

## 3DM Analyst 2.1 Trial at BMA Coal's Goonyella Mine

10th March, 2004

### Trial Overview

The purpose of this trial was to demonstrate the suitability of ADAM Technology's 3DM Analyst product for pit wall mapping and geotechnical data collection by generating 3D surface data of a 700m section of pit wall at Goonyella mine in Queensland.

A series of 108 images were captured from 36 locations using a 5.9 megapixel Nikon D1x digital SLR to test various methods of operation. The camera locations and 52 control points were surveyed by BMA Coal's surveyors using a total station.

Two sets of camera orientation data were obtained. The first set was generated using 3DM Analyst alone, which required the use of all of the control points. This mode of operation is quite common and is also consistent with the approach that BHP Billiton Technology has used in the past, but it does require more surveyed control points in order to be used. The other set was generated using 3DM Analyst in conjunction with ADAM's 3DM CalibCam block adjustment package, which allowed a subset of control points to be used, and demonstrated one method of reducing surveyed control requirements without sacrificing accuracy.

In both cases, 3DM Analyst was then used to automatically generate the 3D surface data of the pit wall, which was then exported as 3D images for use in Vulcan.

### Equipment Used

- 5.9 megapixel Nikon D1x digital SLR for image capture (images resampled to 10.5 megapixels)
- 2.0 GHz Athlon XP 2400+ personal computer with 1GB of RAM running Microsoft Windows XP Professional

### Background Information

3DM Analyst operates on the principles of *photogrammetry* — the science of determining 3D data from pairs of 2D images. In order to do this, the precise location and orientation of the camera at the time the images were captured must be determined.


3DM Analyst supports several methods for doing this. The simplest is to supply the software with the 3D co-ordinates of a series of points in the images, known as *control points*. When the camera positions are not known, at least three control points are required for a single *model* (pair of images). If both camera positions are known then only one control point is required; more are usually recommended, however, because without any redundancy, it is impossible to detect errors in the data.

In addition to the control points, which require surveying, 3DM Analyst requires additional *relative-only* points to help it determine the precise relationship between the cameras. If only one control point is in use, then eight or nine relative-only points will be required. There are no particular requirements on the points chosen other than that a good distribution of them over the images will strengthen the solution. Because the software can use these points to determine the relative orientations of the cameras, there are no requirements on the orientations of the cameras themselves when the images are taken — they can be rotated to any angle without affecting the results.


The method that 3DM Analyst uses to determine the *exterior orientations* (i.e. both the camera positions and angles of rotation) is known as a *bundle adjustment*. This takes into account all of the information supplied to the software — the 3D co-ordinates of the ground control points and (optionally) camera positions; the supplied accuracies of each point in each direction (X, Y, and Z); the 2D image co-ordinates of the ground control and relative-only points; and the supplied accuracies of the image co-ordinate observations in each direction (X and Y).

Usually, when using the above approach, the two images to be used for each model would point at roughly the same location. This is known as *convergent* photography and gives close to 100% overlap in the area covered by the two images (Figure 1). An alternative approach is to capture the images in a *strip*, with each image overlapping the images before and after it in the strip by at least 60% (Figure 2). This ties all of the images together and allows orientation information to be spread along the strip, without requiring each model to have control points. This method is used extensively by the mapping industry for this reason.

To use this approach, 3DM CalibCam is used to determine the camera orientations instead of 3DM Analyst. 3DM CalibCam uses an even more powerful algorithm known as a *block adjustment*. This is basically a bundle adjustment enhanced to simultaneously consider the observations for an unlimited number of images containing an arbitrary number of points (control and relative-only) in order to determine the final orientations. Because of this, not only is 3DM CalibCam able to achieve even higher accuracies than 3DM Analyst is capable of, but also it can do so using far fewer surveyed control points.



