PROCEEDINGS OF SCIENCE

Photogrammetry for Large Structures

M. Kesteven*
CASS, CSIRO
E-mail: michael.kesteven@csiro.au

Photogrammetry is a valuable metrology tool, allowing engineers to measure (and adjust) the reflectors of large radio-telescopes. We describe the process, based on the recent experience with the ASKAP antennas in outback Australia

From Antikythera to the Square Kilometre Array: Lessons from the Ancients,
12-15 June 2012
Kerastari, Greece

*Speaker.

PoS(Antikythera & SKA)025

1. Introduction

Photogrammetry is a valuable metrology tool - allowing one to determine the 3D coordinates of targets on a large structure such as a radio telescope reflector. The theory is basic surveying, known since (at least) the time of the pyramids. The actual implementation, however, is more recent - it really became practical only with the advent of digital cameras and personal computers.

This note provides a description of the technique, with examples drawn from a recent metrology campaign at the SKA pathfinder site in western Australia.

2. The procedure

We place small retro-reflecting targets on the surface - see figure 1 and figure 2. There are few restrictions on their location. We do require that their distribution provides a meaningful sampling of the surface. The targets are visible to a camera if the ray from camera to target is within a cone of about 60 degrees about the normal to the target surface.

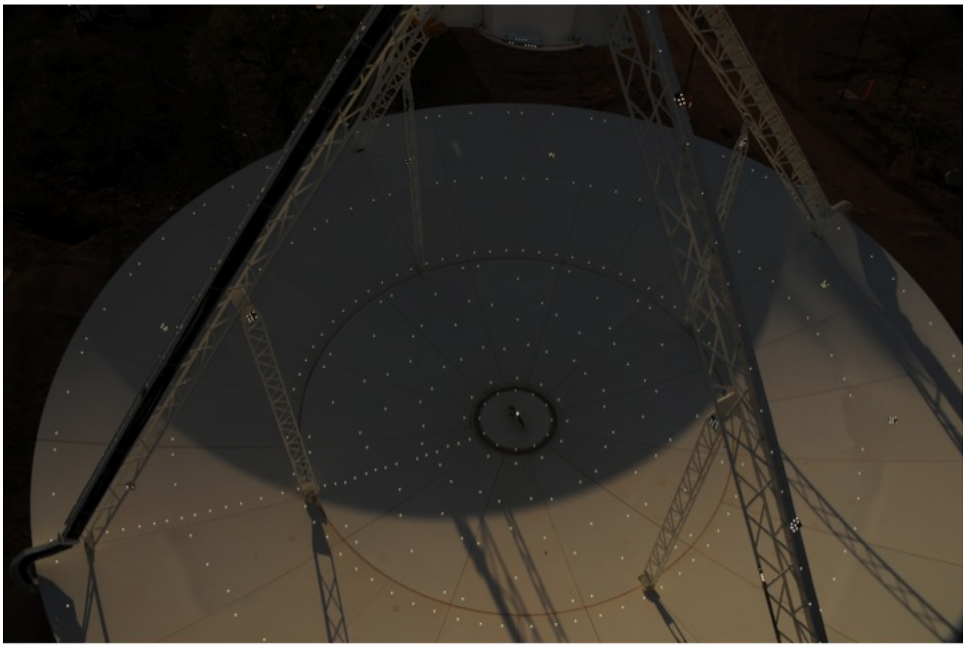

Figure 1: An ASKAP antenna setup for photography. There are about 150 targets on the reflector surface.

We take a number of photographs - roughly 100 for the 12m antennas of the SKA pathfinder. The processing involves three steps:

- We scan every image to extract the plate coordinates of every target (each detected target must meet some modest criteria: compact (occupy a small number of pixels); bright (significantly above the the image noise); circular.
- We search for the special target (a 3-D structure consisting of 5 targets whose relative positions are accurately known).

