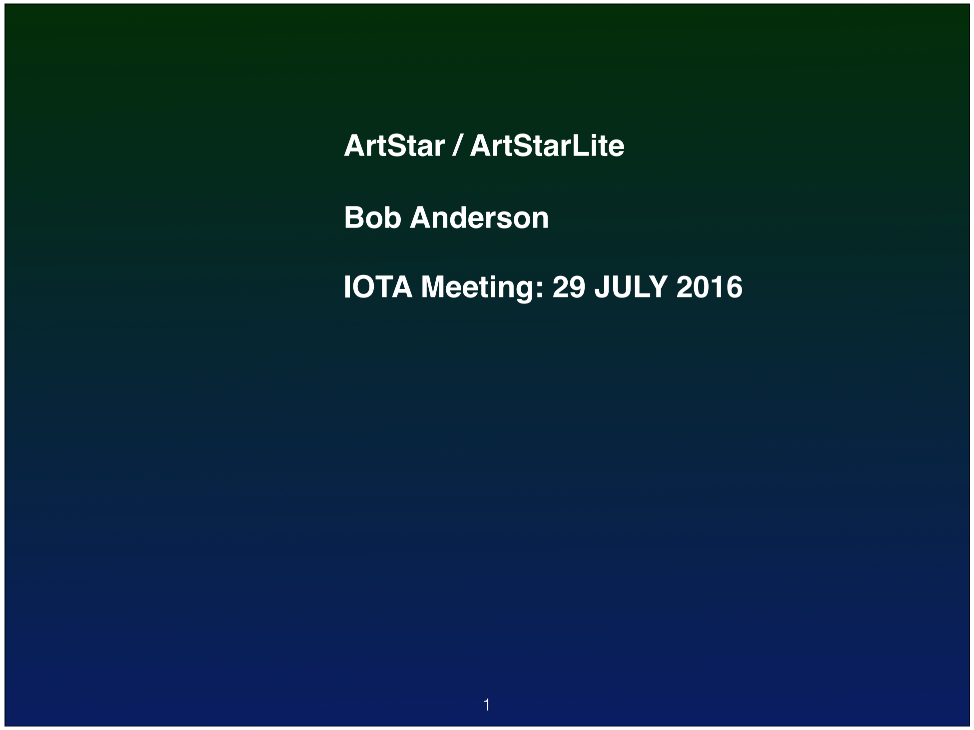

I'm Bob Anderson the programmer behind R-OTE and Occular.

This presentation is about 2 test systems for characterizing video cameras like the Watec910 together the frame grabbers commonly used to capture videos from these cameras.

My motivation for building these systems was to determine by direct measurement whether the mathematics employed by R-OTE is consistent with the actual behavior of typical cameras employed for occultation observations. That goal was quickly achieved, but the project took on a life of its own and has now evolved into a second, simplified system that I call ArtStarLite and is being introduced here for the first time.

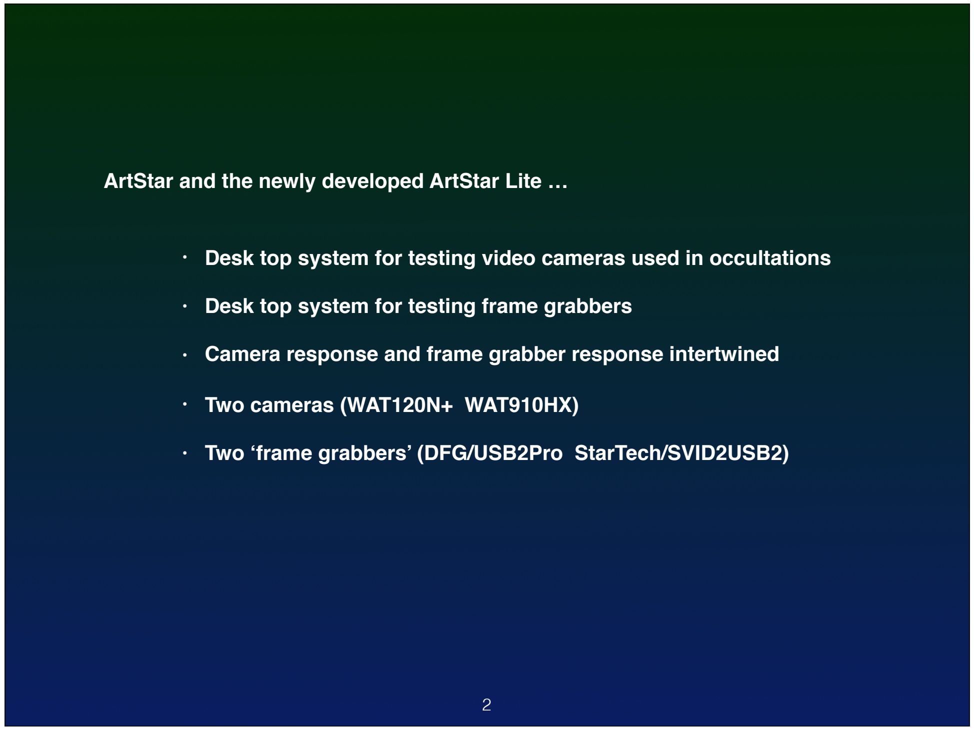

Both ArtStar and ArtStarLite are test systems to be used from the comfort of a desk chair with normal room lighting. **NEXT**

They are designed for testing video cameras like the W120N+ and … **NEXT**

… the frame grabber used to record the video. **NEXT**

I want to stress that in these systems, what is being tested is a camera and frame-grabber combination. So, when some unusual phenomenon is observed, the first question to be answered is — did the camera do this — or did the frame-grabber do this. **NEXT**

Two different cameras — both with Sony HAD chips — were the subject of the tests described in the slides to follow. **NEXT**

Two different frame grabbers were used: the DFG/USB2Pro (by Imaging Sources) is a relatively expensive ($200 or so) frame grabber that features uncompressed output and automatic gain control (AGC). The StarTech SVID2USB2 is a relatively low cost ($50 and change) frame grabber that is in common use, produces mp4 compressed output and does not have automatic gain control. I will talk more about frame grabber automatic gain control toward the end of this presentation. NEXT

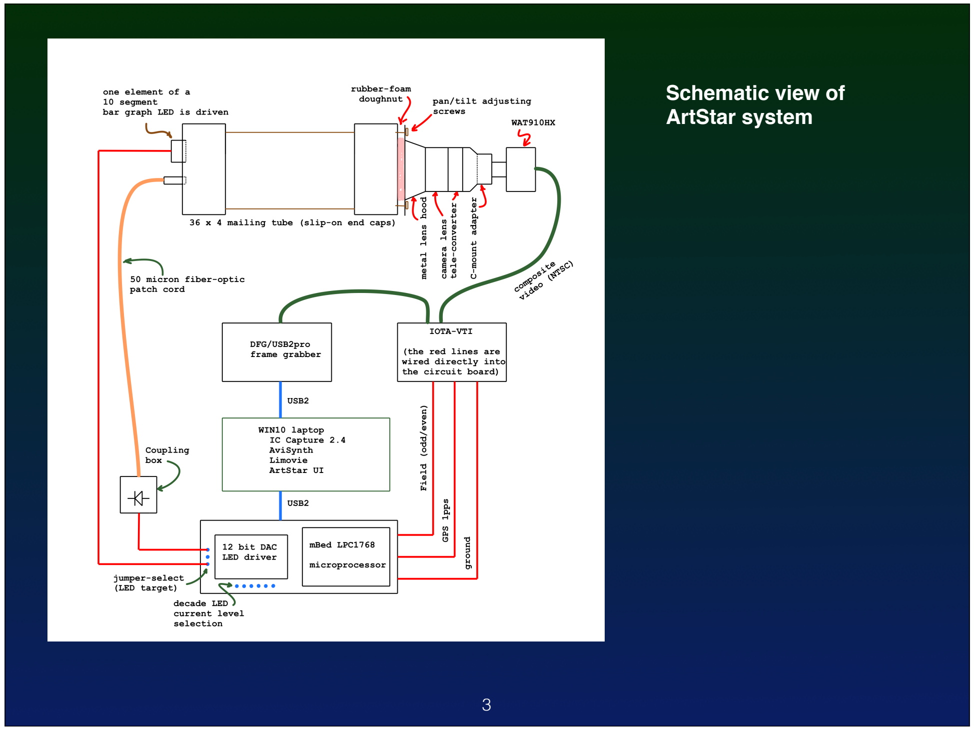

Here we see a diagram of the major components that make up ArtStar.

Across the top and down the lefthand side of this diagram are the mechanical parts. There are a lot of mechanical parts — I stress this because the ArtStarLite system that I’ll describe later eliminates nearly all of them

The remainder of the diagram comprises the electronics portion of the system. The key piece is at the bottom of the diagram. That is where a microprocessor, synchronized to the camera field/frame times, controls the illumination of two LEDs.

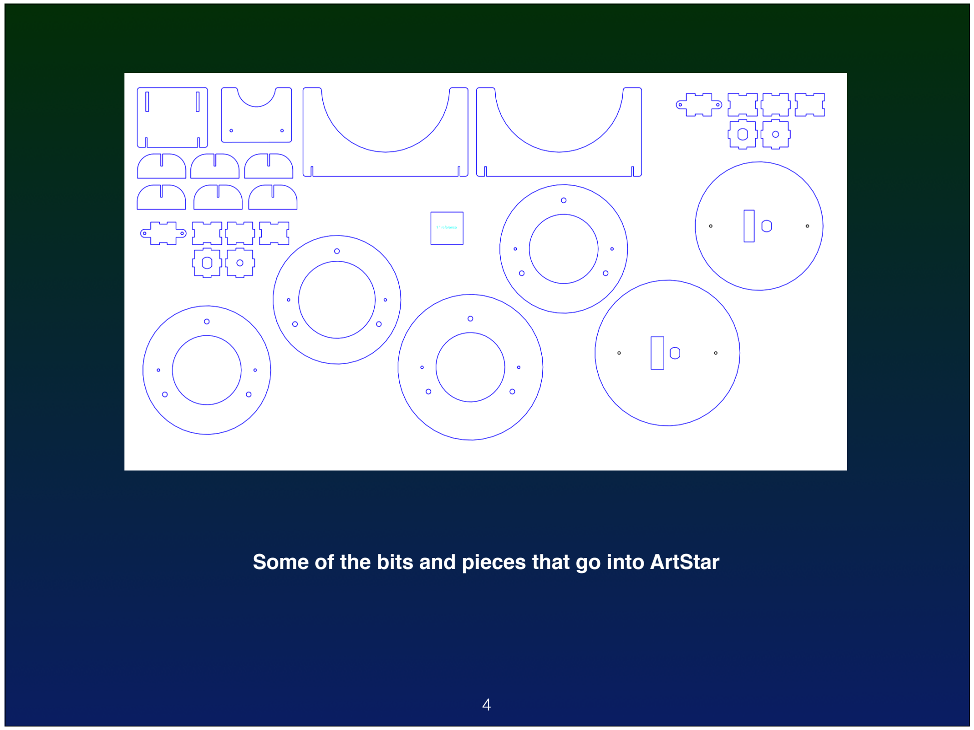

There are lots of bits and pieces that go into ArtStar — these are just the wooden parts. ArtStarLite doesn't need any of these parts.

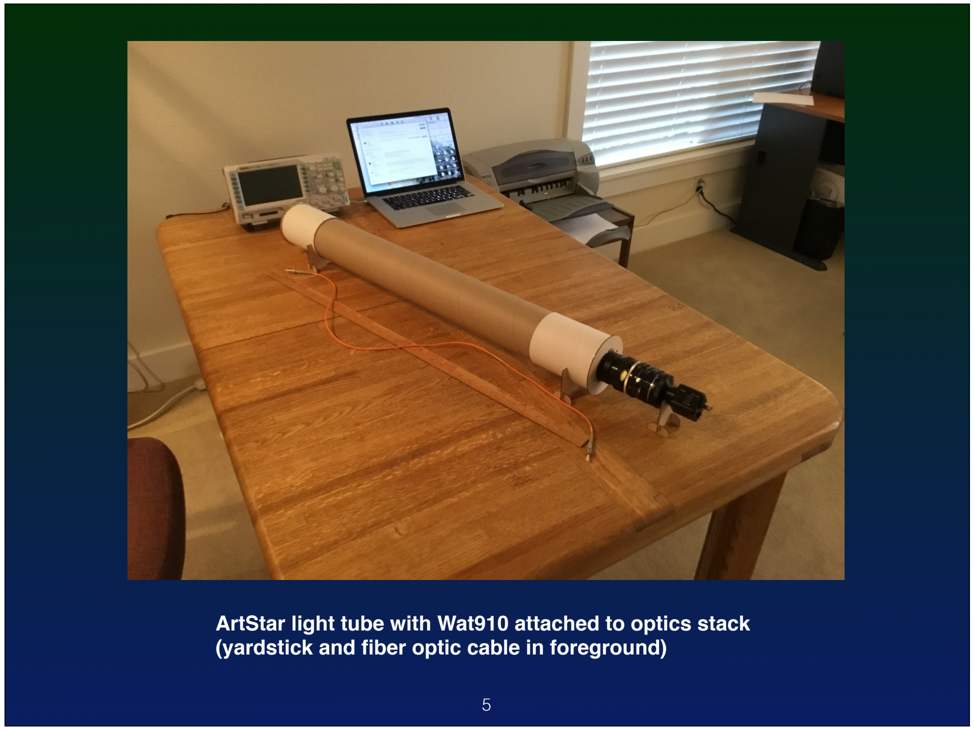

Here we see the mechanical parts of ArtStar assembled on my desk waiting for electronics to be attached.

