© Commonwealth of Australia 1997

Processing of airborne magnetic data

A.P.J. Luyendyk[1]

The processing of aeromagnetic data collected along survey flight lines to a grid of values ready for the application of enhancement techniques and interpretation involves the sequential processes of editing, correction for diurnal effects, removal of the Earth's background magnetic field, the levelling of all data to a common base and, finally, the application of a gridding routine.

Introduction

The aim of this document is to give an overview of the processing of airborne magnetic data, referring to examples of methodologies used by the Australian Geological Survey Organisation (AGSO).

While the general procedures are well defined, it is still not practicable to make all the measurements needed to unambiguously process airborne data. A skilled analyst is still needed to make judgements in the levelling and micro-levelling processing procedures, which can have a significant effect on the overall result. The current industry aim is to level magnetic data to better than 1 nT. Whether or not this has been achieved is often difficult to assess. The measure of quality usually used is that enhanced images of levelled data should reveal a minimum of artefacts attributable to the data gathering or reduction processes. It is still not possible to guarantee that an image without appreciable artefacts can be produced without seriously affecting the integrity of the data. The most likely sources of errors are discussed.

Processing of aeromagnetic survey data

The overall processing of aeromagnetic data involves eight major steps in two phases:

Phase 1—Pre-processing

- verifying and editing the raw data;
- locating the data in x and y.

Phase 2—Processing

- parallax corrections;
- removing diurnals;
- removing the component attributable to the Earth's regional field;
- levelling the data;
- micro-levelling—removal of any residual levelling errors;
- gridding and contouring.

A variety of software packages is available for processing and presenting aeromagnetic data. They vary in the features and methodologies they provide, their style of user interface, which ranges from command line through to menu driven, fully fledged graphical user interfaces, and the level of user feedback. User feedback is often ignored in software assessment, but it can be critical in the production of high-quality processed magnetic data, as some of the major processing phases, such as levelling, rely heavily on the skill of the operator, whose job is made much easier with interactive visual feedback.

No attempt is made here to assess or provide an exhaustive list of available software packages. Some of those available and their developers are listed below.

- ARGUS — Australian Geological Survey Organisation.
- ECS — Software Engineering Computer Services, Bowral, New South Wales.
- GIPSI — Paterson, Grant & Watson, Ontario, Canada.
- TerraTools — TerraSense Inc., California, USA.
- Intrepid — Desmond Fitzgerald & Associates, Brighton, Victoria.

The algorithms and techniques described here have been implemented in a software package called Intrepid, jointly developed by Desmond Fitzgerald & Associates, AGSO and a consortium of Australian mining companies, including BHP Minerals, Pasminco and Stockdale (Fitzgerald & Associates 1996).

The reference level

For aeromagnetics, since we are only interested in magnetic perturbations due to the Earth's crust, the ideal data set is a 'snap-shot' of the magnetic field at all required locations at the same instant of time, but with the Earth's regional magnetic field removed. This ideal data set is the reference surface to which the data are to be reduced.

Sources of errors

Magnetometers

Modern magnetometers give an absolute measurement of high sensitivity, with virtually no drift and for all intents and purposes can be regarded as giving an exact reading. Typical 'noise envelopes' are ± 0.1 nT.

The same cannot be said of the fluxgate magnetometers used as little as 10 years ago. They were not absolute and had to be manually calibrated, had high rates of drift and were sensitive to about $\pm$ 1 nT. Drift curves varied exponentially with time and, even allowing for standard operating procedures, such as turning them on well before the commencement of a day's flying, drift rates up to 10 nT per hour were common.

Aircraft effects

The magnetic signature of the aircraft consists of three components. Its permanent magnetisation, magnetisation induced by the motion of the aircraft through the Earth's magnetic field, and a component due to the flow of electrical currents within the aircraft. These currents fluctuate as the pilot varies any aircraft controls, such as changing the engine speed, adjusting ailerons, etc.

The permanent magnetisation of the aircraft leads to the familiar *heading error* caused by the vector addition of the Earth's field with that of the aircraft, resulting in a base-level shift between lines flown in opposite directions. Higher frequency errors are introduced by aircraft manoeuvring and are referred to as *manoeuvre noise*. The general procedure for removing these effects is called *compensation*.

Modern survey aircraft are compensated in real time by sophisticated modelling packages, which have the aircraft orientation as their input. The normal procedure is to establish the various model parameters by flying the aircraft around a square at high altitude, over a region of low magnetic relief, while executing several roll, pitch and yaw manoeuvres. These are then used to eliminate any induced magnetic effects. There is no direct control on the flow of electrical currents, but the aircraft is always operated in a 'steady state' to minimise any effects. Routine tests are carried out during the survey to

[1] Australian Geological Survey Organisation, GPO Box 378, Canberra, ACT 2601

ensure the compensation has not deteriorated. With these procedures heading error is reduced to ± 0.25 nT and manoeuvre noise to ±0.15 nT.

