The Accuracy of The Munro Society Heighting Surveys

Introduction

During 2009, four mountains namely Beinn Teallach, Ben Vane, Sgurr a'Choire-bheithe and Sgurr nan Ceannaichean were surveyed using differential GPS and their heights measured. The results of these surveys were announced at a Press Conference held on 10 September 2009. Arguably, the most interesting result was the measured height of Sgurr nan Ceannaichean which at 913.43m is 1metre below 914.40m – the metric equivalent of 3000.00 feet. Consequently, this mountain was reclassified from the list of Munros to the list of Corbetts. The result for Beinn Teallach also aroused interest since its height was measured to be 914.60m, only 0.20m above the Munro threshold.

The announcement of these results has generated much interest and comment. Since the press conference, several people have asked to learn more about the techniques we use and the accuracy associated with the measurements. In the following paragraphs we hope to answer these queries by briefly describing the heighting process, and then critically assessing the accuracy of the techniques and measurements.

The Standard for Height Measurements

First, let us examine the standard on which height measurements are based? The standard used for almost all of the UK is the mean sea level at Newlyn in Cornwall taken from tide measurements made between 1915 and 1921. For practical purposes, the reference point is a bench mark near Tolcone School in Newlyn. The Newlyn reference is the STANDARD for height measurement and is NOT necessarily current mean sea level. All maps in the UK are based on the Newlyn reference irrespective of the current (2009) value of mean sea level there. Changing that would mean having to change all the maps of Britain and references to them! Consequently, concerns about changing sea level on the measured heights of mountains are unfounded and of no relevance to the argument

The Height Measurement Process

The measurement of the height of a mountain is a two stage process. The first stage is to find the summit position and the second stage is to measure the height of that position. The techniques used for these two stages in the surveys are different, and so the accuracy of each stage needs to be assessed separately. Once that has been done, the estimated errors in the measurements of the two stages can be combined to give an overall estimate of the error of the height measurement.

Summit Location

To locate a summit position we use a surveyors' automatic optical level, specifically a Leica NA730, and an aluminium staff extendable to 5m. The method is to set up the automatic level on a tripod at an appropriate point, and then take readings from the staff held vertically over the points to be measured. This process has to be carried out in a systematic way, but the lowest staff reading will represent the highest point. The ease of finding the summit position depends very much on the shape of the summit. A steep slope rising to a sharp rock is an easy summit to locate and indeed no equipment is needed in this situation. However, a flattish top with several “domes” will mean quite an extensive survey to eliminate candidates and locate the highest point. There is no defined method for seeking a summit, and each one has to be considered in the field and an appropriate approach worked out.

Accuracy in Summit Location

What are the errors associated with the equipment we used to determine the location of a summit? The staff is essentially an “aluminium ruler” calibrated in centimetres. To determine a summit location we only compare one reading with another, absolute staff-height readings are not required so the only important factor is that the unit-marks on the staff are consistent, which they are to 1mm in 1m. If an actual height difference is required between two points, then it is easily possible to measure this to an accuracy of 3mm or less. The automatic level is a more complex piece of apparatus. However, there is very little that can go wrong mechanically with an automatic level and these faults are easily detected by a number of checks that we carry out routinely. The most common fault is misalignment of the telescope from the horizontal, but again that is easy to detect by a standard calibration check procedure that we carry out regularly. Even if the instrument were incorrectly aligned, it would still have no effect on height difference measurements made with the level set up in a fixed position. Therefore, we can be completely confident that the automatic level and staff are capable of finding a summit position accurately and indeed they are the best equipment available for this task.

So if the automatic level and staff have the capability of finding a summit position what about operator error? Of course it is very easy to take an incorrect reading from the staff. However, in critical situations we swap operators to get a second set of readings to make sure a mistake has not been made. Alternatively we use a process called “triple point measurement” to eliminate error. This means taking readings from the staff from the upper, middle and lower horizontal lines in the level, averaging the three readings and checking that they agree to within 1mm of the middle measurement. If not, the measurement readings are repeated until they do. This is a mistake-proof method of eliminating an incorrect reading.

