Erin's message got me thinking. Thinking about the assumptions that aven, survex, or the surveyors might be making regarding coordinate systems and so on. I got kind of into it and maybe took things too far but that's what I do (that my course is tedious beyond belief might go some way towards an explanation). Maybe this will be interesting to somebody. Bear in mind it's just a bunch of stuff I pulled from the web, combined with a few back-of-the-envelope calculations so there are undoubtedly mistakes....

Apology in advance for use of sexagismal notation - 3' is 3 arcmin, 4" is 4 arcsec.

Assumption: locations can be described by x,y,z coordinates with a Euclidean measure of distance ($d = \sqrt{x^2 + y^2 + z^2}$).

Violated because: General Relativity says that space will be curved in the presence of the earth's gravitational field so this is false, but the size of the effect is about 5mm if you were to survey all the way around the earth, so I think it's pretty safe to ignore this effect.

Violated because: the positions of locations do not stay constant in time. Reasons why not:
1) Tidal effects on the earth. The two sources of earth tides are the sun, causing effects of 17cm amplitude, and the moon, with 36cm (the next most influential body is venus, whose effects are less than 0.05% of the above). However, as long as our measurements cover an area of limited dimensions then the effects will be small. For example, over a 100km distance, the tides should come to less than a cm. This means we must work in a local coordinate system rather than a global one. (Alternatively we could use a global coordinate system and make allowances for tides but there are rather subtle and tricky issues to do with there being a significant (~30cm) zero-frequency component to the combined tide (i.e. you can't remove all the tidal effects merely by time-averaging);

2) Plate tectonics - causes relative shifts of up to 30cm/year. Again a solution is a local coordinate system. For example, for Europe, ETRS is based on the (global reference system) ITRS, but corrected for the overall average motion of the European landmass and is the system in terms of which the UK national grid is defined. This stops the landmass sliding across the map grid, at least in the bulk, though there will still be local shifts in tectonically active areas - the only thing you can do in such areas is to associate an epoch (time of measurement) with each observation and try to model the tectonic distortion (eg. there is the HTDP model for the USA). There are similar (locally stationary, defined in terms of ITRS) systems for North America (NAD83) and Australia and presumably other places too;

3) Other causes of time-variations are glacial rebound (up to 1cm/year), annual groundwater-pumping cycles (as much as 11cm/year in Los Angeles, much much smaller most other places), local subsidence due to mining works, etc. All probably small enough to be safely ignored (or very localized and hard to deal with in any general way).

Assumption: the directions "north" and "down" are constant.

Violated because: the earth is ellipsoidal. At the 10^{-5} level the earth is roughly an ellipsoid. The mean radius is about 6400km. The "down" directions for 2 points on the surface of the earth and separated by 100km differ by almost a degree.

Violated because: there are mountains, local deposits of lighter/heavier minerals, etc, which distort the gravitational field of the earth. So-called "deflection of the vertical" can be of the order of 1' in mountainous areas and can vary significantly over relatively short distances.

Violated because: The axis of the earth's rotation shifts over time. Most terrestrial reference systems are "conventional" in that they take some specified direction for the earth's rotation rather than the true instantaneous direction. However, the direction of north given by astronomical observations or by a gyrocompass, etc, will be derived from the instantaneous rotation. The difference is a small one, primarily due to the "chandler wobble" an incompletely understood effect which has an amplitude of about 0.7". (This seems negligible but in fact has to be taken into account when targeting inertially-guided star-tracking ICBMs or they'd miss by up to a few hundred metres. Certainly this effect is far too small to be relevant to cave surveying with any present or readily conceivable cave surveying technology - even gyrotheodolites (used occasionally in mines surveying) only generally seem to give 30" accuracy at best, from what I understand).

NOTE: because the direction "north" is actually a direction which is perpendicular to the direction "down" rather than just the direction parallel to the earth's axis (conventional or instantaneous) any errors in measuring either the gravity direction or the axis direction will be multiplied by tan(lattitude) when calculating "north". Ie a 1' deflection of the vertical at a place on the 70th parallel would translate to a north error of almost 3' (even assuming perfect instruments). If you happened to be surveying caves close to the poles, then of course directions relative to north would be extremely problematic.

