On the Bench:
Atik 314L+

Craig Stark, PhD

A few years ago, I did a review here of the Atik 16IC (December, 2007). At the time, I called it an "exceptionally clean, solid camera" having been impressed with how well I was able to pull out faint signals from the noise. The electronics seemed quite clean, but the sensor used (Sony ICX-424) wasn't terribly exciting, being a bit low in efficiency at things like the critical H-alpha line and having only 0.3 megapixels. It's a nice camera and I stand by my review completely, but after the review, thoughts about it or other Atik cameras drifted from my mind. I'd heard of some other Atik cameras being released with things like updates to USB2 busses for faster downloads, but to be honest, I didn't think about them too much. Why I didn't think about them is an interesting question in and of itself. Here, on this side of the pond, Atik (http://www.atik-cameras.com) doesn't get a lot of attention it seems. We don't see lots of shots off of their cameras on US-based Internet forums. This might lead you to believe that the cameras are, in some way, inferior – that they're just not up to snuff, so folks who know don't buy them.

Having now tested a current offering, the Atik 314L+, I can unequivocally refute this. This camera has what it takes and has put up some of the best test results I've seen. If you're don't need a really big sensor (and let's face it, a huge number of targets don't need huge sensors), the Atik 314L+ should definitely be on your list of cameras to seriously consider. The only cameras in this class that are less expensive test significantly worse than the 314L+ and the Atik tests at least as well as others that cost a heck of a lot more. Have I got your attention yet?

The Basics
Contents
The camera comes in an attractive box packed with all the accessories you need to get up and going. In addition to the main camera body, you'll find a 1.25" nosepiece that attaches to the camera's built-in T-threads, a USB2 cable, a 12V "cigarette plug" power cable, and a CD-ROM. The CD contains the camera drivers, a PDF of the user manual, a quick-start guide, plug-ins for *Maxim* and *AstroArt*, and the included *Artemis Capture* image acquisition software.

Physical
The camera itself is fairly compact and housed in reddish-copper anodized aluminum. The main portion of the camera is a disk, about 4.25" in diameter and about 1.25"

tall. On the front side of the camera is a smaller black disk (2.25 x 0.75") housing the CCD chamber itself. On the rear of the camera is a small fan, a USB2 port, a ST4 autoguide output port, and a 12V DC power jack. This jack uses a standard coaxial plug and needs a mere 0.8A of current. It tips the scales at about 14 oz (400 g).

Sensor
Inside the CCD chamber, you'll find the Sony ICX-285AL CCD mounted on a Thermo-Electric Cooler (TEC). The CCD sits about 12 mm behind an anti-reflective coated window, specified to be 2 mm thick. The 285 is a well-known and proven chip with 1392 x 1024 pixels measuring 6.45µ on a side. The total image area is 8.78 x 6.61 mm.

Some may balk and call this a "small" chip that isn't large enough for their needs. Certainly, it's not large enough to cover a number of targets, but in this current age of ultra-mega-pixel-lust, it's worth perhaps considering just what does fit on a sensor this size. A lot of very good work has been done on cameras like the Starlight Xpress H9 which uses this same chip. Yes, it's a bit smaller than the KAI-2020M used in cameras like the SBIG ST-2000 series and the QSI 520 series, but a lot does fit. Shown here are some sample FOV pictures I drew a number of years ago for this sensor at various focal lengths.
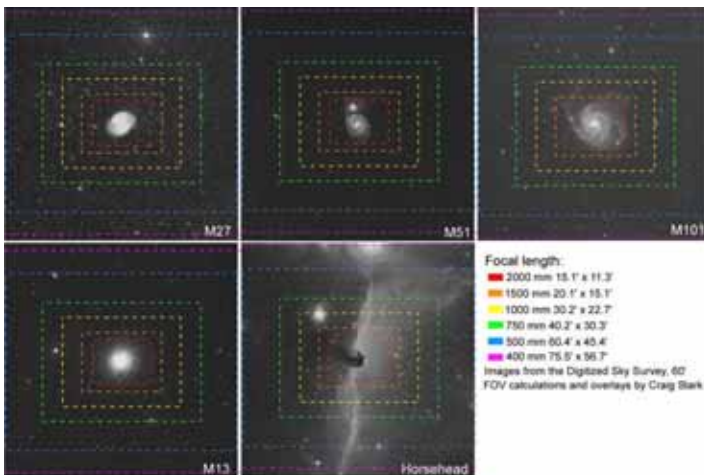
With the 6.45µ pixel size, a resolution of 1"/pixel corresponds to about 1300 mm of focal length. Thus, for typical seeing conditions, running at 1000 mm will be a reasonable focal length for high-resolution work (1.3"/pixel). I personally consider 1.5"/pixel all I ever hope to get out of my skies and 2"/pixel all I seem to be able to actually muster the vast majority of the time (from this I mean that the final image when sampled at this resolution and scaled to something like 1"/pixel looks no worse than if I sampled it natively at 1"/pixel or better). So, if we look at the green or yellow boxes in these sample FOVs, we see that we're covering these targets quite well. Sure, we're