So we are confident the level and staff are working correctly, we have implemented measurement techniques to avoid operator error, but how do we identify the point that is the summit? If we find that the highest point is “natural” ground, how do we know for example that there is not higher ground buried under the adjacent cairn? We only identify summit positions if they have permanence. If the highest point is measured as a protruding rock then it must not be possible to remove the rock. Loose stones for example would not count as a summit feature. If the highest point appears to be ground, then we would not consider the top of heather or grass tussocks since they have no degree of permanence. In the case of Beinn Teallach the highest point was the top of a protruding rock. We could compare the height of this with the adjacent cairn and it was clear that there was no permanent feature buried in the cairn that could be higher. Since the rock was well defined, we would estimate that we had found the summit position to better than +/-0.01m. The summit of Sgurr nan Ceannaichean is at the base of a small cairn and we dismantled this to check the ground underneath. There was no permanent rock beneath so we took several staff readings to find the highest ground. From these readings we were able to estimate that we had found the highest point to an accuracy of +/-0.05m. The base of the main cairn to the North West was measured to be 0.9m LOWER. Although we did not dismantle this well-engineered cairn, it was not tall enough to conceal a >0.9m permanent natural feature. (A comment has been made by the OS that the spot height of 915m on the maps probably arose because the photogrammetric operator who made the map spotted to the top of the cairn on the aerial photographs.)

GPS for Surveying

So we are satisfied that we have a robust method of finding summit positions and we can easily estimate the probable height error introduced in this process. The next stage, using the GPS to measure the absolute height is the more complex part of the measurement and so the following discussion is centred on the accuracy of these measurements. We do not intend here to give a detailed account of how GPS works, only its key features, but will concentrate on the measures taken to make sure results of the highest accuracy are obtained.

GPS units receive radio signals from a system of satellites orbiting 22000km above the earth. The position of each satellite at any particular time is known accurately and is controlled by monitoring stations in different parts of the world. The distance of a point on the earth's surface from a satellite is measured by the time taken for the radio signal to travel from the satellite to the GPS unit. In principle, it is possible to calculate position and height coordinates for a point on earth from the signals received from three satellites. A fourth satellite is needed for a "time correction" because the atomic clocks in the satellites are much more accurate than the clocks in the GPS units. In practice, accuracy of all GPS measurements is improved if more satellites can be observed simultaneously. It is possible to find out from satellite almanacs which satellites will be in view at a particular place and time. We checked satellite availability in the planning of our Munro/Corbett surveys to make sure that we were receiving signals from at least 6 satellites, but usually 7 or more.

A survey grade GPS unit such as the Leica 530 we use, differs from hand-held and SatNav units in that it has a more complex receiver. Each satellite transmits two signals simultaneously, each at a different radio frequency. Hand-held/SatNav units lock onto one of these frequencies only in order to calculate positions. Dual channel receivers, such as the Leica 530, lock onto both frequencies and, using a process of measuring phase shift, calculate positions much more accurately than single-channel units. However, this in itself is not sufficient to yield the accuracy required for survey work.

The real improvement in accuracy comes from Differential GPS, which is the measurement of position relative to a reference point that is itself being measured simultaneously with another GPS unit. With this technique, it is possible to measure heights in the UK to an accuracy of a few centimetres. But how does this work? The Ordnance Survey operates a system of over 100 base stations, whose positions and heights are known very accurately. Each one of these base stations is fitted with a survey grade GPS (indeed many are equipped with Leica 530 instruments identical to ours) that collects GPS signals continuously. This is shown diagrammatically below. Point P represents the base station whose position and height (X, Y, Z) are accurately known. The position and height of Point Q are not known and are being measured by GPS. The GPS units at points P and Q are both receiving signals from the satellites. Let us suppose that due to distortions in the satellite signals at a particular instant P measures its height to be 1.325m higher than its known height. As the distance between points P and Q is small compared with the distances to the satellites, one can assume that the simultaneous measurement at Q is also 1.325m higher than it should be. So a correction can be made to Q to give the "centimetre" accuracy required. As the distortions of the GPS signals changes with time, the corrections necessary for the point Q will also vary with time, mainly caused by the changing satellite positions and variations in the condition of the troposphere/ionosphere.

