

Registering 3D Point Clouds: An Experimental Evaluation

Building and Fire Research Laboratory
Gaithersburg, MD 20899



United States Department of Commerce
Technology Administration
National Institute of Standards and Technology



Registering 3D Point Clouds: An Experimental Evaluation

Christoph Witzgall
Geraldine S. Cheok

May 2001
Building and Fire Research Laboratory
National Institute of Standards and Technology
Gaithersburg, MD 20899



U. S. Department of Commerce
Donald L. Evans, *Secretary*

National Institute of Standards and Technology
Karen H. Brown, *Acting Director*

Abstract

Four separate laser scans of a wooden box, taken from different vantage points, were examined in a laboratory setting. Visual and numerical registration methods, aimed at aligning the individual scan data with respect to a common frame, were explored. The numerical registration method aligns point clouds with triangulated elevated surfaces optimizing residual-based measures-of-fit, as outlined in this report. Essential procedures for data filtering are described including methods for shadow delineation. Data phenomena beyond common noise were observed, particularly, “phantom points” resulting from split signals (mixed pixels). Such points interfere with determining occlusions. Results from four experiments applying numerical and visual procedures are reported.

Key words: Delaunay triangulation, elevated surfaces, LIDAR, mixed pixel, occlusion, registration, residual measure-of-fit, shadow delineation, split signal, TIN.

Table of Contents

1.	Introduction	1
1.1	Organization of the Report.....	2
2.	General Approaches to Spatial Registration	3
2.1	TIN Techniques	4
2.2	Scanning a Wooden Box.....	5
2.3	Outline of Experiments to be Reported	10
3.	Elevated Surfaces and Delaunay TINs	11
3.1	Delaunay Triangulations.....	12
3.2	The Insertion Method for Delaunay TINs	13
3.3	Termination.....	14
3.4	Options for TIN Generation.....	15
3.5	TIN Boundaries; Shadow Delineation	15
3.6	RMS Adjustments.....	18
4.	Point Cloud Against Surface Registration.....	19
4.1	The Transformation Mechanism.....	19
4.2	A Search Procedure.....	20
4.3	Editing Procedures for Point Clouds.....	21
4.4	Remarks about Medians.....	22
5.	Experimental Results.....	25
5.1	Volume calculations.....	26
5.2	Cleaning Procedure	27
5.3	Experiment 1	27
5.3.1	Experiment 1.1	28
5.3.2	Experiment 1.2.....	30
5.3.3	Experiment 1.3.....	32
5.3.4	Discussion of Experiment 1	33
5.4	Experiment 2.....	34
5.5	Experiment 3.....	36
6.	Conclusions	39
	References.....	40

1. Introduction

The need for investigating and implementing methods for “registering” two or more separate 3-D data sets was encountered during the conduct of the *Non-Intrusive Scanning Project* at the National Institute of Standards and Technology (NIST) [8]. The objective of that project was to contribute to the development of tools for automatically monitoring progress at large construction sites. A major component of such a monitoring system would have to be the numerical assessment and computer modeling of construction terrain.

The work in the project led to the investigation of the utility of laser scanners to obtain 3-D surface information in the form of

“point clouds”.

These are collections of x , y , z data points. Since a laser scanner is a line-of-sight instrument, there will, in general, be incomplete information about any object that is entirely or partially shadowed by another object in front of it. To obtain data coverage without occluded areas, scans from different vantage points have to be obtained.

Point clouds from different vantage points have different reference frames since the reference points for the respective point cloud coordinates are located at these different vantage points and since the orientation of the scanning instrument may also vary. The x , y , z coordinates of the data points will therefore relate to different coordinate systems. The task of

“registration”,

as understood in this report, is to apply the necessary rigid coordinate transformations -- rotations and translations -- to two or more point clouds so that their respective coordinates refer to a common coordinate system.

Two registration tasks need to be distinguished:

“spatial registration” and *“temporal registration”*.

Spatial registration refers to the alignment of two or more scans of a fixed scene when several scans are needed to deal with obstructions of view. Once the separate point clouds are referenced with respect to a single coordinate system, they can be combined to provide a data model of the entire scene. Spatial registration thus addresses the generation of a data model for a given fixed scene.

In situations where the scene changes with time, the emphasis is on comparisons between scenes in order to capture temporal developments. Thus temporal registration is needed to tie the models of scenes that have changed over time to a common frame of reference. An example would be a construction site where excavation is ongoing, and where it is desired to determine daily the amount of cut and fill by comparing subsequent models of construction terrain. In that application, temporal registration would, of course, not be necessary if the positions of the scanning instruments

were not to change between the days of interest. This, however, may not be practical because it requires the availability of several laser scanners, the protection of the scanning locations, and the correct anticipation of developing obstructions.

For both registration tasks, specific and readily identifiable targets are commonly used to aid in the registration process. The reference coordinates of such targets are assumed to be known exactly. Three or, preferably, more such targets then permit setting up the correspondence between the scan coordinates and the desired reference coordinates. In spatial registration, any distinctive object or feature that is part of the scene itself would be the obvious choice as a target. In temporal registration, such an object or feature may not be present in all the scenes as they change. There may therefore be a need to set up particular artifacts as permanent targets.

There are occasions which call for spatial registration without the benefit of targets. One such occasion was encountered on the NIST Gaithersburg campus [9]. This effort required a

“target free” or “free form”

registration approach. Such registration without recourse to specific targets will be addressed in this report. Its purpose is to provide an update on ongoing work at NIST on that problem.

1.1 Organization of the Report

In what follows, Chapter 2 briefly discusses the different approaches to registration, the routines used for registration, and a general description of the experiments. Chapter 3 introduces the concepts of triangulated elevated surfaces and Delaunay TINs. Also, an adaptive insertion procedure for constructing Delaunay TINs from a given point cloud is proposed. This is followed by a discussion of options for subsequently editing TINs. These options address surface adjustments for minimizing RMS error and also the important issue of delineating the boundary of data sets in the footprint plane, including determination of shadows, that is, areas of occlusion. Chapter 4 describes the actual registration process implemented as the routine *TINregister* and applied in the course of this work. Experimental results, finally, are reported in Chapter 5, followed by concluding remarks in Chapter 6.

2. General Approaches to Spatial Registration

The registration problem has attracted considerable attention during the past decade as the selected references indicate. Between particular applications, there are critical differences as to methods of data collection, numbers and densities of data points, statistical attributes, such as noise or the prevalence of outliers, and geometrical object characteristics, demanding different methods for registration. Applications may call for registering

- *point cloud against point cloud,*
- *point cloud against surface,*
- *surface against surface.*

In general, three major phases of the registration process may be considered ([17])

- *matching,*
- *registering,*
- *integration.*

The first phase refers to the task of moving separate data into a rough correspondence. Fine-tuning transformations during the second phase aims at achieving the most accurate registration possible. Integration of two registered surfaces into a single combined surface is still a major topic of research, and will not be considered here.

In each of the above phases, several alternate approaches may be followed, sometimes in combination:

- *direct measurement,*
- *visual adjustment,*
- *numerical adjustment.*

Direct measurement is based on a physically established reference coordinate system. It uses mechanical devices such as survey instruments or the global positioning system (GPS) to locate the scanning instrument. It also requires that the scanning instrument be oriented by aiming it at some known point. Location and orientation of the instrument are all that is needed to determine the coordinate transformation which registers the scan coordinates with respect to the reference coordinates. Errors from this type of registration are linked to the initial set-up of the instrument, i.e., how well it is leveled and aimed, and how accurately the distance between the instrument and the GPS antennae has been determined. Direct measurement may thus be considered as a matching process.

Visual adjustment is a manual process by which computer images of point clouds are moved with respect to one another, or with respect to a visual model, so as to achieve visual agreement in areas where they refer to identical portions of object or terrain. The parameters of that movement, that is, the rotations and translations involved, are recorded. They define the rigid coordinate transformations which are expected to achieve alignment. Visual adjustments are labor intensive

and limited by the fact that computer images represent 2-D projections of data clouds. It may be necessary to employ different projections. Visual adjustment, particularly, when following direct measurement, may be considered as a registration process.

Numerical adjustments may also be considered as a registration process. In most instances, numerical adjustments are based on distance measures, say between point clouds, surfaces, or between point clouds and surfaces, often relying on the “iterative closest point (ICP)” algorithm [5]. Such measures are numerically evaluated and optimized again by suitable coordinate transformations. Ideally, the optimization is accomplished automatically by an optimization algorithm. In general, numerical adjustments may be used either in lieu of visual adjustments or as a fine-tuning device following visual adjustment.

2.1 TIN Techniques

This report addresses distance based numerical adjustments for aligning point clouds with surfaces in the context of spatial registration. The registration process will be free-form, that is, without recourse to known targets.

The authors have developed and are evaluating a particular approach to that task, implemented as a prototype FORTRAN routine

“TINregister”.

