GEOSOFT

Technical Paper

A New, Rapid, Automated Grid Stitching Algorithm

Stephen Cheesman, Ian MacLeod and Greg Hollyer, Geosoft Inc., Toronto

Abstract

Regional compilations of gridded geophysical data from disparate individual surveys are playing an ever more important role in resource exploration. A key processing step in such compilations is the merging of overlapping grids to create a single grid. Traditional methods of connecting grids together can produce smooth final products but the process is time-consuming and has difficulty with differences that involve both long and short wavelength errors. A novel, completely automated method addresses several main challenges, such as determining how to select a path along which overlapping grids can be joined. The technique uses Fourier analysis to deconstruct the errors along a suture path into a sum of functions with different spatial wavelengths, and applies corrections that propagate smoothly into the grids by a distance proportional to the individual wavelengths. The result is an almost seamless grid that minimizes distortion from the correction process.

Introduction

The compilation of regional geophysical data sets from individual surveys and smaller compilations has become a significant activity in the support of exploration programs and in the development of national and international earth resource databases. Such compilations involve many processing steps (Barritt et al., 1993; Black et al., 1995; Johnson, 1996; Reeves, 1994; Reford et al., 1997, 1990; Tarlowski, 1996), but a key activity in almost all compilations is the merging of two or more gridded data sets into a single consistent grid.

After removing all possible systematic noise, and bringing grids to a common physical reference level, adjacent geophysical grids will rarely match along their common edges (Fairhead et al., 1997; Black et al., 1995). This is despite the fact that each data set is intended to observe the same physical phenomena. As a result, the creation of a simple mosaic of the gridded data is often not sufficient to deliver the interpretational and processing value of the data.

In all the compilations referenced here, a significant component of the compilation work was involved in sorting out and correcting for the differences along grid edges. Traditional techniques are semi-automated and involve a number of steps, such as manually adjusting values on neighboring grids, leveling to an existing low-resolution grid or using various weighting schemes to merge grids (Black et al., 1995). The professional time required to join two grids can range from hours to days (S. Reford, personal communication). Furthermore, when the differences between grids include a wide range of wavelengths (long to short), the best-effort joins using these techniques can still be poor.

In this paper we describe a fully automated technique for merging two grids. The primary goals of this method are to minimize the professional time involved in the merging process itself, and to resolve corrections with a variety of wavelengths with a minimum in the distortion of real data. A third objective of the method is to provide flexibility for the data processor to choose and customize the techniques and merging parameters to suit the wide variety of problems that may be presented by different data sets.

The compilation of a number of gridded data sets into a single grid first involves the removal of all possible systematic noise (leveling error, IGRF removal in magnetic data, continuations to a common observation plane, de-aliasing filters, etc). The process of compiling the final merged grid can then be reduced to the task of merging, or "suturing" two overlapping grids, which can then be sutured to neighboring sutured grid pairs, and so on until a final grid is produced. The term "suture" is used because this method requires the definition of a suture path in the overlap region, along which data on either side is discarded or modified to produce a continuous function along the suture line.

Grid Preprocessing

Regional compilations must often deal with very large gridded data sets, and every processing step can be time consuming and require large amounts of disk space. The new process is designed to minimize the number of preprocessing steps required before applying the suturing process.

The process that has been implemented will also merge grids that have different cell sizes or map projections. If required, the merging process re-samples one of the grids to match the cell size and projection of the other.

The grids may also contain static offsets or differences in average slope. These are corrected by selecting either a static shift (zero order trend) or a slope difference (first order trend) between the grids. Generally, the correction is applied to give one grid the trend of the second; however, if corrections are requested of both grids, they are treated independently and the trends are simply subtracted from the individual grids. The trends may be calculated using all the points in a grid, just the edge points, or the points in the region of overlap between the two grids. Some care should be taken when using only the overlap points when calculating a slope. In the case where only a few points overlap, or where the overlap occurs along a narrow line, the calculated trend may differ substantially from the slope of the entire grid, and the "correction" may in instead introduce an undesirable slope in the output.

