Accuracy of Terrain Elevation Modelling

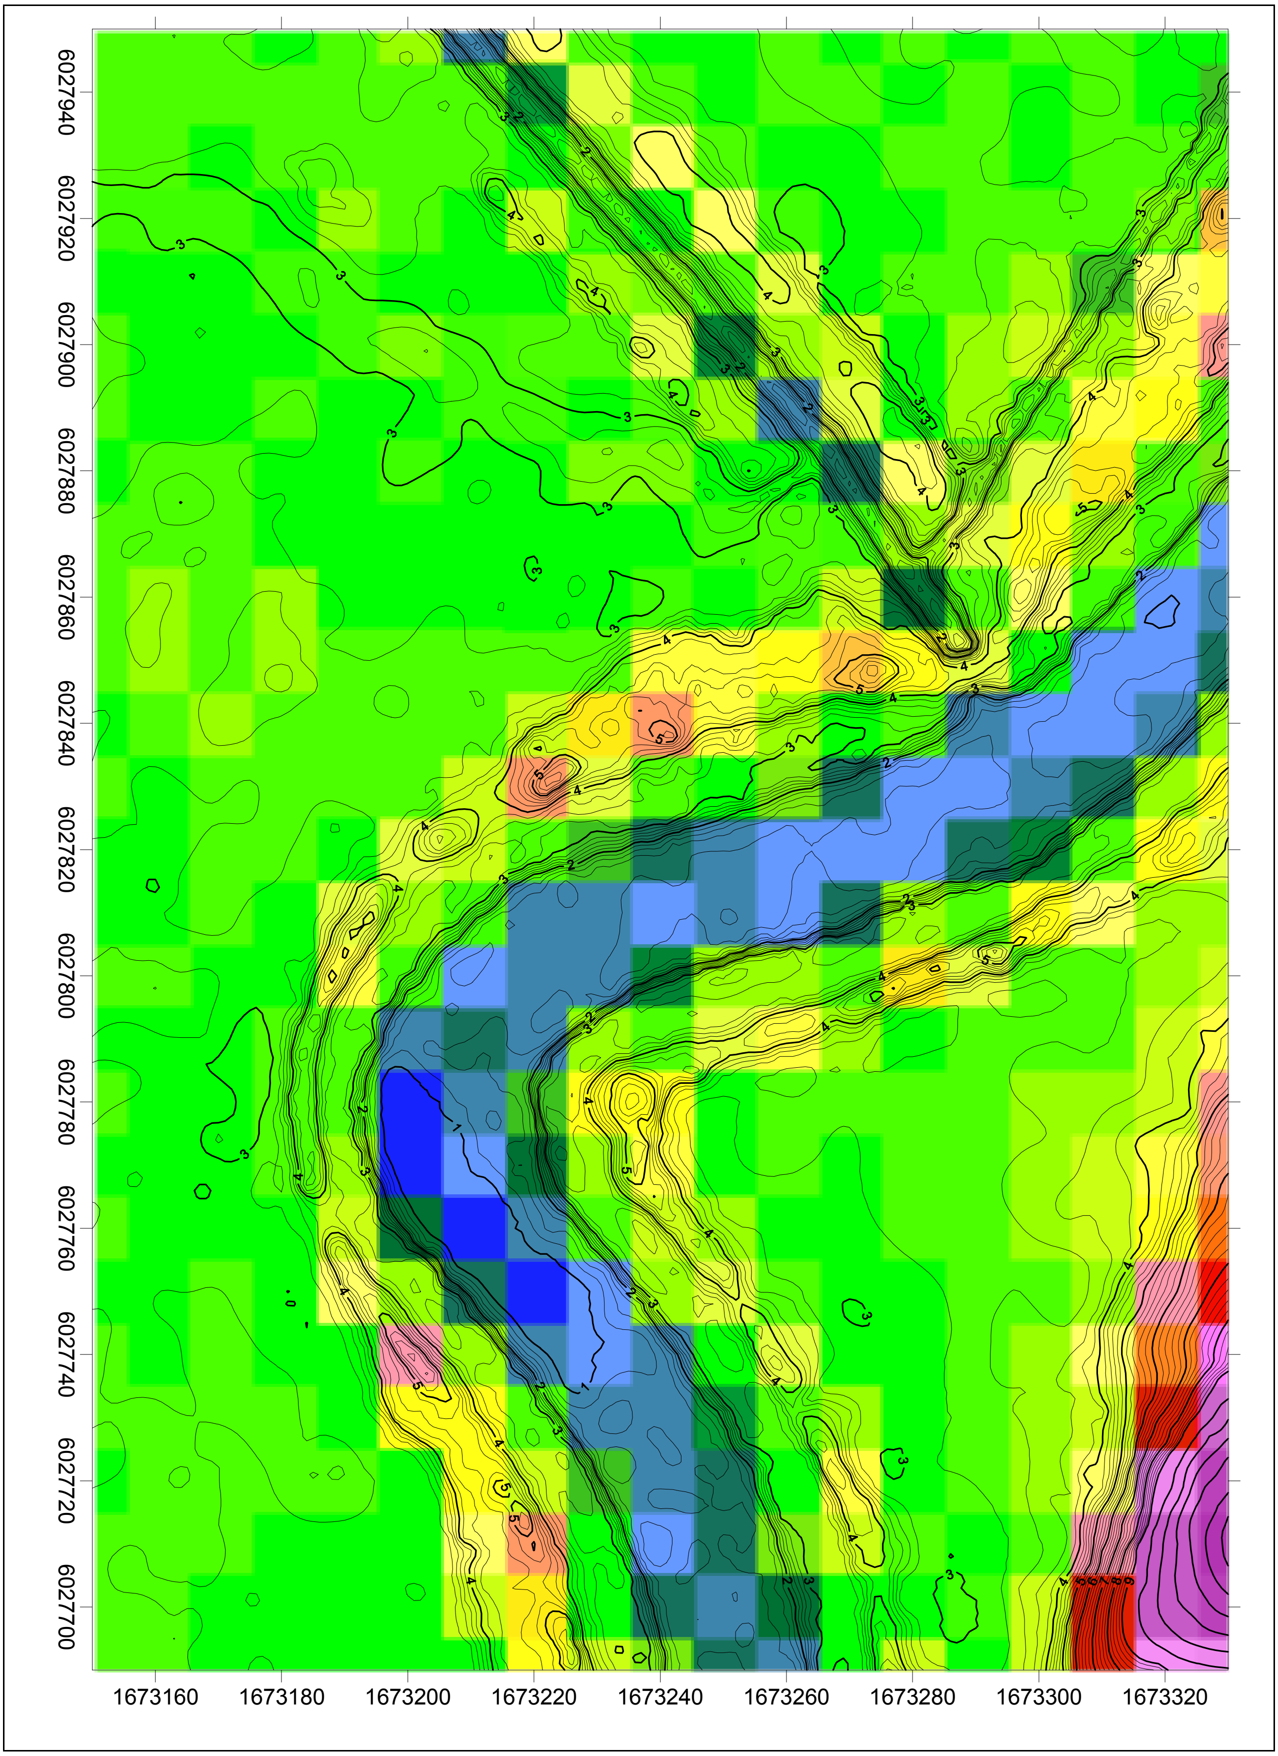

Accuracy of Terrain Elevation Modelling

By A.G. Barnett and H.L. MacMurray

Barnett & MacMurray Ltd

May 2008

Cover Picture: Comparison of Kriging-based Grids at 1m (Contoured) and 10m (Block Filled) Resolution

Accuracy of Terrain Elevation Modelling

1 Terrain Data Input

Primary terrain elevation data is now commonly derived from aerial Lidar survey, with secondary sources from older aerial photogrammetry or ground-based survey. For hydraulic modelling, such secondary terrain data is still essential in all zones hidden from Lidar penetration, including pipes, closed conduits and beds of open channels carrying permanent flow, because Lidar reflects from any water surface. Ground-based survey is also essential for structures featuring steep or vertical walls such as lined cuttings, especially in deeply incised channels (see Figure 1).

