is not possible (although this tends to be the case with most all-round photogrammetric surveys).
- Use of the rotary table for control field and object positioning requires a mounting arrangement.
- Acquisition of one image at a time – as dictated by the framegrabber card – is not ideal, due to temporal changes in the image capture conditions.
- The method of point cloud refinement does not offer a completely secure means of mismatch removal, thus requiring conservative use.
- A need for a reliable method of interpreting the acquired surface data, either in a reverse engineering form (e.g. *AutoCAD*) or in a purely visual form (e.g. *CosmoWorlds*), was not satisfactorily met.

7. Recommendations

The preceding discussion, conclusions and, particularly, the list of the shortcomings of this work, give rise to a number of recommendations for further research.

The method presented in this thesis does not permit *in situ* all-round measurement. This is because the object has to be moved and mounted in the rotary table before it can be measured. Therefore, with this unavoidable reduction in time-efficiency, the need for near real-time data production becomes all the more pressing. As discussed in Chapter 6, the method presented in this thesis does have excellent potential for automation in the object placement, image capture and subsequent data processing stages. A system, incorporating automated object rotation and image capture, with automatic file generation (naming of the image file, and storage on the hard drive), would greatly reduce image acquisition times which, even without automation, could be considered rapid.

A framegrabber which would permit simultaneous image acquisition would, firstly, eliminate temporal changes in image capture conditions and, secondly, reduce image acquisition time yet further. For example, a framegrabber with three channels allowing for simultaneous image capture, in conjunction with automated object positioning and automatic image naming and file storage, would reduce the image capture time drastically.

A further recommendation needs to be made regarding software and processing. Currently, the processing of images and subsequent data is probably the most time-consuming stage in the digital photogrammetric process. This is partly due to the intensive processing required in, for example, the bundle adjustment and the image matching procedures. But, the lack of a single software package, running on a single platform and capable of performing all the necessary processing functions, severely hinders the digital photogrammetric procedure. For instance, in the UCT Geomatics Department's Digital Photogrammetry Group, the bundle adjustment is currently performed on the *DOS* platform, while the image matching is currently run under *Linux*. Software linked into one package, on one platform, would greatly reduce data

preparation and processing times. The need, in the author's opinion, is for a suite of software applications enabling the photogrammetrist to perform image enhancements, thresholding, target identification, DLT, bundle adjustment, edge/feature detection, image matching, point cloud refinement, and possibly data presentation on one dedicated PC.

References and Selected Bibliography

Ackermann, F. ***Techniques and Strategies for DEM Generation.*** Digital Photogrammetry – An Addendum to the Manual of Photogrammetry. ASPRS. Maryland. 1996.

Adams, L. P. ***Industrial Photogrammetry.*** Non-topographic Photogrammetry, Second Edition. Karara, H.M. (Ed.) ASPRS. Virginia. 1989.

Allison, D. *et al.* ***Matrix Methods for Engineers.*** UCT. Cape Town. 1993.

Atkinson, K. B. (Ed.). ***Close Range Photogrammetry and Machine Vision.*** Whittles Publishing. Caithness, Scotland. 1996.

Baltsavias, Emmanuel P. ***Multiphoto Geometrically Constrained Matching.*** ETH-Zürich. 1991.

Budd, Timothy A. ***Classic Data Structures In C++.*** Addison-Wesley Publishing. Massachusetts. 1994.

Chikatsu, H., Anai, T. and Hatano, K. ***Real-time Ortho Imaging and Surface Modeling for Archeological Artefacts.*** Department of Civil Engineering. Tokyo Denki University. 1997

Cignoni, P., Montani, C., Scopigno, R. ***A Merge-First Divide and Conquer Algorithm for 3-D Delaunay Triangulation.*** CNUCE Internal Report C92/16. National Research Council of Italy. 1992.

Ekstrom, Michael P. ***Digital Image Processing Techniques.*** Academic Press. Florida. 1984.

Greve, Clifford W. (Ed.) ***Digital Photogrammetry – An Addendum to the Manual of Photogrammetry.*** ASPRS. Maryland. 1996.

Gruen, A. and Karmen, H. (Eds.) ***Optical 3-D Measurement Techniques II.*** Wichmann. Karlsruhe. 1993.

Haala, N. and Brenner, C. ***Unwrapping of Detailed Surface Models.*** GIM International. Vol. 13 No. 3. March 1999.

