International Hydrographic Review, Monaco, LXIII (1), January 1986

CONTOURS AND CONTOURING IN HYDROGRAPHY PART II - INTERPOLATION

by M.J. CASEY (*) and D. MONAHAN (**)

This paper has already been published in Lighthouse, *Edition No. 30, November 1984 and is reproduced here with the kind permission of the Canadian Hydrographers' Association.*

ABSTRACT

In Part I of this series, the authors discussed those issues which we feel are fundamentally important and which must be addressed by any method which aims to mechanize the drawing of depth contours for hydrographic charts.

In this article we begin the discussion of the How of contouring. In particular, we concentrate on some of the most common methods used in the interpolation of the synthetic surface upon which computed contours will lie.

INTRODUCTION

Mathematical interpolation is the heart of machine contouring — the rest is purely cosmetic. This is the thesis we follow in this paper. Cosmetics are an important issue — but they are secondary in importance. The interpolation algorithm will determine the shape and course of the plotted contours and *this* is what we care about.

(*) Planning & Development, Canadian Hydrographic Service, Department of Fisheries & Oceans, 615 Booth Street, Ottawa, Ontario K1A OE6, Canada.

(**) Canadian Hydrographic Service, Headquarters, Department of Fisheries & Oceans, 615 Booth Street, Ottawa, Ontario K1A OE6, Canada.

Before becoming immersed in the details of interpolation, let's examine the situation at a *higher* level. Figure 1 illustrates the main ideas behind interpolation. Figure 1a shows a sequence of measured sounding profiles. This is the data from which we wish to draw our contour map. One can imagine the contours as a sequence of shoreline snapshots — each one taken with the level at progressively lower elevations. Figure 1b shows the situation at a particular water level — say 10m below datum. The *lower* level problem is this — *how* do we connect-up the protuberances above each water level in a meaningful way ?

In order to draw contours we need to predict the behaviour of the contours *between* the survey lines. To do so we want depth estimates at regular intervals between the observations. The closer together the depth estimates, the smoother the contours. Figure 1c illustrates one popular approach called *gridding*. In this method, one drops a *uniform grid* over the survey area and, at the grid intersections (called *'nodes'*), estimates depths by using the observed depths. How these estimates are made is the crux of the matter.

WHY GRID THE DATA ?

The threading of the individual contour lines through the survey area can be a straightforward procedure if the data is established on a uniform and tightly spaced sampling plan. Contouring a typical field sheet for instance, where the soundings are spaced every 5 mm at scale, is relatively straightforward and a set of rules can be established to define the contouring procedure. When the data is sparse, however, the rules become less meaningful and, as a consequence, more and more judgement is called for. This becomes a case of *interpretation*, not interpola-

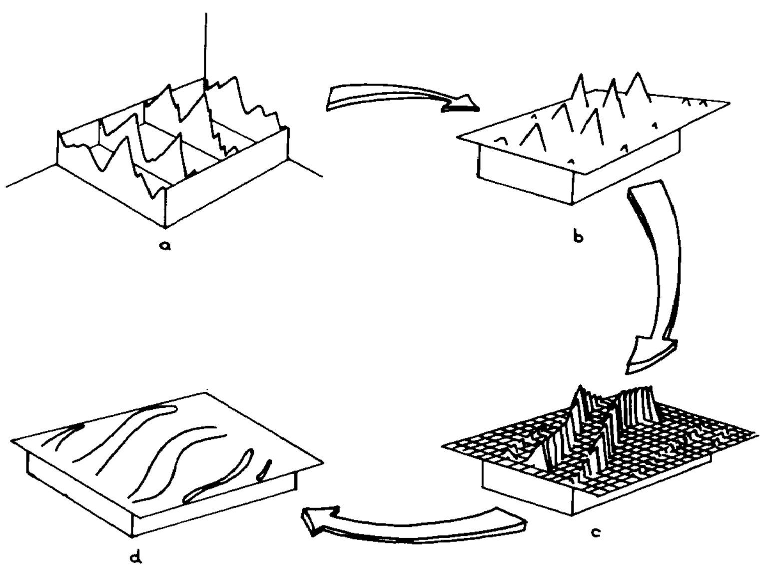

FIG. 1. — Interpolation and contouring.

tion. Such procedures cannot, in general, be mechanized. To overcome the problem of sparse or non-uniform sampling, researchers have found it expeditious to re-cast the survey so that it would appear to have been sampled in a more convenient manner. *Density* and *uniformity* are the two characteristics of the data which make machine contouring more viable. The uniform grid is an obvious choice but others, including triangulation schemes, are used in practice. We concentrate on gridding because that is the technique with the widest usage and because, in the end, the differences between gridding and its alternatives are often academic.

The actual contouring itself is done by threading the individual contours through the grid. Once depth values have been established for each grid node, these grid nodes can be used as gate-posts, allowing or denying access to the *interior* of the grid box. If access is allowed, then progressively finer grids can be established inside the main box in order to guide the course of the contour. In this way the contour's apparent smoothness is governed by the fineness of the gridding. Finer gridding can improve the *smoothness* of the contours and make them appear more "realistic". Appearances can be deceiving, however. The *accuracy* of the contour position is governed by the survey sampling resolution — not by the resolution of the grid.

There are essentially two main methods for making these estimates : *point models* and *area models*. Point models estimate depths at *fixed* points in the area — such as at the grid nodes. Area models, on the other hand, estimate surface *continuously* over an area. That is, a smooth mathematical function (often a polynomial in 2-dimensions) is fitted to the data points within an area. This estimated surface then defines depth estimates at *every* point within the region over which the surface is fitted. Thus the knowledge of the contours' location is continuous — resulting in very smooth contours. Of course, if the surface model is wrong (i.e. if the fit is not very good) then the contours are *also* wrong.

We now examine these two methods in some detail.