Assumption: There are mathematical ways to convert position coordinates from one coordinate system to another. I'm going to break this into 3 issues:

1) grid projections. A grid projection is just a way of representing a (roughly, 10^{-5} level) ellipsoidal earth on a flat piece of paper. There are a number of different grid projections, with different properties for different purposes, and the equations for converting between them are frequently complicated, but they do exist, and are 100% precise. So as long as you know what the details of the projection are (easy if you have a GPS unit, just set it to lat-long or UTM or anything else that's well-known; potentially harder if you're dealing with foreign maps - sometimes the details are printed on the edge of the map, sometimes not) then you can convert from lat-long to UTM to Lambert conformal conic to anything else you wish - or indeed to (x,y,z) earth-centred earth-fixed (ECEF) coordinates, which is possibly the most useful format to a computer - and you can do this reversibly and without loss of accuracy (modulo machine-precision issues). So now we know that it is actually possible to deal with grid projections we should look at what the consequences would be if we do not (this is mostly what Erin's original email addressed). Obviously(?) all projections introduce distortion of some kind into the map. The distortions can be of various kinds and can be minimized in some areas at the expense of other areas - this is the reason for the menagerie of different projections in use. For a conformal projection (like UTM, for example) then the two types of distortion that come into it are convergence - that grid north is not the same direction as "true"

(see comments above) north for all points on the map; and scale error – that one "grid metre" is not the same as one metre "on the ground". To get an idea of these errors, for UTM, the convergence is zero along the central meridian of a zone and reaches up to ±3° at the edges of the zone. For UTM the scale error is zero on two lines spaced either side of the centre of the zone and reaches a maximum of 0.1%. Another possible distortion is a nonuniform stretching, resulting in lines that meet at one angle on the ground meeting at a different angle on the map. This type of distortion is not present in so-called "conformal" projections (this is the definition of "conformal") such as UTM (Erin's comment in her email notwithstanding!), but is relevant to other grids. Finally, with many projections, UTM included, the earth is split into several "zones" each with its own projection and a non-seamless stitchline. With UTM, although it is possible to line up maps from adjacent zones because the scale factor is the same each side of the zone boundary, there will be a 6° shift in convergence as you cross the boundary and the grid coords will jump suddenly too. One solution is to extend a zone slightly beyond its usual boundaries, with an accompanying increased amount of distortion the further you go outside the usual boundary. What to do about all these problems? Well, since you can convert from grid coordinates into Cartesian coordinates (so long as you know what grid was used) just do that and do the survey reduction in Cartesian space (its almost certainly what the survey programme wants anyway). What about displaying the processed data? On aven/caverot, whatever, IMO you want to carry on drawing a real isometric projection of the Cartesian data and the only problem you'll encounter is that there's no unique "north" or "up" directions to align things to (for drawing your compass rose, for when the user presses "p" to get plan orientation). But north and up won't change by much over any given survey so it should be fine just to use the north and up directions at the centre of the survey for example. On the other hand, if you're plotting a plan view to print on paper and overlay onto a topo map, then you'll need to use the same projection as is used on your map. If you just want to print it out on paper but you haven't got any particular map to overlay on, then just use your personal favourite projection and document it in the margin. So much for grid projections, now let’s look at the other issues...

