RECENT POSTS

Open Mike: Readings

Young People, Now People

And Here's What I Think the Pentax Film Project Ought to Be

What Will the Pentax Film Project Camera Look Like?

We Are Lucky

Stellar Seasonal Bargain: Sony A7III

Open Mike: Wasaaaaaaabi! (Wasabi)

A Fictional Fable About a Spottie and a Hottie

David Turnley Christmas Print Offer

The Best Battery Jumper: Update to the Update, and I Promise This is the Last of This

+ SPONSORS

RESOURCES

About Us

The Old Site

Advertising information

Comment Guidelines

Publisher's Statement

Contact

« Random Snaps: Clement Valla's Melted Bridges | Main | The Megapixel Race is Over »

Wednesday, 16 March 2011

Photography at the Speed of Light

By **Ctein**

I was catching up on some of my photonics journals last week and they reminded me that I still haven't answered for myself the question of just how fast digital cameras can get. What is the ultimate limiting ISO?

You know there has to be one. Simple logic. I mean, when you reach the point where the shutter is open just long enough to let one photon get through, you can be pretty confident that you're not going to be able to form an image.

Unfortunately that simple logic doesn't really tell you very much. There are a zillion comments online to the effect that the "laws of physics" limit how fast the camera can get, but it's pretty clear to me that the people saying that don't have any idea what that limit is. Statements can be both true and irrelevant. Much like pointing out that the laws of physics limit how fast a rocket can travel. No argument there, but rocket designers aren't bumping their heads on that speed-of-light brick wall just yet.

Where is ISO's "speed of light" limit? I don't know, but I can tell you of some improvements that push sensor speeds way beyond anything you've imagined. All are either proven technologies or such simple extensions of existing technology that we can be entirely confident they can be implemented. That doesn't mean they will be, but they do constitute a proof of existence.

The current speed champ, the Nikon D3S, provides very decent quality up to ISO 6400, and is usable up to 102K. Not that that would ever be my choice, but if it's the only way I can photograph The Second Coming or First Contact, it'll do. How do we make it even faster with no loss in image quality?

First, install a back-side-illuminated sensor (unless I misread the specs, the D3s isn't using one). BSI sensors are about twice as efficient as conventional ones at gathering light. Large pixels don't benefit as much as small ones. For the D3S I'd be conservative and call it a 2/3 stop gain (for micron sized pixels, it's distinctly more than one stop).

BSI sensors haven't made big inroads yet, although they're in a few cameras. I suspect a combination of cost and the fact that smaller-sensor cameras need them more desperately.

Next, Kodak has this nifty modified Bayer filter array that gains you another stop of efficiency. That's three-year-old technology; not certain why we haven't seen it in a camera yet. Maybe I'm not looking in the right places, maybe Kodak's being a pain about licensing, dunno. Whatever, it's not a physical issue, so far as I know.

The Kodak array isn't the theoretically best one. It's probably possible to squeeze another half to full stop of efficiency out before you've reached the limit. I'll be conservative and say half a stop more.

The new hot area of development is sensitizers for silicon. Silicon detectors, at their very best, run about 50% efficiency. In addition, they are much more sensitive to long wavelengths of light than short ones. That's particularly bad because most often you're using high ISOs under indoor lighting that is generous in the long wavelengths and weak in the short. That's why the red channel in your high ISO indoor photos looks pretty clean, the green channel looks mediocre, and the blue channel looks horrible.

Anyway, an assortment of companies are developing layers and coatings that can be applied to the surface of the sensor that catch nearly all the photons and transfer their energy to the silicon. Some of these technologies are very close to rollout, if they're not already in production for the minuscule cameras that go into cell phones.

Combine BSI sensors, improved filters, and sensitizers and there's a 10X performance improvement. Let me emphasize that this is real improvement, not funny games being played with more amplification and noise reduction. What the Nikon D3S can do at ISO 6400, these technologies would let it deliver at 64,000. And when the flying saucers (or Jesus) descend onto your front lawn, you can crank that dial up to "11" for an unbelievable extended speed of 1,000,000.

Honest to God, I have no idea what you do with a speed of 1,000,000. Make handheld photographs by starlight? Just sayin'—it can be done, and we haven't exhausted the bag of technical tricks so don't assume that even this represents that performance brick wall.

Realistically, these technologies will appear in small-sensor cameras sooner, because they have more need for them. I don't think too many people will be unhappy when digicams offer the image quality of today's full-frame cameras. Trust me, they are going to get there.

PORTALS

Support Sustain Contribute via Patreon

SEARCH

Google

WWW
T.O.P.
Google Search

CATEGORIES

Around the Web (220)
Baker's Dozen (50)
Blog Notes (564)
Bon Mots (10)
Book of the Week (15)
Books (521)
Camera Reviews (138)
Cameras, new (778)
Cameras, old (255)
Carl Weese (23)
Collecting (53)
Computers (91)
Ctein (370)
Digital Techniques (72)
Eamon Hickey (2)
Editing and Portfolios (75)
Exhibits (84)
Film and Darkroom (127)
Followups (147)
From the Archives (38)
Geoff Wittig (2)
Gordon Lewis (15)
Interviews (7)
Jim Hughes (31)
John Camp (8)
John Kennerdell (17)
Ken Tanaka (11)
Kirk Tuck (8)
Legal and social issues (112)
Lenses (316)
Mini-Portfolios (5)
Morning Coffee (32)
Music Notes (106)
News and Occasions (473)
Obituaries (81)
Off-topic posts (489)
Open Mike (337)
Oren Grad (1)
Past Blasts (23)
Peter Turnley (66)
Photo equipment (343)
Photo industry (169)
Photo-tech (221)
Photographers, current (240)
Photographers, historical (163)
Photographic aesthetics (219)
Photography businesses (26)
Photojournalism (85)
Polls (16)
Print Offers (261)
Printers and Printing (169)
Quote o' the Day (95)
Random Excellence (246)
Random Snaps (123)
Reader Prints (7)
Recommended Cameras (14)

Meanwhile, the skies still the limit, folks, and we still don't now how high is up.