POINT MODELS

Point models estimate depths by linearly combining the surrounding observed data points. That is, the depth at any unobserved point can be "guessed" by a weighted average of the observed depths which fall geographically *close* to the unobserved point. Consider Figure 2. In this illustration we see the essence of point modelling. On the x-y plane we have a number of observations. The observations are more-or-less *randomly* located on the plane with no connecting information uniting them. To contour, we want to have values on the uniform grid. This can be accomplished by sequentially marching through the grid nodes and making weighted averages at each grid intersection. In Figure 2b we examine one intersection in detail. The area around the point to be estimated is searched for observations and then these values are used in the weighted average formula. The two issues we should concentrate on are the concepts of *neighbourhoods* and *weighting*.

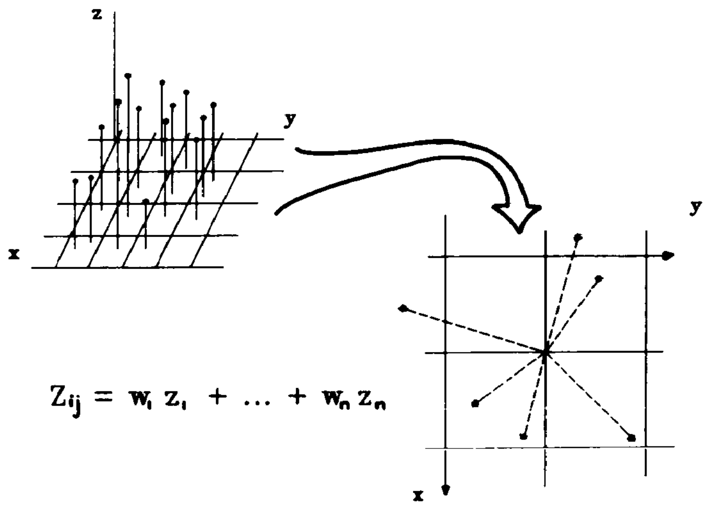

FIG. 2. — Point models.

NEIGHBOURHOODS

A typical field sheet contains about 20,000 soundings. How should we use this vast volume of data ? In general, the soundings in the lower-left corner of the field sheet cannot be used to predict the behaviour of the bottom portrayed in the upper-right corner. In hand contouring, hydrographers examine only those soundings which sit close to the spot where his pen lies. We need some similar mechanism to limit the blind inclusion of excessive and unconnected data.

Neighbourhoods are required to limit the number of data points included in the linear combination. In its simplest form the neighbourhood is defined to be a circular area of user-set radius, centred upon the point to be estimated. Any observations found within this area are included in the computation. Figure 2b shows that six observations were found in the neighbourhood of the central grid node. The observations are found by doing a search of the data record and checking each point to see if it falls within the neighbourhood.

Unfortunately, defining the neighbourhood as some simple circular area surrounding the grid node won't always work. Figure 3 shows some of the drawbacks of using a simple, pre-defined, static neighbourhood.

In Figure 3a we have the problem of sparse data. In order to include a minimum number of observations in the calculation, some programs expand the search circle in increments until either the required minimum number of points is included or some maximum radius is achieved. Alternatively, if there are too many points in the standard neighbourhood, the radius is reduced incrementally until the maximum number of desired observations is achieved or until some pre-set minimum radius is reached. Note that the neighbourhood in these cases is independent of the *value* of the data and is dependent only on its *spatial structure.*

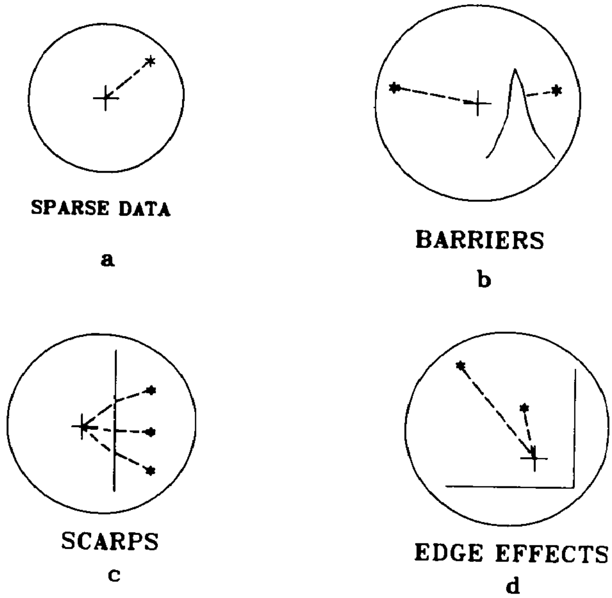

FIG. 3. — Neighbourhood problems.

Figure 3b illustrates another common problem — how to include some *unmeasured* information in the interpolation scheme. Here we have a situation where an observation is geometrically close to a grid point — but on the other side of a *barrier* (in this case a point of land). Clearly, the observation indicated should *not* be used in estimating the value at the grid point — even though it falls in the neighbourhood. To overcome this one could sample, digitize and include the shoreline information as observations. But this method is not foolproof since the shoreline points are considered in the same way as all of the other observations — as elements in the weighted average. The real solution is to have some mathematically impermeable barrier through which interpolation cannot take place. The inclusion of barriers is a feature of many of the more sophisticated contouring packages available on the market.

A similar problem is shown in Figure 3c. Here we have an underwater scarp or cliff. The problem here, again, is that observations made on *one* side of this feature should *not* be included in the estimation of grid point values on the *other* side. In this particular case, only data points on the *lower* level of the scarp should be included in the estimation of grid nodes there.

The problem illustrated in Figure 3d is one of *extrapolation,* not interpolation. This problem is particularly apparent with polynomial surface fitting procedures and is diagnosed as a very "wavy" appearance of the contours along the edge of the sheet. One solution is to ensure that only *interpolation* takes place. This can be done by having the program *only* contour within a window which is bounded externally by observed depths.

WEIGHTING SCHEMES

In the depth estimation process, weights are applied to the observed data to allow observations of higher "quality" to have a greater influence than points of lesser "quality". This quality feature can refer to the *relative* quality of the various depth measurements, but is usually used as a means of ensuring that "closer" observations have a higher weight than observations farther away. In this restrictive sense, quality is a function of the distance between the observations and the grid node to be estimated. The fact that closer observations should have a higher weight than ones farther away is initially appealing, but is *not* universally applicable. This can be seen in Figure 3b where the geometrically closest observation is hidden behind a barrier. Hence a more sophisticated distance-weighting scheme is required.