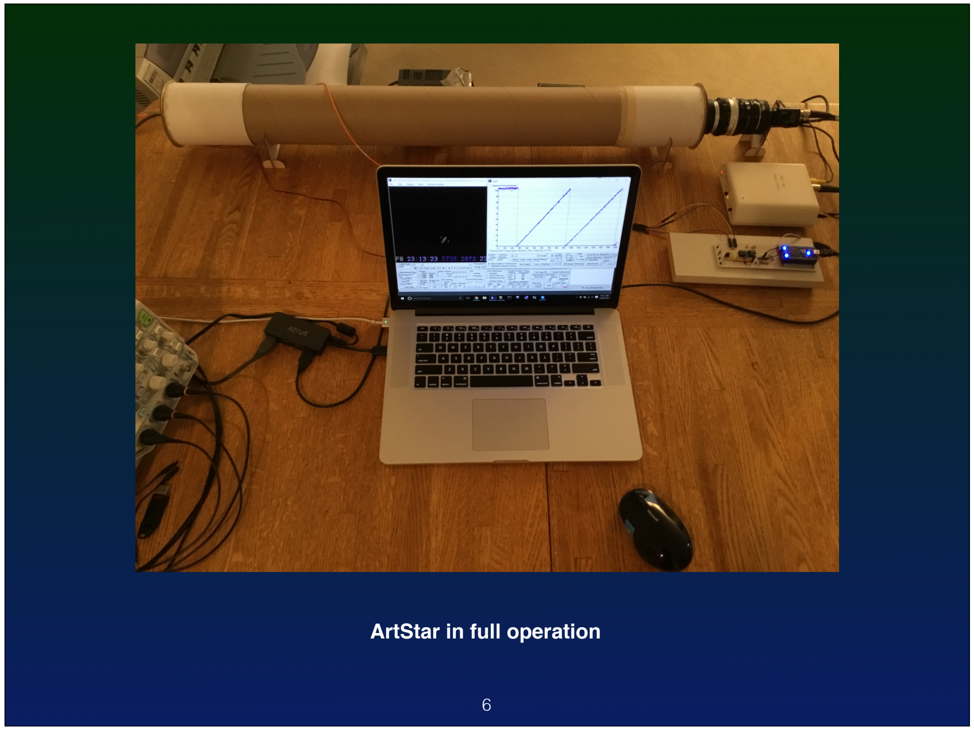

And here we see ArtStar in operation.

The digital oscilloscope peeking in from the left was used during hardware development to confirm pulse timing.

The next slide shows what the camera is seeing as it peers up the mailing tube …

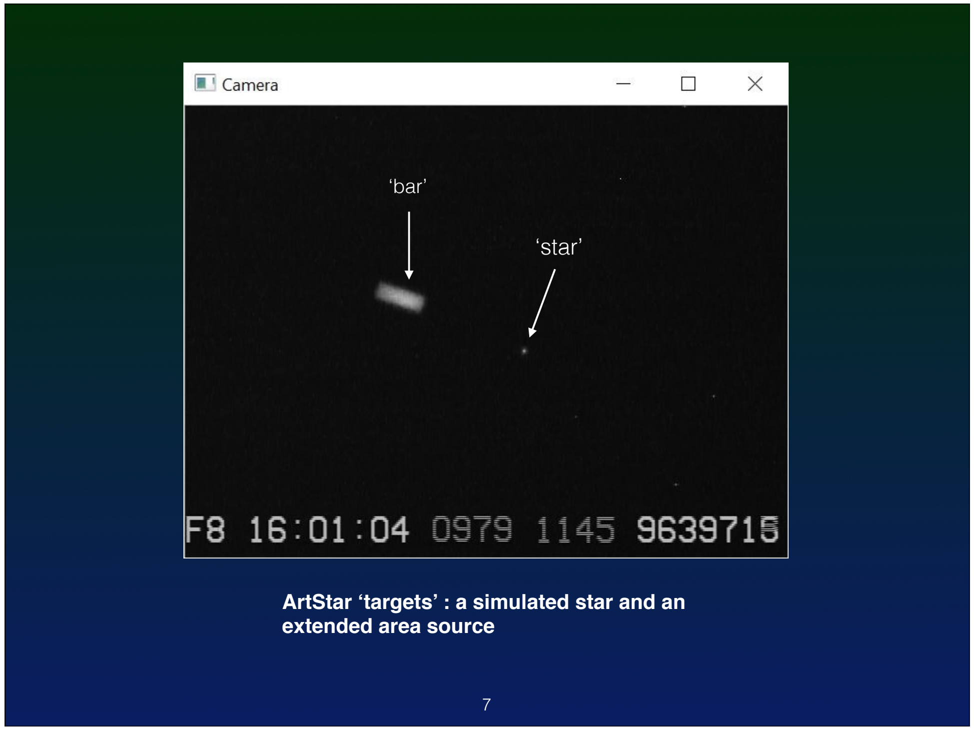

This is what the video camera sees.

The ‘star’ shines from the output end of a 50 micron fiber optic cable. The ‘bar’ is one segment from a 10 segment bar graph LED.

The microprocessor is able to control the brightness of these targets on a field by field basis using the scheme shown in the next slide …

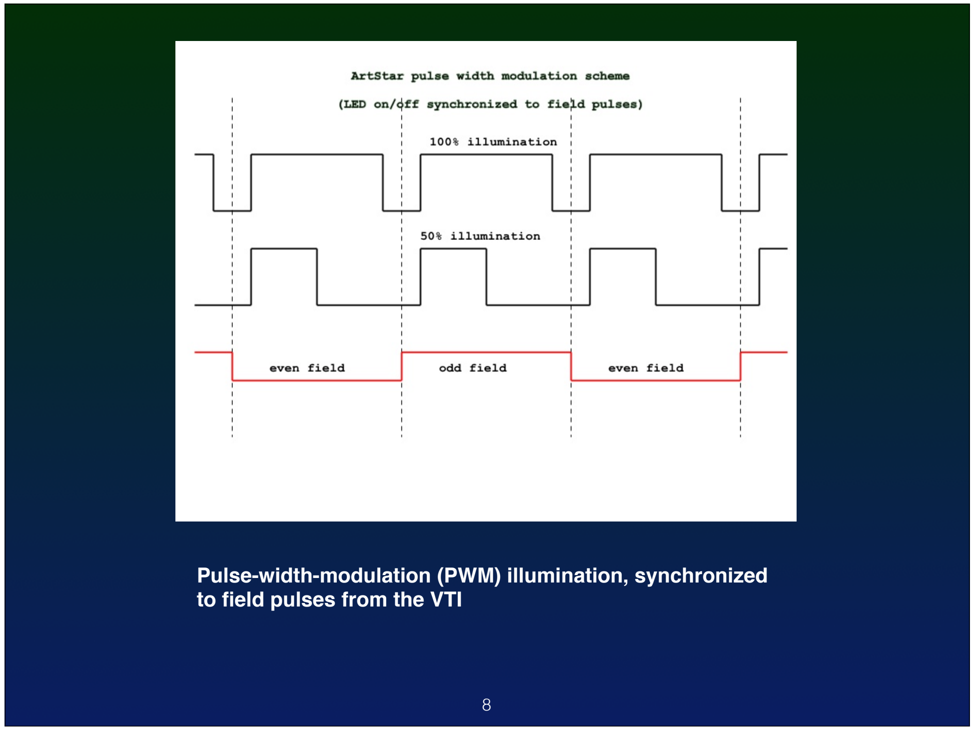

The key idea in controlling target illumination in ArtStar and ArtStarLite (which uses the same electronics) is shown in this slide.

The red pulse train at the bottom shows the field timing waveform extracted from the composite video by the IOTA VTI.

The microprocessor controls when the LED that lights a target will turn on and off. To get half the number of photons delivered to the camera, we simply cut the time the LED is on in half. This is an inherently linear process, provided one can avoid heating effects in the LED, and that is easy to do here because typical currents delivered to the LED are a few microamps and less — the cameras don't need much light, so the LEDs don't need much current.

The next slide shows the response of the W120N+ to LED illumination that linearly increases from 0 to 100% using this scheme …

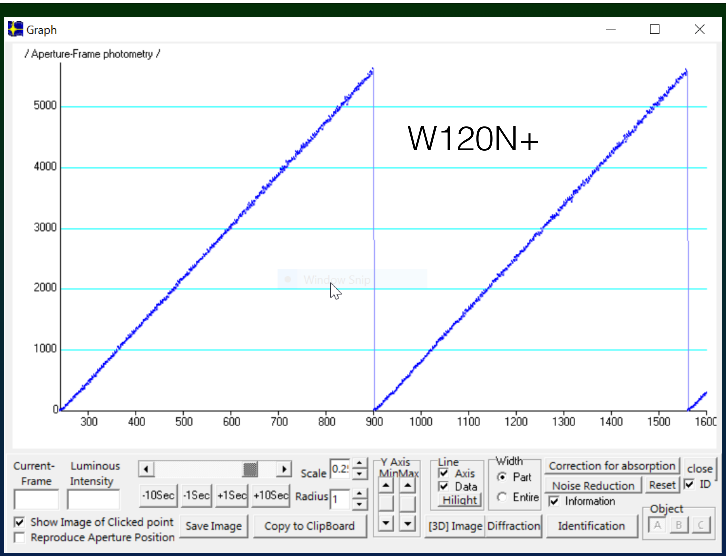

Camera/frame grabber response to linearly increased LED pulse width (ramp in time — LED pulse width starts at zero and incremented at each successive field until max pulse reached, then starts over)

9

The DFG frame grabber is in use here.

ArtStar can also modulate the current to the LEDs rather than on-time.

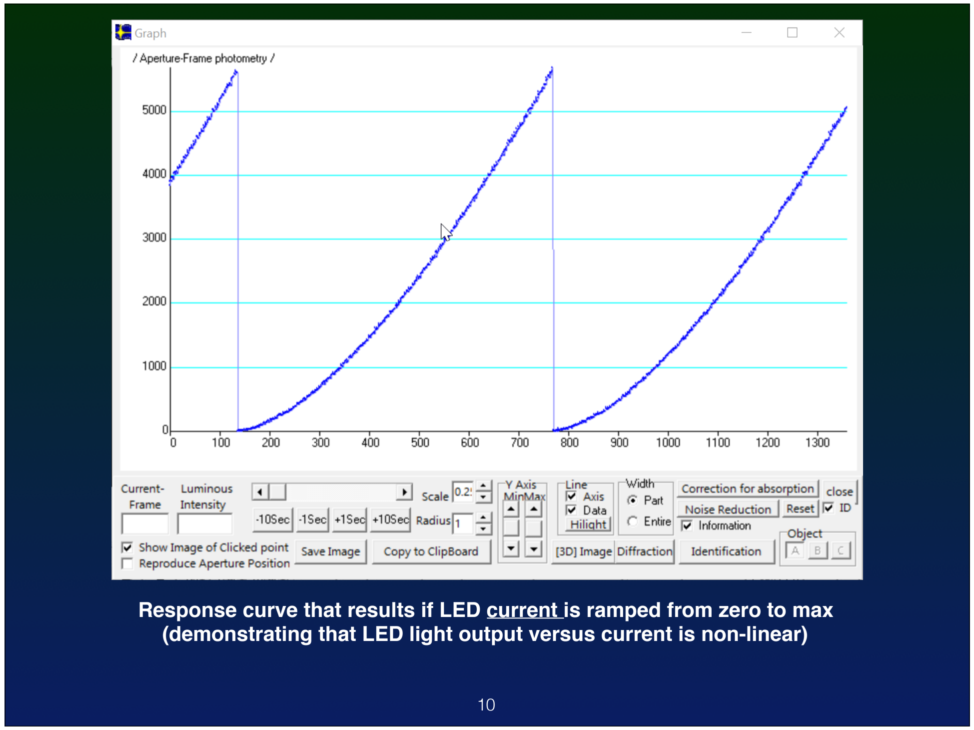

Response curve that results if LED current is ramped from zero to max (demonstrating that LED light output versus current is non-linear)

This was recorded using the W120N+ with the DFG frame grabber and shows that light output of the LED used here is a non-linear function of current.

