ISSN: 1361-8962

Issue 31

COMPASS POINTS

July 2003

BCRA

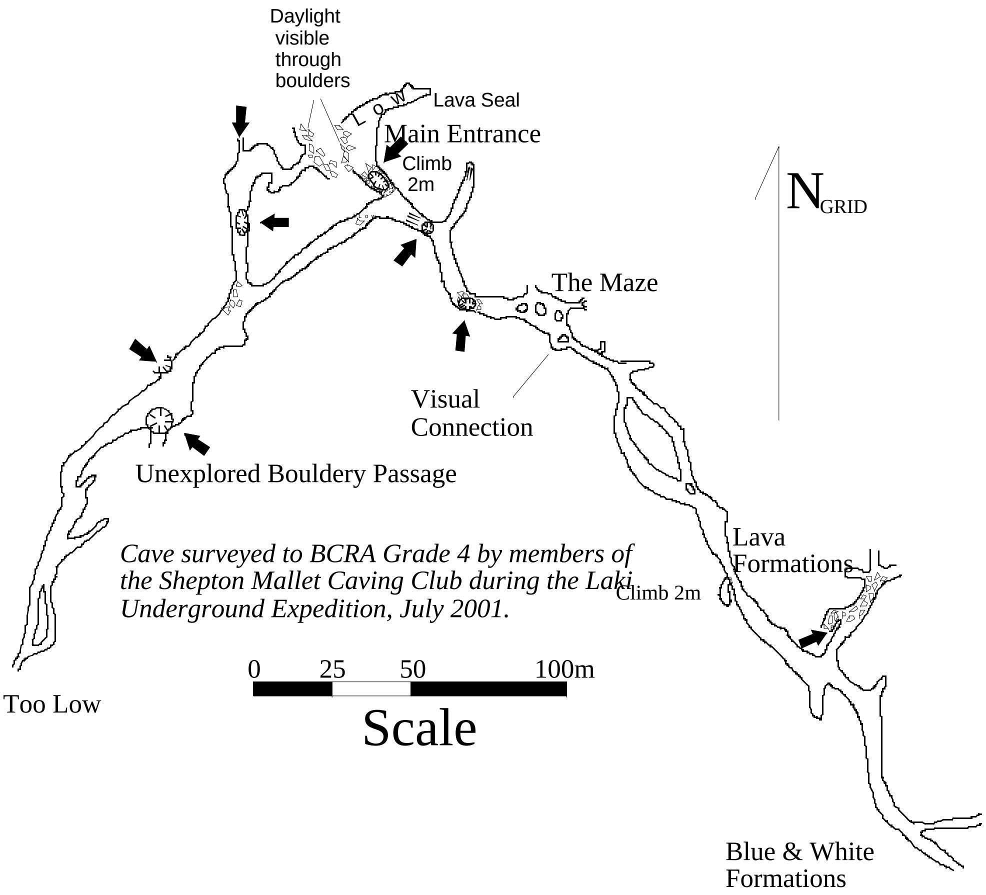

Cave surveyed to BCRA Grade 4 by members of the Shepton Mallet Caving Club during the Laki Underground Expedition, July 2001.

Surveying Icelandic lava tubes

Describing survey quality

Aerial photos from the Millennium Atlas

The Journal of the BCRA Cave Surveying Group

COMPASS POINTS INFORMATION

Compass Points is published three times yearly in March, July and November. The Cave Surveying Group is a Special Interest Group of the British Cave Research Association. Information sheets about the CSG are available by post or by e-mail. Please send an SAE or Post Office International Reply Coupon.

NOTES FOR CONTRIBUTORS

Articles can be on paper, but the preferred format is ASCII text files with paragraph breaks. If articles are particularly technical (i.e. contain lots of sums) then Latex, OpenOffice.org or Microsoft Word documents are probably best. We are able to cope with many other formats, but please check first. We can accept most common graphics formats, but vector graphic formats are much preferred to bit-mapped formats for diagrams. Photographs should be prints, or well-scanned photos supplied in any common bitmap format. It is the responsibility of contributing authors to clear copyright and acknowledgement matters for any material previously published elsewhere and to ensure that nothing in their submissions may be deemed libellous or defamatory.

COMPASS POINTS EDITOR

Anthony Day, 11 Claire Court, St Martins Drive, Walton on Thames, Surrey, KT12 3BW. Tel: 01932 226099
E-mail: `csg-editor@survex.com`

SUBSCRIPTION & ENQUIRIES

Andrew Atkinson, 31 Priory Avenue, Westbury-on-Trym, BRISTOL, BS9 4BZ Tel: 0117 962 3495
E-mail: `csg-secretary@survex.com`

PUBLISHED BY

The Cave Surveying Group of the BCRA. BCRA is a registered charity.

OBJECTIVES OF THE GROUP

The group aims, by means of a regular Journal, other publications and meetings, to disseminate information about, and develop new techniques for, cave surveying.

SUBSCRIPTION RATES (FOUR ISSUES)

U.K. - £4.50 Europe - £6.00 World - £8.00

These rates apply regardless of whether you are a member of the BCRA. Actual "membership" of the Group is only available to BCRA members, to whom it is free. You can join the BCRA for as little as £5.00 – contact BCRA directly for details. Send subscriptions to the CSG secretary (see "subscriptions and enquiries"). Cheques should be drawn on a UK bank and payable to *BCRA Cave Surveying Group*. Eurocheques and International Girobank payments are acceptable. At your own risk you may send UK banknotes or US$ (add 20% to current exchange rate and check you don't have obsolete UK banknotes). Failing this your bank can "wire" direct to our bank or you can pay by credit card, if overseas. In both these cases we have to pay a commission and you should pay 5% extra to cover this.

BACK ISSUES

Past issues of *Compass Points* are available from the secretary (see "subscriptions and enquiries") subject to availability. Cost is £1.25 per issue, plus postage and packing at rates of £0.50 (UK), £1.50 (Europe) or £3.00 (world). Published issues are also accessible on the Web via the CSG Web pages at `http://www.bcra.org.uk/csg/`

DATA PROTECTION ACT (1984)

Exemption from registration under the Act is claimed under the provision for mailing lists (exemption 6). This requires that consent is obtained for storage of the data, and for each disclosure. Subscribers' names and addresses will be stored on computer and disclosed in an address list, available to subscribers. You must inform us if you do not consent to this.

COMPASS POINTS LOGO

courtesy of Doug Dotson, Speleotechnologies.

CAVE SURVEYING MAILING LIST

The CSG runs an e-mail list for cave surveyors around the world. To join send a message containing the word 'subscribe' in the body text to `cave-surveying-request@survex.com`

CONTENTS
of Compass Points 31

The journal of the BCRA Cave Surveying Group

Cover image: Example of a CAD generated survey from the Laki Underground Expeditions (see article starting on page 5).

Editorial

Once again, this editorial is a grovelling apology for the late arrival of an issue of Compass Points. On this occasion, it was 99% complete when I disappeared off on expedition for three weeks - hence the delay. I really will try to be more timely in future, though producing the next issue is likely to be a little fraut (see admin. item) – so if anyone is sitting on any potential articles, please let me know about them in good time.

CSG Admin

The editor is on the move. The editor's postal address in the masthead will cease to work on 13th September, and at the time of going to press I don't know exactly where I will be after that. Please direct any paper correspondence to the "Subscriptions and Enquiries" address after this date. However, email can still be sent to the `csg-editor@survex.com` address.

Forthcoming Events

Hidden Earth 2003

The UK's annual national caving conference will be held at Hanley Castle High School, Upton on Severn (near Worcester) over the weekend 3-5 October. Further details can be found at the conference website, `http://www.hidden-earth.org.uk`

The Arthur Butcher award is judged and presented by BCRA annually at the conference for, broadly speaking, "excellence in cave surveying." There is a cash prize and a trophy to be kept by the winner(s) for a year. To be considered, individuals or caving clubs must bring their work to the attention of the judges. This can be achieved by displaying your work at the conference. If you want other work to be considered – such as a report or publication on a surveying topic, or other more general achievements – then you should contact the judges in advance. Full details of the rules of the award and nomination procedure can be found on the Hidden Earth website.