Haralick, R. and Shapiro, L. ***Computer and Robot Vision, Vol. 1-2.*** Addison-Wesley. 1992.

Hartman, J. and Wernecke, J. ***The VRML 2.0 Handbook.*** Addison-Wesley. 1996.

Hoffa, Pierre G. J. ***The Mensuration of Textured Surfaces Using Close-Range Digital Photogrammetry for Applications in Industry.*** UCT. Cape Town. 1996.

Hoppe, Hugues. ***Surface Reconstruction from Unorganized Points.*** Ph.D. Thesis, Department of Computer Science and Engineering, University of Washington, 1994.

Karara, H. M. (Ed.) ***Non-Topographic Photogrammetry.*** ASPRS. Virginia. 1989.

Kraus, Karl. ***Photogrammetry, Vol. 2.*** Dummler. Bonn. 1997.

Mason, Scott O. ***Expert System-Based Design of Photogrammetric Networks.*** ETH-Zürich. 1994.

Maune, David F. ***Introduction to Digital Elevation Models.*** Digital Photogrammetry – An Addendum to the Manual of Photogrammetry. Greve, Clifford W. (Ed.) ASPRS. Maryland. 1996.

Meyer, R. ***The Present State in Industrial Photogrammetry***. Surveying News. No.17. 1973.

Salas, Saturnino L. and Hille, Einar. ***Calculus: One And Several Variables.*** John Wiley & Sons. New York. 1982.

Schenk, Anton F. ***Automatic Generation of DEMs.*** Digital Photogrammetry – An Addendum to the Manual of Photogrammetry. Greve, Clifford W. (Ed.) ASPRS. Maryland. 1996.

Shammas, Namir C. ***Teach Yourself Turbo C++ In 21 Days.*** SAMS Publishing. Indiana. 1994.

Smit, Julian L. ***The Three Dimensional Measurement Of Textured Surfaces Using Digital Photogrammetric Techniques.*** UCT. Cape Town. 1997.

Smit, Julian L. ***Digital Photogrammetry.*** UCT. Cape Town. 1996.

Swokowski, Earl W. ***Calculus With Analytic Geometry.*** PWS-Kent. Massachusetts. 1988.

Tabatabai, A. J. and Mitchell, O. R. ***Edge Location to Subpixel Values in Digital Imagery.*** Transactions on Pattern Analysis and Machine Intelligence. Vol. PAMI-6, No. 2. 1984.

Van Der Merwe, N. ***Devlopment Of An Image Matching scheme Using Feature- and Area Based Matching Techniques.*** UCT. Cape Town. 1995.

Van Der Vlugt, Graeme. ***The Development Of A Real-Time Photogrammetric System For Patient Positioning In Proton Therapy.*** UCT. Cape Town. 1991.

Van Der Vlugt, Graeme. ***Algorithms And Design Aspects Of An Automated Vision Based 3-D Surface Measurement System.*** UCT. Cape Town. 1995.

Yi-Hsing, Tseng. ***Digital Photogrammetric Approach to Ice-Flow Determination in Antarctica.*** Report No. 422. Department of Geodetic Science and Surveying. Ohio State University. Columbus, Ohio. September 1992.

APPENDICES

Appendix 1: Digital Photogrammetry Group File Formats

All files are text files.

3-D Coordinate / Control Point File

File extension: .crd
File format: *P* *X* *Y* *Z* *Type*
: : : : :

This file contains a 3-D coordinate listing of all control and object points observed. *Type* is a binary value and differentiates between fixed (control) points, denoted by 1, and free (object) points, denoted by 0. *P* is the point number.

Digital Parameters

File extension: .dig
File format: h
v
s_x
s_y

This file contains information about the sensor array, which is necessary for transforming pixel coordinates into metric image coordinates. Here, h and v are the number of columns and rows in the image, and s_x and s_y are the pixel dimensions in the x- and y-directions.

Camera Interior Orientation

File extension: .io
File format: x_p
y_p
c

This file contains the three interior orientation parameters (principal point position and principal distance) of each camera used.

Camera Exterior Orientation

File extension: .eo
File format: Camera number
ω
κ
ϕ
X_c
Y_c
Z_c

This file contains the six exterior orientation parameters and the camera number, which identifies different cameras where more than one has been used.