not getting all of M31 or M45 in there and even M42 would want a two-panel montage until you're down below 500 mm of focal length, but there is a lot in the sky that will fit on this chip in a single frame.

The ICX-285 has a lot going for it as well. In particular, it's got a very nice quantum efficiency (QE) curve. While Sony does not publish QE curves in absolute numbers (they show relative QE curves only), two camera makers have curves listed on their websites for this sensor. Starlight Xpress puts it peaking at 65% and falling off to 55% at 450 nm and 650 nm (the H-alpha line) that continues well into the infra-red.   Apogee shows a slightly less peaked response but that is overall very similar.   This compares quite favorably not only to other Sony chips but also to the the Kodak KAI-2020 or KAI-4022 chips used in many popular cameras. My own reference camera, a QSI 540 uses the KAI-4022 chip which peaks at around 55% and drops to around 30% at the H-alpha line. Photons that aren't converted to electrons don't make it into our image, so having a very solid QE curve like the ICX-285's goes a long way to helping get that image.

Software
The included software, *Artemis Capture* is designed to let you control all of the features of the camera and to - as its name suggests - give you a platform to capture images from the camera. As such, it works well enough for basic use. You can capture individual frames or a continuous stream of frames. You can control binning and sub-frames in addition to exposure duration. You can, of course, also control the TEC and other features of the camera. While functional, it is a minimalist application and there is no provision for image calibration, stacking, etc. For this, Atik is working on another program, *Dawn*. As of this writing, it is still listed as being in an "alpha" stage of development, so it is not reviewed here. Most users will want to use the camera in other software that supports it such as *Maxim DL*, *AstroArt*, or *Nebulosity*.

On the Bench
As regular readers here will know, I place a good amount of weight on how the camera performs on the test bench. It's here that we get to see just how well the camera does its job of converting photons into an image and how much of its own "character" it adds into the image. When I did reviews for audio magazines with my father, we talked about how close an amplifier was to a "straight wire with gain". Well, the same can be asked of cameras here.

System Gain
The first thing I like to get a handle on is the camera's system gain or transfer function. This is the number of electrons in every ADU (Analog Digital Unit - aka intensity value) in the image. I used nine pairs of flats taken with my EL flat panel at various intensities to assess the mean and variance of the difference for each pair. A simple

Transfer Function

$$y = 0.267 * x + 51.7$$
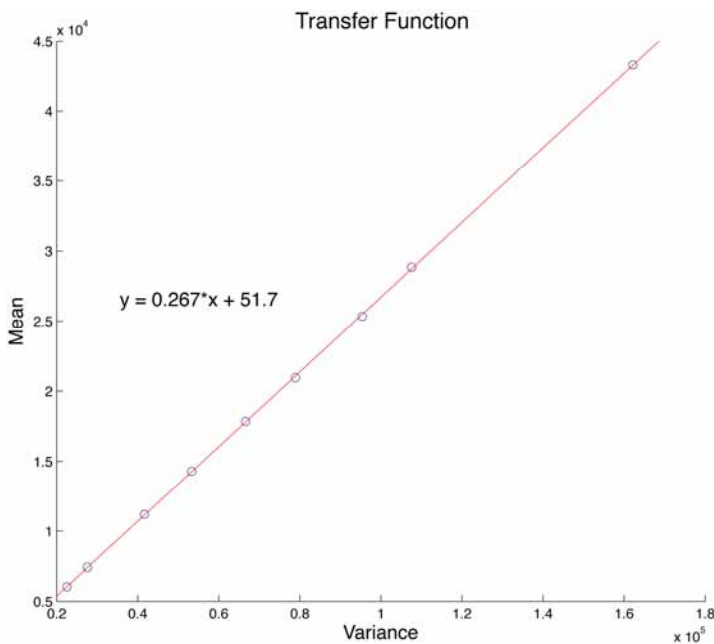
linear regression on the results showed a gain of 0.267 e-/ADU. With this number in hand, we can convert all measurements into electrons rather than arbitrary ADU. We can also determine that the maximum number of electrons that can be faithfully recorded is about 17.5k.