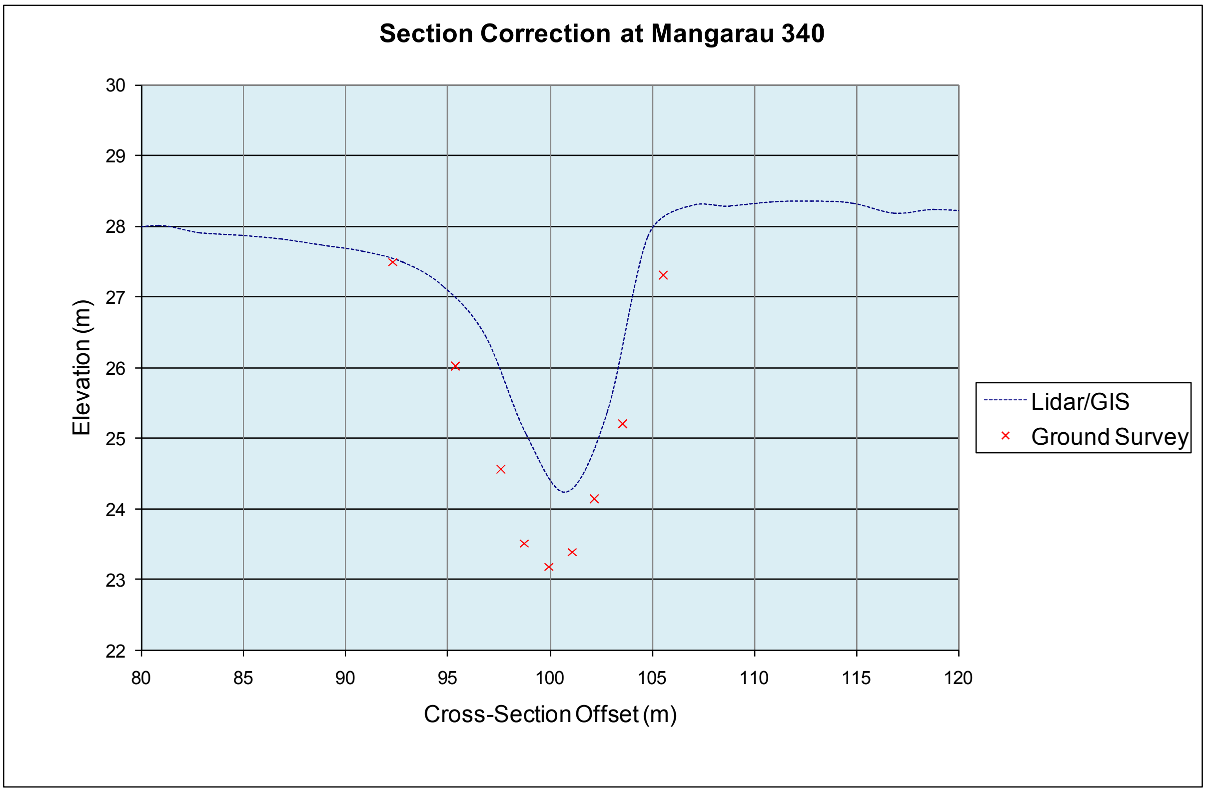

Figure 1: Correction of Typical Incised Cross-section Using Ground Survey

This Figure illustrates the considerable underestimate of flow capacity which would result from sole use of the Lidar survey to estimate the cross-section. As well as the reflective effect of water flowing through the channel at the time of survey, heavy masking vegetation is also common in the channel sides. Also Lidar beams are usually slightly oblique, and therefore cannot “see” right into the bottom of a channel in the shadow of high banks on each side.

A further source of essential secondary data is terrain changes after the date of Lidar survey, including those not yet commissioned but under evaluation for planning purposes. Studies using historical records may also require terrain changes before the date of the Lidar survey to be taken into account.

An advanced hydraulic modelling system must make provision for collating data from all these sources into a single coherent model. Further, such collation must be done efficiently to avoid needless costs, so should keep in view the accuracy of the relevant hydraulic assumptions.

2 Lidar Data "Thinning"

Figure2 shows a typical floodplain survey area, plotting with a small red cross the position of each "raw" data point supplied.

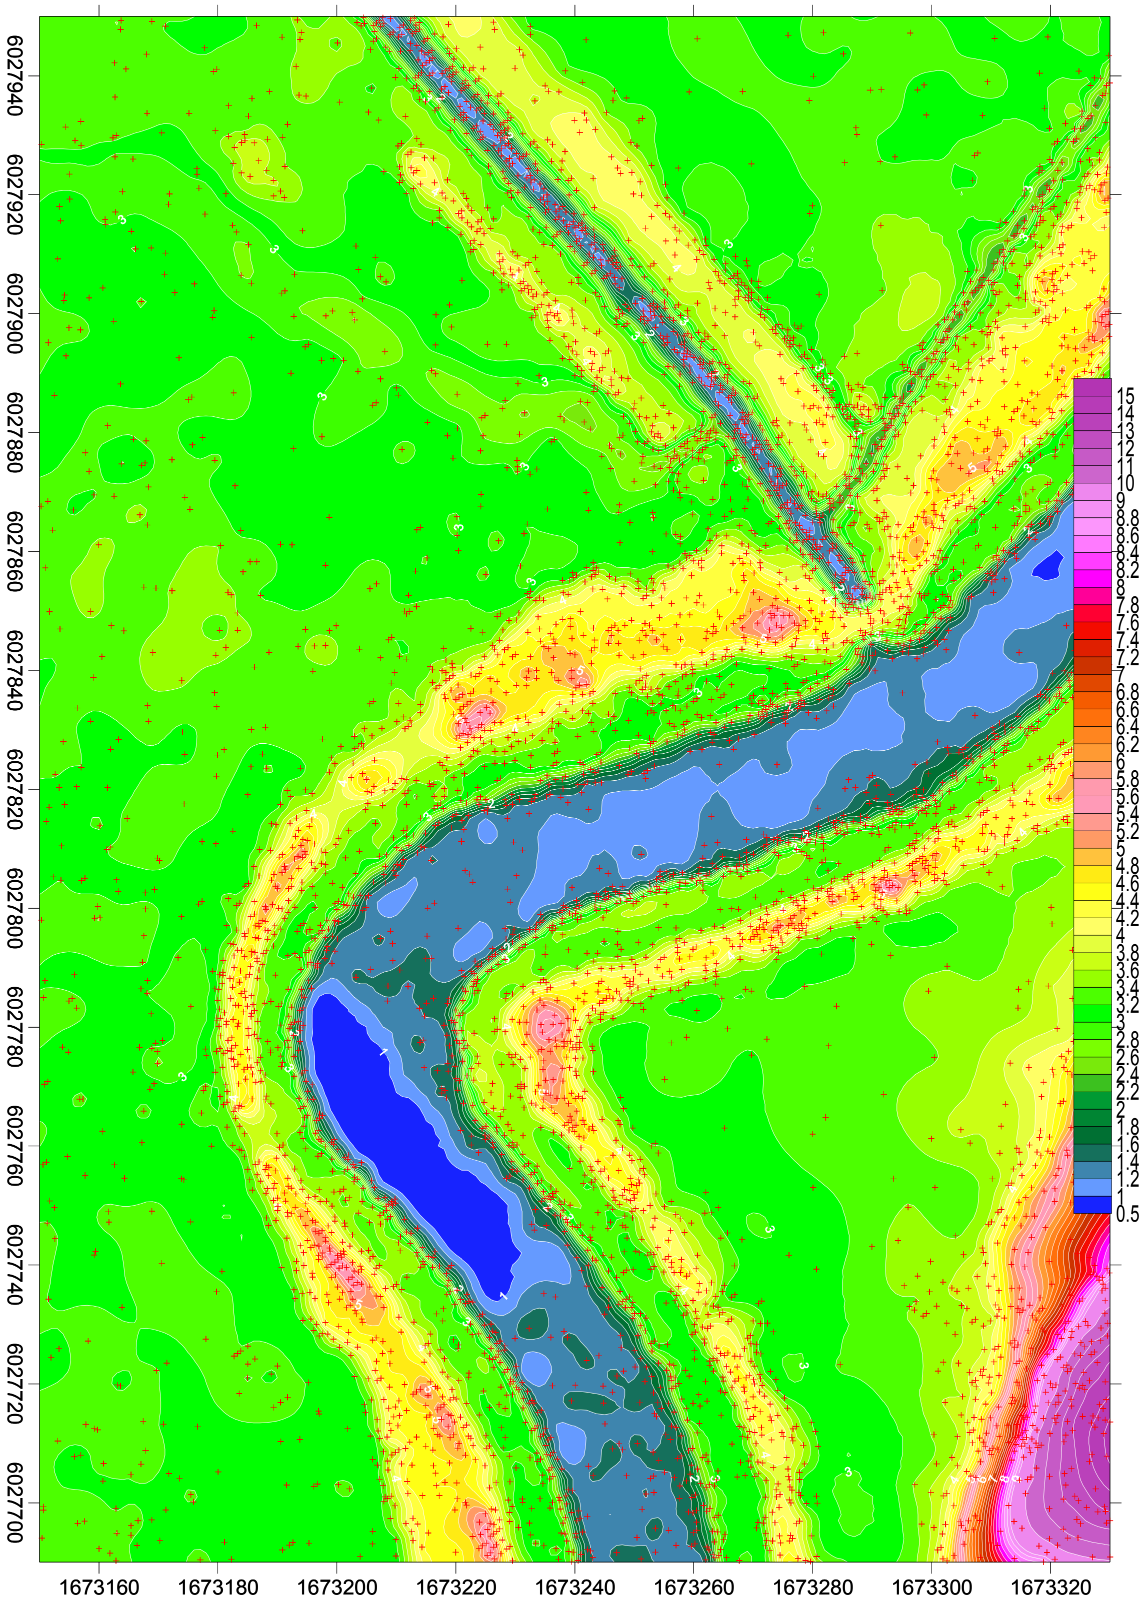

Figure 2. Thinned Raw Data (XYZ format)

Clearly, Lidar data is supplied to territorial authorities in partly processed form, as the density of points is far greater in some areas than others.