Snippets

Instrument Problems at Altitude: Update

Wookey

In the last issue of Compass Points, I described the problems various surveyors have had when using typical cave-surveying compasses at high altitude. In summary, based on practical experience it appears that Silva Clinomaster and Sightmaster can develop bubbles when used at altitudes in excess of ~1700m for a period of more than a week or two, rendering them unusable in extreme cases. Such bubbles arise because the capsule expands slightly in the lower pressure but the volume of liquid inside remains almost constant, hence a bubble forms. Suunto instruments and Silva Type 80 compasses (the plastic-bodied prismatic ones) appear to be less susceptible to this problem. It is also strongly advised that the instruments be carried in cabin baggage on aeroplanes to avoid them being subjected to extremely low pressures in the hold.

Since this article was printed, some more information has been received from the manufacturers. The Silva technical support staff state that the accepted altitude at which bubbles occur is 5000m, and that such bubbles will only become permanent if taken above 12000m (in an unpressurised aeroplane hold, for example). They accept that our experience differs significantly from the design specifications. Suunto have also been contacted, but the only information currently available is for their wristop computers, which should be ok up to 9000m. A request for information about their compasses is working its way through the system. As I said in the last issue, any feedback on your experiences of the susceptibility of various compasses to this problem would be appreciated.

Press Round-up

Compass and Tape issue 52 (Dec 2002)

Reviewed by Wookey

Call for papers at the 2003 NSS convention, Porterville, California.

Minutes of 2002 Survey and Cartography Section meeting

Held at 2002 NSS convention - 28th June. The section has nearly $5000 and 211 members. There were 4 talks at the conference and 14 entries in the Cartographic Salon. They had not produced the intended Cart Salon special issue of C&T as intended. Putting Salon maps online, the south-eastern region salon, and the possible definition of computer-drawn maps was discussed. Next year's convention will have a computer workshop of digital map creation and the possibility of a children's surveying class was considered. Robin Barber was voted secretary.

2002 Cartographic Salon Results - Steve Reames

20 maps were entered - 6 being "display only". 8 maps received awards. Caves of the Snake Well Complex, by Brent Aulenbach won the Best of Show Medal.

Development and testing of three components of the process of transferring Digital Cave Survey Data - Mike Yocum

Mammoth Cave National Park (MCNP) needed to integrate their Flint Ridge Mammoth Cave survey data (collected by the CRF (Cave Research Foundation) over 40 years) with the Arcview GIS system for park/karst management. Thus a project was started to: 1) develop and test procedures for converting CRF survey data to arcview format; 2) provide metadata for this info. (i.e. how accurate it is); 3) develop a framework and procedures for recording the content and status of CRF data as it is conveyed from CRF to MCNP.

A team of experts was assembled and digital data assembled for a test area in suitable formats and datums for Arcview. The survey data was in Compass, CML, SMAPS and Walls using a special CRF co-ordinate system. The test area included both underground and surface benchmarks which provided suitable control points.

Walls exports Arcview Shapefiles directly, Compass files can be converted to Arcview using the CaveTools Arcview extension. SMAPS and CML data need to be converted into Compass or Walls first. The most complete dataset was in SMAPS so it was imported into both Walls and Compass. Versions were also created with the CRF co-ordinate system converted to UTM, NAD27 metres. All four datasets were saved as shapefiles.

Meta data were generated from discussion amongst the team and from the results of registering the data against the Arcview surface data. The Shapefile format was extended to include extra survey attributes. The framework for tracking data is to reference each station to the relevant map sheet in the CRF Mammoth survey.

A discussion of the process and results follows. The most significant being that the SMAPS->Shapefile conversion process via Walls worked OK, but via Compass and Cavetools the historic entrance moved by over 3000ft, due to a difference in the constants used in SMAPS, Compass and Cavetools combined with a rounding error in Cavetools.

It also became clear that the various co-ordinate system conversions and different surface data feature formats (DOQQ, DRG, DLG, DEM) were not in agreement and there were various discrepancies of up to 100ft. You need to be very careful when converting data.

One result of this work has been an improved shapefile format - this is published as an appendix.

Obituary - Mike Yocum

Mike died shortly after submitting the above article. He was a keen cave-surveyor and project caver in Kentucky and Tenessee, being particularly active with the CRF and working for MCNP in the nineteen nineties.

CREG Journal 52

Reviewed by David Gibson

Although it was not billed as a “special” on radiolocation or cave surveying, the latest journal from the BCRA's *Cave Radio and Electronics Group* devotes 18 of its 36 pages to these topics in six articles.

Aerial photography of the entire UK is increasingly available, both electronically and as hard copy. Imagery from other countries is also available, either from aircraft or satellite. In *A Caver's View of Remote Imaging* Mike Bedford looks at what is on offer and the possible benefits to cavers. Unfortunately, a gremlin in the word processor meant that two of the photographs were re-sized and their captions consequently disappeared. However, a corrected version of the article is available at:
`http://bcra.org.uk/cregj/j52-pp0507.pdf`

At the CREG field meeting in March, Wookey gave a talk on *Electronic Cave Surveying Instruments* highlighting, in particular, how CREG members could help to design equipment. His article in CREG journal 52 is an extended version of that talk. It gives some background for non-surveyors, and outlines the pros and cons of several types of electronic instruments. It also gives a “time line” showing who has been working on what over the years, and gives a list of further reading. CREG members are scouting around for projects at the moment, and there would seem to be plenty of ways for people to get involved in a project to help to “automate”cave surveying.

The current state of the art in professional surveying equipment is the 3D laser scanner, described by Mike Bedford in *First Impressions of a 3D Laser Scanner*. The heart of a 3D scanner is a laser-based distance measurement device. However, instead of measuring the distance to single points as defined by the user, a 3D scanner automatically builds up a three-dimensional image by scanning the scene both vertically and horizontally. Data are collected automatically as a so-called 3D point cloud once the area of interest and the scanning interval have been specified by the user. Applications include quarrying and architecture.

An obvious caving application of a 3D laser scanner is surveying large cave chambers. With conventional cave surveying techniques it may be possible to do little more than survey the perimeter. Using a 3D scanner, though, it is possible to build a complete three-dimensional representation. Mike describes how he used the *Laser Ace Scanner* from MDL (Measurement Devices Limited) (`http://www.mdl.co.uk`, `http://www.laserace.com`) in Yordas Cave. The article includes a graphic showing the rendered view of one wall of the chamber.

The MDL instrument is a combined 3D laser scanner and reflectorless “total station” (to use the professional jargon). It costs £20,000, including the software for viewing 3D point clouds. This compares favourably with a typical price of £15,000 for a reflectorless total station and between £60,000 and £100,000 for most 3D laser scanners. Mike also describes MDL's *Laser Ace 300* which is a hand-held device that takes a single reading of range, bearing and elevation. Its use as a cave surveying instrument is self-explanatory, and at around £2400 it is more within reach of a caving group, although not by much.

David Gibson presents a ground-breaking paper [well, it *is* me writing this review!] on radiolocation showing how we can determine the bearing, elevation and distance of a beacon transmitter by making measurements at a single radiolocation receiver station. Conventional radiolocation requires two stations in the same horizontal plane, with one of them being at “ground zero”. Clearly for a generalised 3D method (i.e. a “global positioning” method) we cannot rely on this being possible, and so the single-station method would seem very attractive. In *3D Radiolocation Using a Single Station* David describes the limitations of the method, and the errors that are introduced when skin depth (a measure of the electrical conductivity of the ground) is taken into account. Very few discussions of radiolocation consider skin depth, which can be a significant cause of loss of accuracy. David also explains why “time of flight” methods that rely on timing the arrival of radio waves (i.e. like true GPS) will not work in a sub-surface setting.

The “single station” method described in the above article depends on an accurately calibrated transmitter, unless a ratiometric method is used (David suggests measuring the gradient of the vertical field component). However, this complication can be avoided by taking a second receiver reading. This technique is described by Richard Rushton in *Towards 3D Radiolocation Using Two Stations*, in which he explains how readings from two receiver stations can be combined to give position information *without* an accurately calibrated transmitter, provided that the relative position of the two receiver stations is known and they are in the same horizontal plane.

