Characterising errors in airborne laser altimetry data to extract soil roughness

Article

Accepted Version

Davenport, I.J., Holden, N. and Gurney, R.J. (2004) Characterising errors in airborne laser altimetry data to extract soil roughness. IEEE Transactions on Geoscience and Remote Sensing, 42 (10). pp. 2130-2141. ISSN 0196-2892 Available at https://centaur.reading.ac.uk/4788/

Characterising errors in airborne laser altimetry data to extract soil roughness

I J Davenport, N Holden, R J Gurney

Keywords: airborne laser altimetry accuracy soil roughness

Abstract

Airborne laser altimetry has the potential to make frequent detailed observations that are important for many aspects of studying land surface processes. However, the uncertainties inherent in airborne laser altimetry data have rarely been well measured. Uncertainty is often specified as generally as 20cm in elevation, and 40cm planimetric. To better constrain these uncertainties, we present an analysis of several datasets acquired specifically to study the temporal consistency of laser altimetry data, and thus assess its operational value. The error budget has three main components, each with a time regime. For measurements acquired less than 50ms apart, elevations have a local standard deviation in height of 3.5cm, enabling the local measurement of surface roughness of the order of 5cm. Points acquired seconds apart acquire an additional random error due to Differential Geographic Positioning System (DGPS) fluctuation. Measurements made up to an hour apart show an elevation drift of 7cm over a half hour. Over months, this drift gives rise to a random elevation offset between swathes, with an average of 6.4cm. The RMS planimetric error in point location was derived as 37.4cm. We conclude by considering the consequences of these uncertainties on the principle application of laser altimetry in the UK, intertidal zone monitoring.

I. Introduction

Airborne laser altimetry provides a spatially dense set of spot heights, based on measuring the return time of a laser pulse emitted from an aerial platform towards the ground, and combining this time with a differential GPS system to measure the location of the aircraft, and an inertial navigation system to measure its orientation. For the differential GPS, a base station is located in the vicinity. Variations between systems include variability in the beam divergence, the angular diameter from the emission point beyond which the radiation density drops below 1/*e* of its central level, and different scanning patterns and densities of the laser pulses. Some systems will record the complete received returned intensity profile of each laser pulse with time, others will record only the times of the first and last returns.

Laser altimetry measurements provide a cost effective means of collecting topographic information of use to land surface process models. In the simplest terms, laser altimetry data is usually provided to the user as a set of three-dimensional spatial coordinates, with an indication for most systems that the planimetric accuracy is around 40cm, and the height accuracy around 20cm. However, the way the platform location is determined, and the acquisition pattern and sequence mean that there are different internal uncertainties between the planimetric positions and elevations of points gathered a few milliseconds apart and months apart. For example, two spatially overlapping swathes may have a systematic height difference far greater than the internal swathe height uncertainty, leading to an overlap region of artificial apparent roughness if the data are simply combined.

If we intend to make measurements of structures or changes in topography, understanding the temporal dependence in these uncertainties is crucial in qualifying our results, and getting the most out of the data. In [1] the elevation accuracy of the NASA Airborne Topographic Mapper (ATM) is estimated based on seven overlapping flights over a stretch of open beach on one day, and contemporary differential GPS surveys. The authors conclude that the random within-swathe RMS error is around 11cm, with a "mean" error representing swathe-swathe displacement between 2cm and 13cm.

In this work, we present a more temporally dense and extensive analysis using the Optech Airborne Laser Terrain Mapping (ALTM) system, which is in wide use. The principles apply to both scanning and profiling systems, and so the methods should allow a general way of describing errors of these systems to be implemented. To study the fine temporal detail of elevation accuracy we analyse 37 minutes of laser altimetry data acquired over a very flat 120m x 120m area, and compare it with 408 ground-acquired differential GPS measurements of the surface. This allows us to examine the error between individual points measured with laser altimetry, corresponding to the random error of [1], and the systematic height drift with time. The fine time scale semivariance analysis indicates that the semivariance between points gathered in short time periods is small, which makes possible a surface roughness estimation, and we demonstrate a technique for this. To study the reproducibility of measurements over time scales from hours to months, we acquired a controlled site 38 times over a 15 month period, incorporating an intensive campaign of 28 swathes over 3.5 hours. By studying the elevation of a number of points on a flat region of ground, and the elevation of a flat roof within these data, we estimate the reliability of the elevation measurement over different time scales, corresponding to the mean error estimates within one day cited in [1]. By studying returns from a vertical face, we are also able to estimate the planimetric error inherent in laser altimetry data.

II. Short term (within swathe) height error

A. Objective

We are interested in knowing how consistent height measurements made with an airborne laser altimetry system are, on different time scales. If we make a point height measurement at one time, how will measurements made milliseconds later, seconds later, and minutes later compare to this? In this section we study the relative error between individual elevation measurements. The results will be most appropriate to studying targets and elevation variability on a spatial scale comparable to the instrument sampling spacing.

B. Test site and acquired data

A helicopter was flown repeatedly over an area of a concrete apron in Baginton airport, 120m x 120m in extent, at an altitude of 370m for 37 minutes between 11:23 and 12:00 on 3rd July, 2002, whilst operating an Optech ALTM 1205 laser altimeter at an acquisition rate of 5000 pulses per second. The laser was scanned in a straight line 7 degrees either side of nadir at 15Hz, as the helicopter moved orthogonally to the scan direction to create an area coverage. For this acquisition the beam divergence of the laser was set to 0.0003 radians, which gives a footprint at 370m of about 11cm in diameter. The differential GPS ground station was within 1km of the helicopter for all data acquisition. The resultant dataset consists of just over 8 million data points, each comprising a three-dimensional location for the first and last return, and an acquisition time. Because of the return trigger in this system, the last return is only useful if it occurs at least 3.8m below the first return, which is not the case with this data, so the last returns were discarded. Spatial locations are referenced to Ordnance Survey's OSGB36 datum in metres, and time is measured in seconds. Ground truth was also acquired in the form of 514 measurements of the ground made using a pole-mounted differential GPS. 408 of these locations were within the area of aerial acquisition. The location of these points is shown over a shaded relief rendering the laser altimetry data in Figure 1.

C. Measured height drift over a short time

1) Method

To compare the LiDAR data to the ground-measured heights, the elevation of each of the ground-measured heights was compared to those LiDAR returns occurring within 20cm of it planimetrically. The mean and the root mean square (RMS) of the LiDAR-ground difference were calculated for each minute of the acquisition, and are shown in Figure 2.

2) Results

A total of 41384 comparisons were made between the laser altimetry data and the ground truth. The data for the first two minutes correspond to a time of non-optimal GPS configuration, and consequently show a larger RMS error. Even ignoring these points, there is a long-term drift component to elevation measurements, going from a discrepancy of about +5cm at the start of the acquisition to about –2cm at the end. Comparing consecutive minutes, the maximum difference is 3.3cm in the gap between minutes 16 and 17, and the mean magnitude of change between minutes is only 1.0cm. This suggests two sources of noise, a long term drift with a shape indicated by the envelope of the mean elevation offset, and a random noise of a magnitude indicated by the low RMS error where the offset is zero. The random noise component can be estimated as 3.2cm RMS by detrending the data to remove the drift. Since the overall RMS error has an average value of 4.4cm over the 37 minutes, we can estimate the contribution of the drift as 3.0cm RMS. Whilst the drift is clearly not random on this time scale, if it is, as expected, periodic, then it can be considered random when undersampled on a longer time span acquisition, and this indicates that the drift has an impact on the error budget approximately equivalent to that of the random noise. The behaviour of this drift over a longer time span will determine its overall impact on the error budget, and evaluating this will require a longer acquisition sequence. We address this issue in section III.

D. Semivariance analysis

1) Method

To examine more closely the temporal behaviour of the elevation error, we calculated the semivariance of the LiDAR data along the time axis. The semivariance of the first return height was measured along the time axis using equation 1.

$$\gamma(t) = \frac{1}{2N} \sum_{i=1}^{N(t)} \left(x_i - y_i\right)^2$$

where
$\gamma(t)$ is the semivariance for time lag t
x_i, y_i are the two measurements of height made planimetrically within 20cm of each other, but measured at different times separated by the lag, t
N is the number of height points pairs extracted from the data for time lag t

Semivariance is usually used to measure the variability of a surface in space, to quantify for instance the magnitude of terrain roughness at different spatial scales. By measuring variation in the height of a flat, level surface along the time axis we are observing the variability in the instrument measurement of elevation.

The semivariance analysis consists of finding a number of pairs within the data with a range of temporal displacements, but with negligible spatial displacement, in this case under 20cm. Since we are only concerned with changes in measured height over time, and not consistency with the ground truth, we can use as much of the LiDAR data as is appropriate to the analysis. It is only important in devising the semivariogram that the actual height of the surface being measured at each point in the pair is the same, so that any measured difference in height between the points is solely due to error in the measurement system.

The parts of the apron surface where height varied significantly over the scale of a metre were removed from the dataset for this analysis. Those areas remaining had a maximum slope of 0.013. Additionally, one vehicle was moved in the area during the acquisition, so the data in that region was also removed. Restricting spatial displacement within a pair to under 20cm then translates to a possible slope-induced 2.6mm height measurement error. The beam diameter of 11cm means that a first return could be triggered from up to 5.5cm away from the beam centre, giving rise to a possible error of 0.7mm. The planimetric location error cited for the instrument is 40cm, which will lead to an additional independent error of 5.2mm, making a total position-induced height error of 5.9mm. Concerns over the GPS data quality and acquisition in the first and last few minutes of acquisition resulted in only data between 3 and 34 minutes after the start being used for the semivariogram. After excluding all unsuitable data, 6068051 points remained for the semivariance analysis.