The background map shows the processed terrain data as contoured by the Surfer® (Golden Software Inc.) package at high resolution (see details in Section 3) so that the variations in point density can be related to the terrain. The contour levels are colour coded according to the elevation scale provided, with the river appearing in the blue tones, the low swampy floodplain areas in greens, the flood embankments in yellow to orange, and adjacent hills in purple tones.

As can be seen by inspection, the points are close together in areas of higher gradient such as the sides of embankments and relatively far apart in flat areas such as the low-lying drained swamps which make up the majority of the map. Provided that all significant high-gradient areas are large enough to return at least enough original Lidar points to be detected by the gradient algorithm, it would seem that this "thinning" of surplus original points will greatly reduce the size of raw data files with only a minor degradation in the resolution of features of interest, especially embankments. This improves the efficiency of working with the terrain data, as smaller files are more easily handled by computers.

In summary then, the supplied "thinned raw" (XYZ) data seems likely to contain information about embankments which is of considerably higher resolution than the average point density, with a corresponding reduction in the point density in flat, featureless areas. However, although this higher resolution evidently applies also to channel sides, the XYZ data resolution of channel cross-sections must still be regarded as unreliable on the evidence illustrated by Figure 1.

3 Terrain Surface Reconstruction

Hydraulic models work from cross-sections, which are linear in plan and therefore cannot be drawn by connecting Lidar measurement points which are scattered as shown in Figure 2. A terrain surface must therefore be reconstructed through these measurement points, and a mapping package such as Surfer offers a range of techniques to achieve this. Of these, two are selected for detailed comparison:

1. Kriging. This is the method recommended by Surfer for general purposes, and default settings are rated as "quite effective" with most data sets. The main drawback is given as slow computation because a large number of XYZ data points may be found close enough to affect the Z value fitted at a given set of coordinates X,Y.

2. Triangulation with Linear Interpolation. This is a fast method, as the Z value fitted at any given set of coordinates X,Y is affected only by three adjacent XYZ data points. This method is therefore widely used in GIS packages, but this strength is also a weakness in that information from all other adjacent points cannot be used, even if the excluded points are almost equally close to the given X,Y position.

Both methods can be used to represent the reconstructed surface on a regular grid. This has the advantages both of uniform spatial resolution and of economical file sizes, as only the Z values need to be retained – the X and Y values are implicitly given by the ordering of the data array.

Grid representation can always be improved by refinement of the mesh size, so the two selected methods are compared at a square grid resolution of 1 metre. Taking into account current practice with Lidar data technology, this is considered to be adequately fine to eliminate all likely problems with grid under-representation on the reconstructed surfaces (also see further discussion under Section 4).

The background contour map in Figure 2 is based on the kriged 1m grid version, and Figure 3 plots on this same background the contours of the triangulated surface represented on the same square 1m grid.

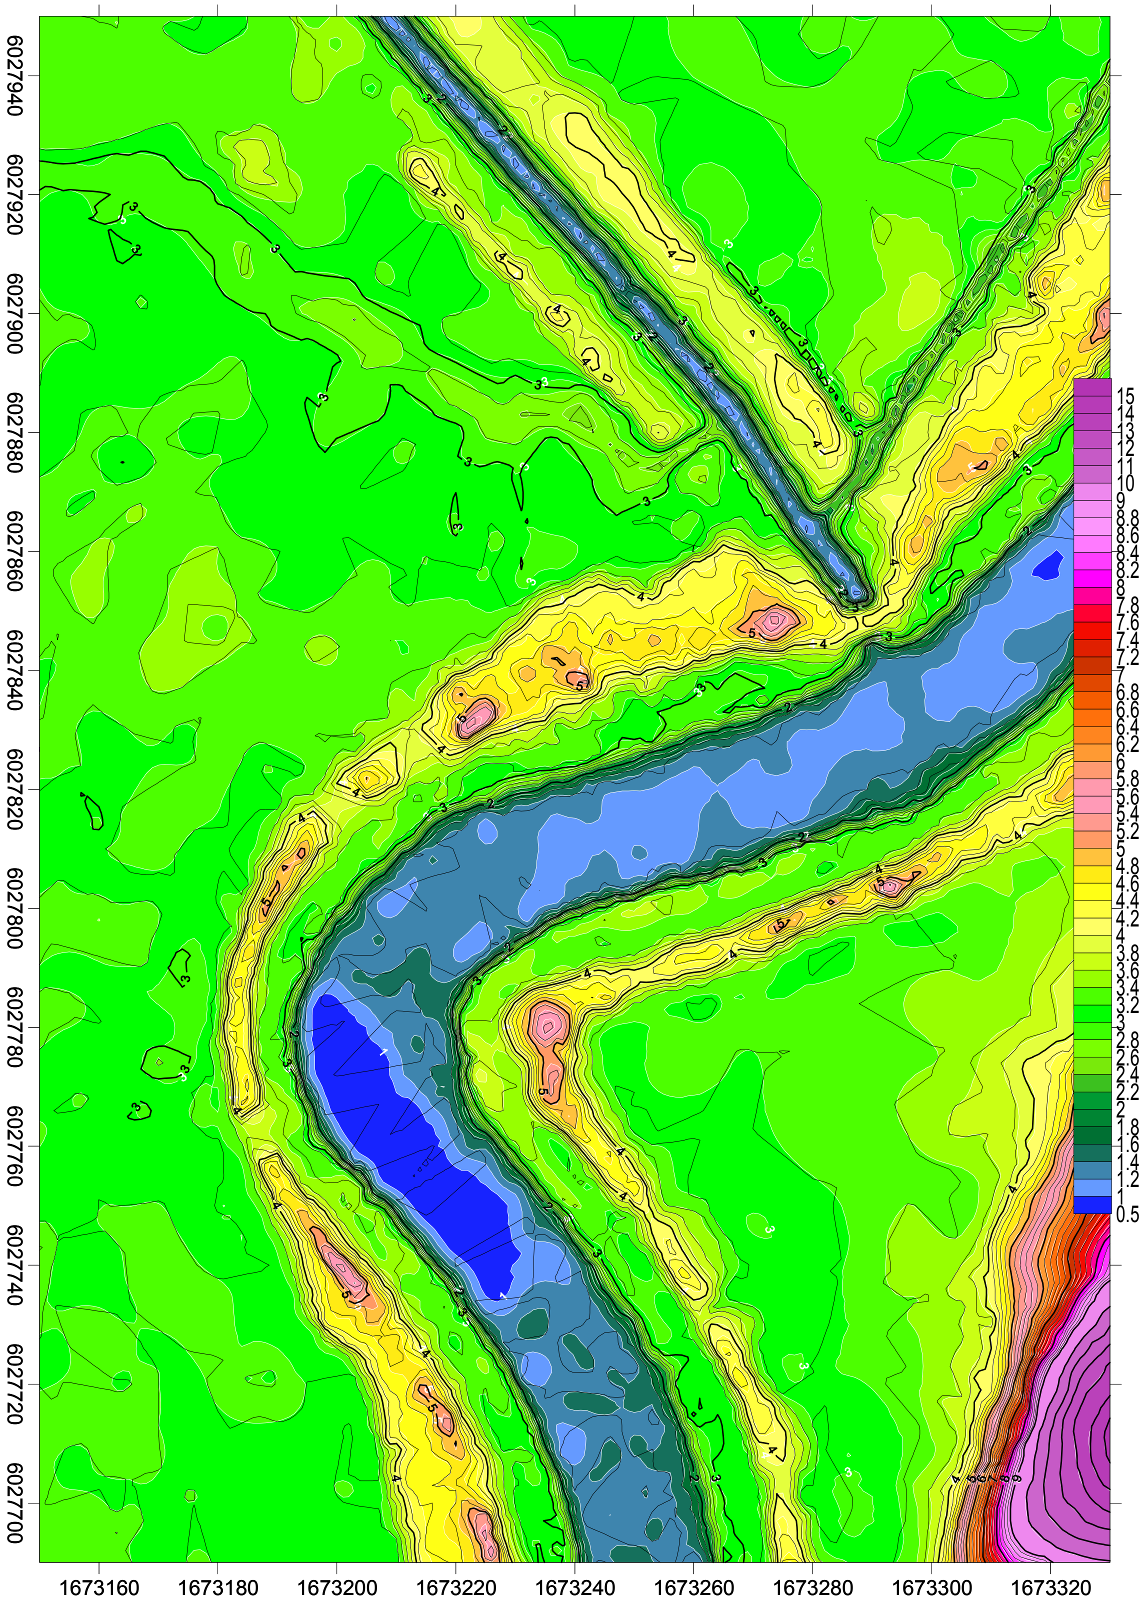

Figure 3. Comparison of Kriging-based and Triangulation-based Grids at 1m Resolution

The kriging-based contours are plotted in white with colour fill as in Figure 2, while the triangulation-based contours are overlaid in black. Significant differences are obvious, particularly in the river bed, where (see Figure 2) the XYZ data is rather sparse. The white kriging-based contours appear intuitively appealing, whereas the black triangulation-based contouring appears to offer rather improbably geometric outcomes. In the hill area (lower right) however, the contours are almost identical.