2) Datum conversions. First, some terminology: the word datum has been misused so I won't use it. I'll make a distinction between a Terrestrial Reference System (TRS) and Terrestrial Reference Frame (TRF). A TRS is the set of conventions which define the system for making locations – stuff like "the origin of the system is the centre of mass of the earth including atmosphere" or "the origin of the system is an engraved mark on a brass plate in London" or "the z direction is the mean direction of the earth's axis for the years 1919 to 1986". The point is that you usually don't have an instrument that can directly tell you your coordinates in the TRS - you rely on benchmarks with known positions and locate yourself relative to a trig point or mountaintop or satellite or whatever. These benchmarks constitute a TRF and the crucial thing is that they are *measured* and can have errors. Two people can agree on a TRS but disagree on the location of a point because they use different TRFs realizing that same TRS. It is possible, in theory, to transform between TRSs - you need 7 parameters to do it (rotations, translations and a scale factor) but they can be arbitrarily complicated functions of time which probably need to be measured - eg. a valid, if bizarre, TRS could have as its origin the centre of mass of Lev Bishop: clearly, accurately translating between this and a TRS which has as its origin the centre of mass of the earth and atmosphere is going to be tricky. A case where the transformation is easy to do is when one TRS is *defined* in terms of another. For example, ETRS, as I mentioned earlier is defined in terms of ITRS but with a published transformation that keeps the european landmass on average stationary (in a plate tectonics sense). When we start to think about TRF transformations the situation gets much worse.

To do the transformation properly you need to take into account the distortions introduced by errors in the measurements which were used in realising the TRS in a TRF. Depending on the size of these errors, the distortion can be significant. For example OSGB36, the TRF used for the UK national grid, contains distortions of up to 5m. TRFs for more difficult or less-carefully surveyed parts of the world could be even worse, TRFs that were realized by satellite surveying should have smaller distortions. As in the case of TRSs, some TRFs are *defined* in terms of each other. For example the OSGB36 TRF which is these days defined in terms of the ETRF which is in turn defined in terms of ITRF. However, in order to keep all the distortions of the original definition of OSGB36 (so coordinates didn't jump when they changed the definition to be in terms of ETRF) the transformation has about a million parameters - this is the sort of problem we're dealing with here. So, basically, there's no general way to convert between TRFs beyond a certain level of accuracy, except by getting out on the ground and surveying relative to benchmarks for both TRFs - ie not converting at all, just measuring your positions in both frames simultaneously - but if you could do that then you wouldn't be looking for ways to convert between frames, would you?

3) Altitude conversions. Even though it would be possible to use height above the ellipsoid as a measure of altitude, that's not what's done (because the surface of lakes would not be constant, rivers might flow uphill, etc). Instead height above the geoid is what's usually used. (The geoid effectively being the level surface you'd get by continuing "mean sea level" into land areas). The RMS deviation between the WGS84 (GPS) ellipsoid and the geoid is about 42m, with a maximum of over 100m, so this is quite a large effect. However, the geoid has been modelled and for example there is the EGM96 geoid model, which is a spherical harmonic expansion of the earth's gravity field to order and degree 360. It has a RMS error of about 40cm but this is rather non-homogenous depending on the number of observations in a region and in remote or rugged areas the error can reach maybe 3m (of course these are the types of places cavers often go to). In addition, despite the huge number of spherical harmonic parameters, this model can only represent variations with a minimum wavelength of about 50km - anything more local than this is missed - obviously a problem in areas with local mineral deposits, or mountains, etc. Finally, the model is optimised for use at the earth's surface and so if you are doing deep caving there could be additional errors. One final complication with heights is that "mean sea level" does not actually describe a level surface - there are permanent ocean currents that give rise to "mean dynamic sea surface topography" of the order of 1-2m so any given country's height datum could deviate from the geoid height by maybe this much. All told I'd guess that in the worst-case situation (perhaps a very deep cave in the himalaya or something) you might be maybe 5m off even after using EGM96 corrections.