2) Results

The computed semivariogram, of which the 0ms – 1000ms range is shown in Figure 3, has a lowest semivariance of about 0.0006m^2 at a temporal lag of 50ms. The data density is not high enough to obtain a reliable figure below this lag. This semivariance corresponding to a lag tending towards zero is referred to as the "nugget" semivariance. A nugget of 0.0006m^2 means that with a normal data distribution, the standard deviation of heights acquired near-simultaneously will be 3.5cm. The semivariance then rises to around 0.00085m^2 at 0.5s lag, which corresponds to a standard deviation of 4.1cm, then more slowly up to 0.002m^2 at a lag of 1800s, or 30 minutes, which corresponds to a standard deviation of 6.3cm. With a normal distribution, 99.7% of measurements will occur within three standard deviations of the mean. If we assume the mean measured height to be close to that which would be measured with ground truth, then only 0.3% of laser altimetry measured points would be more than 18.9cm away from the ground truth, which corresponds to the widely-cited laser altimetry elevation measurement uncertainty estimate of 20cm.

E. Analysis

Three sources of noise are evident from the semivariogram. The nugget semivariance is an indication of the time-independent error. This is caused by the electronic system for measuring the laser return time. The linear increase in semivariance up to 0.5s, and the subsequent plateau indicates that another source of noise is manifesting itself when measurements are separated by time, and reaches a maximum when they are 0.5s apart. The location of the aerial platform is updated every 0.5s or 1.0s from the DGPS data, and the location of the platform between updates is interpolated between these estimates. Therefore, point pairs further apart than the DGPS update interval will be subject to both the timing noise and the DGPS noise, and pairs much closer together than the DGPS update interval will be subject only to the timing noise. As the time lag between measurements increases from nearly zero up to the DGPS update interval, more of the DGPS noise is evident because of the interpolation between successive DGPS height estimates, and the semivariance increases as can be seen between the lags of 0ms and 500ms. On the longer timescale the semivariogram also shows a more gradual increase in semivariance up to 30 minutes, which is caused by the drift in measured height seen in Figure 2.

Thus we can describe the uncertainty sources in this laser altimetry data by three sources of noise.

1. The internal noise inherent in every measurement, due to the timing electronics and the Inertial Navigation System (INS) corrections, in this case with a standard deviation of 3.5cm. We refer to this later as the point timing noise.
2. The random noise from individual DGPS height measurements which becomes apparent when comparing points which depend on different DGPS location calculations, in this case with a standard deviation of 2.1cm. Combined with the timing noise this makes a total standard deviation of 4.1cm.
3. The drift of the measured height over periods of a few minutes, due to drift in the DGPS or the INS system, in this case a drift of about 7cm over 37 minutes. A similar drift has been noted in [1] and [2], on a scale of 6-8cm over one hour, and is attributed to the DGPS system.

F. Conclusions

Within existing datasets, relative height measurements of structures on a scale of a few metres can be made with a greater degree of accuracy than might be expected. Measurements made within 50ms, which would typically include at least 100 points, have a standard deviation of 3.5cm. So for any pair of points, one incident on the ground and one on the roof of a structure, the elevation difference has a standard deviation of under 5cm, which could be reduced by using more points.

From a completely stable platform where DGPS and INS noise can be eliminated, for instance a fixed monitoring station, we could eliminate noise sources 2 and 3 by processing all of the laser returns according to one estimate of the platform location, which would be based on an average derived from acquired DGPS data over a long period. This would, if the drift over time is unbiased, yield a more reliable and reproducible estimate of height. Such a system could be deployed for a number of acquisitions, as long as each acquisition period was long enough to average out the drift, and the instrument mounting point could be replicated adequately on each installation. This system would then only be subject to the point timing noise in this system with a standard deviation of 3.5cm, so 99.7% of points occur within 10.5cm, doubling the precision over a dynamic platform system.

Also, if we are studying spot heights which are acquired close enough together in time that the DGPS/INS drift and noise seen in the data makes a negligible contribution, we will have a set of points which are more closely correlated, and thus may yield information on surface roughness by spot-to-spot variation. We address an application of this in Section IV.

III. Systematic height error estimation from minutes to months

A. Objective

In section II we studied the relative error between points measured milliseconds, seconds and minutes apart. This is clearly important for the study of features comparable in size to the instrument sampling spacing, such as estimating a small building or tree height, or surface roughness. For this kind of measurement the absolute accuracy is less important than the relative error. If we consider extended features such as beaches, or slowly changing topography, we can eliminate short-term random error between adjacent points by resampling the data, and in the case of these measurements we are more interested in how consistent they are on a longer timescale. The 1cm per minute drift observed in section II contributes minimal internal error to a swathe which is acquired at 3km/minute, but if we acquire beach data on two occasions, how does the drift affect these data? How consistent are the swathes, and consequently what level of change can we detect? In section II we could see how the mean measured height drifted downwards about 7cm over half an hour, but is this a linear drift or a snapshot of random noise?

B. Data acquired of Sonning Farm

To measure long term systematic drift, we carried out a series of acquisitions over two time scales, using an Optech ALTM2033 instrument mounted in a Cessna fixed-wing aircraft, with some variation in the range of instrument parameters between swathes. Over a period of 15 months between June 2001 and September 2002 we acquired laser altimetry data of an area of Sonning Farm, near Reading in the UK. This is a farm run by the University of Reading which we were able to monitor and manipulate, ensuring no unexpected changes in the target. Whilst the land surface changed over the duration of the experiment, a flat-roofed mobile laboratory was on site and static over this time.

Over the 15 months, 38 acquisitions of the area were made in total. A differential GPS ground station was set up on site, and was within 1km of all data used in this study. One acquisition was made on 25th June, 2001, one on 23rd August 2001, 28 acquisitions were made on 3rd April, 2003, four were made on 14th July, 2002, and four on 23rd September 2003. This acquisition frequency enables us to examine height drift between swathes gathered only a few hours apart, and over a year apart.

C. Measuring height drift over a few hours by studying bare soil areas

1) Method

On 3rd April 2003, we measured the variability in systematic return height from swathes in two ways. One area of soil 50m x 50m was ploughed, harrowed with a rotary cultivator, and rolled flat to within 1cm in height over a metre planimetrically. Because barrels were deployed within this area for a related experiment, ten specific sites were selected, away from barrels, and the heights of these sites as acquired by the laser altimetry measured in each swathe. The ten heights in the first swathe set the baseline, then the relative heights of these points in subsequent swathes indicate the drift between swathes. For instance, as shown in Figure 4, the points in the second swathe are measured at a mean of 9cm above where they were measured in the first swathe, with a standard deviation of 7cm, and the points in the third swathe are 13cm above the first, with a standard deviation of 4cm. Since the time period over which these points were gathered is 3.5 hours, and there is no obvious trend in the data excepting the points at the start of each acquisition, we can assume that all the acquired data are subject to a DGPS drift with a shorter period than our data acquisition window. If the heights are, therefore, unbiased, it then follows that the mean height of each soil point over 16 swathes should be close to the actual elevation, and that the RMS error for these data can be estimated by subtracting the measured heights for each soil area from the mean height over the 16 swathes. A longer period drift, or a systematic bias will invalidate this value.

2) Results

1. It is an interesting anomaly that all the other swathes have a positive offset above the first.
2. The acquired swathes are clearly divided into two collection periods, and the first one or two swathes acquired in each of the periods have a noticeably different offset to the following swathes, which might be attributable to an operational deficiency such as some element of the system “warming up”.
3. After the first two swathes in each acquisition period, the heights are fairly consistent. By comparing each swathe to every other swathe, we can build up a swathe-swathe offset distribution as in Figure 5. This shows that if we had picked any two of the swathes at random, they would on average be offset by 6.3cm, and by at most 26.0cm.
4. If we eliminated the first two swathes in each acquisition period, then the mean swathe-swathe offset would be 3.8cm, with a maximum of 11.0cm.
5. If we can assume that the LiDAR measured heights are unbiased over 3.5hours, the RMS error for individual points, as calculated from 159 measurements (16 swathes, 10 soil areas, minus one swathe/area missed) is 8.6cm. By using the average offset per swathe as shown in Figure 4, hence reducing the effect of point timing noise, the RMS error due to the drift alone can be estimated as 5.7cm. If we remove from this analysis the two swathes at the start of each acquisition period, then the RMS error due to drift becomes 3.2cm. Since this is the same as we observe in the 37 minute sequence in section II, we infer that this seems to be the maximum impact of the drift on this time scale.

D. Measuring height drift over a few hours by studying a mobile laboratory roof

1) Method

An alternative means of estimating the systematic drift in height as measured by the laser altimeter is to use the roof of the mobile laboratory. It is easily distinguishable, as it is over 2m above the surrounding ground points. With a roof area of 4m x 2m = 8m^2, fewer laser altimetry points will be returned per swathe compared to the bare soil technique. This has the disadvantage that the point timing noise will be more significant with this technique than with the soil area technique. Also, variation in the acquisition pattern, density and coverage between swathes means that not all swathes have hits from the laboratory roof. Returns from the laboratory roof are identified both by their height relative to the local returns from the soil, and their spatial location relative to ground points and other identified roof points. If we can assume that the measured height is unbiased over this acquisition period, then we can, as we did with the soil areas, assume that the mean measured height of the lab roof is close to the actual height, and calculate an RMS deviation from this.

2) Results

Figure 6 shows the estimated mean mobile laboratory roof height per swathe using this technique on the 3rd April 2002 data. The RMS variation of the individual height points about the mean, which will incorporate all noise sources, is 8.2cm.

In both Figures 4 and 6 there are similar patterns in the height trend. There is a low start which stabilizes, and we can see in both cases that the first one or two swathes in each acquisition period are notably different from those following. The swathe-swathe offset distribution using this technique is presented in Figure 5 alongside the same distribution using the soil surface technique. Including all swathes shows a mean offset of 7.8cm, with a maximum of 30.2cm, though if we discount the first two swathes in each of the acquisition periods the mean comes down to 6.7cm with a maximum offset of 24.9cm.

E. Measuring height drift over 15 months by studying a mobile laboratory roof

1) Method

The advantage of the laboratory roof technique over the soil area technique is that since the mobile laboratory was on site over the complete 15 months, the height drift can be analysed over a longer period. We therefore used the remaining laser altimetry data acquired over the 15 months to detect the measured mobile laboratory roof height.