But the previous slide showed that ArtStar can provide field by field linear illumination when run in pulse-width-modulation mode.

At this point, I switched to using a Watec910 as the camera-under-test — and the first surprise appeared … actually, two surprises.

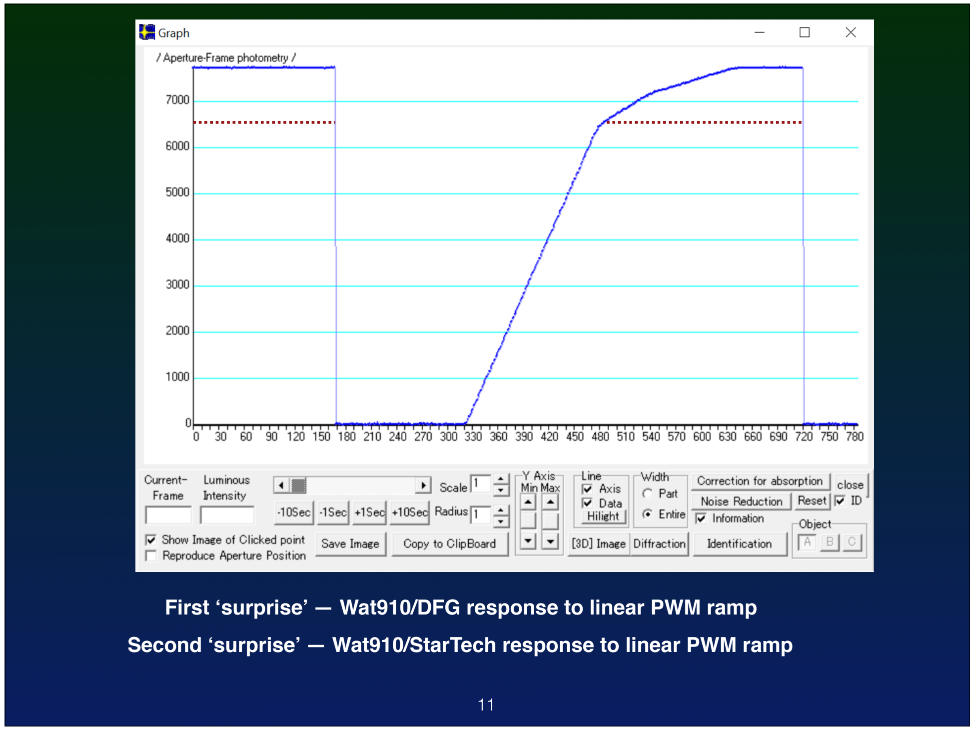

First ‘surprise’ — Wat910/DFG response to linear PWM ramp

Second ‘surprise’ — Wat910/StarTech response to linear PWM ramp

This was a surprise: the ‘kinks’ in the response curve were unexpected — so much so that I was concerned that the frame-grabber might be responsible. So I switched to the SVID frame-grabber and the ‘kinks’ ALMOST disappeared, but not quite (**the brown dotted lines show the light curve recorded with the SVID**). Subsequent measurements of the amplitude of the composite video signal from the Watec910 showed that it was larger than the expected standard 1 volt peak-to-peak: the SVID frame-grabber does not have AGC to compensate for this (a later slide will demonstrate this) and so was clipping the output; the DFG frame-grabber has a very effective AGC and so records the full camera output.

Now the question was: were the ‘kinks’ particular to my camera — or was this response curve typical — time to talk to the manufacturer.

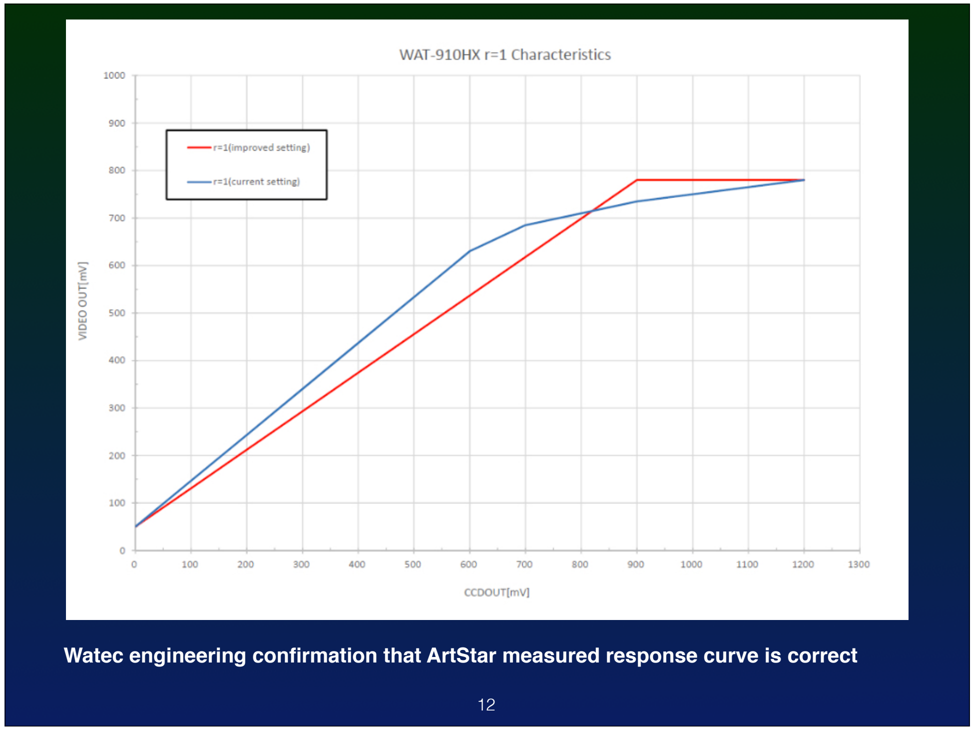

A communication with the vendor resulted in a prompt and helpful response from an engineer involved with the 910 that confirmed that the ArtStar measured response curve was consistent with their design intent. He added that, if a completely linear response curve would be more useful for our application, they could provide the red response curve if desired. This option would have to be selected at order-time and there would be a customization fee involved.

Now I’m going to switch briefly to a frame grabber/compression discussion …

Most frame-grabbers produce a compressed video. In fact, it's quite difficult to find a frame-grabber that will produce an uncompressed video stream.

The DFG/USB2pro frame-grabber does output an uncompressed video.

The SVID2USB2 frame-grabber produces an mp4 compressed video.

One of the important questions to be answered is whether this compression affects occultation timing.

READ THE SLIDE

To help answer the last question, ArtStar was designed to be able to ‘replay’ an actual occultation.

This is shown in the next slide.

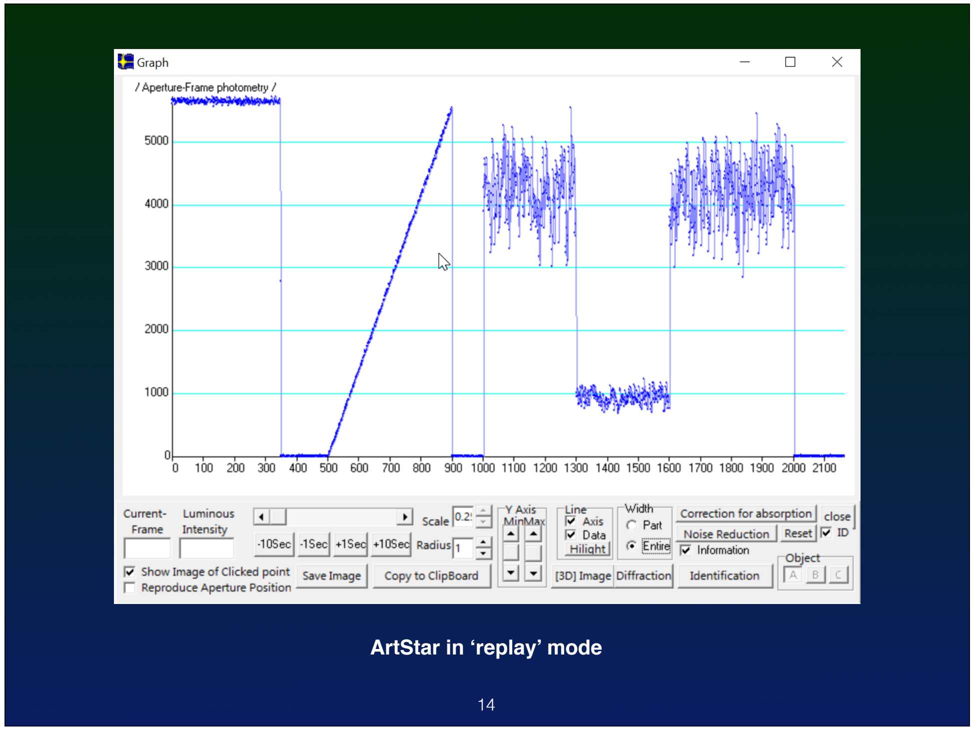

In 'replay' mode, ArtStar begins with maximum illumination. It then plays a linear illumination ramp — this lets one know that the camera is operating in linear mode (or not). After a brief pause, an actual occultation light curve is used to modulate the illumination frame by frame.

When such curves are recorded with and without frame-grabber compression, subsequent D and R estimations from ROTE can be used to determine whether the compression adversely affected D and R times.

Limited testing performed to date indicates that there is no statistical difference in D and R times extracted from mp4 videos versus uncompressed videos.

Which brings me to the next section of this presentation and an area that caused me a lot of struggles: micro-lens effects and what is known in the literature a 'sub-pixel response'.

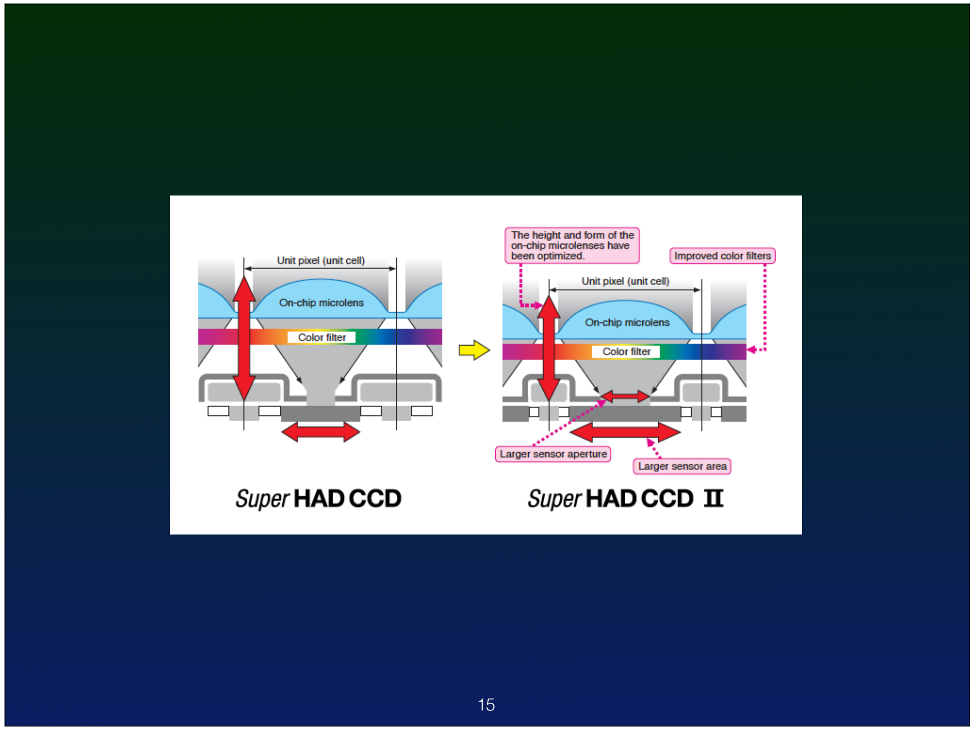

Sony has shown diagrams such as the one shown in this slide (from a white paper published by Sony about their HAD and Super HAD devices).