As a footnote to this article David Gibson then observes that Richard's method can be generalised to any two receiver locations, and he suggests a further logical step, which is to use the “single station” method to survey an entire system, and then to scale the resulting 3D survey in order to fix some known points. This technique avoids the complication of a calibrated transmitter and receiver, and it also avoids the need to fix pairs of stations during the “two station” method.

Both these articles describe a technique that sounds attractive in theory although, in practice, there are limitations on how useful it is likely to be. Clearly further work is needed.

David Gibson's *Bibliography of Cave Radiolocation* is an extended version of the article that appeared in Compass Points 30 and lists 30 references to the subject, most of them by David and most of them from the CREG journal.

The remainder of CREG journal 52 contains several articles on the HeyPhone sub-surface radio system, as well as the usual collection of news, letters and short articles. A contents list for CREG journals is available on their web site at `http://bcra.org.uk/creg/` and back issues can be ordered from:
`http://caves.org.uk/payments/creg/`

Letters

Compass Sight Errors

Bob Thrun

Jos Burgers presented a nicely balanced set of compass reading data (*in CP30 – ed.)* It is easy to separate the effects of compasses and compass readers because everybody read every instrument. The main reason that compass error was much more than reader error was that Compass C was different from the other five compasses by about 3 grads. If you remove it from the mix, there is about as much variation in readers as there is in compasses. Reader 2 had every reading below average. Reader 5 was generally high. Reader 11 was very close to the average and very consistent. Reader 6 was inconsistent. The readers who read to a tenth of a grad showed less variation than those who usually read to integer values.

BCRA Grade Definitions

Peter Cousins & Dave Irwin

We get the strong impression on reading issue Compass Points No. 30 that both John Stevens and the authors of the recent Cave Studies booklet have misunderstood Byran Ellis' use of the term “accuracy” in relation to survey grades.

Although we often felt that Bryan's definition of the difference between “precision” and “accuracy” (the nearness of the result to the true value) sat uncomfortably when used for magnetic cave survey readings; the distinction was nevertheless a useful one and, so far as we know, there is no other word in the English language we could use.

However, we are confident that the distinction which Bryan wished to make in introducing the term “accuracy” for Grade V and VI was simply that the survey instruments had been calibrated so that the readings could be referred to some local absolute standard [1]. It is no coincidence that both Bryan Ellis and Denis Warburton were trained as laboratory chemists and used calibrated instruments and glassware on a daily basis in their work.

In part the introduction of “calibration” for Grade V was also probably a reaction to the low standard of many of the then current Grade IV surveys, which were often published with a simple magnetic North arrow and without any elevations or real altitude data.

Unfortunately John Stevens has also misunderstood the entire grading scheme by suggesting that the grading should refer to the accuracy of the entire survey. It does not. The grade relates to the precision with which the instruments have been read (we are including tape and station position as an instrument here) and additionally, by suffix, to the quality of detail recorded (see below).

Implicitly the survey grading tells us nothing about the frequency of “blunders” or other unpredictable disturbances; although in the higher grades adequate attention to “before and after” calibration will allow the pooling of data from several instruments and largely eliminate the possibility of errors due to magnetic disturbance or instrument damage during the survey (e.g. Cousins [2])

John Stevens seems unaware of the long history of work by (for example) Denis Warburton and Dave Irwin regarding the relationship between observed loop closures and survey grades (e.g. Irwin & Stenner [3]), but we are surprised that he falls into the trap of thinking that it is always desirable to survey every passage to the same grade as the main survey so as to introduce as many loops as possible. The traditional Ordnance Survey practise was to reserve their main 36 inch theodolites for the primary triangulation and fill in the secondary and tertiary networks with smaller, but less precise, 24 or 18 inch instruments.

One problem with all these discussions is that authors rarely, if ever, define what is “the finished survey”. Until the digital age it was the published plan, elevations and other information as that was all that could generally be available. Now it sometimes seems that a table of numbers on some essentially transient website is all that matters. Either way the scale of final plotting largely defines the “accuracy” of the finished map (e.g. Irwin [4]); although we should not forget that until recently the Ordnance Survey only claimed 30m for minor features in rural areas - so even the most accurate cave surveys would not necessarily superimpose correctly on OS mapped surface features.

References

[1] Ellis, B.M. (1976). Surveying caves, [Bridgwater: British Cave Research Association.] 88pp.

[2] Cousins, P.R. (1980). Calibration of cave survey instruments, *Transactions of the British Cave Research Association,* Vol. 5(2), 85-90.

[3] Irwin, D.J. & Stenner, R.D. (1976). Accuracy and closures of traverses in cave surveys, *Transactions of the British Cave Research Association*, Vol. 2(4), 151-166.

[4] Irwin, D.J. (1970). The presentation of cave surveys, *Transactions of the Cave Research Group,* Vol. 12(3), 155-164.

Laki Underground Expeditions to Iceland

Phil Collett, Martin Ellis & Ed Waters

This article describes the surveying activities on two expeditions to Iceland undertaken by Bournemouth University with the assistance of the Shepton Mallet Caving Club. This work earned the Laki Underground team, and Chris Woods and Ed Waters in particular, the Arthur Butcher Award for 2002.

The Laki Underground Expeditions of September 2000 and July 2001 explored and surveyed over 11km of lava tube cave passages in the Eldhraun lava flow of southern Iceland. In addition to this “traditional” rationale for the production of cave surveys, surveying activities were carried out to support some of the specific scientific aims of these expeditions, both in the Eldhraun lava flow and the well known caves of the Hallmundahraun in western Iceland.

A key feature of the work undertaken was the successful integration of underground surveying with GPS measurements. Therefore, this article begins with a general discussion on the practical use of GPS for caving expeditions, before going on to describe the work of the Laki expeditions.

GPS and Caving Expeditions

In the last few years sales of hand held GPS systems have soared. Virtually every high street has at least one shop that stocks units manufactured by the likes of Garmin and Magellan. Typically these systems are credited with an accuracy of about 15m, though this is subject to a degraded performance of 100m when the US Department of Defence Selective Availability (SA) program is running. It is only since SA was switched off in 2000 that hand held GPS became a viable tool for caving expeditions. In addition, many hand held GPS systems have averaging functions that allow them to get a higher accuracy in areas with poor satellite coverage, which is a most valuable function and is highly desirable on units used for cave survey work. In these cases the accuracy of the location improves the longer the unit is left stationary – if the unit can be left for several hours an EPE (Estimated Positional Error – an estimated root mean square error in position) of below 5m can be obtained.

At its simplest, GPS is useful to a caving expedition as a quick and easy way of recording the location of cave entrances and other speleological features as waypoints. Experience has shown, however, that care needs to be taken when using a GPS in this manner. In areas with many cave entrances it is vital that a notebook is used to record what is at each waypoint. Once you have logged 40 or 50 features in a day, it can become incredibly difficult to remember why you recorded number 23 for instance. Noting the waypoint name (a coding for waypoint names can be useful also), location and a description of the feature can help prioritise the next day’s work and also guards against battery failure and accidental waypoint deletion.

As a minimum this technique should be accurate enough to relocate the feature logged at a later date. In cases where the features are so close together that it is difficult to determine which is which, the written description can be a big help. The moral to this story is that, if you do have some features very close to each other (say 20m or less), it is well worth making a more comprehensive set of description notes about them.

GPS track functions can be used to aid determination of the size and shape of surface features. For instance a depression can be quickly measured by walking around its rim. This can be a very useful aid to help drafting a survey, but this technique requires a GPS with a large memory or a laptop to download the information onto. GPS tracking functions can also be used to record paths and tracks not marked on maps, simply by carrying the unit with the party, which can be a useful aid for other parties venturing to the cave in question. However, some accuracy problems were encountered when mapping lava channels in Iceland using the track facility on the Garmin GPS 76. It transpires that the GPS 76 "throws away" track points when the track is saved as a named track (maximum of 128 points in a saved track), but maintains all track information in its current track. The workaround solution for the GPS 76 when using it to survey surface features using tracking is to never save a track using the Garmin facility, but to download the current track to a computer. This may work with other Garmin models, and is certainly the case with the Etrex.

