ADAM TECHNOLOGY

Image Capturing Recommendations

Introduction

This document outlines ADAM Technology's recommendations for image capturing procedures using 3DM Analyst and 3DM CalibCam in a mining environment.

The primary goal is to find the optimum balance between two conflicting aims:

1. To minimise the time and risk associated with capturing images.
2. To maximise the robustness and accuracy of the resulting data.

These recommendations cover control design and placement and three different image capturing methods (independent, convergent models; strips; and image fans).

Control

ADAM recommends the use of both control points and surveyed camera positions, as this will reduce the amount of control required and surveying camera positions will usually be safer and more convenient than surveying control (since the location must be safe in the first place for images to be captured from there). For independent models and image fans, surveying both camera stations reduces the control point requirement to a single point per model (or fan); for strips of models, it is not necessary to survey every single camera station but it is better to do so.

The number of control points required for a strip of models is very flexible: independent trials have shown that even for a wide range of control, the resulting accuracy is not greatly affected. Keeping this in mind, ADAM recommends two or three control points at the start and end of each strip, with an additional control point or two every five to ten images or so. (These figures are not set in stone — if mistakes are rare or the costs of recapturing images not great, feel free to increase the distance between control points.)

Image fans are even more flexible, due to the reduction in the number of unknowns to be solved by the bundle adjustment. At minimum, one control point is required to control the entire project, if all images were captured from two surveyed camera positions. (Theoretically, no control points are needed if three camera positions are used, but this approach is not recommended as it tends to amplify control observational errors.) Additional control points provide redundancy and allow accuracy estimates to be made and are therefore strongly recommended.

Design

ADAM recommends circular white targets placed on a black background (or vice-versa), with a mark in the centre of the target for unambiguous surveying using a total station or GPS receiver. The actual materials used are not important for the software; they just need to be visible, stay in position from the time they are measured until the images are captured, and be stiff enough to handle a bit of wind.

Creating disposable targets is certainly a reasonable thing to do and has the advantage that they only need to be put into position and not necessarily removed.

An alternative approach that works well for long-term control points in stable areas is to paint circular marks on the pit wall. Painting an identifying mark next to the control point (e.g. its number) is also strongly advised.

Placement

Ideally, control points should "bracket" the area of interest, especially when camera stations are not surveyed; a model or strip with control points concentrated in one location will not be as accurate at other locations as a model or strip with control points surrounding the area. (Surveying camera stations decreases this problem somewhat.)

Placing control points several metres in front of the pit wall rather than directly on it to improve safety is not a problem — errors scale linearly with distance, so if a control point is surveyed to 2cm accuracy and placed halfway between the camera and the pit wall, then you would expect the error to be no more than 4cm at the pit wall (not taking into account the accuracy of the surveying of the camera's position). The main issue to be aware of is that the control point must still be visible in both images and not too close to the camera that it becomes blurry from being out-of-focus. (Part of the reason that ADAM recommends always using an aperture setting of f/8 is to ensure a large depth of field (i.e. the distance from the closest point that is in focus to the farthest point that is in focus), providing a fair degree of flexibility in the placement of the control target without compromising its sharpness in the image. In any event, a small amount of blurring is not particularly harmful to the accuracy of the target centroiding algorithm, as long as the blurring is uniform.)

Placing control points along the top of the pit wall as well should be considered, but only if it is completely safe to do so. Using these control points has the advantage that it will highlight any error in the calibrated focal length, but if the lens has been calibrated well this becomes less of a concern. The top-of-wall control points are therefore not critical for general surveying, but they would be highly desirable during camera calibration.

Even if a few control points are lost or damaged, 3DM CalibCam can still compensate for any missing control points using a strip or fan combination and any in-pit control.

Image Capture

There are three primary methods of image capture that can be used with 3DM Analyst and 3DM CalibCam:

1. Independent, convergent models.
2. Strip/block of models.
3. Image fans.

In many cases, the choice between the first two methods is largely determined by the relative importance of the different trade-offs that they entail. The third method is mainly useful over long distances.

Before comparing the first two methods, it is worth noting that the number of models that will be required to capture a given length of pit wall (assuming the entire height of the wall is captured in a single image) is given by

$$Number\ of\ Models = \frac{Total\ Length}{Effective\ Model\ Width}$$

In the comparison below, the effective width of each model will therefore be critical in determining how many models are required for the job.