Image Coordinates

File extension: .obs
File format: *P* *x* *y*
: : :

This file contains the image coordinates of all the observations made at a given camera station. The point number, *P*, corresponds to the point number in the 3-D coordinate file.

Digital Elevation Model Point Coordinates

File extension: .dem
File format: *P* *X* *Y* *Z* *T*
: : : : :

This file is produced by the matching program. *P* is the number of the point, *X*, *Y* and *Z* are the coordinates of the point, and *T* is a binary number referring to the estimated accuracy of the match, where 0 denotes a poor match and 1 a good match.

Appendix 2: All-Round Control Frame

The six faces of the control frame are shown along with the point numbering.

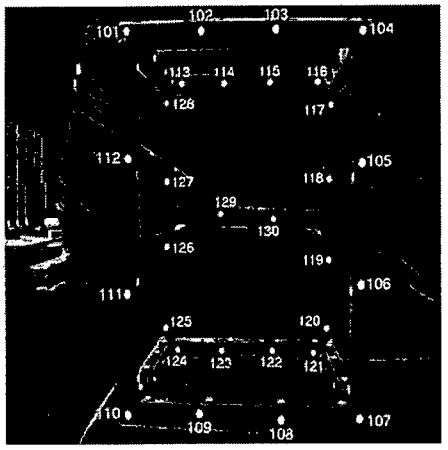

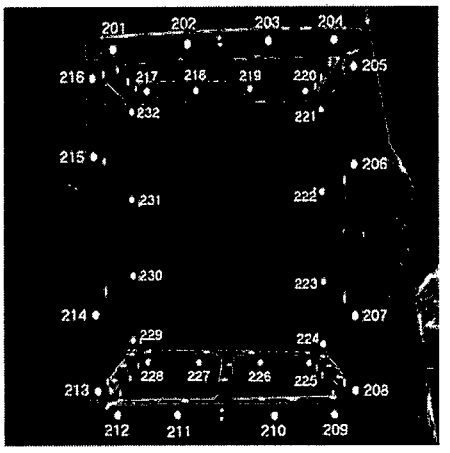

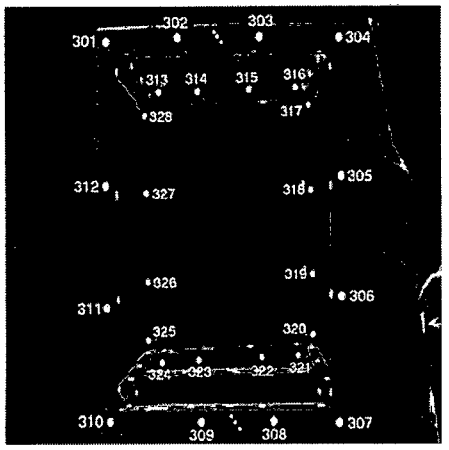

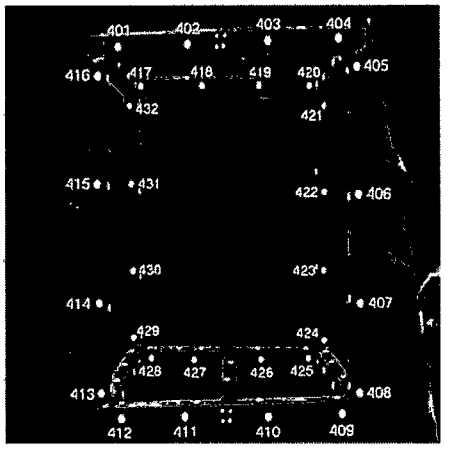

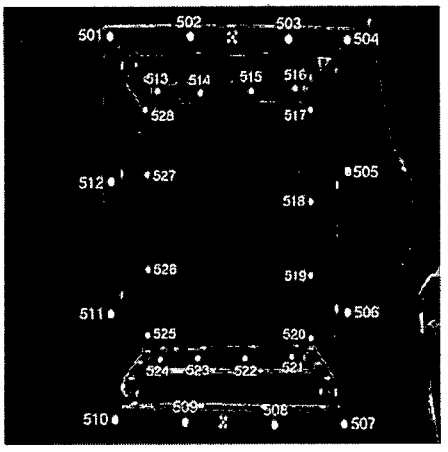

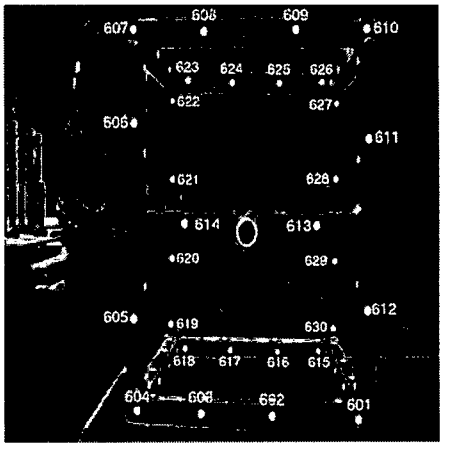

The control point coordinates of the all-round control frame used in this project are given below.

Point Number	X	Y	Z
101	987.38	1344.83	1403.12
102	1097.22	1339.19	1405.04
103	1196.29	1335.04	1406.29
104	1312.47	1327.35	1407.43
105	1306.12	1208.23	1403.79
106	1298.24	1076.05	1399.82
107	1294.4	952.47	1396.08
108	1190.03	956.5	1396.46
109	1079.02	965.95	1397.26
110	972.8	967.32	1398.29
111	978.81	1088.92	1399.97
112	985.32	1217.3	1401.47
113	1021.92	1350.18	1029.65
114	1103.82	1344.04	1030.47
115	1203.72	1340.59	1031.75
116	1291.11	1337.03	1032.82
117	1314.06	1304.33	1032.5
118	1305.9	1191.08	1028.94
119	1301.56	1100.04	1025.81
120	1295.76	990.36	1021.82
121	1265.97	963.02	1020.81
122	1182.79	968.19	1021.73
123	1092.35	970.99	1023.01
124	998.71	974.37	1024.61
125	974.15	1006.98	1025.01
126	981.92	1110.54	1026.25
127	987.04	1222.31	1027.69
128	993.81	1318.23	1029.25
129	1093.63	1161.31	1027.62
130	1197.3	1149.82	1027.53
201	974.93	1313.9	1388.3
202	969.46	1205.79	1388.22
203	964	1091.12	1387.17
204	959.7	997.28	1386.15