Accuracy

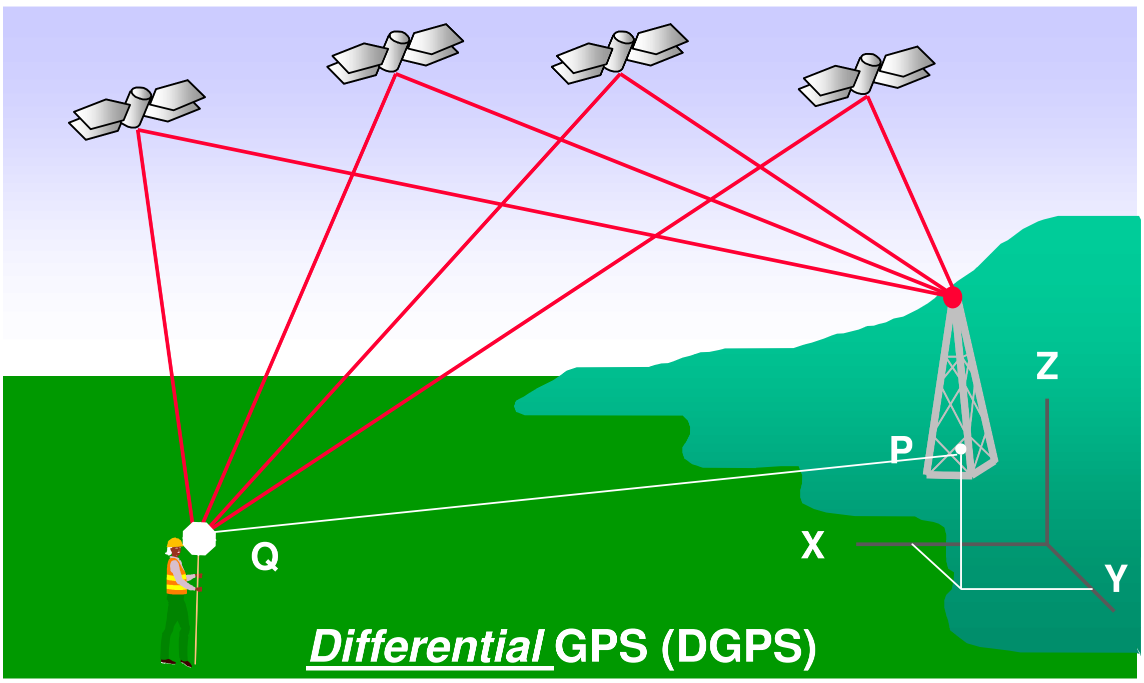

Relative Position… Good

Using GPS for Heighting

So that is an outline how survey grade differential GPS works in the UK. It is widely used by surveyors and is a well proven and trusted technology. A typical setup for making GPS measurements is shown in the next photograph. In fact this was the setup used to measure the height of Ben Vane.

The circular antenna at the top of the GPS is mounted on a pole which is firmly held in a tripod. The antenna has been optically levelled to the top of the rock behind and to the left in the photograph, which had been identified as the summit position on Ben Vane. The tripod legs are held down with rocks to prevent movement. Below the top of the tripod and fixed to the pole is the keyboard and display which allows communication with the instrument. To the left and at the base of the tripod is a small rucksack that contains the "electronics" and battery power supply. Operation of the instrument is quite simple. Once it has been switched on and has locked onto the satellites, a file is named for the survey. The parameters for the survey are set and the "go" button is pushed so the instrument can start collecting data. After the required time the survey is stopped and the data file is saved (automatically).

One important piece of information that has to be added to the survey setup is the "vertical offset", that is we have to tell the instrument how far off the ground the antenna has been positioned. Normally the measurement is made to the point at the bottom of the pole. The pole comes in two sizes 1.000m and 2.000m, so whichever size was used then that figure has to be recorded in the survey setup. Entering incorrect data would make a major error in the survey result. However, we always take a photograph of the setup, so it is easy to see which pole was employed and the "vertical offset" can be displayed subsequently if desired. Should we find later that the photo does not match what was programmed into the instrument, then this can always be corrected subsequently. Moreover, each of us checks the data being entered into the GPS at the time of the survey and this further reduces the chance of error. For Ben Vane we levelled the antenna to the rock so a "vertical offset" of 0.000m was used in this survey.

So, in addition to the measures already described, how do we further ensure that the data collection process is carried out to give maximum accuracy in the final result? To maximise the accuracy of any measurement it is important to take as many measurements as possible, but this is a process of diminishing return and so a compromise is needed. From previous research survey work carried out over the years, scientists at the Ordnance Survey are able to advise on this. They specify a minimum of two hours measurement, but preferably more if practically possible. For Ben Teallach and Ben Vane we collected data for 2 hours and for 3 hours on Sgurr nan Ceannaichean and Sgurr a'Choire-bheithe, thus satisfying OS requirements. During these surveys we set the instrument to take readings every 5 seconds. That means for Sgurr nan Ceannaichean and Sgurr a'Choire-bheithe, the OS was able to use 2160 individual measurements for its height calculations and 1440 measurements for the other surveys. Clearly this is a multiple measurement process and not a single measurement as has been misunderstood by some observers.