2) Results

Using the same technique to measure the height of the mobile laboratory roof in all of the datasets gathered between June 2001 and September 2002 gives us the heights shown in Figure 7. The mean roof elevation of the roof in these datasets is 38.95m. The 3rd April estimated heights are also included on this figure, and make an interesting contrast. It is noticeable that the swathes at the beginning and end of the 15 month period vary less than the swathes gathered in one day. This is because where a single swathe, or a few swathes, are acquired in one day, the acquisition time can be optimised so that the DGPS error in the location of the platform is minimised. In the case of the 28 swathes acquired on 3rd April 2002, each swathe is subject to the prevailing DGPS quality at acquisition time. To show the effect of laser altimetry acquired with optimal GPS data, Figure 8 shows the distribution of height-height offsets for all the swathes, excepting the 3rd April swathes, and demonstrates that the mean offset is not too different from the within-day mean, at 6.4cm, however the maximum difference is notably lower at 16.6cm. Again, because of the need to hit a small target, one swathe has only one hit on the lab roof, and another has only two However, removing these swathes from the statistics has a negligible effect, increasing the mean inter-swathe separation to 7.3ccm, with the same maximum of 16.6cm.

If we assume that the average height of the laboratory roof over the 31 swathes with roof hits is close to what a survey would yield, which is a reasonable assumption if the DGPS system is unbiased, then we can also interpret the measured roof heights in terms of deviation from expectation. In this case, the swathes deviate by a mean of 8.3cm and a maximum of 20.8cm, with 90% of the swathes under 16cm from expectation.

Again making the assumption that the measured height is unbiased, the RMS deviation of the heights from the mean, which we expect to be close to the actual value, is 9.9cm for all swathes with a laboratory roof hit over the 15 month period. This increase is seemingly due to the apparent 15cm discrepancy between the mean roof elevation as measured from the 3rd April 2002 swathes, compared to that calculated from the other swathes.

F. Analysis

Height data acquired over 3rd April 2002 has a mean swathe-swathe offset of 6.3cm, and a maximum of 26.0cm, using the mean elevations of ten flat soil areas.

If we exclude the two swathes at the beginning of each acquisition period, the mean swathe-swathe offset in the 3rd April 2002 data is 3.8cm, with a maximum offset of 11.0cm.

Height data acquired over 15 months has a mean swathe-swathe offset of 6.4cm, with a maximum of 16.6cm. Further work might slightly reduce these numbers, as within-swathe noise may still be making a significant contribution, where for an extended surface this effect would be averaged out.

The relationship between these different swathe-swathe offsets is shown in Figure 9, where the proportion of swathe pairs falling below a given elevation difference is compared.

G. Conclusions

Compared to the accuracy cited in section II, here we are considering the systematic drift in measurement which occurs between swathes acquired minutes or months apart. This error is thus applicable where the subject of study is greater than the instrument sampling spacing, and the point timing noise discussed in section II can be eliminated by resampling.

1. When acquiring a small area in a time period which allows for optimisation of the GPS satellite configuration, but the acquisitions are months apart, the elevation change detection threshold is 16.6cm.
2. Where a large area of data has to be gathered in a limited time, and the GPS configuration cannot be optimised by flight timing, the elevation change detection threshold is 26.0cm.

3. The fact that the maximum inter-swathe error can be reduced significantly by eliminating the first two swathes in each acquisition sequence in the 3rd April 2002 data suggests that, if the source of this error is corrected, a large area comprised of a mosaic of swathes could be gathered with a maximum internal systematic error of 11.0cm. However, change detection of the area by comparing it to a dataset acquired months later would then be restricted by the systematic error observed between months-separated data in Conclusion 1.
4. The 9.9cm RMS error over 15 months compared to 8.2cm over one day, and the differences between the 3rd April heights and all other days as seen in Figure 7 suggests there may be a longer term elevation drift occurring.

IV. An application of reduced short term height error - soil surface roughness estimation

A. Introduction

Soil surface roughness provides an important boundary variable for sediment and nutrient transport models on the field scale [3], and it is also of value in interpreting the reflectance and emittance data provided by remote sensing instruments. A technique to estimate surface roughness at the field or catchment scale without time-consuming and labour-intensive fieldwork would be extremely valuable for investigations of water and nutrient surface flow, and even for land-atmosphere interactions [4], [5].

As we have shown in section II, measurements made in rapid succession have a smaller relative error than those acquired minutes or hours apart. Measurements made closer together than 50ms, for example, will have an instrument-based standard deviation of 3.5cm, and even at the lowest acquisition rate this represents 100 pulses. Determining the magnitude of a source of variability above this instrument noise level may allow surface roughness to be distinguished. To test this, we used the 3rd April 2002 LiDAR acquisitions of Sonning Farm.

B. Description of soil surfaces

A field at the Reading University Sonning Farm, which has a sandy loam soil and was bare of vegetation at the time of the experiment was used. The soil in the test area was divided into four test sites, each approximately 50m square, with areas between 1950m^2 and 2650m^2. Each site received a different cultivation treatment. Site 1 was ploughed, harrowed with a rotary cultivator, then rolled flat. The only local variation in height in this region was due to tractor tracks about 1cm deep. Site 2 was ploughed, then harrowed with a rotary cultivator. The peak-trough height of this surface over a 1m scale was about 3cm. Site 3 was only ploughed, and comprised a surface with a peak-trough height range of about 5cm over 1m. Site 4 was ploughed, then ridged with a potato-ridger, resulting in a regular sawtooth profile in the soil, with a period of 75cm and elevation between trough and peak of about 20cm. Representative profiles of sites 2-4 are shown in Figure 10.

C. Aerial acquisitions

26 of the laser altimetry swathes acquired on 3rd April, 2002 between 11:58 and 15:23 covered sites 1-4, and were used. The swathes were acquired a few minutes apart except for a hiatus of 90 minutes between swathes 12 and 13 when the aircraft was on the ground.

The beam divergence was set at two values for these acquisitions – 0.3mrad and 1.0mrad. The altitude of the acquiring aircraft was in the range 640m to 740m, giving rise to laser pulse footprint diameters on the ground of between 67cm and 69cm for the 1.0mrad divergence, and between 19 and 22cm for the 0.3mrad divergence. The laser pulse rate is the number of spot heights acquired per second, and was set at 12kHz and 33kHz on different swathes. The maximum deflection angle perpendicular to the flight line which the laser scans was set at 9, 13, 18 and 20 degrees, giving rise to elliptical footprints with a maximum eccentricity of 1.06, and swathes between 200m and 500m wide. The different instrument settings used created datasets having between 0.31 and 1.92 points per m^2. This range of settings was selected to represent systems currently in use, and test the applicability of the technique to the existing archive of data, and data gathered for other purposes.

D. Analysis

In height measurements of a featureless, flat, level surface, the standard deviation of returns would be a measurement of the height variation due to the instrument error. The standard deviation of returns over a level, rough surface will have an additional contribution due to each beam being returned from a different part of the surface. To account for surfaces with a significant slope on the scale of a few metres we refine this approach by removing the trend in the height. The points acquired up to 5m before and after each point are used to detrend points individually. The standard deviation of this detrended height over a rough area will be due to a combination of the instrument measurement error and surface variability. A variation on this technique was used to differentiate the heights of different crops using laser altimetry [6].

This processing was applied to each of 26 swathes. The swathes were then divided into three groups. The six swathes acquired with the wide 1.0mrad beam divergence comprise one group. The swathes acquired with the narrow 0.3mrad beam divergence have been divided into those acquired with a high pulse rate, and consequently a point density between 0.94 and 1.92 points per m^2, and those acquired at the lower pulse rate giving rise to a point density of between 0.31 and 0.35 points per m^2. The mean and extreme values of detrended standard deviation of sites 1-4 within each of these groups is illustrated in Figure 11.

To judge the ability of this technique to discriminate consistently between the various surfaces, a simple binary classification was implemented using the 16 swathes with a high point density and 0.3mrad beam divergence. Each swathe was used to calculate a set of thresholds midway between surface types, and the thresholds used to classify the sites in the other swathes, yielding 240 test target pairs for each discrimination test. This method distinguished between rolled and harrowed surfaces correctly for 56.9% of the sites. The rolled sites were distinguished from the ploughed correctly in 92.5% of the sites, and the potato-ridged sites were distinguished correctly from all other types in 100.0% of cases. Training using the swathes on the other side of the hiatus from the target sites yielded an 87.9% distinction between rolled and ploughed, while training using swathes on the same side as the target data gives 97.1% distinction.

E. Conclusions

By studying the standard deviation of detrended height returns of LiDAR acquisitions made with a narrow (0.3mrad diameter) beam, we can distinguish between surfaces with peak-trough height variation of the order 1cm, 5cm and 20cm, with only 7.7% confusion between the 5cm and 1cm surfaces, in areas as small as 50m x 50m. The fact that training with data from the same side as the 90 minute hiatus as the target decreases this confusion to 2.9% suggests that reacquiring calibration data every few hours may improve the discrimination accuracy.

While the results acquired with the wide (1.0mrad diameter) laser footprint show an upward trend with increasing soil roughness, the scatter of results from the different treatments overlap to an extent that distinguishing between soil roughnesses solely on the basis of the laser altimetry analysis would be unreliable. The spot size on the ground, between 67 and 69cm, is approximately the same size as the coarsest scale of horizontal changes in the potato ridging (75cm). This result indicates that care is needed in choosing the instrument parameters depending on what is being observed.

This technique has been designed such that first order trend, or slope, will have no effect on its validity, however a second order trend or undulation on the spatial scale under 10m may disturb it. If the sensor spatial acquisition frequency were high enough, this potential shortcoming could be circumvented by either performing a second order detrending, or reducing the first order detrending window width such that the points within it show no significant second order trend.