205	958.24	970.53	1358.65
206	957.51	973.92	1260.27
207	956.42	976.3	1109.41
208	956.25	977.7	1034.59
209	960.03	1007.55	1010.58
210	965.46	1093.9	1011.92
211	974.47	1233.37	1015.56
212	980.09	1321.45	1016.4
213	981.25	1349.27	1041.55
214	980.61	1349.08	1120.42
215	979.11	1345.88	1280.85
216	978.57	1345.35	1360.5
217	1298.67	1294.26	1392.84
218	1293.75	1198.07	1390.15
219	1288.35	1092.3	1388.41
220	1283.05	985.09	1382.26
221	1281.95	955.44	1355.64

222	1282.09	957.41	1242.68
223	1282.54	957.78	1119.79
224	1282.64	960.45	1034.77
225	1283.53	988.65	1009.62
226	1288.65	1084.43	1012.87
227	1295.9	1204.85	1016.25
228	1301.5	1306.1	1019.28
229	1305.65	1334.46	1049.77
230	1304.36	1330.43	1139.36
231	1302.91	1330.12	1243.97
232	1300.54	1326.11	1365.59
301	1313.39	1338.46	1393.44
302	1213.74	1343.92	1394.37
303	1101.85	1349.79	1392.09
304	992.59	1355.25	1389.81
305	993.31	1358.55	1257.13
306	994.22	1360.74	1139.59
307	997.68	1362.81	1014.8
308	1089.05	1358.68	1018.19
309	1189.42	1353.67	1018.62
310	1318	1346.46	1018.88
311	1318.46	1344.57	1133.15
312	1316.73	1341.4	1253.78
313	1268.22	967.8	1386.43
314	1189.39	972.03	1384.92
315	1088.04	976.72	1385.34
316	998.73	981.37	1387.12
317	971.78	983.47	1361.49
318	970.6	986.61	1244.27
319	969.52	988.52	1125.97
320	969.41	990.8	1040.28
321	998.92	989.75	1011.16
322	1070.82	986.25	1009.76
323	1195.76	979.47	1007.81
324	1270.13	975.68	1005.02
325	1295.88	973.61	1036.87
326	1294.83	971.14	1120.28
327	1295.04	969.14	1246.08
328	1295.44	967.7	1353.71
401	1309.2	984.54	1383.38
402	1313.9	1083.34	1385.78
403	1319.31	1196.73	1390.54
404	1324.26	1296.75	1395.24
405	1325.75	1323.73	1368.01
406	1328.17	1326.88	1241.42
407	1330.05	1331.39	1131.79
408	1331.12	1333.26	1041.02
409	1327.71	1308.85	1019.11
410	1321.6	1202.91	1015.08
411	1314.54	1082.96	1013.12
412	1308.97	990.76	1008.39
413	1308.8	961.37	1034.32
414	1308.82	959.07	1126.46
415	1308.52	955.57	1247.52
416	1308.12	955.37	1354.45
417	984.78	989.49	1387.53
418	990.62	1108.34	1388.41
419	996.17	1218.75	1389.95

