SPECULATIONS ON THE MEGAPIXEL RACE OR "DO I NEED A + 32 MEGAPIXEL SENSOR?"
By S.Saragosa a.k.a. ***tjshot***

Recent announcements of 24 Mpxls APS-C and rumors of 36 Mpxls full frame sensors seem to have changed the scenario for digital camera enthusiasts, triggering opposite reactions: some of us are out of the skin to get the new babies (I'm among them) while others seem to question the effectiveness of such upgrades.
We had similar doubts in the past when upgrading to higher density sensors and I've read many articles on the internet stating that you don't really need more than about 20 Mpxls full frame sensor (due to lense diffraction/aberration, sensor sampling performance etc.)

I hardly agree.

At the level of development we have reached today a small increase of pixel pitch (effective sampling rate -> sharpness and resolution boost) in the sensor will induce a much higher (square growth law) megapixel count, thus requiring much larger in-camera buffer memory, faster data saving channels, more powerful computers etc.

At the same time noise control becomes increasingly difficult (square growth law) due to higher amplification of the signal needed to achieve the same ISO rating; on the other side smaller pitch sensors seem to produce a much more even and defined noise, partly counterbalancing the degradation in quality (think about the terrible "pepper" grain in Canon EOS 5D compared to the smoother Canon EOS 7D).

Many additional elements factor in, resulting in final image quality: A/D converter performance (12bit to 14bit seems to produce a very high boost in dynamic range, noise performance etc.), image processor, firmware etc.

We'll eventually get to a break-even point where the pros in increased sampling capability will match cons in noise and data rate, when it comes to final image quality and system usability.

I believe we have not reached this point yet.

My take on the matter is that for full frame bodies 32-36 Mpxls still makes sense in terms of performance improvement; 32 Mpxls is probably a better trade-off for the advanced user.

Lets consider a few factors.

1) PRINT SIZE

With + 32 Mpxls it's possible to print to A3 + format (13x19 inches at + 350 dpi) without interpolation, which is the upper (easily usable) size limit for most advanced home printers (I'm assuming you are strict about colour management and want to directly print your work); if you regularly need more then you are probably a fashion pro and should consider (could afford) digital medium format.

Interpolation to higher/lower sizes will always cause a loss of resolution and quality (upscaling being the worse case); I'd rather shoot with a resolution that exactly fits my highest requirements than having to pay (in terms of noise, image size and ... money) for an overhead that I'll hardly ever use.

2) SHARPNESS AND RESOLUTION

That's where flame wars usually start on forums.
The short story: 32-36 Mpxls on a Bayer matrix sensor is still "low enough" to squeeze significantly more information out of most of your lenses.
I see 65 Mpxls as the practical limit for an excellent prime at F 8.

Probably no much sense in going over 50 Mpxls anyway (due to noise and dynamic range).

The long story: sharpness and resolution in a camera system depend mainly on lense performance and sensor sampling rate (pixel pitch in digital).

Lenses for 35 mm usually reach top performance between F 5.6 and F 8; wider apertures are affected by aberrations (colour fringes, halos, micro distortion like "coma" etc.), smaller apertures start to be affected by diffraction (an unavoidable physical property of light).

I'll keep a qualitative approach and assume all the definitions and technicalities are known; many sources are available on the internet, Norman Koren's site being the one offering the most comprehensive photography-related background. A worth reading if you are math oriented.
(Address: http://www.normankoren.com/Tutorials/MTF.html)

Diffraction is a property of light and sets the top obtainable performance of a lense for a given F-stop: for each F-stop and for a given wavelenght of light, there is a corresponding limiting resolution due to diffraction; a real lense can only perform lower than this limit, due to "losses" in the glass, by a margin depending on its quality.

With actual technology it's relatively feasible to produce a nearly diffraction limited lense in the range F11 - F16; very few manufacturers seem to be interested in producing a commercially available lense (i.e. a profitable design) which is diffraction limited at wider F-stops (probably only Leica).
Best lenses differ from average ones in their ability to control the aberrations that limit performance at wider apertures; to achieve a good result at extreme apertures (think of the best Leica, Canon, Zeiss fast primes with maximum F 1 - F 1.4 stop) without raising the price to unprofitable levels (and dropping light transmission), a few design compromises are often required, which involve some degradation of performance at more typical F-stops.