Ctein

TOP's weekly column by Ctein (pronounced "kuh-TINE," and yes, that's his entire legal name) appears on Wednesdays.

Posted on Wednesday, 16 March 2011 at 10:17 AM in Cameras, new, Ctein, Photo-tech | Permalink

Unique contents copyright © 2011 by Ctein and the commenters

Comments

Where does temperature fit into this? I thought that CMOS was better because it ran cooler and Live View was bad because it heats up the chip. That implies, to me, that we are bumping up against a kT limit.

Is the temperature of the chip the fundamental physical limit? If so, will we see active (pelltier) coolers on the chips to lower the noise floor?

Posted by: KeithB | Wednesday, 16 March 2011 at 10:42 AM

I'm reminded of Stephen Johnson (http://www.sjphoto.com/) back in the day, chimping under a dark cloth on a Powerbook G3, being astonished at just how much further through the distant haze his scanning, 8x10 camera back could see.

Also, images by starlight aren't that unusual these days. I've no idea what the "ISO speed" of the Photonis XD-4 tube in my binoculars would be, but it is definitely very sensitive to photons and can make images by starlight (which are monochromatic, low-res, and noisy). Attaching it to my K-5 seems like a step backwards in every area *except* sensitivity!

Posted by: MarkB | Wednesday, 16 March 2011 at 10:46 AM

Stop teasing us. :)

Posted by: Robert Roaldi | Wednesday, 16 March 2011 at 10:53 AM

The limit is like "how far in space can we go?" If there is a commercial reason to develop the low light much further than it is, technology will be developed to meet the need. If it isn't commercially viable, it won't.

I'm always amused when people talk about the laws of physics limiting this or that. The laws won't change but technology figures workarounds.

Posted by: Mike Spivey | Wednesday, 16 March 2011 at 11:07 AM

It seems to me that something doesn't quite match up...

According to the data from the sensorgen.info site, green channel quantum efficiency for most current sensors is in the 30-50% range. If I understand correctly, that's computed based on the DxO measurements and includes all the losses from the AA filter, color filter, etc. (I think they're using SNR to get the well capacity at saturation, and the saturation ISO to get the number of photons needed to fill the well.) That says that there's at most a factor of 2x or 3x available before you're capturing every available photon. So where's the other 3x-5x?

Is the sensorgen or DxO data wrong, or (more likely) is my understanding of what their quantum efficiency figure represents wrong?

Posted by: david | Wednesday, 16 March 2011 at 11:29 AM

I can think of a use for ISO of a million.
Last winter I pointed my now ancient D70 at the night sky and did a thirty second exposure at ISO 1600. As expected it was noisy but M42 was a lovely red arc across part of the frame.
It leaves me wondering how a D3s with some older high speed long glass, say a 400 2.8, would do on some of my favorite night sky objects?

Posted by: Mike Plews | Wednesday, 16 March 2011 at 11:37 AM

It seems to me that the technological limits for light gathering/resolving will ultimately be somewhere far beyond what anyone could care about, and that cameras will reach the "who cares" limit fairly soon. Compare the cameras of 5 years ago with the current generation and extrapolate out another 5-10 years and it's not hard to imagine ISO equivalents in the 128,000+ range that are perfectly usable.

I predict ISO will go the same way as the megapixel race - at some point well before technological optimization people will decide it's good enough and stop being willing to pay a premium for more.

Now optics, on the other hand, is where there are legitimate physical limits - though liquid lenses are intriguing.

Posted by: Evan | Wednesday, 16 March 2011 at 12:17 PM

I thought back in the D80 times when I could shoot digital at ISO 3200 better than film it was great... and since then I've shot a lot of indoor, low existing light photographs. And looked at a ton of others.... And overall, I don't think any of them are all that good.

Reflections (198)
Reruns (19)
Roger Overall (2)
Sarge (1)
Satire Alert! (25)
Shooting Skills and Camera Handling (6)
Shooting techniques (157)
Software (75)
Sunday Support Group (73)
TOP contests (6)
TOP's Own Books (10)
Visual Culture (270)
Websites and links (141)
Weekend Exercises (5)
Workshops (19)
Your Turn (9)

ARCHIVES

December 2023
November 2023
October 2023
September 2023
August 2023
July 2023
June 2023
May 2023
April 2023
March 2023
More...

When I get down to it, ISO 800 seems to be the threshold of decent photography. The high ISO stuff is like most HDR images: technological stunts without substance.

Posted by: Frank P. | Wednesday, 16 March 2011 at 12:52 PM

you go up to mulholland drive at night and take a picture of los angeles, only you can see the stars, too. no barndoors or photoshop nonsense, maybe graduated neutral density.

Posted by: aizan | Wednesday, 16 March 2011 at 01:28 PM

Good One !
Thank goodness you weren't lecturing us about "punctuation" !
(last week really was the only boring thing I've ever read by you)

Posted by: Kenneth Voigt | Wednesday, 16 March 2011 at 01:33 PM

Ctein- give up the bayer array and use 3 back-illuminated chips for each color gets you to 90%+ quantum efficiency. Work on the amplifiers to get read noise down, or use the electron multiplying methods so the read out noise becomes non-important. All these methods are currently in use in scientific cameras, so you should be able to put them into use if you want to spend the money. I'd guess that you could implement 3 back-illuminated electron mulitplying chips into a camera for maybe 100K$ street price in todays dollars. But 4 years from now it would cost $375, right in the P&S range.

Posted by: Phil Allen | Wednesday, 16 March 2011 at 01:50 PM

1. As I hear and suppose, some kind of organic material based sensors would do a lot to increase both resolution, sensitivity and most importantly: dynamic range. I can imagine a new sensor with stochastically arranged (color)sensitive material, which would resemble to film. It would't have to use a Bayer filter or an AA filter or whatever. Or we could put 2 or 3 of them on top of one another: one for the highlights, one for the shadows and one for the middle tonal range.