V. Planimetric error estimation

A. Objective

In section III we looked at the error sources inherent in the elevation measured using laser altimetry. By using flat targets in the case of the soil surface method (Section III.C) and a well-defined target in the case of the lab roof method (Section III.D) we were able to ignore errors in the spatial location of spot heights. However, such errors will clearly have an effect on measurement. Planimetric errors in point location will be a result of error in the DGPS-derived location of the instrument platform, the INS-derived platform orientation and the horizontal component of the laser return timing. If we were attempting to measure the change in a sand bank using two swathes, and there was no physical change or elevation measurement error, but the two swathes had a spatial offset, then a false change would be indicated.

We therefore need to know…

1. How consistent is planimetric location is within each swathe, as this will affect the apparent dimensions of structures measured within a swathe? What is the relative error between points within swathe?
2. How consistent is planimetric location between swathes acquired nearly concurrently? If two swathes are measured in succession, how well will the location of an object in one swathe correlate to its location in the other? This is important in measuring extended objects larger than a swathe, by combining swathes.
3. How consistent is planimetric location between swathes acquired months apart? This is important for change detection, and if the DGPS system is unbiased, will give an estimation of the absolute accuracy of laser altimetry.

B. Method

Examining the data acquired of the mobile laboratory used in Section II.D and II.E enables us to estimate the elevation of the roof and the surrounding ground. Since the laser altimeter scans not only straight down, but also from side to side, when the aircraft passes to the south and east of the laboratory, the laser will occasionally return pulses from the southeast vertical face of the laboratory. Returns in this area which have an elevation more than 50cm above the ground, and more than 50cm lower than the roof can be reasonably assumed to have been incident on this face. No ground survey was carried out of the location of the face, however assuming the DGPS and the laser altimetry data are spatially unbiased, the line describing the face can be estimated by fitting a line through the laser return points. The distance between each laser return point and this line then represents the 1D planimetric error.

To determine the within swathe consistency, we need to estimate the apparent location of the face within the swathe, and find the displacement of points from this line. The bearing of the face is the same as estimated from all of the laser altimetry data, constraining the fit to the face points.

C. Results

The estimated location of the face using this method, and the positions of all the points in the 38 2001-2002 swathes incident on the face are shown in Figure 12. Seven of the 28 swathes acquired on 3rd April 2002 had returns from the face, none of the 2001 data contained face points, and individual swathes within the 14th July 2002 and 23rd September 2002 had face returns. While there are significantly more face hits from swathe 23 on 3rd April 2002 than any other swathe, eliminating these points from the face line fit would displace the estimated face location by less than 10cm. The mean displacement of points within each swathe from the face has an average of 36.0cm, and a maximum of 69.7cm. The RMS displacement from the face, averaging over swathes, was 37.4cm.

Figure 13 shows the distribution of 1D planimetric offset of all the gathered face points, indicating for example that 78% of the points were within 40cm of the face.

To analyse the within-swathe consistency, the points within two of the 3rd April swathes, swathes 23 and 28, containing 20 and 4 face points respectively, were analysed. In each case, an estimate of the face location within the data was made by fitting a line of the same bearing to the swathe data. The distance between each point and the face line is then calculated, and the results added to the cumulative histogram in Figure 13. The mean within-swathe displacements were 11.8cm and 9.7cm for swathes 23 and 28 respectively, with maxima of 40.1cm and 14.6cm.

D. Analysis

It is notable that the July and September 2002 points are closer to the face than some of the 3rd April swathes, with displacements of 31cm and 13cm respectively. Compared to the mean displacement over all swathes of 36cm, this suggests that the effect of a gap of months between swathes does not have a substantially greater impact on planimetric accuracy than a gap of a few minutes.

Whilst our face points are more consistent within swathe than between swathes, it is likely that this is simply because the platform is stable over the 0.1s time it takes to acquire them. The aircraft will only have had translational motion interpolated in this time, and this will describe the motion accurately, unless a high amplitude vibration is present, or a turning manoeuvre is underway. It seems likely that over the longer timescale that it takes to acquire a complete swathe, the errors due to DGPS drift and the INS measurement of variation in pitch, roll and yaw will lead to a spatial error within swathe closer to the error between data acquired in different swathes. However, since we measured only one localised object per swathe, we cannot be sure this is the case.

E. Conclusions

Using any swathe, on average a point measured using the laser altimetry system is displaced 36cm from the mean (which we assume to be the "truth") in each dimension, in 78% of cases is 40cm or closer, and in our tests, is never more than 70cm out. The consequences of this on elevation measurement accuracy will depend on the topography under study. Assuming the error in each axis is independent, this puts points on average 51cm from their specified location. While points planimetrically close have a higher relative accuracy, this would probably be only of use in measuring small buildings.

VI. Conclusions - interpreting the error estimates

Combining the height and planimetric errors derived above, we can make some estimates as to the usefulness of the laser altimetry system for some applications.

A. Planimetric uncertainty

We know that an individual laser point is returned from a point on average planimetrically 51cm from its given location, 36cm in each axis. Data acquired up to 15 months apart fits into this estimation as well as data gathered minutes apart, suggesting no significant time element to this measurement. This is an indication of the combined planimetric error contributed by the DGPS system in locating the platform, the inertial navigation system in determining the orientation of the platform, and the instrument measurement of the angle of deflection of the mirror used to scan the laser across-swathe. The distance to the target as measured by the laser travel time also makes a contribution to this error, but since the beam is at most 20 degrees off-nadir, a maximum of one third of the timing error contributes to planimetric error. Although this location error shows some local self-correlation, this is unlikely to extend further than a few metres, so for most purposes this error must be considered randomly applied to each point.

B. Height uncertainty

1) Systematic swathe-swathe error

In comparing heights between swathes, if the area is being measured in two swathes, the difference in height due to instrument error is on average 6.3-6.4cm, seemingly due to the DGPS-induced drift. The maximum possible systematic difference between swathes depends on GPS optimisation, but the worst case in our data was 26.0cm. If the cause of the anomalous swathes at the start of the acquisition periods can be established, or eliminated by ignoring early data, then an improved correlation could be established between data acquired as part of a contiguous area in a short time period. In this case the mean swathe-swathe error could be reduced to 3.8cm, with a maximum in our data of 11.0cm. This would be of use in acquiring one-off data, but for change detection applications, in the case of swathes acquired months apart the mean swathe-swathe displacement in our data was 6.4cm with a maximum of 16.6cm.

2) Random point timing error

In addition to the systematic swathe-swathe error, there is a random error contribution to each measured point, due to the measurement of the laser return time and INS/DGPS system random noise. This imparts a normal distribution noise with a standard deviation of 4.1cm, or 3.5cm for points closer than 50ms. This could be eliminated on large scale data by collecting points into say 3m x 3m cells, and using the mean of the 10 or so points that would fall within each cell using typical acquisition parameters. For fine scale relative measurements, where the inter-swathe systematic error is unimportant, the relative heights of these individual points can be useful in determining land surface properties.

C. Volume change

A potential application for laser altimetry in the UK is the detection of change in the coastal zone, for example measurement of volume changes in the beach to detect erosion or deposition. Typical beaches in the UK have slopes between 1:500 and 1:30.

1) Example: a 1:30 slope beach

The tidal range of the UK is typically 10m, which means that an inter-tidal zone with 1:30 slope will be about 300m wide. This could be acquired with current systems as a single swathe, allowing each acquisition time to be optimised, giving a maximum swathe-swathe systematic elevation displacement of 16.6cm. Since the beach only slopes in one direction, a mean 36cm planimetric error for points in each axis will impart an additional 1.2cm of height error to each swathe. Two random errors of 1.2cm will contribute a total error of 1.7cm, which added to the 16.6cm maximum elevation error, yields a change detection threshold of 18.3cm. This means that if we did a simple subtraction of one swathe from another, a mean difference of more than 18.3cm would definitively indicate a change. We can also see from Figure 9, that in data acquired optimally, and months apart, the systematic error is 10cm or below in 67% of pairs, so if the mean difference exceeds this added to the planimetric-derived error of 1.7cm, a total of 11.7cm, that there would be a 67% probability of change having occurred.

In the case of a single swathe-width beach, it is also possible that the inter-swathe displacement could be removed by identifying an unchanging feature on the land with a flat top about 10 sq.m. in area, which could be acquired in the same swathe as the beach. This would leave the residual DGPS drift that occurs within-swathe between the reference area and the target area, which from Figure 2 and section II.C.2 seems to be an elevation random walk of order 1cm per minute. With an aerial speed of about 50m/s, this would give a maximum elevation error of 1cm for a 3km swathe with a reference area at one end. The point timing error could then be removed by resampling as mentioned above.

2) Example: a 1:500 beach

With a tidal range of 10m, 1:500 slope beaches in the UK can be up to 5km wide. This requires multiple swathes for each acquisition, making an optimised acquisition less likely. If we assume the swathe-swathe errors in our one-day data exist within the acquisition sequence, then the maximum swathe-swathe difference will be 26.6cm. Even a maximum planimetric error of 70cm would translate to an insignificant height error of 1.4mm, leaving the 26.6cm detection threshold. However, if the problem seemingly related to the first couple of swathes per acquisition were resolved, then the swathe-swathe within-day offset distribution would be lower than the month-month distribution, and the detection threshold would be 16.6cm. So if the point timing error is removed by resampling, and the difference between any pair of resultant points within the data spatially co-located but temporally separated is greater than 16.6cm, then a change is indicated.

If we are interested in detecting an overall change rather than the specific location of change, then a different detection criterion could be used. For a 5km wide beach, more than ten overlapping swathes would be needed for complete coverage on each occasion. From the within-day swathe-swathe offset plot in Figure 7, each swathe would have a different elevation offset due to DGPS drift. If this is random on the acquisition time scale, and we can assume that there is no long term drift between "before" and "after" acquisitions, then the mean difference between the "before" and "after" datasets over all the swathes should sum to zero. A significant deviation from zero in this difference would indicate a change, though a more detailed statistical analysis in each case would be required to determine the change detection threshold.

Acknowledgments

This work was funded under NERC grant F60/G6/12/03. We are very grateful to the Environment Agency National Centre for Environmental Data and Surveillance for acquiring the laser altimetry and ground truth, to Paul Conrad of Optech for useful discussions on the LiDAR system, and to Steve Jones at the University of Reading Sonning Farm for carrying out the soil surface treatments.

