18th World IMACS / MODSIM Congress, Cairns, Australia 13-17 July 2009
http://mssanz.org.au/modsim09

Stitching Fine Resolution DEMs

Gallant, J.C. [1] **and J.M. Austin** [1]

[1] *CSIRO Land and Water, Black Mountain Laboratories, Action ACT 2600*
Email: John.Gallant@csiro.au

Abstract: High resolution DEMs are being collected at various places across Australia according to need and funding. In some cases the gaps between existing data sets are being filled in and there is a growing need to join these DEMs together. A coherent mosaic provides much more value than an overlapping set of unconnected DEMs. This paper describes the techniques developed to stitch several DEM and bathymetric data sets together to create a continuous coherent DEM of the Murray River corridor from the South Australian/Victorian border to the Murray mouth, including the Lower Lakes and the Coorong.

A variety of challenges were encountered in this process: some features were represented well in one DEM and not in the overlapping DEM; some parts of the DEM and bathymetry did not represent real data and needed to be removed; one DEM had non-constant height errors requiring modelling of the difference. Resolving those issues required a substantial amount of manual effort. In one case, bathymetry data collected for the same area at two different times showed significant differences in channel form highlighting the dynamic nature of channel and lake bed bathymetry.

The stitching of disparate DEMs could be automated to a substantial degree if two conditions are met. Firstly, the DEMs should overlap rather than abut; this allows analysis of the differences between the DEMs in the overlap area, detection of differences and an assessment of the quality of the join. Secondly, information on the quality of the surface is needed so that informed choices can be made about which data set provides the better quality information in different places; at a minimum it is necessary to know where a DEM is supported by measurements and where it is interpolated.

We recommend that data purchasers consider these requirements when specifying data products so that new DEMs can be connected with existing DEMs, creating maximum value from new data acquisitions.

Keywords: *Digital elevation models, laser altimetry, bathymetry, stitching, acquisition*

1. INTRODUCTION

High resolution DEMs are being acquired for a variety of purposes across Australia, typically using airborne lidar or photogrammetric methods. These acquisitions are in many cases adjacent to or overlap similar data acquired at different times and by different methods. There is a clear benefit in combining these disparate data into a single coherent DEM that covers substantial areas seamlessly.

In an ideal world, combining such data sets would be simple matter of merging adjacent and overlapping data using simple GIS tools. In reality, a simple merge results in obvious artefacts at the boundary due to differences in detail or heights, and due to the different characteristics of the various methods of data collection. Creating a seamless DEM with no obvious joins requires careful analysis of the differences between DEMs in overlapping regions, adjustments to bring DEMs together and in some cases removal of inaccurate data.

Joining bathymetry (underwater topography of rivers, lakes and ocean) to high resolution terrestrial DEMs presents additional challenges due to the substantial differences in horizontal resolution and vertical precision typically encountered. Much coastal bathymetry data is referenced to a vertical datum that is different to the height datum for terrestrial data, requiring an additional adjustment to a common datum before stitching into a single data set.

There is very little published material describing the methods for stitching DEMs and bathymetry. Many articles mention joining or stitching DEMs but most do not describe the methods used to achieve this. One notable exception is Gesch and Wilson (2002), which describes a process of interpolating a surface at the boundary of the terrestrial and bathymetric layers using data from both sources to produce a seamless product.

This paper summarizes the methods used to stitch together six terrestrial DEMs (from lidar and photogrammetry) and four bathymetric models into a single seamless elevation model covering the South Australian section of the Murray River floodplain, Lakes Alexandrina and Albert, and the Coorong (Figure 1). The resulting DEM includes about 450 km of river corridor and 600 km of lake shoreline. Many issues were encountered that were treated using a variety of methods, most of which required manual manipulation of data.