The distance weight can be simply the inverse-distance between the observations and the grid node to be estimated but, usually, the *inverse-square* of the distance is used. This ensures a faster drop-off of the influence with distance. This inverse-square weighting is commonly referred to as the "Gravity Model" — the relevance being the inverse-square relationship between two bodies in Newton's Universal Law of Gravitation. Note, again, that the weighting is independent of the *value* of the observations but is based on the *spatial* relationship alone.

Figure 4b illustrates another problem associated with simple distance weighting — *trends* in the data. The data points on the left hand track will clearly have a greater influence on the estimate than those on the right hand track. Suppose that there is a left-right linear trend to the data with the soundings on the left considerably deeper than those on the right. Then the estimated depth at the

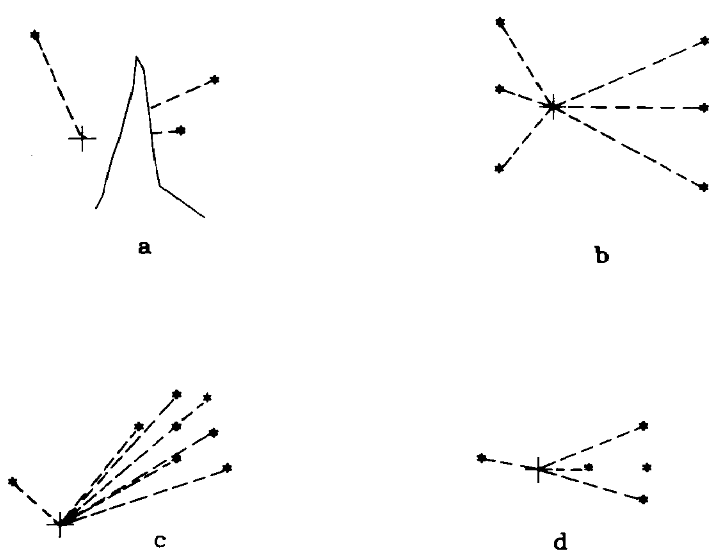

FIG. 4. — Weighting problems.

point indicated will be *deeper* than it should be. Weighted averaging systems have no *explicit* way of handling data which has clear trends in it. Such data sets are better handled by contouring algorithms which model these trends.

Clustering of data points can also place emphasis on the *wrong* data. This situation is shown in Figure 4c. The data points to the upper right will have a greater effect on the estimate than the lone point to the left — yet this point *should* be included in any estimate. Why ? Because, firstly, it is *closer* than the other points, and secondly, it acts as a means of *determining* trend.

Data *shielding* is also an important consideration. If one were to hand-contour the data shown in Figure 4d, one would *not* consider the values of the data points in the line to the right of the right-hand sounding. This would be inappropriate since the right-hand data point *shields* the line data. For instance, if the track data was considerably deeper than the other two soundings, then the interpolated depth would be influenced by the nearness of these deep soundings. This would result in a depth estimate which is *too deep* — thus showing a depression where one probably does not exist. To get around this problem, one can apply a second level of weighting — *directional* weighting. In this case, the algorithm must seek out observations within the neighbourhood which are shielded by other observations. This can be accomplished by examining the spatial relationships of the observations vis-a-vis the grid node, determining the associated angles and de-weighting any observations which fall within the shadow cast by closer observations. Many modern contouring programs feature such directional weighting automatically.

Another of the consequences of simple weighted averages is the fact that each observation is considered as a *local extremum.* That is, each observation is considered to lie on either the peak of a local hummock or the pit of a local depression. This is a direct fall-out from the use of the weighted average. The grid estimates will *always* be bounded by the observations. One cannot estimate a value deeper than the deepest observation within the neighbourhood and neither can one estimate one shallower than the shallowest. The effect of this is most apparent when a rugged area is contoured at a close contour spacing. The observations are all ringed by contours. This might make sense for topographic surveying where observations are taken upon the local extrema, but it would *never* make sense in hydrography where we *never* see the surface we are mapping and consequently the chances of occupying a local peak or pit are slight.

Including some slope information is the way to get around this particular problem. But slopes are not observable in hydrography, so they have to be *inferred.* Geomorphologists use *external* information on the surrounding geology and geomorphology to help them create models for the unseen surface. If this information is not available, then the observations alone must be used to estimate the slopes. Essentially this involves calculating the slopes from the differences in depth of the observations. Several of the contouring packages available commercially offer slope estimation as a program option. Once slope information is available, it can be used to predict extrema other than at the observation points.

AREA MODELS

To overcome some of the above limitations — particularly those which deal with *trends* in the data — methods have been derived which specifically *exploit* these trends to make estimates. Such methods assume that the surface can be expressed as an analytical function — usually a polynomial. Deviations from these surfaces would be classified as noise. Figures 5, 6 and 7 show some examples of analytical surfaces used by surface-fitting programs.

Having the surface expressed as a mathematical function has certain advantages. Depth estimates can be calculated at *any* position. Once the surface has been fitted, grid estimates at any density can be calculated quickly and easily. The *shape* of the surface is also, to some degree, predictable. A polynomial surface of the first order would exhibit a constant slope. A quadratic surface (Figure 7) would show concavity — either upwards or downwards, depending on the data. A Fourier surface (Figure 6) would appear periodic. This ability to predict the shape of the surface has some attraction because we are often faced with data which has clear trends which could be exploited by such surface-fitting techniques. On the other hand, if we *force* a surface onto data which does *not* exhibit such a trend we could introduce artificial features into the surface — for instance, more hollows than actually exist.