The routine is based on the TIN (Triangulated Irregular Network) technique for generating and representing *elevated triangulated surfaces* (see Chapter 3). The first part of the routine prepares a TIN surface either from a data cloud or a model description. It then considers a separate point cloud which is to be registered to that TIN surface. It performs specified rigid coordinate transformations, that is, rotations followed by translations, and then measures how well the transformed point cloud aligns with the surface. Three “measures-of-fit” based on vertical deviations or,

“residuals” $z_i - \hat{z}_i$,

of the points (x_i, y_i, z_i) in the data cloud are examined. Here \hat{z}_i is the elevation of the TIN surface at the footprint point (x_i, y_i) . The three measures are:

- “maximum deviation (MAX)”: The size -- or absolute value -- of the largest residual.
- “root mean square (RMS)”: The square root of the mean of the squared residuals.
- “average size deviation (ASD)”: The average -- or mean -- of the absolute values of the residuals.

The transformations are specified manually, following repeatable prompts. *TINregister* evaluates quantitatively the quality of a specified alignment of a data cloud with a specified TIN surface where each alignment is given by values for the x, y, z coordinates and for the three angles, ϕ, ϵ, θ ,

which represent rotations around three coordinate axes (see Section 4.1 for angle definitions). Based on that information, the user may choose new values for these six transformation parameters.

As it stands, the routine can be used for manually improving alignments. Automating the process has been deliberately deferred pending a better understanding of the mechanisms involved and the relationship of the selected measures-of-fit to the desired quality of the registration.

A related FORTRAN routine,

“TINvolume”

also creates a TIN surface model for a point cloud by essentially the same process as does *TINregister*. But whereas *TINregister* proceeds to compare transformations of a second point cloud against that model, *TINvolume* generates displays and calculates volumes.

Each volume to be calculated is the volume of the solid body bounded above by the TIN surface and below by a horizontal plane of specified elevation, the

“floor level”.

Volume calculations are sensitive to this elevation, which may be obtained from direct knowledge of the elevation of the scanning instrument or be derived by data analysis (see Section 5.3).

2.2 Scanning a Wooden Box

The work reported here describes several preliminary approaches that were tried in order to generate a 3D TIN surface model by scanning, in a laboratory setting, a

0.914 m × 1.219 m × 1.524 m (3 ft × 4 ft × 5 ft)

wooden box from four separate vantage points and, subsequently, re-determine sizes and volume of the box from the scan data.

Examining the resulting point clouds revealed data phenomena of significant impact beyond the inevitable data inaccuracies and anomalies, such as noise, that will, of course, influence the quality of registration and volume determination.

The scanning instrument measures, for a particular direction, the distance between the laser and the object by measuring the time-of-flight of the signal (see Ref. 2 for description of instrument). The instrument points along a direction specified by a “planar” angle and a “polar” angle (analogous to the geographical quantities of “longitude” and “latitude”).

While angle specifications are certainly subject to inaccuracies, range measurements appear to be the main source of noise and other anomalies. The extent of noise in the range measurement and, eventually, the pointing direction, are currently the subject of an instrument calibration project.

Questions in need of resolution are how standard errors depend on distance, surface consistency and angle of incidence.

Since the laser signals have a certain width, they are liable to be split so that one portion of the signal is reflected from a closer distance than the remaining portion. This split signal effect, commonly referred to as

“mixed pixels”,

has been observed early on [14]. When scanning the box, signals were split by both the horizontal and vertical edges of the box. Split signals may be responsible for an observed data anomaly referred to (see [1, 10]) as

“phantom points”.

Here two partial returns of the signal appear to be combined to produce a weighted average of the two different range measurements, resulting in a single data point on the line of sight somewhere between two object locations. Figure 1a illustrates the generation of phantom points when a signal is split by a horizontal edge of the box. Such phantom points can be observed in Fig. 1b which depicts a side view of the visually aligned point cloud of the box. Phantom points generated along the vertical edges of the box can be seen in Fig. 9 below. Phantom points present a major problem because they appear, by their generation, at locations where there is no solid object, and because their automatic recognition and removal is as yet unresolved. They thus may cause false extensions of the generated TIN surfaces (Fig. 1c).

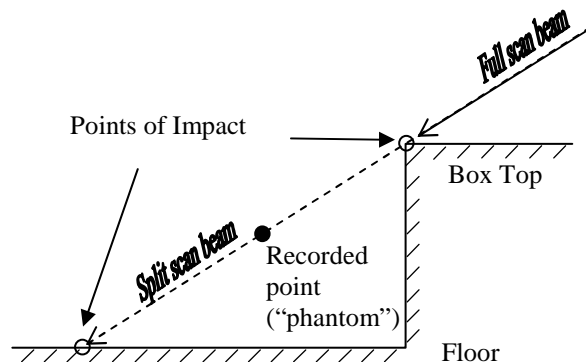


Figure 1a. Split signal at horizontal edge of box giving rise to a phantom point.

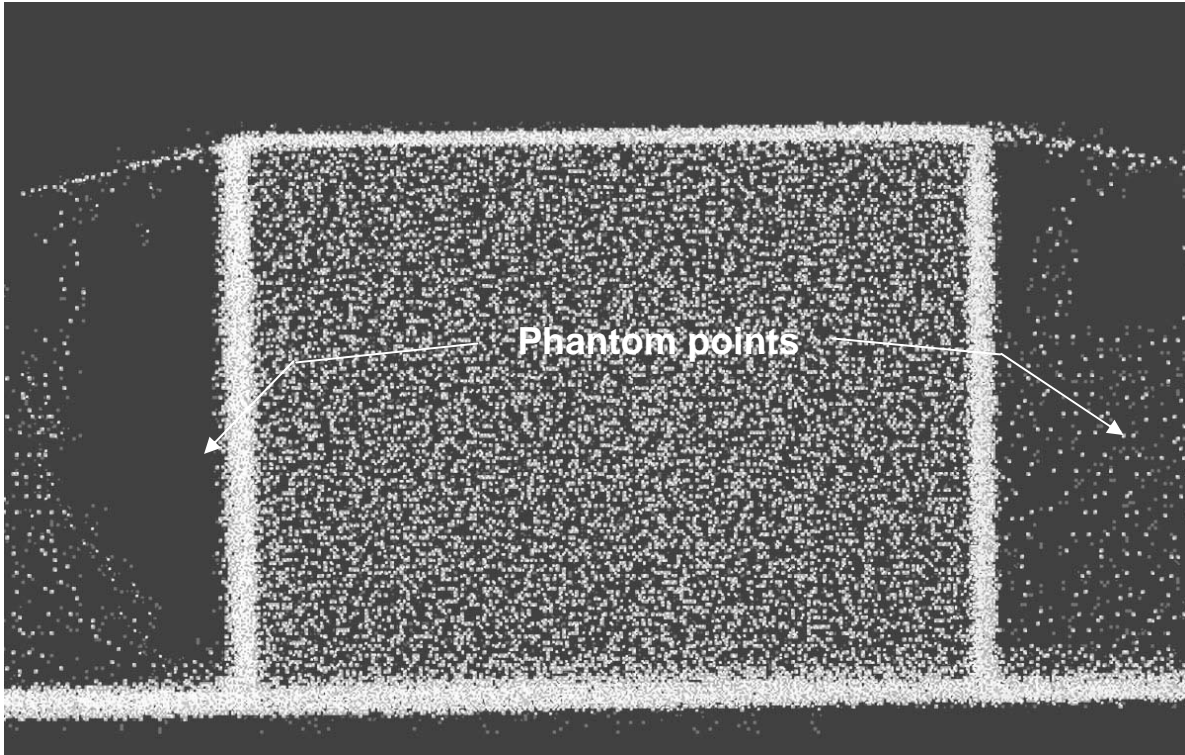


Figure 1b. Side view of box registered point cloud.




Figure 1c. Surface extensions due to phantom points (mixed pixels).

Another anomaly, caused by the vertical meshing procedure of the TIN, might be called

“vertical noise”.

Noise of measurements aimed at vertical surfaces is emphasized if residuals are defined as vertical deviations. In the case of the wooden box, a shortfall in range measurement of, say, 2 cm may show up as a deviation by 90 cm -- roughly the height of the box -- if referenced against the laboratory floor (Fig. 2). The virtual reality model language (VRML) rendering of a TIN surface portrays such points as clusters of spikes in the neighborhood of vertical faces of the box (Fig. 3).




Figure 2. Schematic of vertical noise for elevated TIN surfaces at vertical surfaces.




Figure 3. Spikes at vertical walls due vertical noise.

It is illuminating to also examine a footprint triangulation, which is a TIN of the point clouds determined by the laser scans (Fig. 4). The scanned vertical faces of the box show up as bands of high density triangles. This feature permits the assessment of the quality of registration, and may even be utilized for visual registration. Noise tends to elongate the box top by recording range overshoots, which are points that cluster just beyond the apparent end of the box top. The

footprints of what apparently are phantom points can be observed well inside the shadow area of the box.




Figure 4. Triangulation of Point Cloud from One Location; Footprint of Box Superimposed.