**Figure 1. Pairs of convergent images**



**Figure 2. Strip of images**

## Method of Operation

The basic workflow using 3DM Analyst is:

1. Capture images and control information.
2. Determine camera orientations (using 3DM Analyst or 3DM CalibCam).
3. Generate the Digital Terrain Models (DTMs).
4. (Optional) Digitise additional vector data, including breaklines.
5. Export the results for processing in other packages (e.g. 3D images, DXF files, etc.).

### Image Capture

The first step is to plan the photography and lay out the control points to be surveyed. In this case, one constraint was that the pit was only about 50m wide, and the shortest focal length lens available was 24mm, so the ground coverage of each image was only about 40m. This meant that more images were needed than were required in order to achieve the desired accuracy for this trial.

Two sets of images were captured:

1. A series of convergent models with at least one control point visible in each model.
2. A single strip of images with at least 60% overlap.

The purpose of the first set was to demonstrate the results that can be obtained by 3DM Analyst using the same methods currently in use by BHP Billiton Technology. The second set was to demonstrate how 3DM Analyst and 3DM CalibCam could be used together to minimise the amount of control that needed to be surveyed while still maintaining high accuracies.

To cater for both cases, a total of 36 camera positions were chosen, averaging 20 metres apart. Eighteen control points were also marked, one directly in front of every second camera position. Each camera position was used for one image pointed directly towards the pit wall, one image looking forward, and one image looking back, resulting in a total of 108 images. (The convergent images were actually captured with the intention that only half of them would be needed. For example, the images taken from camera positions 1 and 3 were designed to be used for a single model, while the images taken from camera positions 2 and 4 were designed to be used for a backup model in case something went wrong.) Normally,

once a method of image capture had been chosen, only one of the two sets of images would be captured. Also, if a block adjustment was to be performed on a strip of images, far fewer control points would be used.

### ***Orientation Determination***

The next step is to precisely determine the camera's exterior orientation when each image was captured. Note that both 3DM Analyst and 3DM CalibCam operate directly with the original images at this stage — no pre-processing is required to remove lens distortions, etc., from the images.

### ***Independent, Convergent Models***

Using this method, 3DM Analyst is used by itself and each model requires at least one control point in addition to the two camera positions. The process of determining the orientation is straightforward:

1. Pairs of convergent images are loaded into 3DM Analyst, together with the camera's calibration and control point information.
2. The control point(s) are digitised using either the target centroiding tool (capable of locating the centre of circular targets to within  $1/10^{\text{th}}$  of a pixel) or the natural target tool (which will match an existing observation using Least Squares Matching to within  $1/3^{\text{rd}}$  of a pixel).
3. Relative-only points are digitised using either the dual cursor digitising tool (which will search in a window around the two points the user has digitised looking for the best match to within  $1/3^{\text{rd}}$  of a pixel) or the natural target tool.
4. The camera positions are specified by selecting the corresponding control point ID from a drop-down list for each image.
5. A bundle adjustment is performed, and the results checked for accuracy. If there are points with significant errors, they should be checked and removed if found to be incorrect.

For most models, this entire process takes between two and five minutes. Using labelled circular targets for the control points improves the speed and accuracy of this process dramatically, as locating and clicking on a target only takes a few seconds per target. For other control points it can take 15–30 seconds to decide exactly where to digitise in order to maximise accuracy.

### ***Strip of Images***

To use CalibCam to perform a block adjustment the following steps are required:

1. All images are loaded into CalibCam, together with the camera's calibration and control point information.
2. The control point(s) are digitised using either the target centroiding tool or the natural target tool, as for 3DM Analyst.
3. Relative-only points are digitised using either the dual cursor digitising tool or the natural target tool, as for 3DM Analyst.
4. The camera positions are specified by selecting the corresponding control point ID from a drop-down list for each image.
5. A resection and then a block adjustment is performed, and the results checked for accuracy. If there are points with significant errors, they should be checked and removed if found to be incorrect.
6. The project is saved and CalibCam exited.
7. Pairs of neighbouring images are loaded into 3DM Analyst to form models. Camera calibration and control point information is not needed as the images are fully oriented when loaded.

