Beyond HD:
How Ultra can UHD resolutions go?

A TELAIRITY WHITE PAPER

Harlan McGhan
TELAIRITY | 3375 SCOTT BLVD., SUITE 342, SANTA CLARA, CA 95054

Table of Contents

A Deep Dive Into
4K Technology

Just when it might be thought that so-called “Full HD” resolution (or “2K”) was the absolute cutting edge, “Ultra-High Definition” (UHD) made an appearance and changed the equation. Actually 4K UHD technology has been in the news since 2010, but 2015 saw a steep drop in the prices of devices supporting 4K and, judging by the latest sales figures, the market has warmed considerably to the new resolution standard. But, just as 4K is starting to replace HD, some broadcasters are starting to promote 8K as the natural and inevitable successor to 4K—an upgrade that apparently cannot arrive too soon, given its advantages over 4K.

What in the world is going on? The SD standard lasted over half a century. Now, a new resolution standard (HD, 4K, 8K) seems to be appearing every decade or so. Is 8K to be followed by 16K, which in turn will be superseded by 32K? Is there any end to this ceaseless upgrade cycle of digital resolution standards?

1 What is Digital 4K Technology

1.1 First, Resolution Is About Pixel Count

In the old analog world of CRT display screens, screens were measured by the number of scan lines they supported. A standard SD screen had either 480 or 576 visible scan lines (depending on which standard, NTSC or PAL, was used in a region). The shift to HD, with its much higher resolutions, precipitated a shift in the underlying technology, from analog to digital.

In the new digital world of TV technology, display screens are no longer measured by the number of scan lines they support, but rather by the number of pixels they display. Digitally speaking, the highest resolution SD screens are now pixel arrays 720 wide x 480 or 576 high. Full HD screens are pixel arrays 1920 wide x 1080 high. Rounded off, HD resolution is about 2000 x 1000, which gets shortened to “2K” in digital-speak. UHD or 4K simply doubles each of these HD (2K) dimensions, to about 4000 x 2000 (in exact terms, to 3840 x 2160). In round numbers, then, HD is about 2 million pixels/screen, while UHD is about 8 million pixels/screen, or 4X the resolution of HD.

1.2 Much Remains Outside Pixel Count

It is easy to get confused here, because pixel count is only one aspect of the technology used to manufacture displays. Another critical aspect is the technology used to render pixels (whatever their number). This is where you encounter words like "Plasma" and acronyms like "LCD" (Liquid Crystal Display), "LED" (Light Emitting Diode), "OLED" (Organic LED), "QLED" (Quantum Dot LED), and so on. Rendering technology controls the maximum darkness (black) and lightness (white) of a screen (its contrast ratio), as well as how bright and vivid colors appear. Yet another issue has to do with screen shape (curved or flat), and the effect this has on the viewing experience.

Obviously, if you change multiple aspects of a display at once, the total impact on the viewing experience can be far greater than the impact of any one change in isolation. No doubt, shifting from a flat LCD HD display to a curved OLED UHD display will dramatically transform one's viewing experience. But what part of this transformation is specifically contributed by the change in pixel count, i.e., the shift from HD to UHD? And what by the other independent new technologies for displays (like rendering technology and screen shape) now coming into commercial use?

1.3 Second, Resolution Specifies Data Rates

The important point about digital bitmap formats like HD and UHD is that they fix the number of pixels to be displayed in each video frame, independent of screen size. Every "full HD" screen is an array of 1920 x 1080 pixels, whether the screen measures 30" or 70" or some other number. Similarly, every 4K UHD screen is an array of 3840 x 2160 pixels, regardless of how large or small the UHD screen.

The standard also specifies the format of individual pixels in the array, i.e., whether the pixels are "8 bit", "10 bit", or some other number of bits. Multiplying these two numbers together—pixels-per-frame x bits-per-pixel—will generate a bits-per-frame number. For example, using "8-bit" pixels, a full HD frame is 1920 x 1080 = 2,073,600 pixels x 24 bits/pixel[1] = 49,766,400 bits.

The last piece of the data rate puzzle is for the standard to specify a frame rate per second, or fps value for video. The long-time motion picture standard is 24 fps, but both the NTSC SD and 1080i "full HD" standards use 30 fps, while the PAL SD standard uses 25 fps. The 720p HD standard uses 60 fps, the same as the 4K UHD standard.

