StarZap: A Method of Removing Unwanted Stars
John Menke
Mar. 2005

Abstract.. While performing photometry on moving targets (e.g., asteroids), the object may come close to or even over a star on the image field. Later measurement of the brightness of the object will then be in error because of the star. StarZap is a program that will remove the offending stars from a sequence of images, allowing higher quality photometry to be performed.

The major problem that makes asteroid photometry more difficult than stellar photometry is that the asteroid moves through the FOV. Depending on the distance and orbit of the asteroid, it may move only a few a-min in a night, or it may move 20-30 a-min (or much more, if it is a near earth asteroid). In any case, the asteroid is almost always in a different field every night. This requires different reference stars every night, and creates the problem of making accurate night to night matchups of the data.

But the other problem from the moving asteroid is that it, in its journey across the FOV, it seems to seek out every possible star. As the asteroid moves across, or near, a star, the background photometry annulus will first cross the star, causing an apparent decrease in the asteroid brightness (background is subtracted from the central photometry annulus). Further on, the central annulus covering the asteroid may cross the star, thus increasing the apparent brightness. Then it decreases again as the asteroid leaves the star, looking for its next victim.

Naturally, Murphy's Law requires that these brightness variations always fall at the points on the photometry curve that are most critical (e.g., the missing section of the curve, where it matches to the previous night, where you think there was evidence of a satellite, etc.). Obviously, the easy thing to do is to throw out the data near the star, leaving a data gap. However, this loses valuable data. And it seemed to me that it should be possible to correct for the problem at least some of the time.

The basic idea is to subtract an appropriate star image from the combined asteroid-star image, leaving the asteroid by itself (and to do this for all the affected images in the sequence). There are two obvious sources of subtracting star images

- The subject star taken from a different image in the sequence in which the asteroid is not involved
- Some other star in the image of concern, suitably normalized to match the subject star which is mixed with the asteroid.

Of course there are pros and cons of each method. The first method is in some ways easier to implement, and in some way "better" in that it uses the image of the subject star (especially important if the subject is a double star or has a nearby star). On the other hand, variations in transparency during the sequence of images will require some sort of normalization be applied to the subtracting image. However, experiments with both approaches seem to show that the second method, to use a star image from the subject image, appears preferable. This is because the largest contributor to error in the process (other than limited statistics) is the variations in seeing and guiding between images. That is, on average it is better to use a different subtracting star (appropriately normalized) from the same

image than the same star from a different image. If the offending star is in fact a close double, rather than an isolated star, then in that case it would likely be better for the subtraction to use a copy of the offending star (suitably normalized) as taken from another image. This is easy to arrange.

Using first manual experiments, then a VB program, I have experimented with this general scheme using MaximDL. Maxim is suitable because it contains many pixel math functions that can be accessed by scripting or VB programs. Other programs may also be suitable hosts for this technique.

Experiments showed that if I use MaximDL to align a sequence (e.g., 50 images), then measure the variation in where the centroids of the resulting star images are, I find that the variation tends to be in the 0.03 pixel range as measured by the variation in the x/y positions of the star centroids. Thus, I can automatically align all the images at the beginning of the process, identify stars of interest (i.e., that the asteroid comes close to) by their x/y coordinates one image, and use that information throughout the analysis.

All but one of the methods this program script uses are standard in MaximDL. The exception is the current MaximDL automatic star alignment method will not work with only one star in each image. This feature is needed to align the ref star to be subtracted to the exact position (<< 1 pixel) of the target star determined from the REF image. At my request, Doug George had Hilderic Browne modify the system to permit one-star automatic alignment.

Images used in this test were taken using the ST7E at f6 on the C11 giving about 1 a-s per pixel. The setup is mounted on an AP1200 operated without PEC and without guiding over the 180 sec exposures. The average P-P tracking error is about 3 a-s on a 7 min. period, so there can be 1-2 a-s of image movement during an exposure. The FWHM of the star images is typically in the 3.7-4.3 pixel range (worse if poorly focused). I normally use a photometry aperture radius of 8 pixels, with an additional 3 pixel radius for the background annulus.
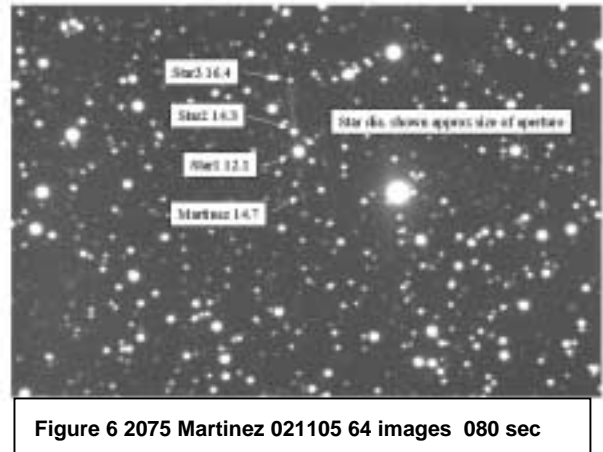