Assumptions: compasses read true north/you can correct magnetic north to true north/compasses read magnetic north. Well, presumably most surveyors know there is a difference between magnetic north and true north, but maybe not all of them realise that correcting for the deviation is more than just reading the correction off the corner of the map. For a start you really have to calibrate your compass because you don't know that it even reads magnetic north. Now the complications come because you do the calibration at a different place and different time to where you do your survey. The grid convergence varies from place to place; the magnetic deviation varies from place to place and time to time. For example, a worst case scenario - lets say you calibrate your compass by taking sightings to a mountaintop and reading the grid direction to the mountain from your map, but you're on the edge of a UTM zone and you cross into the adjacent zone to do your actual surveying - your 'calibrated' compass now reads 6° in error from grid north. Or if you're unlucky enough to go surveying during a big geomagnetic storm, you could have magnetic field fluctuations of 10's of degrees (this would be a big storm, but on a 'quiet' day, in a 'quiet' part of the world (away from the geomagnetic poles & equator), you still have magnetic field direction fluctuations as low as 2', and if there are days to months between calibration and surveying then 1° variation is typical (this is all for geomagnetic lattitudes less than 60°, more nearer the poles). This is why I'm notreally sure about Erin's comments about compass corrections based on grid convergence only. I know that Erin's been surveying in an area where short-term field fluctuations are relatively small (neither close to the geomagnetic poles nor to the geomagnetic equator) , and in in fact she's in an area where the magnetic deviation is very close to zero, but she still has to calibrate her compasses, and if she calibrates to grid north for the same part of the same grid that she's surveying in, then her compass will (more or less) read grid north already, without any further corrections. This will be the case, for example, for calibration by sighting to a nearby mountaintop or somesuch. I'm presuming Erin doesn't have a true north instrument to calibrate against (gyrocompass, star sightings, etc). Anyways, there are geomagnetic field models of the earth, such as the IGRF, a spherical harmonic expansion of the field out to n=10. But the IGRF only attempts to describe the long-wavelength variation of the field, won't account for variations in magnetic crustal rocks and minerals (typically half a degree but up to 10s of degrees in extreme situations, and can vary over as little as a few miles). The IGRF models variations on a timescale of years and therefore doesn't account for geomagnetic storms, etc. Perhaps you could get more precise corrections by having a magnetic field logging device left somewhere stationary on the surface near the cave, and recording the exact time of each compass reading (this would be easier with an automatic, data-logging electronic compass rather than a suunto). Such surface data-logging devices (generally proton magnetometers) are in fact manufactured for almost this purpose - use by geophysical mineral prospecting by magnetic field surveying (you leave one at base and take one surveying with you), although I expect they're probably quite expensive. Alternatively, there is a network of geomagnetic observatories around the world and if there is one sufficiently nearby (no I have no idea how close "sufficiently" is) you could download the data for any given time from the web (see, www.intermagnet.com) and maybe correct compass readings from that.

So that's all the factors I can think of that might produce systematic errors in a cave survey. I suppose now I should go through which ones are actually relevant and which ones are too small to deal with, and try to suggest ways of dealing with the important ones. First we need to decide how big an error we can tolerate for different purposes, because all the different error sources become important in different situations, so I'll just go through the different sources quickly and summarize the dependences.

First some effects that seem they will almost always be completely negligible:

General Relativistic Effects: negligible (~5mm for a survey around the globe).

Chandler wobble and other variations of the earth's axis: negligible (~1").

Groundwater pumping, glacial rebound, subsidence effects: negligible, or localized and impossible to deal with in a systematic way.

Now effects that relate to using fixed points. If we don't have any fixed points then these effects are irrelevant. I'll assume that we don't care about sources that cause less than a metre or so of error, as our lower accuracy limit. Its also true that many times our surface fixes themselves will be only about 10m accuracy (eg., if reading off a large scale map or using standard GPS technology) so in these cases we can be more lax:

1) Points from different coordinate systems - if all fixed points come from the same system (eg. all from GPS or all from surface surveys to known points in a given grid) then these effects will not be relevant:

a) failing to convert datums at all - definitely relevant - causes errors of up to several hundred metres.

b) using only a TRS conversion - borderline relevant - causes errors of around 5m, possibly more.

c) using a TRF conversion, where available - residual errors should be negligible.

d) using ellipsoid height versus geoid height - definitely relevant, errors of up to 100m. To fix this, use a geoid model such as EGM96 (I believe most handheld GPS units contain a (simplified) geoid model anyway, so not likely to be an issue unless you're doing more complex GPS work, in which case you probably know how to do these corrections).

e) Geoid model errors - generally negligible (EGM96 has RMS error of 42cm), but possibly starting to become relevant in extreme topography in areas of poor survey coverage (maybe as much as 5m).

f) Mean Sea Surface topography issues, ie issues of the "zero" or "mean sea level" - almost negligible (1-2m).

g) Fixed points in different UTM zones. Clearly relevant - you simply can't proceed without doing some kind of conversion. Solution: do the (exact, so long as you know the grid (easy if its UTM) and ellipsoid (not always so easy) parameters) conversion to ECEF cartesian coords.