Dark Current and Stability
Next up is the average dark current and the stability of the dark current. Here, I took 1 m, 5 m, and 10 m darks with the TEC regulation set to 0º C and calculated the median value in each image. I plotted these along with the average bias signal and the four points were nicely co-linear. Thus, dark current was linearly related to exposure duration (as it should be). The linear regression showed an average of 0.111 ADU per minute, which corresponds to 0.0005 e-/sec. This is astoundingly low and might as well be zero (yes, it went up 1 ADU in 10 minutes). This is equivalent to one electron of dark current every 34 minutes. The shot noise contribution from this is therefore below one electron for any exposure you're likely to do.

Dark Stability - Atik 314L+
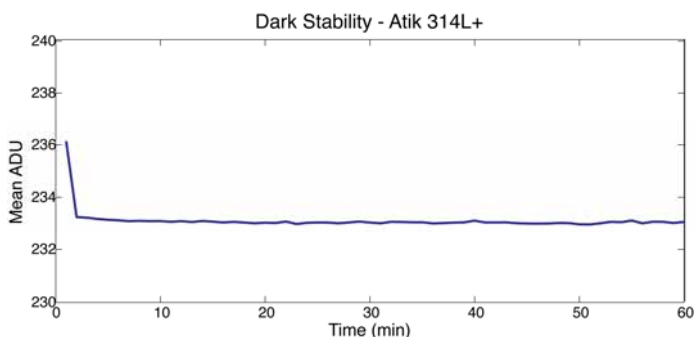
The temperature regulation worked very well also. At an ambient temperature of 77º F (25º C), I was able to get the TEC to read as low as -4.8º C for almost a 30º C

drop. Atik's specifications for this are a 25º C drop, meaning the camera was within spec. To give the camera some room for regulation, I setup for a total of 60 1-minute dark frames at 0º C and calculated the mean in each frame. Within a minute or two, the camera had gotten the TEC down from ambient (75º F) to the set-point of 0º C and it showed no variation from there (the first file read 5.27º C at the end of the exposure and all others read 0.04º C). As the figure here shows, the dark current was perfectly stable throughout this time.

In short, you can set the TEC's regulation to something like 0º C, take a library of darks, and use them across many sessions without worrying about matching dark current. The dark current is so low, however, that you may-well be better off simply taking a large stack of bias frames once to use for bias correction and use a simple bad pixel map for hot pixel removal. If you go for this option, I should note that while cooling is needed, regulation of the TEC isn't actually needed as you're not trying to match dark current.

Read Noise
I calculated the read noise both using *Nebulosity* for image capture and using the included *Artemis Capture* to ensure that nothing done in my handling of the camera would be causing any problems on this critical test. In both programs, stacks of 250 bias frames were created and stacked. The standard deviation of the difference between each individual bias frame and the bias stack was calculated and averaged across samples. Both programs turned in read noise values of 13.8 ADU, corresponding to an exceptionally low 3.7 e-.

Some cameras I've tested have turned in read noise specs on par with this, but not many. None of the ones that have been this low have had "clean" read noise. While the value itself is low, there has been a structure to the read noise, taking the form of a fixed-frequency (but variable location) pattern in the image.

People sometimes assume all patterns in a bias frame are bad. If it repeats the same way from frame to frame and is small enough to not make for variable shot noise across the image, a fixed pattern to the image is nothing to worry about. On the Atik 314L+, if one examines a large stack of bias frames, you can just pick out such a faint pattern in the image. Over the course of the first 400 pixels in each row, there is a slight modulation of the signal, almost like a ringing effect that goes for four cycles before fully dampening out. This is a very slight modulation. The first one is about 6 ADU (1.6 e-), the second, about 3 ADU, etc. It's so small that even with such a low read noise, it's pretty much impossible to see visually in a single bias frame (looking at the average row you can see it). It's also perfectly repeatable which means that it's easily removed with either darks or bi-

ases. *Subtracting a single bias from this bias stack showed a perfectly clean response*. This is the kind of fixed-pattern noise you want to have if you're going to have it. It's very small and it's entirely repeatable. All this means is that if you're going stretch the image like crazy (and who doesn't?), you'll want to use either darks or biases to remove this pattern. Do that, and don't worry about it ever again.