Older surveys were not flown with such sophisticated techniques. The 'compensator' consisted of an array of three orthogonal electrical coils acting as magnets. By varying the current through each coil the permanent magnetisation of the aircraft could be effectively neutralised. The standard procedure involved rotating the aircraft on the ground over a fixed spot and varying the currents until a constant reading was obtained. Time-varying effects were not compensated for, and the procedure was prone to error, as it was basically a manual operation. Heading errors of ± 5.0 nT were common.

Navigational effects

The advent of global positioning systems (GPS) in the late eighties has greatly improved the quality of navigational data. Positional accuracies of about ± 5 m are now routine (Boyd 1992). When comparing data between lines, and especially line/tie intersection points, the effect of navigational errors is obviously related to the gradient of the local field and will have a varying effect. It should be remembered that the GPS data include X, Y and Z data. The Z values, the height above the ellipsoid, can be put to good use when processing aeromagnetic data, specifically, when transforming the positional data between datums, in the calculation of the Geomagnetic Reference Field (GRF), and in combination with ground clearance data to produce digital elevation models.

In most systems the primary navigation data are adjusted to reflect the location of the most position-sensitive recording instrument—invariably the magnetometer. If this adjustment is not made, parallax errors result with readings from adjacent lines being offset by twice the distance between the magnetometer and the actual navigation reference point along the axis of the aircraft. The distance between the navigation and magnetometer recording instruments is referred to as the *cable length*.

Before the advent of GPS, radio-beacons were often used for navigation. More commonly, the location of the aircraft was determined by comparing strip film or video recordings made of the ground during flight with aerial photography and topographic maps. In the latter case, an optimistic accuracy of the order of ± 50 m was claimed for each of these *fixes*. Three to four fixes were determined at the start and end of each line and then at regular intervals along it, usually chosen to coincide with tie crossover points. Intermediate points were infilled with additional navigational aids. At AGSO, Doppler infill was used. Where such facilities were not available, extra fixes were needed and the flight path was approximated by straight line segments between neighbouring fixes. Variations in ellipsoidal height were ignored, the mean height of the survey region above sea level being used whenever such values were required.

In regions with few distinguishable topographic features, such as deserts, errors of the order of 500 m were often detected during processing. Clearly, in many cases the navigation was not very precise and was considered the greatest source of levelling errors before the introduction of GPS.

Time variation in the magnetic field

The Earth's magnetic field varies with time (Table 1). The variations can be random or cyclic, varying from the 11 year cycle of sunspot activity down to geomagnetic pulsations with periods of the order of seconds (Parkinson 1983). Some relate to local time, for example diurnals, but others, such as magnetic storms, relate to universal time and can be considered synchronised to less than a minute worldwide. Not only can these variations be out of phase, but their amplitude can also vary significantly with position. Typically, for a 1:250 000 map sheet area, the average variation peak to peak along a line is of the order 3 nT, and the total variation over the whole survey is about 30 nT.

Table 1. Time variations in the Earth's magnetic field.

Class	*Description*
Pulsations	Period 1–300 s; magnitude generally below 10 nT with the amplitude of pulsations decreasing with frequency; occur at random.
Magnetic storms	Sudden onset with decay lasting from hours to weeks; magnitude up to tens of nanoteslas; caused by solar flares leading to a rapid change in the flux of cosmic rays interacting with the atmosphere.
Diurnal	Period of 24 hours; magnitude up to 50 nT; due to the rotation of the Earth relative to the sun.
Lunar	Period of 27 days; due to the orbit of the moon about the sun.
Solar	Period of 1 year; due to the orbit of the Earth about the sun.
Secular	Gradual change with a variable rate across Australia within the range of –20 to 30 nT per year.

Ground clearance variation

The amplitude of local magnetic anomalies varies with distance from the recording instrument, i.e. with respect to the ground clearance of the aircraft. The rate of change increases as the wavelength of the anomaly decreases. This is most noticeable where the plane flies across an escarpment, since it needs to climb gradually as it approaches it from one direction, but may descend much more rapidly in the other. A typical herringbone pattern can be observed in a grid of the magnetic data, attributable to the variations in ground clearance between adjacent lines. For example, AGSO's Aero Commander has a survey speed of 70 m/s and a climb rate at this speed of 50 m/km. To clear a 50 m obstacle climbing must commence about 20 s beforehand to give a margin for safety. Ground clearance variation can be expected to cause errors of tens of nanoteslas over a typical survey area.

Altitude variation

The Earth's magnetic field varies with height above the ellipsoid. Typically, the rate of change with height is 0.025 nT/m. Since rapid changes in altitude are usually associated with changes in ground clearance, these effects are usually masked by ground clearance variations.

