

Measuring Skyglow with Digital Cameras

These instruments are designed to take pretty pictures, but you can also do science with them. *By Tony Flanders*

AMATEUR ASTRONOMERS have lots of opinions about light pollution but very little hard data to back those opinions up. Traditionally, we assess conditions by noting the faintest star that's visible to the unaided eye. Unfortunately, the results vary tremendously depending on the observer's eyesight, skill, dark adaptation, and level of exhaustion. Also, the visibility of any given star is ultimately a subjective judgment, creating the hazard of unconscious bias — especially with something as emotionally charged as light pollution.

Until Unihedron introduced the Sky Quality Meter (see the box on page 104), no standard commercial light meter was sensitive enough to measure the brightness of the night sky. Skyglow *can* be measured with instruments that vary the brightness of an illuminated surface until it matches the sky (*S&T*: February 2001, page 138), but these have proved difficult to standardize and calibrate, and they still rely on a subjective judgment.

A few years ago I started to wonder if a point-and-shoot digital camera could do the job. I had been thinking of “going digital” anyway, for terrestrial snapshots, and this seemed like an ideal excuse to take the plunge. I bought a Canon PowerShot A80 and discovered that it can measure urban and suburban skyglow quite accurately despite the fact that its exposures are limited to 15 seconds. More recently, I've

been using the Canon Digital Rebel XT, a digital SLR that is sensitive enough to quantify fine variations of brightness in a nearly pristine sky.

It's easy to measure skyglow with any digital camera that can take reasonably long exposures under full manual control. All you have to do is snap a photo with a standard set of exposure parameters, download it to a computer, and measure the pixel levels in the area of interest with an image-processing program such as *Adobe Photoshop*. I've done numerous experiments indicating that these correlate well with the brightness of the subject and are affected very little by extraneous factors such as temperature, battery charge, and random chance.

The color information is potentially interesting, but to avoid being swamped by data, I decided early on to measure all my skyglow shots in grayscale only. When using my point-and-shoot camera under reasonably dark skies, I use a few photo-processing tricks to get more precise results, as outlined ahead. But even the crudest methods work fine for typical suburban skies.

What's It Good For?

As an astronomy writer, I find objective, quantitative skyglow measurements immensely helpful. For instance, if I



Amateur astronomers have attempted to measure light pollution for decades, with widely differing results. *Sky & Telescope* associate editor Tony Flanders has developed a simple technique using standard digital cameras to measure sky brightness to an accuracy of $1/4$ magnitude. This shot — one of 21 used to create a full-sky brightness map — shows the Sagittarius Milky Way merging into the glow of New York City 80 miles south of the observing site.




Any digital camera that can take multisecond exposures under full manual control can measure suburban skyglow. The angle scales on this pan/tilt head make it easy to take the sequence of shots needed for full-sky coverage, but the author frequently meters the zenith just by setting the camera on the ground with its lens pointing up.

■ astrophotography

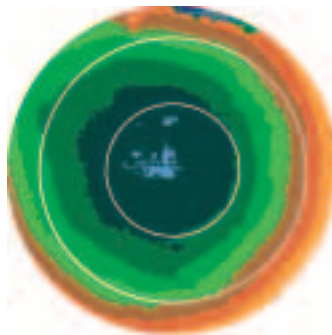
Magnitude
per square
arcsecond




Arlington, MA
Aug 25, 2004, 4 a.m.




Westford, MA
May 13, 2005, 1 a.m.



Springfield, VT
Aug 7, 2005, 1 a.m.



Kancamagus Pass, NH
Aug 4, 2005, 1 a.m.