In order to gain an appreciation of this problem let us consider the surface shown in Figure 7. This is a bivariate (i.e. a 2-variable), second-order polynomial — a common function used in surface fitting. Fitting functions to data is usually accomplished by using a numerical technique called *regression* or, more commonly,

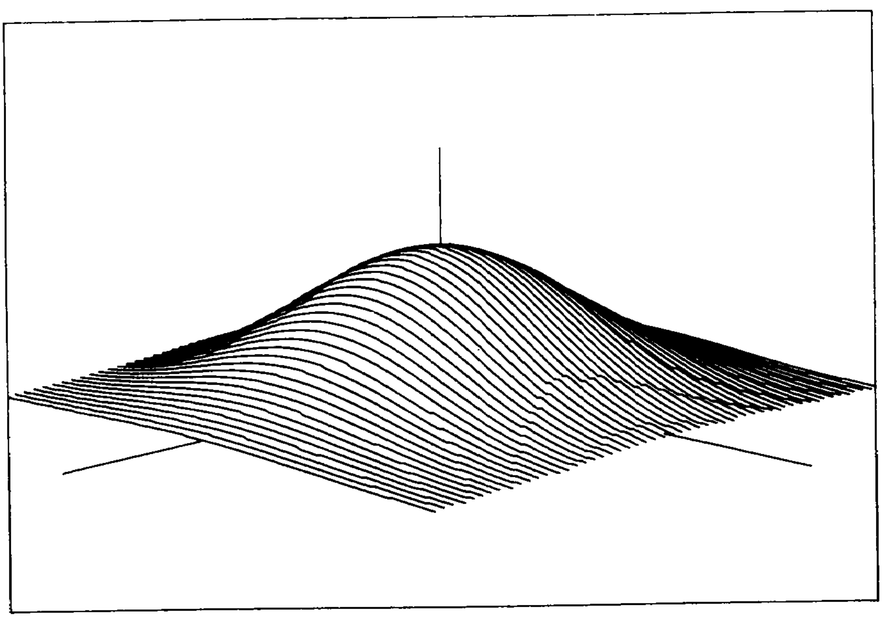

FIG. 5. — An exponential surface.

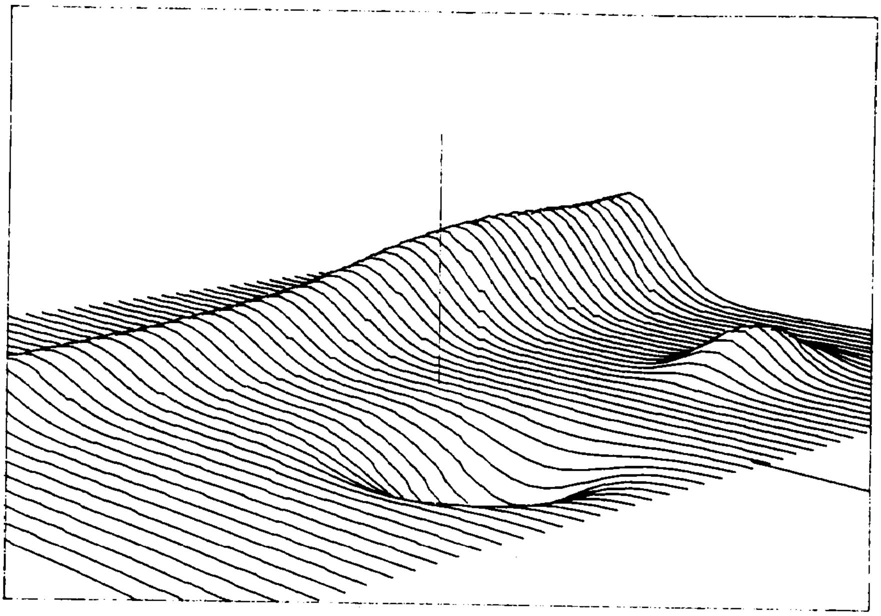

FIG. 6. — A Fourier surface.

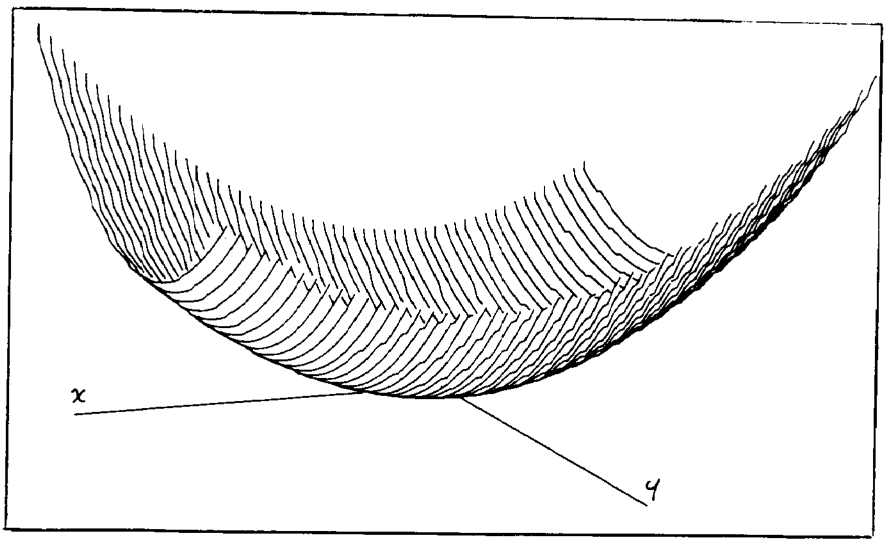

FIG. 7. — A quadratic surface.

least squares. The details of the method can be found in any text on regression, such as *Applied Regression Analysis* (DRAPER and SMITH, 1981).

To fully see the effects that the fitting of such a surface has, it is far easier to examine the graph in *two* dimensions. Let's see what happens in the quadratic case. In other words, we will take a slice (or section) out of a surface like that of Figure 7.

In Figure 8a we have some data taken from a depth profile to which we have fitted a quadratic curve. We can see that the choice is not a good one. The data does *not* exhibit a significant trend — yet a quadratic one has been imposed. This is important, since the contours will be determined by this *artificial* surface — *not* the one defined by the observations. Figure 8b illustrates another problem. In this case the data *does* exhibit a trend — a *linear* one, whereas, again, a *quadratic* one has been imposed. Depths substantially *deeper* than measured have been estimated. Alternatively, Figure 8c shows a case where the quadratic does fit well.