Wave noise

Over large bodies of water, surface waves can produce detectable variation in the magnetic field. Waves represent a conductor moving through the Earth's magnetic field resulting in a secondary induced magnetic field (Weaver 1965; Ochadlick 1989). Wave noise from ocean swells of 1.5 nT has been detected by AGSO's survey aircraft at a flying height of 80 m.

Pre-processing

Verifying and editing raw data

Raw data must be visually inspected for spikes, gaps, instrument noise or any other irregularities in the data. This is most easily performed via an interactive editor where the data are displayed as a continuous trace as a function of time. Automatic procedures to detect spikes are well documented, including fourth difference analysis and Naudy filtering (Naudy & Dreyer 1968) to extract any short-wavelength noise.

As a rule of thumb, the depth to the source of a narrow magnetic anomaly is given by the width of the anomaly at half its amplitude. With this in mind, any data with a wavelength less than that of the flying height of the survey aircraft can be regarded as noise and removed via low-pass filtering. Such filtering is only applied in limited cases, such as when high-frequency wave noise is apparent.

Ideally, this verification phase should be carried out in the field so that any errors attributable to data acquisition can be corrected in the survey aircraft as quickly as possible.

Locating data in X and Y

Once the navigation data have been checked and edited they can be merged with the magnetic data in order to locate the magnetic data.

Processing

Parallax correction

The distance between the navigation reference point along the axis of the aircraft and the magnetometer recording instrument is referred to as the *cable length.* If the cable length is not zero a parallax correction must be applied to synchronise the magnetic and navigational data. Since the navigation data vary smoothly and are not the primary data set, they should be interpolated to coincide with the location of the recording instrument rather than interpolating the data to coincide with the navigation. A parallax correction is applied by calculating the velocity of the aircraft at each x,y point from the navigation data. This velocity is then used to adjust the navigation data to allow for the cable length.

Diurnal variation corrections. It is common practice to collectively refer to all the time variations of the Earth's magnetic field by the misnomer *diurnals.* This practice is continued here. The diurnal variation is monitored by a base station at a fixed location on the ground. It serves two purposes.

- It is used to monitor the short-term rate of change of the field. Flying should be curtailed if the rate of change exceeds a specified cut-off, typically 2 nT per minute. The main purpose is thus to identify periods of magnetic storms, but the readings can also indicate periods of high pulsation activity.
- The base station data are time-synchronised with the aircraft data, and can be subtracted to give a residual which is a function of position only. This assumes that base station variations are fully representative of temporal variations over the whole survey area. This is not strictly true. Significant variations in phase and amplitude are known to occur over distances of 50 km or more. There is a risk of introducing high-frequency errors that cannot be removed by subsequent levelling procedures.

Applying diurnal corrections introduces an arbitrary base-level shift into the magnetic data. In general, this is of no consequence as we are only interested in local anomalies and not the absolute value of the magnetic field. A constant representative of the mean magnetic field value over the survey area can be added to the survey data. Moving the base station or the use of more than one base station will lead to different base-level shifts which are automatically removed by the standard levelling procedures, described below. In addition, the noise envelope of the base station magnetometer is added to that of the aircraft, but this is insignificant for modern magnetometers in comparison to other errors.

If base station data are not subtracted, diurnals with a period less than twice the flying time between ties cannot be removed by subsequent levelling procedures. A typical survey aircraft speed would be 70 m/s, giving a flying time of 70 s between ties separated by 5 km. The base station magnetometer should obviously be sampled at a rate high enough to effectively monitor the temporal variations—typically 1 s or less—with interpolation being used for intermediate values. (Note that low sampling rates would render the high-frequency micropulsations invisible.)

The debate is continuing on the usefulness of subtracting diurnals (Barton & Johnson 1988). At AGSO it is standard practice to apply this correction after filtering the diurnals to remove all wavelengths with a period of less than 10 s. Experiments are currently being conducted, using a high-precision base station magnetometer, sampling at 0.1 s and accurately synchronised to Universal Time, to try to resolve the problem.

GRF—removing the Earth's regional magnetic field

This processing step merely involves subtracting a well-defined model of the Earth's regional field, the so-called **G**eomagnetic **R**eference **F**ield (GRF), from the data. Standard models of the Earth's regional field are based on satellite and ground observations. They provide estimates of the field as a function of position, including height above the ellipsoid, and time. The secular variation is very slow and is ignored with a mean date for the survey being used to calculate the GRF.

GPS navigation provides the (X,Y,Z) data that should be used to calculate the GRF as a function of position. It is standard practice to replace the Z value with the mean height of the aircraft above sea level for a particular survey. The errors introduced by this approximation are small, as the rate of change of the Earth's magnetic field with height is of the order of 0.025 nT/m. These errors can be eliminated altogether by using the GPS Z data when available.

The field varies slowly with position, the change being insignificant over the sample interval of the magnetometer. To speed up calculation of the GRF, it is normal practice within the industry to calculate values at an interval within the range 100–500 m, with intermediate values being interpolated from their neighbours. As this operation only has to be applied once, the additional overheads to calculate the field for every point are insignificant.