In the following I'll propose some sharpness/resolution calculations, further developing some assumptions made in an article "Do sensors outresolve lenses?", published on Luminous Landscape that seems to have received some attention.
(Address: http://www.luminous-landscape.com/tutorials/resolution.shtml)

Based on practical experience and some math speculation, a Bayer sensor, with proper sharpening procedures, is able of reproducing a signal with good accuracy when its sampling rate is at least 3 pixels per cycle (1 cycle= 1 line pair).

In the aforementioned article "required megapixels" calculations are based on a 2 pixels per line pair assumption, which implies a perfect alignment between signal (line pairs) and sample units (pixels) plus a sensor resolving straight down to its Nyquist limit; in practice the two will often be out of phase (misaligned) and anti alias filters, required for Bayer matrix technology, will degrade performance at higher frequencies (lpm), thus involving an oversampling requirement.
4 pixels per line pair seems to deliver optimal results with Bayer sensors.
3 is ok for most real scenarios, if proper sharpening is applied.

About definition of "proper sharpening" there seem to be different opinions; I'm partial to integer radius values (mostly 1 and 2 radius) as I find they involve less software-induced aliasing.
(Tech note: radius 1 sharpening is good when you have a sharp image to start with, as it generally boosts response at Nyquist frequency i.e. in the very high lpms near the resolution limit of the sensor(however it depends on sharpening algorithm used by software); radius 2 may be better for not-so sharp images, as it generally boosts response at lower frequencies (half Nyquist); for actual sensors with 4000 to 5000 dpi it falls in the 38 - 48 lpm range, corresponding to the combined MTF 50% value (perceived sharpness) of the sensor plus a good - excellent lense.

In other words radius 2 sharpening, with a very good lense and a 21 + Mpxls sensor, lets you recover image information in the "lpm zone" where our eye-brain system perceives sharpness. Very effective.)

Following the 3 pixels per line pair sampling assumption, a 36 Mpxls Bayer matrix full frame sensor is able of registering most of the relevant information available from a nearly diffraction limited lense at F11: this is a pretty common performance both for modern primes and quality zoom lenses.
In the following I'll propose some simulations obtained using real data from DXOmark site (Address: http://www.dxomark.com)and a Matlab function proposed by Norman Koren.

To keep things simpler we'll consider only lense + camera, but printers are also involved in practice.

As a reference for interpretation of following data, let's consider that:
- The impression of sharpness is often related to the MTF 50% value of the imaging chain system: that's the frequency where contrast drops to half the value of the original target.
In other words that's the value that suggests how much "WOW factor" the picture delivers due to sharpness.
- The resolution is often related to the MTF 10% value of the imaging chain system: that's near the maximum perceivable detail a system is able of delivering before grain and noise start to prevail and mask information.
You won't probably notice that detail unless enlarging the image very much (40 x with actual sensor pitch) on a monitor (or under a loupe film) and it would not be "clean enough" to be pleasant to the eye; so it's often assumed that it does not contribute much to the "WOW factor".
The more you'll enlarge, the more you'll be relying on low MTF information to build up your final, printed image.
Just consider these as general guidelines, as perception and quality standards are also a subjective matter.
Just to show the basics of the simulation, I'll start with a scenario that is well known to many users owning a current top DLSR: a "20 Mpxls range" full frame sensor (Canon 5D Mk II, but it can be extended to other top line bodies from Nikon and Sony).
Values for such sensor with a perfect (diffraction limited) lense at different F-stop would be:

	MTF values for a diffraction limited lense			**MTF values for a diffraction limited lense + 21 Mpxls Full Frame Sensor**				
	Wavelenght of light = 0,00055			Wavelenght of light = 0,00055				
				No Sharpening		Sharpened		
F-stop		MTF 50% (lpm)	MTF 10% (lpm)	MTF 50% (lpm)	MTF 10% (lpm)	MTF 50% (lpm)	MTF 10% (lpm)	Sharpening Parameters
2,8		247	532					
4		173	373					
5,6		123	266	54	78	76	76	K=0,35 r=1
8		86	186	49	78	70	78	K=0,4 r=1
11		63	135	44	78	65	78	K=0,45 r=1
16		43	93	35	63	60	78	K=0,4 r=1
22		31	68	28	52	48	78	K=0,7 r=1

Where:

Wavelenght of light is set about in the middle of the visible spectrum (Green-Yellow zone) as it's a common situation in daylight shooting; shifting towards blue will increase resolution, toward red will decrease it.