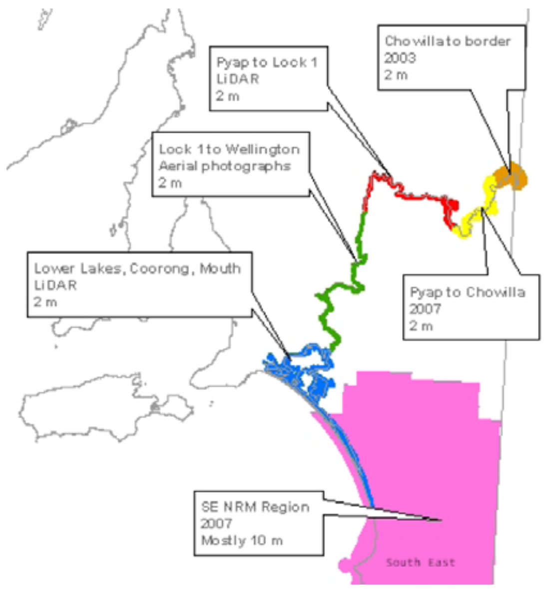

Figure 1. The six terrestrial DEMs stitched together, which were supplemented by bathymetry in the Lower Lakes and Coorong and in the river channel below Lock 1.

2. PREPARATIONS FOR STITCHING

Before DEMs can be stitched together they must of course be in the same coordinate system using the same datum and height references, and at the same resolution with the same grid origin.

In the case of the Murray River project no spatial shifts were identified. In some instances the grid origin needed to be changed by half a grid cell, which was achieved by subsampling the 2 m resolution DEM to 1 m resolution then resampling back to 2 m with the correct grid origin.

A useful first step is to calculate differences in height in the area of overlap to assess the degree of consistency of the two DEMs. An exact match is not expected, but the differences should be near zero on average and most of the differences should be comparable to the expected height accuracy of the two DEMs. Examples of the difference surfaces are shown in each of the cases described below. Where the difference is not acceptable, the source of the differences need to be identified and a treatment chosen to bring the DEMs into agreement.

3. SIMPLE STITCHES

The simplest case is where consistent elevations are found throughout the overlap area. The DEMs simply need to be combined in a way that avoids any visible edges. The *MOSAIC* function in ESRI ArcInfo Grid performs this function by weighting the relative contributions of the two DEM grids in the overlap area to provide a smooth transition.

4. STITCHING WITH NO OVERLAP

If the DEMs abut rather than overlap, differences cannot be calculated directly but the consistency at the join still needs to be assessed. In the one case where this was encountered in the Murray River project this was assessed by joining the DEMs then constructing contours at fine vertical resolution (0.3 m). Inconsistencies either side of the join appeared as irregularities in the shapes of the contours. An adjustment surface was created spanning 5-10 cells either side of the join and the less reliable (photogrammetric) DEM was modified to match the more reliable (lidar) DEM.

Gaps between adjacent DEMs can be filled by interpolation or substitution of data from another source (e.g. Grohman *et al.*, 2006) but it would be misleading to fill gaps wider than two or three grid cells, given the level of detail typical of fine-scale lidar and photogrammetric DEMs. Larger gaps are better left unfilled unless there is a clear operational requirement to have a void-free DEM, in which case an ancillary data layer should be provided indicating where gaps have been filled.

5. STITCHING WITH CONFLICTS IN ELEVATION

Where conflicting elevations exist in the two overlapping DEMs the relative reliability of each DEM needs to be assessed and decisions made about how much to adjust each data set. One DEM may be considered more reliable than the other, in which case the less reliable DEM can be adjusted and the more reliable data unmodified. More complex situations can require treating specific parts of the overlap in different ways.

5.1. Conflicts over particular features

The overlap between the photogrammetric and lidar DEMs around Lock 1 on the Murray River revealed two areas of inconsistency appearing in Figure 2 as the purple coloured areas.