Suture Path Selection

The suture method requires the definition of a join path between the two grids, along which, and around which changes to the grids are made to create a smooth transition. Only the difference between the two grids along this path is taken into consideration, as opposed to the traditional blending or feathering method, in which the grids are added together in varying fractions over an area to create a smooth transition.

By definition, the suture path can only be chosen through regions where the two grids overlap. Moreover, the end points are by necessity where the edges of the two grids cross. Figure 1 shows various options for defining paths.

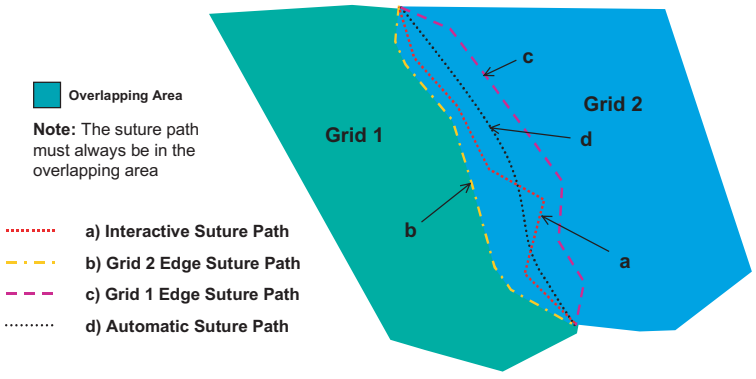

Figure 1.
VARIOUS PATHS THROUGH THE OVERLAP REGION: A) USER-DEFINED, B) EDGE OF GRID #2, C) EDGE OF GRID #1, D) AUTOMATIC (BISECTION).

The geophysicist may interactively define a path joining the two fixed end points through the overlap region. For example, it may be desirable to include or exclude certain features occurring in one of the grids, and it is possible to do so by drawing the path around them.

Two "external" paths are the edges of the grids themselves. Often, one grid is of higher quality or resolution than the other, and choosing its edge for the suture path ensures that as much as possible of its contribution will be included in the joined grid.

The path may be chosen automatically to bisect the overlapping region. This is accomplished by calculating, for every position in the overlap region, the distances to the nearest edge point in each grid. The suture path is derived naturally as the collection of points where the distances to both grid edges is most nearly equal. This works even where there are "holes" in the data.

The path is created from the selected points by joining them together – much as beads on a necklace. This process can run into difficulties, especially with jagged grid edges or breaks in the data. Ambiguities can arise over which point gets joined to which, and it is not uncommon for the suture path to be broken into a number of sections of varying length, spanning several overlapping areas. The method works best with a relatively straight path, and careful attention to geometrical details is necessary, not only to produce a valid path, but get the best possible results in the suture process. For instance, when three points are found to join at right angles in a corner, the "corner" point is eliminated. This not only produces a smoother transition between a vertical and a horizontal path, but also allows diagonal paths to be defined without a "stair-step" pattern.

Decomposition of the Difference Function

The difference or error function is defined as the difference between the Grid #1 and Grid #2 values along the suture path (Figure 2).

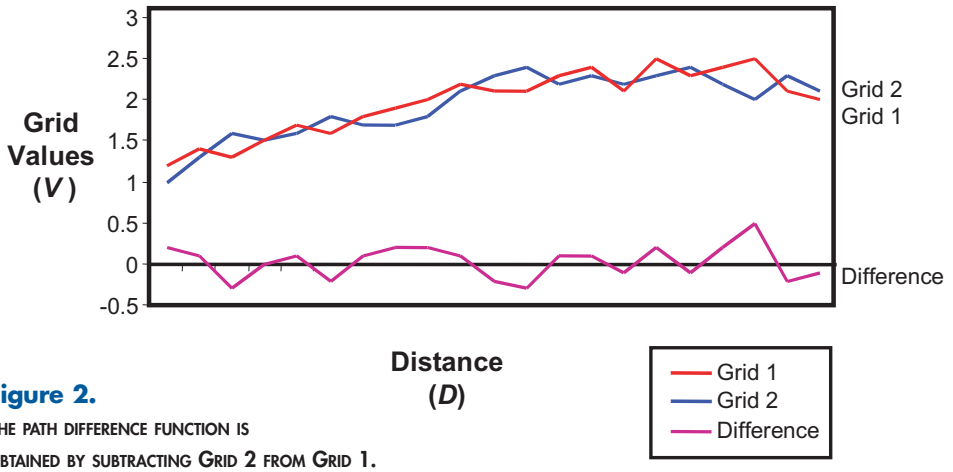

Figure 2.
THE PATH DIFFERENCE FUNCTION IS OBTAINED BY SUBTRACTING GRID 2 FROM GRID 1.