When scanning a floor or a horizontal surface such as the box top, range shortfalls will cause overstatement of height, just as overshoots will result in underestimation. When least squares estimation is involved, the pattern may mimic a slight slant, say, of the floor plane. It is not clear whether there is evidence for such a

“slant effect”

in the course of registration, as will be discussed later. Figure 5 illustrates a 2D slant effect for an artificial data set.




Figure 5. Schematic of a slant effect: upper points are right shifted. The slanted line is the linear regression line that minimizes RMS.

2.3 Outline of Experiments to be Reported

Three Experiments were conducted. Each experiment involved several registration applications corresponding to variations of the same basic approach.

Experiment 1 employed an exact surface model of the box, and registered each of the four point clouds against that ideal surface manually using *TINregister* with some guidance from wire models of the TIN itself, that is, the triangulation of the footprint plane. After registration, the four point clouds were combined into a single encompassing point cloud. The final model was then developed using *TINvolume*, which provided displays and volume calculations.

The purpose of this experiment was to explore the extent to which an accurate registration could be achieved under ideal circumstances, where accuracy was measured by the degree to which the correct box volume was reproduced. Such registration against a known model may also be required, if a data cloud has to be aligned with specified targets, or if a known shape is to be identified as part of a scene.

In *Experiment 2*, a partial TIN model of the visible portions of the box was created from one of the four scans. A second point cloud was then registered against that surface. A third point cloud was registered against a model of the combination of the previous point clouds, and so on, until all four point clouds were combined in reference to the coordinate system of the point cloud selected first. The final TIN model was then based on that combined data cloud.

For the purpose of comparison, *Experiment 3* was based on visual adjustments. Multicolored displays of the respective point clouds were manually moved into alignment by planar rotations and translations. Horizontal and vertical projections provided guidance. Again, the combined data cloud was processed to yield a TIN model and associated volumes.

Prior to attempting registration adjustments, direct measurements had been used to achieve an approximate matching of the four scans.

3. Elevated Surfaces and Delaunay TINs

A brief discussion of TIN surfaces is given in this section, and the reader is referred to [8] for more details.

A TIN characterizes a class of data structures on which to base the representation of surfaces, in particular, terrain surfaces. Such surfaces are almost always of the form

$$z = z(x, y),$$

in other words, each footprint point $p = (x, y)$ in a planar

“footprint area”,

has a unique elevation z associated with it. For surfaces of this kind, overhangs, arches, and similar terrain features are ruled out. Such surfaces will be referred to as

“elevated surfaces” (see Fig. 6).

The term *“parametric surfaces”* is also found in the literature (e.g.[5]), as is the term *“2.5-D”*.

A TIN, in essence, represents a

“triangulation”

of the footprint area (Fig. 6), that is, a covering of that area by triangles t_k , $k = 1, \dots, \ell$, with vertices $v_j = (\hat{x}_j, \hat{y}_j)$, $j = 1, \dots, m$. Those triangles are not permitted to “overlap”. That is, two triangles can meet only in a single vertex or an entire edge of both triangles.




Figure 6. Triangulated Footprint area and the Corresponding Elevated Surface.

Given elevations z_j at each vertex v_j , then each footprint triangle is elevated into x, y, z - space. If two footprint triangles have an edge in common, then so have their elevated counterparts. The elevated triangles thus assemble into a piecewise triangular surface, which will be referred to as the

“TIN surface”.

3.1 Delaunay Triangulations

Given data points $P_i = (x_i, y_i, z_i)$, $i = 1, \dots, n$, a frequent goal is to construct a TIN surface which in some sense -- to be defined -- well represents a given point cloud. There are, however, many ways of triangulating a set of planar points and, consequently, there are many TIN surfaces for any given specification of vertices v_j . The quality of the representation depends very much on the choice of the triangulation. In particular, the occurrence of very long edges and associated skinny triangles in the interior would clearly be undesirable, as it would distort the appearance of the surface.

It is therefore important, that for every set of planar points, there exists a triangulation, usually unique, which satisfies the following

- *Empty Circle Criterion:* No circumcircle of a triangle t_k of the triangulation contains any vertex v_j of the triangulation in its interior.

Such a triangulation is generally called a

“Delaunay Triangulation” or *“Delaunay TIN”*.

The Delaunay TIN tends to avoid long skinny triangles and has therefore been the generic method of choice for most TIN procedures. In fact, the term “TIN” is often used synonymously for “Delaunay TIN”. Both *TINvolume* and *TINregister* use Delaunay TINs for their surface modeling.




Figure 7. Delaunay Triangulation.

3.2 The Insertion Method for Delaunay TINs

For their construction of Delaunay TINs, *TINregister* and *TINvolume* proceed by successive refinements based on the paradigm known as

“insertion method”.

As in most constructions of elevated TIN surfaces, the footprints of the given data cloud have been enclosed to start within a rectangular area whose sides are, typically, parallel to the coordinate axes of the footprint plane. In that case, this rectangular area will be referred to as the

“map”,

of the TIN. The map may have been generated by a cropping procedure.

Once *TINregister* and *TINvolume* have generated a TIN in such a specified map, options for further adjustments are provided. They are discussed in Section 3.5.

For a description of the insertion mechanism, that is, the procedure of updating the TIN to include the new vertex while maintaining the Delaunay property, we refer the reader to [8].

Obviously, what will be called here

“duplicate points”,

namely, data points with the same footprint, cannot be included simultaneously into the desired TIN. A TIN which includes all points, except some that are duplicates of included points, will be called a

“complete TIN”.

It is important to realize that, for many kinds of applications, it is neither necessary nor desirable to create a complete TIN. Instead, a subset of the data points may be selected as

“critical points”,

and used to generate vertices of a TIN surface. The selection of such critical points can be made a part of the insertion process. In that case, the insertion process is often called “adaptive”.

Consider thus a data point $P_i = (x_i, y_i, z_i)$ which is not yet part of the TIN created so far. Then this data point has a

“residual” $z_i - \hat{z}_i$,

which expresses the difference between the given elevation z_i and the elevation \hat{z}_i that is predicted by the surface corresponding to the current TIN at the footprint location (x_i, y_i) . Those residuals of

remaining data points can be used to guide the selection of critical points as part of the insertion process. For each triangle in the current TIN, a

“key point”,

which maximizes the size of its residuals, is chosen from among the data points P_i whose footprints are located in the triangle.

The triangles and their key points are then ranked by a suitable criterion and maintained in a sorted list accordingly. The highest ranking triangle is located at the top of that list. Its key point is first in line to be selected as the next critical point to be inserted, unless the process is terminated.

3.3 Termination

Termination is simply by reaching a specified

“budget”

that is, a number of allowable vertices in the TIN. Other approaches would terminate, for instance, if a specified accuracy of representing the remainder of the point cloud to be TINned has been reached.

The process for generating TIN surfaces outlined above can be categorized as a

“bottom-up”

procedure. It differs from other adaptive TIN methods which employ a “top-down” strategy, in which a complete TIN, containing all possible data points as TIN vertices, is computed first, and selected vertices are deleted afterwards.

In those cases where only a fraction of the data points are selected as TIN vertices, it is possible to determine the residual-based measures-of-fit, that is, quantify the deviations of all data points from the established TIN surface. The three measures-of-fit,

MAX, RMS, ASD,

described in Section 2.1, are therefore reported. There are several options for conducting the process of generating a TIN surface which shape the triangulation, unless the goal is a complete TIN. Those options are discussed below. The accuracy of the representation will also be affected. In Section 3.6, however, additional adjustments for improving the RMS measure-of-fit will be described.

3.4 Options for TIN Generation

The ranking of the key points and their associated triangles determines the sequence in which key points are selected as critical points. *TINregister* provides two alternate options for ranking triangles. They are both keyed to the value of the residuals at their respective key points.

- “*Ranking by maximum deviation*”: rank triangles t_k according to the size, that is, absolute value of the key residuals.
- “*Ranking by volume deviation*”: rank triangles t_k according to the product of the size of the key residuals times the area of the triangle.

The first ranking criterion is most commonly used. It is geared towards an aggressive reduction of the maximum residual error. As to the second ranking criterion, note that the product between size of the residual at the key point of the triangle and its area describes -- up to a factor of 1/3 -- the volume change that an insertion of the key point into its triangle would cause, if the TIN surface were considered as bounding a volume. For this reason, the term “ranking by volume deviation” was chosen. When using that criterion, the resulting triangulation tends to be more locally homogeneous as to triangle size, because a large triangle may reach top rank, -- and be broken up -- even if its associated residual size is relatively small.

The computational effort of the selection procedure to generate a Delaunay TIN is biggest in the early stages of the procedure when many data points need to be examined in order to determine key points in a triangle. For this reason it is recommended not to start with too sparse an initial triangulation. An additional reason is that a triangulation represented by only a few data points does not provide good guidance for the selection of key points and, subsequently, critical points. Of the many ways an initial selection of data points can be achieved, the following has been implemented.

- “*Initial Binning*”: divide the reference area into a grid of almost square bins. In each bin that contains data footprints, select the data point P_i whose footprint $p_i = (x_i, y_i)$ lies closest to the center of the bin and interpolate this subset of data points to achieve an initial Delaunay TIN.