For GPS-derived entrance fixes to be of any use in the cave survey, they must be properly tied in. In order for this to occur the best method is to make the place that the best GPS fixes are obtained into an external survey station and then treat it exactly the same way as you would treat a survey leg underground. For multi-entrance systems, the GPS fixes can be useful as a check that the underground survey is not grossly in error since it is quite likely that a competent party will produce survey data that is more accurate than the entrance fixes. For example, a 1% error (a good rule of thumb indicator for a fair cave survey) over 1000m of cave passage gives a positional uncertainty of 10m whilst a best possible GPS EPE of about 5m also gives a positional uncertainty of at least 10m. Similarly, if two or more cave entrances are close together, say less than 300m apart, it will be more accurate to survey between these entrances than to use GPS readings at each entrance. If suitable computing facilities are available on expedition, these techniques can expose errors in the field as they occur which allows them to be rectified if necessary or the dodgy surveying redone.

In cases where the cave entrance is in a cliff or in undergrowth (both are very common in the authors' experience), it is often well worth taking the GPS reading at the closest point to the cave that the minimum EPE can be obtained. The exact location of the entrance can then be determined by the use of normal cave surveying techniques.

In the opinion of the authors it is best practice to mark a permanent external station on the surface. Not only is this available for use in surface surveys or repeat sections, but it allows each party entering the cave to pause to take a GPS fix. If these fixes are recorded a series of them can give a better idea of the accuracy of the surface coordinates.

Finally, if the GPS data are to be integrated with surface maps, it is important that you record details of the datum in which your GPS co-ordinates were recorded and the datum of the surface map to assist you in any co-ordinate transformation that may be required. These issues have been discussed in detail in recent issues of Compass Points and will not be repeated here.

Differential GPS

Even if Selective Availability is turned off errors in indicated position can be caused by several factors, the greatest one being disturbances to the GPS signal by the ionosphere. Errors in the GPS system, including SA, can be greatly reduced by using a differential GPS (DGPS) system. A fixed base station GPS and a roving GPS are used. The same satellite errors apply to both GPSs, so the difference of position between the two stations is known to an accuracy that eliminates most error. A DGPS system based on the technology employed in hand-held units would typically be expected to give an accuracy of between 1 to 3 metres.

One advantage of the cooperation with Bournemouth University during one of the expeditions to Iceland was access to a survey quality DGPS system, the Leica GPS system 500, together with people skilled in its use. This system is able to process and store the GPS signals in a far more sophisticated way than hand-held units. Although having a precision an order of magnitude greater than hand held units (measurements to the nearest centimetre should be possible), these systems are many times more expensive. As with other differential systems the accuracy is relative to the base station. It is important to be able to replace the base unit in exactly the same position to within one centimetre. A typical Leica system would consist of a base unit on a tripod (Figure 1) and a roving unit with a battery pack in a small rucksack and the GPS unit fixed to a survey staff. In addition to very accurate location of cave entrances the Leica was used to lay out grids for geophysical surveys, for example to search for lava tubes using the magnetometer. The Leica can be used to make very accurate 3D representations of small areas in very great detail. For example a strange lava doughnut formation about 15m diameter and 1.5m high was measured and converted to a 3D image with little problem. To complete the task with conventional means to the same accuracy would have required a metal structure to have been erected to use as a reference point to measure to.

Figure 1: A DGPS Base station in use. As can be seen the base station is a large tripod mounted piece of equipment. The mobile unit is of similar size but fitted with carrying straps. [photo: Phil Collett]

Surveying to Record the Caves

Like all caving expeditions, the primary aim of the majority of the cave surveying undertaken on the Laki expeditions was to make as accurate a record of the caves that were explored. The techniques used for the actual surveying were pretty standard. Equipment used underground was Suunto/Silva compasses and clinos with Fibron tapes used to measure leg lengths. One novel piece of equipment used was a Leica Disto-Pro laser distance measurer. This was not generally used for leg length measurement, but was ideal for the collection of passage cross-sectional data, and for continuations.

The techniques used were generally in accordance with BCRA Grade 5, though the instruments were not calibrated in the field. Grade 4 is thus the highest grade claimed for any of the expedition's surveys.

A major potential problem in the surveying of lava tubes is that considerable local magnetic deviations can be present. In one incident an expedition member noticed a c.60° swing of a compass needle whilst moving only 5m. This should not have come as a surprise as the Shepton expeditions to Iceland in the 1970s had been aware of this problem, and carried out non-magnetic surveys to compensate. Sometimes we fail to learn from those who go before us!

One method of mitigating this problem would have been to assume that any magnetic deviation was constant at a point, and take forward and back bearings at each survey station. This would have allowed an "included angle" to have been calculated between each leg. Unfortunately the problem with this approach is that there would have been no certainty of the direction of the initial leg on which to base the subsequent angles on, without correcting to grid via a surface resection exercise. With hindsight it may have been possible to determine the bearing of this initial leg by having it on the surface, making it very long and using GPS to calculate the co-ordinates of the two stations.

Instead the approach taken was to use the GPS logging of entrances to provide confidence in the survey's accuracy and correct them to Grid North. This was an ideal situation for this technique as the majority of the caves have multiple entrances and Iceland is just about ideal GPS territory with its flat treeless lava flows.

The standard practice used was to plot the cave survey using traditional techniques and then place it over a piece of graph paper with the entrance locations obtained from GPS marked on it. The plotted survey was then rotated until the entrance positions tallied with the GPS locations. In the southern parts of the flow it was possible to use a differential GPS system to log cave entrance locations, whilst in the more remote areas hand-held units had to suffice.

In most cases the fit was astonishingly good, in the case of the cave Blámi with eight entrances, it was possible to rotate the survey so that all of the entrances were with 5m of the GPS plotted positions. This is of similar accuracy to the hand held GPS used to record the entrance locations. Overall it proved possible to close survey traverses to less than 2% error with these techniques, the closed traverse in Iðrafossar having an error of 8m over a 400m closed traverse.

This use of GPS also showed that the underlying magnetic deviation was not constant over the entire flow, but changed locally, with magnetic North varying from location to location. However, the data obtained over the wide range of the flow allowed reasonable estimates of local magnetic variation to be made for those caves with single entrances (such as Rauðsteinshellir).

A final "proof of the pudding" was made on return to the UK by overlaying the combined surveys of the Eldhraun caves with an aerial photograph of the lava flow. The cave survey data along with GPS entrance co-ordinates were used to construct an area map with a cheap CAD package (Total CAD, costs c.£10 from the cheap shelf of most computer shops). It was then necessary to rectify the photograph to a regular scale. This was done by identifying a number of known features whose location had been recorded by GPS. The photo was then scaled onto a set of points in the CAD package and the cave area map overlaid.

In the field it had been noticed that some of the shallow caves had inflated the lava flow surface above them, so producing a humped ridge above their passages. When the map was overlaid on the photo it became apparent that these ridges were visible on the photo. This means that this photograph provides a very powerful endorsement of the techniques used, as it was possible to match the surveyed passage to the ridges overlying the passages.

Of course some (minor) discrepancies were noticed, primarily the location of Þjönappahellir. Since the camera doesn't lie these aspects of the survey were corrected to the photograph. However, the correlation of the photo and the map did show some remarkable correlation. Perhaps the most impressive was the discovery that the chamber at the far end of Rauðbogahellir, from which daylight was visible through a boulder choke, lay directly beneath a depression on the surface.

Surveying for Scientific Aims

As well as the exploration, the expeditions had a series of scientific aims. One of these was to relate the caves to the morphology of the Eldhraun lava flows. Another was the use of geophysical techniques (principally magnetometry) to detect the presence of lava tube caves from the surface, both on the Eldhraun and on the Hallmundahraun lava flow in western Iceland.

For both of these aspects it was vital to accurately relate the caves to the surface. Some of the aims of the former aspect were satisfactorily dealt with by the combination of cave surveys and aerial photographs. However, for some of the more detailed work other techniques had to be used.