For hydraulic modelling, though, the most important difference is the contrasting treatment of undulations in the stopbank crest, in particular the low points left centre in the plot. An enlargement of this area is presented in Figure 4.

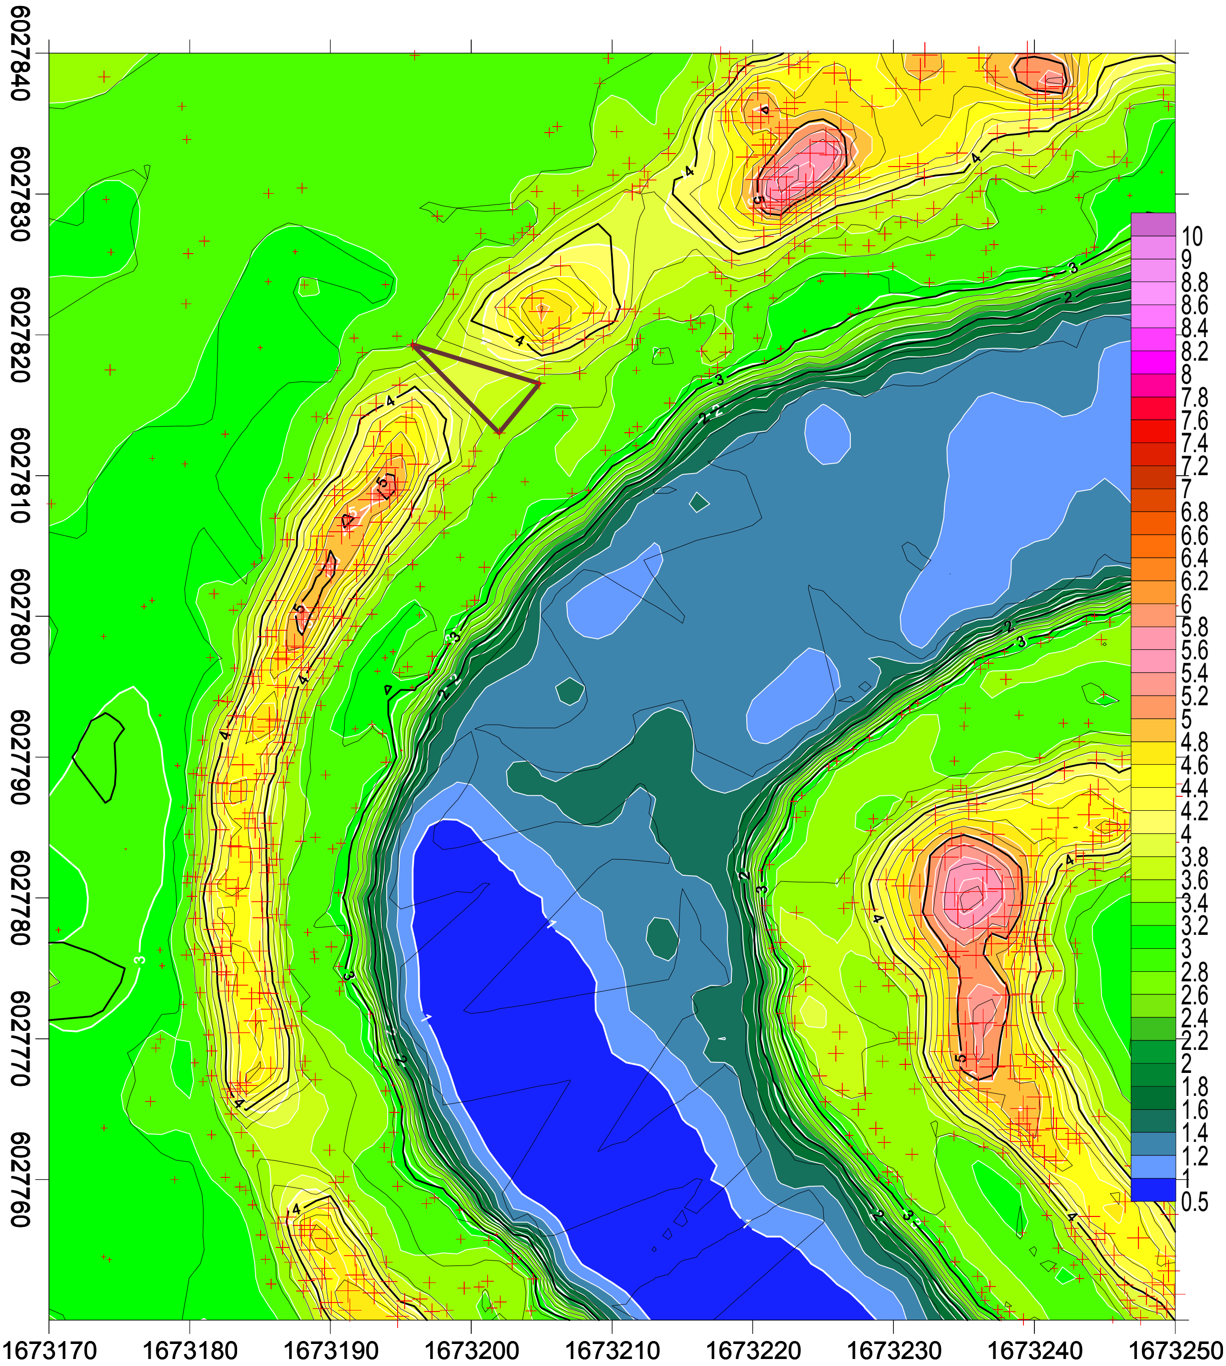

Figure 4. Enlargement of Kriging/Triangulation Comparison Showing Original XYZ Data

In addition to enlarging the data of Figure 3, the original XYZ data has been included. This is plotted as red crosses as in Figure 2, but each cross is now plotted to a square root scale starting from zero for Z=3m. This removes all XYZ data for Z<3.0m, and gives small crosses for Z<4.0m, rapidly enlarging above this value. As would be expected, the large crosses are confined within the 4m contours, and measured XYZ elevations between 3.0m and 4.0m can readily be assessed by eye by comparing cross sizes. The embankment height differences between the two grid construction methods are clearly shown by the differences between the kriging-produced contours marked by white and colour filled, and the triangulation-produced contours marked in black.

Where there are gaps in the XYZ data, significant corresponding gaps through the embankment are predicted by triangulation, whereas kriging merely predicts a saddle with lowest point slightly below the adjacent crest levels. The reason for this difference is illustrated by the triangular element marked in brown in the central gap in the picture: triangulation treats this as a plane between the relevant XYZ points at the vertices, which all have elevations of around 3.5m. The effect of using this plane to construct the surface through the gap can be seen from the obvious influence on the nearest contours running parallel to the triangle sides.

Since the two lower vertices of this triangle have elevation readings of around 3.4m-3.5m and are on either side of the embankment, the gap will also run through the embankment at this level according to triangulation, as compared with the minimum crest elevation of around 3.9m indicated by the kriging model. This gives a difference approaching 0.5m in minimum crest level!

Ground survey data along the bank crest is available as plotted in Figure 5, with the sag point under the triangle in Figure 4 marked with a red line. This data samples the crest level approximately every 50m, and supports the adjacent Lidar peak levels (see Figure 4) of 4.5-5m. However this ground survey resolution is too coarse to clearly support a minimum level at the sag point of either below 3.5m or about 3.9m, although the latter would seem more likely.