This difference function shares much of the character of the grids that define it, with both short and long wavelength features. One of the drawbacks of current grid-joining algorithms that define a join path is their relative inflexibility in determining how far away from the path to "feather" corrections. Large-wavelength errors should be corrected over a larger region than short wavelength differences.

The basis of the new method is the recognition that, as with any linear function, the difference function can be decomposed using Fourier analysis into a number of simple sine-type functions of different amplitudes and phases. Corrections are then applied for each single-wavelength function, and summed. In each case, corrections are applied at a distance from the suture path no greater than one-quarter of the wavelength under consideration. The result is that the total solution automatically adjusts itself along the path to do as much, and no more than, what is necessary to provide a smooth join.

For the purposes of this calculation, the points are considered to be equal-spaced, even though there will likely be "diagonal" as well as horizontal and vertical pairs. This considerably eases subsequent calculation, and the distortion introduced by this dodge is minimal. Corrections are applied in one sense to Grid #1, and the opposite sense to Grid #2, in order to make up for the difference. The user may weight the correction toward either of the grids as desired, for instance where one grid is considered to be of better quality.

A Fast Fourier Transform (FFT) is used, and the number of points is increased to at least the next power of 2 to allow padding and the use of a maximum-entropy prediction filter to prevent ringing.

Static Corrections

The "longest" wavelength in the decomposition of the difference function is of course the static or zero frequency term. It makes no sense to apply this correction everywhere, since there is usually more than one suture path, and each will have a unique static term. Instead, the static term is deemed to have a "cut-off" distance equal to one-quarter of the length of the path. At any individual grid position within this range, the distance to the nearest path point is determined, and a correction is applied using a cosine taper function that applies the full static correction to points along the path, and no correction beyond the cut-off distance.

Long Wavelength Corrections

Consider Figure 3. Grid #2 with a sinusoidal undulation must be joined to a flat grid #1. The pertinent question is: Assuming that all the correction is to be done to grid #1, what kind of correction makes the transition as smooth as possible? To first order, the suture method adds semi-circular "dome" corrections with the same cross-section as the sinusoid, and hence with alternating sign, along the path (see Figure 4).

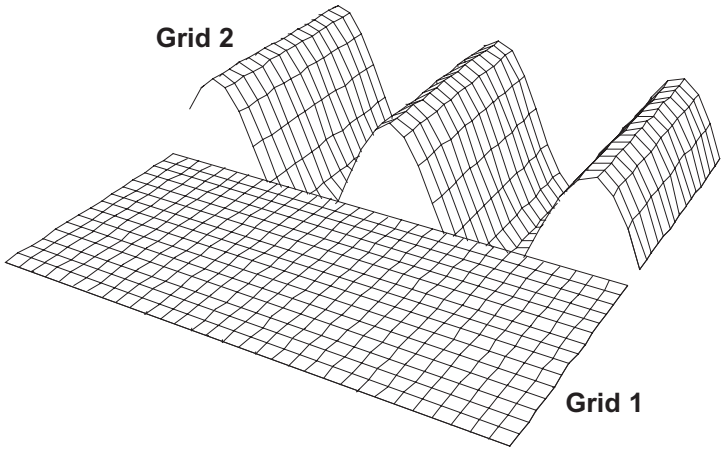

Figure 3.
THE SINUSOIDAL DIFFERENCE AT THE JOINT BETWEEN TWO GRIDS. THE METHOD BREAKS DOWN THE ACTUAL JOIN INTO A SUM OF SIMILAR PROBLEMS.