In the case of geophysical experiments it was vital to ascertain whether anomalies found were due to the caves below. For this to be the case the relationship of the cave to the surface survey traverses had to be known with some confidence. This was carried out with detailed DGPS survey areas marked out on the surface, which in turn were tied into detailed non-magnetic surveys of small parts of the underlying Surtshellir/Stefansellir system.

The data for the survey were taken from that obtained by the 1972 Shepton Mallet Caving Club Expedition. Although this dataset was complete it had never been drawn up or published. However, for the purposes of the experiments carried out it was invaluable as it was generated using non-magnetic methods, in contrast to all other surveys of these caves (before and since).

The equipment used in 1972 comprised simple tripod mounted cave theodolite, tripod mounted sighting targets and 30m fibron tapes. Backsights and foresights were taken at each station and the instruments read to the nearest ¼° in azimuth and ½° in elevation. Analysis of these data showed a misclosure of about 1% on the 490m closed traverse in the central section of Surtshellir. This error was shared between all survey legs for drafting.

During the 2000 expedition the 1972 data was supplemented with more accurate and frequent measurements of passage cross-sectional details, taken along a simple traverse line generated by compass and clinometer. These measurements were taken with a Leica Disto-Pro laser measurer. These data were later corrected to the line of the 1972 survey.

The line of the relevant cave passages were then located on the surface by repeating the 2000 survey traverse data from points of origin of the underground survey (using plumb bobs from the surface down into the entrance collapses). Though this work was carried out using magnetic equipment it provided a "good enough" starting point for the geophysical work.

Geophysical data were then collected along traverse lines and grids that were geo-referenced and located relative to the cave entrances with the use of a Leica GPS System SR530. This instrument was also used to plot the relationships between the traverse lines/grids, the edges of the surface collapses and the points of origin of the cave survey line traverses (plumb bob positions).

The care taken with the measurements in the field then allowed the geophysical data and cave surveys to be overlaid. This showed unambiguously that the underlying cave passages had been detected by the magnetometry techniques. In addition the survey showed that there is more cave passage, as yet un-enterable, upflow from the terminal lava seal in Stefánshellir, as shown in Figure 2.

This article is based on a longer article by the same authors entitled "Developments in Expedition Cave Surveying" which is due to be published in the forthcoming Shepton Mallet Caving Club Journal.

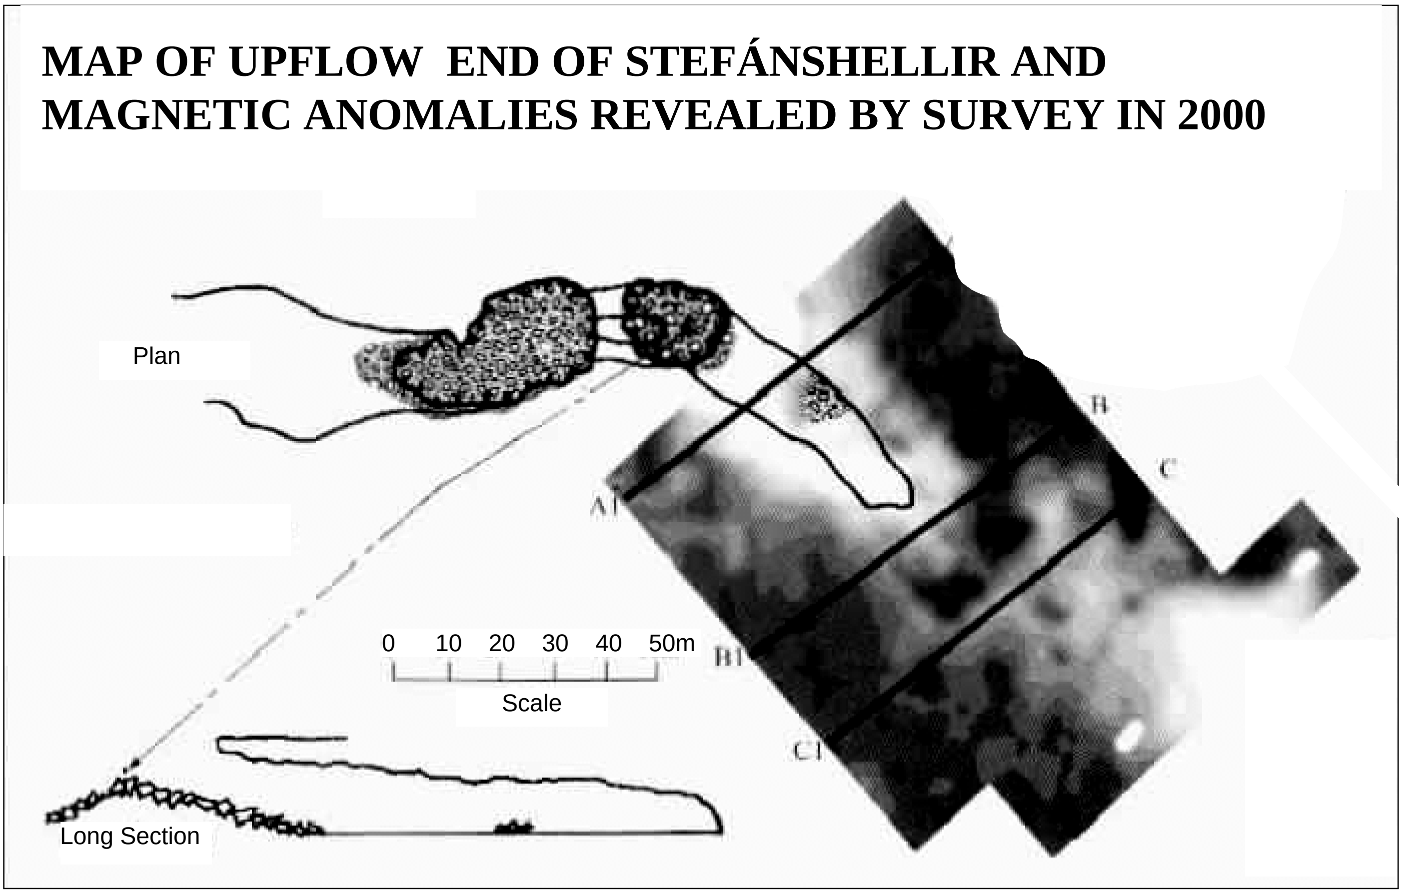

Figure 2: A survey of part of Stefánshellir with magnetometry survey results overlaid. The survey shows the magnetic anomaly associated with the known cave, the lava seal that terminates the cave, and an anomaly consistent with further cave passage beyond the lava seal. [figure courtesy of Chris Woods]

Describing Survey Quality

Bob Thrun

Statistics

In his article, "Proving Survey Accuracy", John Stevens overlooked some aspects of the way that random errors accumulate. The random error in the sum of measurements tends to increase as the square root of the number of measurements. This occurs because the errors sometimes add and sometimes subtract from each other. It can be hard to find this explicitly stated in a statistics book. This has been pointed out by some caving authors: Heinz Schwinge, Denis Warburton, Mike Luckwill, and Irwin and Stenner. Many cave surveys have a large loop where the percentage error is particularly low and the surveyors all brag about it. I found that many of my short loops, with only three or four shots, had large percentage errors. Both of these situations are simply to be expected.

Schwinge showed that, when there are only length and compass errors, there is an optimum survey shot length that minimizes the error when surveying a given distance. Schwinge did not have clinometer or station position errors in his derivation. The minimum error occurs when the error due to the angle measurement in each shot is the same as the error due the length measurement. For the 1976 BCRA Grade 5, this shot length is 5.73 meters. Station position error would favor somewhat longer shots. The average shot lengths in the loops from the Ogof Draenen survey that Stevens presented are close to the optimum.

Some authors calculate a standard deviation for a survey and equate it with "probable error". The term should be "most probable error" since any size of error has some probability associated with it. The definition of standard deviation involves only one variable, so it is one-dimensional. The most probable error in one dimension is zero. Standard deviation refers to the width of the probability distribution.