2. Non of the digital cameras of today with a single sensor have the latitude of color negative film. You can't get back the highlights of a digital image (+2 or 3 stops misfired) that's burned into a RAW file. At least not with 12-14 bit A/D converters that are standard today. And if you could, or expose carefully: it still won't look like conventional film. Even those 250.000+ dollar Hollywood HD über-motion-picture-cameras cannot do that, you've got blown-out highlights and find yourself spending a whole lot of post production money to make it watchable and still got crappy blueish-black shadows and red-and-green face-colors. Blah. :-)

Posted by: Gabor Metzker | Wednesday, 16 March 2011 at 02:03 PM

What also happens with high ISO's in low existing light, is that the camera sees color where our eyes don't. This might be a factor in the unnaturalness ('technological stunts without substance') that Frank P. mentions. Of course special fields like astrophotography (M 42, the Orion Nebula) are another matter. But apart from that, black & white photography is the natural thing to do in low light. (Nothing against 'unnatural' experiments with color in the dark, though.)

Posted by: Hans Muus | Wednesday, 16 March 2011 at 02:38 PM

What I find interesting about ISO is that despite all the differences between film and digital, the micro lenses, etc. - digital cameras still operate at a base ISO of roughly where we were with film (for best quality). I suspect it is as much practical as it is the laws of physics. If cameras started using a base ISO of 800, who could shoot outdoors with wide apertures?

Posted by: Brad Calkins | Wednesday, 16 March 2011 at 03:00 PM

Frank P, I think it depends on one's tastes. I am willing to put up with a pretty "thin" exposure if the composition is right.

That's my own photograph, but the published work of Roy Decarava contains many comparable images.

Posted by: Ben Rosengart | Wednesday, 16 March 2011 at 03:04 PM

Super-high ISOs would have plenty of practical application right now. Suppose that with ordinary, mainstream cameras and lenses, you could routinely shoot in typical indoor light with a shutter speed of 1/10,000? Joe Schmo would never have to worry about blurry pictures again, and expensive image stabilization systems would become unnecessary.

Of course, that 1/10,000 speed would require another technological step up -- such as the "global shutter", another technology that forum posters write about as if it were totally realizable RIGHT NOW, and which yet seems to be taking a long time to appear (kind of like antigravity neckstrap lugs...)

Posted by: Ranger 9 | Wednesday, 16 March 2011 at 03:28 PM

Dear Keith,

kT is a physical limitation, it's not a physical limit. In other words, it's purely an engineering issue. As you note, there are ways to cool chips. There are also some really cool tricks out there for distinguishing between kT electrons and photo-excited electrons. Maybe in some future column.

~~~~~~

Dear Mark,

Yes, image amplifying tubes do not produce what we would call “quality” photographs, but they do demonstrate that we already know how to turn the knob on the amp up to 11. The angle I tackled in this column was what got you more speed without degrading quality.

~~~~~~

Dear Robert,

Nuh uh. In fact, here are links to my previous columns under future tech designed to tease and frustrate you:

"Keeping the 'X' in 'Xmas'"

"Shapes of Some Things Long to Come"

I am so mean...

~~~~~~

Dear Mike,

Not exactly. How far in space we can go is essentially an engineering problem. Heck, you can traverse the entire visible universe in a lifetime if you can figure out how to wrangle 10 billion tons of antimatter and engineer sufficient shielding and streamlining. It's not a physical limit, like going faster than the speed of light.

There are physical limits to how fast a camera can get. I just haven't been able to calculate it, yet.

pax \ Ctein
[Please excuse any word-salad. MacSpeech in training!]
==========
-- Ctein's Online Gallery http://ctein.com
-- Digital Restorations http://photo-repair.com
==========

Posted by: ctein | Wednesday, 16 March 2011 at 03:38 PM

Dear David,

I think the confusion lies in what the efficiency is being measured relative to. For example, I can immediately account for a factor of two by noting that they are looking at front-side-illuminated sensors. Most of the gain in BSI sensors is purely geometric–– less shadowing and shading of the sensor that prevents photons from getting to it. That reduces your discrepancy to 1.5 X-2.5 X. I also talked about more efficient filter array designs being able to squeeze out another half stop, maybe more.

I don't think it's worth digging deeper. The thing to remember is that unless someone is talking about genuine ab initio calculations, they aren't talking about fundamental limits but engineering efficiencies. Which are always relative to some particular, not necessarily theoretically ideal, case.

An ab initio calculation of ultimate speed is possible. If I could relate ISO to number of photons per second per square centimeter at the sensor plane, I could tell you, within a factor of two or three, the ultimate physical speed limit for a camera as a function of image quality. Unfortunately, I'm not quite up to doing the photometry. I don't have confidence that I haven't missed a factor of pi somewhere or something like that.

If there's anyone reading this is actually comfortable doing photometric calculations, I'd love their help. If not, well, if I'm close to being able to calculate this it's certain there are plenty of people out there who already have. Problem is, I haven't seen their papers, so I remain “in the dark.”

~~~~~~

Dear Evan,

The legitimate physical limits of optics are not quite as cast in stone as we thought they were. Look into meta-materials and computational imaging. There are some very interesting end-runs possible there.

~~~~~~

Dear Frank,

Depends on what you're doing and what you need. I can get by with ISO 800. I'd really prefer to be able to go to ISO 3200 with high image quality. I have genuine needs, not technological stunts. Good and high ISOs also open up photographic possibilities we didn't imagine before. Easy for me to say I can't imagine why anyone would need ISO 60,000 for “serious” work, but I wonder how an original Kodachrome photographer would feel about us arguing over ISO 800?

Art will expand to fill the technological space allotted to it, of this I am confident. Probably won't be the art you or I do, but someone will.
~~~~~~

pax \ Ctein
[Please excuse any word-salad. MacSpeech in training!]
==========
-- Ctein's Online Gallery http://ctein.com
-- Digital Restorations http://photo-repair.com
==========

Posted by: ctein | Wednesday, 16 March 2011 at 04:08 PM

Dear Kenneth,

Still leaves you my Biggest Fan.