3.5 TIN Boundaries; Shadow Delineation

Once *TINregister* and *TINvolume* have generated a TIN in a specified map, further adjustments are, in general, needed. Two such adjustments -- quite unrelated to each other -- are discussed below. They address issues of what constitutes the footprint area of a data set, as well as how to improve the approximation of the point cloud by the TIN surface of the previous section.

Scanning is inevitably subject to occlusion. Thus there are shadow areas in the footprint plane. For most purposes it is important to delineate those shadow areas in the footprint plane as well as the general extent of scan coverage. Since the TIN generation discussed here takes place within a rectangular map, and since points at the corners of the map are formally specified rather than representing actual data points, triangles in the TIN that contain corner points are not

“relevant triangles”.

for the surface that needs to be created. The same goes for triangles that cover shadow areas. Usually it is necessary to identify and delete such triangles from the TIN.

For this purpose, one observes that triangles which cover shadows are typically characterized by very long edges or, alternatively, by circumcircles of large diameters. The following choices are therefore offered for designating relevant triangles or, equivalently, deleting irrelevant ones:

- *all triangles are relevant*
- *only triangles meeting map corners are deleted*
- *triangles containing map corners or with edges longer than a specified edge length are deleted*
- *triangles containing map corners or with circumcircles of larger than specified diameter are deleted.*

The boundaries of TINs delineated in this fashion, however, are typically very irregular. The following flaws are common (see Fig. 8). There may be

“boundary spikes”,

that is, consecutive boundary edges which form a very acute angle. Another problem is the occurrence of

“multiple boundary points”.

Here four or more boundary edges, instead of the usual two, are incident to a same boundary vertex. Triangles may have been deleted in the interior of a contiguous set of triangles. This represents a legitimate feature if the deleted triangles describe a shadow area. It may, however, not be desired if only a small number of adjacent triangles forming a

“hole”

is concerned. Conversely, it may not be desirable to have small clusters of triangles isolated from the bulk of TIN triangles.




Figure 8. Boundary Irregularities.

Consider two triangles in the TIN connected if there exists a chain of triangles, any two subsequent ones sharing an edge, which connects the two original triangles. Consider a subset of TIN triangles a “connected component” if every two of its triangles are connected by a triangle chain, and if the set cannot be enlarged without losing connectivity. In many cases, the TIN triangles form one large connected component, comprising the bulk of the triangles, with isolated small clusters,

islands,

as the remaining connected components.

For those occurrences, *TINregister* and *TINvolume* provide the following editorial options:

- *boundary spikes* with an angle smaller than a specified lower bound are deleted, that is, their respective triangles are marked “irrelevant”
- *multiple boundary points* are disconnected by effecting a bypass. This is also achieved by selective adding and deleting of one or more triangles
- *holes* are “filled”, if their respective numbers do not exceed a specified lower bound, by resetting the designation of the triangles in such holes to “relevant”

- an option is provided for restricting the TIN to its largest *connected component* thereby eliminating *islands*

The above operations are not independent in that each corrective action, except the last one, may cause the need for additional corrective measures. Those editing procedures are, therefore, applied repeatedly.

3.6 RMS Adjustments

The TIN surfaces constructed by the procedures outlined in this chapter interpolate the critical points selected as TIN vertices. That is, the approximation of the point cloud by the TIN surface is exact as far as that sample of data points is concerned. However, the remaining points in the point cloud will in general deviate from the surface. It appears to be in some sense “unfair” to single out a subset of data points for zero error, that is, zero residual.

A more natural approach would be to adjust the TIN surfaces in such a way that nonzero residuals are permitted to occur for all data points, including those whose footprints identify TIN vertices, so that the overall RMS error is minimized.

One way to implement this idea is to admit TIN vertices (x_i, y_i, z_i^*) which are not data points but share a footprint with the TIN vertex (x_i, y_i, z_i) , and to use the resulting adjusted TIN surface to represent the point cloud.

It is important to realize, that this does not mean that the data point (x_i, y_i, z_i) has been tampered with. This data point remains a member of the point cloud. It is just that the TIN surface is no longer forced to pass through that point, which now may incur a residual of size $|z_i^* - z_i|$, and this error contributes in egalitarian fashion to the overall RMS error.

Both *TINregister* and *TINvolume* offer the option for an

- *RMS adjustment*

which minimizes the RMS error for the given triangulation. The authors are experimenting with a corresponding ASD adjustment.

The RMS adjustment is determined iteratively and requires specification of the maximum number of acceptable iterations as well as of a termination tolerance. The adjustment process will terminate when either the maximum number of iterations is reached or when all vertex adjustments during one iteration fall below the specified tolerance.

The RMS adjustment is restricted to relevant triangles.

4. Point Cloud Against Surface Registration

The procedure *TINregister* constructs a TIN surface representing a first one of two given point clouds along the lines described in the previous section. It then repeatedly prompts for six transformation parameters with which to transform the coordinates of the second of the two given point clouds. At each instance, the three measures-of-fit, MAX, RMS, ASD, are calculated from the residuals of the points in the transformed point cloud with respect to the TIN surface. Transformed points whose footprints are external to the footprint area of the TIN surface -- as determined in the previous section -- will be ignored.

4.1 The Transformation Mechanism

The transformation consists of a translation defined by three parameters

$$\Delta x, \Delta y, \Delta z,$$

and three axis rotations,

$$\begin{aligned}\Delta\phi &= \text{rotation about the } z\text{-axis (yaw)} \\ \Delta\epsilon &= \text{rotation about the } x\text{-axis (roll)} \\ \Delta\theta &= \text{rotation about the } y\text{-axis (pitch)}.\end{aligned}$$

The sequence, in which the translation and the three axis rotations are carried out, matters. In this work, translation follows rotation according to the formula

$$\begin{bmatrix} x' \\ y' \\ z' \end{bmatrix} = \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta z \end{bmatrix} + M \begin{bmatrix} x \\ y \\ z \end{bmatrix}$$

Here M denotes a general orthogonal matrix $M^T M = I$, describing rotations. The transformation matrix M is determined as the product of the three axis rotations:

$$M = \begin{bmatrix} \cos \Delta\theta & 0 & \sin \Delta\theta \\ 0 & 1 & 0 \\ -\sin \Delta\theta & 0 & \cos \Delta\theta \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \Delta\epsilon & -\sin \Delta\epsilon \\ 0 & \sin \Delta\epsilon & \cos \Delta\epsilon \end{bmatrix} \begin{bmatrix} \cos \Delta\phi & -\sin \Delta\phi & 0 \\ \sin \Delta\phi & \cos \Delta\phi & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Again the sequence of the multiplication matters. It corresponds to executing the $\Delta\phi$ rotation first, to be followed by the $\Delta\epsilon$ rotation and the $\Delta\theta$ rotation in that order. The $\Delta\epsilon$ rotation is affected by the $\Delta\phi$ rotation since the latter may have changed the direction of the x -axis. The direction of the

y-axis may have similarly changed as a result of the previous two rotations¹. The angles $\Delta\phi, \Delta\epsilon, \Delta\theta$ are known as “Euler angles” (e.g., formula 2.70 in [11] in a different sequence).

As *TINregister* allows for repeated specifications of transformations, it permits simple manual trial-and error schemes in the search for a best fit, as described in the following section. Once an optimum has been found, the associated transformation is the desired registration transformation. The repeated transformations are not cumulative, that is, each specified transformation acts on the original coordinate system of the point cloud.

4.2 A Search Procedure

The routine *TINregister* reports for each set of six registration parameters, the ASD, RMS, MAX deviations of the transformed point cloud from the registration surface. In this work, the ASD measure was primarily used for guidance. In those cases where a small ASD increase was compensated by a substantial RMS decrease, the RMS measure was chosen. The MAX measure was ignored.

The following search procedure for minimizing ASD/RMS measures, was followed manually in most of the experiments. It is clear that, in order for the routine to be viable in practice, search methods need to be automated to a large degree and need to draw on the arsenal of advanced efficient optimization algorithms. It should be reemphasized, however, that the overriding concern is that of finding suitable minimization criteria.

1. *Start with registration parameters $\Delta x^o, \Delta y^o, \Delta z^o, \Delta\phi^o, \Delta\epsilon^o, \Delta\theta^o$*
2. *Minimize the first two parameters by exploring an eight point neighborhood, defined by the following perturbations,*

$$\begin{array}{lll} -\delta, -\delta, \dots, & -\delta, +\delta, \dots, & -\delta, 0, \dots, \\ 0, -\delta, \dots, & 0, +\delta, \dots, & \\ +\delta, -\delta, \dots, & +\delta, +\delta, \dots, & +\delta, 0, \dots, \end{array}$$

where in this work, δ is chosen mostly as 0.1 cm.