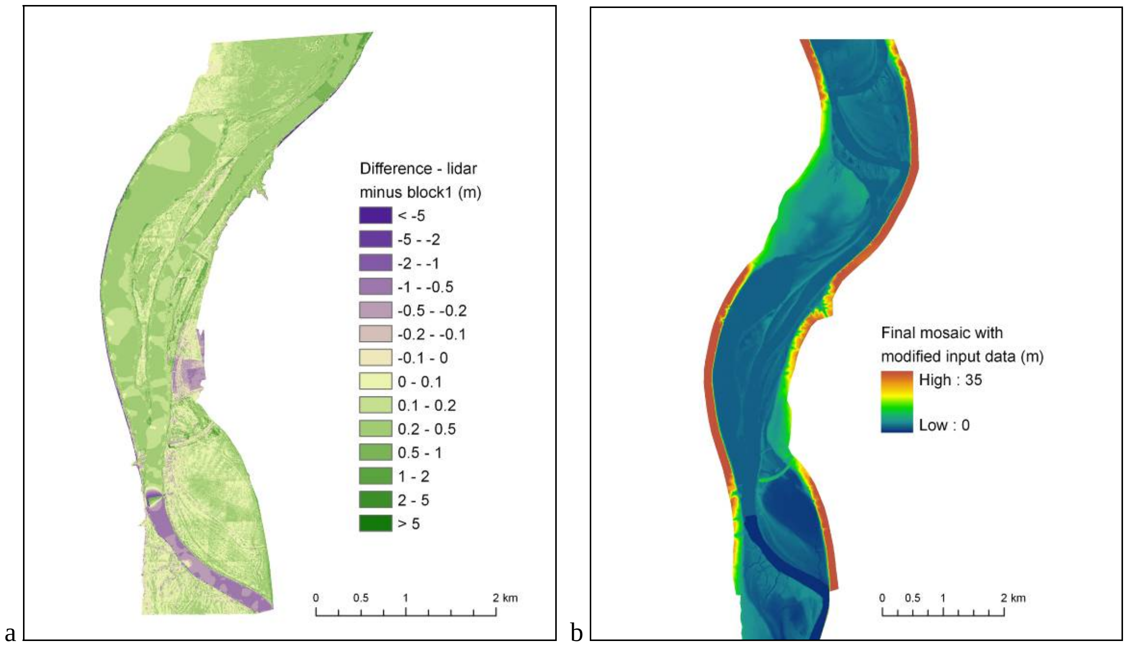

Figure 2. Stitching at Lock 1. (a) Elevation differences between photogrammetry block 1 and lidar DEMs. (b) Final mosaic of modified DEMs

Investigation of the differences determined that the area near the centre of the overlap was due to incorrect heights in the photogrammetric DEM while the river area at the southern end of the overlap was due to incorrect heights in the lidar DEM, particularly around the lock area where the height changes abruptly but the lidar did not acquire heights from the water surface.

The two areas were manually delineated and the less reliable data removed from each of the source DEMs. The modified surfaces were then merged to create the composite DEM in the overlap area that is continuous with both DEMs at the opposite ends of the overlap.

5.2. Systematic offsets in elevation

The overlap between the Chowilla and Katarapko (Pyap to Chowilla) lidar data sets showed significant height differences with the Chowilla data being 10 – 40 cm higher. The Katarapko data was found to be consistent with the DEM covering Pyap to Lock 1, and was collected more recently than the Chowilla data. On this basis we concluded that the Chowilla data were more likely to be incorrect, and this was later confirmed by other analysis showing that the Chowilla data were based on an erroneous height reference. The height difference varied systematically from north to south as shown in Figure 3(a); there were also variations in the east-west direction but not with any clear pattern. We chose to model the pattern of height variation in the north-south direction (Figure 3(d)) and create an adjustment surface (Figure 3(b)) that matched the north-south pattern in the overlap and tapered to zero over several kilometres on the Chowilla side of the overlap. Figure 3(c) shows the differences between the surfaces after modification of the Chowilla DEM.

Following adjustment of the Chowilla surface the two DEMs were combined as for a simple stitch.