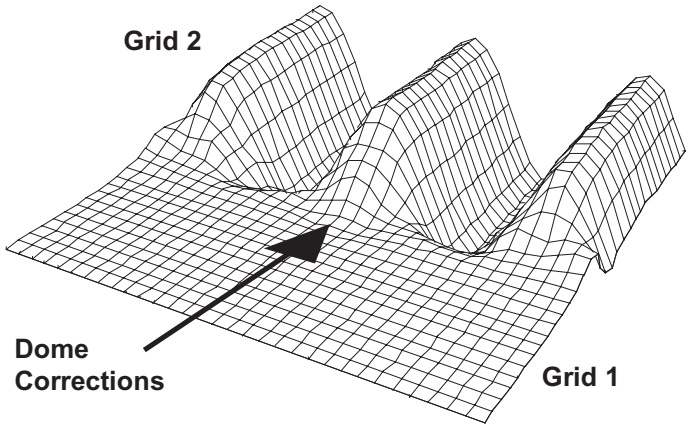

Figure 4.
THE DOME CORRECTIONS HAVE A SEMI-CIRCULAR SHAPE WITH A SINUSOIDAL CROSS-SECTION. THE NEGATIVE CORRECTIONS TO GRID 1 ARE THE INVERSE OF THE POSITIVE CORRECTIONS THAT CAN BE SEEN IN THIS FIGURE. NOTE THAT THE SHAPE OF THE DOME CORRECTION IS SIMILAR TO THE SHAPE OF THE SINUSOIDAL FEATURES OF GRID 2.

An astute observer will note that in cases where the correction is split between both grids (the default case), the added "semi-circle" will subtract from and distort the regular sinusoidal pattern of grid #2. However, the causative feature in grid #2 which produced the sinusoid along the difference path is just as likely to be circular, or irregular, or a shape which will not so notably reflect the nature of the correction. In fact, because this method is designed not to consider features away from the suture path, the result represents the most general case. Therefore, circular-type anomalies are arguably as likely to be encountered as the undulations in the contrived case shown here.

Corrections beyond the ends of the path must be treated specially. Because of phase shifts, a "dome" may have its center point well beyond the last path point, and it would be incorrect to assume that the correction takes on its full value over the entire area covered by the dome. Instead, the cosine taper is employed again, this time using the end point of the path as a reference, and a cut-off distance of one-quarter wavelength. Figure 5 shows in profile, along the path, the end correction for a single wavelength, both with the full dome (a) and with the cosine taper applied (b).

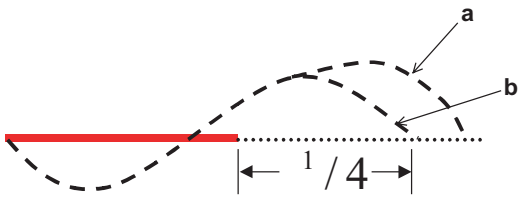

Figure 5.
COSINE TAPERING OF CORRECTION OFF THE END OF A PATH SECTION. THE CORRECTION FROM THE DECOMPOSED FFT WAVENUMBER (A) IS REDUCED TO ZERO AT A QUARTER-WAVELENGTH DISTANCE FROM THE END OF THE SUTURE PATH SHOWN AS A SOLID HORIZONTAL LINE.

A drawback of the dome solution is that it does not blend smoothly into the grid away from the path and produces an "edge" (along the path negative and positive domes meet to form the smooth sinusoid). To help reduce this effect the radial cross-section of the dome is modified somewhat at angles measured away from the path, from a sinusoid to a pure cosine taper. This transition occurs with about 10 degrees of the path, so that the area where edges are introduced is kept to a minimum (see Figure 6).

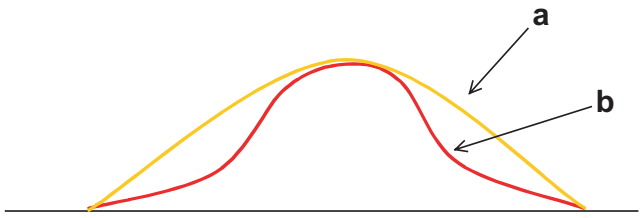

Figure 6.
ALTERATION TO THE DOME TO PRODUCE A SMOOTH TRANSITION INTO THE GRID AT ANGLES AWAY FROM THE PATH. A) CROSS SECTION ALONG THE PATH (HALF OF A SINE WAVELENGTH). B) CROSS SECTION AWAY FROM THE PATH (COSINE TAPERS JOINED BACK-TO-BACK).