Once all the field work has been completed as described, the next part of the process is to calculate the position and height for the survey point from the raw data-set. Our agreement with The Munro Society was that the definitive result would be calculated by scientists at Ordnance Survey who use leading-edge Bernese software. We have no influence on this process. The OS is able, from the data, to give an estimate of the accuracy of the final measurement from the spread of results obtained for the individual data measurements.

Through our own software, Leica GeoOffice 6, we are also able to analyse the data-set and arrive at an independent result. If the results are very similar, we can be confident that they are robust. In the case of Ben Vane and Beinn Teallach, Leica Geosystems also calculated the results for us as a third check on accuracy! Whoever carries out the calculations will use the same methodology which is common to all those surveyors in the UK who carry out static Differential GPS Surveys. This is now described below.

As mentioned earlier, the data-set we collect needs to be corrected by data from the relevant Ordnance Survey base stations. So how does one get the correction data for this calculation? The first step is to decide which OS base stations you want to use. This is a fairly simple decision and will usually be the nearest five base stations to the point which has been surveyed. On the relevant page of the OS website, you can feed the four figure grid reference for the survey point and you will be told the nearest base stations and their distances from the survey point. The diagram below shows the five nearest base stations to Sgurr nan Ceannaichean and their relative positions. It is important that the base stations are spread around the survey point as intuitively this will eliminate any directional bias. It is also considered good practice to choose base stations not more than 90km from the survey position.

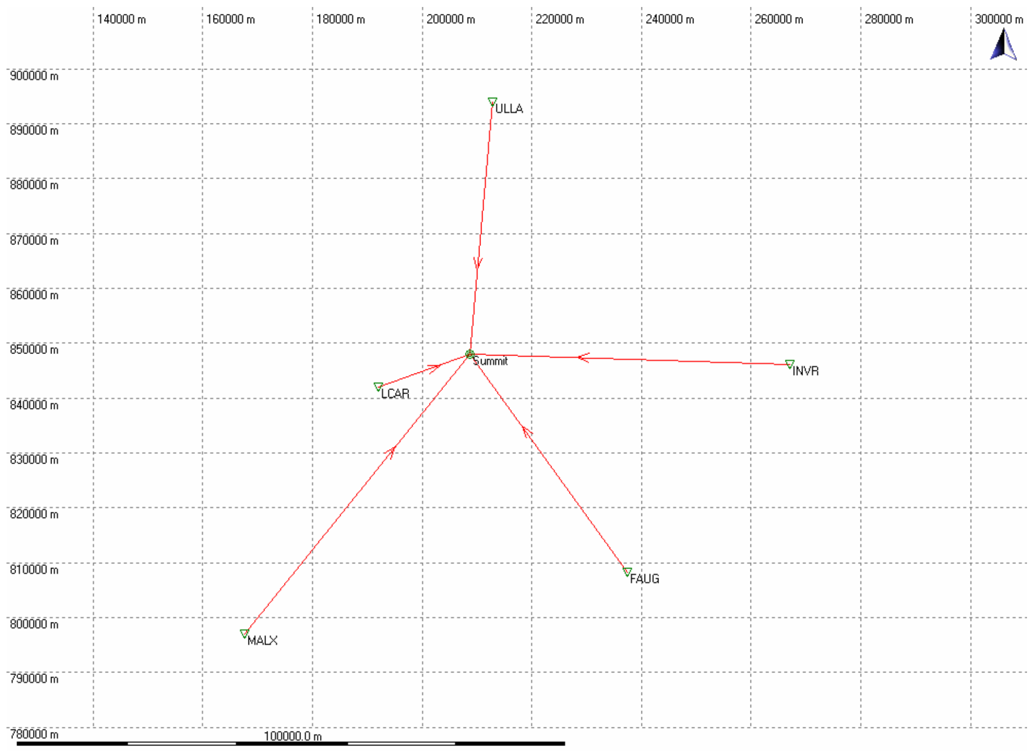

The next step is to obtain the correction data for the survey calculations. Again on the same page of the OS website, you feed in the date and exact time span of the survey and then the appropriate correction files (RINEX files) can be downloaded into the Leica GeoOffice software so that the calculations for position and height can be made. The only input to the software is to specify the tropospheric model used. For situations where there is a large difference in height between the base stations and the survey point, and data has been collected for 2-3 hours, the appropriate model is the "computed" one. The results for Sgurr nan Ceannaichean are shown in the table below.