The workload for each additional image is slightly more than that for each image in 3DM Analyst because more relative-only points are usually digitised in order to strengthen the solution and compensate for the lack of control points. Running the block adjustment also takes longer, but overall it should take less time than the sum of all of the individual model bundle adjustments combined.

## DTM Generation

Once 3DM Analyst has a completely oriented model (either from performing a bundle adjustment itself, or from loading camera orientations generated by 3DM CalibCam), generating the DTM is very straightforward — clicking on the “Go” button in the toolbar is all that is required for 3DM Analyst to generate a DTM. Generally this takes less than five minutes.

If a large number of models need DTMs generated, the DTM Generator batch-processing program that comes with 3DM Analyst can be very useful. Its purpose is to allow DTMs to be generated for a number of projects in a batch-processing mode, without requiring human intervention (Figure 3).

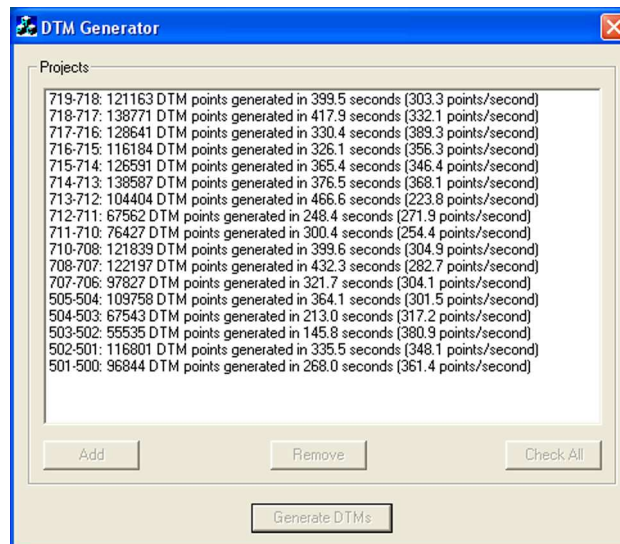


Figure 3. DTM Generator

In order to use DTM Generator, the user first creates a project in 3DM Analyst and then performs a bundle adjustment to derive the orientations before saving the project. Any number of projects can be set up in this way. Then the user can set up DTM Generator to load each of those projects in turn, resample the images, generate the DTM, and save the project with the updated information.

Since this software is designed to operate unattended (e.g. overnight), DTM Generator also includes the ability to check all of the supplied projects for potential problems (missing images, etc.) at the click of a button (“Check All”) before starting the DTM generation process.

## Digitising

If 3D vector data is required, 3DM Analyst offers three methods for digitising polylines:

1. Stereo View digitising
2. Single Image digitising
3. 3D View digitising

The Stereo View is the only view that does not require a DTM in order to digitise, but it does require good stereo perception on the part of operator. The software can help by automatically tracking the surface as determined by the DTM if one is available.

Single Image digitising allows the user to digitise the polyline directly on one of the two images in the Images View, while the software automatically locates the corresponding point in the other image. It uses the DTM as an initial approximation, then refines the point using Least Squares Matching (the same algorithm used to identify 3D points for the DTM).

3D View digitising allows the user to digitise the polyline directly on the 3D model of the DTM, viewed from any angle. The software locates the 3D point by intersecting the digitised location with the DTM.

All three modes allow the user to digitise polylines of different user-defined types and various attributes (Figure 4).