References

[1] Sallenger, AH, Krabill, WB, Swift, RN, Brock, J, List, J, Hansen, M, Holman, RA, Manizade, S, Sontag, J, Meredith, A, Morgan, K, Yunkel, JK, Frederick, EB, Stockdon, H, "Evaluation of airborne topographic lidar for quantifying beach changes", Journal of Coastal Research vol. 19, no. 1, pp.125-133, 2003

[2] Stockdon, HF, Sallenger, AH, List, JH, Holman, RA, "Estimation of shoreline position and change using airborne topographic lidar data", Journal of Coastal Research vol. 18, no.3, pp. 502-513, 2002

[3] I J Davenport, M Silgram, J S Robinson, A Lamb, J J Settle, A Willig, "The Use of Earth Observation Techniques to Improve Catchment-scale Pollution Predictions", Physics and Chemistry of the Earth, 28 (vols. 33-36) 1365-1376.

[4] J L Monteith, M H Unsworth, *Principles of Environmental Physics, 2nd edition*. Butterworths, London, 1990.

[5] R L Bras, *Hydrology: An Introduction to Hydrologic Science*, Addison Wesley, 1990.

[6] I J Davenport, R B Bradbury, G Q A Anderson, G R F Hayman, J R Krebs, D C Mason, J D Wilson, N J Veck, "Improving Bird Population Models using Airborne Remote Sensing", *Int J Rem. Sensing*, 21 (vols. 13&14) pp. 2705-2717, 2000

Figure Captions

1. Acquisition site for LiDAR data (shown as relief shaded height data) with ground truth locations overlain.

2. Mean recorded first return relative to ground truth, and the RMS error between the LiDAR measurements and the ground truth, calculated and presented per minute for clarity, over 37 minutes.

3. Semivariance of height with time of LiDAR measured height, temporal lag below one second.

4. Drift of elevation measurement over 3.5 hours estimated by measuring the mean height of ten areas of flat soil.

5. Distribution of swathe-swathe systematic elevation difference based on the lab roof and soil areas in one day.

6. Drift of elevation measurement over 3.5 hours estimated by measuring the elevation of the roof of a mobile laboratory.

7. Laser altimetry measurements of mobile laboratory roof elevation, over 15 months.

8. Swathe-swathe offset distribution for data acquired up to 15 months apart, using mobile laboratory roof, ignoring 3rd April 2002 acquisitions.

9. Cumulative inter-swathe height error for different techniques and time periods.

10. Profiles of the three non-flat surfaces. (a) Harrowed with a rotary cultivator, (b) ploughed, (c) ploughed with a potato-ridger.

11. Mean and extremes of detrended height standard deviation over four treatment type sites (1=rolled, 2=harrowed, 3=ploughed, 4=potato-ridged) for the three acquisition groups. (a) 16 swathes with the narrow beam footprint and high acquisition frequency, (b) 4 swathes with the narrow beam footprint and low acquisition frequency, (c) 6 swathes with the wide beam footprint and low acquisition frequency.

12. Individual returns from the south-east face of the mobile laboratory, and the face location estimated from the returns.

13. Measured displacement from face for points identified by their elevation as being returned from the face

Biographies

Ian Davenport received the Ph.D. degree in Physics with Astronomy from the University of London in 1993. His research at the NERC Unit for Thematic Information studies and the Environmental Systems Science Centre has concentrated on the applications of novel sensors in environmental science.

Nick Holden is a Science Manager working in the Environment Agency's Technology Group based at Bath. He gained a degree in Applied Chemistry at Portsmouth and a Masters at Chelsea. He has 29 years experience in the environmental monitoring field, 10 of which have been devoted to remote sensing, and the need to address the problems of spatial and temporal robustness inherent in single point sampling has been the main driver.

Robert Gurney received the Ph.D. degree from the University of Bristol in 1975. He has worked at the Institute of Hydrology in the UK for five years, at the Hydrological Sciences Branch of NASA GSFC for eight years, two as Head, one year at the Dept. of Civil Engineering, University of Maryland, five years as Director of the NERC Unit for Thematic Information Systems, and from 1995 to the present as Director of the NERC Environmental Systems Science Centre. His main research interest is global environment change.