---

[1]Why are pixels 24 bits long described as "8-bit"? It's because "8-bit" refers not to pixel length, but rather to "channel" length, or the number of bits used to encode each of the 3 primary colors (Red-Green-Blue) that make up a pixel. Adding the 3 8-bit primary color "channels" together gives the overall total of 3 x 8 or 24 bits/pixel.

Note that while the pixel array (H x W) and the length of a pixel is fixed by the resolution standard, the standard does not fix pixel size/shape. These "local" pixel attributes are determined by the display screen receiving the data, and will vary widely from small "retina" displays to big screen TVs. Even specific color specifications (that is to say, the particular 8+8+8 RGB data for individual pixels), will render differently on different types of displays.

Continuing with our HD example, to get the total bits per second (bps) data rate, we have to multiply the roughly 50 million bits/frame of full HD by 30 fps, which yields 1.5 billion bps (in round numbers, or precisely: 1,493,292,000 bps).

1.4 Pixel Size and PPI Fixed by the Display

Obviously, with a fixed number of pixels—roughly, 2M in a 2x1 "2K" HD bitmap, 8M in a 4x2 "4K" UHD bitmap—what must happen as an HD or UHD screen gets larger or smaller is that the individual pixels in the array must grow or shrink accordingly. This brings us to yet another critical metric for displays, known as ppi or pixels-per-inch. Although an old idea (familiar to anyone who has ever bought a raster printer as dpi or dots-per-inch), this metric was first popularized for displays by Apple, with the term "retina display", meaning a display where the pixels are too small to be individually distinguished by the human eye, even on relatively close handheld viewing. In ppi terms, pixels get too small to be seen (by all but the most eagle-eyed on very close scrutiny) somewhere just short of the number 300, so a "retina display" is any screen with a ppi number of 300 or greater. Note that 300 is also the classic dpi number for quality printing/scanning.

1.5 "Recommended Viewing Distance"

The notion of ppi, in turn, brings us to our final critical metric, viewing distance. Even the largest pixels can be made too small to be individually distinguished by the human eye, by the simple expedient of moving the eye further away from the display. This is the principle behind "Jumbotron" displays, that have pixels the size of playing cards (or bigger), but are designed to be viewed from hundreds of feet away.

If TV screens were built to retina display standards, intended to withstand close scrutiny from a foot away, they would be disappointingly small. An HD screen built to the "retina display" threshold of 300 ppi would be smaller than 7 x 4 inches (about the size of many current smartphone screens). Even with 4K UHD resolution, a "retina display" would be less than 13 x 8 inches, or about the size of a tablet screen.

The reason TV screens of 50" and more are common is simply that TVs are not designed for close up "retina display" viewing. As screens (pixels) are made bigger, the adjustment made by display manufacturers is simply to increase the recommended viewing distance (thereby maintaining a constant apparent pixel size in the eye of the viewer). Conversely, as screens (pixels) get smaller, viewers are allowed to move closer, following recommended viewing distance guidelines, again with no change in the apparent pixel size.

1.6 Basic Difference Between HD and UHD

In a nutshell, then—ignoring, for the moment, differences in pixel length ("8-bit" vs. "10-bit") and frame rate (30 vs. 60 fps)—the whole technical difference between an HD display and a UHD display reduces to relative pixel/screen size. Since UHD formats cram 4X the number of pixels onto a screen as HD, for screens of the same size, UHD pixels are ¼ the size of HD pixels; conversely, for pixels of the same size, UHD screens have 4X the viewing area of HD screens.

Thus, in basic terms (holding the pixel length and frame rate parts of resolution standard constant), the whole viewing difference between an HD display and a UHD display comes down to just one point: bigger screens with no loss of visual quality—where "visual quality" is measured by the single metric of apparent pixel size. It makes no difference whether you replace your old display with a new UHD display of the same size and move 4X closer to it; or keep the same viewing distance, but replace your old display with a 4X bigger UHD display. In both cases, the effect is the same: the screen looms 4X larger in your visual field.

1.7 A More Immersive Viewing Experience

The ability to increase apparent screen size with no loss of visual quality is not everything, but it is not nothing, either. The apparent size of a screen in our viewing area is a key factor in what is generally called viewing *immersion*; indeed, the illusions of virtual reality are created largely by covering one's entire visual field with a screen.