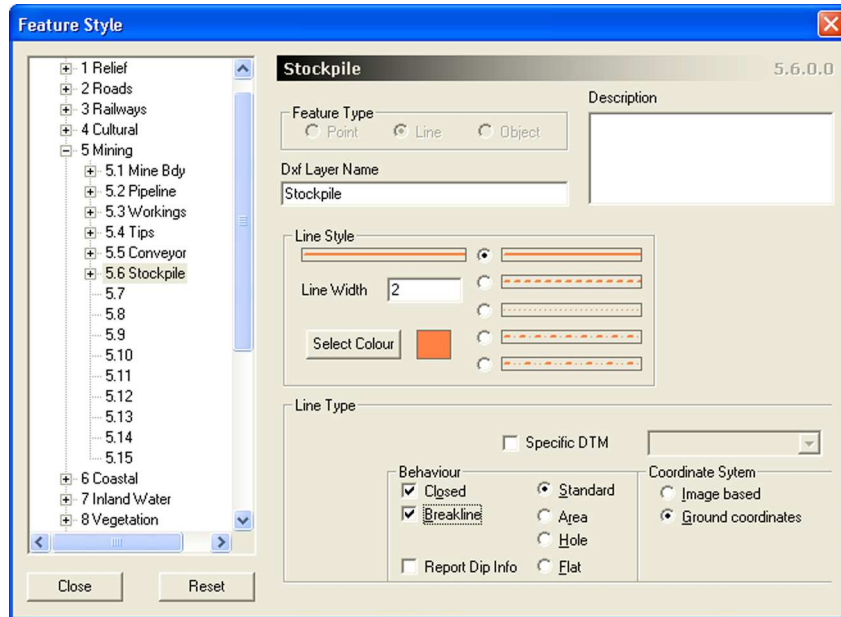


Figure 4. Feature Style dialog

The vector data can be imported and exported in DXF format, using the layer name specified in the feature’s definition.

Apart from simply representing features on the ground, polylines can also be used to control the generation and triangulation of the DTM according to the behaviour specified in the feature definition. For example, if the polyline is defined to be a *breakline*, then triangles are not permitted to form across the polyline. This can be used along crests, for example, to ensure points on one side of the crest aren’t inadvertently connected to points on the other side. If any *area* features are present, then only points within an area feature will be generated. *Hole* features are used to isolate areas that shouldn’t have DTM points (e.g. a pond, river, etc.).

In addition, statistical information about the polylines can be obtained — the number of points in the polyline, the length of the polyline, and the area enclosed within the polyline (Figure 5). If *Report Dip Info* is enabled, the dip and dip direction of the plane that best fits the digitised points will also be reported.

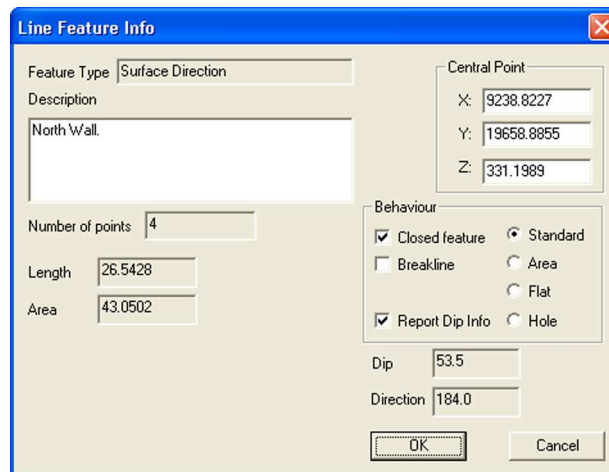


Figure 5. Line feature information dialog

In addition to these operations, 3DM Analyst supports merging of multiple DTMs, volume calculations either with a single DTM against a base or the difference in volume between two DTMs, and filtering of DTMs.

## ***Export to Vulcan***

Exporting 3D Images to Vulcan is straightforward: the user simply clicks two points to select the top-left and bottom-right points of a rectangle around the area to be exported and 3DM Analyst generates a TIFF file containing the image and a text file containing the surface information for Vulcan to import.

## **Results**

We have generated two primary sets of data — one set derived from the convergent images using at least one control point per model, the other set derived from the strip of images using 3DM CalibCam for the orientations.

Each set consists of .ori files (one for each image, containing the orientation information); .vwr files (3DM Analyst projects that can simply be loaded and used); and .dtm files (the actual surface generated by the software).