420	1001.5	1316.04	1390.65
421	1004.86	1345.1	1363.66
422	1005.27	1346.7	1247.9
423	1006.33	1346.89	1140.27
424	1007.47	1350.55	1042.68
425	1005.9	1319.24	1017.05
426	999.96	1227.61	1014.45
427	991.34	1096.45	1011.66
428	985.98	1012.06	1012.27
429	981.59	976.82	1039.92
430	982.07	975.45	1133.77
431	982.72	971.5	1254.16
432	983.63	967.47	1360.03
501	968.08	955.95	1386.46
502	1078.85	950.75	1385.32
503	1213.35	944.8	1382.52
504	1295.04	941.41	1382.64
505	1294.79	943.79	1255.1
506	1296	946.53	1117.69
507	1292.89	949.62	1005.97
508	1195.37	954.03	1005.94
509	1070.77	960.11	1008.41
510	972.36	965.21	1009.86
511	968.58	962.88	1115.79
512	969.59	959.69	1245.4
513	1012.34	1328.82	1392.53
514	1096.85	1324.45	1389.28
515	1197.55	1318.95	1392.58
516	1284.4	1314.16	1395.9
517	1314.2	1311.31	1365.44
518	1315.34	1315.79	1238.32
519	1315.65	1318.65	1134.19
520	1317.23	1320.9	1045.31
521	1279.13	1323.01	1018.46
522	1186.37	1328.08	1017.29
523	1092.29	1332.96	1016.79
524	1017.62	1336.68	1014.43
525	992.5	1336.86	1049.04
526	993.88	1335.28	1142.92
527	992.15	1332.65	1276.23
528	988.13	1330.68	1366.11
601	1307.75	1337.74	1007.32
602	1187.22	1342.06	1005.81
603	1088.14	1347.63	1004.61
604	996.55	1351.95	1003.66
605	987.64	1259.8	1002.52
606	976.65	1068.52	1000.08
607	969.89	976.73	999.42
608	1069.71	975	997.84
609	1195.33	966.93	995.88
610	1291.63	960.04	994.57
611	1300.51	1064.64	998.95
612	1311.05	1228.92	1004.35
613	1235	1151.47	1001.74
614	1053.59	1161.38	1001.81
615	1285.47	1328.02	1382.03
616	1203.77	1334.49	1381.24
617	1109.34	1339.37	1379.73

618	1018.73	1343.71	1378.13
619	987.96	1310.26	1377.5
620	984.6	1217.04	1376.5
621	978.79	1106.15	1375.19
622	971.87	997.36	1373.31
623	1001.73	966.9	1372.73
624	1090.22	965.1	1371.99
625	1183.04	960.61	1371.28
626	1266.38	954	1370.41
627	1296.65	982.6	1371.41
628	1301.77	1085.34	1375.08
629	1307.95	1198.97	1378.93
630	1310.02	1292.38	1371.25