On this basic analysis, then, the advantage of UHD over HD is primarily its ability to create a more immersive viewing experience, by allowing the viewer to get closer to screens of the same size, and view larger screens at the same distance, with no loss in visual quality. This is presumably a good thing, at least when we want to be more immersed in what we are viewing. But, like many good things, UHD has its own trade-offs.

1.8 The Problem with Digital Video

The most obvious trade-off for UHD is simply the cost quadrupling the number of pixels per video frame, from about 2 million to about 8 million. As a viewer, you might think that doesn't matter, as long as advancing display technology makes new 8-million pixel UHD screens available in the same price range formerly paid for comparable 2-million pixel HD screens. Like an iceberg, however, the implications of multiplying pixels run far deeper than the visible surface of a UHD screen.

Digitally speaking, every pixel is a number, specifically a binary number that represents a specific color shade. For each pixel, the display reads its number, and generates the colored block appropriate for that number in the location appropriate for that pixel in a size appropriate to the resolution format for a display of the given dimensions.

The pixel numbering standard in common use today for broadcast television is so-called "8-bit" color, which generates a 24-bit 3-channel RGB binary number for each pixel, sufficient to enable a total palette of over 16 million colors. Since 16 million is more color shades than even the most discerning human eye can distinguish, 8-bit color (24 bits/pixel) is sometimes called “true” color, as the first and simplest digital color scheme to enable everything the human eye can see (and more).[2]

The problem created by digital imagery in general, and HD and UHD television in particular, isn’t that digital technology is inferior to older analog technology, or that it is inadequate to express the full range of our senses. It is simply that digital technology able to provide a high-quality experience takes a lot of bits, and improvements in quality take even more bits.

Specifically, an HD picture composed of 2 million pixels, each corresponding to a 24-bit number, requires 48 million bits to express.[3] And that is just for a single frame. Full HD plays out at 30 frames a second, meaning a total bit rate of nearly 1.5 billion bits every second. This is not just a large number; it is an overwhelming number. It is impractical to store 1.5 billion bits for every second of HD video captured, let alone transmit bits at that rate.

2 Importance of Video Compression

Fortunately, there is a powerful remedy for the proliferation of bits required by digital rendering technology, namely digital compression technology. Compression technology is especially powerful for video, where standards like H.264 allow the elimination of 299 bits out of every 300, reducing 1.5 billion bits a second to a much more manageable 5 million bits a second (or even lower, in some cases). But what happens to data rates when the television industry shifts from HD to still higher UHD resolution standards?

2.1 How Data Compression Works

Broadly speaking, digital data compression is simply a process of eliminating information in inverse order of importance. More exactly, digital compression reduces to two logical tasks. The first and most straightforward task is simply to eliminate everything redundant in the data being transmitted, i.e., to find the smallest number of bits that can be used to encode any given amount of data with no loss of information. Technically, this is called *entropy* encoding.

The second task begins with a ranking exercise, where the non-redundant or meaningful data is sorted in order of its interest or importance. Data is then eliminated starting at the bottom, with the least significant (most uninteresting), and working up through the ranks of data in order of increasing significance. This second type of compression stops either when a set size target or a specified level of importance (quality standard) is reached.

The first type of compression, which retains all information and eliminates only redundancies, is known as “lossless” compression; the second type of compression, which also eliminates the least important or most uninteresting parts of the original data, is known as “lossy” compression. For images, lossless compression alone is generally inadequate; i.e.,

[2]For further discussion of digital color coding, see Appendix 1: Binary Color Coding.

[3] 48 million is the “round number” approximation of 2 million times 24. Multiplying out the actual HD numbers (2,073,600 pixels/frame x 24 bits/pixel) yields an actual answer much closer to 50 million (49,766,400).

unable to achieve the kinds of very substantial reductions needed to cope with the overwhelming number of bits generated by digital imaging technology. For this reason, all standard image compression technologies (JPEG, MPEG, etc.) are lossy.

The good news for video is that the first lossy reductions are imperceptible. Image sensors capture color information the human eye is not able to discriminate, so the first step in video compression is simply to eliminate what we cannot discern. If the compression process stops here, there is absolutely no degradation of visual quality. Everything the human eye can actually see is still present in the remaining data.

If still more compression is needed, the next step is to eliminate what can be seen with perfect vision, but whose absence most people will not notice. From there, the process proceeds to losses that may be more widely seen, but will be generally viewed as insignificant, and so on.