In addition to generating this data, we have also performed various tests with varying numbers of control points in use in the strip of images to give some indication of what effect the number of control points used has on the accuracy.

## ***Times***

1. **Data capture.** The total time taken to set up the total station, mark out the control, capture the images, and survey the camera positions and control points was approximately five hours. The time to set up the camera at a new camera position, capture three images, survey the camera position, and walk to the next camera position was five to six minutes.

These times would drop dramatically in normal operation. A large number of images were captured for the trial in order to test various scenarios — in normal operation, only the images required by the method chosen would be needed. A large number of control points were also marked and measured for accuracy testing — especially when using a strip, far fewer control points are actually needed. Finally, for repeat surveys of the same area, the camera positions and control points would not need to be surveyed again at all — someone could simply walk along the pit wall, stopping to take images at each pre-marked camera location.

2. **Orientations.** In the case of the independent, convergent models, the time required to perform a bundle adjustment for each model ranged from two to ten minutes. The quickest times were obtained when circular targets were used; the longest times when other control point markers were used, making it difficult to judge exactly where to digitise (see Figure 6 — trying to achieve sub-pixel accuracy when the head of the arrow is six pixels across takes time), and when mistakes were made when labelling control points or camera positions.



**Figure 6. Arrow used for control point**



The time taken to digitise the control points and relative-only points were similar for the block adjustment, and the issues were the same. The time taken to actually perform the block adjustment, generating a solution for all 38 images, was 2 minutes 20 seconds, although this was iterated a number of times to locate and remove bad observations.

3. **DTM generation.** Each model took an average of 4.4 minutes to generate an average of 91,000 points. (The DTM shown in Figures 6–8 consists of 96,186 points and took 3 minutes and 17 seconds to generate.)
4. **Exporting.** Exporting the data for a single model for use in Vulcan takes less than ten seconds.

## Accuracy

### *Independent, Convergent Models*

If two camera stations are used, then the minimum amount of control required for individual, convergent models is a single point. Two control points provide some redundancy and help in detecting a bad observation. If both points are the same distance from the cameras, however, they will not help the software to detect a bad camera calibration. A small error in the calibrated focal length would cause stretching or compressing of the data closer to and farther from the cameras than the control points. For this reason, we would recommend at least two control points at two different distances from the cameras (e.g. on different benches) if independent, convergent models are to be used and accuracy is a concern, but we would also expect the software to be at least as accurate as the current methods if only one control point per model was used.

The accuracies for the independent, convergent models in this trial, measured by checking control point values, were 4.4cm on average, or slightly worse than the strip case using 9 control points for the entire strip.