Appendix 3: Point Cloud Refinement Routine – *C++*

```
/*PROGRAM EXCL: Excludes D.E.M. points that are located on
surfaces forming too small an angle with the object ray...
The centre of the network of cameras used for the capture of the
images which produced the matched points is estimated from the
relevant exterior orientation files. */

#include <iostream.h>
#include <stdlib.h>
#include <stdio.h>
#include <math.h>
#include "matrix.h"      //matrix class allowing matrix operations

double Xn, Yn, Zn,       //network centre -> needed later!
x, y, z;                 //current point
const Size = 15000;      //expected max. no. of matched points
const Pi = 3.1415927;    //constant
double demx[Size],       //declare arrays for x, y and z
demy[Size],
demz[Size],
residual[Size];          //array of residuals
int no_of_points, p, pass;
char demFile[40];        //name of dem file

//************************************************
//READING THE EXTERIOR ORIENTATION FILE PARAMETERS

void eoRead()
 {
 FILE *eo;                          //file stream for .eo file
 double omega, kappa, phi,          //local vars -> not needed later!
 Xc, Yc, Zc;                        //persp centre of each camera summated
 int cameraNumber, numStations;     //not needed later
 char eoFile[40];                   //filename

 cout << "\nNetwork Centre Location";
 cout << "\nEnter the number of camera stations : ";
 cin >> numStations;

 Xc = Yc = Zc = 0;       //this block of code calculates the network centre
 for (int station = 0; station < numStations; station++)
  {
  cout << "\nEnter name of exterior orientation file No. " << station+1 << " : ";
  cin >> eoFile;
  eo = fopen(eoFile, "r");
  cout << "\nReading file ... \n";
  if (! eoFile)
   {
   cout << "\nCannot open file to read!!!\n";
   exit(1);
   }
  fscanf(eo, "%d", &cameraNumber);
  fscanf(eo, "%lf", &omega);
  fscanf(eo, "%lf", &kappa);
```

```
 fscanf(eo, "%lf", &phi);
 fscanf(eo, "%lf", &Xn);
 fscanf(eo, "%lf", &Yn);
 fscanf(eo, "%lf", &Zn);
 fclose(eo);
 Xc += Xn;
 Yc += Yn;
 Zc += Zn;
 }
Xn = Xc/numStations;
Yn = Yc/numStations;
Zn = Zc/numStations;
cout << "\nCamera network centre is at " << Xn << ", " << Yn << ", " << Zn  << "\n";
};

//*****************************************
//READING THE x, y, z COORDS FROM THE D.E.M.
void demRead()
{
FILE *dem; //file stream for .dem input file
double res; //current residual
int pointNumber, tag;

cout << "\nMatched Point Cloud";
cout << "\nEnter name of the file containing matched points: ";
cin >> demFile;
cout << "\nEnter number of matched points in " << demFile << " : ";
cin >> no_of_points;
dem = fopen(demFile, "r");
cout << "\nReading file ... \n";

if (! demFile)
 {
 cout << "\nCannot open file to read!!!\n";
 exit(1);
 }

for(p = 0; p < no_of_points; p++)
 {
 fscanf(dem, "%d", &pointNumber);
 fscanf(dem, "%lf", &x);
 fscanf(dem, "%lf", &y);
 fscanf(dem, "%lf", &z);
 fscanf(dem, "%d", &tag);
//  fscanf(dem, "%lf", &res);
 demx[p] = x;
 demy[p] = y;
 demz[p] = z;
//  residual[p] = res;
 }
fclose(dem);
};

//*************************************************************
//SORTING THE D.E.M. FILE POINTS; OUTPUTTING THE NEW D.E.M.

void demSort()
```

```
{
FILE *out;                      //file stream for output of sorted points
char outFile[40];
const unsigned N = 5;           //No. of points used in L.S. procedure
double thetaThreshold;          //incidence angle threshold = degrees

thetaThreshold = 100.0;
while (thetaThreshold < 0.0 || thetaThreshold > 90.0)
 {
 cout << "\nEnter incidence angle threshold for retaining points";
 cout << "\n0 (lenient) - 90 (strict): ";
 cin >> thetaThreshold;
 }

cout << "\nRefined Point Cloud";
cout << "\nEnter name and extension of output file : ";
cin >> outFile;
out = fopen(outFile, "w");
if (! outFile)
 {
 cout << "\nProblem opening output file !\n";
 exit(1);
 }
cout << "\nProcessing file ... \n";

matrix A(N, 3);     //declaring matrices for LS plane generation
matrix At(3, N);