Indeed, the rest of the story of declining bitrates should be familiar to anyone who has watched much streaming video. As the information being discarded becomes more and more important, the picture begins to soften up. Eventually, the edges of objects blur and run together. Finally, the objects themselves become unrecognizable, dissolving first into a mosaic of colored blocks, then into nothing at all.

2.2 The Limits of Video Compression

The point to this story of progressively degrading image quality with increasing levels of compression is simply that there are limits to what even the best possible compression technology can achieve without significant degradation. Great video compression technology, like H.264, may be able to eliminate 299 out of every 300 bits in the original image, with little or no perceptible loss in visual quality. Even better compression technology might be able to eliminate still more bits, leaving only 1 in 400 or 500. But no compression technology can eliminate all the bits, and no compression technology can eliminate bits from the hard-residual core of important information (however large or small it may be) without significantly degrading the quality of what remains.

There is an important moral here about what can be expected from the next step in video compression technology: the move from H.264 (MPEG-4 or AVC, for Advanced Video Coding) to H.265 (MPEG-5 or HEVC, for High Efficiency Video Coding). We will return to this issue later.

3 UHD Data Rates

3.1 The Basic Calculation

For now, however, the point is simply that, since UHD has 4X the pixel count of HD, and pixels are nothing but 24-bit bundles of color data, the starting point for any discussion of upgrading HD to UHD *without* degrading picture quality (cutting more of the *significant* information from the bit stream) is how to move 4 times as many bits across our broadcast networks. 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). [4]

In truth, shifting from HD to 4K UHD at the cost of moving just 4X the amount of information would be the good news. The bad news is that the amount of additional data to be moved for 4K UHD is likely to be far higher.

3.2 The Complete Calculation

The data challenge posed by 4K UHD is worse than a mere 4X increase in data rates, because the UHD standard encompasses more than just doubling and redoubling pixel count. The UHD standard also allows a shift from 8-bit to 10-bit color[5], 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 to 10-bit color channels increases pixel length from 3 x 8 = 24 bits to 3 x 10 = 30 bits. Which is to say, the new UHD 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 10-bit 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. Rather, that is merely its beginning. In addition, the UHD standard makes each pixel 25% longer, to carry new color information, and also

---

[4] In brief, image compression is very much like floating a boat that threatens to sink in heavy seas under the weight of too much cargo. You begin by throwing overboard everything extra, and follow by throwing overboard everything dispensable, until a critical core of functionality is reached past which the boat ceases to be viable. Nothing in this process of lightening the load, however, changes the size of the boat. If you leave port aboard the UHD supertanker, you must arrive aboard the UHD supertanker, and not on some trim little yacht.

[5] For further discussion of 8-bit vs. 10-bit color, see Appendix 1: Binary Color Coding.

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 data 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.

4 Who Benefits from UHD

4.1 Display Manufacturers Get a New Market

Clearly, display manufacturers are the group that is most enthusiastic about the new "4K" UHD standard. Given the ability to manufacture a screen with 8 million pixels, each able to display a broader range of "10-bit" HDR[6] colors at 60fps, the advantage to embracing the new UHD standard for this group is obvious. Just as HD converted a mature, no-growth market for SD screens into a high-growth market for HD screens a decade ago, 4K UHD now has the same promise to revitalize an increasingly stagnant market for HD displays.

Of course, the responsibility borne by display manufacturers for the 4K standard ends with providing a 4K UHD-capable screen. Is there any content available able to take advantage of 4X higher resolution, 2X faster frame rates, and a 1.25X increase in color data? For display makers, at least, content delivery is not a problem—or, at least, not beyond the need to help foster the conviction that UHD content is coming, as necessary to encourage screen sales. The importance of UHD to a display manufacturer is not that UHD, when fully implemented, helps create a more immersive viewing experience. Rather, it is simply that the new standard makes all HD TVs obsolete, hopefully precipitating a whole new buying cycle for UHD TVs.

There is a cautionary lesson here. Once the 2K-to-4K transition is largely over, the incentive for display manufacturers is to begin promoting whatever new standard advancing display technology will then make possible. From this viewpoint, the good news is that a new standard is already on the books: the "8K" UHD resolution standard. 8K is a doubling of 4K L X W numbers, to arrays of 7680 x 4320 pixels, or roughly 8K x 4K = 32 million pixels. This quadruples 4K UHD resolution, in exactly the same way 4K quadruples 2K HD resolution.