K is the edge contrast boost ratio, usually set in % value by the slider in the sharpening software.

R is the radius of sharpening, also set in the sharpening software.

In previous table, sharpening parameters were set to what I consider the highest safe values still allowing a reasonable control of halos or artefacts.
Values for F 2.8 and F 4 are not reported as they are of no practical use for comparison with real lenses.

Let's see now how the sensor behaves with the simulation of a real excellent prime lense (think of a good sample of a Canon 85mm F 1.8 ; it's easy to match the simulation with real data from DXOmark site):

	MTF values for excellent real prime			MTF values for for excelent real prime + 21 Mpxls Full Frame Sensor				
	Wavelenght of light = 0,00055			Wavelenght of light = 0,00055				
				No Sharpening		Sharpened		
F-stop		MTF 50% (lpm)	MTF 10% (lpm)	MTF 50% (lpm)	MTF 10% (lpm)	MTF 50% (lpm)	MTF 10% (lpm)	Sharpening Parameters
2,8				37	75	63	78	K=0,5 r=1
4				37	75	63	78	K=0,5 r=1
5,6		65	195	41	78	65	78	K=0,45 r=1
8		53	159	37	75	63	78	K=0,5 r=1
11		45	135	34	70	61	78	K=0,55 r=1
16		38	91	31	62	55	78	K=0,6 r=1
22		28	68	25	52	49	78	K=0,7 r=1

For the simulated lense MTF 50% values were chosen to match the empirical data from the DXOmark site; MTF 10% were calculated using a simulation of “how the lense looses resolution” towards the higher lpm range(tech note: an MTF simulation curve with an order 2 rolloff was used for the F 5.6 – F11 range, i.e. the lense is performing well under the diffraction limit; an order 2.5 rolloff was used for F 16 – F22 as the lense behaves closer to the diffraction limited lense model in that range)

A real lense does not reach the peak MTF 50% value of the ideal diffraction limited one but on the other side it looses resolving power more gradually towards the highest frequencies.

Modern lense design for 35mm is optimized for performance at wider apertures: quality lenses will peak at F 5.6, a few even at F 4, and gradually loose performance at smaller apertures, where the quality is uniform across the whole field.

At wider apertures and up to F 2.8 they are generally able to keep top performance in the centre of field while degrading towards corners; different designs result in very different performances in this range.
Proposed values for F 2.8 – F4 are to be intended as top performance in the center of field.

Due to lower MTF value of a real lense I pushed sharpening parameters a bit compared to the diffraction limited scenario, as halos on the low frequency edges are less probable to occur.

Now let's simulate the same scenarios for a 36 Mpixels full frame Bayer Sensor:

Diffraction limited scenario:

	MTF values for a diffraction limited lense			MTF values for a diffraction limited lense + 36 Mpxls Full Frame Sensor				
	Wavelenght of light = 0,00055			Wavelenght of light = 0,00055				
				No Sharpening		Sharpened		
F-stop		MTF 50% (lpm)	MTF 10% (lpm)	MTF 50% (lpm)	MTF 10% (lpm)	MTF 50% (lpm)	MTF 10% (lpm)	Sharpening Parameters
2,8		247	532					
4		173	373					
5,6		123	266	66	105	91	105	K=0,35 r=1
8		86	186	58	102	89	105	K=0,5 r=1
11		63	135	49	88	72	105	K=0,5 r=1
16		43	93	38	70	61	97	K=0,6 r=1
22		31	68	29	56	49	86	K=0,7 r=1

Real lense scenario:

	MTF values for excellent real prime			MTF values for for excellent real prime + 36 Mpxls Full Frame Sensor				
	Wavelenght of light = 0,00055			Wavelenght of light = 0,00055				
				No Sharpening		Sharpened		
F-stop		MTF 50% (lpm)	MTF 10% (lpm)	MTF 50% (lpm)	MTF 10% (lpm)	MTF 50% (lpm)	MTF 10% (lpm)	Sharpening Parameters
2,8				41	88	80	105	K=0,5 r=1
4				41	88	80	105	K=0,5 r=1
5,6		65	195	47	95	80	105	K=0,5 r=1
8		53	159	41	88	80	105	K=0,5 r=1
11		45	135	37	82	77	105	K=0,65 r=1
16		38	91	33	70	67	105	K=0,7 r=1
22		28	68	26	56	60	100	K=0,8 r=1

At a first glance, with no sharpening applied, it seems that for a real lense at smaller apertures there is no performance gain in using a denser sensor: from F 16 on, both a Canon 5D MKII (diffraction limited at F16 according to previous assumptions) and a 36 Mps full frame sensor will perform at the same level with only marginal differences (a few may occur from A/D converter and image processor performance).
The same effect can be seen today comparing the MTF graphs from DXOmark at F 11 and F 16 for a Canon EOS 20D (same pitch than 5D MK II) and Canon EOS 500D (about same pitch of a 39 Mpxls full frame sensor).
At F 11 the 500D still has an edge, at F 16 they have the same performance.
But the real difference comes out when sharpening is applied: the denser pitch is able of recovering much more of the low MTF information coming from the lense , delivering details up to its higher Nyquist resolution limit.
Moreover a denser pixel pitch will also allow higher sharpening parameters than a lower pitch, at similar Signal to Noise Ratio (that's the catch: hard to keep up the SNR while increasing pitch)
In other words the effect of sharpening is boosted by the increase in sensor pitch.
The difference will be more noticeable for big enlargements, digging into the low MTF zone.
So more megapixels actually mean more information, but it becomes clear only after proper processing of the image.

Now let's get further and simulate the same scenarios for a 50 Mpixels full frame Bayer Sensor:

Diffraction limited scenario:

	MTF values for a diffraction limited lense			MTF values for a diffraction limited lense + 50 Mpxls Full Frame Sensor				
	Wavelenght of light = 0,00055			Wavelenght of light = 0,00055				
				No Sharpening		Sharpened		
F-stop		MTF 50% (lpm)	MTF 10% (lpm)	MTF 50% (lpm)	MTF 10% (lpm)	MTF 50% (lpm)	MTF 10% (lpm)	Sharpening Parameters
2,8		247	532					
4		173	373					
5,6		123	266	74	115	102	115	K=0,35 r=1
8		86	186	63	115	90	115	K=0,45 r=1
11		63	135	52	95	78	115	K=0,5 r=1
16		43	93	40	74	68	115	K=0,7 r=1
22		31	68	30	58	50	94	K=0,75 r=1

Real lense scenario:

	MTF values for excellent real prime			MTF values for excellent real prime + 50 Mpxls Full Frame Sensor				
	Wavelenght of light = 0,00055			Wavelenght of light = 0,00055				
				No Sharpening		Sharpened		
F-stop		MTF 50% (lpm)	MTF 10% (lpm)	MTF 50% (lpm)	MTF 10% (lpm)	MTF 50% (lpm)	MTF 10% (lpm)	Sharpening Parameters
2,8				44	96	90	115	K=0,65 r=1
4				44	96	90	115	K=0,65 r=1
5,6		65	195	50	105	90	115	K=0,55 r=1
8		53	159	44	96	90	115	K=0,65 r=1
11		45	135	39	89	88	115	K=0,7 r=1
16		38	91	34	74	74	115	K=0,75 r=1
22		28	68	26	58	69	115	K=0,85 r=1

All the evaluations made for the 36 Mpxls sensors are still applicable.
After sharpening the 50 Mpxls sensor we get for every F-stop a + 10 lpm difference to the 36 Mpxls one and a + 20-25 lpm difference to the 21 Mpxls one. Sounds pretty good for a camera that is supposed to perform the same as the other two from F 8 down to F 22 due to diffraction.
Probably however the smaller stops F 16 -F 22 will suffer from excessive sharpening at the proposed parameters and most user would settle with safer values.
All the mayor sensor producers seem to actually own the technology to produce a 45-50 mps full frame sensor; a full frame body with the canon EOS 7D pitch would be about 50 mps. This is close to the Mpxls needed to sample a real excelent prime lense at F 5.6 - F 8.
Given the performance of modern pro line primes, there is still enough information delivered to the sensor to make a big difference at F 5.6 , F 8 and maybe F 4.
Furthermore, with such pixel density it would probably be possible to drop the antialias filter (Leica had done it with the M9!) an have a further boost in noise and sharpness (none in resolution), having the sensor resolve down to the Nyquist frequency with far less contrast loss, in a way similar to actual Foveon sensors (slightly less effectively).

