Telairity Deep Dives Into 4K Technology – Part 6

Since the whole resolution difference between a "4K" UHD image (composed of roughly 8 million pixels) and a "2K" HD image (composed of roughly 2 million pixels) boils down to a 4X increase in pixel count, and a pixel is nothing but a string of 24 1s and 0s ("bits" in digital speak), the obvious conclusion is that the shift from HD to 4K UHD involves moving or storing 4 times as many bits per unit of time as was needed for HD. More precisely, the uncompressed HD figure of about 1.5 billion bits per second (1.5Gbps) jumps to about 6 billion bits per second (6Gbps) for UHD; while the compressed HD figure of about 5 million bits per second (5Mbps)—barring any substantial improvements in compression technology—jumps to about 20 million bits per second (20Mbps).

In fact, however, the challenge posed by UHD is worse than a mere 4X increase in data rates, because the UHD standard encompasses more than just a 4X increase in pixel count. The UHD standard also allows a shift from "8-bit" to "10-bit" color, and a shift from a rate of 30 frames per second (30 fps) to a rate of 60 fps. Each of these changes further boosts the data rates required for UHD video.

The move from 8-bit "true color" to 10-bit "deep color" channels[4] increases pixel length from 3 x 8 = 24 bits to 3 x 10 = 30 bits. Which is to say, the new UHD "deep color" standard requires a further 25% increase in bit rates. This means raw data rates will not be just 6Gbps, but rather 7.5Gbps, to accommodate the 6 extra color bits for each pixel.

And that increase is not the end of the story, given the UHD standard also allows doubling the HD frame rate. Transmitting twice as many frames a second requires a further doubling of the bit rate, so the 7.5Gbps need to transmit "deep color" UHD at 30 fps becomes instead a data rate of 15Gbps, to transmit the new UHD data at 60 fps.

So the complete UHD story, at least in terms of the increased bitrates required, does not end with simply quadrupling HD resolution. That is merely its beginning. In addition, the UHD standard makes each pixel 25% longer, to carry new "deep color" information, and also provides for a doubling of frame rates. The cumulative effect of all these changes is that the 4X penalty in bit rates imposed by the increased pixel counts balloons into a 4 x 1.25 x 2 = 10X increase in total bit rates.

In short, the shift from full HD to the full 4K UHD standard requires an order of magnitude jump in data rates. At the level of uncompressed data, that's a leap from1.5Gbps to 15Gbps. At the compressed level—again, barring significant improvements in compression technology—it's a jump from 5Mbps to 50Mbps.

Is an order of magnitude change in bit rates a problem? It all depends on who you are and where you're located in the video chain. For some, UHD is all gain. For others, it is a massive headache. To be continued in the next part of this series.

[4]You may recall from the note on binary arithmetic in part 4 of this series that 8-bit color channels are called "true color" because coding a color with 8 bits means distinguishing 256 separate shades of it. 256 shades of a color are not only more shades than the human eye can perceive, but, when the 3 primary shades (Red-Green-Blue or RGB), are mixed, 256 x 256 x 256 multiplies out to over 16 million possibilities—a total that includes many, many millions of colors that are indistinguishable to the human eye.

So what is the advantage of a new 10-bit "deep color" standard? The binary logic of 10 bits, at least, is straightforward. Each new bit added to a binary number doubles the number of combinations previously possible, so the move from 8-bit to 10-bit color first doubles the number of possible color shades from 256 to 512 (at 9 bits), and then redoubles that number from 512 to 1024 (at 10 bits). 1024 x 1024 x 1024 multiplies out to over a billion (1,073,741,824) possible RGB "mixed" colors, or 64X (4 x 4 x 4) the number of 8-bit colors. But, if we can't perceive as many as 16 million colors, what is the advantage of increasing that number by a factor of 64? Isn't is just adding over a thousand million more colors no one will ever see?

There are two answers to this question. The first answer, regrettably, is that the marketing of "deep color" often relies on a simple "more is better" argument. If 8 bits are good, then 10 bits must be even better. The name says it all. With 10-bit coding, colors will have to be somehow "deeper". Why settle for a mere 16 million colors when you could have a billion?