matrix Z(N, 1);
matrix X(3, 1);
matrix Points(N, 3);
matrix Vector(3, 1);
matrix Normal(3, 1);
double distance, VdotN, cosTheta, lengthV, lengthN,
tempOut, tempIn, ptTempIn[3], ptTempOut[3], min_dis[N-1];
int i; //to search points for closest
pass = 0;

for (p = 0; p < no_of_points; p++)
 {

 for (int n = 0; n < N; n++)
  {
  min_dis[n] = 1000.0 + n*1.1;          //set arb min dists betwn pts
  }

 x = demx[p]; Points.poke(0, 0, x);
 y = demy[p]; Points.poke(0, 1, y);     //centre point placed in first location
 z = demz[p]; Points.poke(0, 2, z);
 for (i = 0; i < no_of_points; i++)
  {

//Beginning IF statement, repeated for all points...
  if (i != p) //precludes finding dis between point and itself i.e. 0
   {//IF -> BEGIN
        distance = sqrt( pow((x-demx[i]),2) + pow((y-demy[i]),2)
        + pow((z-demz[i]),2) );
        tempIn = distance;
```

```
        ptTempIn[0] = demx[i];
        ptTempIn[1] = demy[i];
        ptTempIn[2] = demz[i];
        int flag;
        flag = 0;
        for (unsigned long q = 0; q < N; q++)
         {
         if (distance < min_dis[q] && flag == 0)
          {
          flag = 1; //i.e. distance IS < min_dis[q]!
          for (unsigned long r = 0; r < N-q-1; r++)
           {
           tempOut = min_dis[q+r];      //shuffling minimum distances
           min_dis[q+r] = tempIn;
           tempIn = tempOut;
           //**********************//
           for (int s = 0; s < 3; s++)
            {
            ptTempOut[s] = Points.peek(q+r+1, s); //shuffling N closest pts
            Points.poke(q+r+1, s, ptTempIn[s]);
            ptTempIn[s] = ptTempOut[s];
            }
           }
          }
         }
 }//IF -> END
}//CLOSE -> for (i = 0; i < no_of_points; i++)

for (unsigned long j = 0; j < 5; j++)  //setting matrix elements
 {
 A.poke(j, 0, (Points.peek(j, 0)));     //assigning x's
 A.poke(j, 1, (Points.peek(j, 1)));     //assigning y's
 A.poke(j, 2, 1.0);                     //assigning 1's
 Z.poke(j, 0, (Points.peek(j, 2)));     //assigning z's
 }

 At = !A;                               //transpose calculated
 X = ( ~(At*A) * At ) * Z;              //matrix operation...

 for (j = 0; j < 2; j++)
 {
 Normal.poke(j, 0, (X.peek(j, 0)));     //creating Normal vector
 }

 Normal.poke(2, 0, -1.0);
 Vector.poke(0, 0, (x - Xn));           //creating object ray vector
 Vector.poke(1, 0, (y - Yn));
 Vector.poke(2, 0, (z - Zn));
 VdotN = dot(Vector, Normal);           //dot product calculated
 double termV, termN;
 termV = 0;
 termN = 0;

 for (j = 0; j < 3; j++)
 {
 termV += pow(Vector.peek(j,0),2);
 termN += pow(Normal.peek(j,0),2);
```

```
    }

    lengthV = sqrt(termV);
    lengthN = sqrt(termN);
    cosTheta = VdotN/(lengthV*lengthN);          //-1 to 1
    double theta;
    theta = (acos(cosTheta))*(180.0/Pi);         //arccos = in degrees

    if (theta < 90.0)
     {
     theta = 90.0 - theta;
     }
    if (theta > 90.0)
     {
     theta = theta - 90.0;
     }

//Must be that Theta < 90, or how did camera "see" point?
    if (theta >= thetaThreshold) //incidence angle >= threshold
     {
     pass += 1; //counter for points that pass the "Theta test"
     fprintf(out, "%-7d" " " "%-10.3lf" " " "%-10.3lf" " " "%-10.3lf" " " "%-5.3lf" "\n",
     pass, x, y, z, theta);

/*[-] = pad to the right with blanks
  [10] = total width of output, incl blanks
  [.3] = no. of decimals
  [d] = integer, [lf] = double*/
     }

  }//CLOSE -> for (p = 0; p < no_of_points; p++)
 fclose(out); //output stream closed
 cout << "\nNumber of points excluded from " << demFile << " = "
 << no_of_points - pass << "\n";
};

/*********************************************************************/
void main()
 {
 eoRead();              //Ext. orientation parameters are read ...
 demRead();             //Matched points are read and stored ...
 demSort();             //Points are searched and excluded ...
 }
/*********************************************************************/
```