2) Issues affecting all fixed points, the same coordinate system or not:

a) Tidal effects: definitely negligible for local TRFs, effectively negligible even for global TRFs (order of a metre);

b) Tectonic effects: could possibly be relevant for long-term projects. Drift rates of up to 30cm/year for coordinates expressed in a global TRF. If you have fixes in a global TRF then transform to a locally stationary TRF defined in terms of the global one (eg., ETRF for europe). At the very least, record the dates of the fixes so somebody has a hope of doing the transformations at a later date. If you're in a tectonically active area (eg. western USA) then there may be models of tectonic activity (and other sources of variations), such as the Horizontal Time Dependent Positioning (HTDP) for the USA, which includes tectonic models as well as models of effects in the vicinity of the largest few earthquakes in recent history.

Next there are error sources that affect direction measurements. For these it is harder to decide what size of error is acceptable, because most of them can be considered systematic errors. So for example if you are working to BCRA grade 5 angle accuracies of 1°, for a really big survey project with 100,000 legs the random errors will be dominated by systematic errors unless the latter are kept to below 10", a tough prospect indeed. Even for a more typical survey project with 1000 legs, the systematic errors will dominate over the random errors unless they are less than 2'. Another way to look at it is that for a large survey project with an overall linear extent of 10km, an error of 3' will introduce an overall error of 10m. By linear extent I mean the overall length eg. from the furthest north point to the furthest south, or the side of a box containing the survey, not the total surveyed length.

Errors resulting from not taking into account ellipsoidal shape of the earth: depends on lattitude and is proportional to the the square of the linear extent. For example at 70°N, a medium survey with 1km linear extent gets an error of 50cm, a larger one with 10km extent gets 50m, and a really big project with 100km extent gets 5km error. Alternatively you can think in terms of angles, in which case the 1km survey gets 30" error, the 10km gets 5' error, and the 100km gets almost a degree. Whether these errors are significant depends on the details of the situation, but in the case of the 100km it seems pretty obviously hard to claim they're negligible. I don't have any data on the linear extent of survey projects so I don't know what the maximum size that people deal with is. Easy to remove all these errors by treating the earth as ellipsoid rather than a plane. It doesn't even matter particularly which ellipsoid we use, even for a 100km linear extent - all the commonly-used ellipsoids are the same to the 10^{-3} level.

Deflection of the vertical: errors of maybe 3' (at lattitudes less than say 70°). Can correct for them either by looking at the slope of the geoid model (such as EGM96) or (better) using a special deflection of the vertical model, such as DEFLEC99 (which covers the USA), which differs from simply looking at the geoid slope by extrapolating out to the earth's surface and correcting for curvature of the plumb line (up to 3"/km), by using Digital Elevation Models, etc. Of course most caving isn't done at the geoid level or at the surface, so neither is necessarily right, but a deflection of the vertical model will have more short-wavelength detail so should be better. In bad cases I'd guess that deflection modelling will only remove maybe $\frac{2}{3}$ of the error, but generally it'd probably do much better than this and I can't see it making things worse.

Issues with magnetic compasses including calibration issues. Clearly these can be very significant, and in fact if these aren't dealt with carefully, then there really isn't any point in dealing with **any** of the other sources of angular error. Certainly it's necessary to calibrate compasses close to the surveying, both spacially and temporally, to have any hope of getting good accuracy. Survex could definitely help with some of this, by having specific support for different types of calibration (before and after the trip, grid north vs. true north, differing positions for calibration and surveying (different grid convergence, deviation), etc, etc). Maybe it would be profitable to allow exact times to be recorded for each reading (for the case of electronic instruments) and allow for corrections from mag field logging instruments or observatories. One can imagine the use of gyrocompasses instead of magnetic ones but realistically I can't see that happening anytime soon.