For the Australian continent there is a choice of two GRF models, the International Geomagnetic Reference Field (IGRF) and the Australian Geomagnetic Reference Field (AGRF) (IAGA 1991; Barton 1988). Both are satisfactory, although the AGRF, which applies to the Australian region only, does reflect the long-wavelength field over Australia more accurately. For our purposes, with survey areas of only 100–200 km, the choice is academic. As applicable AGRF models were not available until 1985, and the IGRF is widely recognised, AGSO has always used the IGRF. It should be emphasised that the reference model and input parameters used should be recorded so that in the future it is possible to adjust the data to another GRF model.

Levelling

Aeromagnetic surveys are flown in a particular pattern, or grid, designed to give duplicate measurements at so called *crossover points.* The magnetic reference to which the magnetic data are levelled is time invariant and, therefore, any discrepancy at a crossover point represents an error. A typical flight path consists of *lines,* which give the primary coverage for the survey, and *ties,* which are flown at right angles to the lines and are used for control. The intersections of the ties with the lines give the crossover points.

Levelling is the procedure by which the discrepancies between the readings at each crossover point, the *intersection errors,* are reduced by systematically proportioning them between the ties and the lines. Several methods are in common use (Green 1983; Yarger et al. 1978; Foster et al. 1970). They can be grouped into three categories.

Minimising intersection errors by adjusting navigation data. If an estimate of the accuracy of the navigation data is known, then this can be used to define a circle about each intersection point in which the intersection point can be freely moved. Since the gradient of the magnetic field is known at an intersection point, its true position within this circle is assumed to be the one that minimises the intersection error. This procedure is repeated at each crossover point, giving a set of corrections to the navigation data which can be interpolated to all the navigation data.

This type of methodology can be applied with various levels of sophistication, generally aimed at determining that proportion of the intersection error to be attributed to errors in the navigation. Thus Green (1983) considered the eight immediate neighbours of an intersection point and deemed the error attributable to the navigation to be that which minimises the closure errors of the loops formed with these neighbouring points.

The main objection to adjusting intersection points by such methods is that it treats each intersection point in isolation, or only considers some effect from immediate neighbours, and will always reduce the errors at every intersection point. In general it is only used in combination with other levelling procedures.

Loop closure methods. The grid of intersection points of lines and ties can be considered a network of closed loops and standard procedures for network adjustment in geodetic surveying can be applied. For a more detailed analysis see Green (1983). Though mathematically attractive, in that an immediate solution is available with no manual intervention, it has been found to work well only in regions of low magnetic gradient where the range of intersection errors is small. Because of this, it is standard practice to apply navigational adjustment, as described above, to minimise intersection errors before applying this method.

Polynomial Levelling. In its simplest form, this method involves fitting a polynomial to the intersection errors along a flight, line or tie as a function of elapsed time, by the method of least squares. These polynomials are then subtracted from the original data, reducing the intersection errors. Various implementations of this method have been described (Yarger et al. 1978; Foster et al. 1970). The main disadvantages of the polynomial levelling method are:

- Problems can occur at the end of lines where the polynomial may not be well controlled and it diverges from what would be regarded as an acceptable fit to the data.
- A skilled operator is required to get the best results, it being necessary to judiciously select the degree of the polynomials fitted to the data. To this end, it is essential that interactive visual feedback be provided so that the operator can rapidly check the fitted polynomials and adjust them accordingly.
- When the intersection errors are irregular the simple polynomial fits may not follow the variations well enough. In such instances a series of localised polynomials is required to get good results.

Since this type of levelling is the one chosen by AGSO, the implementation devised by AGSO is described in some detail in Appendix 1.

Micro-levelling

Micro-levelling is a general term that refers to the removal of any apparent residual errors in airborne geophysical data after standard processing and the application of the more rigorous levelling techniques described above. Micro-levelling adjustments are necessary because quite minor data errors become clearly visible when grids of data are displayed as enhanced images. Not only does this make the image unattractive, but subtle features may be masked.

Geophysical exploration companies regard their own micro-levelling techniques as proprietary and little information is available on particular processes. In general, automated methods involve filtering a grid of data to detect residual errors. These errors are then subtracted from the original point-located data. Data are also manually edited, but this can be very time consuming. It involves correlating errors from images or grids of data with flight-path information to estimate the magnitude of any observable errors. These errors are then subtracted from the original point-located data.

AGSO currently uses a grid filtering technique for micro-levelling, the details of which are described in Appendix 2.

Gridding and contouring

Gridding point-located data is the first step in AGSO's micro-levelling method and provides the best method of data quality control via their display as enhanced images. It is essential that a grid honours the original point-located data and also provides a smooth continuous surface.

