A Successful Observation and a Curiosity

by Bob Buchheim (Altimira Observatory)

Those of you who were at the 2004 SAS Symposium will remember that Dr. Tim Castellano challenged us to observe the approximately 1.8% brightness drop of HD209458 when its planet ("b") transits the stellar disk. I am pleased to report that it can be done from your backyard, just as Tim said, if you do a little advanced planning, and are very careful with your photometry.

This measurement presents three interesting challenges. First, it is a very subtle effect – less than two hundredths of a magnitude – so it demands high signal-to-noise ratio. Second, the star is very bright – visual 8th magnitude – so that you'll use short exposures to avoid non-linearity in your CCD images of the target star. Third, the transit event is visible only approximately once a week, and you'll need to gather a continuous sequence of images for at least 2 (and preferably 4 or more) hours in order to have a good record of the pre- and post-entry brightness.

In order to get properly prepared, I took a series of images of the target star on "non-transit" nights, using various filters and exposures. My objective was to verify the linear range of my imager (SBIG ST-8XE, NABG), and find a combination of filter and exposure that would provide good Signal to Noise Ratio (SNR), reasonably long (!) exposure time, and good "efficiency" of collecting photons. This advance planning is particularly important because of the brightness of the target star: saturation of the image could easily swamp the tiny photometric signal of the transit. I used a Johnson-Cousins "B" filter, and 30 second exposures with my NexStar-11 operating at F/6.3. The "B" filter presents relatively low sensitivity (compared to "V" or "R"), permitting longer exposures. The 30-second exposure is beneficial for two reasons. First, it tends to average out seeing-induced scintillation. (Theoretical models show that at exposures shorter than 30 seconds, the effective SNR is dominated by scintillation, and that the effect can be really serious for ~5-second exposures that would be required with my set-up using "V" band). Second, the 30-second exposure seemed like a better use of photons than the shorter "5-second V band" alternative. This "efficiency" effect is based on the following:

action	"B" band	"V" band
science data exposure	30 sec	5 sec
image download	15 sec	15 sec
pause to allow autoguider to stabilize after image download	15 sec	15 sec
total cycle time	60 sec	35 sec
imaging efficiency = (science exposure/cycle time)	50%	14%

That is, with the 30-sec B band sequence, about half of the time is spent collecting science data frames; the "5 sec V band" alternative, only 14% of the time is spent collecting science data. My formal SNR was about 800 on each image.

Autoguiding is important – Tim's recommendation (validated by my experiments) is that even after good flat-fielding, a stellar image wandering around on the pixels can generate a fraction of a percent of counterfeit fading or brightening.

My data reduction was pretty standard: dark frame and flat fields, then photometry using MPO Canopus with 17 pixel diameter measuring aperture. The resulting raw data was pretty good, showing random scatter of less than 1% RMS. In order to smooth the results further, I ran a "top-hat" running average over it, using the equation:

$$\mathrm{avg}_i = (\mathrm{raw}_{i-2} + \mathrm{raw}_{i-1} + \mathrm{raw}_i + \mathrm{raw}_{i+1} + \mathrm{raw}_{i+2})/5$$

The net result of running that 5-point average is shown in Figure 1:

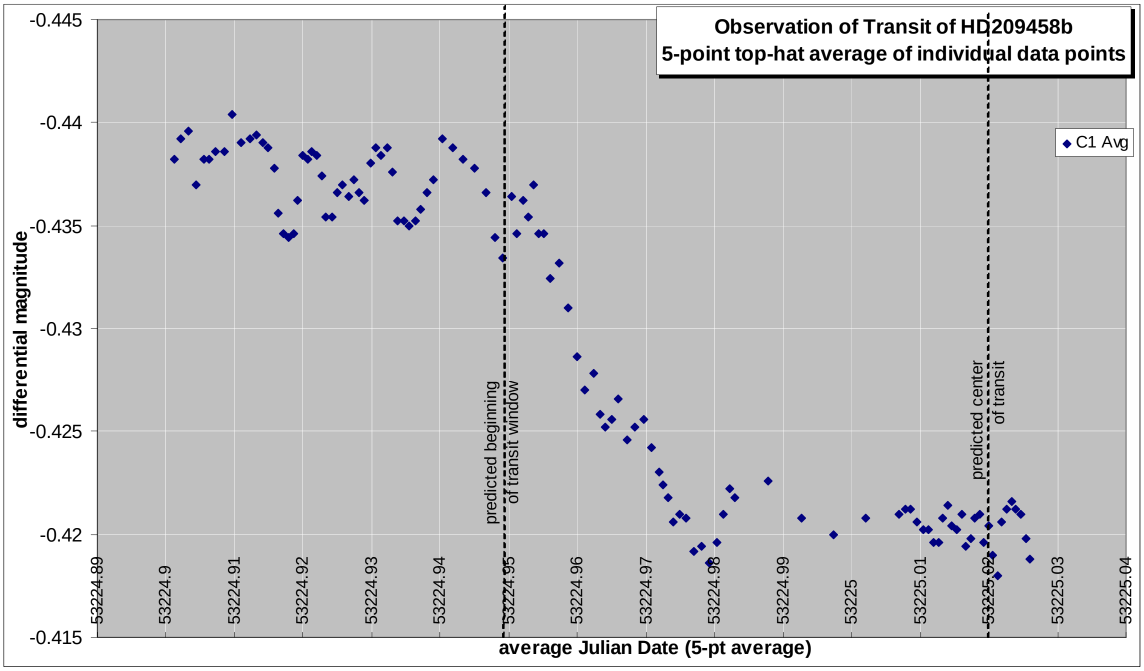

The 1.8% fading caused by the transit is plainly visible, and occurs at exactly the correct (predicted) time.

That's the successful observation. The curiosity follows –

I noticed a curious effect in the raw photometry of the target and comp stars. If I plot instrumental magnitude vs. time, I get the following:

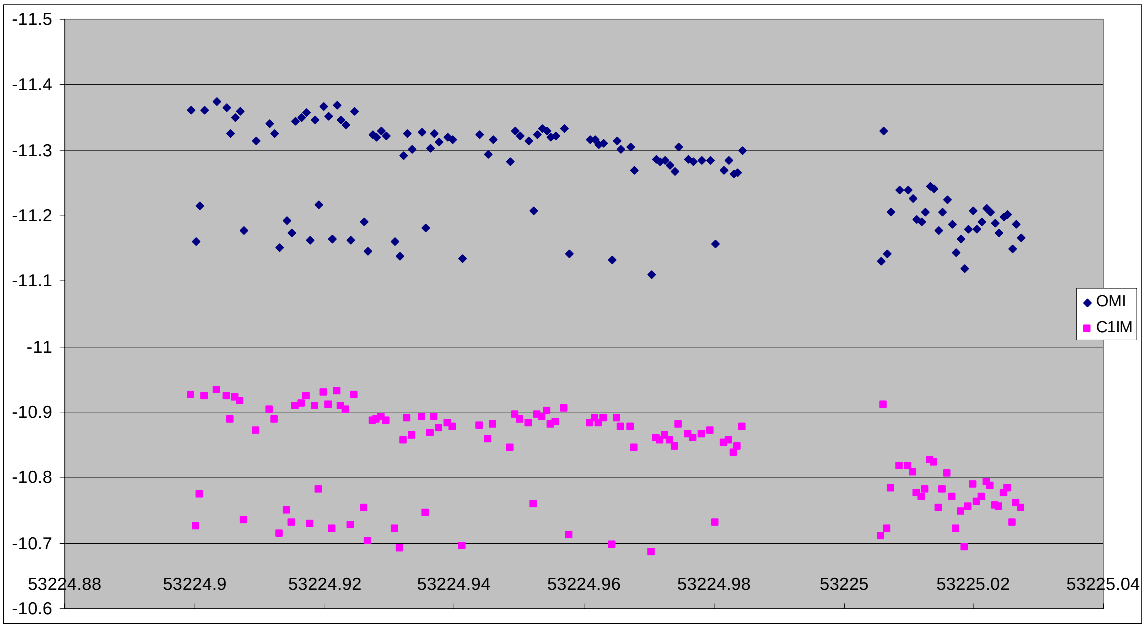

Both the target star and the comp star show the expected slight brightness reduction as they move away from the zenith, due to atmospheric extinction (the night seemed clear, and seeing was good, but I did not specifically measure the extinction coefficient). And the two stars move up and down in brightness in perfect synchronicity, as expected. Nevertheless, the curves are curious – each is almost a bi-modal distribution, in which the star spends most of the time at its “normal” brightness, but about 20% of the data points lie 0.1 magnitude below the main “normal” trend line. This doesn’t seem like noise or scintillation, because it’s so clearly bi-modal (noise or scintillation would randomly “fill in” a continuous random distribution). Ordinarily I would blame this sort of thing on wispy clouds contrails passing through the FOV, but it is hard to accept that theory when the brightness bounces up and down by 0.1 mag in one minute.

I suspect that the effect is caused by a slight variability in the exposure time (i.e. inconsistency in shutter timing in the camera). According to SBIG, their shutter should be accurate to about 5 msec (i.e. negligible effect). The best guess is that my PC (an old Pentium I) is so heavily worked running CCDSoft, autoguiding, and exposing images, that it occasionally skips a beat, resulting in slight variation in the exposure time.

This is another lesson that when you’re attempting photometry at the one-percent level, “everything matters”.