And now errors of length/scale:
Grid scale differing from real scale. Up to 1 part in 1000 for UTM. This is negligible only for small surveys. Eg, its equal to random error from measuring to the nearest 10cm for a survey with less than 100 legs, assuming 10m average leg length. Can be fixed by working in an ECEF coordinate system rather than using grid coords.

(One further source of error that I won't discuss because its always going to negligible for cave surveying is confusing great circle distances and straight line distances. Not important unless you take very long legs, like 10's of km).

I'll try to be brief as to what I think might be done to survex. Most everything can be implemented as a helper app without changing survex/aven much at all:

First, I think that since survex already accepts cartesian coordinates for fixed points, it would be easiest to implement grid projection, geoid modelling, TRS/TRF conversion, tectonic modelling, etc – everything relating to fixed points - outside of survex itself in some kind of helper wrapper application that takes lists of fixed points in some kind of tabular format (grid, TRS, TRF, coords, epoch, location ID, etc) and spits out something like a fixedpoints.svx file that can be *included into the actual survey files. The helper wrapper app should be highly configurable (maybe some sort of macro language?) and at the least should come pre-configured to do UTM and lat-long to ECEF cartesian (and other grid conversions should be easy to add), should know the parameters of the commonly-used ellipsoids (and the list should be editable), should be able to do 7-figure TRS conversions and know the parameters for the TRSs listed in NATO document ftp://164.214.2.65/pub/gig/datums/NATO_DT.pdf (and again this list should be editable) should be able to wrap around OSTN02 and OSGM02 the OSGB horizontal and vertical TRF converters and other similar things for other TRFs, should be able to wrap around EGM96 and HTDP and other specialized models, and ideally should be able to take a list of user-measured control points and construct some kind of rubbersheeting transformation for situations where the grid/TRF are unknown but the surveyor has located sufficient control points by GPS or other means to define his own transformation. All the transformations performed by this helper wrapper thing should be exactly invertible so that you can also use it to reproject survex.3d files into a suitable grid for producing map overlays. This helper app is not strictly required but it would make life much easier for the average cave surveyor who has no interest in geodesy. No changes would be required to survex or aven themselves, except that the up is no longer the +z direction but rather the "r-hat" ($\hat{\mathbf{r}}$) direction, and similarly for north. Actually, there is one potential problem that I can see: with the centre of the earth is taken as the origin all the coordinates will have large offsets from 0 and could there be numerical precision/rounding issues, maybe?

Second, I think that instrument calibration also ought to be implemented as another helper app, something again that takes a set of calibration data and spits out a "calib.svx" that can be *included into the survey file, as well as producing some statistics on the quality of the calibration, perhaps. If more fancy things are used by anyone, involving datalogging magnetometers or downloading magnetic observatory readings with which to correct compass bearings, then I feel safe in assuming that this will only be done in conjunction with an electronic datalogging compass and so the corrections can be carried out in conjunction with the downloading of the data from the compass and conversion into survex format - ie no need to make a general-purpose programme for this. Well perhaps, a programme that pulls the survey dates out of the files, then goes on the web and looks up the space weather and raises a warning if extreme geomagnetic storms were in effect at any time... that might not go amiss.

With these 2 helper apps and the slight change to survex and aven regarding the definition of "up" and other directions, all the errors I flagged as being potentially significant can be eliminated except for 2, namely errors arising from considering the earth planar, and from neglecting deviation of the vertical. The 2 are closely tied, in the sense that if you bother to take care of one you might as well deal with the other at the same time. Unfortunately, to do so would require fairly significant changes to survex, because the north and up directions would change depending on the station position and so an iterative approach to survey adjustment would be necessary (2 iterations should be plenty, however). Since both error sources are at the limit of being important (both cause errors of the order of 3' (for linear extent less than 10km, lattitudes less than 70°) my feeling is that there are more important things to be worrying about than this.