The gridding technique used by AGSO is an algorithm developed by Briggs (1974), as implemented by the software package Intrepid. It takes the randomly distributed survey data and interpolates it onto a regular grid. The method represents the surface locally with splines, whose total curvature is minimised. If a grid point has an original observation falling within half a grid cell size of it, the surface is further constrained to pass through that point. A direct solution is not available and the final solution is obtained iteratively from initial values determined for each grid point by interpolating from the three closest original observations.

Though producing an excellent result, the technique converges very slowly to the final solution, the rate of convergence for a given grid cell decreasing the further it is from an original data point. To obtain a visually acceptable grid, this effectively restricts the cell size of the grid to a minimum of about one-fifth of the line separation.

Because the data sampling interval along the flight-line direction is typically hundreds of times greater than the line spacing, aliasing problems are often present in gridded aeromagnetic data. To alleviate the problem, the line data should be low pass filtered so the frequency content of the data in the flight-line direction is comparable to that perpendicular to the flight-line. Likewise, typical grid cell sizes are an order of magnitude greater than the data sampling interval along the line. Data with a wavelength less than twice the grid cell size can only be regarded as noise and, again, should be removed before being sampled for input to a grid. In general, because the end users of gridded data want the highest frequency content possible, line data are not filtered before being gridded. Once a grid has been produced it can be displayed as an image or a contour map.

Conclusions

The methodology and rationale for the processing of aeromagnetic data has been reviewed, with the techniques used by AGSO for levelling and micro-levelling being given in some detail in Appendixes 1 and 2.

The data cannot, as yet, be unambiguously processed, the results especially depending on the levelling and micro-levelling techniques employed and the judgement of the analyst using them. Research is still needed to develop non-invasive, effective micro-levelling techniques, as the current options remove a significant proportion of the true signal.

Diurnal removal is another contentious issue. On the one hand, tie-line spacings generally used in surveying are insufficient to allow high-frequency diurnal variations to be removed, while, on the other, current techniques used to monitor the Earth's field, using base stations, are inadequate.

Ground clearance variations are generally ignored and a standard methodology for their removal needs to be developed.

Appendix 1. AGSO polynomial levelling

Aeromagnetic surveys are flown in a particular pattern, or grid, designed to give duplicate measurements at so-called *crossover points*. The magnetic reference to which the magnetic data are levelled is time invariant and, therefore, any discrepancy at a crossover point represents an error. A typical flight-path pattern consists of *lines,* which give the primary coverage for the survey, and *ties*, which are flown at right angles to the lines and used for control. The intersections of the ties with the lines are referred to as crossover points, and the difference between the magnetic measurement made along the tie and that along the line at the crossover point is known as the *intersection error.*

The basic assumption underlying the method is that the errors in the magnetic readings associated with each line or tie vary slowly and can be well approximated by a set of local polynomials as a function of time. These polynomials are referred to as *drift curves* and the task is to reconstruct them from the observed intersection errors. The intersection error at a given crossover point will be made up of two components, one from the line and one from the tie. The method attempts to separate the errors into those due to the lines and those due to the ties. Any errors due to navigation are regarded as noise.

Note that flights typically consist of several lines. Lines recorded in the same flight can normally be related to each other by time and, therefore, drift curves that apply to a flight can also be considered.

There are four basic procedures. These are listed below in the sequence they should be applied.

- Level the ties: calculating the tie drift curves.
- Drift the lines by flight to the ties: calculating the drift curves for individual flights of lines.
- Drift the lines individually to the ties: calculating drift curves for individual lines.
- Drift the ties individually to the lines: calculating residual errors for individual ties.

Principal tie

The first problem is that an absolute reference is required. No such reference is available and an arbitrary tie from the survey is chosen as the absolute reference, the so called *principal* tie, which is assumed to have zero drift. Obviously the principal tie should be chosen carefully. To enhance the reliability of the principal tie, it should be one that was flown during a period of quiet diurnal activity, be located over a region of low magnetic relief, and be located approximately at the centre of the survey area.

Levelling ties

Here, the task is to calculate the drift curves for each tie, effectively levelling them to each other. Firstly, a tie is selected to be the principal tie. It acts as the absolute reference to which the whole survey is levelled. Ties are classified as either *levelled* or *unlevelled.* In the first instance, the levelled set of ties consists of the principal tie only. Each tie in the survey is levelled in turn by adjusting it to the current set of levelled ties.

At this stage, each intersection error involving the unlevelled ties consists of a component from the tie and one from the line at the corresponding crossover point. The first task is to estimate the drift, or error in the magnetic reading, attributable to a line at its crossover with the tie being levelled. Once it is known, it can be subtracted from the corresponding intersection error, giving the error at the crossover point due to the tie alone.

The estimate of the drift of a line is determined by inspecting its intersection errors with the current set of levelled ties. Since the drift for the levelled ties is assumed to be zero, the intersection errors between them and a line consist of the component due to the line alone. In actual practice, a better estimate of the drift curve of a line can be obtained by considering the whole flight containing the line. Since there are more crossover points along a flight compared to those for a single line, a higher degree polynomial can be fitted to the intersection errors, giving a better estimate of the drift curve.