The point to be made here is that the unit cell that we think of as a pixel is a complicated structure.

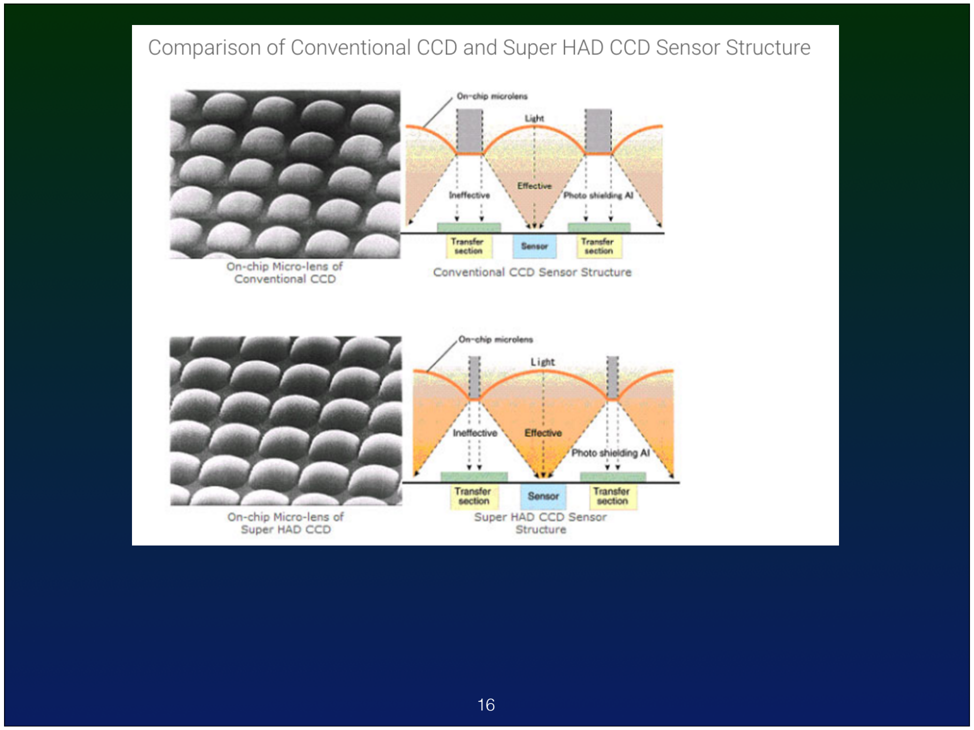

This diagram foretells the phenomenon known as 'sub-pixel response'. By 'sub-pixel response' is meant the variation in electrical response to an incoming photon based on precisely where that photon hits the pixel — some areas are completely unresponsive — others are partially responsive in comparison to the response to a photon that hits squarely in the center of the photo-sensitive region.

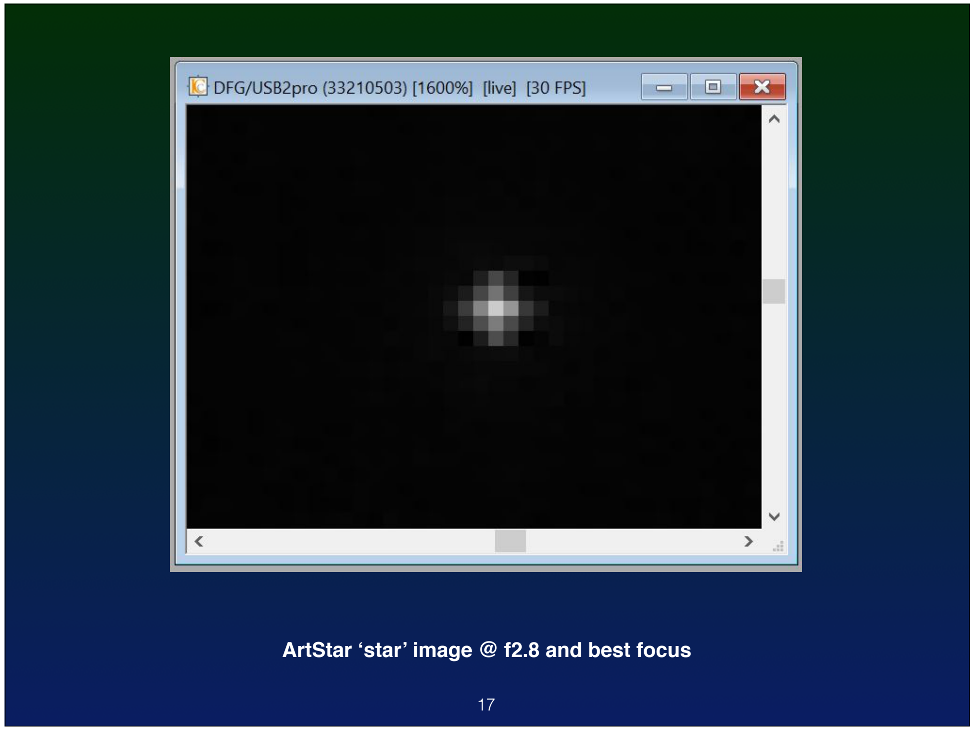

Here we see an image recorded of a best-focus star. Calculation of the 'star' size using the ArtStar optics show that the 50 micron diameter 'star' projects onto the ccd array as a 4 micron spot.

As the unit pixel in the W120N+ and the W910 is 8.4 microns by 9.8 microns, it should be possible to light a single pixel — but it isn't — as this slide shows. This is more likely a micro-lens effect rather than due to charge diffusion (which does affect back-lit ccd arrays).

Under these conditions, only a single pixel should be lit. We should have gotten something that looked more like a hot pixel. Instead, neighboring pixels are quite well lit. I'm guessing that this is a scattering effect caused by the micro-lens structure, but I could find no way to test that hypothesis, so at the moment that conjecture must remain an opinion.

To investigate micro-lens structure phenomena, ArtStar is equipped with positioning screws that have fine enough pitch that the spot can positioned with sub-pixel precision. Those positioning screws are shown in the next slide..

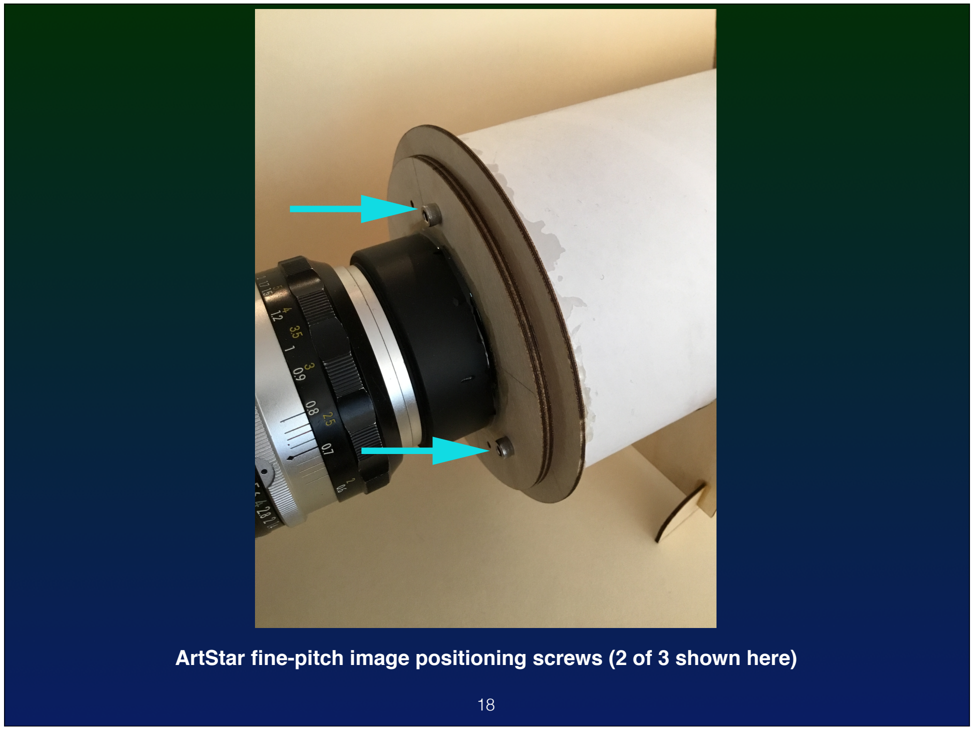

With these positioning screws, it is possible to move the ‘star’ image slowly across the pixel array. When one does that, you obtain a light curve like the one shown on the next slide…

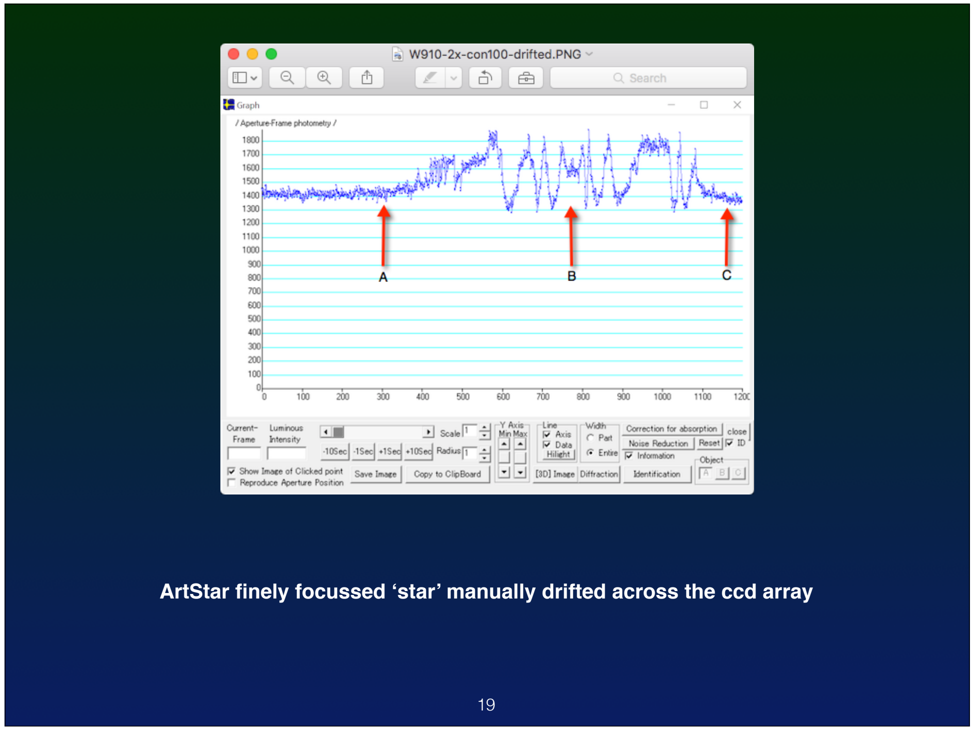

The procedure that produced the light curve on this slide was this:

The recording was started (unsaturated star image — illumination is in linear range) and at point A I began slowly twisting one of the positioning screws.

At point B, I reversed the twisting of the position screw until I reached the original starting position at point C

This result caused me to conclude that my artificial star was <u>**too**</u> artificial — in field observations through a telescope, the psf of a star image (whether caused by the telescope optics or ‘seeing’) extends over multiple pixels. That mitigates the effects of this ‘sub-pixel response’.

A natural question is: what should the FWHM (full-width-half-maximum) of a psf be to minimize the effects of this ‘sub-pixel’ response characteristic. I did the following little study to get some feeling for this effect …

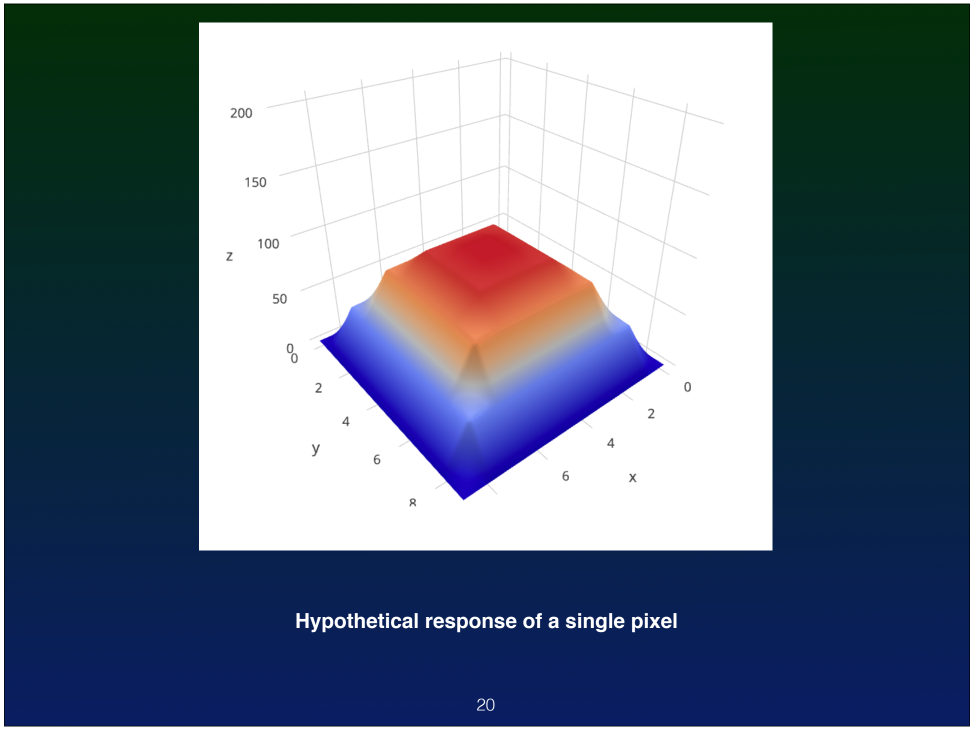

I guessed at the sub-pixel response based on the electron microscope pictures shown in an earlier slide.