Independent, convergent models

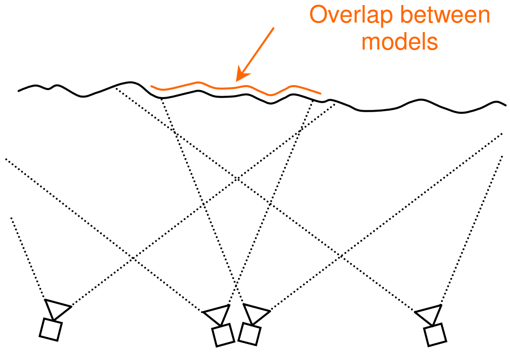

Figure 1. Independent, convergent models.

The key characteristic of this method is that close to 100% of each image is used in a single model, and there is very little overlap between models. This means that if an image captures 50m of the base of the pit wall, then the effective model width will be quite close to 50m, and therefore this method requires only half the number of models (and hence 3D images) required by the strip method. (In practice it would be wise to plan for a 10–20% overlap between models.)

The downside of this method is that each model needs to be fully controlled, and so the surveying requirements are more onerous. (There *is* some scope for passing control information between models by using 3DM CalibCam and taking advantage of the overlap between them (Figure 1), but the accuracy will be worse than if the model was controlled unless the overlap is quite large. Exactly what overlap would be required to allow, for example, a strategy of one control point for every second model to be adopted without sacrificing too much accuracy is an interesting question that can only be answered by experiment.)

This method is most desirable when the overhead of working with twice as many 3D images (as required by the strip case) outweighs the costs and risks of placing and surveying control for each model.

Other advantages of this technique are:

1. There is a great deal of flexibility in the distance between the camera stations. The ratio of the pit wall distance to the camera separation can be freely chosen between about 2:1 and about 10:1, with the choice being largely determined by the desired depth accuracy of the model.
2. This method can be used with cameras of any focal length operating over any distance.

Strip of models

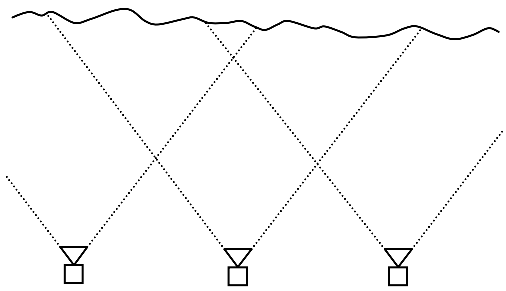

Figure 2. Strip of models.

In this method, a series of parallel images with large overlap (typically 60%) are captured and processed in 3DM CalibCam. The total number of images used for any given job will be about the same as the number of images required with the convergent models method, but the number of models will be about twice as high, due to each model only occupying about half of each image.

The key advantage of this technique is that the large degree of overlap between images allows orientation information to be reliably and accurately passed between models, drastically reducing the number of control points required for a given job without sacrificing accuracy.

This method is best used when mapping a long stretch of pit wall with a short focal length lens (e.g. 28mm), as was done during training. Given the lens, the distance to the wall, and the area to be mapped, the Object Distance spreadsheet installed with both 3DM Analyst and 3DM CalibCam can be used to plan exactly how far apart each camera station should be, and it will calculate exactly how many images will be required. The desired image overlap should be set to 60% to determine the nominal camera positioning, and 50% to determine the absolute maximum distance allowable between camera positions before the strip will be broken.

For example, using the Canon EOS 1Ds camera with a 28mm lens, completely mapping a 1km section of pit wall from a distance of 50m would require 41 images (40 models) captured 25m apart with a maximum separation of 32m. (Note that if the cameras are not parallel, the maximum separation will be less, which is one reason for building in a safety margin.)

The Object Distance spreadsheet calculates that the image width (perpendicular to the view direction) in this case is 64m. If this same section of pit wall were to be captured as independent, convergent models, this width should be reduced by 20% to provide overlap between models for safety and to compensate for perspective distortion, resulting in an effective model width of around 51m, requiring a total of 20 models to capture the entire wall. This would require a total of 40 images, or about the same as the strip technique. As noted, however, the convergent method does require only half the number of models and therefore generates half the number of 3D images (although they will be much larger). Which technique is more productive depends, therefore, on the relative cost (and safety) of placing and surveying control vs. the cost of working with twice as many (smaller) 3D images.