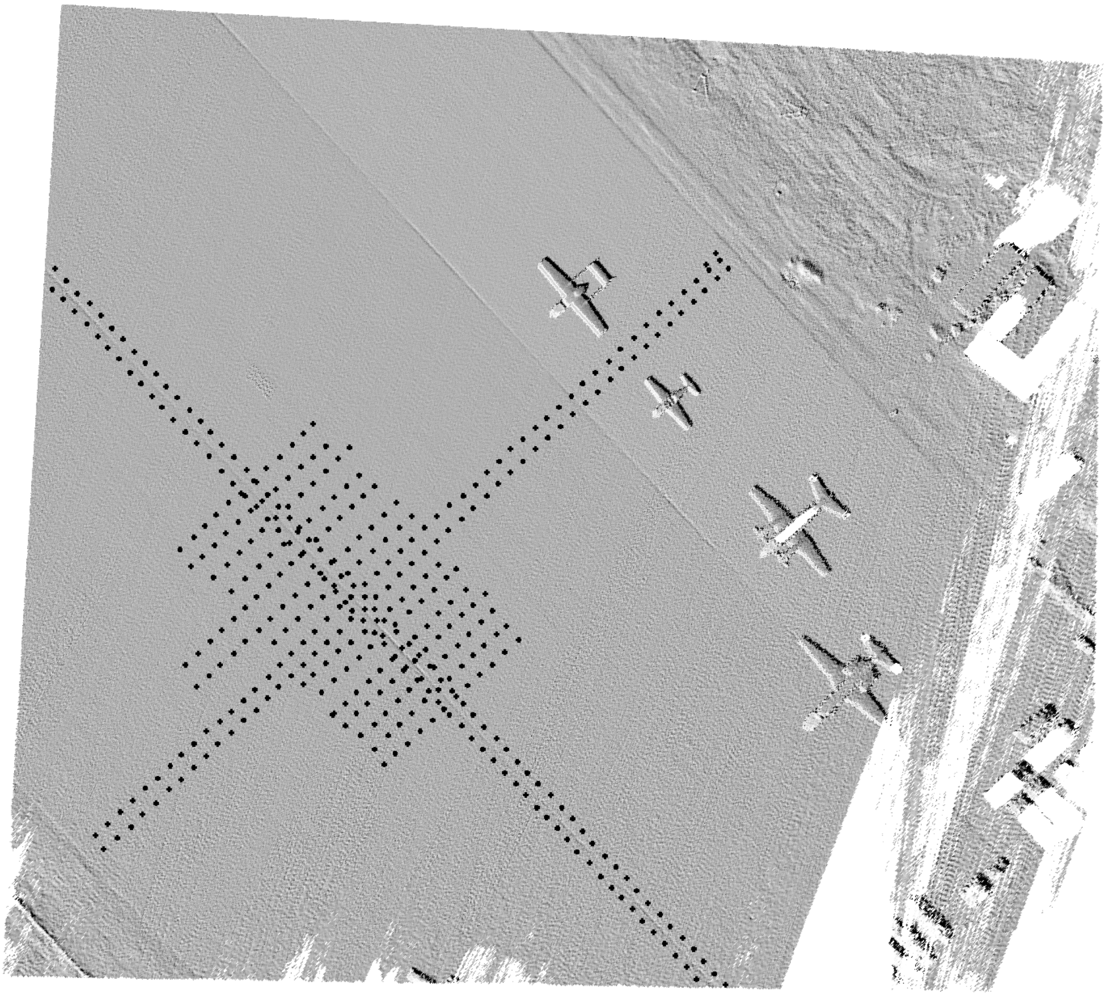

Figure 1

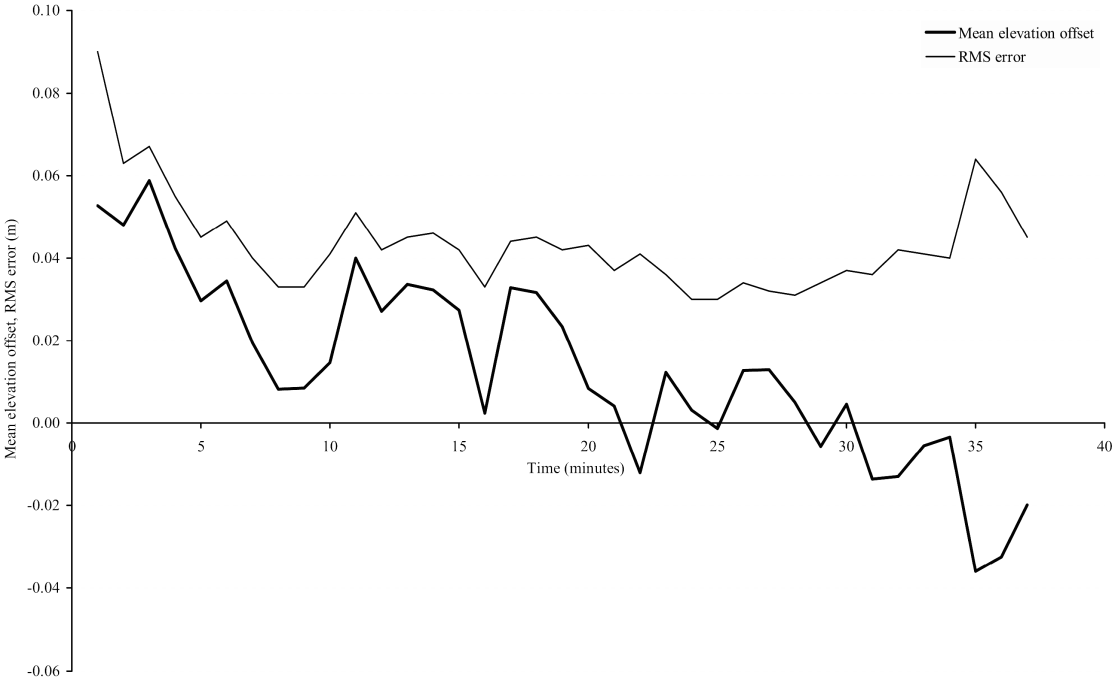

Figure 2

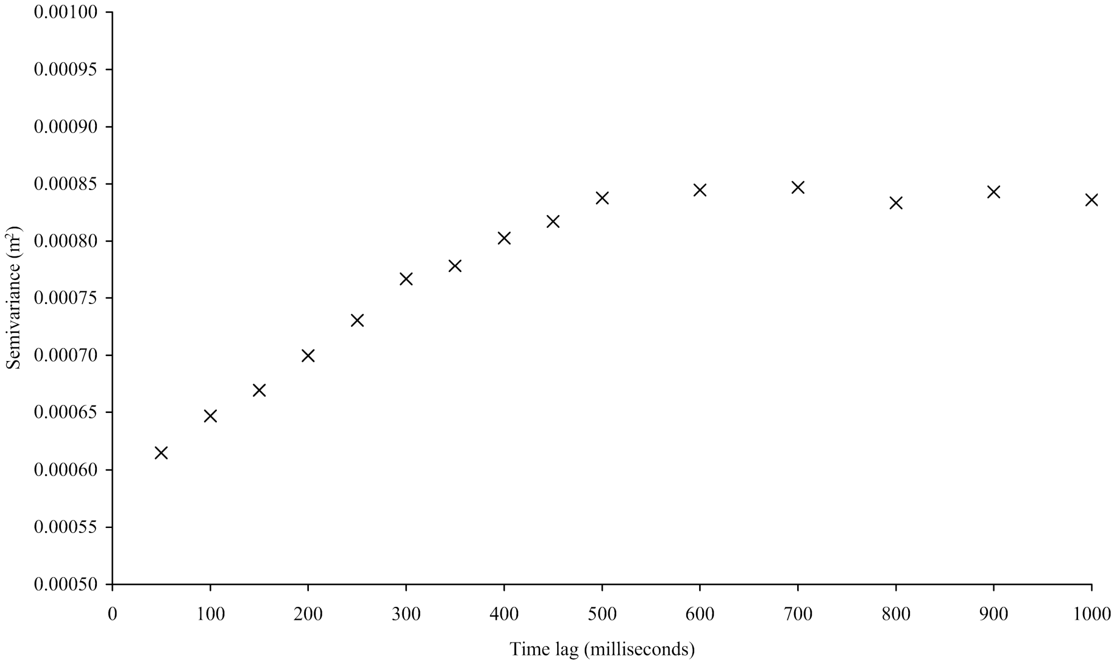

Figure 3

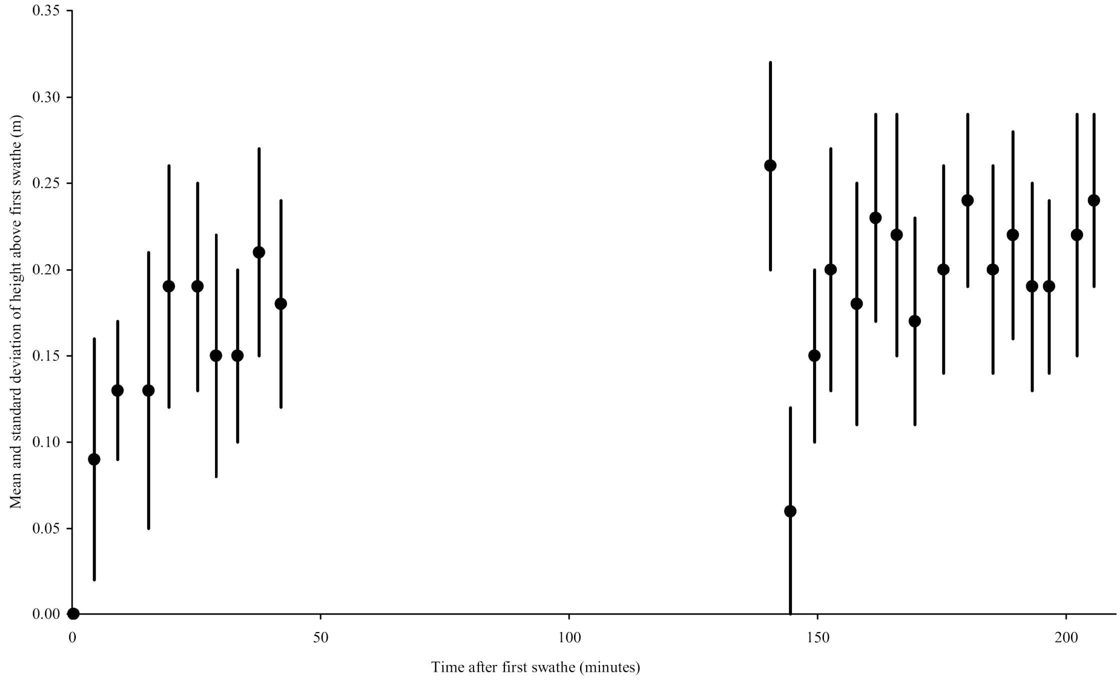

Figure 4

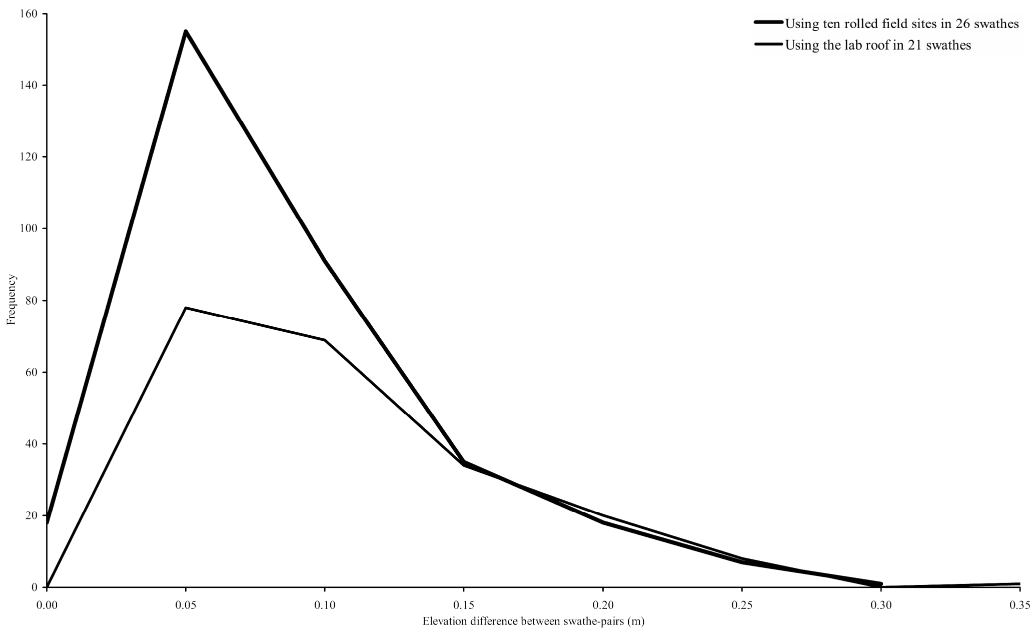

Figure 5