Some authors also confuse the one-dimensional and the three-dimensional error distributions. If there are independent and equal probability distributions in more than one variable, then we have a chi-square distribution, which has a non-zero most probable error. The errors in a single survey shot of optimum length are an example. If the errors on a traverse are the same in all directions, we have a good approximation of a chi-square distribution. The chi-square distributions in two and three dimensions have the special names of Rayleigh and Maxwell distributions. The chi-square distributions can be expressed in units of the one-dimensional standard deviation. These distributions are shown in Figure 1. Since we are concerned with the absolute value of the error, the values for the one-dimensional normal distribution are doubled. The cumulative probability, the probability that an error is less than some amount, is shown in Figure 2. The maximum, median, and mean values for the three distributions are given in Table 1. In Figures 1 and 2, and in Table 1, it is assumed that the errors are independent and the same in all directions.

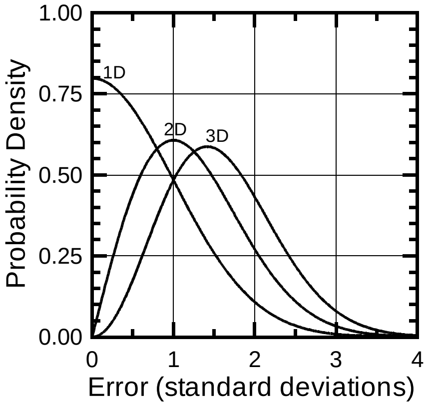

Figure 1: Probability densities in one, two, and three dimensions

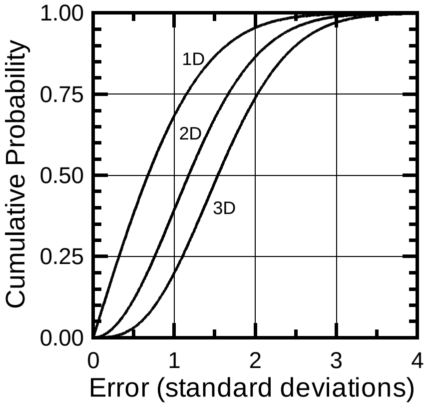

Figure 2: Cumulative probability distributions in one, two, and three dimensions.

	Maximum	Median	Mean
1-D	0.000	0.674	0.798
2-D	1.000	1.177	1.253
3-D	1.414	1.538	1.596

Table 1: Maximum, median, and mean values for one-, two-, and, three-dimensional probability distributions.

I did computer simulations for the loops that Stevens showed in his Table 3. I constructed loops with the same length and number of shots that he had. On each survey shot I added random errors from uniform distributions as specified in the 1976 BCRA Grade 5 standard. I did 1000 simulations for each set of loop conditions and made bar graphs showing the distributions from the simulations. On each of the bar graphs I added a symbol showing the closure error from the actual survey. These bar graphs are shown in Figures 3 to 11. For most of the loops, I constructed N-sided regular polygons, giving nearly circular loops. This ensured that the errors were the same in the two horizontal directions. If I did more simulations for each set of loop conditions, the bar graphs would more closely resemble the 3-D distribution in Figure 1.

I tried loops where half the shots went out in one direction and then doubled back along the same route. The doubling back had a slight effect on the overall error distribution of the loop. One of the doubled-back loops is shown in Figure 4.

All the survey shots in the simulated loops were of equal length. If both the total length and the number of survey shots are specified, the minimum error occurs when all the shots are the same length. The worst case is one long shot and a bunch of very short shots. My simulations predict less error than would more realistic simulations with varying length shots and the proper angles. Better simulations could be done by adjusting the loops and then adding random errors to the adjusted measurements.

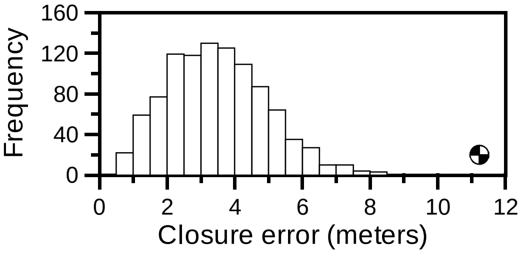

Figure 3: Probable loop closure error distribution for a circular loop 5083 meters long with 614 shots.

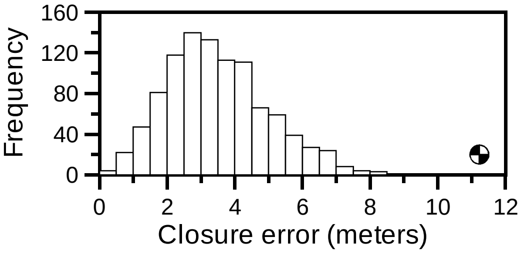

Figure 4: Probable loop closure error distribution for a doubled-back loop 5083 meters long with 614 shots.

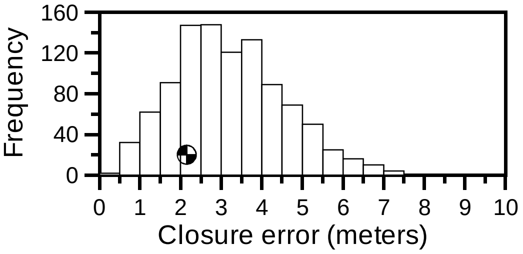

Figure 5: Probable loop closure error distribution for a circular loop 4158 meters long with 481 shots.

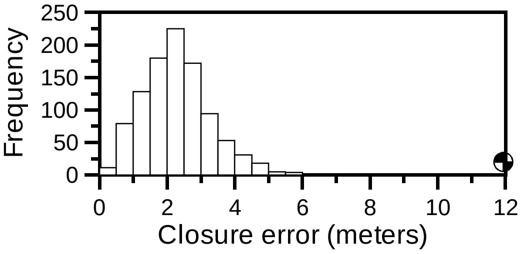

Figure 6: Probable loop closure error distribution for a circular loop 2197 meters long with 292 shots.

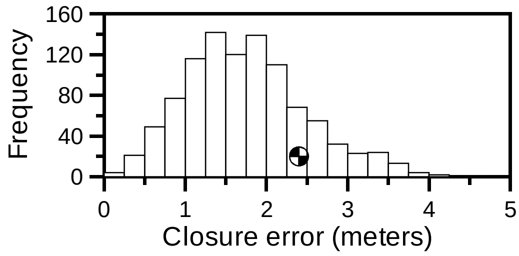

Figure 7: Probable loop closure error distribution for a circular loop 1314 meters long with 159 shots.

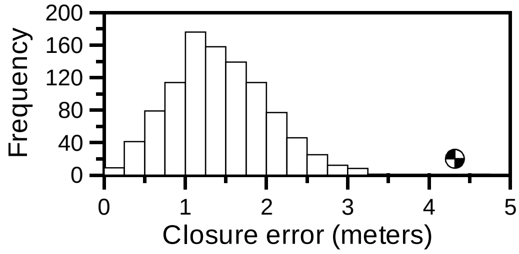

Figure 8: Probable loop closure error distribution for a circular loop 850 meters long with 139 shots.

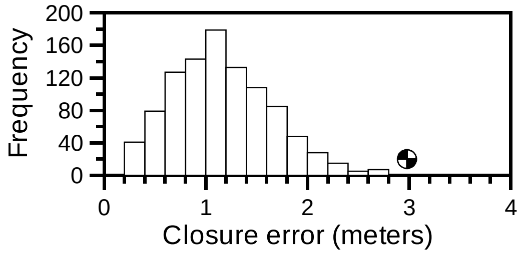

Figure 9: Probable loop closure error distribution for a circular loop 570 meters long with 102 shots.

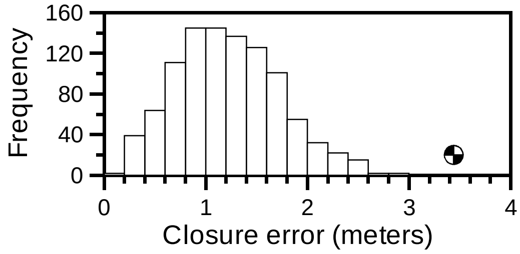

Figure 10: Probable loop closure error distribution for a circular loop 610 meters long with 71 shots.