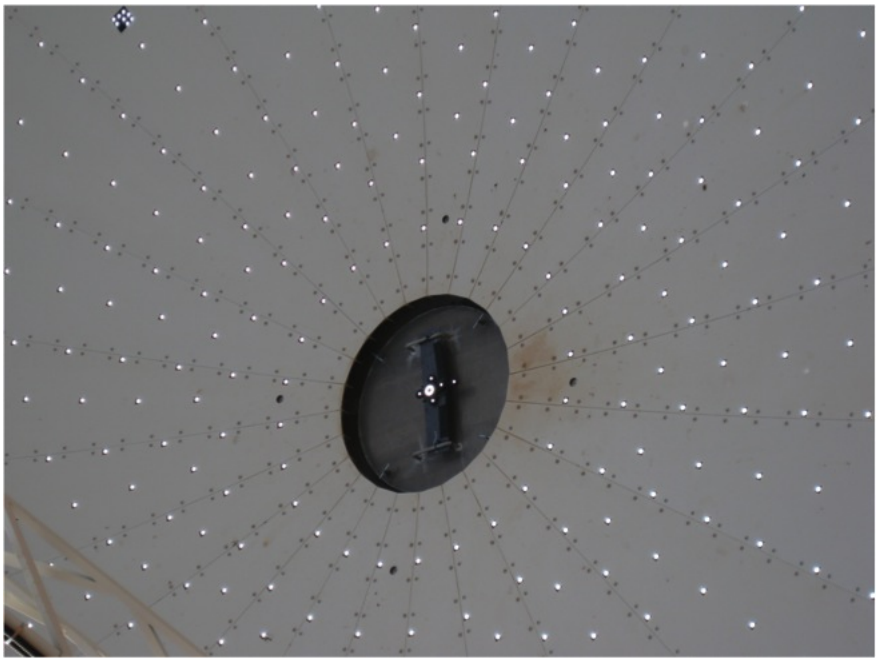

Figure 2: Antenna targets in more detail. (This is the 12m Patriot antenna at Parkes). Note the special target at the centre of the reflector

- We use the special target to provide a initial estimate of the camera location, relative to the special target.
- We iterate refining the target locations as well as the camera position and orientation for each photo.

The end product is a table of the (x,y,z) location of every target (figure 3)

3. The theory

The key to the operation is the recognition that a camera is essentially a very versatile theodolite: each pixel in the camera's image plane corresponds to a direction - in effect, azimuth and elevation. But we do need to know where the camera was when the picture was taken, and we need to know where the camera was pointing (the orientation of the camera axis).

We simplify the problem by limiting the number of unknowns - we place retro-reflecting targets on the surface of the object. The problem is manageable if the targets are representative of the surface at each target location. Each photograph will gather 2N new bits of information for a cost of 6 new unknowns (camera location and orientation).

Given N_t targets, and N_p photographs we can solve for the camera locations provided that :