Short Wavelength Corrections

The upper half of the spatial frequencies are short enough that one-quarter wavelength does not span a single grid spacing, with the result that corrections need only be made to the path points themselves. The corrections can therefore be summed together, and applied to the grids in one step.

Limiting Path Length

Initially, there was no limit to the path length, with the result that corrections could extend to cover much of the grids at low wavelengths. This was also rather slow and inefficient, as grid values over large areas of the grid had to be repeatedly accessed, both for the local static corrections and the lower, longer wavelengths. Both to improve efficiency, and to limit the corrections to closer to the path itself, the path length is now divided into sections no more than 64 points in length. This limits corrections to within 16 points (or one-quarter of the maximum wavelength) of the correction path, and speeds up processing considerably.

"Interference" Between Path Sections

When a point in a "dome" is located, the method applies a correction based on the distance from the center of the dome, and of negative or positive sign, depending on which grid the point belongs to, and in the pre-determined proportion based on the grid weighting. This works well for path sections in isolation but produces a discontinuity across a second join if the join crosses the field of the dome. A similar problem occurs when the path is split into smaller sections with a maximum of 64 points. Corrections tend to "interfere" with each other. The problem is resolved by employing yet another cosine taper (see Figure 7).

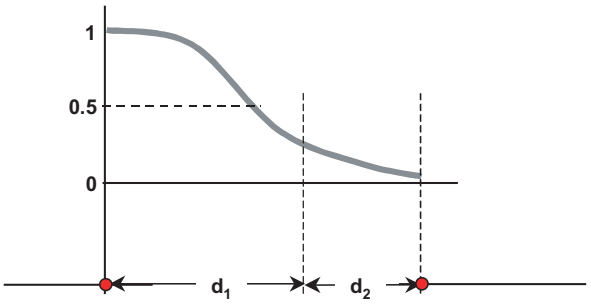

Figure 7.
"Interference" filter: The correction applied from the section on the left is reduced to zero as it approaches the section on the right to prevent interference, using the factor indicated by the cosine taper and the distances to the nearest points in both sections.

Both the distance to the closest point in the current section, d1, and the distance to nearest point in the nearby section, d2, are determined. If d2 is less than one-quarter of the current wavelength, then the correction is multiplied by a factor which varies like a cosine, from 1.0 when d1 = 0 to 0.5 when d1 = d2, to 0.0 when d2 = 0. This ensures that the effects of any one section's corrections do not unduly interfere with another's, and that the transition between the two is smooth. This has been found to work satisfactorily for both local static and low-wavelength corrections.