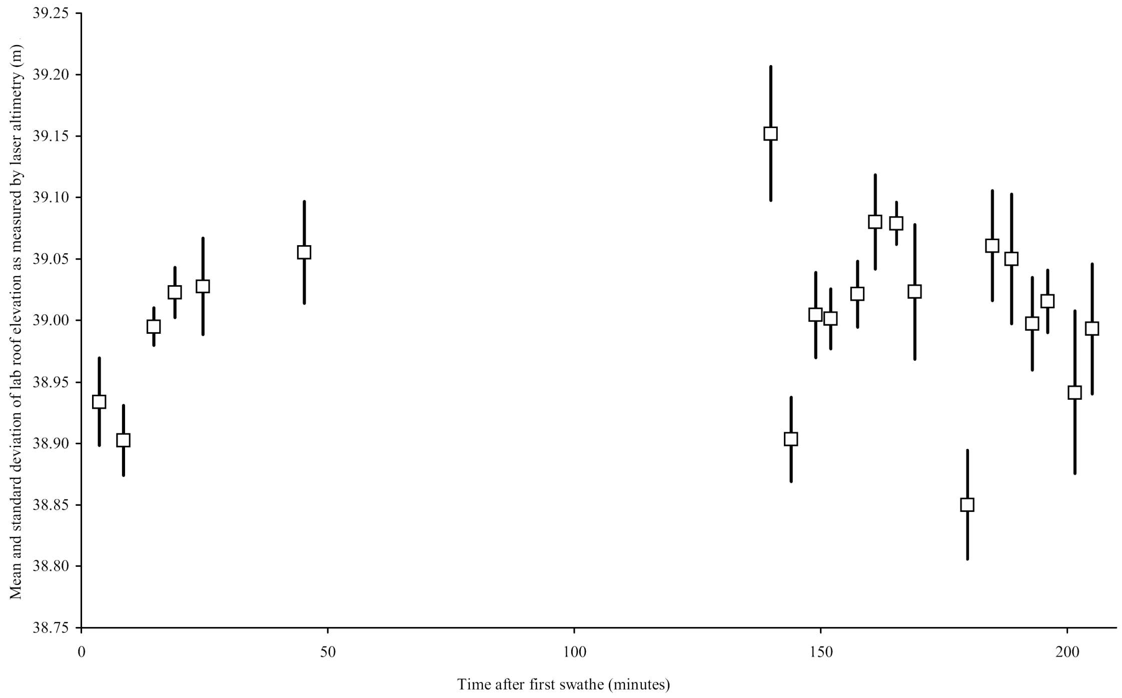

Figure 6

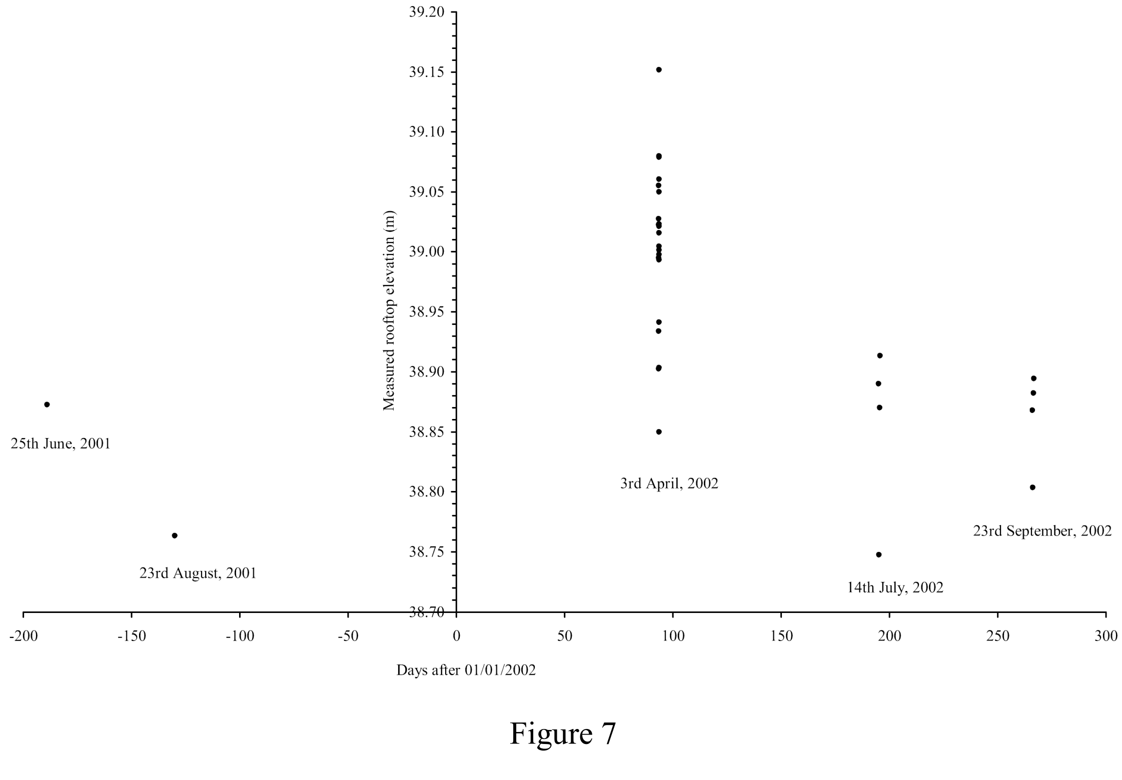

Figure 7

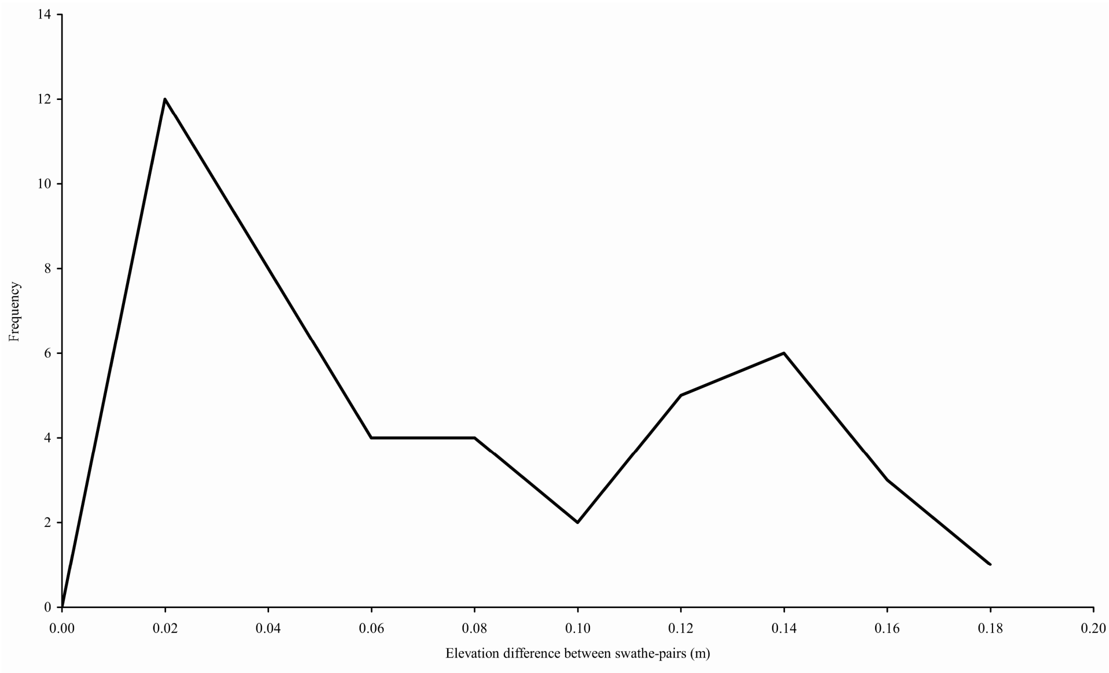

Figure 8

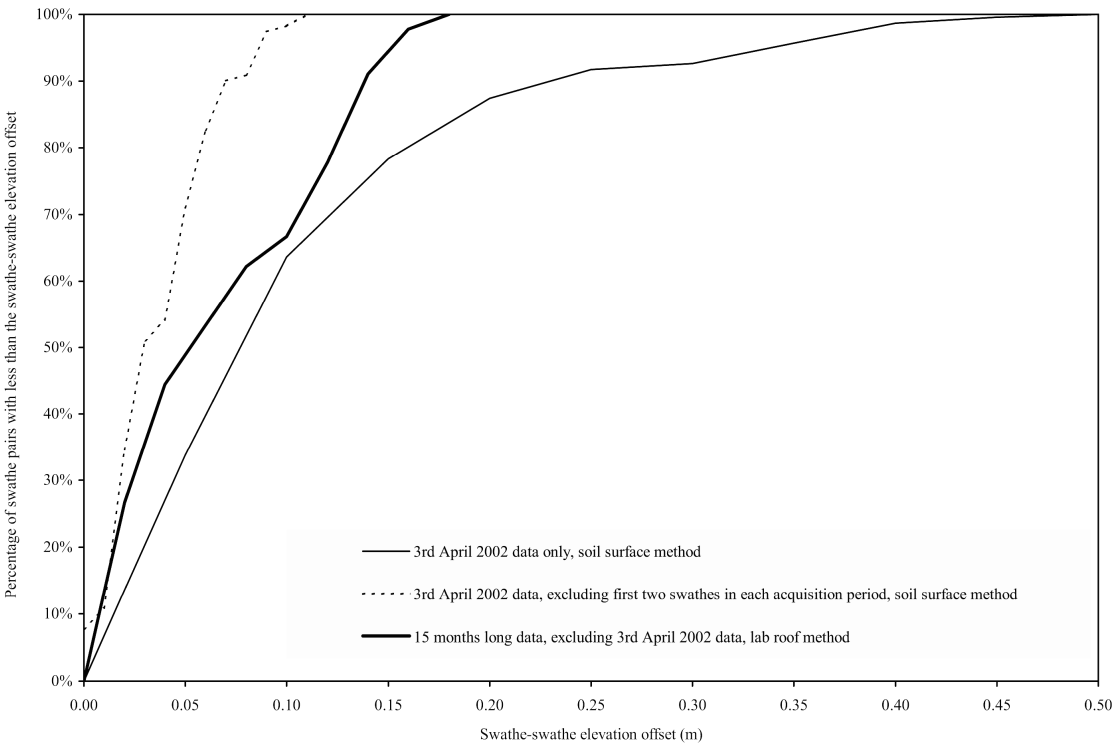

Figure 9

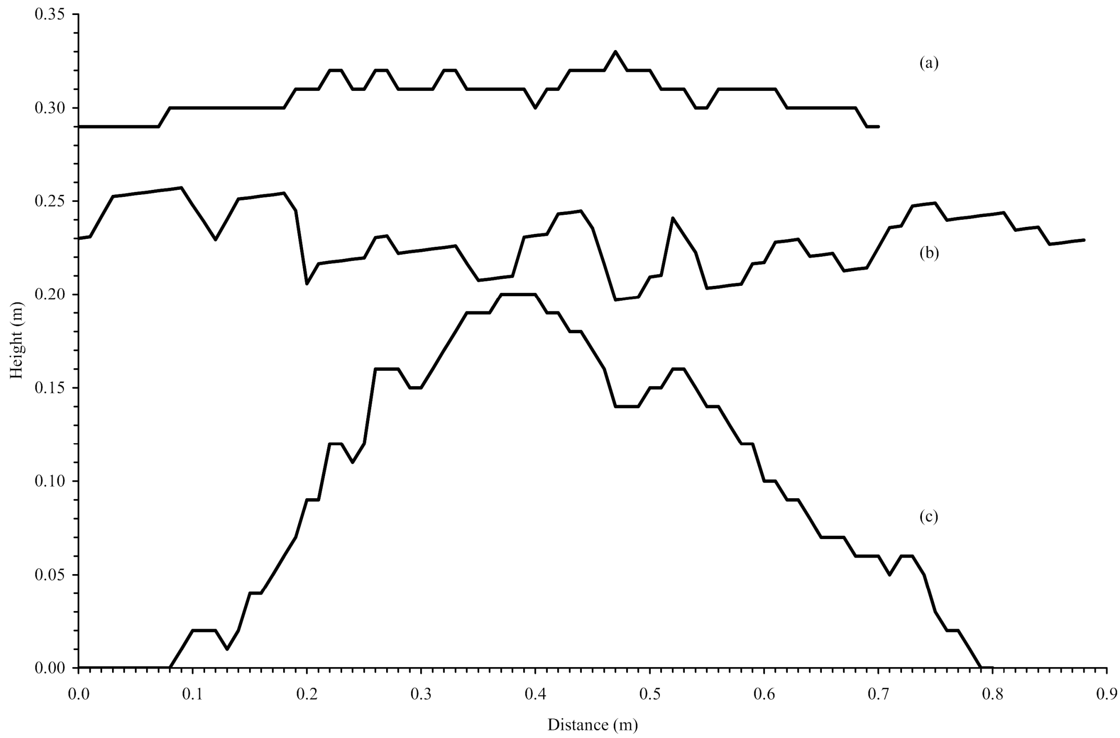

Figure 10