$$\underset{(data)}{2*N_t*N_p} > \underset{(unknowns)}{(3*N_t+6*N_p)}$$

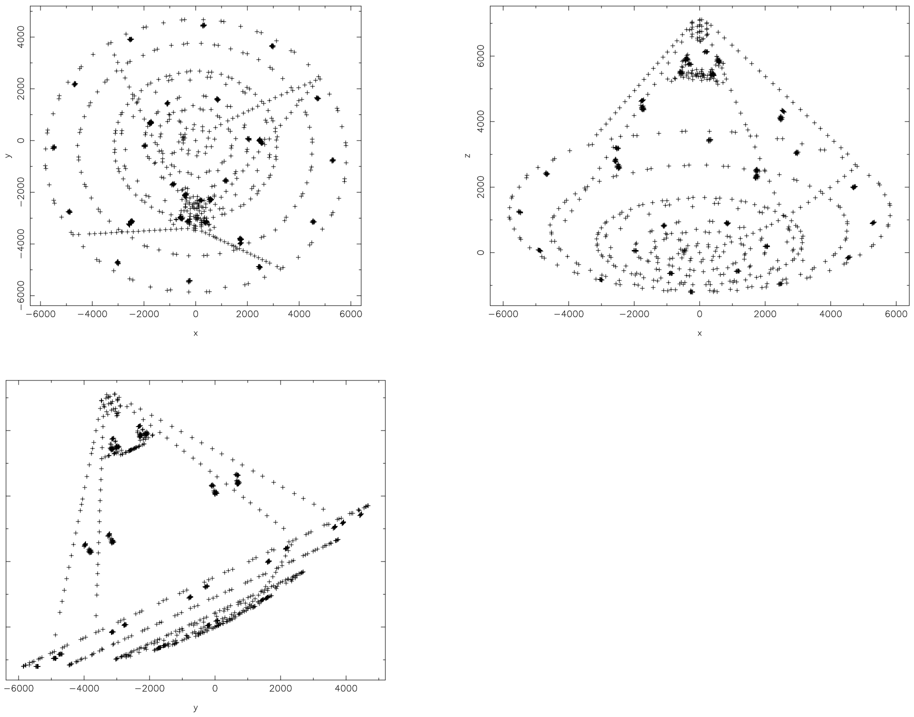

Figure 3: The 3-D target solution

The problem is easily over-determined; we can solve for the targets, the camera locations and obtain an error estimate for each target.

The critical issue is a matter of registration - the ability to identify each target in the various images. The relationships are shown in figure 4.

Each object (reflecting target) will appear in a number of photos. Registration is the identification process of correctly labelling every detected target in all the photographs. The processing then does an iterative loop to determine the location and orientation of the camera for each photograph, as well the location of every target. The position error assigned to each target is set by the convergence of all the rays around the target. (In the ideal world all the rays for a given target would converge on the target; in the real world they converge around the target).

3.1 Accuracy Issues

There is an automatic estimate of the accuracy that comes from the metric used in the iterative algorithm : for each target we have the rms perpendicular distance to the rays in its bundle.

We typically find $\sigma \sim 0.03$ mm for each axis. (for the 12m askap antennas).

The camera calibration is one contributor to this error. This is the mapping from pixel to an angle relative to the camera body. This calibration is refined as part of the iterative solution.

The distribution of rays at the target can be an issue : ideally the rays should be isotropic about each target, so some care is needed in taking the photos.

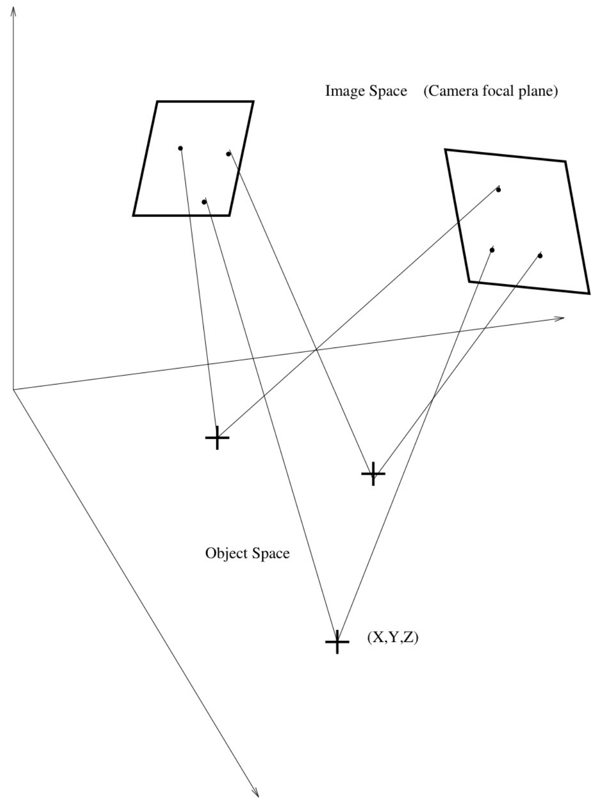

Figure 4: The registration issue - relating targets to their location on the camera images

Our experience is that the system is remarkably simple and robust.

The process gives us the location of each target. Whether or not the target is a representative of its immediate surroundings is a different matter. In effect, we require the scale size of panel defects to be larger than the target spacing.

References

[1] Clive S. Fraser, "Microwave Antenna Measurement" *Photogrammetric Engineering and Remote Sensing*, Vol. 52, No. 10, October 1986, pp1627-1635