Flanders shot all-sky panoramas at four locations in New England and used custom software to compile them into color-coded brightness maps. Everything within 15° of the horizon has been cropped, to minimize the effects of trees, houses, and lights. The white circles correspond to 30° and 60° above the horizon. The panoramas from Vermont (shot at the Stellafane telescope makers' convention) and New Hampshire show the Milky Way running prominently overhead. To make it easy to correlate the light domes with cities shown on terrestrial maps, the pictures are shown with north up and east to the right — the reverse of what you would see if you were looking up at the sky.

want to describe how an object will appear in a bright city sky when it's overhead, I can observe it four months earlier, when it's low in the sky, from a good suburban site. I know that these conditions are comparable because I've measured the skyglow in both cases.

I take most of my measurements on nights of good transparency, when there is no obvious haze or cirrus cover. Near cities, poor transparency shows up as a dramatic in-

crease in nighttime sky brightness. Under rural conditions, where much of the skyglow comes from sources above the atmosphere, the reverse happens and poor transparency tends to make the sky darker.

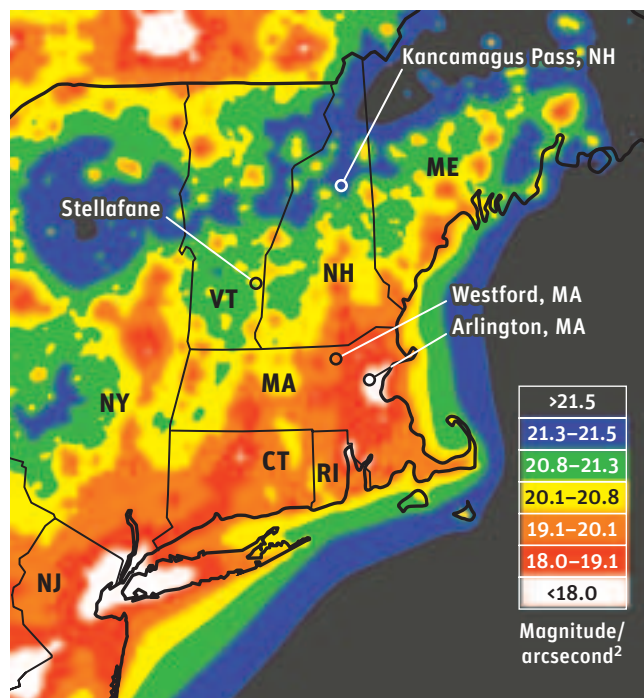
The variation in artificial skyglow by season and time of night is surprisingly large. Using a device similar to the Sky Quality Meter, Dan McKenna (University of Arizona) has found a difference of 1/2 magnitude between early evening

and 3 a.m. in Tucson, Arizona. I've obtained similar results around Boston, Massachusetts. In the eastern US, which is blanketed by deciduous forests, artificial skyglow increases about $\frac{1}{4}$ magnitude when the leaves fall from the trees, and fresh snow can amplify that effect tremendously.

Using custom software to stitch images together while correcting for the camera's distortion and vignetting, I've made numerous full-sky brightness maps, some of which are shown on the facing page. Each requires 21 separate shots — eight at altitude 15° , eight at 45° , four at 75° , and one at the zenith. The results are easiest to visualize when I color-code each pixel to represent the brightness of the sky at the corresponding altitude and azimuth.

The full-sky maps confirm what we already know: the best single measure of light pollution is the brightness at the zenith. This doesn't characterize a sky completely, but it's about as good as you can do with a single number. In a suburban setting, the very darkest spot in the sky is skewed slightly away from the city center, but it's typically within 10° of the zenith and no more than 0.1 magnitude darker. But 15° above the horizon, skyglow can vary by more than a full magnitude from one direction to another.

The Milky Way is faintly visible on some of my suburban panoramas, and it's a major "pollutant" at any rural site. In the Northern Hemisphere, this is particularly true in late summer, when the Cygnus star cloud (roughly magnitude 21 per square arcsecond) passes overhead. At such times, the darkest spot may be quite far from the zenith.