3. *Vary the planar angle parameter $\Delta\phi$ in steps of 0.05° , re-optimizing $\Delta x, \Delta y$ at each such step*
4. *Optimize the parameter Δz by a straightforward search in steps of, say, 0.1 cm while keeping the three planar parameters fixed. Alternatively, the three spatial parameters $\Delta z, \Delta\epsilon, \Delta\theta$ are optimized jointly by trial and error*

During the planar adjustment described in steps 2 and 3 above, the ASD and the RMS directions of change were rarely contradictory. During the vertical adjustment process described in step 4, the

¹ The terms yaw, roll, and pitch are more often used for rotation around fixed coordinate axes.

minimization of the ASD and the RMS errors, respectively, tended to yield vastly different results. In one case during the Experiment 1.1 described in Chapter 5, the discrepancy exceeded 19 cm, and the RMS error minimum was clearly unrealistic as far as registration was concerned.

The optimization problem encountered here is unfortunately of the kind where there are “local” minima which do not correspond to “global” minima. These local minima, moreover, tend to vary significantly from the desired solution. For the numerical adjustment procedure described here to succeed, it is therefore necessary, that the search procedure starts with a good first guess. In other words, the procedure presupposes that an initial matching, say, by direct measurement and/or visual adjustment, has already been achieved.

4.3 Editing Procedures for Point Clouds

An important preliminary task is to edit the data sets. Figure 9 illustrates the importance of such a task. Note the prevalence of vertical spikes due to phantom points and vertical noise.




Figure 9. Combinations of Raw Versions of Four Registered Point Clouds Without Cleaning.

Three editing procedures developed and implemented by the first author have been used to “clean” data point at various stages of the TIN modeling process.

- *bincull*
- *TINscreen*

- *TINfilter*.

The procedure *bincull* subdivides the map area into a grid of approximately square cells or

“bins”

of user specified size. For each bin that contains data a

“virtual data point”

is constructed, whose footprint (x, y) is the center of the bin, and whose elevation is the median of the elevations found in the bin. The data point closest to the virtual data point is then selected. All other data points are discarded. If there is a single point in the bin, then this point will be automatically selected. The selection procedure differs from the binning procedure for selecting initial triangulation points, which is described in Section 2.4. Also the purpose of both binning procedures is totally different. The purpose of *bincull* is to reduce excessive densities in parts of a point cloud rather than define an initial triangulation. It will also avoid the occurrence of duplicate points.

The procedure *TINscreen* constructs a complete TIN for the given point cloud: each vertex of the TIN is a data point, and -- with the exceptions of duplicate points -- all data points are represented as vertices. For each TIN vertex v , its immediate neighbors, that is, those TIN vertices that share an edge with v , are examined. The median of their elevations is compared with the elevation of vertex v . If the latter elevation exceeds that median by more than a specified upper tolerance, or if it falls short by more than a separately specified lower tolerance, then the data point is discarded. The purpose of *TINscreen* is to eliminate outliers.

The procedure *TINfilter* operates similarly to procedure *TINscreen*. It also constructs a complete TIN and, for each vertex v , determines the median of its adjacent elevations. Contrary to *TINscreen*, where the data point corresponding to vertex v is removed if it is found out of tolerance from the median, *TINfilter* resets every such data point to the respective median value. More precisely, the medians of the adjacent elevations are stored for each TIN vertex. After completion of that determination, the elevation of each vertex is reset at the associated stored value. Medians are thus always calculated with respect to the original values. Finally, the set of TIN vertices with their adjusted elevations is output as the filtered data set. The purpose of *TINfilter* is to eliminate spikes at the vertical surfaces while not reducing the number of data points.

4.4 Remarks about Medians

The median of a set of numbers may not be unique. In that case, -- it happens only for even families of numbers, where it is a common occurrence, -- there exists an

“upper median” \bar{m}

and a smaller

“lower median” \underline{m}

such that the interval $[\underline{m}, \overline{m}]$ is the set of all medians. Furthermore, both the upper and the lower median are members of the family of numbers for which they are determined, and they are the only members in the interval $[\underline{m}, \overline{m}]$.

Several tie-breaking strategies are commonly considered in this case. One may settle for the midpoint

$$\frac{\underline{m} + \overline{m}}{2}$$

of the median interval as the desired median. One may wish to select the mean of all numbers whenever that mean is also a median. Otherwise, choose \overline{m} , if the mean is above the upper median \overline{m} , or \underline{m} , if it is below the lower median \underline{m} . This selection rule, just like the midpoint rule, may not yield a median which is one of the numbers for which it is determined. If the latter property is desired, then one may choose from among \overline{m} and \underline{m} a median closest to the mean. That is the tiebreaker used in procedure *bincull*.

In procedures *TINscreen* and *TINfilter*, the elevation z of the vertex v under consideration is selected as the median of the adjacent elevations if that value z is already a median of the latter, that is, if z falls between the upper and the lower median. Otherwise the median closest to it is selected.

5. Experimental Results

A plywood box, painted white, with height 0.914 m (3 ft), width 1.219 m (4 ft), and length 1.524 m (5 ft), was scanned from four different

locations C, D, E, F,

approximately 4.5 m from the box (see Fig. 10). The locations were chosen so that two vertical faces of the box were clearly visible.




Figure 10. Box and Relative Scanner Locations

Assuming that the measurements are correct within 1.58 mm (1/16 in), worst-case error bounds can be estimated for the following quantities.

$$\begin{aligned}\text{base area} &= 1.858 \, m^2 \pm 0.25 \% \\ \text{volume of the box} &= 1.699 \, m^3 \pm 0.4 \, \%\end{aligned}$$

For the purpose of the experiments to be described, the four point clouds were roughly aligned using direct measurements. The TIN displayed in Fig. 11 gives an idea of that starting alignment.




Figure 11. Initial Alignment of Four Scans.

Subsequently, TIN models of the box are to be generated from the given four point clouds by registering, combining and cleaning these data. Volumes of those models are calculated and compared to the actual volume of the box. The degree to which agreement is reached is considered an indicator for registration accuracy. Similarly, dimensions such as the height of the box as derived from the TIN model are in some cases compared to the actual values. The visual quality of the surface representation of the box by the model provides yet another indicator of accuracy.

5.1 Volume calculations

As pointed out in the Section 2.1, volume calculations are carried out using routine *TINvolume*. This routine determines the “cut” volume between a horizontal plane below and a TIN surface above. *TINvolume* prompts -- repeatedly -- for the elevation of that horizontal cutting plane.

The TIN surface covers part of the floor around the box. Other parts represent the faces, but not the base of the box. For a complete model of the box and, in particular, for volume computation, the base, that is, the floor level has to be determined. Volume computations are sensitive to floor level determination. A change of the floor level by 1 cm results in a volume change of about 1 %.

The floor level may be ascertained directly by measuring the elevation of the scanning instrument above the supporting floor. In the local coordinate system of the scanner, the negative of that elevation provides the z coordinate associated with the floor level. It may, however, not always be possible to determine the correct center of the instrument and, in addition, there may be differences in floor elevations at the scanner location and at the box location.

The following procedure is therefore used for defining the floor using information provided by the point clouds themselves. The procedure is based on the fact that the rate of the dependence of the cut volumes on the specified cut elevations changes abruptly as the cut elevation passes the floor level. To capture that transitional elevation and to determine the associated floor level, three steps were followed:

1. *determine a sequence of cut volumes* at cut elevation increments of 0.2 cm.
2. *determine the first and second differences* of this sequence of volumes.
3. *select as floor level* the cut elevation associated with the second difference of maximum absolute value

An analogous procedure may be used to determine the top elevation of the box. The difference between floor level and top level so determined can then be compared against the known height (0.914 m) of the box.

Since the elevation increments are constant, the first differences track the rate of volume change, whereas the second differences indicate change in the rate of volume change. (The first and second differences when divided, respectively, by the increment or its square, approximate first and second derivatives).

5.2 Cleaning Procedure

In what follows, a cleaning process will be mentioned repeatedly. In all instances, this cleaning process will first apply the routine *bincull* with bin size 2 cm. This is to be followed by routine *TINscreen* which deletes points which are either more than 5 cm above or 5 cm below the median of adjacent elevations. The routine *TINfilter* will finally replace every elevation by the median of adjacent elevations.

5.3 Experiment 1

In this experiment, each of the four point clouds is separately registered against an exact TIN model of the box, called the Box Model henceforth. The coordinate axes of this Box Model run parallel to edges of the box. The base elevation of the Box Model in its coordinate system was determined at -1.934 m by manual measurement of the instrument height at location C. The faces of the Box Model deviate slightly from true verticality: the base extends by 1 mm on each side beyond the top, which has the specified dimensions. This slight deviation from verticality is necessary to allow representation as an elevated TIN surface.

Under these conditions, several scan-based surface models of the box were considered for varying assumptions and parameter settings. In each such case, *TINregister* was used manually to select transformations of the respective point clouds to the coordinate system of the box.

5.3.1 Experiment 1.1

In this experiment, only four of the six transformation parameters were varied during registration: angle corrections $\Delta\epsilon$ and $\Delta\theta$ were considered zero throughout the registration process. The results from *TINregister* are reported in Table 1.