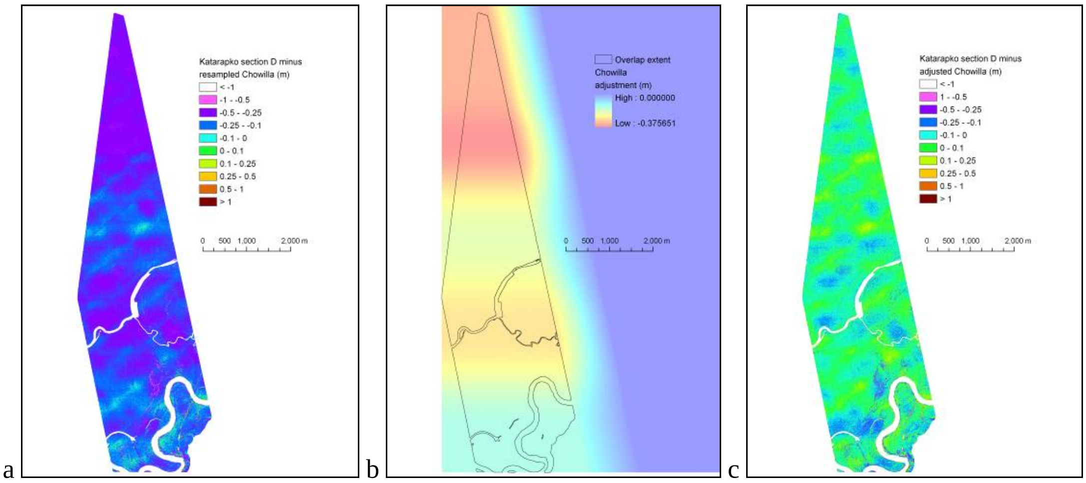

Figure 3. (a) Difference between Katarapko and Chowilla DEMs showing 15-40 cm height difference. (b) Adjustment surface for Chowilla. (c) Difference after adjustment. (d) Differences measured along north-south line through the surface shown in (a), and the model of those differences used to create the adjustment surface shown in (b).

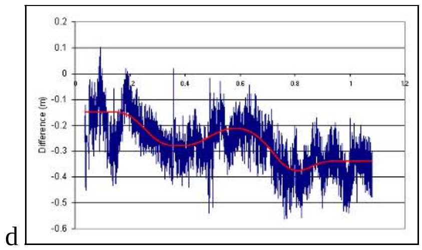

6. STITCHING DEMS TO BATHYMETRY

While the stitching of terrestrial DEMs required some careful adjustments and editing, the stitching of bathymetry to other bathymetry and to DEMs caused the greatest difficulties in this project. Three types of stitches were required: channel bathymetry to adjacent terrestrial DEMs; channel bathymetry to lake bathymetry; and lake bathymetry to terrestrial DEMs.

6.1. Channel bathymetry to adjacent DEM

From Lock 1 to Lake Alexandrina, about 200 km of river channel, bathymetry of the channel was available to join to the DEM covering the floodplain. The bathymetry was collected using sonar mounted on a jet-ski equipped with GPS positioning. This method produced usable data in water depths greater than about 1 m and the data were provided as a grid with 5 m horizontal resolution. The stitching was complicated by the existence of a nominal water height in the channel area of the terrestrial DEM that needed to be removed. Due to the higher accuracy of the photogrammetric DEM (both positional and vertical) compared to the bathymetry, the photogrammetric data were accepted wherever there was an overlap.

Figure 4 shows a preliminary join of the bathymetry and DEM and the final version after further editing and interpolation to fill gaps.