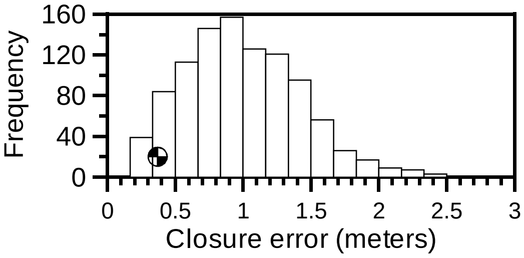

Figure 11: Probable loop closure error distribution for a circular loop 395 meters long with 38 shots.

I examined about 40 cave surveys for a talk I gave at the 2000 NSS Convention. None of these surveys met BCRA Grade 5 standards, though some came close enough that I feel there might be a Grade 5 survey somewhere. Stevens presented eight loops that he claimed to be Grade 5 or 6 quality. Of these, only three of the four claimed Grade 6 loops meet Grade 5 quality, while none of the claimed Grade 5 loops do. Two of the loops look like they might meet Grade 6 quality. Or they might just be some better-than-average loops from a lower quality survey. We would need to examine many more loops to be sure.

Finding Loops

For the first cave survey data reduction program I wrote, I had to specify the route of each loop. I found that even a small network has a large number of possible loops. As an example, consider a simple network with nine small squares making up one large square. If we count just rectangular loops, we get the numbers of loops shown in Table 2. There are other possible shapes of loops and many more loops in this network. It is easy to get over a hundred loops.

Number of loops	Size of loop
9	1x1
12	1x2
6	1x3
4	2x2
4	2x3
1	3x3

Table 2: Counts for some loops in a simple network, consisting of nine small squares making up one large square.

Irwin and Stenner described a survey adjustment method where they pick some key junctions in the survey and then average a few different routes from the entrance to each of these junctions. This method is best suited for a person working with a calculator. It would be hard to program a computer to find just a few distinct routes.

Larry Fish, in describing his COMPASS program, said it finds all the loops in a survey. I puzzled over this because I realized that a very large number of loops can be constructed. COMPASS finds a minimal set of loops that are called fundamental cycles of a graph in graph theory. Any loop may be constructed by adding cycles from a set of fundamental cycles. Where parts of two cycles coincide, they cancel out. There are many possible sets of fundamental cycles for any graph. Figure 12 shows the same network with two different sets of fundamental cycles. The set that is found by COMPASS is determined by the order in which the data are presented. WinKarst, by Garry Petrie, attempts to find small loops. Other cave survey programs use least-squares and do not explicitly find loops.

Figure 12: Two different sets of fundamental cycles for the same graph.

The problem of finding all the loops resembles the classic Traveling Salesman Problem. However, the salesman can go directly from any city to any other city, while the cave surveyor can go to only a few other survey junctions from any given junction. I have come to the conclusion that if there are N junctions in a cave survey network, there are between 2N and 3N loops in the network. Consider two points with N junctions on the route between them. At each junction there are 2 or 3 route choices. You can come to the same conclusion by considering the effect of an addition to a network. The exact number of loops will depend on the network. It will usually be more than I want to count, even with a computer program. I once suggested a random-walk program to gather loop statistics.

Assigning Grades

Bryan Ellis and others contended there were three basic types of surveys: the rough sketch, the fast survey, and the proper survey. They wanted to reduce the number of grades to three, but put in the non-favored grades because some cavers did not want their surveys downgraded from a 5 to a 3! I can't think of any instruments graduated in 5-degree increments to use for a Grade 3 survey. If the same surveyors use the same instruments for the two types of survey, they should get the same accuracy. The typical surface survey would be an example of a fast survey. The differences between the two types of survey are in the level of detail and thoroughness in going down every passage.

Surveys done with compasses that are marked in one-degree or half-degree increments have errors larger than would be indicated by the precision of the markings. I would not be surprised if compasses marked in two or five-degree increments produce less error than the precision of marking indicates. I see very few British newsletters. Are there any Grade 3 surveys done? If so, could somebody check the closure errors on some of these? It may be that Grade 3 is obsolete because no Grade 3 surveys are done.

All the surveys I looked at for my 2000 talk would have to be classified as 1976 BCRA Grade 4. Leaving out the very worst surveys, which had a lot of blunders, there were factors of 2 or 3 between the best and worst surveys, which brings us into Grade 3 territory. There was no obvious place to divide the surveys into good and bad categories. The spread in the values of the individual closure adjustments within any one survey was at least as wide as the error simulations I show here. It looks to me like there is no Grade 4 gap between Grades 3 and 5. Instead, they all merge into one broad grade.

I have a copy of the Ogof Draenen Grade 2 survey data. The compass is always read to the nearest degree and the distance is read to 0.1 meter. The 0.1 meter accuracy would qualify for Grade 5 in the 1976 standard, but not the 2002 standard. There would be little difference in the overall accuracy between 0.1 and 0.01 meter accuracy because the larger errors due to angle measurements dominate. A declination correction is made, although the compasses and readers are not individually calibrated. The main difference between the data I have and a Grade 5 survey is the lack of clinometer readings. Perhaps someone could compare the Grade 2 and 5 maps.

Some surveyors publish maps with a magnetic north arrow. Some get the declination from a map or a magnetic field model. Some get a combined declination and compass correction by a surface shot. Some do surface shots in multiple directions to get a correction for compass eccentricity error. I will leave it to others to argue if calibration or the lack of it is enough for a separate grade.

The wording of the BCRA grades gives the impression that baseline accuracy is the most important aspect of a cave survey. We all try for accuracy in our surveys, but does it really matter? For most uses, the amount of detail and the completeness of the map are more important. An inaccurate survey is adequate as a road map as long as there are no blunders. On the other hand, even the most accurate magnetic survey is not adequate for drilling a new entrance far from the original entrance.

William E. Davies produced many maps for his book, Caverns of West Virginia. He often mapped small caves on consecutive days, so they must have been compass-and-pace maps. In comparison with modern maps, his angles may be off by up to 10 degrees and his distances may be off by 30 percent either way. The only way to tell is to overlay one of his maps with a new map. He was a good observer and his maps often have as much passage detail as the modern maps. He did not push down every small passage.

The BCRA detail gradings are worded awkwardly. Rather than being based on how much detail there is in a map, they are based on whether or not the details shown were measured. A map with less detail could have a higher detail grade than a map with more detail, provided that all the details shown were measured. It should be obvious to anyone looking at a map how much detail it has. And what about small versions of large maps?

British cavers are accustomed to grades and regard them as a measure of quality. Others use BCRA grades with the mistaken idea that they are scientific. The use of a grade is an implicit claim of accuracy, but there have been very few efforts to determine if a survey meets the specification. There were Irwin and Stenner, myself, and now John Stevens. Any others? Most British cave surveyors simply called a Compass, Tape, and Clinometer survey (a clearer term that I prefer) Grade 5. The 2002 revision brings the BCRA grading system in line with actual practice and changes the system from a failed attempt to quantify the accuracy of a survey to jargon for the sake of jargon.

If you want to describe the quality of a survey, using a system that is clear, meaningful, and unambiguous, I suggest the following:

- No map.
- Sketch with no measurements.
- Compass and Pace.
- Compass and Tape
- Compass, Tape and Clinometer.

Something is either measured or not measured in this classification. And there is no need to use code. This classification says very little about the accuracy of the survey, but the BCRA grades don't say much either.

Some cavers may want to assign a numerical value to the accuracy of their surveys. Two possible measures for a single loop are $Error/\sqrt{N}$, where N is the number of survey shots in a loop, or $Error/\sqrt{length}$. The first of these will be roughly proportional to the error per shot. The second will be roughly proportional to the angular error on the longer shots. Neither of these measures takes into account the lengths and directions of the individual survey shots. For a given survey, there will be a mixture of good loops and bad loops, long loops and short loops. It is possible to display the values as a point cloud, similar to the way I showed least-squares adjustments in my 2000 talk. I can't think of a good non-pictorial way to summarize the statistics for an entire survey. To use either of these two loop measures, it is necessary to explicitly find loops.

I can't think of any way of evaluating the survey accuracy that takes into account the survey shot lengths and directions that does not make a comparison to an assumed standard. If that is done, there is still the problem of summarizing the results.