Table 1. Transformation Values – No pre-cleaning of individual data sets prior to registration.

	Translations (cm)			Rotations (degrees)			Error (cm)		
	x	y	z	φ	ϵ	θ	ASD	RMS	MAX
C	6.3	-10.6	-0.3	2.5	0	0	19.3907	34.0236	94.1
D	6.9	8.5	0.6	2.45	0	0	25.2441	38.6948	95.4
E	19.7	-11.7	0.1	2.15	0	0	25.1995	38.0722	97.6
F	23.9	9	0.1	2.15	0	0	24.4846	38.3320	93.9

The point clouds, after having been registered separately to the Box Model, are then combined into a single point cloud, which then is cleaned using *bincull*, *TINscreen*, and *TINfilter* (see Section 4.3 and 5.2). Based on the resulting data, three separate TIN surfaces were created with 5 000 vertices, 8 000 vertices and 16 000 vertices, respectively.

In all cases, 50 x 50 bins were specified for the purpose of generating an initial triangulation (see Initial Binning, Section 3.4), and only triangles adjacent to map corners were deleted. *TINvolume* was used for that construction and also for determining a sequence of cut volumes for determining floor and box top levels. These sequences are displayed in Table 2. Figure 12 exhibits the TIN surface for 5 000 vertices. The evenness of both top and floor areas indicates good registration.




Figure 12. Surface Model with 5 000 Vertices – Combined Registered Scans Using Model Box. Optimized four transformation parameters, no pre-cleaning of individual data set but post-cleaning of combined data set.

Table 2. Experiment 1.1 – Floor and Box Top Determination.

	5 000 Vertices				8 000 Vertices				16 000 Vertices			
Floor (cm)	Vol. (cm ³)	1 st Diff.	2 nd Diff.		Vol. (cm ³)	1 st Diff.	2 nd Diff.		Vol. (cm ³)	1 st Diff.	2 nd Diff.	
-195	1969028	-23448			1969846	-23465			1969923	-23465		
-194.8	1945580	-23449	-1		1946381	-23466	-1		1946458	-23466	-1	
-194.6	1922131	-23448	1		1922915	-23466	0		1922992	-23467	-1	
-194.4	1898683	-23441	7		1899449	-23457	9		1899525	-23459	8	
-194.2	1875242	-23432	9		1875992	-23450	7		1876066	-23449	10	
-194	1851810	-23405	27		1852542	-23421	29		1852617	-23425	24	
-193.8	1828405	-23322	83		1829121	-23320	101		1829192	-23337	88	
-193.6	1805083	-23089	233		1805801	-23056	264		1805855	-23090	247	
-193.4	1781994	-22477	612		1782745	-22344	712		1782765	-22408	682	
-193.2	1759517	-21289	1188		1760401	-21116	1228		1760357	-21171	1237	
-193	1738228	-19494	1795		1739285	-19378	1738		1739186	-19401	1770	
-192.8	1718734	-17353	2141	← Floor	1719907	-17309	2069	← Floor	1719785	-17333	2068	← Floor
-192.6	1701381	-15255	2098		1702598	-15295	2014		1702452	-15343	1990	
-192.4	1686126	-13443	1812		1687303	-13527	1768		1687109	-13524	1819	
-192.2	1672683	-12055	1388		1673776	-12159	1368		1673585	-12092	1432	
-192	1660628				1661617				1661493			
-98	0	0			0	1			0	1		
-98.2	0	1	1		1	3	2		1	3	2	
-98.4	1	5	4		4	13	10		4	12	9	
-98.6	6	15	10		17	42	29		16	42	30	
-98.8	21	90	75		59	104	62		58	110	68	
-99	111	218	128		163	257	153		168	260	150	
-99.2	329	519	301		420	558	301		428	551	291	
-99.4	848	1015	496		978	1041	483		979	1058	507	
-99.6	1863	1676	661	← Box Top	2019	1685	644	← Box Top	2037	1704	646	← Box Top
-99.8	3539	2463	787		3704	2388	703		3741	2394	690	
-100	6002	2922	459		6092	2861	473		6135	2876	482	
-100.2	8924	3041	119		8953	3021	160		9011	3007	131	
-100.4	11965	3064	23		11974	3058	37		12018	3047	40	
-100.6	15029	3073	9		15032	3066	8		15065	3054	7	
-100.8	18102	3076	3		18098	3068	2		18119	3056	2	
-101	21178				21166				21175			
Floor (cm)	-192.7				192.7				-192.7			
Box Top (cm)	-99.7				-99.7				-99.7			
Box Ht. (m)	0.93				0.93				0.93			
Error Bx. Ht. (%)	1.71				1.71				1.71			
Volume (m ³)	1.7101				1.7112				1.7111			
Error Vol. (%)	0.65				0.72				0.71			

5.3.2 Experiment 1.2

In this version of Experiment 1, all six transformation parameters are optimized manually using *TINregister*. Table 3 shows the results of that registration process.

Table 3. Transformation Parameters – No pre-cleaning of individual data sets prior to registration.

	Translations (cm)			Rotations (degrees)			Error (cm)		
	x	y	z	ϕ	ϵ	θ	ASD	RMS	MAX
C	6.3	-10.6	-1.0	2.5	-0.4	-0.9	18.0402	32.2359	93.3
D	6.9	8.5	1.1	2.45	1.3	-1.4	23.5852	34.5002	96.1
E	19.7	-11.7	-1.7	2.15	0.4	0.7	24.3214	37.5852	96.1
F	23.9	9.0	-1.7	2.15	1.0	0.6	21.7738	35.0529	94.0

Subsequently, the same procedure as in Experiment 1.1 is used to construct a TIN surface with 5 000 vertices. Figure 13 shows that the resulting TIN surface is not at all even, and that contributing point clouds are clearly discernable. The quality of the registration has apparently suffered from the more “exact” registrations of the original point clouds. This is also borne out by an increase in volume error (Table 4).




Figure 13. Surface Model – Optimized Six Transformation Parameters, no pre-cleaning of individual data set but post-cleaning of combined data set.

Table 4. Experiment 1.2 - Floor and Box Top Determination.

Floor (cm)	Vol. (cm ³)	1 st diff.	2 nd diff.		Floor (cm)	Vol. (cm ³)	1 st diff.	2 nd diff.	
-195.0	1712345	-13823			-99.0	5	4	1	
-194.8	1698522	-13058	765		-99.2	9	5	1	
-194.6	1685464	-12241	817		-99.4	14	10	5	
-194.4	1673223	-11648	593		-99.6	24	13	3	
-194.2	1661575	-11068	580		-99.8	37	19	6	
-194.0	1650507	-10495	573		-100.0	56	25	6	
-193.8	1640012	-9685	810		-100.2	81	33	8	
-193.6	1630327	-8756	929		-100.4	114	41	8	
-193.4	1621571	-7743	1013	Floor ←	-100.6	155	51	10	
-193.2	1613828	-6687	1056		-100.8	206	62	11	
-193.0	1607141	-5716	971		-101.0	268	75	13	
-192.8	1601425	-5021	695		-101.2	343	92	17	
-192.6	1596404	-4660	361		-101.4	435	114	22	
-192.4	1591744	-4480	180		-101.6	549	145	31	
-192.2	1587264	-4289	191		-101.8	694	184	39	
-192.0	1582975	-4316	-27		-102.0	878	239	55	
-191.8	1578659	-4250	66		-102.2	1117	305	66	
-191.6	1574409	-4161	89		-102.4	1422	408	103	
-191.4	1570248	-4249	-88		-102.6	1830	584	176	
-191.2	1565999	-4074	175		-102.8	2414	911	327	Box Top
-191.0	1561925				-103.0	3325	1410	499	←
					-103.2	4735	1688	278	
					-103.4	6423	1865	177	
					-103.6	8288	1991	126	
					-103.8	10279	2098	107	
					-104.0	12377	2200	102	
					-104.2	14577	2301	101	
					-104.4	16878	2384	83	
					-104.6	19262	2472	88	
					-104.8	21734	2553	81	
					-105.0	24287			
Floor (cm)	-193.3				Box Top (cm)	-103.0			
Volume (m ³)	1.6177				Box Ht. (cm)	90.3			
Error Vol. (%)	-4.79				Error Bx. Ht. (%)	-1.20			

5.3.3 Experiment 1.3

In this experiment, the four point clouds are subjected to the same cleaning procedure that the combined registered point clouds were subjected to in Experiments 1.1 and 1.2. In this experiment, cleaning thus precedes registering. The results of registering the cleaned data sets, using only four transformation parameters as in Experiment 1.1, are recorded in Table 5.

Table 5. Transformation Parameters – Pre-cleaning individual data sets prior to registration.