The actual levelling of an unlevelled tie proceeds as follows (note that all the other unlevelled ties are ignored at this stage). Each flight of lines is considered in turn. A polynomial is fitted to its intersection errors with the set of levelled ties, as a function of time along the flight, giving an estimate of the drift for the flight. (In the case of the first tie being levelled, the only control is where a particular flight crosses the principal tie. As the procedure progresses, the pool of levelled ties becomes larger, improving the control for the estimate of a flight's drift curve.)

The tie to be levelled is now inspected. At each of its crossover points with the current flight, the component of the error due to the line at the crossover is calculated from the estimate of the drift curve for the flight. This is a simple matter, since the flight time at the crossover point along the line is known. Subtracting it from the current intersection error gives an estimate of that part of the intersection error attributable to only the unlevelled tie at that crossover point. These components of the intersection errors due to the unlevelled tie are saved.

Once all the flights have been processed, a complete set exists of estimates of the component of the intersection errors attributable to the unlevelled tie, to which a polynomial is fitted as a function of time along the tie. This polynomial gives the required drift curve for the unlevelled tie. The drift curve is used to update the magnetic value at each crossover point along the unlevelled tie, which is then considered to be levelled and added to the set of levelled ties. The procedure is then repeated for the next tie to be levelled, ultimately giving a complete set of levelled ties.

Drifting lines by flight

Once the ties are considered levelled to each other, their drift is assumed to be zero and any remaining discrepancies at the crossover points are attributed to errors along the lines alone. Thus, the recalculated intersection errors can be used directly to give an estimate of the drift that occurred during each flight.

Lines are processed in flights at this stage to give the widest continuous time base available. Longer period trends in the intersection errors can, therefore, be detected than if the lines were processed individually.

Each flight is considered in turn and its drift curve estimated by fitting a polynomial to its intersection errors. This estimated drift curve is then used to update the magnetic value at each crossover point for each line in the flight.

Drifting individual lines

The procedure is identical to that for drifting lines by flight, except the individual drift curves are calculated for each line and used to update the corresponding magnetic values. This will remove effects that are line specific, such as heading error.

Drifting individual ties

Once the ties have been levelled and the lines drifted by both *flight* and *line*, the survey is essentially levelled. Any residual errors in the ties can be removed by drifting them to the lines. Several iterations of drifting lines and drifting ties can be performed, effectively 'shaking down' the survey, but have little if any effect after the first iteration.

Applying corrections to the full data set

At this stage only the magnetic values at each crossover point have been adjusted. It is a simple matter to calculate the corrections applied by subtracting the original values from the levelled ones. These corrections are interpolated along each line or tie by the local polynomial procedure described by Akima (1970) and subtracted from the original magnetic values, producing a levelled survey.

Polynomial fitting

All drift curves are approximated by polynomials fitted to the underlying data by the method of least squares. The success of the AGSO polynomial levelling method is critically dependent on the quality of these fits. Polynomials are used since the data being fitted can be well approximated by polynomials as a function of time over a localised region.

Least squares techniques are used to fit polynomials to the data to allow for random noise. That is, the intersection errors will always have a noise component due to navigational errors and any temporal variations with a wavelength less than twice the spacing between the ties and lines. The least squares technique minimises the discrepancies between the function being fitted and the actual data, but it does not require that the data be exactly honoured. In addition, data points can be weighted in proportion to their gradient, thereby recognising the fact that errors in navigation will have greater effect in areas of high gradient.

There are several problems associated with fitting polynomials to noisy data by the least squares technique.

- The degree of the polynomial should be small in comparison to the number of data points being fitted.
- The fit may be poor at the ends of lines/ties, since there is no constraint on the polynomial beyond the limits of the data and they may diverge from the expected result. The problem becomes more severe as the degree of the polynomial is increased.
- Data points should be evenly spread along the data interval being fitted, to ensure that the polynomial is well behaved between the data points. Irregular groupings of data points can lead to the polynomial diverging markedly from the 'expected' result between the groups of data points.
- Data points that are in error may have an undue influence on the result of the fit.

To force some control over the first problem, restrictions are automatically applied on the maximum degree of the fitted polynomial in comparison to the number of data points. For example, at least 3 points are required to fit a degree 1 polynomial, 5 for a degree 2, and so on.

The fourth problem is minimised by fitting the polynomials in a two-step process. An initial polynomial is fitted and 'rogue' points lying outside a few standard deviations of the resultant fit are rejected. The final polynomial is then calculated from the remaining points.