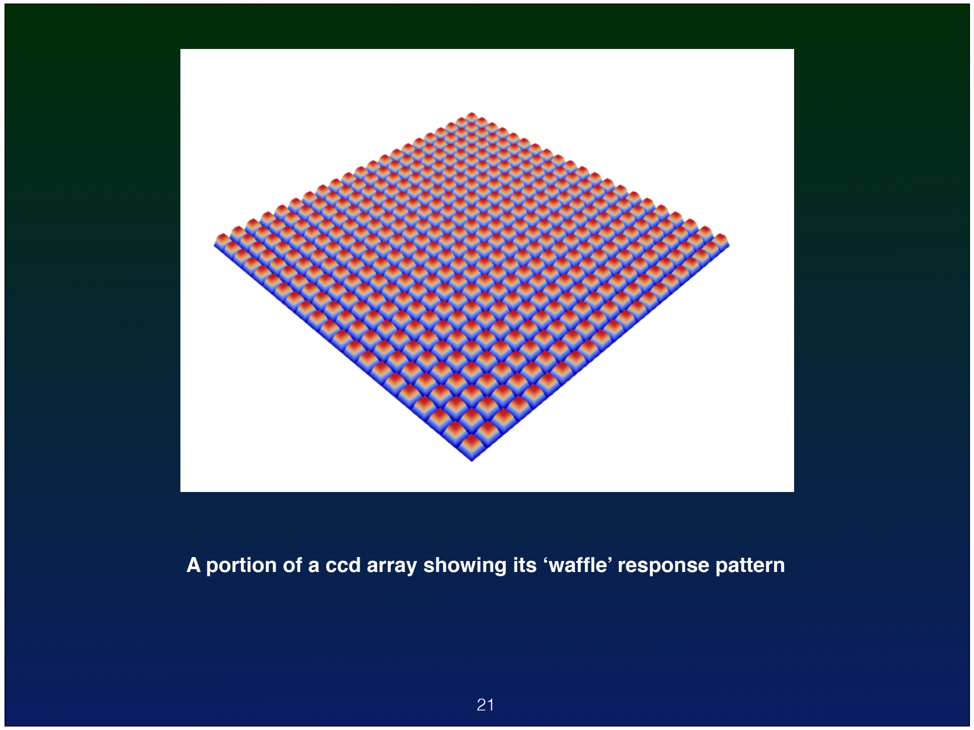

This is just a pretty picture showing my hypothetical pixel replicated in a 21 by 21 array. Actually, it wasn’t just a pretty picture — the plan was to project a point spread function onto this response pattern to see what the effect of the response variations across a pixel would have on the cumulative observed response.

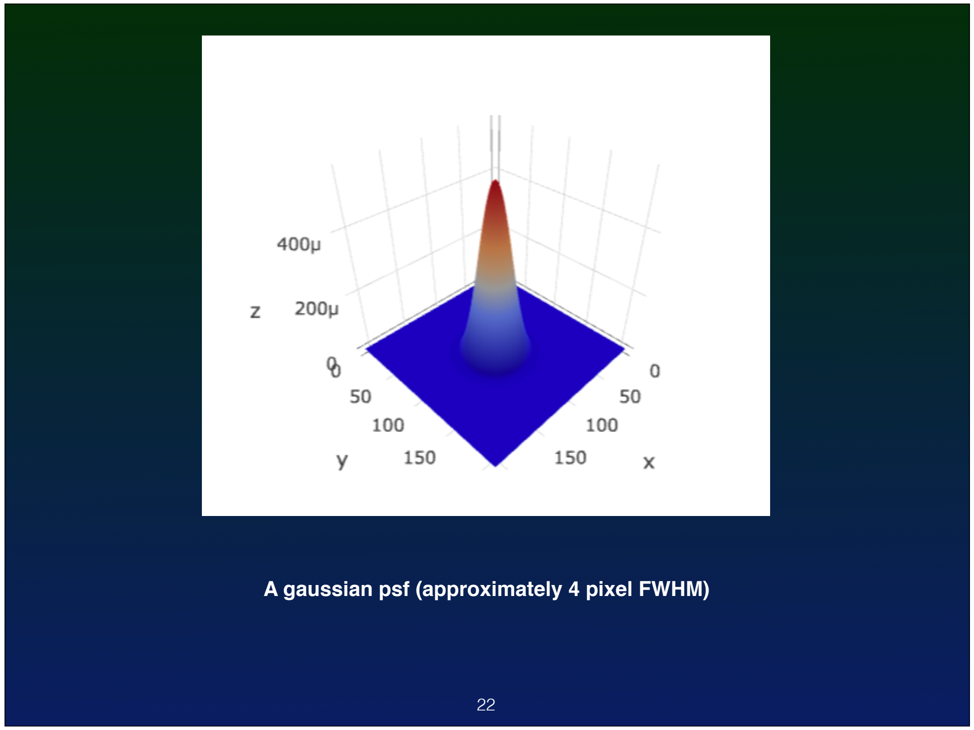

Here is a plot one of the point spread functions that I used.

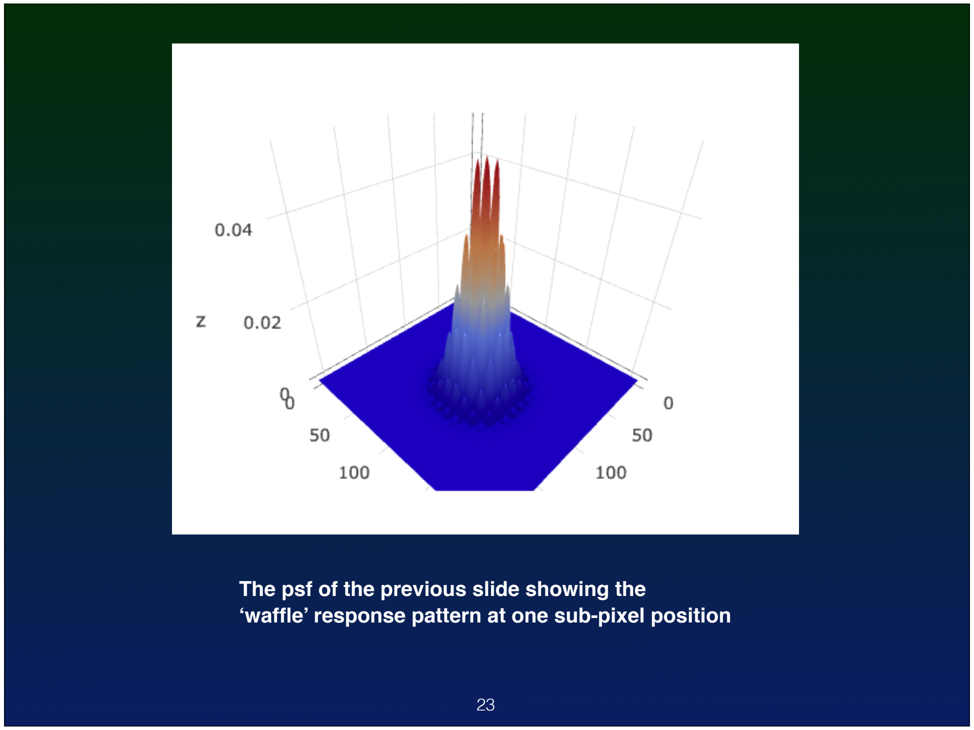

And this is what happens when you convolve the point spread function with pixel array responsivity.

The next step is to position the center of the psf over a unit pixel divided into 100 regions —at each sub-pixel position, a numerical integration of the entire psf is performed to simulate the response to this width point spread function. From those 100 values, the minimum, maximum, and average were recorded.

This process was repeated over a range of FWHM sizes and plotted as shown in the next slide.

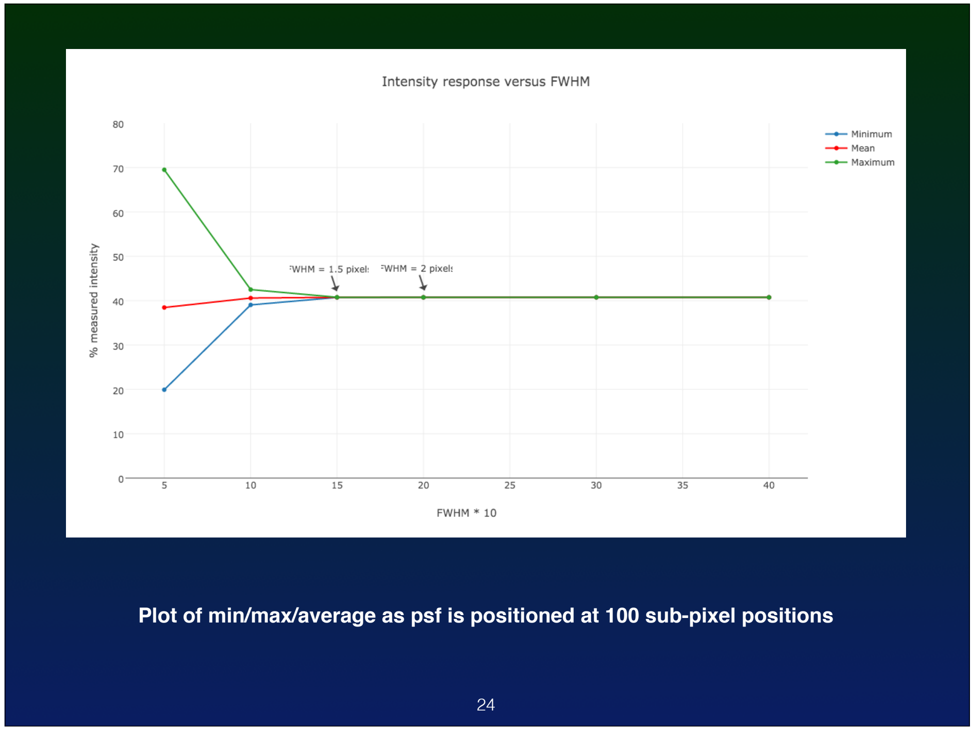

This slide suggests that sub-pixel response effects do not affect camera output whenever point spread functions are 2 pixels or larger in half-maximum width. This is commonly the case with telescope, camera, focal reducer, and seeing combinations encountered in occultation observations.

Note that this simulation does not model effects of the micro-lens structure. I personally believe that the micro-lens structure produces a scattering effect that is not well captured by the simple model used in this simulation. Even so, this plot went a long way toward explaining the light curve variability shown in the earlier slide where the positioning screw was used to traverse the star image across the pixel array.

I concluded that something had to be done about the ArtStar 'star' — it was just too small and well-focussed.