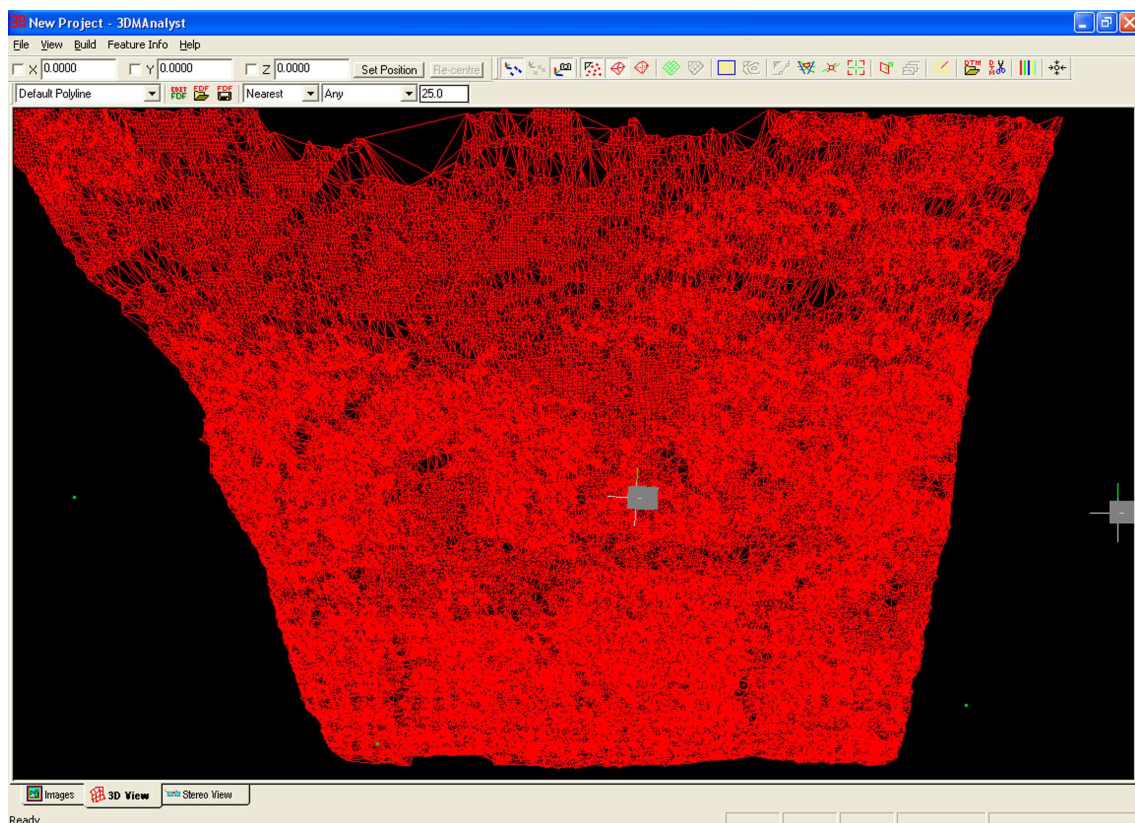


Figure 7. Wireframe view of a DTM

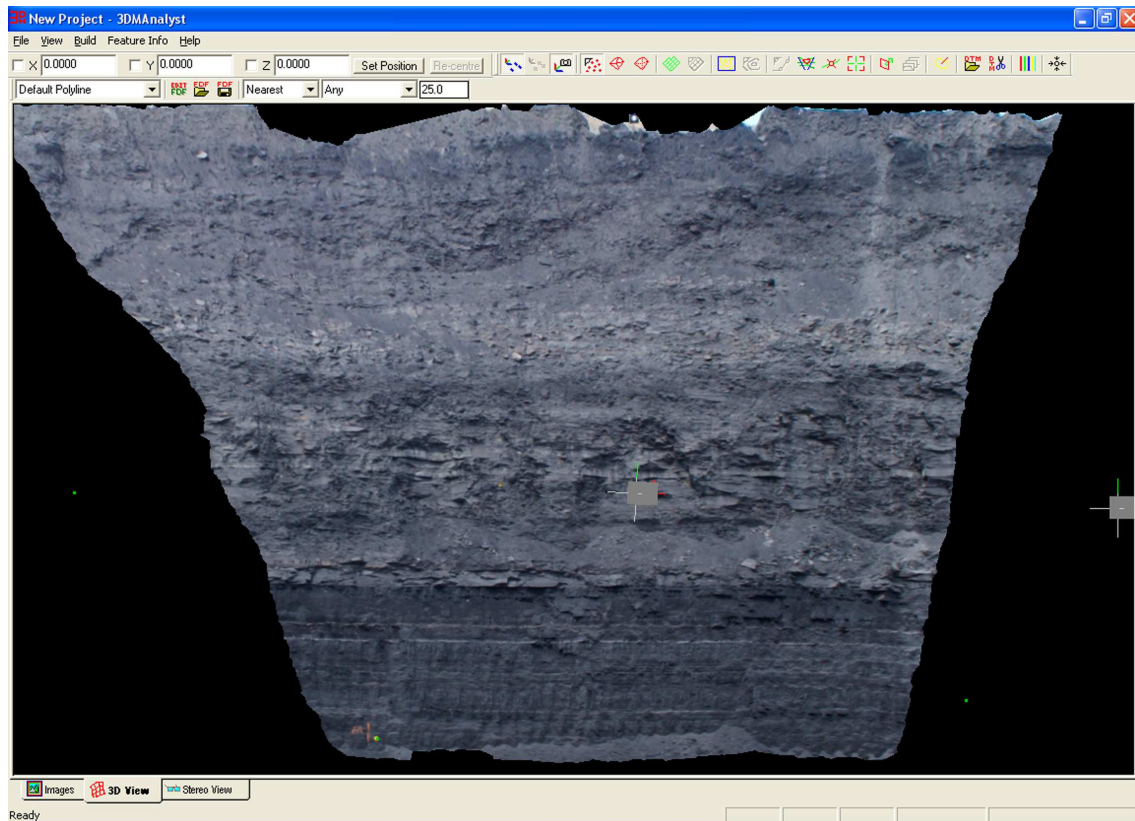


Figure 8. Textured view of a DTM

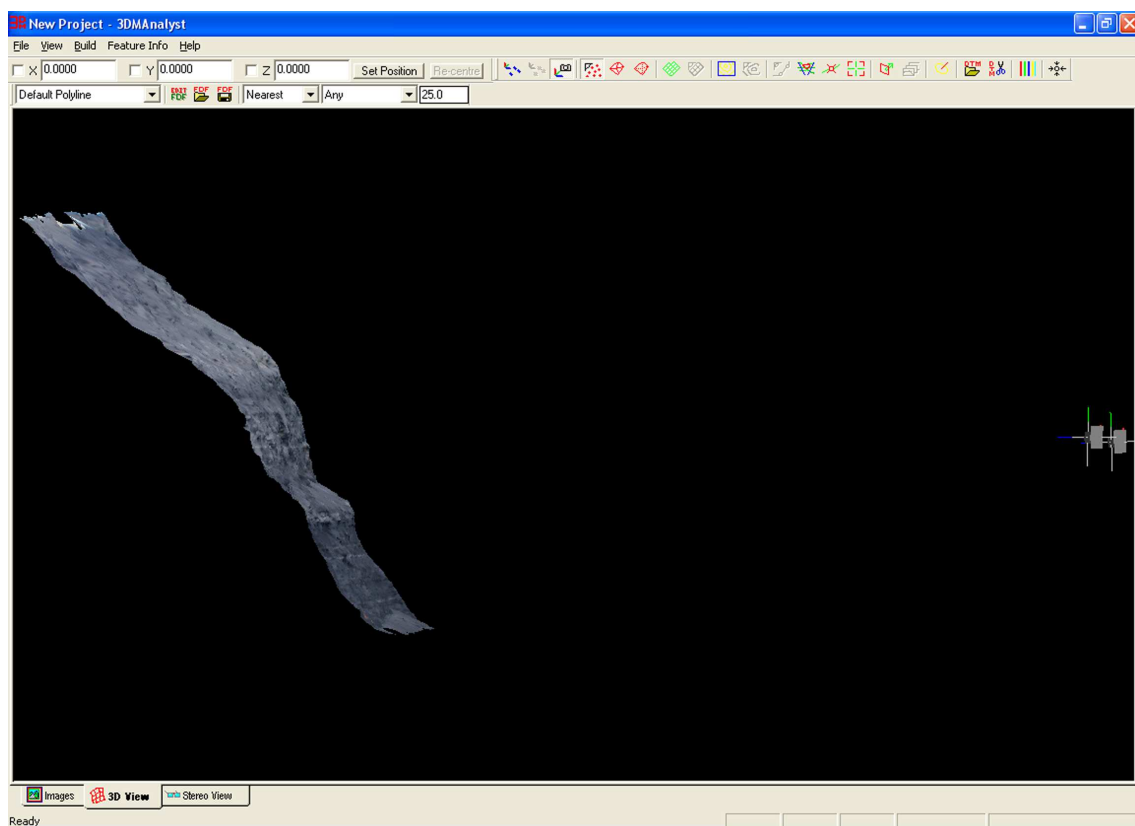


Figure 9. Textured view of a DTM from the side



### Strip of Images

In order to check the accuracy of the results, the 3D co-ordinates of the control points generated by the software were compared to the surveyed values. The RMS values reported in the table are the sample standard deviations to give an unbiased estimate of the expected accuracy. The figures for each co-ordinate axis are reported separately as well as the actual distance. All figures are in centimetres.