What is problematic are patterns that are left over *after such a subtraction.* These arise when there are patterns to the image that change location from frame to frame. There may be bands that appear in each image equally spaced from each other, but in a different position in each frame. This kind of fixed-frequency noise is problematic as it cannot be removed by calibration with darks or biases. The best way to see such patterns are with Fast Fourier Transforms (FFTs) on the read noise frame (the difference between one bias frame and the stack of bias frames).

An ideal two-dimensional FFT plot will show a bright dot in the center (representing the "offset" of the image) and nothing else but simple noise. There should be no other bright spots or lines in the FFT. If there are, these indicate that there is significant energy at certain spatial frequencies in the image. Here, I'm showing the FFT of both a single bias frame (left) and of the read noise frame (right) using the same scales I have used in other reviews.

What we can see here if you look closely is a slight vertical line in the single-frame FFT at the center. This stems from that slight fixed pattern I described above. Mind you, if this were a read noise frame FFT, this would be quite good in and of itself as there is very little going on here. Many cameras don't have read noise frames as good as this single bias frame. What we care about, though, is the FFT of the read noise frame. The FFT of the read noise frame is, for all intents and purposes, perfect. It shows a bright dot in the middle surrounded by nothing but noise. You just can't get any better than this.

One final way we can look at the read noise frame is to examine the histogram of the read noise. In particular, we can look at the histogram when plotted on a logarithmic scale (many programs do this automatically when looking at histograms, but some do not). What we want is for the noise to be purely random and not biased towards some values more than it should be. For example, if there is "salt and pepper noise", very bright and very dark pixels come up a lot more often than they should. This will show up as in the histogram as "tails" on either side that shouldn't be there. Here, I'm showing the histogram of the read noise for the Atik 314L+ superimposed with the histogram of a synthetic image made entirely of random, Gaussian noise of the same standard deviation. In short, the Atik's read noise frame is wonderfully clean. Not only is the read noise very low overall (3.7 e-), but it is very clean. The histogram for its read noise frame is indistinguishable from that of a purely synthetic noise

Single Frame Total Noise: 2D FFT - Atik 314L+

Read Noise: 2D FFT - Atik 314L+

Pure noise
Atik 314L+

Log(count)

Intensity difference (ADU)
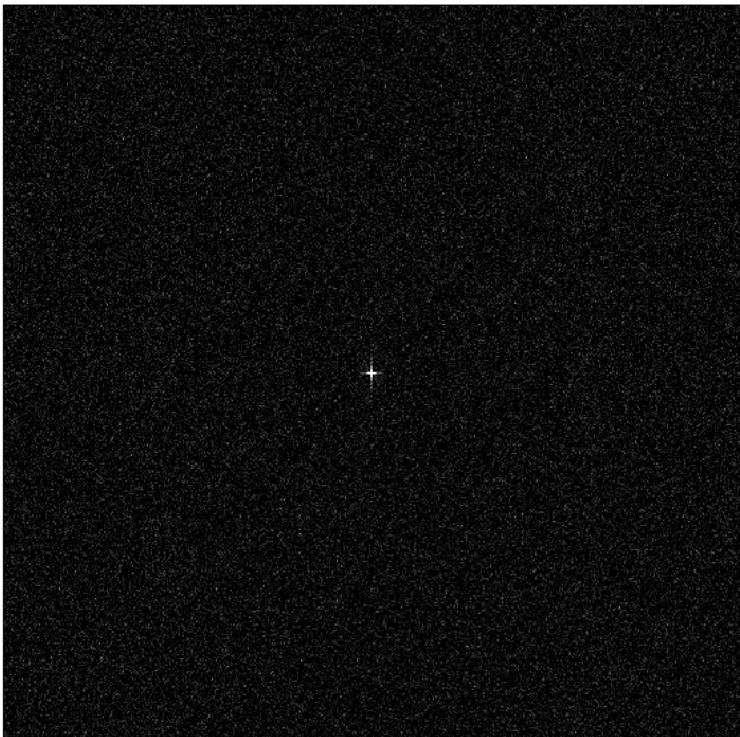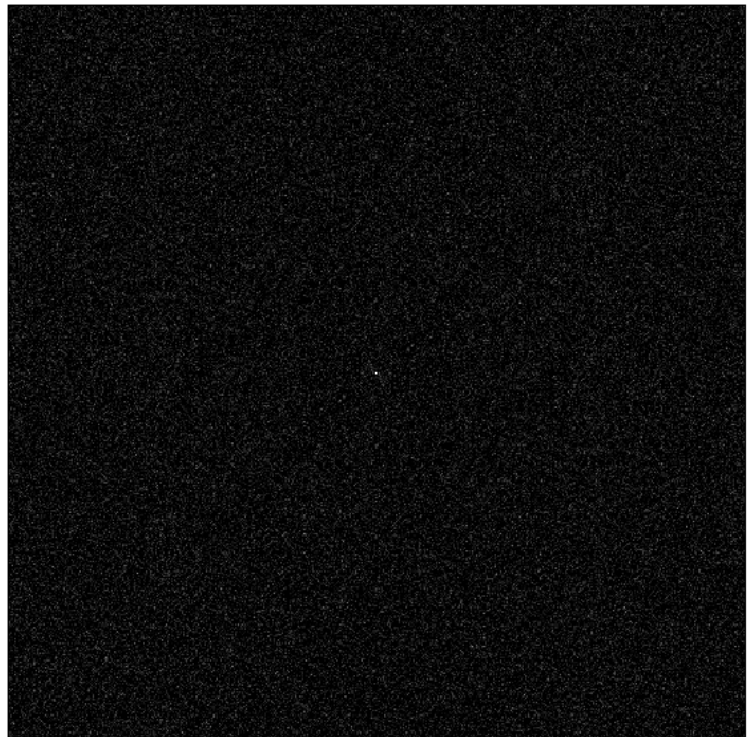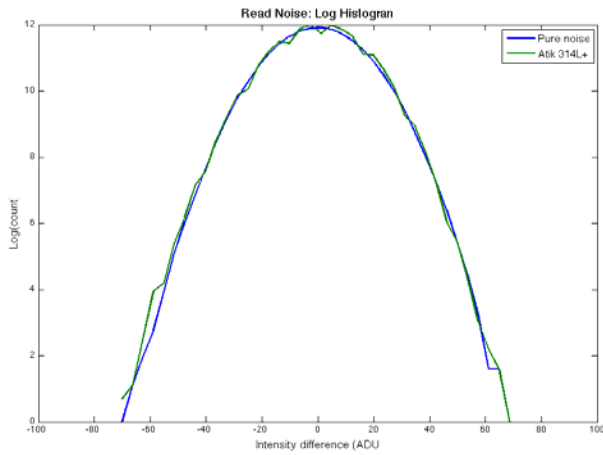
frame. This is doing very well as a camera stand-in for that "straight wire with gain" I looked for in audio amplifiers.