	Translation (cm)			Rotation (degree)			Error		
	x	y	z	ϕ	ϵ	θ	ASD	RMS	MAX
C	6.0	-10.5	-1.1	2.4	0	0	1.7126	8.4063	92.7000
D	3.4	4.0	-1.2	2.30	0	0	2.5990	11.9420	93.0000
E	22.9	-10.0	-3.0	2.35	0	0	2.3536	11.1109	93.0000
F	22.6	9.1	-1.1	2.7	0	0	1.9470	9.4150	92.6000

Surprisingly, the registration parameters vary significantly, especially in the Z-translation, from those in Experiment 1.1, even though the only difference is the pre-cleaning of point clouds. After combining and, again, cleaning, a TIN surface with 5 000 vertices was constructed, following the same procedure as in Experiment 1.1.




Figure 14. Surface Model – Registered Scans Using Model Box, Optimized four transformation parameters, pre-cleaned individual data set and post-cleaning of combined data set.

The box volume determination is given in Table 6, and the reader notices a loss of volume accuracy. The reasons for the unexpected loss of accuracy, as compared to Experiment 1.1, remain unclear. The cleaning process resulted in a loss of most data points along the faces of the box. It is possible that this loss has in some fashion adversely affected the registration process.

Table 6. Experiment 1.3 – Floor and Box Top Determination.

Floor (cm)	Vol. (cm ³)	1 st Diff.	2 nd Diff		Floor (cm)	Vol. (cm ³)	1 st Diff.	2 nd Diff	
-195	1882884	-23536			-98	0	0	0	
-194.8	1859348	-23479	57		-98.2	0	0	0	
-194.6	1835869	-23255	224		-98.4	0	0	0	
-194.4	1812614	-22574	681		-98.6	0	0	0	
-194.2	1790040	-20855	1719		-98.8	0	0	0	
-194	1769185	-17846	3009	Floor	-99	0	0	0	
-193.8	1751339	-14393	3453	←	-99.2	0	0	0	
-193.6	1736946	-11314	3079		-99.4	0	0	0	
-193.4	1725632	-8811	2503		-99.6	0	0	0	
-193.2	1716821	-7487	1324		-99.8	0	0	0	
-193	1709334	-6658	829		-100	0	0	0	
-192.8	1702676	-5909	749		-100.2	0	0	0	
-192.6	1696767	-5501	408		-100.4	0	2	2	
-192.4	1691266	-5094	407		-100.6	2	144	142	
-192.2	1686172	-4579	515		-100.8	146	762	618	
-192	1681593	-4391	188		-101	908	1895	1133	Box Top
-191.8	1677202	-4233	158		-101.2	2803	3014	1119	←
-191.6	1672969	-4297	-64		-101.4	5817	3250	236	
-191.4	1668672	-4112	185		-101.6	9067	3294	44	
-191.2	1664560	-4181	-69		-101.8	12361	3200	-94	
-191	1660379				-102	15561	3404	204	
					-102.2	18965	3305	-99	
					-102.4	22270			
Floor (cm)	-193.8				Box Top (cm)	-101.1			
Volume (m³)	1.7513				Box Ht. (cm)	92.7			
Error Vol. (%)	3.08				Error Bx. Ht. (%)	1.42			

5.3.4 Discussion of Experiment 1

The comparison of the volumes determined in Experiment 1.1 using 5 000, 8 000, and 16 000 adaptively selected critical points, respectively, for the generation of the corresponding surface models shows a surprising insensitivity to the number of those points (see Table 2). The differences between those values appear to be solidly in the range of noise and round-off. The resulting accuracy of around 0.7 % is surprisingly good, only about double the uncertainty of the known box

volume. Note also that the floor and box top elevations in Table 2 agree in all three instances. This indicates that the process used for determining floor and box top elevations tends to be robust.

A counterintuitive effect is observed when comparing the results of Experiments 1.1 and 1.2. While leaving the three planar transformation parameters determined in Experiment 1.1 unchanged, both the ASD and the RMS errors were reduced in Experiment 1.2 by adjusting the z coordinate along with the roll and pitch angles $\Delta\epsilon$ and $\Delta\theta$, thus adjusting all six transformation parameters as opposed to just four of those in Experiment 1.1. One would expect that improving both measures-of-fit would also improve the quality of the registration. However, using the same procedure as in Experiment 1.1 to construct a TIN surface with 5 000 vertices results in a surface (see Fig. 13) that is not at all even, and whose contributing point clouds are clearly discernable. The quality of the registration has clearly suffered rather than improved by the additional adjustments. This is also borne out by the increase in volume error from 0.65 % for 5 000 vertices in Experiment 1.1 to -4.79 % in Experiment 1.2.

An explanation for this result has not been found. It is faintly possible that the slant effect described in Section 2.2 is substantial enough to affect the registration by mimicking nonexistent slants of point clouds with respect to the floor and the box top.

The ASD and RMS errors of the cleaned and registered point clouds in Experiment 1.3 appear to be drastically reduced compared to the values from the previous experiments. This is to be expected because of the removal of outliers and filtering. Part of that reduction is also attributable to the fact that fewer points are fitted. The reader should keep in mind, however, that both ASD and RMS measures are averages of errors, and thus designed to be somewhat insensitive to the number of points. Again, the reduced errors do not translate into improved registration. While the resulting surface model is smooth, the volume error of 3.08 % for 5 000 vertices exceeds the error of 0.65 % for the same number of vertices in Experiment 1.1. Note also that the transformation parameters found in Experiment 1.3 are quite different from those found in Experiment 1.1.

It was unexpected that cleaning the point clouds prior to registration would result in a less accurate volume determination. The cleaning process does remove most of the data points along vertical faces of the box. The quality of the registration was perhaps adversely affected by this loss.

5.4 Experiment 2

In this experiment, point clouds are registered against each other rather than against an exact model. For instance if scans of unknown terrain, say, a construction site, may have to be combined.

The procedure then is to construct a TIN surface from a cleaned version of point cloud C, and to register point cloud D against it using *TINregister*. The original -- uncleaned -- point cloud C and the registered, uncleaned point cloud D are then combined, cleaned, and used as basis for a TIN surface of the two combined scans. Point cloud E is then registered against this TIN surface and combined with point clouds C and D. The combination of the three point clouds is then cleaned and represented by a TIN surface, against which point cloud F is registered. The combination of

the registered but uncleaned point clouds is then cleaned, and a final TIN surface is constructed from it.

Similar to Experiment 1.1, only four of the six transformation parameters were varied during registration – angle corrections $\Delta \varepsilon$ and $\Delta \theta$ were set to zero. The registration parameters and box volumes are given in Tables 7 and 8, respectively. As expected, the error in the volume calculated using this registration method, registering the point clouds against each other, is greater than when an exact model is used.




Figure 15. Surface Model – Registered Using Scan C as Basis – Pre-cleaning of individual data set and post-cleaning of combined data set.

Table 7. Transformation Parameters – Pre-cleaning of individual data sets prior to registration.

	Translation (cm)			Rotations (degrees)			Error		
	x	y	z	φ	ε	θ	ASD	RMS	MAX
D	-4	18.5	-0.1	0.05	0	0	16.2638	26.1474	93.0750
E	17.5	0.5	-2	0.15	0	0	16.8945	28.3762	96.9536
F	14	21.5	0.15	0	0	0	18.0913	26.9568	94.1651

Table 8. Experiment 2 - Floor and Box Top Determination

Floor (cm)	Vol. (cm ³)	1 st Diff.	2 nd Diff.		Floor (cm)	Vol. (cm ³)	1 st Diff.	2 nd Diff.	
-195	2066735	-23586			-98	0	0		
-194.8	2043149	-23613	-27		-98.2	0	0	0	
-194.6	2019536	-23560	53		-98.4	0	0	0	
-194.4	1995976	-23586	-26		-98.6	0	0	0	
-194.2	1972390	-23585	1		-98.8	0	0	0	
-194	1948805	-23581	4		-99	0	0	0	
-193.8	1925224	-23544	37		-99.2	0	6	6	
-193.6	1901680	-23403	141		-99.4	6	91	85	
-193.4	1878277	-22995	408		-99.6	97	593	502	
-193.2	1855282	-21954	1041		-99.8	690	1676	1083	Box Top
-193	1833328	-19871	2083		-100	2366	2795	1119	
-192.8	1813457	-16884	2987	Floor	-100.2	5161	3229	434	
-192.6	1796573	-13623	3261	←	-100.4	8390	3312	83	
-192.4	1782950	-10745	2878		-100.6	11702	3330	18	
-192.2	1772205	-8794	1951		-100.8	15032	3336	6	
-192	1763411	-7514	1280		-101	18368	3341	5	
-191.8	1755897	-6588	926		-101.2	21709	3345	4	
-191.6	1749309	-5860	728		-101.4	25054			
-191.4	1743449	-5311	549						
-191.2	1738138	-4971	340						
-191	1733167								
Floor (cm)	-192.6				Box Top (cm)	-99.9			
Vol. (m ³)	1.7966				Box Ht. (cm)	92.7			
Error Vol. (%)	5.74%				Error Bx. Ht. (%)	1.42			

5.5 Experiment 3

Of the three experiments, Experiment 3 was the most subjective as it relies mainly on the operator's ability to visually align point clouds. In Experiment 3, the four point clouds were each assigned a different color and were plotted. Each of the point cloud was then transformed interactively to achieve the best visual alignment of the point clouds. The point cloud at location C was used as the "reference" in this case, i.e., no transformations were applied. Besides being very subjective, the inability to differentiate in the "depth" direction when viewing the combined point cloud in 2-D is a major contributor to alignment errors. A surface model of the box is shown in Fig. 16.