Final comment: My guess is that almost everything I've said so far is actually completely irrelevant because my feeling is that real cave surveys are actually dominated neither by random nor by systematic errors but rather by blunders/transcription errors. But maybe a few select people survey carefully enough, with backsights and other precautions to remove blunders, resurvey their loops when the closure error is too high, to make this message of more than purely academic interest.

See further comments interspersed below.....

On Tue, 5 Nov 2002, Erin M. Lynch wrote:

> On Tue, 5 Nov 2002, Olly Betts wrote:
>
> > On Tue, Nov 05, 2002 at 02:54:32AM -0800, Erin Lynch wrote:
> > > In the long run, it'd be dead handy if survex incorporated the maths so
> > > that users can just enter the coordinates they get out of their GPS. But
> > > it's worth having a think about whether it would be better to
> > > change the coordinate system of the fixed points to an arbitrary local
> > > grid where north is True North, or to leave the fixed points in the
> > > UTM/UPS grid and rotate/stretch the cave survey so it is in UTM/UPS
> > > coordinates (which would be useful for map overlays and surface
> > > prospecting).

IMO: You want to convert fixed points to ECEF XYZ coords for processing
and display. You want to convert everything back to some suitable grid for
printing plan views for map overlay. I guess a straight isometric
projection is fine for elevations.

> > I think this is the hardest part actually. The maths required isn't
> > hard to find - there's probably even GPL-ed C code suitable for
> > assimilation. But you need to decide what grids users will want to be
> > able to transform it into, and how to specify this clearly and
> > succinctly.

Transforming grids should be easy to do and there probably is GPL-ed code
out there for many, although there are quite a lot of grids in existence.
Transforming TRSs is also very easy to do (a simple matrix equation) but
you have to know the parameters in question (lots of them, not all easy to
find, though many are listed in NATO documentation available online).
Transforming TRFs (what you really want to do) cannot be done in general.
One way to go about it would be to GPS the locations of a few fixed points
(mountain tops etc) that you can read off the map, and then define your
own rotate, translate, scale, rubber-sheet distort, etc, transformation
between map coords and GPS.

> For the time being, I'd settle for having a way to record the zone, datum
> (and maybe EPE?) of a fixed point, so when the mathy bits are implemented
> I don't have to go through and change all of my old data files.

I agree it's a good idea to standardize some way of specifying TRF, grid,
zone, time of measurement, vertical datum, error ellipsoid, (anything else

relevant that I haven't mentioned?)

> > For your purposes, simply converting coordinates in one zone into the
> > extension of coordinates for the neighbouring zone is probably good
> > enough.
>
> No, unfortunately it's not :(
>
> > Grid north and true north for a UTM zone are presumably only exactly
> > aligned at the centre of each zone. What's the divergence of the two
> > at the edge of a zone? If it's major, then we need to think about
> > compass corrections for grid north...

3 deg at the edge of the zone for UTM. Don't forget there is a scale error
too. Best not to work in grid coords for best accuracy.

> The compass correction varies slightly depending on the direction of your
> bearing,

No, it doesn't! Not for UTM at least, which is a conformal projection.

> but to give you an idea of the size of the error, I've calculated
> the correction for a northwards bearings in a few caving areas in China.
> Houping (on the edge of the 48R Zone) is the worst at 1.5 degrees. Nandan
> is 1.1, Tianxing -1, Youyang -1, Leye 0.6 etc. The problems this can
> cause are best illustrated in Houping, where a 700m surface survey
> produced a 26m "loop closure error" between two GPS points if we
> assumed the GPS points and survey were in the same grid. Once the
> GPS points were converted to the grid of the survey the error was only 7m.

I want to know more about the compass calibration before I accept that the
"corrected" survey is actually any better than the uncorrected one.

Phew! Sorry that was so badly organised, but I couldn't face working it
into a more logical order. You can go back to ignoring me now.

Lev