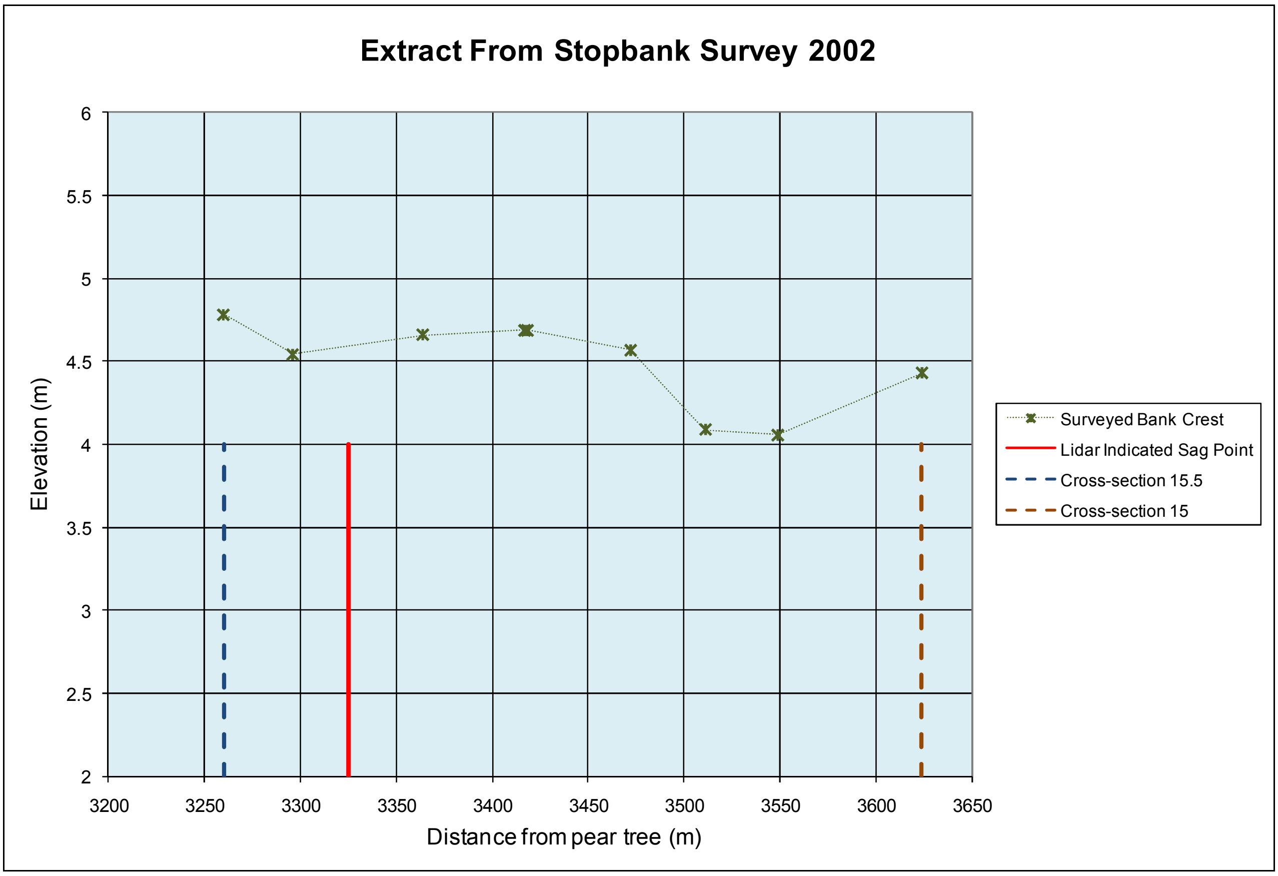

Figure 5. Available Crest Levels from Ground Survey

While no actual Lidar data remains within the marked triangle in the thinned XYZ data, it seems unreasonable to suppose that a straight line joining points measured on each side of the embankment is likely to provide a good interpolation of the crest height, so the kriging contours intuitively seem more appealing. Also, in objective terms the ability of kriging to take into account the measured elevations at all adjacent points must be favoured over the use by triangulation of only three rather arbitrarily chosen adjacent points. Therefore the Surfer recommendation of kriging as the most accurate general method of interpolation between irregularly spaced points is endorsed by the sample analysis presented here.

4 Effect of Grid Resolution

As mentioned above, kriging does have the disadvantage of slow execution, so it is unsuitable for direct use for overlaying a cross-section line on a terrain surface, an operation which must be repeated as many times as a cross-section is defined. However this disadvantage becomes unimportant if the kriging is performed only once, after which the interpolated terrain surface is represented on a closely spaced grid of points. If this grid is fine enough, interpolation between grid points can be performed by simpler methods, which can be expected to be at least as fast as the triangulation method, but with virtually the full accuracy of the kriging method.

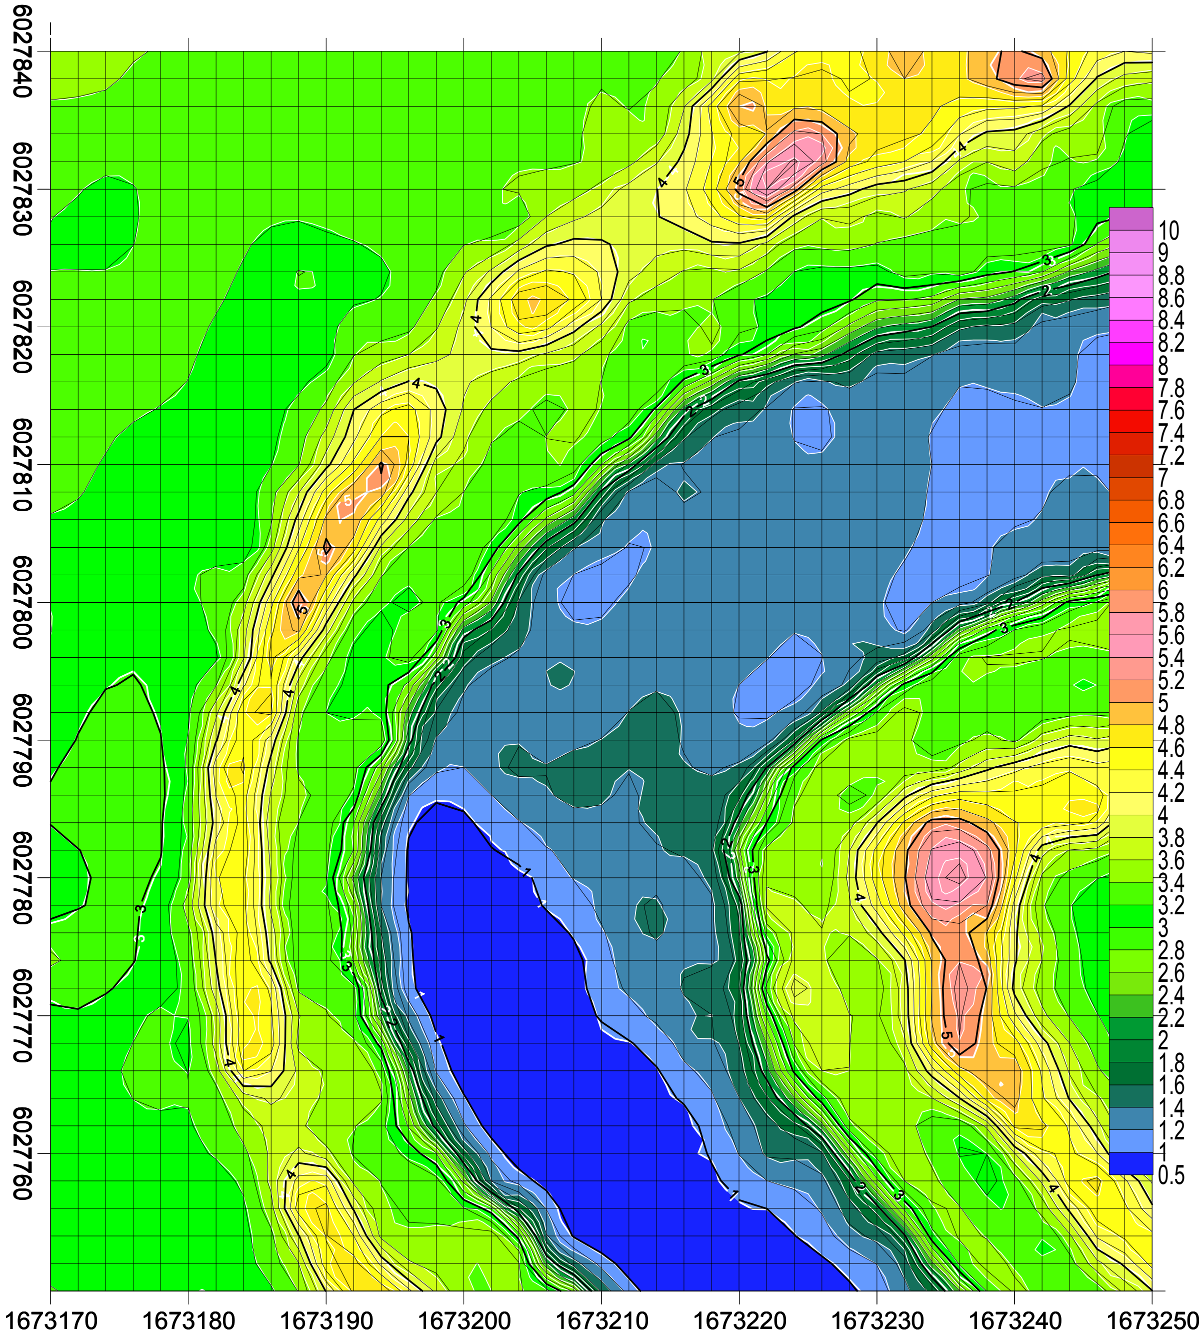

Figure 6. Comparison of Kriging-based Grids at 1m and 2m Resolution, with 2m Grid Overlaid

A square grid is attractive, as a relatively simple interpolating surface - the hyperbolic paraboloid - is always mathematically available within squares. This exactly fits all four corner elevations, as well as allowing for a trough or ridge crossing either diagonally or aligned with the sides. A square grid also gives the same resolution in both x and y directions (important where channels may cross in any direction), and is very economical on computer storage. This is because the grid point locations are implicitly defined by the order of listing of the data, allowing this location (X,Y) information to be discarded, leaving only the elevation data (Z) explicitly recorded. Many other processing options (such as contouring) can also be optimised for elevation data on a square grid, with corresponding improved convenience of use.