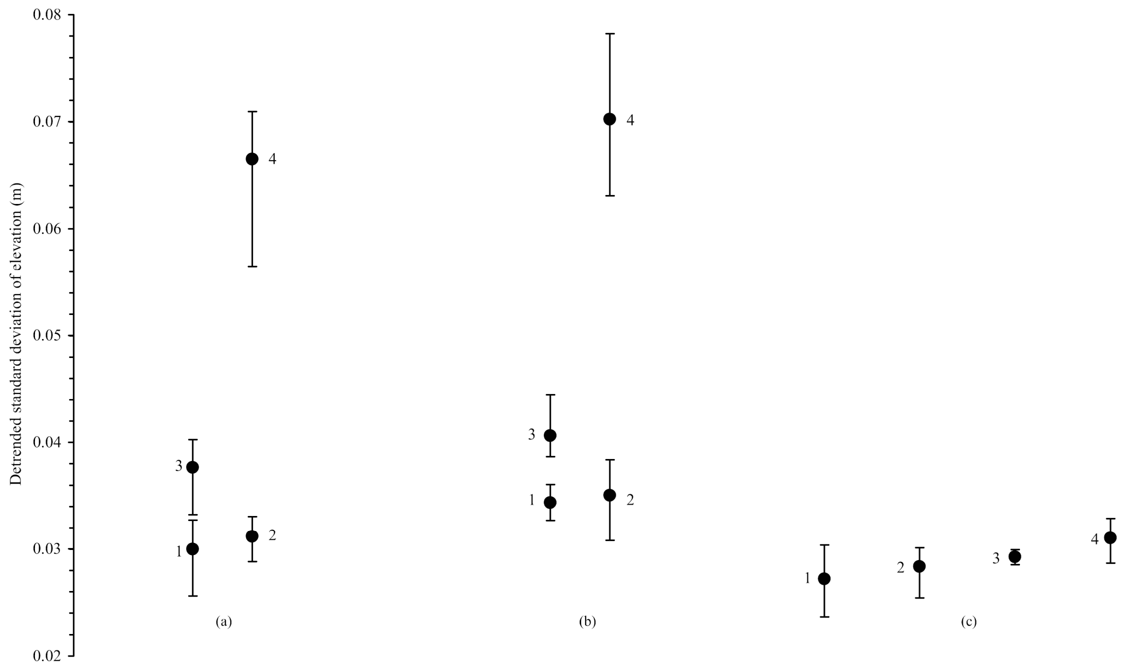

Figure 11

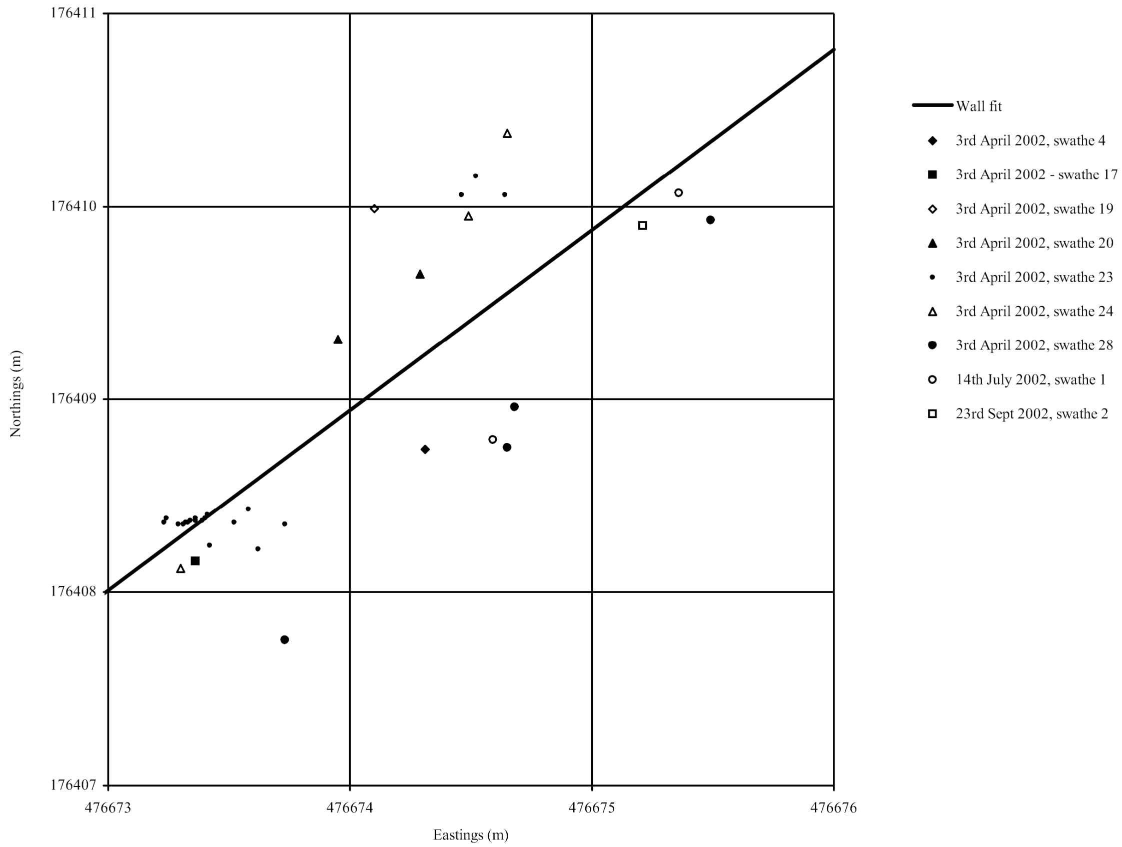

Figure 12

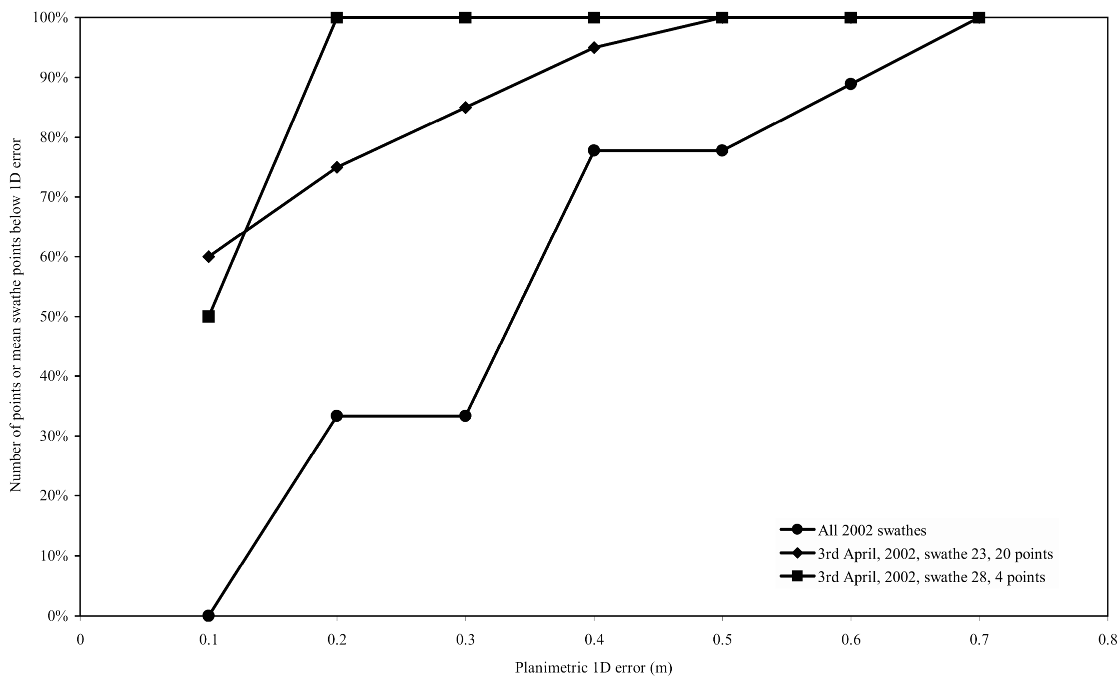

Figure 13

Author photo – Ian Davenport

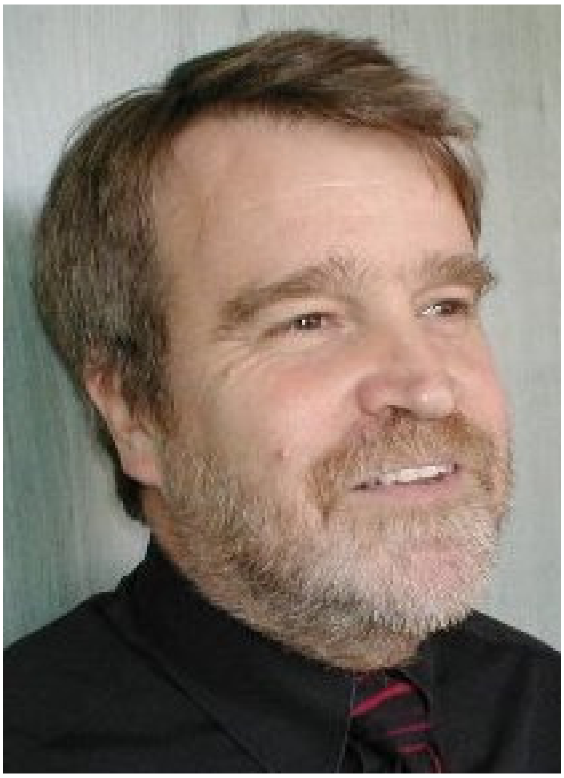

Author photo – Nick Holden


Author photo – Robert Gurney