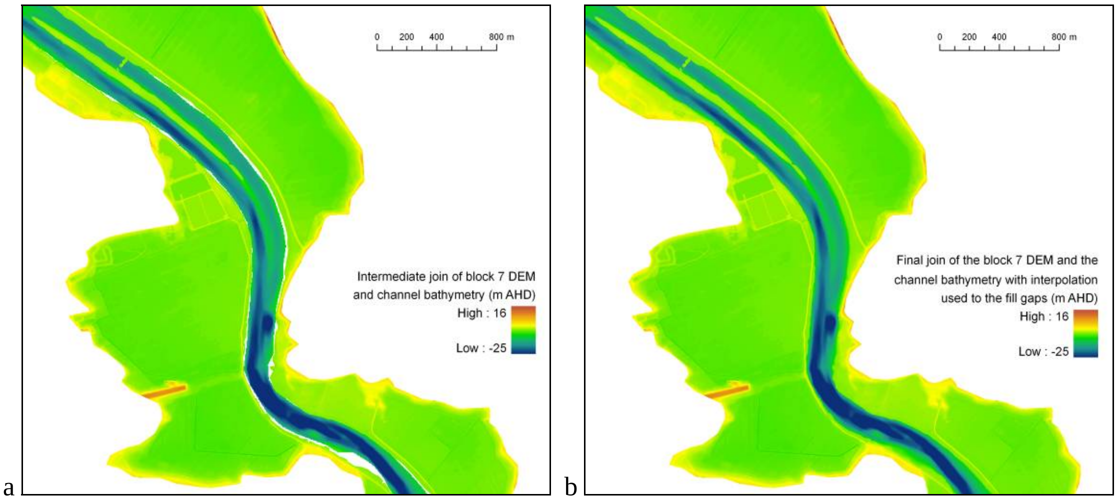

Figure 4. Stitching of bathymetry and terrestrial DEM. (a) Preliminary join after removal of spurious bathymetry data and deletion of water surface from terrestrial DEM. (b) Final stitched product.

6.2. Channel bathymetry to lake bathymetry

The region where the Murray River discharges into Lake Alexandrina was covered by several overlapping sets of data (Figure 5): lidar DEM around the river and lake (2 m resolution); bathymetry in the river channel (5 m resolution); lake bathymetry that extended partway into the river channel (100 m resolution); and a second lake bathymetry layer ("Pomanda", 50 m resolution) covering just the uppermost section of the lake, considered to be more accurate than the rest of the lake bathymetry.

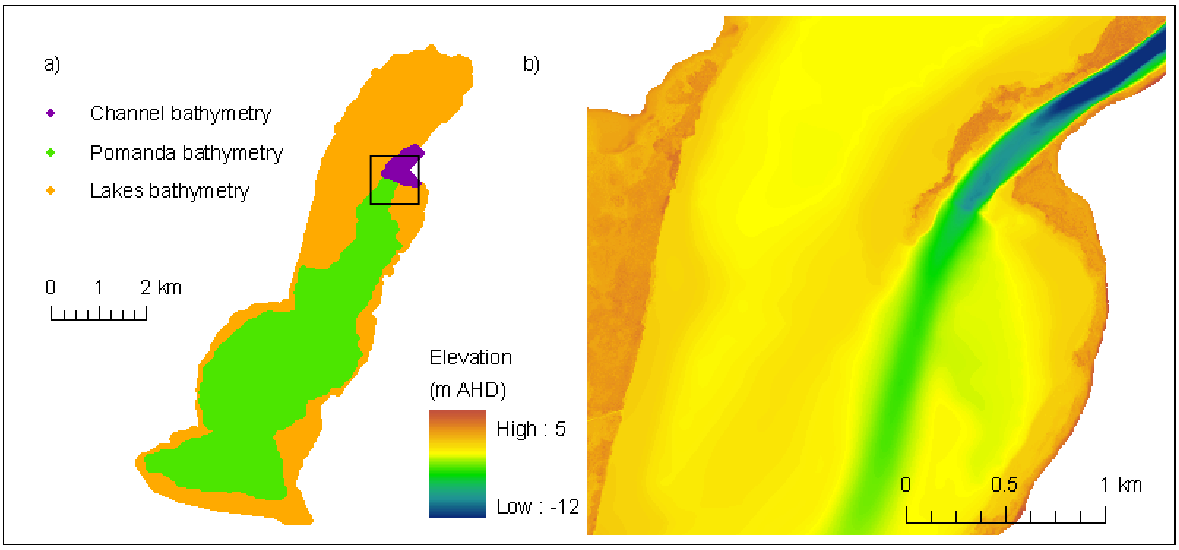

Figure 5. (a) Overlapping bathymetry data in the upper lakes area. (b) Stitched product of three bathymetry layers and the terrestrial DEM.

The multiple representations of the southern end of the channel were inconsistent to varying degrees. Some of the differences appeared to be due to actual changes in the channel bed during the time between the bathymetry acquisitions, while other differences could be attributed to the varying resolution of the data sets. Combining the bathymetry data sets to produce a coherent surface required some judicious choice of boundaries to clip the source data and some interpolation to produce a sufficiently smooth join. Figure 5(b) shows the product at the end of the river channel; the join between the more and less detailed channel bathymetry is visible but this is considered to be an acceptable result given the limitations of the source data.