Assuming a fine square grid is preferred, the remaining question is: "How fine is fine enough to maintain full kriging accuracy?" Surfer advises "To increase the likelihood that your data are honored, you can increase the number of grid lines in the X and Y direction."

The XYZ thinned data used in this example was reported by Surfer to have a spacing (nearest neighbour distance) mean of 1.21m and median of 0.89m. This suggests that a square grid spacing of 1m should offer similar average resolution, being better in sparsely populated areas and poorer in densely populated areas. If the Surfer advice is followed, a good way to test whether the "data are honored" is to compare the contours obtained using a 2m grid with those using a 1m grid. If the differences are insignificant, the same can be expected to apply to further grid refinements, so the 1m grid (and indeed the 2m grid) will be "fine enough".

The results of this comparison of contours from kriging using a 2m grid and the previous contours from kriging using a 1m grid are therefore presented in Figure 6. As in Figures 2, 3 and 4 the kriging-produced contours for a 1m grid are plotted in white with colour fill, but this time the black contours are those produced by kriging on a 2m grid. That 2m grid is overlaid to allow comparison between the grid point positions and the contours derived.

By inspection the contours match closely, with the main discrepancies the tops of the narrow ridges which are less than 2m wide, so are sometimes spanned by a single grid point with corresponding slight underestimation of level. However, where data points are sparse the kriged surface is well represented at 2m grid resolution, so the embankment crest sag point contours are virtually identical, and these are the critical overflow points for hydraulic modelling purposes.

Use of a 2m grid is therefore justified for embankments of the type analysed here, because compared with a 1m grid the 4 times reduction in file sizes and processing times imposes little cost in reduced accuracy.

5 Further Grid Enlargement

Further enlargement of the grid size leads to rapid degradation in accuracy however, as demonstrated by Figures 7, 8 and 9 using kriging-based grids at 5m, 10m and 20m grids respectively for comparison with the base 1m grid contours.

The cover picture also uses kriging-based grid data at 10m resolution, but the difference is that Figure 8 uses contours while the cover picture plots the same data using a block-based colour fill (on the same colour scale as Figure 8) to represent the grid-based terrain data. This has the advantage that the resolution limitations of this relatively coarse grid are explicit in the presentation. The cover picture also uses black for the base 1m contours because in this case the colour fill is used for the 10m data.

For the 5m grid (Figure 7) the embankments are still recognisably continuous, and the kriging has still been able to preserve reasonable levels through the sag points.

For the 10m grid (Figure 8 and cover picture) the kriging has failed to preserve the embankment integrity on the outside of the prominent stream curve in the lower-middle part of the plots, although the flat floodplain areas are still reasonably well represented.

For the 20m grid (Figure 9) the whole stream/embankment/floodplain structure is starting to break down.