Test Number	Number of Control Points Used	RMS x	RMS y	RMS z	Overall RMS
1	41	1.2	1.3	1.2	2.1
2	9	2.4	1.9	1.9	3.6
3	4	3.7	3.6	2.8	5.9
4	1	3.0	3.6	2.4	5.3
5	1	3.5	3.4	3.4	5.9
6	0	6.6	4.7	16.0	18.0

#### Notes:

- Test 1. All camera stations and control points were used in this test. (Only 41 control points of the 52 surveyed were visible in the images, and only 33 were present in at least two images; residuals are reported for those 33 points only, as the software cannot generate 3D co-ordinates for points visible in only one image. All 41 were used for the orientations, however.)
- Test 2. Control points 500, 501, 607, 616, 617, 713, 718, 719, and 6110 were chosen for this test. Four of the control points lie at one end of the strip, three at the other end, and two in the middle. This represents quite a good distribution of control without much impact on the final results.
- Test 3. Control points 500, 607, 617, and 719 were chosen for this test. Two of the control points lie at each end of the strip. This would be a reasonable option if control must be minimised or accuracy isn't as much of a concern.
- Test 4. Control point 501 was used for this test. The lower residuals compared to Test 3 indicate that at least one of the control points used in Test 3 was either surveyed or digitised poorly, affecting the orientation.
- Test 5. Control point 6110 was used for this test. The change in residuals compared to Test 4 highlights the dangers of using too few control points — the entire solution depends on the points being surveyed accurately and digitised accurately.
- Test 6. This test is reported for completeness. We do not recommend it be used in practice. If all of the camera stations happened to be collinear then the result would be arbitrary — at least one control point is required to find a solution in that case, and even then there is no redundancy. It happens to work in this case because there is enough deviation from a straight line for 3DM CalibCam to determine the approximate orientation of the cameras. Note the dramatic increase in the Z-axis residuals — this is because rotating the cameras up or down does not have a huge impact on the camera station residuals, so the solution is not particularly strong in this direction.

These tests reveal that very good results can be obtained using very few control points, if required. The software is quite flexible and able to accommodate any surveying restrictions the customer has.

## Company Profile

ADAM Technology was founded in Perth in 1986 to develop and manufacture systems based on the principles of photogrammetry — the science of deriving 3D information from photographs. The company has ten employees and a global network of 12 distributors, resellers, and agents in 11 countries. All manufacturing and software development takes place in the Perth office.

In 1995 ADAM embarked on a major research effort to develop the 3DM (3-dimensional, dynamic measurement) technology. 3DM Analyst is one of the products of this research effort and it was released

in January of last year as a replacement for the MPS-2. 3DM Analyst Professional, which is able to handle images in excess of 750MB in size, is due to be released shortly.

In addition to analytical stereoplotters and the 3DM Analyst products, ADAM has also developed custom monitoring and control systems that are based on the same technology, two of which — 3DM Bunker and 3DM Gob — have been sold commercially in Japan. The 3DM Bunker system is designed to monitor the amount of rubbish in large pits at garbage incinerator plants, while 3DM Gob is used in glass bottle factories to control the size of the bottles produced.

### **Contact Details**

Telephone: +61 8 9479 5575  
1800 672 553 (free call in Australia only)

Facsimile: +61 8 9479 5585

Email: [adam@adamtech.com.au](mailto:adam@adamtech.com.au)

Web: <http://www.adamtech.com.au>