It could be possible to apply less sharpening.

With diffraction limited lenses the MTF at Nyquist would still be too high to allow dropping the antialias filter: color artifacts would occur when shooting very fine and regular detail (think of the old Kodak DCS Pro 14n full frame camera).
With a real lense however, the natural rolloff of resolving power gradually cuts out the higher frequencies that would induce aliasing near the Nyquist limit: the lense itself provides the necessary "antialias filter" for the sensor.

Noise performance and dynamic range with a 14bit A/D converter should still be good (think of Canon EOS 7D), just a small step (10-15%) behind Canon EOS 5D MKII.
Sounds pretty usable to me.
I probably would not need the extra file size for printing, but still it makes sense.
I don't think actual sensor technology is worth pushing over 50 mps, which I consider being the break even point considering pros and cons.
Technology is already implemented but production costs would be much higher than Canon EOS 7D due to yield in sensor manufacturing.
Given my typical use of camera, I'd be better served with a decent dynamic/noise range 32 Mpxls full frame sensor, still getting the most out of my old lenses; that would match the practical resolution of the top medium format film systems, while retaining the wider DOF of 35mm which helps a lot when doing landscapes.
A 16-20 Mpxls full frame sensor, even with a 14 fps shooting rate, an extraordinary wide dynamic and noise range, simply do not appeal me much as I'm no sports or wildlife shooter.

3) FILE SIZE AND PROCESSING POWER

I don't see real problems here, at least outside the in-camera buffer.
Actual 64 bit CPUs and Operative Systems allowed us to get rid of computing limitations experienced with big images on 32bit systems.
I usually work with 300 Mb to 600 Mb files obtained scanning medium and large format slides; I frequently make panos in the 2-3 gb size range.
A core i7 920 processor with 6 Gb ram is more than enough to handle this workflow; it's a pretty common and inexpensive configuration nowadays and computing times are acceptable in most of the cases.
For large Gb images I simply crop a relevant smaller portion to make adjustments on and then save the workflow as a batch to be executed on the whole image; it's common practice for many photographers and allows to evaluate adjustments in real time even on large files.

A final consideration about APS-C and smaller formats.
Given my max print size of A3 +, such a small sensor simply does not match my requirements: even with very high pixel densities (32 to 50 mps) it would take a n almost perfect lense in the F 4 to F 5 range to deliver appropriate resolution to the sensor. There are many good APS-C lenses, but no one performing to that level.
Noise would be hard to cope with and tonal separation would be considerably worse than full frame.
Furthermore you lose the possibility to have a shallow DOF and shoot in available light conditions: there are a few nice full frame F 1 - F 1.4 lenses, but no equivalent F 0.6 - F 0.9
It's a no go for me.

The best upgrade pattern for an APS-C shooter not interested in printing extra-large would be moving to an equivalent full frame resolution sensor; a cheap (1000-15000 USD), fast 18 Mpxls full frame body would deliver a lot in terms of quality and usability and win the favour of many actual APS-C users.
Make it a Foveon-like (multilayer) sensor, drop mirror/antialias filter and it's just the perfect recipe.

In my opinion, full frame DSLRs seem to have become what medium format once was (plus a much faster and easier workflow): the goldilocks of photography.

Final note on the biggest improvement I've seen in the market recently: the Sony/Fuji Electronic View Finders are an incredible feature that make the camera much more usable. An example to follow. Hope Canon finds its way to drop away that mirror...

The scenarios I describe offer me the general guidelines for future system upgrades, with a very probable step up to a + 32 Mpxls and an eventual move to 50 Mpxls if and when they'll become accessible; I don't see a good point in moving on, unless a totally new, cheap technology is developed.