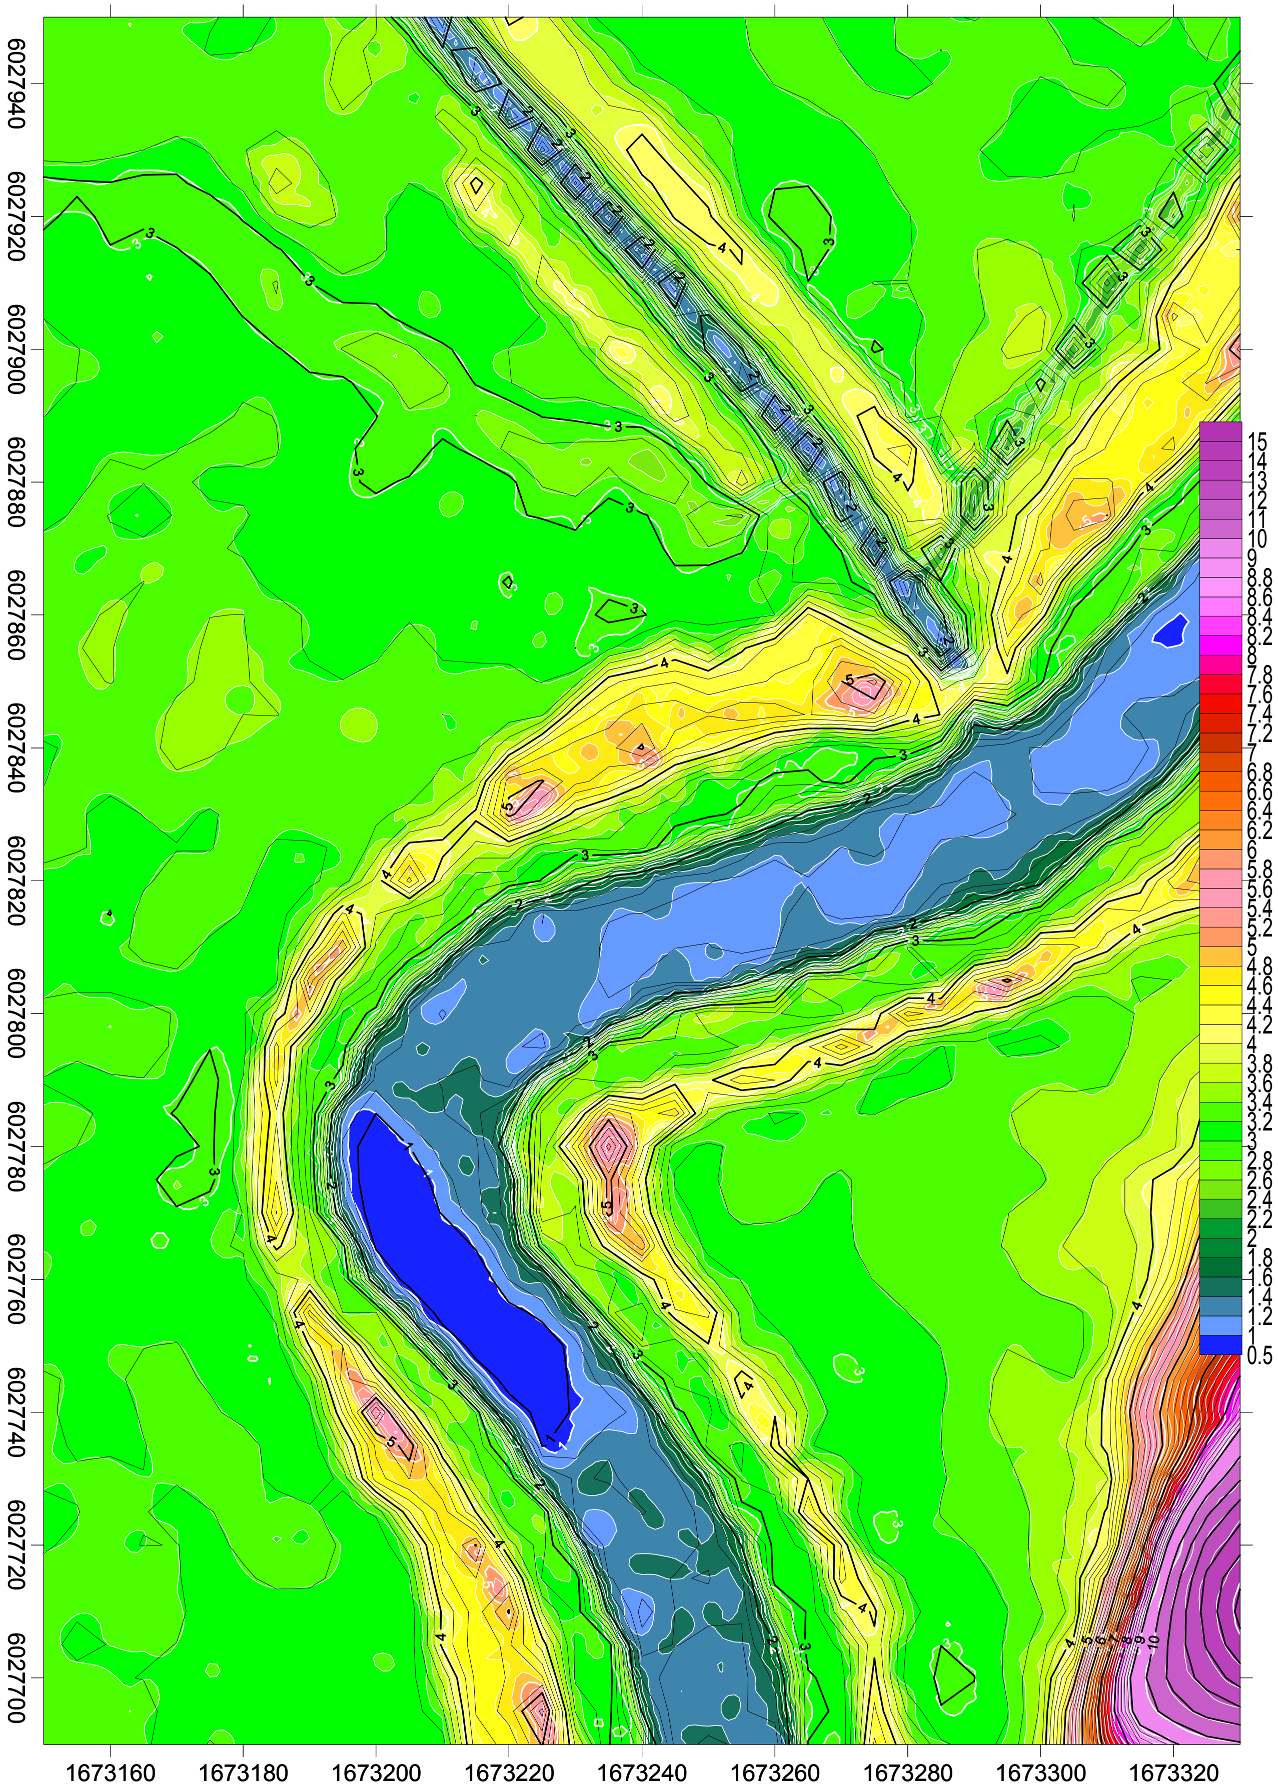

Figure 7. Comparison of Kriging-based Grids at 1m and 5m Resolution

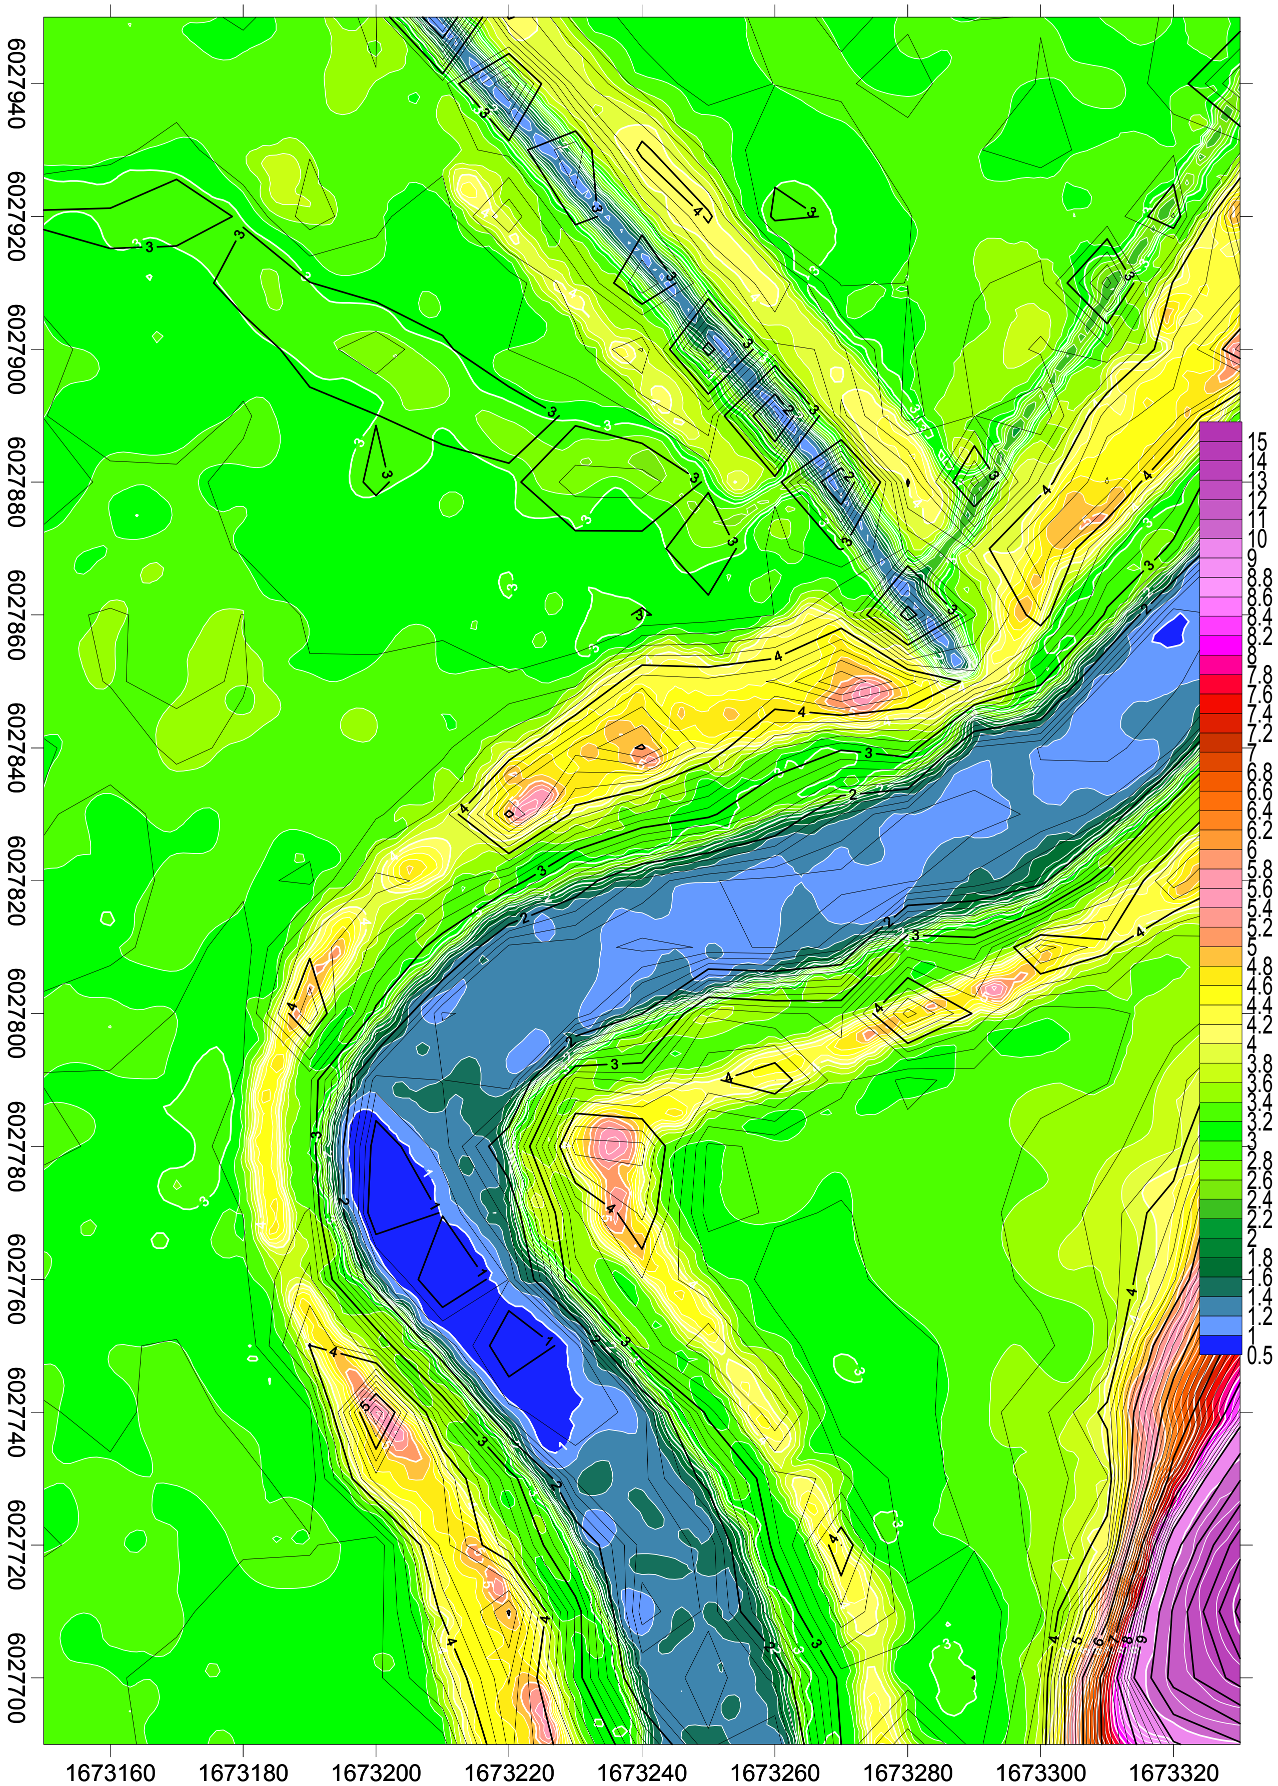

Figure 8. Comparison of Kriging-based Grids at 1m and 10m Resolution

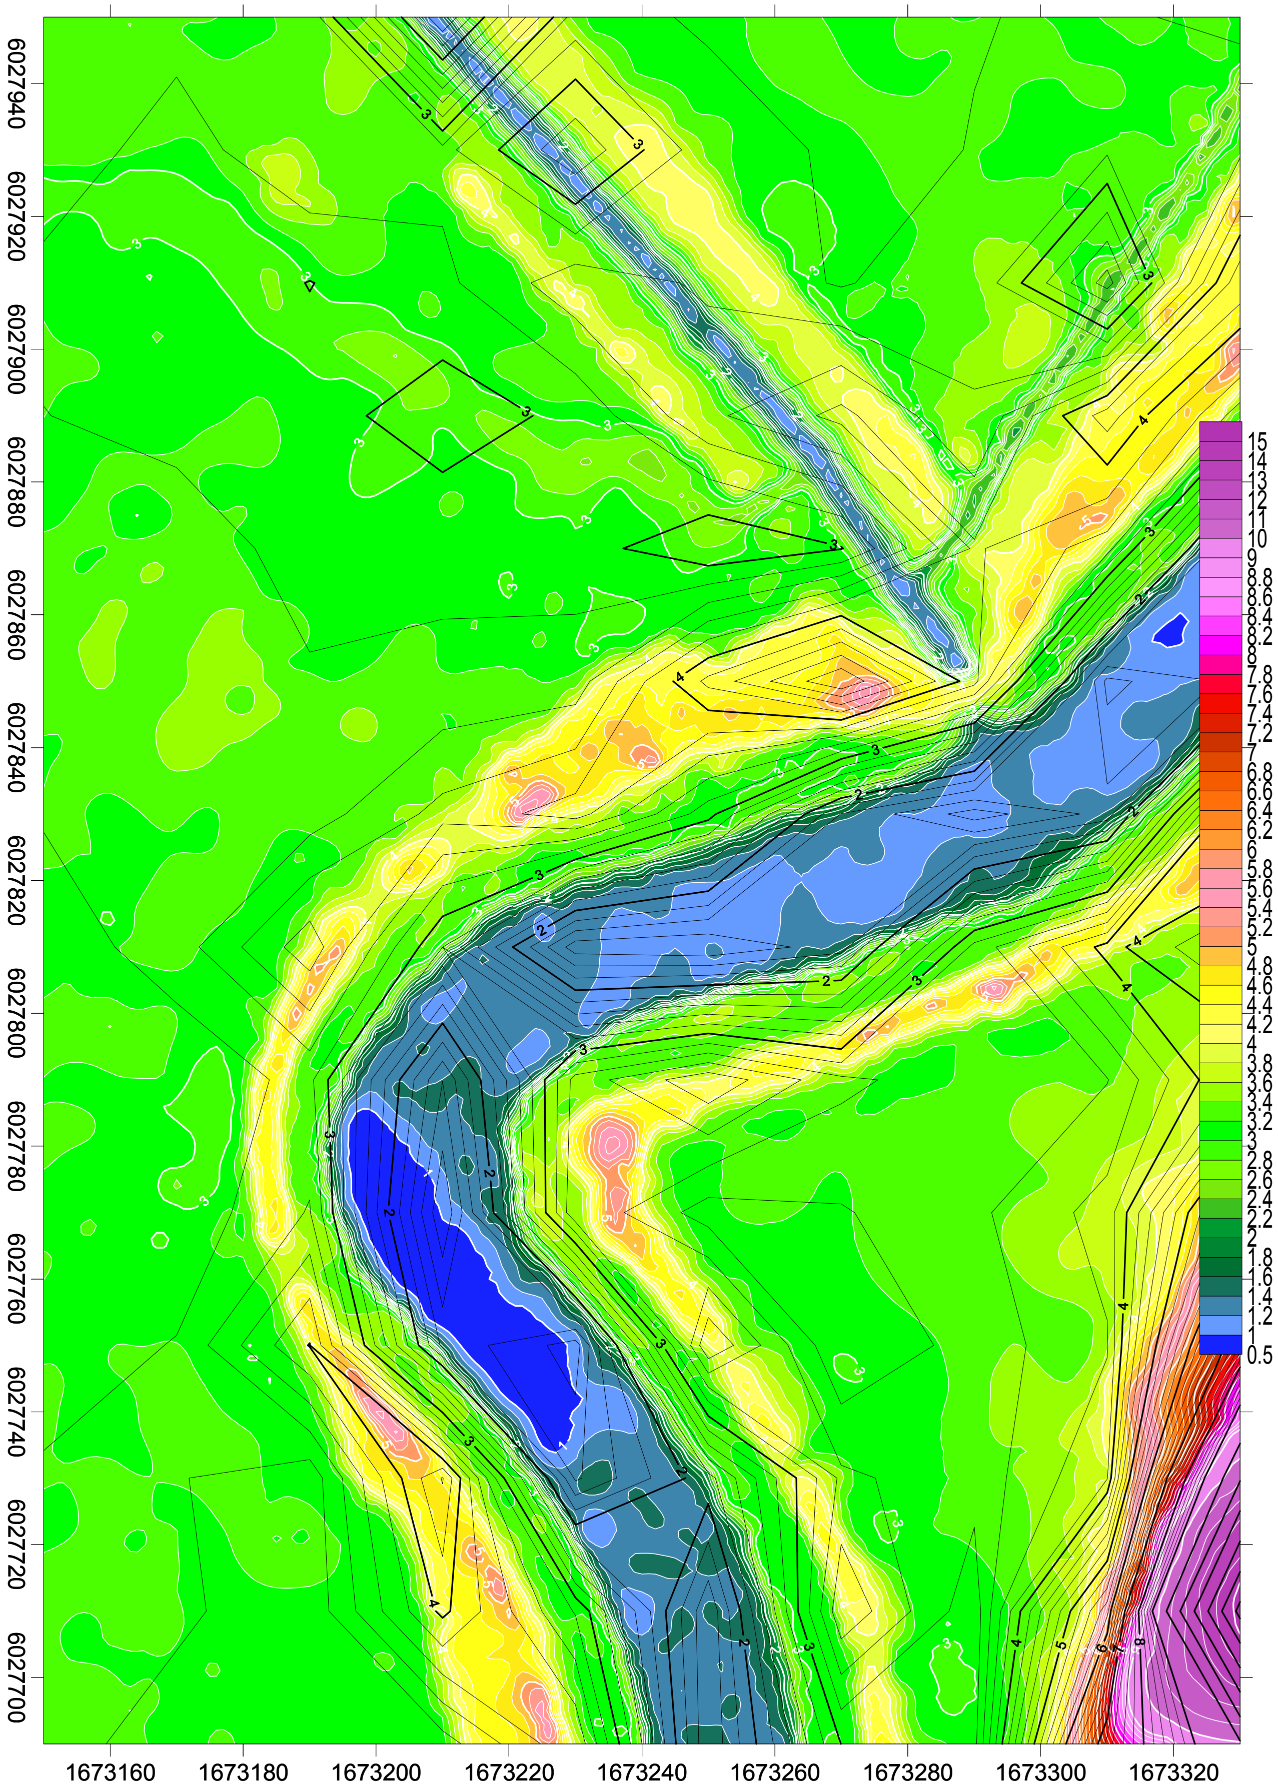

Figure 9 Comparison of Kriging-based Grids at 1m and 20m Resolution

6 Conclusions

1. Terrain survey is increasingly based on aerial Lidar technology.
2. While this is a great advance, supplementary survey data from secondary sources is still essential in many cases. For hydraulic modelling, this includes all zones hidden from Lidar penetration, including pipes, closed conduits, and beds of open channels carrying permanent flow.
3. Reliance on Lidar alone to estimate channel cross-sections usually considerably underestimates the flow capacity. Ground survey checks of channel inverts and structure dimensions are therefore required.
4. Terrain changes before or after the date of Lidar survey must always be taken into account, including earthmoving under consideration for planning purposes.
5. Data from all these sources must be collated efficiently into a single model.
6. Lidar data is supplied to territorial authorities in partly processed form, and this data "thinning" appears to remove points regarded as surplus in flatter, featureless areas.
7. Hydraulic models work from cross-sections, which require a terrain surface to be reconstructed through the thinned measurement points.
8. Various techniques are available to perform this reconstruction, and the resulting surfaces for the two most common (kriging and triangulation) can be compared at high resolution.
9. Where the thinned data is dense, similar results are obtained, but where the thinned data is sparse the triangulated surface is unsuitable, tending to create unrealistic gaps through embankments. Therefore the kriged surface is clearly preferable for hydraulic modelling.
10. Kriging uses many more data points than triangulation, so there is a cost disadvantage in slow computer execution. However this can be largely negated by performing the kriging only once, then accurately representing the resulting surface using a sufficiently fine square grid.
11. The required square grid resolution can be established by empirical trials. These show (for the typical Lidar dataset tested) that grid resolutions of 2m and 1m produce no significant difference in surface contours. This establishes that a 2m grid representation of a kriged surface adequately resolves embankments, especially at the critical sag points along the crest.
12. Further enlargement of the grid size leads to rapid degradation in accuracy, with 10m grids suitable only for flat floodplain areas and 20m grids failing to differentiate adequately between streams, embankments and floodplains.