A burden of notoriety you're going to have to bear up under.

~~~~~~~~

Dear Phil,

Ummm, no.

Well, BSI, yes, but otherwise all you're doing is increasing the total area of the sensor. Haven't improved filter or silicon efficiency one bit. What makes you think otherwise?

~~~~~~~~

pax / Ctein

Posted by: ctein | Wednesday, 16 March 2011 at 04:18 PM

@Gabor Metzker:

have you seen RED's new "HDRx", and "magic motion" which captures both a standard image, and a longer exposure. This aloows them to combine both motion blur and sharp images that have about 13 stops of dynamic range.

Posted by: Nico Burns | Wednesday, 16 March 2011 at 04:30 PM

I have recently come back into amateur astronomy after a long hiatus and one of the surprising new tools for telescope users is highly sensitive *video* cameras that can take "almost real time" exposures of the very dim objects in the sky while attached to a telescope. By almost real time we are talking a minute or two. And by dim I mean things that people used to expose on film for hours.

Strangely, CCD-based astronomical imaging does not seem all that interested in these tools, as you still see people out in the dark exposing their CCDs for hours.

Anyway, my point in all of this is: maybe the ISO 1,000,000 sensor would be sensitive enough to be able to do "almost real time" images over the larger areas of camera chips instead of the small chips the video cameras are using.

Posted by: psu | Wednesday, 16 March 2011 at 05:04 PM

The more I think about it the more I wonder: is it really necessary? Just because you can do something, do you really need to? There was a time when 100ISO was a revelation, to say nothing of 400. I'm often thankful to be able to set my camera to ISO 1600 and know I'm going to get a workable image. But I also wonder if all of these enhancements really just lead to a lazy group of shooters as opposed to an innovative group making better pictures.

Posted by: mike | Wednesday, 16 March 2011 at 05:11 PM

Mike/Ctein,
I can give you a plausibe answer to this question - extrapolated upwards from my own practical experience at high ISO's using wide open primes, allowing for the 1/focal length rule, lower 'shaky hand' tolerance expected with digital, and the need to produce a reasonable DOF in many cases:
Witha a standard lens, 1/125th at F8 in a dark hotel function room, when the bride and groom are dancing. Alternately night time seascapes hand held at a minimum of 1/160, f 8.. Fred Parkers ultimate Exposure Computer gives these at around EV-2. I make that 200,000 ASA at a very rough guess...

Posted by: Ger Lawlor | Wednesday, 16 March 2011 at 05:16 PM

Ctein,

I believe I have an answer for you. I used to work in a vision laboratory in the University of Western Australia, and we had a photometer which counted light in photons. We used this to measure the intensity of light. We needed to measure contrast on our experiments so really well calibrated equipment was necessary.

Our best photometer was a PhotoResearch Spectra Pritchard (which I forgot the name of but easily remembered when I googled it).

http://www.photoresearch.com/current/pr1980b.asp

The rated sensitivity when cooled is: 6.7 x 10-5 candelas/m2... which is about 0.000067th of a candle. That's not a lot!

We also had some really cool colorimeters, I think we had this one (which is now discontinued).

http://www.photoresearch.com/current/pr650.asp

All of this equipment was extremely extremely expensive, but very fun to play on! if you perhaps ask nicely to one or two vision scientists around, I'm sure they'd be more than willing to have a discussion with you and show you their lab.

Cheers, Pak

Posted by: Pak-Ming Wan | Wednesday, 16 March 2011 at 06:04 PM

Dear Ctein,
You wrote:*Honest to God, I have no idea what you do with a speed of 1,000,000. Make handheld photographs by starlight?*

Well, I'd like to photograph a field full of fireflies so they appear as I see them. But, by my calculations, I'd only need a clean ISO of 32 768.

Showing my work:
I've seen light-trail photographs of them shot at f/8 60sec ISO400 (film). I think translates into EV -2, so getting that up to 1/125 or 1/500 would mean an additional 13 or 15 EVs sensitivity. I would assume that f/4 would give more than adequate depth of field on a APS-C or 4/3 sensor, so let's say 15 stops. So, 2^15 = 32 768. Bring on the new technology!

A little more seriously, I don't understand any objections to ISOs above 800. An awful lot of life is lived between EV6 and EV2, and some there are some beautiful landscapes to be found down in EV-2 through EV-6. You know, moonlight. It would be great to have non-tripod required shutter speeds there.

Will

Posted by: Will Frostmill | Wednesday, 16 March 2011 at 06:42 PM

Er, I think I goofed my math there. I meant, 2^15*100= ISO 3 276 800.

So no, not fast enough yet. :)

Posted by: Will Frostmill | Wednesday, 16 March 2011 at 06:46 PM

From a few simulated photon shot-noise limited images, I get something around ISO750K for monochrome, a bit grainy but comparable to Ilford Delta 3200. The limit depends on your acceptable noise criterion and subject matter but within an order of magnitude, ISO1M is about right.

Posted by: Daniel | Wednesday, 16 March 2011 at 09:13 PM

Dear Kenneth,

Oh, BTW, the final exclamation point goes INSIDE the closing quotes, not outside.

[g,d,&r]

pax / punctuated Ctein

Posted by: ctein | Wednesday, 16 March 2011 at 09:14 PM

Dear Will, et.al.,

OK, now you've got me thinking about what I would do with high ISO.

Y'know, that firefly photo is one I've wanted to make my whole life, too. Forgot it was on my "long list."

Sanity check on your ISO-- ISO 400, f/2.8 and 10 seconds records everything the human eye can see and a bit more. At ISO 40,000, that becomes 1/10th sec at f/2.8. ISO 400,000 gets it to 1/100th sec at f/2.8.

Yeah, I'd want at least that to have any chance, even with a fast lens, and ISO 1,000,000+ would be awfully welcome.

Less commonly, I'd like to make some photos of aurorae that really looked like aurorae. IOW, captures the delicacy (meaning low noise and good gradation) and the very fine filamentary structure (means high shutter speeds). Real ISO in the hundreds of K's again.