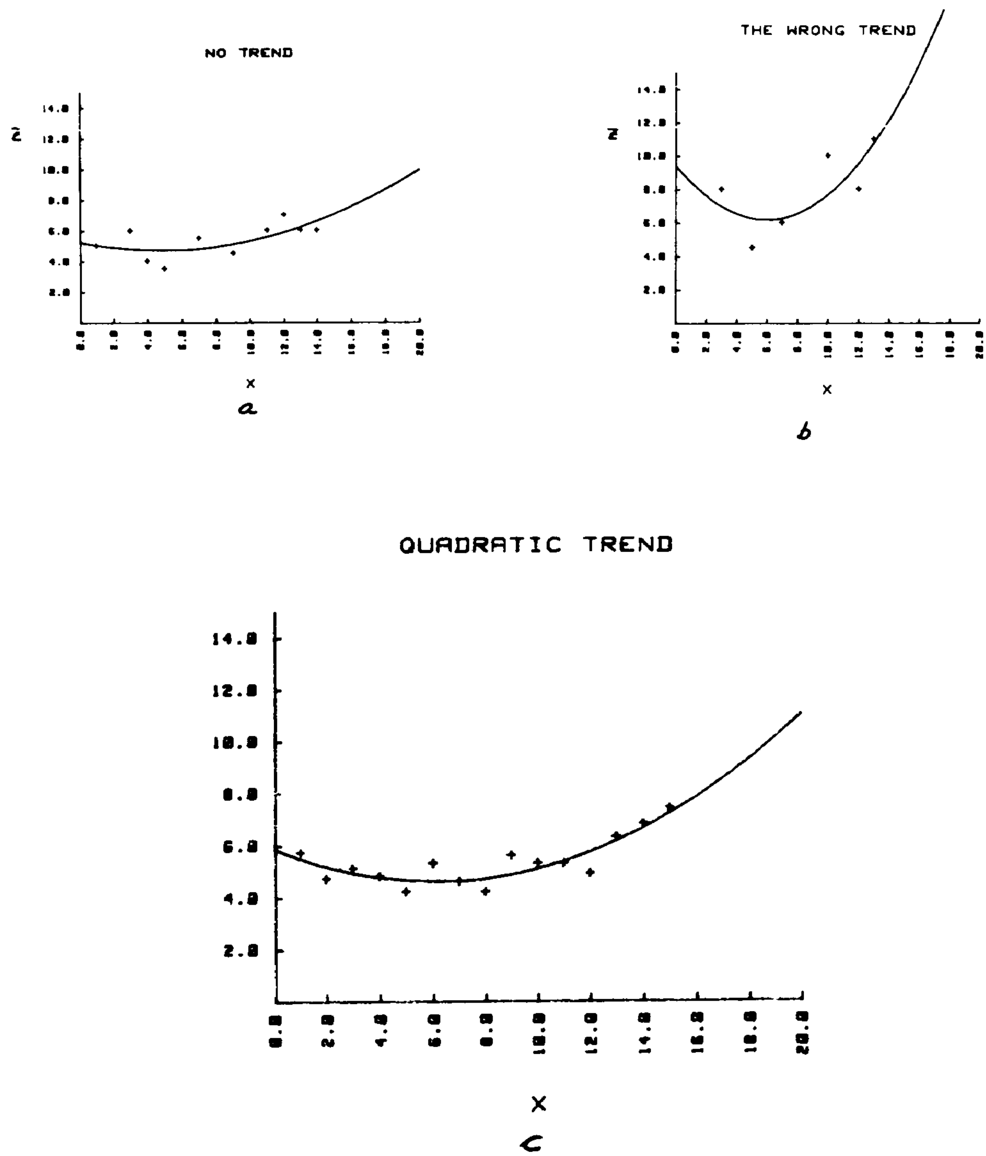

FIG. 8. — Trend fitting.

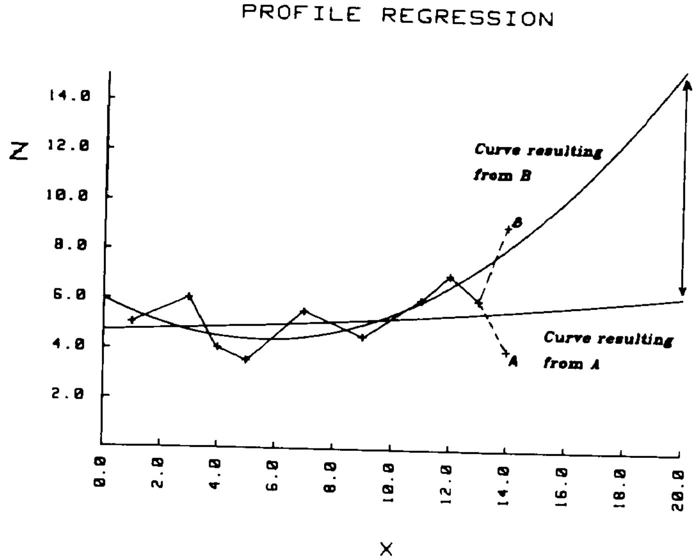

FIG. 9. — Problems in extrapolation.

Figure 9 illustrates the problem of sensitivity in extrapolation when fitting a polynomial. In this case only the *last* point has been changed. Note the drastic change in the value extrapolated at X = 20. Extrapolation is *safer* with distance-weighted averaging since the extrapolated values are bounded by the extreme values in the neighbourhood.

Clearly, area-modelling methods *also* suffer some drawbacks. All too often the surface is far too *random* to be approximated, even locally, by an analytical function.