As another example, mapping the same section of pit wall with the Canon EOS 1D Mark II camera and the same 28mm lens would require 50 images (49 models) captured 20m apart with a maximum distance of 25m. The effective model width (with the same 20% safety margin as before) is 41m, requiring a total of 25 models (or 50 images) to capture the entire wall. Again, the amount of photography is about the same, but the strip method requires fewer control points while generating about twice as many 3D images.

When using strips, the camera separation calculated by the Object Distance spreadsheet should be used. This restricts the distance:separation ratio to the calculated value, making this technique less desirable for longer focal length lenses.

Image Fans

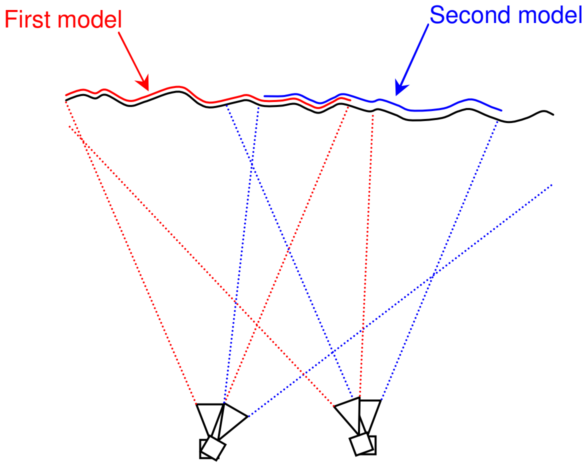

Figure 3. Image fans.

Image fans are similar to independent, convergent models, except that a series of images are captured from each camera location (Figure 3). Ideally, the images should be captured with a small overlap (at least 10%) to reduce the chance of gaps in the models, and provide the option of sharing orientation information so each model does not necessarily need to be individually controlled.

A key advantage that image fans have over the independent, convergent models is that because multiple images were captured from each location there are far fewer unknowns to be determined by the bundle adjustment. This improves the strength of the solution, makes the bundle adjustment run faster, and reduces the minimum number of control points required to find a solution down to one for the entire image fan.

Another advantage is that the next release of 3DM CalibCam supports *image merging*, where any number of images captured from the same location can be merged into a single, high-resolution image to sub-pixel accuracy, making these merged images suitable for use in 3DM Analyst as a substitute for the original images (Figure 4). The practical benefits of this are:

1. A cheaper and lower-resolution camera can be used instead of a much more expensive high-resolution camera to produce the same results, reducing the capital expense of the camera at the cost of slightly increased labour costs (rotating the camera multiple times to capture multiple images).
2. Images can be captured to build up the fan without regard for precisely where they are pointing. All the operator needs to do is ensure there is about 10% overlap between adjacent images captured from the same location.

The only drawback with this technique is that the merged images can become very large if many images are merged together — the comfortable limit for 3DM Analyst on a PC with 1GB of RAM is 30–40 megapixels. (3DM Analyst Professional can easily handle images in excess of 250 megapixels.) To alleviate that, the next release of 3DM CalibCam allows the user to pick out sub-sections of each merged image to create an output image from, allowing the user to generate 30–40 megapixel images containing the corresponding sections of the pit wall in the comfort of the office rather than trying to make sure each corresponding image overlaps 100% onsite when capturing the images from different locations.

Image fans are ideal when longer focal length lenses are used over large distances. Model strips are undesirable with longer focal lengths because they increase the object distance to camera separation ratio (reducing depth accuracy) and place restrictions on the camera locations that may be difficult to meet in practice. (For example, a Canon EOS 1Ds with a 28mm lens achieves a distance:separation ratio of 2:1 in strip mode. With a 50mm lens, this increases to 3.5:1, and with an 80mm lens it increases to 5.6:1. It is the narrowing of the view angle with longer focal lengths that forces this reduction in the relative separation of the cameras.)

Apart from the fact that multiple images are captured from each camera location rather than a single, low-resolution image, there really isn't any difference between image fans and independent, convergent models, so all of the other attributes of the latter apply to this method as well.