Figure 6 2075 Martinez 021105 64 images 080 sec

Figure 6 shows a typical sequence as the asteroid moves through a rather crowded field. As you can see, the asteroid comes very close to several stars in its path, one of which is nearly 10x brighter than the asteroid itself.

If we want the subtraction to be good to .03 mag (3%) then you need about 0.5% precision in your subtraction—a pretty tall order as it implies .005 mag precision in the subtraction process. In some ways, the situation is worse, because if the reference star has a different profile (PSF) from the star, even if the total brightness subtracted is correct, there may be substantial under-subtraction (residuals) or over-subtractions

in different parts of the star/asteroid image. On the other hand, the closer the asteroid path to the center of the star, the more these problems will tend to average out as the photometry aperture measures the total brightness over a relatively large area.

In a normal photometric analysis, we load the sequence into MaximDL, do dark subtract and perhaps flat field, and then measure the brightness of the asteroid compared to selected reference stars (I usually use four or five) in each image of the sequence. I export the raw data into Excel, then visually inspect the reference stars for variability, then usually choose just one as the reference star (usually one of the brighter ones). The images comprising Figure 6 yields the upper curve in Figure 7. The effects of the interfering stars are obvious.

To demonstrate how the StarZap works, let's go back to the image sequence. As noted, I first align the sequence. I then create a summed image that shows the path of the asteroid among the stars as in figure 6. Using this image, I note which stars are close to or on the asteroid path and identify these **target star(s)** by their locations (X,Y) on the image (I do this by just clicking on the stars in the StarZap program). I also identify an appropriate **reference star** I will be using for subtraction. This **reference star** will be a reasonably bright star that I can use both to measure the relative sky transparency appropriate to each image in the sequence, and as a good quality (good statistics) star image for use in the subtraction process.. The program uses the reference image to measure the brightness of each **subject star** relative to the chosen **reference star**.

After choosing the stars, I then start the program which follows the steps shown in the flow chart below. The whole process takes about five minutes total.
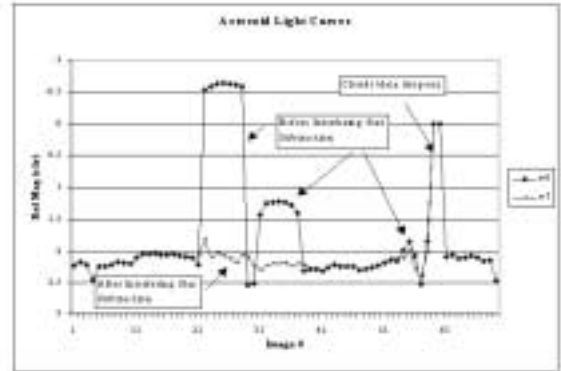

Figure 7 Asteroid Light Curves

Once the program finishes subtracting the various subject stars, I have a sequence of images that are the same as the original sequence except that they have been aligned (which actually makes the photometric analysis go faster) and have had the interfering stars removed. Once complete, I perform the usual photometry on the modified image sequence that has had stars removed.

How good a job does this do? Figure 7 also shows the resulting light curve for the sequence. The improvement in the light curve compared to the original curve is obvious. The apparent errors remaining in the curve after the subtraction are in the range of <0.1 mag (<10%) even for the 1000% error from the brightest star--a very substantial reduction. However, to evaluate the precise errors will require extensive comparisons of data, which I have not yet completed.

StarZap Summary

At this time, I have only just begun using the program as a standard tool in my data reduction so it has not yet been tested on a

wide variety of qualities of data. However, because the results presented here involved NO special settings or fiddling with parameters and yet still produced such good results, it would seem this approach has the potential greatly to improve our ability to measure light curves in crowded star fields. Anyone who wishes to experiment with this program or further develop it is welcome to contact me.

1. Acknowledgements

I would like to thank Doug George for his willingness to modify MaximDL to support this experiment. As always, I also thank Brian Warner for his incredible encouragement of fellow asteroid photometrists. I would also thank my son Greg for his patience and commiseration as I rant about scripting and VB, and for his help in learning how to overcome the hurdles.

StarZap Star Removal Method