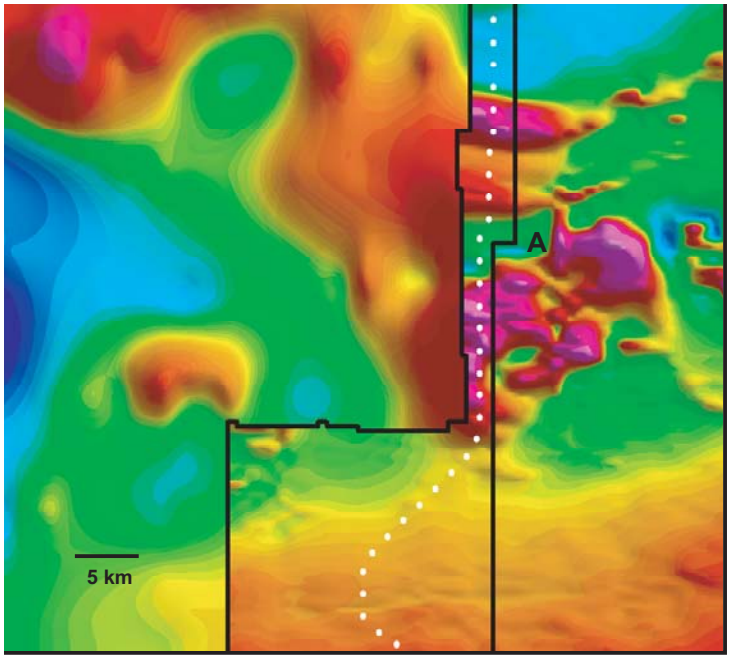

Figure 8.
Two total field magnetic grids from the N.U.R.E. study in Alaska. The grid on the left was surveyed using east-west lines, spaced 9600 metres apart at an altitude of 122, and has been upward continued to 300 m. The grid on the right was surveyed using 1920 m lines at an altitude of 300 m. The black lines show the overlapped edges of the two grids respectively, and the white dotted line shows the automatically chosen suture path.

Example

As an example, we have chosen two total magnetic field grids that were part of the U.S.G.S. and State of Alaska aeromagnetic compilation (Meyer and Saltus, 1995). Paterson, Grant & Watson Limited were contracted to compile 25 separate airborne survey grids into a single seamless grid. Using semi-automated techniques, this process required, on average, one day for each join. The pair of grids chosen for this example were particularly challenging and required a full two days to complete an acceptable join (Racic, personal communication).

Figure 6 shows part of the two grids together with the grid outlines and the automatically determined suture path. The most obvious challenge is the difference in sampled resolution that arises from the very different flight line separations. Also, at the centre of the figure, there is a pronounced low anomaly on the right-hand grid, which has simply not been sampled on the left-hand grid. Figure 9 shows profiles of each grid along the suture path.

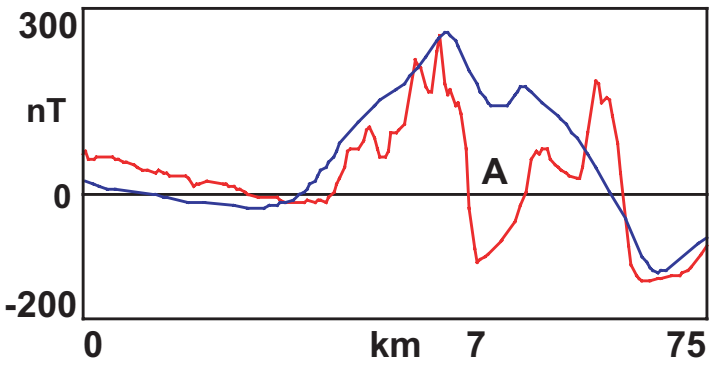

Figure 9.
PROFILES OF THE TOTAL MAGNETIC FIELD SAMPLED SOUTH TO NORTH ALONG THE SUTURE PATH FOR EACH GRID. THE BLUE LINE IS THE LEFT-HAND GRID IN FIGURE 8, AND THE RED LINE IS THE RIGHT-HAND GRID. THE SIGNIFICANT DIFFERENCES CAN BE ATTRIBUTED TO THE VERY DIFFERENT SURVEY SPECIFICATIONS. THE PROFILE AFTER SUTURING WILL BE EXACTLY THE AVERAGE OF THESE TWO PROFILES.

The new suturing algorithm was run on these grids with the following parameters:

- A first-order slope was calculated along the suture path and removed from each grid.
- The automatically chosen suture path was selected.
- The error was distributed evenly between both grids.