Under the Stars
Of course, the final test for any camera is to see how well it performs under the stars. My urban skies don't let me get beautiful LRGB images without incredibly long stacks.

While the camera was here, getting out to dark skies was out of the question and besides, we should be judging how well the camera performs, not the processing skills of the reviewer (good thing for Atik!). But, we do need to know whether the camera faithfully records stars or whether such harsh point-sources yield artifacts. We'd also like to know how well you can pull faint bits out of the noise.

Fortunately, even from my urban yard, shooting in H-alpha lets us get at both. I used a Custom Scientific H-alpha filter on the Atik 314L+ to block out my skyglow as I imaged M57 with *Nebulosity* on a Celestron C8 SCT equipped with a Celestron 0.63x reducer / flattener. With my external focuser (Starlight Instruments Feathertouch), this rig runs at about f/7 for a sampling rate of just under 1" / pixel. All gear rode atop my Takahashi EM-10 and Telescope Stability Systems tripod. Guiding was done with a Starlight Xpress Lodestar, a mini-Borg Guide Scope (with Hutech X-Y stage), and *PHD Guiding*. A total of 26 5-minute images were collected, and processed and stacked in *Nebulosity.* Two stretches of the
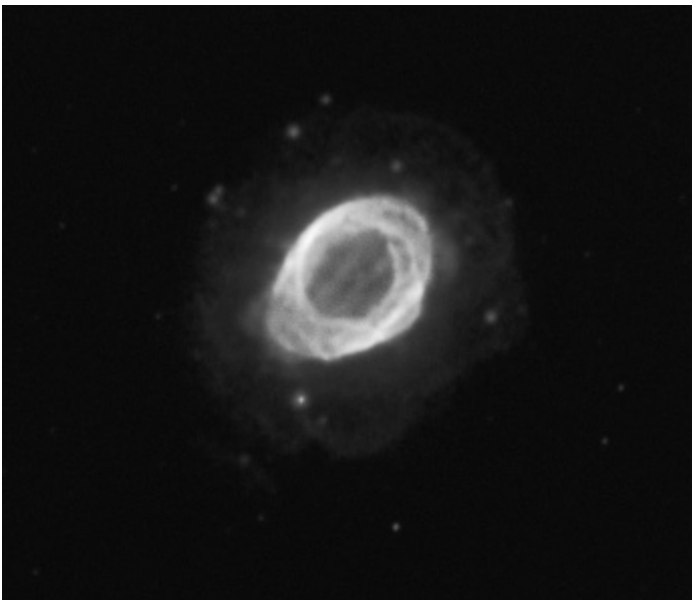
images were made in *Nebulosity* (one standard and one for the halo) and further stretched and layered in *Photoshop*.