AREA VERSUS POINT MODELS

It will come as no surprise that one method is not clearly superior to another. Point models offer safer interpolation in areas where the bottom undulates randomly about some near-constant level whereas area models are safer where the bottom has a clear trend. Area models are cosmetically cleaner and more efficient in storage space and in execution — but these are, by and large, irrelevant issues for hydrographers who have the time and equipment to do the job right. One can't help but feel uneasy about fitting smooth functions to surfaces which are by nature rugged and unpredictable. Point models can handle the ruggedness but are defeated by trends. Clearly some combination of the two techniques might be the ticket; an area model to detect and model the trends and a point model to work on the residual surface which rides on top of the trend. This is the essence behind *universal kriging* — a topic beyond the scope of this particular paper. There are two other techniques, however, which do, to some degree, incorporate features from both area and point models, namely *triangulation* and *parallel line* techniques. We investigate those now.

TRIANGULATED NETWORKS

Many, many techniques have been developed to generate the grid estimates upon which the contour placement will be based. Literally dozens are available — each one considered optimal in one way or another by its author. Some are designed to exploit features inherent in the data itself or in the geometric structure of the sampling program. For instance, some surfaces are very smooth and slowly changing — such as our perception of a gravity surface. Others, including bathymetry, can be very *rugged*. Some surveys are very *dense* and *regular* — like that of Gestalt Orthophotography — while others, like borehole surveys, are often *sparse* and *irregular*. A technique developed particularly for one type of data will not necessarily perform as well on another, radically different, type of data. A technique developed for a large main-frame computer will not perform well — or at all — on a medium-size mini-computer. *Storage* and processing *speed* are two of the chief considerations which many designers hold supreme.

With such a variety of programs and techniques available, it is not surprising that a certain amount of controversy is apparent in the literature as to which technique is *really* the best. One of the most frequently debated items is the *grid versus triangulated irregular network* (TIN). The grid, as we have already discussed, is a square *mesh* applied over the measured surface with the non-regularly-spaced observations being used in some interpolation procedure to derive estimates at the mesh nodes. A TIN, on the other hand, *joins together* the observations into a triangulated network (see Figure 10) and *then* interpolates. New estimates inside the network are interpolated by sub-dividing the network triangles into a series of smaller *clone* triangles. Depth estimates are then made at each of the new vertices. TINs are particularly appealing to hydrographers because they *honour* the data.

We have previously argued (MONAHAN and CASEY, 1983) that honouring is an important issue. We included it in our "Musts and Wants" list as a *must*. We can now differentiate between *two* kinds of honouring : *strict* honouring and *weak* honouring. By strict we mean that the observed data points are honoured in the

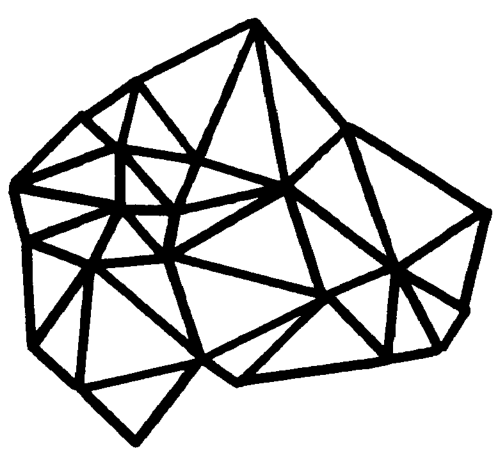

FIG. 10. — A triangular irregular network.

strict *mathematical* sense. The observations lie *on* the interpolated surface. Weak honouring implies that the observations do *not* lie on the surface but lie *so close* as to appear, for all intents, to be *on* the surface. This might seem like a minor academic point but, in fact, is quite important. The important card triangulation proponents have to play is that their technique strictly honours the observations and most other systems do not. We do *not* differentiate between strict and weak honouring, since to do so would be to unfairly categorize systems on what we feel is a minor technicality. Most gridding systems do *not* honour the observations in either sense and so, for hydrographic purposes, are suspect.

Interpolation in triangulation can be relatively straightforward. In its simplest guise, the three data points at the vertices define a planar surface. Contour location on this planar surface is then linearly interpolated (See Figure 11). The survey area can be imagined as being built-up with a network of triangular facets, like, say, a geodetic dome. The contours drawn on such a surface have a very characteristic "angular" look. This is a consequence of the discontinuity of slopes which occurs at each triangular boundary. Much more elaborate and sophisticated techniques are also available to overcome some of these limitations.

One of the most sophisticated of the triangulation algorithms is due to AKIMA (AKIMA, 1978). This algorithm applies a 5th order bivariate polynomial to the *surface* of each of the observation triangles. First and second partial derivatives are computed at each of the vertices and directional derivatives computed along each side to ensure the smoothness within the triangles and along the triangle's borders. These measures ensure that the strong contour angularity, which is so characteristic of triangulation schemes, is minimized. AKIMA's algorithm is also designed to suppress the unsubstantiated undulations or *ringing* effect which the fitting of polynomials usually involves.

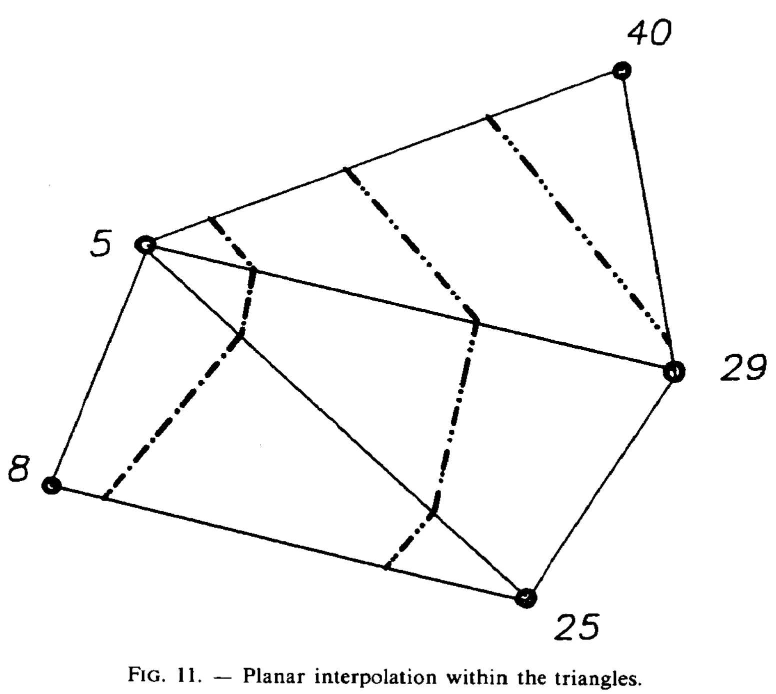

FIG. 11. — Planar interpolation within the triangles.