Of course, the lurking 8K UHD standard does not stop with quadrupling the pixel count of 4K UHD, any more than 4K stopped with quadrupling the pixel count of 2K HD. It also pushes

[6]For further discussion of HDR (High Dynamic Range) color, see Appendix 1: Binary Color Coding.

on from 10-bit to 12-bit color channels (36-bit pixels), and allows another doubling of the frame rate, to 120 fps. As a result, the increase over full 4K data rates for full 8K is another 3-term multiplication problem: 4 x 1.20 x 2 = 9.6X. Piled on top of the 10X increase required by full 4K, full 8K, then, represents an increase of nearly 100X the data rate of full HD.

But even a 100X increase in current HD data rates is not a problem for display companies, since their business stops at selling the display. Figuring out how to transmit two orders of magnitude more data to new 8K screens is someone else's problem. As long as technology supports the manufacture of still higher resolution displays, and customers looking to buy a new TV set believe a 4K UHD screen is better than a 2K HD screen, and an 8K UHD screen is better than a 4K UHD screen—and can be persuaded to buy the latest resolution standard—there is no reason for display manufacturers to want the drive towards a new, higher resolution standard every decade or so to ever stop. ("16K" anyone?)

4.2 OTT Providers Get a New Marketing Tool

From display manufacturers, let's turn to a group that's actually involved with 4K content: so-called "Over The Top" (OTT) video providers, like Netflix, Amazon, Hulu, You Tube, etc. OTT companies are in the business of acquiring video content, which they then make available over the Internet, either freely or to paid subscribers (according to their business model). Netflix has also led an OTT movement into content creation, with original series like "House of Cards".

Netflix, in particular, has been active in promoting UHD, even pledging to create all its new original content to the 4K standard. More, it has raised its own 4K ante, by producing the 3rd season of House of Cards in an upscale "6K" format (6144 x 3160 pixels). This is not a trivial investment. Not only are 6K cameras expensive, but post-production editing costs also go up considerably with higher resolutions. In addition, the resulting source files are huge—the 6K master copy for a 55-minute House of Cards episode is said to be a whopping 5.5 TB.

Still, in the overall budget for a high-profile TV series—relative to salaries, location costs, marketing, etc.—the added production costs for 4K and higher formats are doubtlessly a relatively minor line item. And, in an era when a TB of storage can be had for less than $30, even the cost of storing a TB for every 10 minutes of video is trivial. OTT companies can doubtless write off the entire added cost of UHD production as a line item under marketing expenses, in an effort to gain competitive advantage for their program lineup.

However, the real nub of the crises created by new UHD resolution standards is not the cost of producing, editing, or storing UHD video, but rather the challenge of being able to actually deliver UHD programs, like House of Cards, in a true 4K format. And, here, Netflix is on solid ground, because (like other OTT companies) they don't deliver video.

It may be helpful to keep this point in mind the next time you hear or read glowing comments about the imminent advent of UHD video from Netflix or other OTT providers. Certainly, they are ready to acquire and store UHD content. They may even be eager to produce their own original programs in UHD formats. But their responsibility for UHD ends with embracing the new technology, for whatever marketing advantage may be conferred by offering more and better UHD content.

Of course, OTT companies must also provide the video they offer. Netflix does this either via mail, in a stored media format (DVD, Blu-Ray, Ultra Blu-Ray), or by streaming it over the Internet.[7] But OTT providers are not responsible for anything to do with either of these delivery ecosystems.

Thus, the Blu-ray Disc Association has a long list of contributors to the Blu-ray standard, but this list does not include Netflix. In other words, since companies interested in stored media have developed Ultra Blu-ray for 4K content, Netflix is happy to rent these disks. But, if no stored disk format is developed for, say, 8K UHD resolutions, well, that's not Netflix's call.

Much the same can be said of Internet delivery. True, Netflix is easily the largest consumer of Internet bandwidth and, together with fellow OTT traveler, You Tube (a distant second in bandwidth consumption), now accounts for an astonishing share of over one-half of all Internet traffic during prime loading periods. But OTT companies (unlike Al Gore) do not claim any credit for creating the Internet, and (like Al Gore) they take no responsibility for either maintaining or upgrading it.

If the Internet bogs down and OTT streaming stops thereby, well, your OTT provider sincerely regrets the interruption in your service, and suggests you try again later. And, if that doesn't fix the problem, you are, of course, free to call your local Internet service provider to complain.