The fallacy of this argument is not hard to spot. In fact, 10-bits does not make any 8-bit color "deeper", rather, it simply divides each 8-bit shade into 4 separate sub-shades, all of them indistinguishable to the human eye. This point is worth dwelling on because there is a deeper truth lurking here, just below the surface.

The term "deep color" actually embraces any color depth past the 8-bit "true color" standard, including 12, 14, and even 16-bit colors. So, if deeper is really better, there is no good reason to stop at 10. Rather, we should move on to add color bits as rapidly as technically feasible. Why not 64-bit color? 1024-bit color? Why should the ambition for more color bits ever end?

The deeper truth here is that, with digital coding, "more is better" is not a good argument, precisely because it implies code lengths should be expanded indefinitely, always to the limits of current technical feasibility. This is just silly. The real question when it comes to digital code lengths is not: Can we make it bigger? – since the answer to this question is always "yes" – rather, the real question is: How big does a code need to be? Or, alternatively: How big is big enough?

An example may help. The first microprocessors were just 4-bit machines. But, while 4-bits is good for simple tasks, over the next 25 years, microprocessors expanded to 8 bits, then 16 bits, then 32 bits, and finally to 64 bits. But, since reaching 64-bits about 20 years ago, there has been no clamor for yet another upgrade to 128 bits, and no such demand seems to be looming just over the horizon. The reality is that 64 bits are not only enough for any ordinary computing task, they are enough for virtually any supercomputing task. Throwing vast resources at creating a whole new generation of unneeded 128-bit machines would be not just silly but insane.

Coming back to color depth, there is no problem with expanding bit lengths. Instead of millions, we could have billions or trillions or quadrillions of colors. All it takes is more bits. But none of this expansion changes the human eye. If the question is: How many bits are needed to exceed the capacity of the human eye? – the answer is exactly 8. Add as many colors as you like past the 8-bit "true color" space, it remains certain that no human will ever discriminate even one of these possibilities.

But, if "more is better" is a bad argument for expanding past 8-bit color, and any such expansion of digital color codes adds absolutely nothing to the range of our human color perception, then why does the UHD standard propose 10-bit color? Surely the experts behind this standard must have *some* good reason for wanting to increase color channel depth. In fact, there are two sensible reasons for wanting more than 8-bits/256 shades per primary color.

The first reason pertains to capture devices, like cameras and scanners. As we mentioned earlier, when discussing digital compression, the technologies used for image capture are often far more sensitive to different wavelengths of light than the human eye. The fact that the human eye tops out at perhaps 150 or so color shades does not limit our image sensors, which may discriminate far more than that number. For these superhumanly sensitive capture instruments, recording all of information they make available could require 10 bits, 12 bits, or still deeper color channels. Similarly, old color photos contain subtleties the eye does not see. Hence, to digitize all the information available in an old family photo album, it's a good idea to shop for a deep color scanner that advertises, for example, "48-bit color" (16-bit channels).

The second reason has to do with the "lossy" nature of digital image compression. Since images are compressed by throwing away information, if an image is to undergo several rounds of editing and compression before reaching final form, it helps to begin with a lot more information than is needed at the last step. Remember that 8-bit "true color" stops at the very first bit past human perceptual limits. But if the result of editing the millions of colors in an 8-bit image is to reduce them to mere thousands of colors, the losses very likely will be perceptible. In this sort of case, it is a good idea to begin, not with millions, but rather billions of colors, since reducing billions to mere millions is unlikely to result in any perceptible degradation.

Neither of these reasons, however—which have to do with image capture and image processing, respectively—provides any reason for transmitting 10-bit "deep color" to displays. Displays are all about human perception and that, to repeat, stops at 8 bits. Nonetheless, there is a reason why displays might want to decode 10 bits of information per pixel. That reason, though, is ***not*** about showing new mixed RGB colors that no human eye will ever see. Rather, it has to do with a color technology known as High Dynamic Range (HDR). But HDR color is another story.