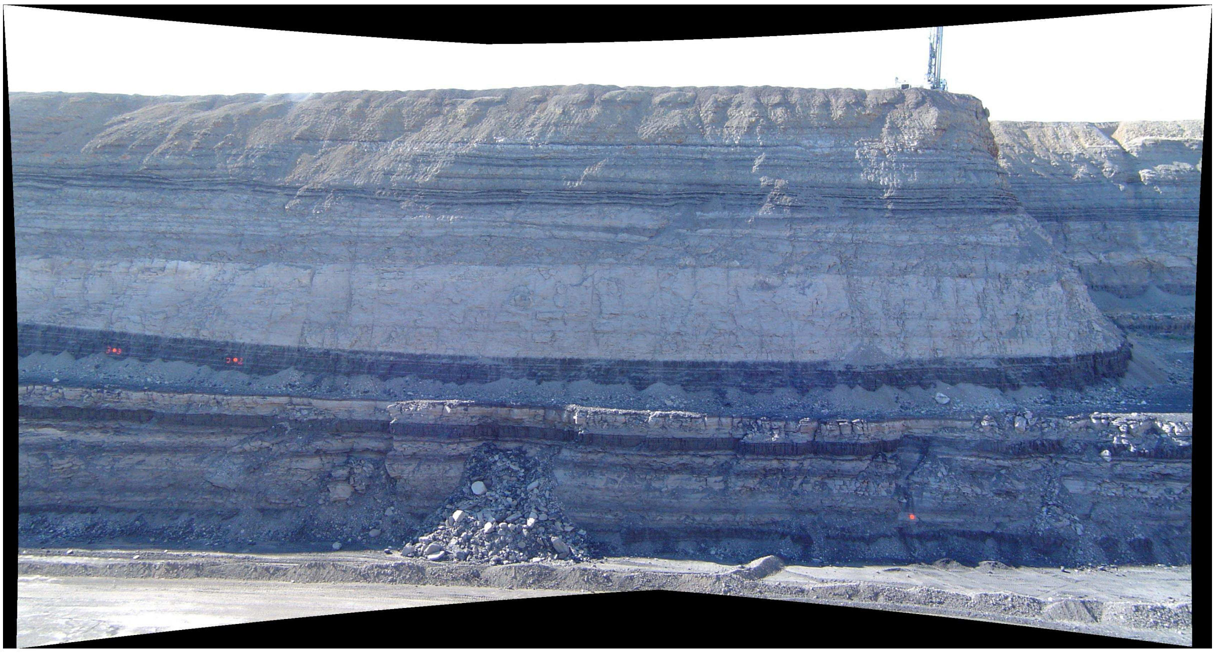

Figure 4. Pair of images captured from the same location merged into a single large image by 3DM CalibCam. Black areas around the image show how far pixels were moved to remove distortions.

Summary

The following table characterises the different image capturing methods and their recommended uses. In each case, the number of images required to capture the area of interest (which depends on the image size on the ground and not on the method chosen) is assumed to be *N*.

Method	Camera stations to be surveyed	Control points to be surveyed	Number of 3D images generated	Key Advantages	Recommended Situation
Independent, convergent models	N	> N	N/2	Flexible; low number of models/3D images	When surveying control points is safe and inexpensive compared to the cost of working with more 3D images
Strip of models	N + 1	4 + N/5	N	Minimal surveying	When mapping pit walls from close range and surveying control points is dangerous or expensive
Image fans	2 or more, depending on whether the entire scene can be seen from two locations	At least one per pair of camera stations[1]	Between 1 and N/2, depending on whether image merging is used	Flexible; low number of models/3D images; Minimal surveying	Long range image capture with long focal length lens and a clear view of the area of interest from a small number of locations

[1] Image fans captured from two camera positions can be thought of as equivalent to a single, independent, convergent model with really high-resolution cameras. Each one of these “mega-models” can also be overlapped with neighbouring image fans to share control information if desired, theoretically reducing the control requirement below one point per pair of camera stations, but additional control is advised for redundancy.

Thus, for a 1000m strip of pit wall to be captured at a distance of 50m using the Canon EOS 1Ds and a 28mm lens, the number of image (N) would be 40; the number of control points recommended would be at least 40 for the convergent models and 12 for the strip; and the number of 3D images generated would be 20 for the convergent models and 40 (narrower images) for the strip.

Note that it is entirely acceptable to switch between different methods for different sections of a pit wall depending on the circumstances. 3DM CalibCam supports combining all of the above techniques into the same project if desired.