Like Ger, I do a fair amount of photography of performers in dimly-lit hotel rooms. I could get by with a really clean ISO 3200, but I don't need your depth of field. Still, I'd not turn up my nose at another factor of 10.

~~~~~~

Dear mike,

Read this:

Lazy? Does Not Compute!

pax / Ctein

Posted by: ctein | Wednesday, 16 March 2011 at 09:28 PM

What are the artistic photographic possibilities in scenes in which the light that is too dim to see by with the naked eye? Other than astronomy and surveillance I cannot say.

But what if your camera's electronic viewfinder lets you see as well in dim light as the sensor can? Do we walk around taking pictures on a moonless night, of people who cannot see us?

But then again, what about a world where everybody's vision is thus enhanced, making night into day without the use of artificial illumination?

You could go on...

Posted by: Mani Sitaraman | Wednesday, 16 March 2011 at 10:04 PM

Would it be sensible to return to the Deutsches Institut für Normung numbers? ISO 6400 would be DIN 39, 12,800 would be DIN 42, etc. Wouldn't that be handier than the rather clumsy proliferation of digits we're getting into when talking about ultra-high speeds?

Mike

Posted by: Mike Johnston | Wednesday, 16 March 2011 at 10:15 PM

Dear Mike,

Ya got my vote, but I'll bet the CaNikon marketing folk won't be happy talking about their next generation of high-speed cams going from DIN 51 to DIN 54 instead of ISO 102,400 to ISO 204,800.

Maybe they can just write "ISO 400,000" in magic marker on the
~~~~~~

ISO control.

pax / Ctein

Posted by: ctein | Wednesday, 16 March 2011 at 10:51 PM

Dear Mike,
I like shoving the last two 00 over, or otherwise deemphasizing them. E.g.
ISO 1 00, 2 00, 4 00, 8 00, 16 00, etc. I would love to see people drop the last two digits as a matter of course:
1,2,4,8,16,32,64,128,256,512,1024
Considering that there's not that much precision with specified ISO vs measured, that seems pretty reasonable. I'm thinking of the nice Dx0 graphs showing how much high ISOs recede when measured, compared to the names manufacturers give them. (Actually, it would be kind of nice to set it to read ...15,30,60,125,250,500,1000. You know, consistency.)

Dear Ctein,
I am delighted to have made such a connection with you. Thank you for the sanity check on my math. That is quite a straightforward rule of thumb, I think I will find it quite useful.

Will

Posted by: Will Frostmill | Wednesday, 16 March 2011 at 11:17 PM

"I'd guess that you could implement 3 back-illuminated electron mulitplying chips into a camera for maybe 100K$ street price in todays dollars. But 4 years from now it would cost $375, right in the P&S range."

The back-illuminated EM-CCD camerain my lab (90% QE @ 525 nm, <0.1 e- read noise per pixel, 14 bits DR) was purchased more than 5 years ago.

Similar cameras coming out today are a bit less expensive but no less bulky due to the requirement for a peltier cooling device that holds the sensor and read circuitry at -80°C. Of course, we image at ambient temperature, so it's necessary to house the whole assembly in a vacuum enclosure to prevent condensation.

I'm not expecting to see this tech in a usable DSLR any time soon.

Posted by: Semilog | Thursday, 17 March 2011 at 01:46 AM

Honest to God, I have no idea what you do with a speed of 1,000,000.

Available light portraits of black cats in coal cellars at night.

Posted by: Steve Smith | Thursday, 17 March 2011 at 03:06 AM

@ Ctein,

this is going to be educational for me. I'm not a scientist!

"Heck, you can traverse the entire visible universe in a lifetime if you can figure out how to wrangle 10 billion tons of antimatter and engineer sufficient shielding and streamlining. It's not a physical limit, like going faster than the speed of light.

I'm intrigued. Let's take it that your streamlined anti-matter fuelled rocket ship is a viable vehicle, and all other variables such as the start point and definition of "lifetime" are agreed. If the universe is constantly expanding, light moves at around 300,000 km/sec (and you are not exceeding that speed), how do you ever reach the end of "visible"? The observable universe has a radius of 47 billion light years centred on you or me, but surely as your rocket ship approaches the boundary of the visible universe as observed from home, looking forward you are going to be seeing more and more universe that you cannot see from earth? Isn't it an endless journey?

Oh, and how do you get back? ;)

Posted by: James | Thursday, 17 March 2011 at 04:45 AM

I would attack the problem from the other direction. Assuming a perfect sensor (each and every photon that enters is detected and counted), how many photons per sample are needed for an “acceptable” image? This becomes a statistics problem of determining the minimum sample size of a binary event needed to detect rates different by the smallest step that will give us the desired image quality (dynamic range, resolution and noise). This would give us the theoretical highest possible ISO. Real life sensor efficiency, filters, lenses etc. would all lower this number.

All else being equal, larger pixels will yield higher theoretical maximum ISOs. Since all else is never equal the limiting factors will be the size of the camera and the cost of the lenses needed to feed the sensor.

Posted by: Speed | Thursday, 17 March 2011 at 06:12 AM

I have a couple of optical devices which seem pretty good at coping with starry skies at night, even with a moderately brightly lit town foreground, and many other very low light levels. They've been pretty reliable for the last 70 years or so and hopefully will continue in like manner. Furthermore, they are coupled to an inboard computer with a workflow which has only changed slightly over the years and is excellent at storing those images which I feel are really important and, generally, recalling them at will. The only problem I ever had with them was the need for prescription spectacles when I was quite young...

Posted by: Len Salem | Thursday, 17 March 2011 at 07:29 AM

Thought I might just throw in here that the best book of night photography I have was first published in 1933. I'm just saying.

With Ctein's ASA 1,000,000 and Strobist's "thyristors everywhere," we ought to have the darkness beat....

Mike

Posted by: Mike Johnston | Thursday, 17 March 2011 at 08:09 AM

The sky IS the limit!
Say good-bye to tracking hardware and software for astrophotography.
Imagine a beautiful photo of a nebula at...what... less than 1 second?
Count me in!
JB

Posted by: Joseph Blankier | Thursday, 17 March 2011 at 08:54 AM

Statement of belief: Nothing that the human eye can see is inherently and absolutely out of bounds for artistic use. (There are many things *I* have no idea how to use artistically, but there are many artists in the world.) (In fact, I'll go further; lots of things the human eye can't see directly are still of artistic interest; I've seen photos made using scientific equipment for artistic purposes, from the Hubble Space Telescope down to scanning electron microscopes and x-ray machines.)