Triangulation systems are not without their problems. A re-examination of Figure 10 will show that many *other* networks could be established from the *same* data points. Early generation triangulation schemes did not address this problem to any great extent, so the impression has grown that these schemes do not have unique ways of defining the network. If different networks are used, the resulting surfaces can look *very* different. Primitive triangulators used the *order* that the data was *entered* as a guide to establishing the network. If the data was re-ordered, then a *different* set of contours would be derived. This is clearly distressing. Such incongruities have not helped the proponents of machine contouring in promoting the use of computers in, what is for many, the final and most visible outcome of their work. Fortunately, many researchers have been working on this problem. The result is that there now are standards for the definition of triangulation networks. The de-facto standard is known as Delaunay Triangulation (SIBSON, 1978).

One popular method for achieving Delaunay triangles is to first form a set of polygons (called Thiessen Polygons) from which the triangles can be formed. This is known as Dirichlet Tesselation (GREEN and SIBSON, 1978). See Figure 12.

TINs and Thiessen networks have attracted much attention in the recent literature and are the subjects of continued study but are relatively rare in commercial usage (SALLAWAY, 1981).

Can triangulation schemes do the job in hydrography ? Yes, in some specific cases. The vast amounts of data produced by area-mapping systems such as the *Navitronix* or the *Larsen* could be successfully contoured using a triangulation scheme *if* (and *only* if) sufficient controls were added to ensure a bias for safety. This could be feasible.

If the object was to contour an existing digital field sheet using *only* the data presented there, then a triangulation scheme would be superior to a gridding scheme. But there is a far better way to contour digital hydrographic data and that is by using *all* of the observed depths — not just the ones portrayed on the field sheet. We investigate one method now which does just that.

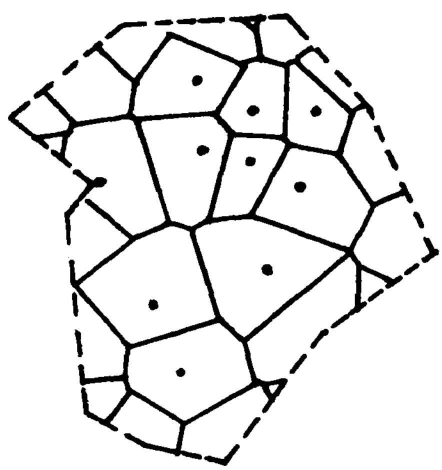

FIG. 12. — Dirichlet Tesselation.

PARALLEL-LINE DATA

All the contouring techniques we have discussed above assume that the data is *irregularly* and *randomly* distributed throughout the survey area. Indeed, most users of machine contouring packages have data which is in this form. *Hydrographic* data, on the other hand, is blessed with a very important characteristic — *continuity* of information along the sounding lines. This feature can, and should, be exploited.

Most users of contouring packages have to be satisfied with interpolated contours. In our case, however, the position is *known* along the sounding lines because we *measured* it there. This measured position can then be used to *anchor* the contour's position as it intersects each sounding line. These points, commonly known as *contour intercepts,* are as well known as any of the soundings we normally plot on our field sheets. Thus only the contour's path across the inter-line zone has to be interpolated. The use of contour intercepts is the modus operandi of geomorphologists who create bathymetric charts such as edition V of GEBCO. Its use in conventional hydrography can be traced to QUIRK (QUIRK, 1966).

We can also exploit the parallel line nature of our sounding lines in determining the contour. To control the contour's position *between* the lines requires an interpolated grid of one kind or another. The sounding lines can be used to establish this grid. Consider the situation shown in Figure 13. In this case we have a series of sounding lines crossed by a set of uniformly spaced parallel lines. These lines will form the *column* lines in the regular grid we are about to construct. At each point where the column lines intersect the sounding lines the digital data record is searched for the appropriate depth associated with that position. These cross-over points are called the *intersection nodes.* These nodal soundings are then used in the generation of the grid estimates.

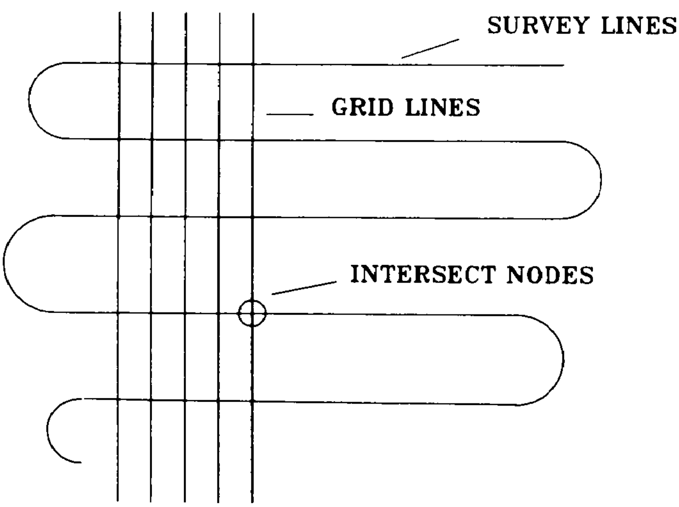

FIG. 13. — Gridding from parallel survey lines.

Both contour intercept honouring and parallel-line gridding are incorporated in a new contouring package developed by the Pacific Region of the CHS in co-operation with Barrodale Computing Services of Victoria, B.C. The package is known as the *Hydrographic Contouring System* (HCS). An overview of its interpolation procedure is shown in Figure 14.

In Figure 14a a grid has been placed over the sounding lines and the soundings extracted at the intersection nodes. For simplicity, we are showing one grid line between each sounding line. In practice this is a variable. Typical values would range from one to four.