6.3. Lake bathymetry to DEM

The stitching of the lake bathymetry to surrounding lidar DEM was the most troublesome part of the project, due to the large disparity in resolutions, differences in precision of the data and extension of the bathymetry data into shallow areas where no actual data was acquired. There were also some problems with the lidar data in this area: the initial data supply suffered from random areas of missing data, while a re-supply of data fixed the missing data but included areas of smooth interpolated surface that were clearly not based on measured surface heights. The raw point data were used to create a buffer which was then used to clip the DEM and remove spurious values.

Most of the lake bathymetry was provided at 100 m resolution, although some of the more detailed areas around channels and islands were at 50 m resolution. Although the bathymetry could only be acquired reliably in 1 m or more of water, the bathymetry data was extended to a nominal shoreline. The lidar was acquired at a time of low lake water levels (mid 2008) and included areas of exposed former lake bed that were also represented in the bathymetry data.

The inconsistency between the lidar and bathymetry data required removal of bathymetric data that was not based on measured data. Bathymetry data was removed where elevations were above -0.5 m AHD, or within 30 m of the lidar data. The 30 m gap was enforced to provide sufficient space to create a smooth join between the lidar and the much smoother bathymetry data; considering the 100 m resolution of the original data, this did not involve much loss of information. Figure 6 shows the input data and combined elevation model, and the difference between the combined model and the original bathymetry.

7. DISCUSSION AND CONCLUSIONS

The product produced by the stitching is in most places quite satisfactory with no visible joins except where changes in data resolution are visible. The project aimed to produce, at each join, a surface that was at least as good as the worst data involved in the join, and this objective was achieved. Minor artefacts remain in some places that could be improved with more manual editing or collecting better quality source data.

A large number of detailed processing steps have been omitted in the above descriptions, including removing spurious data, editing to fill gaps, conversion of grids to points for interpolation, selecting reliable elevation points from bathymetry data and construction of polygons to define buffer and overlap areas.

The process was manually intensive requiring a great deal of intervention, judgment, analysis and editing. Contrary to our early expectations, there was little scope for automation as every join had different characteristics; the project took about twice as long as initially expected. The manual editing and manipulation of data around the shores of the lower lakes created the greatest difficulty.

Many of the difficulties encountered in this project could be avoided by specifying the data collection and processing in a way that would support stitching to adjacent data. The key steps are:

- ensuring that adjacent data acquisitions overlap with existing data, rather than abutting or leaving gaps; and
- identifying where a DEM is supported by measured data and where it is interpolated (invented); one way to achieve that is to ensure that source data (such as lidar point data or bathymetric soundings) are provided in addition to gridded products

Eliminating those two problems, particularly the second, would greatly facilitate the stitching process in subsequent projects.