Base Station	**Code**	**Distance from Survey Sgurr nan Ceannaichean km**	**Height m**	**Height difference from mean m.**
Loch Carron	LCAR	17	913.46	-0.005
Ullapool	ULLA	46	913.45	0.010
Fort Augustus	FAUG	49	913.45	0.013
Inverness	INVR	59	913.47	-0.015
Mallaig	MALX	65	913.47	-0.010

The mean result measured for the height was 913.46m with a variation of result of +/- 0.01m. The Ordnance Survey calculations via the Bernese software and a slightly different tropospheric model gave 913.43m. The two different methods of height calculation have given a difference of only 0.03m!

The above summary of the results of height measurements for Sgurr nan Ceannaichean show the results to be very consistent. That's fine, but how do we know that the instrument was functioning correctly on that day the measurements were taken? Would we get the same result, within a few cm. if we measured the mountain on a different day? To check that the instrument is giving correct results, we frequently measure the height of an Ordnance Survey Fundamental Bench Mark (FBM). The one nearest to where we live is at Daresbury in Cheshire. The results of the survey we carried out on 18 July 2009, three days after surveying Sgurr nan Ceannaichean, are tabulated below.

Feature	Survey	Height m.
Daresbury FBM	OS Data	73.24
Daresbury FBM	Leica 530 Measurement	73.23

The result is typical of previous measurements and is excellent agreement with the published height and proves that our Leica 530 is measuring heights correctly. In addition to the measurement of the Daresbury FBM, we frequently check the instrument against a more accessible control point for which we have our own results.

One question we have investigated experimentally is "how reproducible are survey heights for repeated measurements at the same point?" We have carried out 2-hour surveys of the height of a single point 20 times on different days and under different weather conditions. The Leica GeoOffice software also allows you to select the model of the troposphere that is used in the calculation (The troposphere has a significant effect in distorting the signals from the satellites and therefore this effect has to be modelled by the software to obtain accurate height measurements). We calculated each of the 20 datasets using a Hopfield model and a Computed model. The Hopfield model is usually chosen for short surveys where the height difference between the survey point and the base stations is small. Also given for comparison is the result for a continuous 12 hour survey. This represents the "ultimate" in static surveys since the results are computed from a complete cycle of satellite configurations. However, for mountain surveys it is totally impractical to consider taking measurements for 12 hours. The results are tabulated below.

Tropo. Model	Survey Time hrs		Height m.
Computed	2.0	Mean	136.43
Computed	2.0	Lowest value	136.46
Computed	2.0	Highest value	136.40
Computed	12.0	Mean	136.43
Hopfield	2.0	Mean	136.41
Hopfield	2.0	Lowest value	136.39
Hopfield	2.0	Highest value	136.43

The results for the mean of the 2-hour surveys calculated with the computed model is the same as for the 12-hour survey. The important point to note is the reproducibility of the results. Repeat measurements are falling in a band of +/-0.02m. The quite different Hopfield model gives a mean result which is 0.02m lower with a similar variation. So the results of twenty surveys calculated with different tropospheric models only give a difference of 0.04m; this demonstrates just how robust the technique is. Consequently, these results show that if we were to repeat the mountain surveys we would expect to see only a difference of 0.04m in the worst case, as far as the GPS component to the total error were concerned.

Summary

In this article we have covered in detail a description of the sources of error that can occur in the measurement of height and what we do to make sure that the results are accurate. We trust the reader will appreciate the care and attention to detail that we employ in order to ensure our survey results are of the highest accuracy.

So are we sure that Sgurr nan Ceannaichean is NOT a Munro and that Beinn Teallach IS a Munro? If we sum all the possible errors that were involved in the measurements of these mountains, we would quote the heights of each with a maximum estimated error of +/-0.1m. In statisticians' parlance, these errors represent three standard deviations and there is a 99.8% chance that the height of Sgurr nan Ceannaichean lies between 913.3m and 913.5m. Similarly there is a 99.8% chance that Beinn Teallach lies between 914.5m and 914.7m (statisticians never talk of 100% confidence). Where would you put your money?

J Barnard and G V Jackson, 30 November 2009.