LIGHT POLLUTION SCIENCE AND TECHNOLOGY INSTITUTE

The locations where Flanders shot the panoramas are shown on an excerpt from the World Atlas of Night Sky Brightness, available online at www.lightpollution.it/dmsp/index.html. The calibration of colors in magnitude per square arcsecond is from the Clear Sky Clock Web site (www.cleardarksky.com).




S&E: TONY FLANDERS

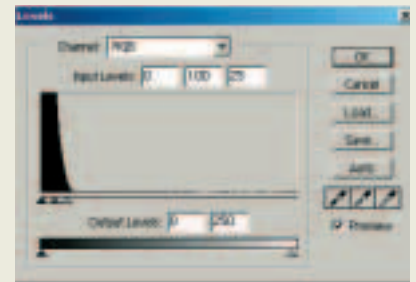
A 15-second shot with a Canon A80 barely shows the skyglow at a reasonably dark site in April 2004. Note Venus above and left of the Pleiades and Mars above the Hyades.

The Response Curve

The sensors inside digital cameras have inherently linear response up to about 75 percent of saturation, and high-end cameras allow you to read these values out directly in RAW-format images. But when photos are converted to JPEG format —the only option for most point-and-shoot cameras — the internal 10- or 12-bit values are mapped onto



Using the Levels function in *Adobe Photoshop*, Flanders stretched the histogram by a factor of 10 (with a gamma of 1) to bring out the skyglow strongly — together with lots of random noise.





Converting to grayscale and resampling the image to 80×60 pixels, Flanders obtained a reading of 109 at the point of interest. So the average pixel level was 10.9 before the histogram was stretched.

the range from 0 to 255 in a nonlinear fashion. That doesn't matter if your only concern is to rank observing sites, saying that one is darker than another, but it *is* a problem if you want to compare them quantitatively.

Fortunately, it's quite easy to reverse-engineer the camera's level-mapping algorithm by photographing a single subject at different shutter speeds and observing how the pixel lev-

els respond. Numerous tests indicate that shutter speeds on both my cameras are quite accurate. The f/stops are less reliable, so I always shoot with the lens wide open. A light-polluted sky makes an excellent test subject, as does a sheet of white paper illuminated by indirect daylight.

JPEGs made with both cameras exhibit nearly linear response for pixel levels up to 50. Knowing the response curves, it's quite easy to correct the pixel levels to what they would be if the linear response continued above 50. But it's still best, whenever possible, to plan the exposures to yield pixel levels between 10 and 100. This avoids some erratic behavior at very low pixel levels and strong nonlinearity at high levels.

Calibration

Once I determined my cameras' response curves, I could make quantitative comparisons between observing sites, and between different parts of the sky at a single site. But I also wanted to publish my readings in absolute units, like magnitude per square arcsecond.

Professional photometrists calibrate their instruments by photographing stars of known brightness and adding up the levels of the pixels in the star's image. I had serious doubts that this method would work with a digital camera. My initial tests with the A80 indicated that two successive measurements of the same star frequently differed by as much as 25 percent. Still, calibrating by star brightness seemed worth a try.

One of my skyglow shots with the A80 captured Cassiopeia when it was near the zenith in a reasonably dark sky. As usual, I was shooting at 15 seconds and f/2.8 at ISO 200. (This yields much lower noise levels than ISO 400.) To minimize disk space and processing time, I save my skyglow photos at the lowest possible resolution (640 × 480), and by photographing a ruler at a known distance, I determined that this yielded about 160 pixels per square degree in the center of the frame.

I computed the total pixel levels for Cassiopeia's 20 brightest stars and extrapolated them as if each star had been boosted to magnitude zero. (For instance, I would multiply the reading for a 5th-magnitude star by 100.) That yielded values between 7,000 and 11,500, with an average of 8,700. Extinction near the zenith at sea level is roughly 15 percent, so the average magnitude-zero star would have a reading around 10,200 if the atmosphere were not present. If that were spread out over the 160 pixels in a square degree, each pixel would have a level of 64. So a pixel level of 64 corresponds to magnitude 0 per square degree, and a pixel level of 1 corresponds to magnitude 4.5 per square degree — or 22.3 per square arcsecond.