But of course, beyond that, there are lots of uses of photography that are not basically artistic (even if they have artistic elements). The shots made at a wedding reception, for example, while judged partly on artistic criteria, exist primarily to document the occasion. The same is true for the shots made in Tiananmen Square. And beyond that, there are uses in surveillance and scientific photography.

I do get the impression sometimes that there are people out there with a depressingly narrow view of "photography".

Posted by: David Dyer-Bennet | Thursday, 17 March 2011 at 09:41 AM

The link to the Kodak Bayer filter array
makes fascinating reading. Can you imagine
a Panasonic Lumix FZ100 (25 to 600 equiv.)
with some future sensor that matches today's
full frame digital? The sky is truly the limit.

Posted by: daugav369pils | Thursday, 17 March 2011 at 09:51 AM

Low light photography is more a matter of graphical presentation of scientific data than making a photograph look as we would see it. Someone notes above that at low light levels human eyes have almost no color sensitivity; consequently, astronomical photography is essentially false color, often also presenting infrared, ultraviolet, or radio waves. Electron microscopy, not even using light, is also a matter of graphical, false-color, presentation.

Insofar as we accept this, the number of photons in the human eye's frequency range is not the last word on photographic usefulness. It's sometimes useful not to forget ultraviolet, infrared, Terahertz photons, etc. Furthermore, the skies the limit if we allow outrageous capture and computational strategies for getting an image. No more than a ridiculous example, but suppose we illuminate a scene with broadband infrared (it's still dark, right?), infer chemical/biological surface properties from infrared radiation/reflectivity, then false color the image so that it looks as it would under a sodium lamp (we do a pale imitation of this already, it's called "white balance", adjusting for the properties of light sources).

We could also consider, amongst other properties, the depth of the field of focus of the human eye (surely it's not equivalent to f16 in a full frame camera?). Do I want to see the world as it is with or without my glasses on? Even, how does the smell change my visual perception?

Graphical presentation of scientific data at its best is surely worthy of the term "art", though not so much when at its worst. From this point of view, the Bayer array or the Bayer 2.0 array might be too close to the human eye's functionality. One wants to optimize the data we gather to the range of artistic choices we expect to make, not to the human eye. Alternatively, we can change the data we gather so that we can imagine a wider range of artistic choices. Automating artistic choice becomes much harder, of course, so this is perhaps not for the faint of heart.

I'm not sure this can be said to be a peaceful approach to the world, however.

Posted by: Peter Morgan | Thursday, 17 March 2011 at 11:25 AM

Dear James,

Sorry for the minor confusion; in astronomy, “visible universe” refers to the universe we can see from here. In other words, a volume of space about 27,000,000,000 ly across. It is not presumed that is the full extent of this universe; it's just a part of it we can see.

It's not really physically meaningful for me to talk about traversing the entire visible universe, because the stuff at the very limits is already receding from us at the speed of light. Doesn't matter how fast we go; when we get there, it will be somewhere else. In other words, my remark should be only read to mean that given sufficiently clever (for some extremely large value of “clever”) technology but no new Physics, we could traverse many billions of light years in a human lifetime.

~~~~~~

Dear Speed,

That is EXACTLY the kind of ab initio calculation I've been trying to do. Problem is I don't know how to compute the starting point; the surface radiance of a standard object. If I could do that, I could convert that energy to numbers of photons and get from there to a flux reaching the sensor, and after that it's all very simple statistics. Trouble is that I'm not very good at the photometric conversions. In principle, it's very simple algebra to go from irradiance, the sunlight coming in, to reflected energy. In practice, if you don't get the geometric factors of 2pi and all that other stuff straight, you can be off by an order of magnitude without even knowing it. I don't trust myself at all.

The answer wouldn't be terribly precise; as you know it really depends on the kinds of information you're collecting from each photon in the first place, and there are all sorts of wonderful classical and quantum mechanical tricks for trading off information you don't need in exchange for information you do. But at least we'd have a number in the ballpark. At the moment, I have no idea if the ballpark is in the range of a real ISO of 100,000, 1 million, or 10 million. Or maybe even lower, and what we've been discussing is the ISO equivalent of “false magnification.”

pax \ Ctein
[Please excuse any word-salad. MacSpeech in training!]
==========
-- Ctein's Online Gallery http://ctein.com
-- Digital Restorations http://photo-repair.com
==========

Posted by: ctein | Thursday, 17 March 2011 at 01:14 PM

I swear to you, I have never read this phrase anywhere before: "I was catching up on some of my photonics journals last week ... "

Posted by: Bill Rogers | Thursday, 17 March 2011 at 02:20 PM

Great title! Doesn't photography happen at (or near) the speed of light?

Posted by: toto | Thursday, 17 March 2011 at 02:30 PM

Ctein,

It looks like the Kodak array at least made it into sensors as of March 2010.
http://pluggedin.kodak.com/pluggedin/post/?ID=2995926

BTW, the post claims the 3 photos were taken on different chips. I can clearly see differences in angle of view between the monochrome shot and the 2 colour shots, but can't see any difference between the one labeled Bayer and the one labeled Trusense.

Martin

Posted by: Martin Bennett | Thursday, 17 March 2011 at 05:22 PM

Ctein,

you and I have a common understanding of the"visible universe", although mine is informed only by Google and Wikipedia, for what that's worth. Yours is informed by being a scientist, which trumps me. There's a small difference of my reading saying it is 95 odd billion light year across, and your view of 27 billion. But here on TOP I'm sure the proprietor will overlook a matter of 68 billion light years.

I was trying to explain to my 11 year old daughter a few of these scientific concepts of couple of weeks ago. I stated that the "visible universe" was 95 billion light years across, and to think of it as a tennis ball lit up internally. That represents the universe as we humans can sense it with all of our magical telescopes and radio-measuring devices. We're in the middle of the lit-up tennis ball. The real question is: how much other universe is there outside of the tennis ball that we cannot yet sense? Is it like a tennis ball inside a shoebox, or a tennis ball inside a racquets court, or a tennis ball inside a stadium? I have no idea. If we're in the tennis ball inside a stadium, are we on the centre-spot of the pitch or somewhere up in the far upper bleachers, in the cheap seats? I tried Wiki'ing the answer and was rapidly defeated by concepts such as "2D elastic sheet" and "ant on an elastic string". Thus chastised, I'll stick to being a marketing manager and she went back to her Justin Bieber fanzine.

She did however wonder how, if there was "nothing" before the Big Bang, how come there is so much "stuff" in the universe. Where did it all come from? I have no answer, and regrettably Professor Brian Cox (possibly the most accessible scientist for we normal people, and one-time bassist for D:Ream) has not yet addressed that in his otherwise magisterial series now playing on the BBC iPlayer.
~~~~~~