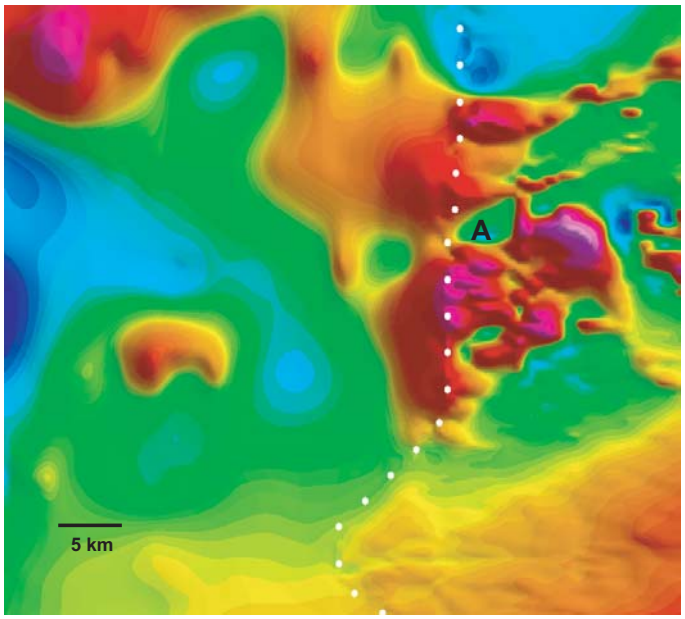

Figure 10.
THIS SHOWS THE RESULTING SUTURED GRID, WHICH HAS EFFECTIVELY DEALT WITH THE TRANSITION BETWEEN A WELL-SAMPLED AND HIGHER RESOLUTION GRID ON THE RIGHT, AND A POORLY SAMPLED, LOWER RESOLUTION GRID ON THE LEFT. THE DOTTED LINE INDICATES THE LOCATION OF THE SUTURE PATH.

The resulting grid is approximately 800x800 points and the suture process was completed in less than 60 seconds on a 266 MHz Pentium computer. The result shown in Figure 11 was produced on the first attempt.

Conclusion

The result of this automatic and rapid algorithm is a merged grid that is almost seamless across boundaries – regardless of the different wavelengths of features along the join path – and that maintains the integrity of the original data as much as possible. The technique of separating corrections into separate wavelength components and extending the required corrections to a minimum distance appropriate for each wavelength is the key component of the algorithm. By fully automating the process, while still providing control over the many parameters involved, this technique significantly reduces the time required to compile gridded data sets into a single grid.

Acknowledgements

The authors wish to thank Stephen Reford of Paterson, Grant and Watson, Ltd., and Louis Racic of Geosoft Inc. for their assistance in providing the background information and data used in the example. In addition, Chris Mussleman of Geosoft Inc. made an important contribution in preparing many of the figures that appear in the paper.

References

Barritt, S.A., Fairhead, J.D., and **Misener, D.J.,** 1993, The African Magnetic Mapping Project, ITC Journal, p. 122-131

Black, P.A., Green, C.M., and **Reford, S.W.,** 1995, A pragmatic approach to continental magnetic compilations, Annual Meeting Expanded Abstracts, Society of Exploration Geophysicists, p. 773-774.

Committee for the Magnetic Anomaly Map of North America, 1987, Magnetic Anomaly Map of North America, 4 sheets, scale 1:5 000 000, Geological Survey of America.

Fairhead, J.D., Misener, J.D., Green, C.M. Bainbridge, G., and Reford, S.W., Large Scale Compilation of Magnetic, Gravity, Radiometric and Electromagnetic Data: The New Exploration Strategy for the 90s, Proceedings of the Fourth Decennial International Conference on Mineral Exploration, p.805-815.

Johnson, A. C., and **von Frese, R.R.B.,** 1996, Report of the Working Group on the Antarctic digital magnetic anomaly map, British Antarctic Survey.

Johnson, A.C., R.R.B. von Frese, and **The ADMAP Working Group (1997)**: Magnetic Map Will Define Antarctica's Structure. Eos, Transactions of the AGU. 78, 185.

Meyer, J. F., and **Saltus, R. W.,** 1995, Merged Aeromagnetic Map of Interior Alaska: U.U. Geological Survey Map GP-1014, scale 1:500,000.

Reeves, C.V., and **Erren, H.,** 1994, AAIME, Aeromagnetics of Arabia, India and the Middle East, International Institute for Aerospace Survey and Earth Sciences, unpublished.

Reford, S.W., Gupta, V. K., Paterson, N.R., Kwan, K.C.H., and **MacLeod, I.N.,** 1990, The Ontario master aeromagnetic grid: a blueprint for detailed magnetic compilation on a regional scale, Annual Meeting Expanded Abstracts, Society of Exploration Geophysicists, p. 617-619.

Tarlowski, C., Milligan, P.R., and **Mackey, T.E.,** 1996, Magnetic Anomaly Map of Australia (2nd edition), scale 1:5 000 000, Australian Geological Survey Organisation.