We now examine a one-column *section* in the next diagram (14b). Here the actual sounding values are plotted as continuous straight lines. We now fit a special type of polynomial to these values. This function, known as a *cubic spline*, fits the observed data *exactly* and yet retains smoothness, throughout the curve. Estimates are then made, using the spline, at the uniform interval corresponding to the intersection of the *row* line crossing the *column* lines in the grid. These estimates are indicated as dashed lines in the section. This spline fitting and interpolation is then carried out for each of the column lines of the grid. In this way the grid estimates are established.

The next diagram (14c, d) illustrates how the course of the contour is controlled *within* each grid cell. Recall that, on the sounding lines, the contour's position is fixed. It is only in the inter-line zone where position estimates are required. Here we show a block of 16 depth estimates corresponding to 9 grid cells. We use the 16 values to fit a *surface* from which we will guide the contour in the *central* cell. Again we fit a spline — this time a bi-cubic spline since we are now working in *two* dimensions. A much finer grid (14e) is then established on the bicubic surface within the central cell and depth estimates made for this fine grid. This procedure is continued throughout the entire grid. The contour is then strung through each of the grid cells by allowing or denying access to the cell's interior.

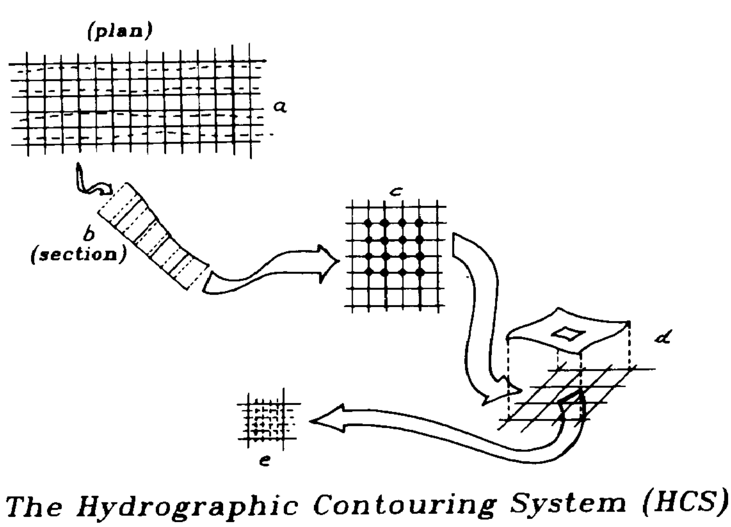

FIG. 14. — The HCS interpolation scheme.

The contour is constrained to pass through the contour intercepts. The data is honoured in the *weak* sense. In fact, the honouring is to the nearest fine-grid intersection. Since the fineness of this grid is a variable, the degree of the honouring is also variable. In practice, the grid fineness is about 1 mm. More information on the HCS is given in CASEY *et al.*, 1984.

SUMMARY

We have given here an overwiew of how most of the available contour packages work. There will be detail differences that we have left out for simplicity, but we have tried to capture the *essence* of each technique.

Hydrographers and cartographers are usually reluctant to work with what they perceive as *computed* depths (*). The feeling is that these computed depths are very much second class compared to measured depths. This is true — they are. Nevertheless, the charts we produce show a continuity of information despite the fact that we seldom make continuous measurements throughout the survey area. So, whether we like it or not, some form of depth computing is built-in to our procedures. The fact that such interpolation is done in the head of some experienced professional does not sanctify the result — a guess is still a guess. To mechanize this guessing — that is the objective of machine contouring. If care is taken in the selection of the interpolation algorithm and in the principles laid out for its implementation, then perhaps contouring of hydrographic data can be mechanized. We will find the answer to this question only after an extensive period of experimentation.

One point cannot be over-emphasized. The most sophisticated and elaborate contouring system ever devised cannot improve a poor survey. The ground rules of hydrography will not change — a poor survey will produce a poor chart every time. No amount of post-survey data manipulation will ever change that.

ACKNOWLEDGEMENT

The authors would like to acknowledge Jim VOSBURGH (CHS, Pacific Region) and Dr. Pam SALLOWAY (Barrodale Computing Services, Victoria, B.C.) for their assistance in explaining the workings of the Hydrographic Contouring System.

(*) This same reluctance apparently does not hold for positions, which are smoothed; tides, which are modelled; speed of sound, which is averaged; or heave, which is filtered.

REFERENCES

AKIMA, H. (1978) : A method of bivariate interpolation and smooth surface fitting for irregularly spaced data. ACM Transactions on Mathematical Software, Vol. 4.

CASEY, M.J., VOSBURGH, J. & MONAHAN, D. (1984) : Automatic contouring for hydrographic purposes. Proceedings of "Hydro '84", the NOS 1984 Hydrographic Conference, Rockville, Md, USA.

DRAPER, N.R. & SMITH, H. (1981) : Applied regression analysis. John Wiley.

GREEN, P.J. & SIBSON, R. (1978) : Computing Dirichlet Tesselations in the plane. *The Computer Journal,* Vol. 21, No. 2.

MONAHAN, D. & CASEY, M.J. (1983) : Contours and contouring in hydrography — Part I : The fundamental issues. *Lighthouse, The Journal of the Canadian Hydrographers' Association,* Edition No. 28. Also in : *The Intern. Hydrogr. Review,* Vol. LXII (2), July 1985.

OLEA, R.A. (1974) : Optimal contour mapping using universal Kriging. *Journal of Geophysical Research,* Vol. 79, pp. 695-702.

QUIRK, A. (1966) : Accenting the contour. Unpublished manuscript, Canadian Hydrographic Service.

RIPLEY, B.D. (1981) : Spatial statistics. John Wiley.

SALLAWAY, P. (1981) : A review of digital terrain modelling applied to hydrographic charting activities. *Lighthouse, The Journal of the Canadian Hydrographers' Association,* Edition No. 24.

SIBSON, R. (1978) : Locally equiangular triangulations. *The Computer Journal,* Vol. 21.