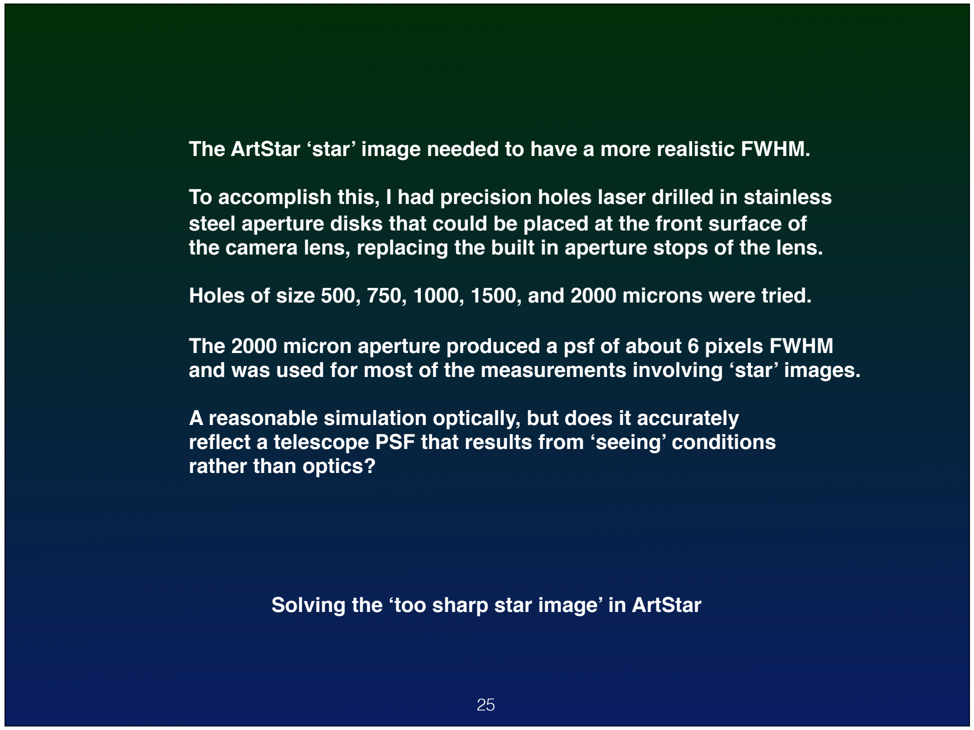

READ THE SLIDE

The following slide shows psf achieved with the 500 micron aperture at the front of the camera lens.

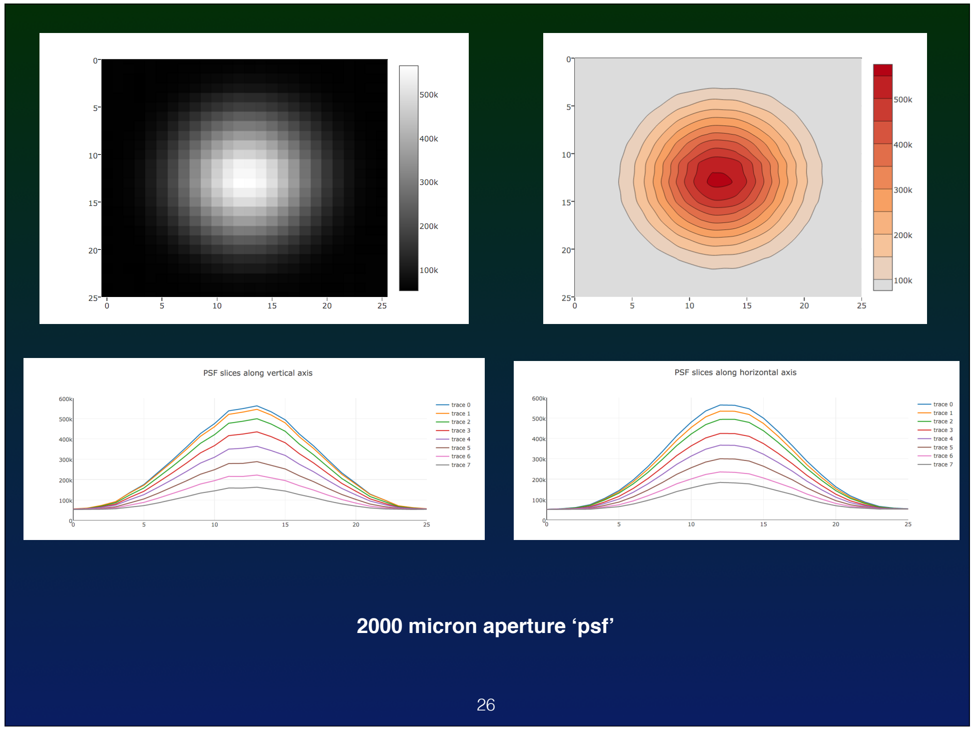

Please note: this slide is mislabelled as being from the 2000 micron aperture test. It should say 500 micron instead. The FWHM for this pdf is about 13 pixels.

With a 2000 micron aperture in place, a star image with a psf in the 6 pixel FWHM range results. Such a psf looks more like that obtained in field observations.

As a reminder, this is the star image before the installation of the pin-hole aperture system …

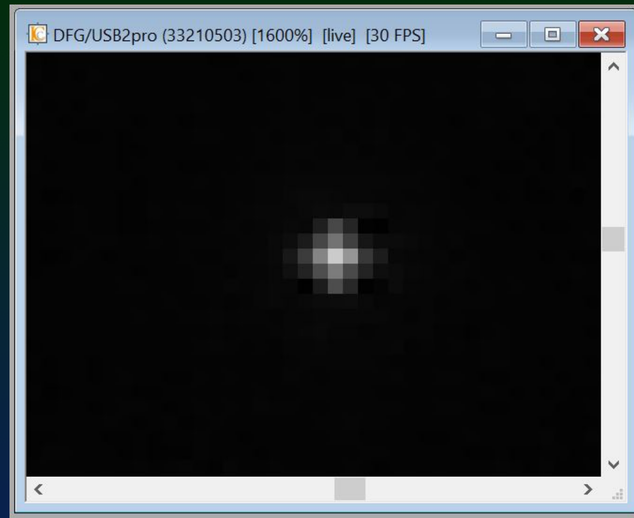

ArtStar ‘star’ image @ f2.8 and best focus

27

A more detailed report on ArtStar is available at …

yahoo.groups.com

IOTAoccultations

file folder: Light-curve-analysis-papers (andersonbob1024)

file: ArtStar-device-report.pdf

If you don't have access to that group, I can email you a copy.

My email is: bob.anderson.ok@gmail.com

ArtStarLite simplifies the mechanics, retains/reuses the electronics ArtStar, and is able to make all of the measurements that ArtStar does **except** for those that depend on the specifics of a star-like target.

Here are the mechanics of ArtStarLite ..

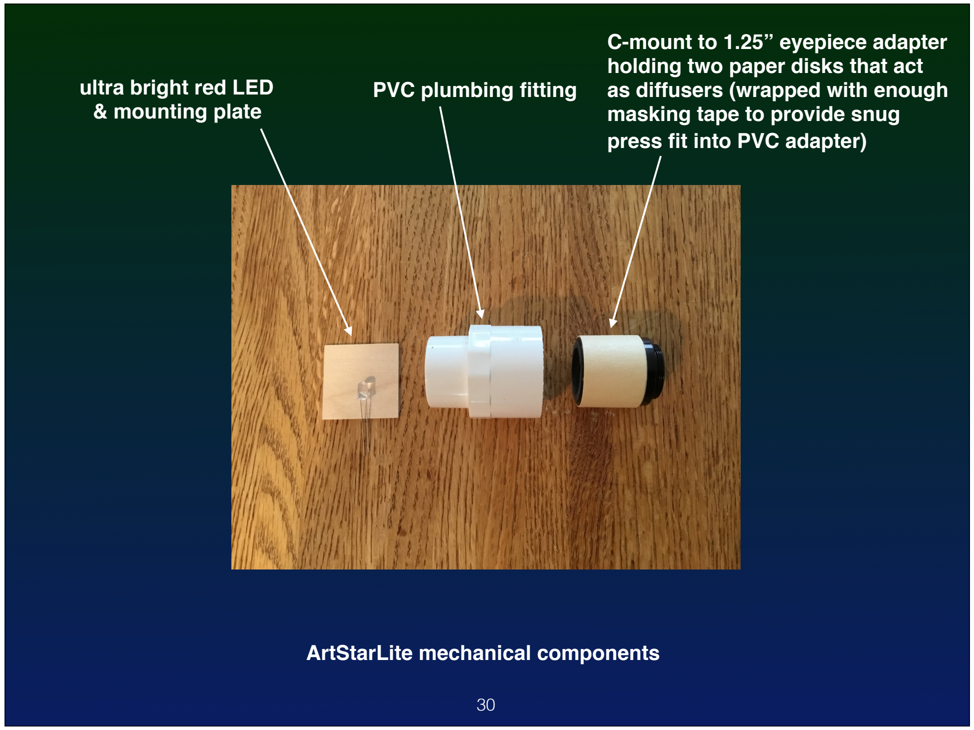

Just a few easily obtainable parts make the mechanical aspect of ArtStarLite.

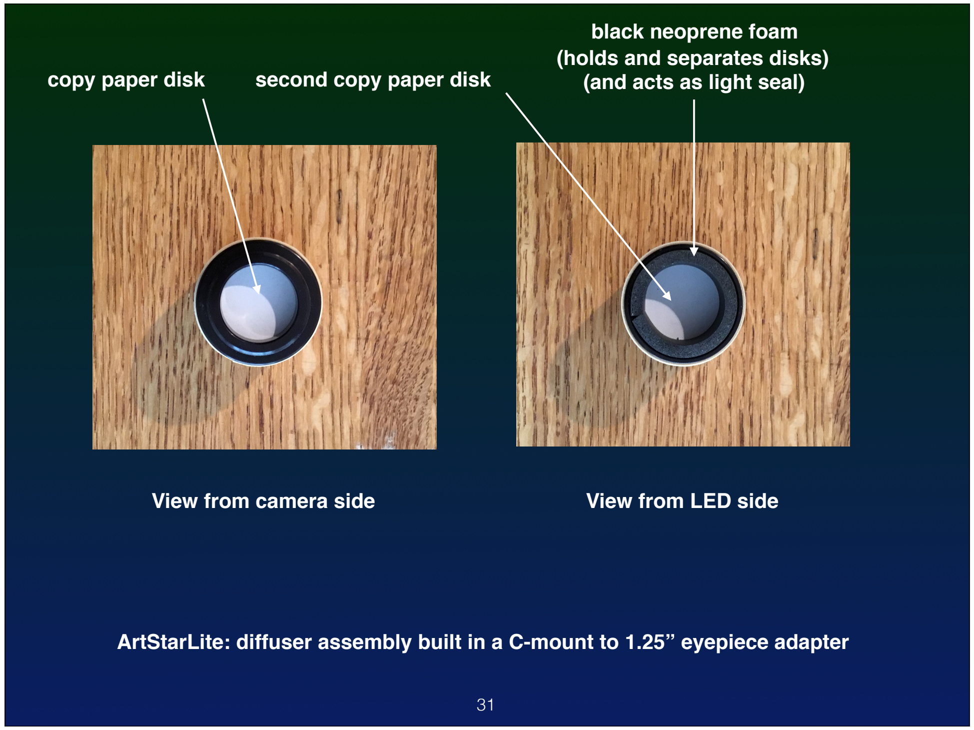

The C-mount adapter was turned into an effective light diffuser using two disks made out of copy paper.

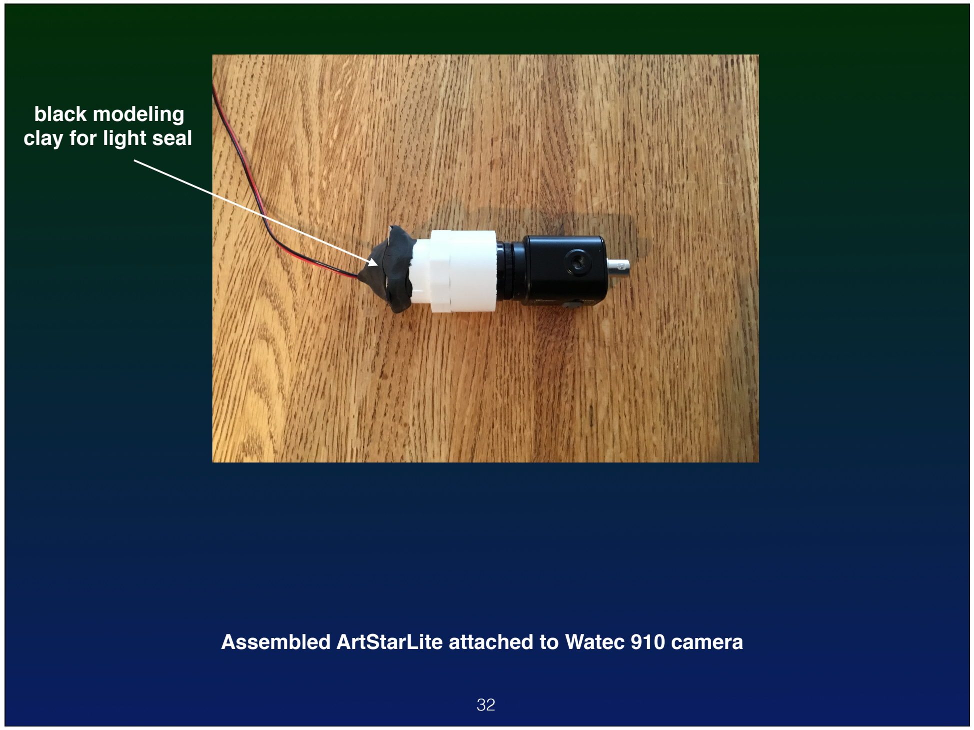

I'd like to note in passing that ArtStarLite could even be mounted so that it shines through the lens cover of a telescope. That would result in a very fuzzy blob being imaged by the telescope onto the ccd array, but the same measurements would be possible except now they will be in-situ measurements. This might be important if there are other optical elements (like filters maybe) that need to be tested simultaneously and don't otherwise fit the mechanical model of either ArtStar or ArtStarLite.