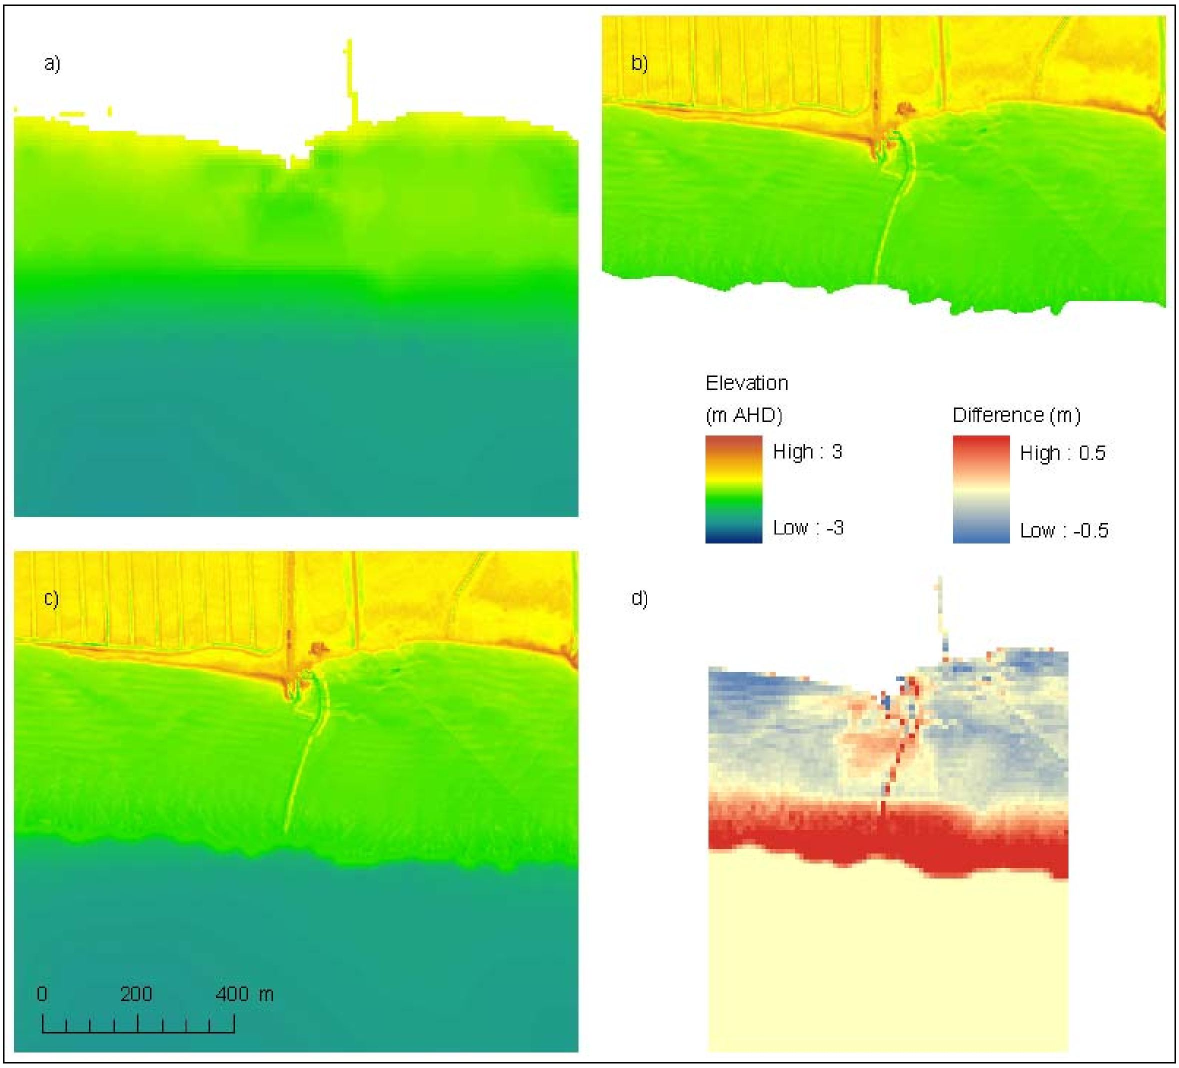

Figure 6. Stitching of lake bathymetry (a) with lidar (b) to produce combined product (c). The area of overlap between the bathymetry and the lidar is essentially replaced by the lidar data, leaving the lidar data unaltered. The difference (d) between the stitched product (c) and the supplied bathymetry (a) is zero in the deeper areas and is mostly between -0.5 and +0.5 m.

ACKNOWLEDGMENTS

This work was funded by the South Australian Department of Water, Land and Biodiversity Conservation. The project was overseen by Russell Flavel (DLWBC) with project management support from Mark Thomas (CSIRO Land and Water) and Susan Stovell (Department of Primary Industries and Resources of South Australia).

REFERENCES

Grohman, G., Kroenung, G., and Strebeck, J. (2006) Filling SRTM voids: The delta surface fill method. *Photogrammetric Engineering and Remote Sensing* 72 (3), 213-216.

Gesch, D. and Wilson, R. (2002) Development of a seamless multisource topographic/ bathymetric elevation model of Tampa Bay. *Marine Science and Technology Society Journal* **35** (4), 58-64.