This analysis rests on many untestable assumptions. But conveniently enough, McKenna had lent me a prototype light meter that could measure skyglow directly, and its readings indicated that a pixel level of 1 corresponds to magnitude 22.4 per square arcsecond — the value that I've used in my analyses ever since. The near agreement of the two results is comforting, but it's probably an accident that they're as close as they are. I don't really trust my calibration to better than 1/4 magnitude.

If ISO ratings were completely reliable, one could deduce the sensitivity of *any* digital camera from this single data point. For instance, the deepest exposure I can take easily with my Rebel XT and its stock lens is 30 seconds and f/3.5 at ISO 1600. Some simple arithmetic indicates that the sensitivity ought to be 10.24 times that of the A80 at 15 seconds, f/2.8, and ISO 200. But experiments show that the sensitivity is only about 6.5 times higher, so that a pixel level of 1

corresponds to magnitude 24.4 per square arcsecond.

The discrepancy between theory and practice isn't surprising. Several people have reported that the A80 is quite a bit more sensitive than its ISO ratings indicate, while the XT's ratings err in the opposite direction.

Future Directions

If you own a digital camera that can take multisecond exposures under full manual control — as many amateur astronomers do — it's easy to set the camera on the ground, pointing at the zenith, and take a skyglow measurement at the end of an observing session. Even if you don't interpret the results right away, the raw data will be available as long as you keep the photo.

Because it's so hard to calibrate different models of cameras accurately, a device like the Sky Quality Meter probably provides a more uniform way to measure light pollution. But it's hard to compete with a digital camera's true strength: the ability to measure millions of spots in the sky simultaneously. That's what makes it possible to produce my full-sky brightness maps.

My current technique for making these maps requires far more photographs than is theoretically necessary. A fisheye lens would make it possible to capture the entire sky in a single photograph while simultaneously minimizing vignetting. Shooting the sky's reflection in a polished metal sphere would accomplish the same thing at a lower cost, but it would also require some complex software to map the photometry of the photograph back onto the celestial sphere.

Another unusual feature of digital cameras is their ability to meter simultaneously through three different color filters. Starlight, airglow, scattered moonlight, and various kinds of artificial lights have different spectral signatures. It should be possible to tease out a huge amount of information about these by comparing the readings of the red, green, and blue pixels. Anybody interested in the challenge? *

Sky & Telescope *associate editor* TONY FLANDERS *does much of his observing under moderate to severe light pollution.*

The Sky Quality Meter

Until recently, no standard commercial light meter was sensitive enough to measure the brightness of a pristine sky. But now Doug Welch, professor of astrophysics at McMaster University, and Canadian amateur astronomer and tinkerer Anthony Tekatch have teamed up to build such a device. It's called the Sky Quality Meter, and it can be ordered online at <http://unihedron.com> for \$119.95.

The Sky Quality Meter complements a digital camera perfectly. Digital cameras have extremely fine

angular resolution, but they're hard to calibrate, and their readings can be obtained only after the images have been downloaded to a computer. The new light meter is factory calibrated, and it provides nearly instant readouts. Extensive experiments indicate that the results are highly repeatable, and they vary by less than 1/10 magnitude among four independently purchased units. They also agree with the results from my digital cameras within 2/10 magnitude.

The unit's greatest flaw is that it

“sees” essentially the whole sky at once. It's most sensitive within a cone 80° across, but light as far as 60° or 70° off axis can affect the readings. That means that it's completely reliable only when it's aimed near the zenith and when your site has no major obstructions or bright lights in any direction. The Milky Way alters the readings when it's high in the sky, but less than one might think. That's because the field of view is so huge that even the Milky Way covers only a small part of it.

— T. F.



S&T: CRAIG MICHAEL UTTER