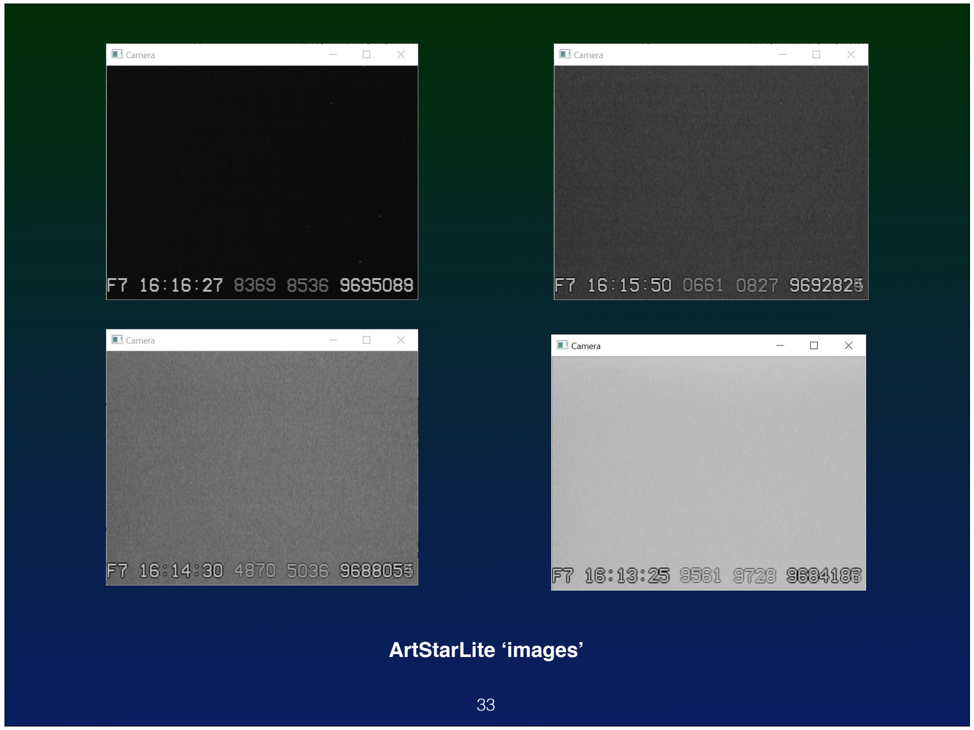

The images of ArtStarLite are just 'flats' whose intensity is under microprocessor control on a field by field basis exactly like ArtStar.

The next slide shows another 'surprise' that I encountered in the testing of this device ..

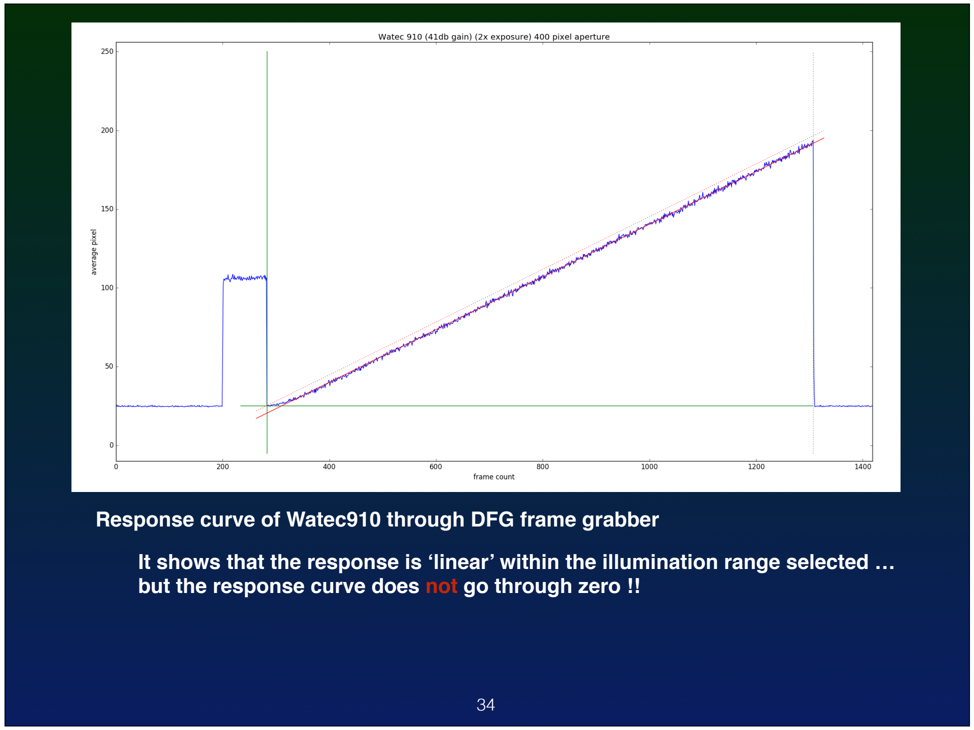

The max illumination was selected so that the 910 remained in its linear range during the ramp portion of this test. But the surprise is that the response curve does not go through zero.

One would expect that photometry is going to be affected by the fact that the response curve does not go through zero.

In effect, there is a small amount of light that reaches the camera but does not produce a corresponding output —it is ‘thresholded’ out of existence.

I argue that this characteristic of the W910 response curve explains something previously reported by Hristo Pavlov, but was mysterious at the time…

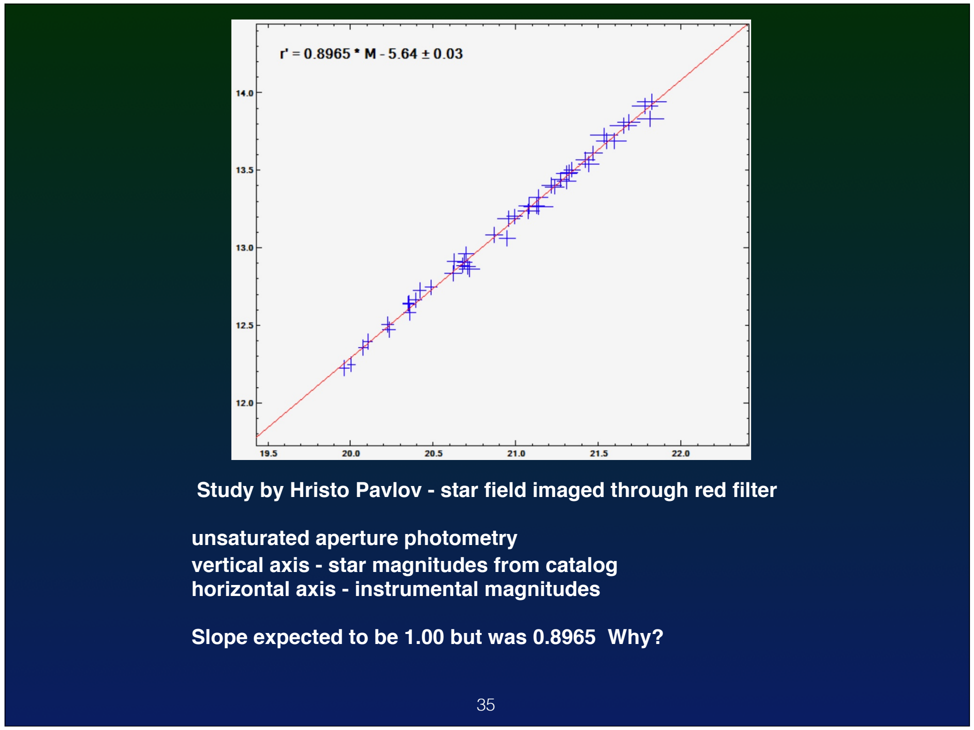

DESCRIBE HRISTO'S TEST (red filter used so that star catalog magnitudes could be legitimately compared)

The 'mystery' was: why does the straight line above not have a slope of 1.00?

What the straight line in this slide demonstrates is that if the star catalog shows a magdrop of 1.00 between two stars, the instrumental magDrop records a 1.115 magDrop between those same two stars.

This effect can be simulated using ArtStar with the 2000 micron aperture (6 pixel FWHM) as shown in the following slide…

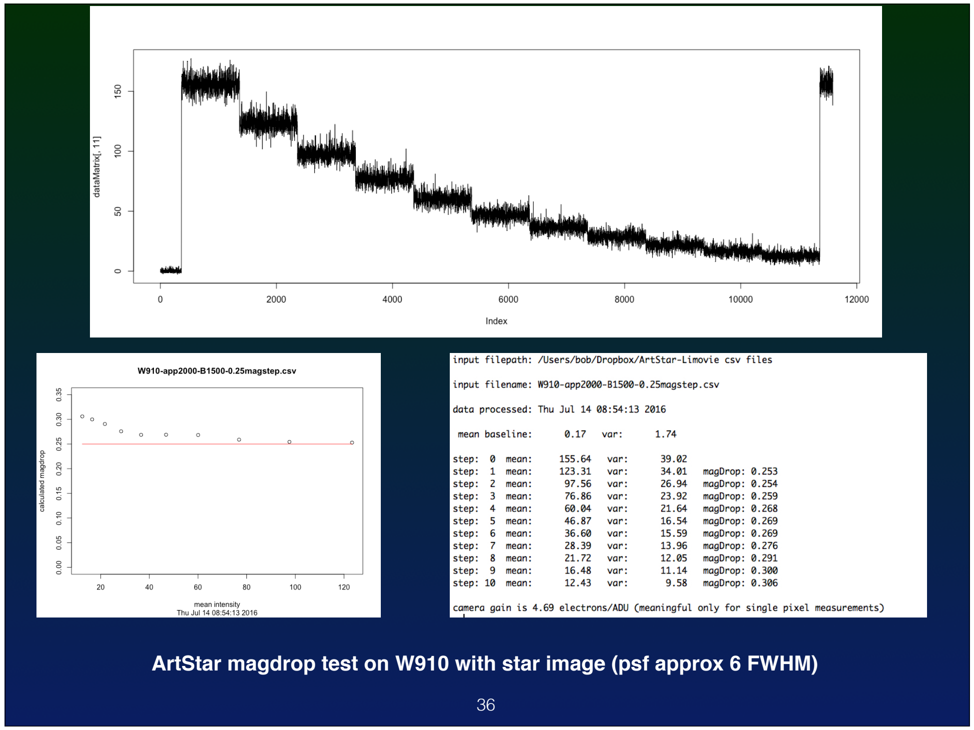

The top plot shows the light curve resulting from 10 steps of 0.250 magDrop each. The initial intensity was adjusted to be in the linear range of the W910.

Note that Hristo's plot covers a 2.0 magnitude change while this test covers a slightly larger 2.5 magnitude drop.

The text box below and to the right shows the instrumental magDrops calculated from the light curve.

The plot to the left shows graphically the measured magDrops compared with the expected value shown by the red line.

This is the same phenomenon depicted in Hristo's plot: when the catalog magDrop is 1.000 (that is, the ArtStar illumination is changed equivalent to a 1.000 magDrop), the measured magDrop (instrumental magDrop) is always greater.

This is due to two things: first, the psf having wing pixels with intensities close to zero and second, the W910 response curve not going through zero.