References

Davies, William E., Caverns of West Virginia (2nd. ed.), West Virginia Geological and Economic Survey, Vol. XIX A, July 1958

Ellis, B. M., "Changes in the System of Grading Surveys", Cave Research Group of Great Britain Newsletter, No. 132, January 1973, pages 15-17

Ellis, Bryan M., "Changes in the System of Grading Surveys", Cave Research Group of Great Britain Newsletter, No. 134, June 1973, pages 4-6

Ellis, Bryan, Surveying Caves, British Cave Research Association, 1976, 88 pages

Fish, Larry, "The Problem With Least Squares Loop Closures", Compass & Tape, Vol. 13, Issue 1, No, 41, April 1997, pages 15-18, also in Compass Points 15, March 1997, pages 8-10, and at `http://www.fountainware.com/compass/compart.htm` (mentions "sorting all of the loops in the cave")

Irwin, D. J. and R. D. Stenner, "Accuracy and Closure of Traverses in Cave Surveying", Transactions of the British Cave Research Association, Vol. 2, No. 4, December 1975, pages 151-165

Luckwill, Mike, "A Brief Review of the Theory Available to the Cave Surveyor", Belfry Bulletin, Vol. 23, No. 4, pages 44-51

Schwinge, Heinz T., "The Accuracy of Cave Survey", The National Speleological Bulletin, Vol. 24, Part 1, January 1962, pages 40-47

Stevens, John, "Proving Survey Accuracy", Compass Points 30, March 2003, pages 11-13

Thrun, Robert, "How Accurate are our Cave Surveys", (abstract only), Journal of Cave and Karst Studies, Vol. 63, No. 1, April 2001, page 54

Warburton, Denis, "The Accuracy of a Cave Survey", Wessex Cave Club Journal, No. 89, Vol. 7, April 1963. pages 166-181

Printable Mapping from the Millennium Atlas

Andy Waddington

You may have heard about the "Millennium Atlas" - an effort to provide a consistent set of aerial photography of the UK at a high resolution. By "consistent", they mean that all the photography has been flown in similar weather conditions and from a similar height, such that photographs from distant areas may be compared like-for-like. The result of all this is that you can buy (for really quite a lot of money) a photograph of (almost) any 500x500m area in England and Wales at a resolution of 0.25m per pixel on the ground. There are also various spins-off like large-scale maps of popular tourist areas with grid, contours and other information superimposed, and a big expensive Atlas on paper.

What is perhaps less well-known is that, at 2 metres per pixel resolution, all this photography is available for free on the web at the `GetMapping.com` site, provided as part of the interface used for choosing the photos they would like to sell to you.

Why might you want aerial photography rather than a good old-fashioned Ordnance Survey (OS) map with all the advantages of names for features, definitive symbols for things that you might not recognise from the air and so on? You will still want maps for all the information that aerial photos don't contain, and which the OS, and others, have kindly gathered by surveying "on the ground", such as rights of way, altitudes, names and "tourist information". However, aerial photos show much detail that is not translated into mapping, some of which is particularly useful to the caver producing surface surveys, searching for new cave, or simply walking the area. Things marked on the map simply as "area of shakeholes", for example, appear in detail as individual shakeholes on an aerial photo. Whilst OS maps show broad-brush vegetation types, there is much more detail on an aerial photo, and things like sheep paths through heather or bracken may show up, which can be useful. More subtle changes in colour of vegetation can indicate underlying soil type, wetness or even the geology - the limestone shale boundary is not marked on topographic maps. Norman & Waltham suggest that, in their study, a limestone/shale contact shows up as a change in colour or tone in 30% of cases, a change of vegetation in 12% of cases, and a similar proportion as a change of texture. 36% of cases showed a step change in level, which can be very clear in the right lighting.

When you go to the `GetMapping.com` web site, you are required to enter a grid reference or postcode from which you get back a 500x500m map which you can scroll N/S/E/W in 167m increments. This is the 2m free dataset, which is presented as 83x83 pixel tiles, each covering an area of 167m square. The tiles are jpegs, compressed down to typically 1 to 1.5 kb, which is rather severe compression, so there are quite a lot of compression artefacts visible on relatively featureless areas like water. However, features like trees, buildings and shakeholes show up pretty well, and good footpaths are clearly visible. Some of the rural areas seem to have been mapped from a single flight, such that small areas are obscured by cloud, or have gaps in coverage. However, by mid-2002, pretty much everywhere I have looked in England and Wales seems to have pretty good cover, though hardly any of Scotland seems to be mapped.

The original air photos have been scanned, scaled and tweaked to be a fairly good and consistent map. In feature-rich areas, the geo-referencing seems to be pretty good, though in moorland areas, less care has been taken (less reference points are available), and in my own local area (Teesdale), for example, I can see the same patch of burnt heather in two adjacent tiles, suggesting something like a 20m inaccuracy in the geo-referencing of the two adjoining photos. To be fair, however, this doesn't seem to be very common and I suspect that for the most part the georeferencing is good to 5m. Feeding this mapping to a 1200 dpi colour printer at 125 pixels/cm and relying on the printer driver software to deal with the mismatch between pixel sizes in the image and on the paper produces quite respectable composite mapping at 1:25000, on which things like shakeholes can clearly be seen, and grid references read off.

The mapping is copyrighted, of course, and among various restrictions placed on your use of the data, republication on the web appears to be specifically prohibited. However, for personal study and other "fair use", the various copyright acts give you a fair bit of leeway.

The catch, of course, is that with tiles this size, you need 36 tiles to make up a kilometre square, and a lot of files need to be downloaded to make up a reasonable mapping area. Doing this by hand soon becomes tedious, but this is exactly the sort of thing that computers are good at doing, and a simple set of scripts can very quickly be lashed up to utilise common cross-platform tools like `wget` and `ImageMagick` to download the tiles and glue them together into rather bigger and more convenient units.

The tiles have URLs like

```
http://www2.getmapping.com/isapi/gettile.dll?
Dataset=2mFree&level=0&i=<i>&j=<j>
```

where the numbers <i> and <j> are derived from the full grid reference. By "full grid reference" I mean a grid reference which includes not just the two-digit numbers along the edge of your map, but also the small leading digit which can be seen at the corners of the map, and which are normally represented by using the grid letters instead. As a concrete example, take Gaping Gill, at SD 751727. The km grid square is SD 75 72, and looking at the corner of the map you can see that this translates into 375 472. The numbers <i> and <j> are obtained simply by multiplying by six, in this case giving i=2250, j=2832. That is the tile at the SW corner of the grid square, and, with six tiles per kilometre, the tile in the NE corner is i=2255, j=2837.

If you montage 36 83x83 pixel tiles into a kilometre square, the result is a 498x498 bitmap. I have found it convenient to put a one-pixel blue border round each such kilometre square, making it 500x500. Stacking these up together then gives a neat 500 pixels per kilometre, and a kilometre grid with lines two pixels thick, which looks about right.

What is needed is a small bit of code that will take a specification for an area you want mapped, then fetch the tiles, assemble them into kilometre tiles with a border, then block these up into either a bitmap for your whole area which you can print, or perhaps as a webpage with the kilometre squares displayed in a table, which is convenient for desktop viewing. Since the tiles are already a bit degraded by being jpegged, its probably best to avoid highly compressing the kilometre squares, causing further degradation. I have used png for the kilometre squares and larger blocks, as these render quickly, but they do use up a lot more disc space than the original tiles (300-500 Kb per kilometre square, as opposed to c 50k for the 36 tiles).

As with all programming tasks, a multitude of solutions are possible. I have hacked together a set of small programs which generate scripts to accomplish the task. This has a somewhat baroque architecture as a consequence of an interesting mixture of hardware on this site and a desire to be able to split the task across two machines so that the actual downloads can go on in background without impacting my desktop machine. However, the software hangs together on (a) Linux box(es), and could serve as an example for someone wanting to write something better. I'm reluctant to put the stuff up on a web page in case getmapping.com decided that this is misusing their resource, so if you would like a copy of the software, please email `mapping@pennine.demon.co.uk`.

Reference

Norman, J.W. & Waltham, A.C. (1969). Aerial Photography and the investigation of karst features, in Symposium on Cave Photography, *Transactions of the Cave Research Group of Great Britain*, Vol. 11(4), 245-254.