Posted by: James | Thursday, 17 March 2011 at 09:04 PM

Martin - Copying the images into Photoshop, doing some alignment, and comparing does show a definite perspective difference between the pictures; It is very slight though, and much easier to see when you can stack the images together.

As far as those array types in the article go, here's some more recent information on an entire family of Kodak sensors that was announced last fall. Including one up to 29 megapixels!

It looks like they are all CCD, and there's not much mention of what their use in consumer cameras may be.

Ragarding the sensors mentioned in the article you posted, they are very low res CCD sensors, which implies cheap. And since they say they are still making the monochrome and standard bayer versions, I assume there some non-negligible extra cost associated with making the new array type (or perhaps it just desirable enough that they can charge more and gain market share).

Either way, how that translates to the higher megapixel consumer camera sensors remains to be seen. It's possible that the potential drawbacks of the new array type - lack of color accuracy, difficulty in demosaicing, etc. - make it worse than the bayer array for those kinds of camera. Just speculation, and it's equally likely that we'll see TRUESENSE cameras later this year or next.

Posted by: David Bostedo | Thursday, 17 March 2011 at 11:59 PM

Dear James,

Okay, truth is that if you're in a hyper-relativistic spaceship, I'm not sure if the static or the co-moving distance is correct. So you could be right about the 95 billion light years. As you say, a minor discrepancy.

On to your daughter's question. (As an aside, if she's advanced enough, she might enjoy my two holiday columns about dark matter and dark energy.) Basically she wants to know how you can get something from nothing. Here's half the answer. What she's talking about is what physicists call a Law of Nature, in this case the Conservation of Mass and Energy. It says that the total amount of mass and energy in the universe stays the same; you can move the stuff around, even transform it into different forms, but you can't actually add or remove any. In other words, you can't get something for nothing.

Now here's the thing about Laws of Nature. We don't actually know whether they're fundamentally true or not. They're just things that we've always observed. But there is no physical theory or model that says that matter and energy have to be conserved. In fact, sometimes we change the Laws-- There used to be separate Laws of "Conservation of Mass" and "Conservation of Energy." By our crude observations, it appeared that each of them remained constant, separately. Then Einstein figured out that mass and energy were interconvertible, and we looked closer and what do you know he was right, so we threw out the old Laws and made a new one.

There even used to be a theory for the universe that explicitly threw out this Conservation Law. It was called the Steady State theory. It propose that even though space-time was expanding, matter got created very slowly out of the vacuum (like one atom of hydrogen per cubic meter every so often, something we humans would never notice) so the new universe never got diluted, it just got bigger and bigger. Physicists and astronomers didn't have any problem with that theory in principle. Turned out to be wrong, but if it happened we just would've thrown out another Conservation law.

So, something for nothing isn't forbidden, it's just what we are used to seeing.

Here's the other half of the answer. There isn't really “nothing” anywhere. Absolutely empty space-time, a so-called perfect vacuum, still has stuff going on. Particles are bubbling in and out of existence on the subatomic level. We've observed this; it's not just unsupported theory. So you start out with what you think of as a big hunk'o'nothing, pre-Big Bang, and it's not really nothing. The kettle is simmering, little blips of mass and energy, very small, happening on very small scales. The blips occur randomly, sometimes they're smaller and sometimes they're bigger. If a big enough blip occurs, you can get a kind of run-away something-for-nothing chain reaction and instead of shrinking it swells and all of a sudden you've got a brand-new universe. Big Bang!

Is this any help at all?

pax \ Ctein
[Please excuse any word-salad. MacSpeech in training!]
==========
-- Ctein's Online Gallery http://ctein.com
-- Digital Restorations http://photo-repair.com
==========

Posted by: ctein | Friday, 18 March 2011 at 01:29 AM

Dear Bill,

Good Lord, do you mean everyone else stays caught up?!

I feel so inadequate...

~~~~~

Dear Toto,
~~~~~

Actually, it happens at the speed of electrons.

But that's not as cool a title.

pax / Ctein

Posted by: ctein | Friday, 18 March 2011 at 01:40 AM

Just out of curiosity Mike, would that be Brassai's Paris by Night?

Posted by: Mani Sitaraman | Friday, 18 March 2011 at 02:49 AM

Coming at this from the other direction, a Rapatronic camera can produce quite usable images with a shutter speed of ~10ns. That's ten nanoseconds, or approximately one-one-hundred-millionth of a second.

I think I first saw a reference to Rapatronic cameras here, but here's a link:

http://www.wikipedia.org/Rapatronic_camera

While Rapatronic cameras are pretty much the definition of non-low-light instruments (they're used to photograph atomic explosions using the light from the explosion itself) it also provides a data point for a discussion of "how few photons can you get away with?"

As for me, my interest in high-speed low-light photography has a usability border roughly equivalent to what I can see with my own eyes. I don't need a camera that can produce an image in an environment where I can't see to frame my shot. I am much more concerned about noise and depth of field, but that's because I mostly shoot people. Even Goth clubs have enough light to see by: if I can get a clean, well-focused image out of the light available in such a place, I'm happy.