It is necessary to use high-order polynomials to fit the data well, but this exacerbates the problems associated with their use. To minimise these problems the concept of *piecewise* polynomials was introduced in an attempt to limit the degree of the polynomial, but at the same time achieving a close fit to the data. Instead of fitting a single polynomial over the entire data set to obtain an interpolated value at point N, a subset of the M closest points to point N is used. M is called the *window width.* A single polynomial is fitted to the M points, from which the interpolated value for point N is calculated. To obtain an interpolated value at point N+1 the window is moved along by 1 point and the procedure repeated. By reducing the window width, the data can be honoured as closely as required since, in the degenerate case of M=1, the data will be honoured exactly. Conversely, if the window width is chosen to be the entire data set, normal polynomial fitting results.

One disadvantage of the use of piecewise polynomials is that the resultant set of interpolated points will not necessarily vary as smoothly as required. This results because the polynomials fitted to successive subsets may be very different, even though successive subsets differ by only two points. Consequently, the interpolated points are filtered to smooth out this high frequency 'noise'. Another disadvantage is that interpolated values can only be directly derived for positions that coincide with the original data set. This necessitates another level of interpolation for intermediate points—a software consideration only.

For the polynomial fitting procedure to be fully effective, it is critical that the software in use allows the user to vary the parameters defining the piecewise polynomials and provides interactive visual feedback of the resultant fit. The user can then iteratively vary these parameters until a visually satisfactory fit is obtained at each step of the levelling process.

Tie sequencing when levelling ties

The success of levelling the ties is dependent on the sequence in which the ties are used. The earlier a tie is used, the greater effect it has on the overall result. This is because only the current set of levelled ties is used to predict the drift along a flight (see section ***Levelling ties***). Early on in the process there are not many crossover points being used, they may be poorly distributed along the flight, and any that are inaccurate may have a marked effect on the predicted drift. To minimise this effect, the ties should be sequenced according to the following factors:

- Ties should extend as close as possible across the full width of the survey area. Note that only those flights that intersect both the currently selected set of ties and the next tie in the list can be used. This may lead to regions of the unlevelled tie being uncontrolled if their geographical extent is beyond that of the currently levelled ties.
- To minimise the effects of navigation errors, preference should be given to ties over areas of low magnetic gradient.
- Preference should be given to ties flown during periods of quiet diurnal activity.
- The ties should be sequenced in such a way as to best define the drift curves for each flight in the tie levelling process. The wider the spread of points in time along the drift curve, the better. Thus, the first tie should be chosen as close to the centre of the survey as possible, the second towards one boundary of the survey, and the third close to the other boundary. Subsequent ties should be sequenced to lie as close as possible at the centre of the gaps across the survey area formed by the previously selected ties.

Appendix 2. AGSO micro-levelling

The AGSO micro-levelling tool is based on the procedure developed by Minty (1991). The technique is based on the assumption that residual errors in the data are characterised by being elongated along the flight-line direction and confined to individual lines. That is, visually, they would appear as streaks in a grid of data which can be theoretically detected and removed from the grid by the application of directional filters. The filtered grid can then be used to correct the original point-located data. The other constraint on the residual errors, especially for aeromagnetic data, is that they are expected to have a small dynamic range about zero.

The micro-levelling method is not rigorous, and cannot distinguish between levelling errors and real elongate anomalies parallel to the flight-line direction. Likewise, neighbouring lines may have errors of similar magnitude and it will only be the difference between the errors that can be detected.

The AGSO procedure for micro-levelling is as follows. If a grid of data is produced, then the residual errors will be evident as spurious elongate anomalies characterised by the following:

- a wavelength in the flight-line direction greater than the tie-line spacing;
- a wavelength perpendicular to the flight-line direction of precisely twice the flight-line spacing;
- a relatively small dynamic range.

If one grid axis is parallel to the flight-line, then the spurious elongate anomalies can be removed from the grid by applying one-dimensional filters in turn to the rows and columns of the grid. The procedure is as follows:

- Create a grid from the original point located data, say grid A.
- Apply a high-pass filter to grid A in the direction perpendicular to the flight-line and store the result in grid B.
- Inspect grid B to restrict its dynamic range by setting all values that fall outside two user defined-limits to those limits.
- Apply a low-pass filter to grid B in the flight-line direction and store the result in grid C. Grid C should now contain only the elongate anomalies we wish to remove.
- Inspect grid C to restrict its dynamic range by setting all values that fall outside two user-defined limits to those limits. Grid C now contains the required residual errors or corrections to be applied to the survey data.
- Subtract grid C from the original grid, A, to get the final grid with all spurious elongate anomalies removed.
- Once a visually well-levelled grid is obtained, the corresponding grid of residual errors from step 5 is subtracted from the actual point-located data, thus completing the micro-levelling process.

Though conceptually simple, the implementation is difficult, owing to the mathematical limitations of digital filters and the subtle nature of errors compared with the large dynamic range of the data. The filters must detect all the residual errors, but at the same time exclude as many real anomalies as possible and introduce no artefacts, that is they must be designed to correspond with real anomalies as little as possible—they must be long and narrow. Therefore, the object of the exercise is to extract residual errors with the longest possible wavelength along the line; the shortest possible wavelength perpendicular to the lines; and the smallest dynamic range such that the procedure still produces a visually well-levelled grid. This is achieved by trial and error by varying the cut-off wavelengths and the allowed dynamic range of the various filtered grids.