With the H-alpha filter in place, shot noise from skyglow was minimal. While my bias frames read about 230 ADU (61 e-) on average, the background in these H-alpha images after 5 minutes was only 340 ADU (91 e-). Unlike imaging in LRGB, with only 30 e- worth of glow here, shot noise from the skyglow won't dominate the image (an L frame here would be about 100x as bright and have 10x the shot noise). These frames therefore do a good job at letting the camera's signature come through.

What I found was entirely consistent with the excellent bench tests. Stars were nice and clean with no sign of any artifacts and the background noise was minimal and clean. It was so clean that even in this fairly short stack of H-alpha data (just over 2 hours), I was able to pull out the faint halo that surrounds M57. This halo is very faint and is not easy to record (I know, I've tried many times). While it wasn't easy to resolve here, it was clearly there and, with careful stretching, it came out in the image.

This is what having a really clean camera can do for you. You can go deeper into the image and pull out faint bits that would elude you under other circumstances. Under dark skies (or with solid skyglow filtering and the right targets), you have a chance at those really faint bits that would otherwise be obscured by skyglow's shot noise. But, to get them, you'll be looking for every bit of signal to noise (SNR) you can muster and you'll still probably be working your noise reduction skills considerably during processing. With a clean camera like the Atik 314L+ here, you'll pull in more than with a camera whose noise characteristics are routinely getting in the way (rather than going away).

Conclusions

At the outset, I said that this camera has what it takes and has put up some fantastic test results. If a sensor of this size is in your thinking (and have a look again at that

FOV figure if you think it's not), I strongly suggest that you look at the Atik 314L+. This is a serious camera that doesn't cost serious money.

Is there anything "wrong" with it? What's not to like? Well, there are a few things that I think should go into someone's purchasing decision here since this is still not pocket change for most of us. Mind you, when I'm drawing comparisons here, it's to cameras that cost over $1000 more than the Atik 314L+. So, keep that in mind.

First, while the camera's form-factor will work well for use in something like a Hyperstar setup from Starizona, it could be a bit better. The Starlight Xpress H9 is only about 2.5" in diameter and won't even add to the central obstruction of an 8" SCT if mounted up there. Here, you'll have a bit of extra obstruction even on a C11. Second, a number of cameras in this chip-class (e.g., Apogee Ascent 285, SBIG ST-2000, and QSI 520) come with options for integrated filter wheels. Atik does have a motorized filter wheel that can be attached, but this will extend the camera's size and isn't as nicely integrated as some of the other configurations. Third ... I'm sorry, honestly I can't even think of a third.

Regular readers here may remember the last camera I was mightily impressed with - the QSI 520. I like the QSI so much, I went out and bought its bigger brother, the QSI 540. At this point, readers might be wondering just how they compare and whether I'd switch over to the Atik. Well, no, I'm not giving up my QSI 540 here. It's chip is a lot bigger (21 mm diagonal vs. 11 mm diagonal) and I'm a sucker for wide-field work. I also like the integrated filter wheel and off-axis guider options, and it runs on my Mac. These would be real things to give up and real things that lead to the 2.5x difference in price, making this an apples-to-oranges comparison. I've got no complaints with the QSI at all as it's a fantastic camera that I'm not looking to replace. It's still my go-to camera. But, if I had the Atik and QSI here on the shelf and was looking at some H-alpha imaging of a target that would be framed well on the Atik's chip, the added H-alpha QE there on the Atik, coupled with its also-excellent noise characteristics could easily have me grab it instead. Yes, this is a great camera. ♦

By day, Craig Stark, Ph.D. is a professor whose research involves trying to pull faint signals out of noisy, moving images of people's brains. By night, he is an amateur astrophotographer and operates Stark Labs. Stark Labs provides software to help users pull faint signals out of noisy, moving images of the heavens. Craig also writes the <u>Fishing for Photons</u> *column for Cloudy Nights.*
<u>http://www.stark-labs.com</u>