To demonstrate that shape matters, I used ArtStarLite to repeat the test, measuring the response from a 16 pixel block. Of course, these pixels are uniformly lit, so there are no 'wing' pixels. The following slide shows the results of this test…

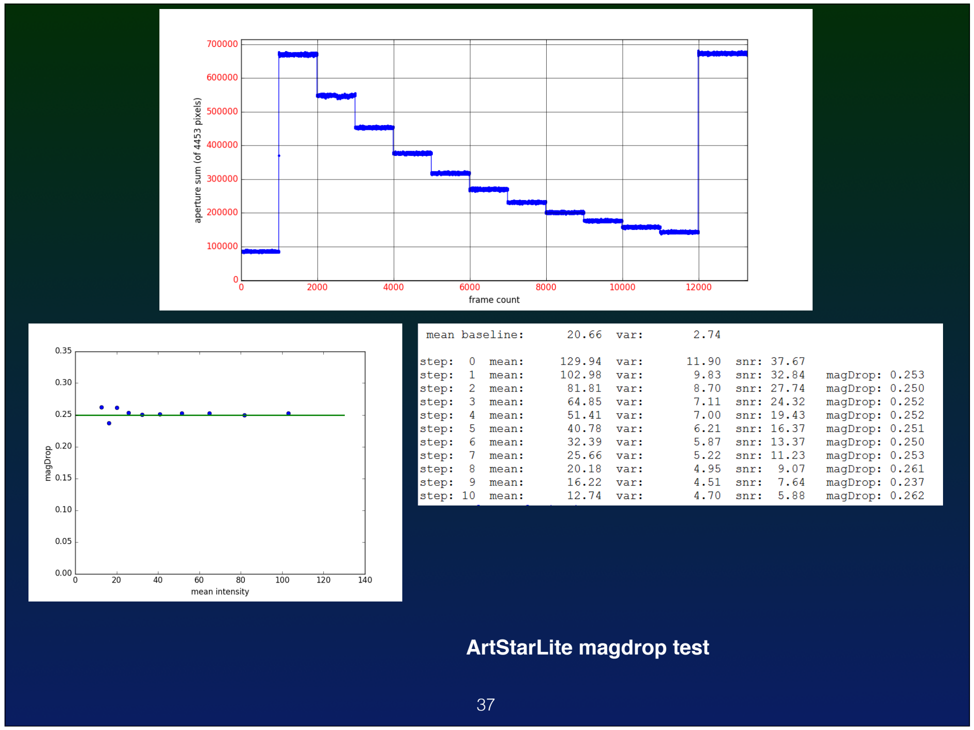

So, if we’re not doing photometry on a star image, magDrops in the range of 0 to 1.75 are very accurate. It is not until magDrops exceed 1.75 that the effect of the response curve not going zero begins to show itself.

After a brief discussion of frame-grabber automatic gain control, I’ll wrap up.

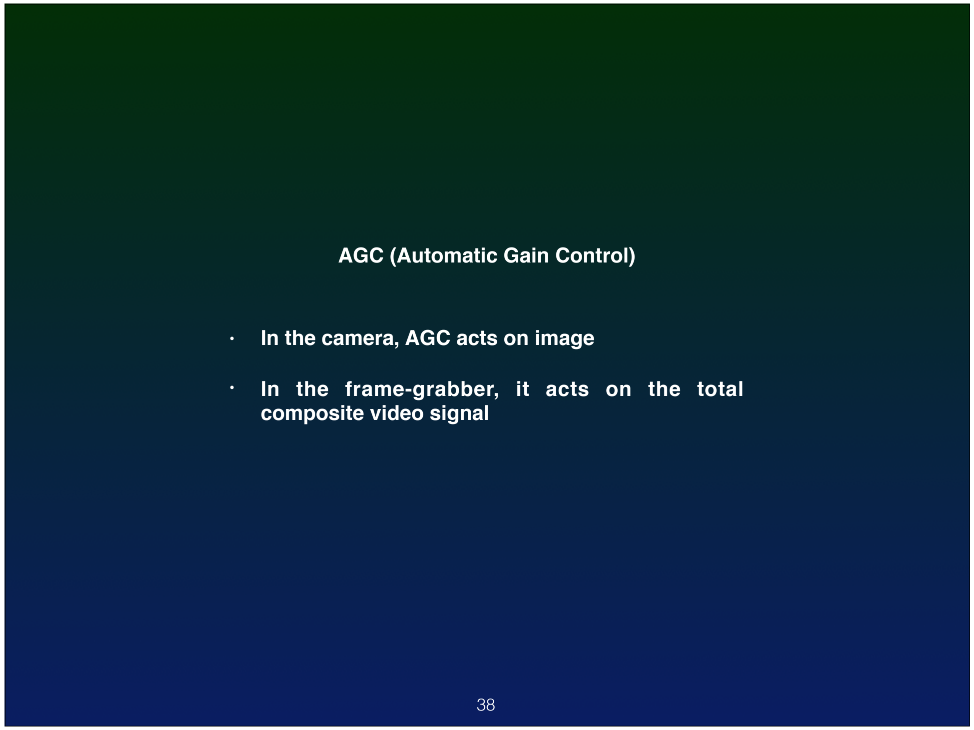

Automatic gain control (AGC) acts differently in a frame-grabber than it does in a video camera …

NEXT (read first bullet)

NEXT (read second bullet)

Frame grabbers may or may not have automatic gain control. When present, the AGC is there to compensate for non-standard amplitude composite video signals. Such out-of-spec signals can arise from the camera electronics and/or cable impedance mismatches. To understand how automatic gain control works when present, we need to look at the following diagram ..

5. Video Levels

Back to Top

The video levels define the levels and ranges for the different parts of the video signal. The unit used to define video levels is the IRE (Institute of Radio Engineers). The blanking level refers to 0 IRE and the white level refers to +100 IRE. The blanking level, which is the reference level for the video signal (usually 0 V), is different from the black level if a setup is applied to the signal as shown in Figure 6.

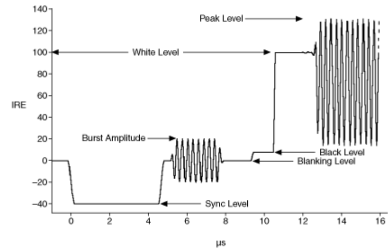

Figure 6. Video Levels

For NTSC, a setup of 7.5 IRE is usually applied, moving the black level to +7.5 IRE. For PAL and SECAM, the black level is aligned with the blanking level at 0 IRE.

Table 3 shows the different video levels depending on the video format.

Video Format	Sync Level	Blanking Level	Black Level	White Level	Peak Level	Burst Amplitude
NTSC	–40 IRE	0 IRE	+7.5 IRE	+100 IRE	+120 IRE	20.0 IRE
PAL	–43 IRE	0 IRE	0 IRE	+100 IRE	+133 IRE	21.5 IRE
SECAM	–43 IRE	0 IRE	0 IRE	+100 IRE	+130 IRE	N/A

Table 2. Video Levels by Format

The analog composite video signal is defined as a voltage source with an output impedance of 75 Ω. The sync-to-white level is normally 1 V_{pk-pk} when loaded with a 75 Ω resistance. Therefore, the unloaded signal is nominally 2 V_{pk-pk}.

Composite video diagram (for color TV signal)

39

The lefthand part of the waveform, which starts at 0 and drops to -40, is standardized and should be a fixed number of volts. When a frame-grabber implements AGC, this is the part of the composite video signal that it examines in order to select the proper gain. The gain needed to adjust this part of the signal back to the correct voltage swing is applied to the total signal.

The video data is on the righthand portion of the waveform — for our cameras, the video data appears in the region between Black Level and White Level. This diagram does show signals above white level, but that is only present in a color TV composite video.

The next slide shows the response of two frame grabbers to a composite video signal that has been deliberately modified in peak to peak amplitude by passing it through an attenuator.

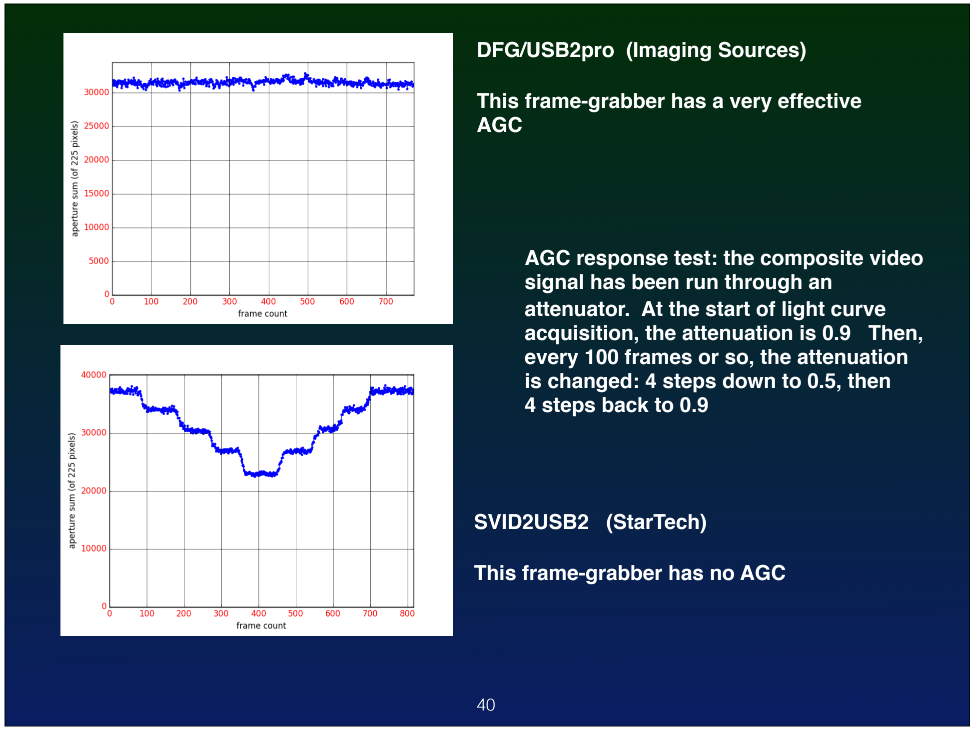

DESCRIBE THE TEST (composite video tun through attenuator — knob controlled attenuation from 0.9 to 0.5)

It can be seen that the DFG frame grabber has a very effective AGC system. There are little ‘spikes’ where I was twisting the attenuator knob, but they settle very quickly.

The StarTech frame grabber has no AGC, so if the composite video amplitude is reduced, this appears as a reduced light curve from the camera.

Does this matter? — well, it was the source of the clipping that we saw in one of the very early slides and effectively reduced the dynamic range of the camera and hid a camera response characteristic of some importance. But if the composite video signal had not been out-of-spec (too big), this wouldn’t have happened. I still use the StarTech grabber, but I make sure that the video amplitude is at the correct level.

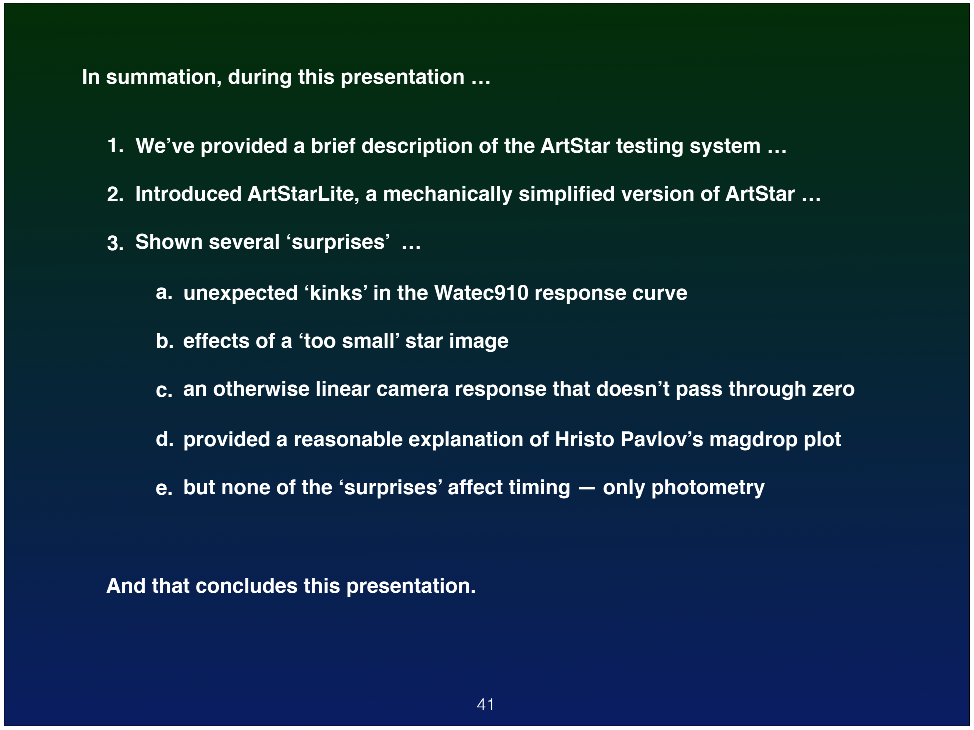

READ THE BULLETS — THEY APPEAR ONE AT A TIME

Thank you.