Typically, good results are achieved with the following parameters:

- a long wavelength of the order of 1–3 times the tie spacing;
- a short wavelength of 2 times the line spacing;
- limits to the dynamic range of 5–10 nT.

To give some idea of the numbers involved, micro-levelling in a recent survey adjusted 30 per cent of the data by amounts of 1–10 nT and only 23 per cent of points by less than 0.1 nT, i.e. very significant adjustments were made.

Although the micro-levelling technique is conceptually simple, the design of adequate filters to implement it is not. Various filtering techniques have been investigated with the following combination producing the best results:

- Naudy non-linear filter (Naudy & Dreyer 1968) for high-pass filtering;
- Fuller band pass convolution filter (Fraser et al. 1966) for low-pass filtering.

The Naudy filter is not a true frequency filter, but operates by detecting anomalies of wavelength shorter than the defined cut-off. Such anomalies are then replaced by extrapolation from neighbouring data points. It is particularly effective as a high-pass filter, normally being used to detect high-frequency noise. Points to note are:

- it is designed to completely remove all anomalies of width less than a given cut-off, while those above the cut-off retain their original shape;
- the anomaly-detecting algorithm is not foolproof and some anomalies may not be removed;
- the interpolation algorithm will leave some data of wavelength shorter than the intended cut-off.

The convolution technique of Fuller is less than perfect: the response of the filter at the cut-off wavelength is not abrupt; for low-pass filters it acts as a smoothing operator where the area under the profile remains constant; and data extrapolation is required to allow filtered values to be determined up to the data boundaries. Therefore:

- leakage occurs around anomalies whose wavelength is comparable to the cut-off wavelength of the low-pass filter;
- when a high-amplitude, short-wavelength anomaly occurs, the base level of the filtered data about this anomaly is raised, producing an artefact similar to the residual errors that are to be removed.

For both types of filters, the data must be extrapolated beyond the edge of the grid to allow filtered values to be determined up to the grid boundary. These extrapolation techniques often lead to unwanted edge effects, especially where large anomalies occur along the grid boundary.

References

Akima, H., 1970. A new method of interpolating and smooth curve fitting based on local procedures. Journal of the Association of Computing Machinery, 17(4), 589–602.

Barton, C.E., 1988. Global and regional geomagnetic reference fields. Exploration Geophysics, 19, 401–416.

Barton, C.E. & Johnson, B.D., 1988. Workshop report and discussion paper: geomagnetic reference fields; base-level and diurnal corrections to aeromagnetic data. Exploration Geophysics, 19, 447–452.

Boyd, T.J., 1992. An annotated outline of GPS. ASEG Preview, 15–26.

Briggs, I.C., 1974 Machine contouring using minimum curvature. Geophysics, 39(1), 39–48.

Fitzgerald & Associates, 1996. INTREPID geophysical processing system and visualisation tools. Desmond Fitzgerald & Associates, Unit 2, 1 Male Street, Brighton, Melbourne, Victoria 3186, Australia.

Foster, M.R., Jines, W.R. & Van Der Weg, K., 1970. Statistical estimation of systematic errors at intersections of lines of aeromagnetic survey data. Journal of Geophysical Research, 75(8), 1507–1511.

Fraser, D.C., Fuller, B.D. & Ward, S.H., 1966. Some numerical techniques for application in mining exploration. Geophysics, 31(6), 1066–1077.

Green, A.A., 1983, A comparison of adjustment procedures for levelling aeromagnetic survey data. Geophysics, 48(6), 743–753.

IAGA, 1991. International Association of Geomagnetism and Aeronomy, Division V, Working Group 8, International Geomagnetic Reference Field, 1991 Revision. Journal of Geomagnetism and Geoelectricity, 43, 1007–1012.

Minty, B.R.S., 1991. Simple micro-levelling for aeromagnetic data. Exploration Geophysics, 22, 591–592.

Naudy, H. & Dreyer, H., 1968. Non-linear filtering applied to aeromagnetic profiles. Geophysical Prospecting, 16(2), 171–178.

Ochadlick, A.R.Jr, 1989. Measurements of the magnetic fluctuations associated with ocean swell compared with Weaver's theory. Journal of Geophysical Research, 94(C11), 16,237–16,242.

Parkinson, W.D., 1983. Introduction to geomagnetism. Scottish Academic Press, Edinburgh.
Weaver, J.T., 1965. Magnetic variations associated with ocean waves and swell. Journal Of Geophysical Research, 70(8), 1921–1929.
Yarger, H.L., Robertson, R.R. & Wentland, R.L., 1978. Diurnal drift removal from aeromagnetic data using least squares. Geophysics, 43, 1148–1156.