Figure 16. Surface Model of Visually Aligned Scans – Pre-cleaning of individual data set and post-cleaning of combined data set.

The box volumes are given in Table 9. As shown by the comparison of the error in volumes, this experiment, of the three experiments, yielded the greatest error. The error in predicting the box height is, however, on the same order as in the previous experiments.

Table 9. Experiment 3 - Floor and Box Top Determination.

Floor (cm)	Volume (cm ³)	1 st Diff.	2 nd Diff.		Floor (cm)	Volume (cm ³)	1 st Diff.	2 nd Diff.	
-2	2159947	-23677			91	44646	-3398		
-1.8	2136270	-23677	0		91.2	41248	-3397	1	
-1.6	2112593	-23677	0		91.4	37851	-3394	3	
-1.4	2088916	-23678	-1		91.6	34457	-3392	2	
-1.2	2065238	-23677	1		91.8	31065	-3391	1	
-1	2041561	-23677	0		92	27674	-3389	2	
-0.8	2017884	-23676	1		92.2	24285	-3386	3	
-0.6	1994208	-23670	6		92.4	20899	-3384	2	
-0.4	1970538	-23653	17		92.6	17515	-3381	3	
-0.2	1946885	-23607	46		92.8	14134	-3375	6	
0	1923278	-23357	250		93	10759	-3347	28	
0.2	1899921	-22605	752		93.2	7412	-3095	252	
0.4	1877316	-21147	1458		93.4	4317	-2250	845	Box Top
0.6	1856169	-19032	2115		93.6	2067	-1351	899	←
0.8	1837137	-16475	2557	← Floor	93.8	716	-562	789	
1	1820662	-13909	2566		94	154	-142	420	
1.2	1806753	-11668	2241		94.2	12	-10	132	
1.4	1795085	-9768	1900		94.3	2	-2	8	
1.6	1785317	-8187	1581		94.4	0	0	2	
1.8	1777130	-7205	982		94.6	0	0	0	
2	1769925	-6469	736		94.8	0	0	0	
2.2	1763456	-5745	724						
2.4	1757711	-5108	637						
2.6	1752603	-4750	358						
2.8	1747853	-4456	294						
3	1743397	-1743397							
Floor (cm)	0.9				Box Top (cm)	93.6			
Vol (m ³)	1.8289				Box Ht. (cm)	92.7			
Error Vol. (%)	7.65				Error in Bx. Ht. (%)	1.42			

6. Conclusions

The task of registering separate spatial data sets is as thorny as it is essential. There are many different variations of that task, each calling for a different approach. There are no magic bullets. Typically, combinations of different approaches will be necessary.

This report addressed ongoing work on a problem of registering point clouds gathered by laser scans of terrains and objects. The approach studied here uses residual-based distance reduction of point clouds from elevated TIN surfaces without the use of targets.

The effort has so far resulted in an improved understanding of the nature of point cloud data sets collected by a laser scanning device, including the role of “phantom points” and “vertical noise” (see Section 2.2). In implementing the particular registration methods, remarkable differences were found in the underlying ASD and RMS measures-of-fit. Also noted was the high degree of sensitivity to change in the angular transformation parameters -- presumably holding for registration methods in general.

Determination of the volume of a scanned box from registered scans, and comparison to its known volume was used as the main criterion of accuracy of the processes involved. While the degree of accuracy achieved was encouraging, there are counterintuitive effects which still lack satisfactory explanation (see Section 5.3.5):

Further work will

- continue to investigate some of the unexplained observations;
- systematically compare the quality of guidance provided by ASD and RMS measures-of-fit;
- automate the optimization procedure;
- examine and compare alternative methods of registration which may apply to the processing of the available scan data, such as registration based on:
 1. closest point determination rather than vertical residuals
 2. “projective” residuals, that is, residuals determined in the direction of the scan
 3. comparison of density patterns.

While the first approach is common practice, the remaining approaches have not been found in the literature. The third approach addresses the problem of registering point clouds against point clouds.

Particularly for visualization purposes, it would be more natural to mesh data points according to their proximity in an “angle-angle-distance” frame centered at the scanning instrument. The meshed surfaces, now no longer elevated surfaces in the sense of Chapter 3, would then be registered against each other and would then be integrated into the final surface.

References

- [1] M. D. Adams and P. J. Probert, “The detection and removal of phantom range data from A.M.C.W. light detection and ranging sensors,” tech. rep., Department of Engineering Science, University of Oxford, UK, 1996.
- [2] R. Benjema and F. Schmitt, “A solution for the registration of multiple 3d point sets using unit quaternions”, *Fifth European Conference on Computer Vision*, 34-50, Freiburg, Germany, 1998.
- [3] J.-A. Beraldin, L. Cournoyer, M. Rioux, F. Blais, S. F. El-Hakim, and G. Godin, “Object model creation from multiple range images: acquisition, calibration, model building and verification”, *Proc. Int. Conf. on Recent Advances in 3-D Digital Imaging and Modeling*, 326-333, Ottawa, May 12-15, 1997.
- [4] Fausto Bernardini and Holly Rushmeier, “Strategies for registering range images from unknown camera positions”, IBM T. J. Watson Research Center, Yorktown Heights, NY. 10598.
- [5] Paul J. Besl and Neil D. McKay, “A Method for Registration of 3-D Shapes”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, 239-256, Feb. 1992.
- [6] G. Blais and M. D. Levine, “Registering multiview range data to create 3D computer objects”, *IEEE Trans. Pattern Analysis and Machine Intelligence*, 17 (8):,820-824, 1995.
- [7] Y. Chen and G. Medioni, “Object modelling by registration of multiple range images”, *Image and Vision Computing*, 10 (3), 145-155, 1992.
- [8] G. S. Cheok, R. R. Lipman, C. Witzgall, J. Bernal, W. C. Stone, “NIST Construction Automation Program Report No. 4: Non-Intrusive Scanning Technology for Construction Status Determination”, *NISTIR 6457*, National Institute of Standards and Technology, Gaithersburg, MD, Jan. 2000.
- [9] G. S. Cheok, R. R. Lipman, C. Witzgall, J. Bernal, and W. C. Stone, “Field Demonstration of Laser Scanning for Excavation Measurement,” *Proceedings of the 17th IAARC/CIB/IEEE/IFAC/IFR International Symposium on Automation and Robotics in Construction*, Taipei, Taiwan, September 18-20, 2000, pp. 683-688
- [10] R. Collier, “Characterization of a Range Scanning System Utilizing a Point Laser Rangefinder,” Masters Thesis, The University of Tennessee, Knoxville, August, 1998.

- [11] J. J. Craig, *Introduction to Robotics, Mechanics and Control*, 2nd ed., Addison-Wesley, Reading, MA, 1989.
- [12] C. Dorai, J. Weng, and A. K. Jain, "Optical registration of multiple range views", *12th Int. Conference on Pattern Recognition*, A569-571, Jerusalem, Israel, 1994.
- [13] S. F. Hakim, J.-A. Beraldin, F. Blais, G. Godin and P. Boulanger, "Two 3-d sensors for environmental modeling and virtual reality: calibration and multi-view registration", *Proc. Int. Arch. of Photogrammetry and Remote Sensing*, Vol. XXXI, Part B5, 140-146, Vienna, July, 1996.
- [14] M. Herbert and E. Krotkov, "3-D Measurements from imaging laser radars: How good are they?", *IEEE/RSJ International Workshop on Intelligent Robots and Systems*, pp. 359-364, 1991
- [15] D. F. Huber and M. Hebert, "A New Approach to 3-D Terrain Mapping", *Proc. of the 1999 IEEE/RSI International Conference on Intelligent Robotics and Systems (IROS '99)*, pp. 1121-1127, October, 1999.
- [16] D. Huber, O. Carmichael, M. Hebert, "3-D Map Reconstruction from Range Data", *Proc. of the 2000 IEEE International Conference on Robotics & Automation*, San Francisco, CA, 891-897, April 2000.
- [17] A. E. Johnson, "Spin-Images: A Representation for 3-D Surface Matching", PhD Thesis, The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, August, 1997.
- [18] B. H. Kirkwood, J.-Y. Royer, Theodore C. Chang, and Richard G. Gordon, "Statistical tools for estimating and combining finite rotations and their uncertainties", *Geophys. J. Int.*, A. J. ol 137, 408-428, 1999.
- [19] X. Pennec, "Multiple registration and mean rigid shapes; Application to the 3D case", *16th Leeds Annual Statistical Workshop*, 178-185, Leeds, U.K., 1996.
- [20] A. J. Stoddart and A. Hilton, "Registration of multiple point sets", *13th Int. Conf. on Pattern Recognition*, B40-44, Vienna, Austria, 1996.