Posted by: MarcW | Friday, 18 March 2011 at 08:58 AM

Quote Ctein "Dear Phil,

Ummm, no.

Well, BSI, yes, but otherwise all you're doing is increasing the total area of the sensor. Haven't improved filter or silicon efficiency one bit. What makes you think otherwise?

~~~~~~~~

pax / Ctein"

After the non-sensing elements on the chip, the biggest loss of photons is from the filtering needed to generate color information. A 525nm photon (green) that falls spatially over the red or blue sensor is lost forever. So while the silicon may be 90% efficient at converting photons to electrons, that's only true if the photon hits the sensor. Splitting up the image into component colors on different sensors minimizes this problem. This is one reason commericial video uses 3 chip cameras, in addition to the effectively greater surface area

Quote Semilog:"
"Quote Phil Allen .... (as above)."

The back-illuminated EM-CCD camerain my lab (90% QE @ 525 nm, <0.1 e- read noise per pixel, 14 bits DR) was purchased more than 5 years ago.

Similar cameras coming out today are a bit less expensive but no less bulky due to the requirement for a peltier cooling device that holds the sensor and read circuitry at -80°C. Of course, we image at ambient temperature, so it's necessary to house the whole assembly in a vacuum enclosure to prevent condensation.

I'm not expecting to see this tech in a usable DSLR any time soon.

end quote"

I'd be surprised if you camera had less than 0.1 e-/event read-out noise. Dark noise yes, readout no. Current scientific CMOS cameras on the market today are getting down to about 1 e-/pixel read-out noises. As for dark noise, it falls by ~50% for every 7 degrees of cooling below ambient. So cooling to 0 or -20, which is easy these days usually drives the dark noise down to <0.01 e-/pixel/second.

Phil

Posted by: Phil Allen | Friday, 18 March 2011 at 09:44 AM

Ctein,

thanks for your answers! I'll try out the columns you mention as well. It isn't easy being a dad of an inquisitive daughter, I'm discovering, particularly as my own areas of knowledge are of no interest to her.

I like your second half answer in particular (ie I can begin to understand it sort of intuitively), even if the concept of the odd random bubbling up of little particles pre Big Bang is vaguely unsettling. This was happening before time existed, and before space existed, neither of which I knew were possible.
~~~~~~~~

The piece about Einstein discovering that mass and energy are interconvertible seems obvious to me, and something that anyone looking at a burning log on a fire could not fail to see once mankind knew about atoms and the solid/liquid/gas states. Log gets converted into heat, not especially efficiently because there is some embers left. Similarly, the fire can boil water, turning closely packed water into free-flowing steam, again with not much efficiency as heat is also present in the air and in the metal of the cooking pot. {*Brief Wikipedia break on "Atom" and "Mass-Energy Equivalence"*} Aaah - it was Newton who first proposed it in Query 30 of *"The Opticks"*, but it took until 1905 for Einstein to prove it.

I am indeed educated by reading this blog!

Posted by: James | Friday, 18 March 2011 at 10:34 AM

Dear MarcW,

Ummm, definitely not a dearth of available light in this situation (understatement).

The question being faced was not one of "how few photons" but "how few nsec."

~~~~~~~~

Dear Phil,

No, false logic. You've tripled the total area of the sensor, but left the number of output pixels the same. That's the gain in apparent sensitivity. It's not a gain in efficiency, just more collection area. A good thing, but you can always get more ISO by making pixels bigger. The question of the day is how much ISO you can get out of a given size pixel?

If you made each of the three chips smaller so that the individual cels were the same size as in the Bayer array camera, it wouldn't perform one bit better.

Not that three-ship cameras don't have advantages. It is indeed a way to pack more sensor area into the same format and to eliminate chroma aliasing. Not dissing them. But each chip still has the losses I described in this article.

Semilog is talking about a special scientific camera that's been designed to catch and use just about every photon. He's pointing out to lay people that I'm not just being hypothetical when I talk about approaching 100% efficiency. We've known how to do that for a long time.

What I'm writing about this column is how far the cameras you'd use for conventional photography are from this limit.

Hope this makes things clearer.

pax / Ctein

Posted by: ctein | Friday, 18 March 2011 at 11:56 AM

Dear Phil,

A P.S. Maybe it'll be clearer this way--

Take your BSI RGB-filtered three chip camera.

Replace the RGB filters with the Kodak filters (it's the spectral bandpasses that give most of the gain, not the geometric arrangement). You gain a stop in speed/efficiency.

Add a sensitizer layer to eliminate the inherent silicon losses. Another stop for you.

Can't get that if your camera was running near theoretical 100%.

pax / Ctein

Posted by: ctein | Friday, 18 March 2011 at 12:05 PM

"It isn't easy being a dad of an inquisitive daughter, I'm discovering, particularly as my own areas of knowledge are of no interest to her."

Amen, brother. My brother is the doctor, and my own kid's questions were 90% medical and scientific. What I wouldn't have given for, "Daddy, I'll take arts and literature for 600, please."

Mike

Posted by: Mike Johnston | Friday, 18 March 2011 at 03:19 PM

In fiber optics, when the small # of photons issue cropped up, systems went coherent. Next there is the " entangled photon" idea. Never say never...

Posted by: Jim Hayes | Saturday, 19 March 2011 at 02:08 AM

Dear Jim,

Coincidentally, I had been toying with ideas for making mixed incoherent-coherent systems that could take advantage of entanglement and/or squeezed light while writing this column.

Nothing's jelled, so I've not yet convinced myself it's possible. But not convinced it isn't, either.

So, yeah, sqrt(N) noise may not be the lower limit. Might be able to squeeze out another factor of two or so, beyond even that.

Since you're borrowing from Peter to pay Paul, informationwise, you couldn't push much further without degrading the image pretty severely.

Pure speculation, at this point.

pax / Ctein

Posted by: ctein | Sunday, 20 March 2011 at 03:15 PM

The comments to this entry are closed.
~~~~~~~~