4.3 Real Time Data Delivery is the Real Crises

In summary, for display manufacturers, UHD is pure opportunity. For forward-thinking video providers, UHD is the next step in the evolution of video technology; a step they may not only be anxious to prove they are ready to take, but also eager to been seen as actively leading. The challenge of providing UHD content on stored media has already been solved for 4K UHD by the Ultra Blu Ray optical disk format, and presumably can be solved for 8K UHD as well, should that format supersede 4K UHD. The real crisis posed by UHD only appears when we turn to the task of actually moving the massively increased amounts of data generated by UHD programs to UHD receivers in real time across remote networks.

5 How to Deliver UHD Data in Real Time

The height and breadth of the real-time roadblock preventing the dawn of the rosy 4K future being promoted today by TV set manufacturers and OTT content providers can be glimpsed by returning to a figure quoted earlier in this series: Netflix and You Tube together account for over half of all peak Internet traffic today. Replacing even 10% of that traffic

[7]For further discussion of Internet streaming, see Appendix 2: Adaptive Bitrate Streaming.

with 10X more expansive UHD data would generate *more traffic than the entire Internet handles today*. Leaving no capacity for anyone not on Netflix and You Tube to do anything at all, not to mention no room for the remaining 90% of the Netflix/You Tube audience to move up to UHD.

5.1 HD to SD No Precedent for HD to UHD

If that reflection is not sufficiently sobering, consider our one and only historical precedent for a transition from HD to UHD, namely, the transition from SD to HD. Serious research into new HD formats for commercial television started back in the 1960s, with attempts to demonstrate viable HD systems beginning in the 1970s and continuing through the 1980s. The pertinent point for this discussion is that all these efforts, based on established over-air analog technologies, ultimately failed for one and the same reason: there was no practically feasible and generally available way to move four to six times as much data for a new HD format across existing SD channels.

This 30-year impasse was finally resolved in the 1990s, by the development of new digital technology. Critical for HD was the fact that, among other advantages, digital imaging allows the use of data compression. Of course, data compression did not actually remove the bandwidth bottleneck that had stifled earlier HD initiatives; rather, it choked the data that needed to be transmitted down to a size that would fit through new digital pipes. Indeed, data compression was so successful at reducing bandwidth needs that, during 2005-2010, the FCC actually reassigned bandwidth, formerly allocated to broadcasters, to Sprint/Nextel for cellular use.[8]

The bad news for the HD to UHD transition is that the analog to digital conversion is a one-trick pony. There is no similar technical legerdemain now waiting in the wings to enable an upgrade to UHD. In today's all-digital world, the only relevant difference between HD and UHD is the difference between some bits and an order of magnitude more bits.

[8]The so-called "2 GHz relocation" of 7 Broadcast Auxiliary Service (BAS) channels. These were 17/18-MHz wide analog channels, located between 1990 and 2110 MHz. This radio-frequency band was allocated to broadcasters for internal communication between studios and outside locations (aka "backhaul" channels, to fixed transmitter sites, fixed field cameras, or mobile field reporters). The Sprint/Nextel initiative converted these 7 analog channels to 7 narrower 12-MHz wide digital channels, located in the upper 70% of the old analog band, between 2025.5 and 2109.5 MHz. The excluded frequencies, in the lower 30% of the band between 1990 and 2025, were then reassigned to Sprint/Nextel, to expand its adjacent Personal Communications Service (PCS) band (which ran between 1850-1990 MHz). As incentive for surrendering 30% of their analog BAS bandwidth, Sprint/Nextel bought new digital equipment for all the broadcasters who, to continue using auxiliary services, were forced to convert to the reassigned 12 MHz digital channels.

The power of digital compression can be seen in the fact that the old analog BAS channels, with 17-18 MHz of bandwidth, were unable to accommodate even a single HD channel; while, using advanced MPEG-4 compression (H.264/AVC) over the narrower 12 MHz digital band, broadcasters were easily able to accommodate 2 HD channels at the same time (each at a relatively robust data rate of 6 Mbps). As a consequence, in the 2005-2010 timeframe, pioneering local stations who upgraded their broadcast studios and operation centers to HD, were able to conclude an "all HD" news upgrade, also shifting their field reporting (aka "ENG" for Electronic News Gathering) to the new HD standard, using digital HD cameras, MPEG-4 encoders, and digital transmitters and receivers purchased with Sprint/Nextel money.

5.2 The HD Model, Doubled and Redoubled

The optimism that underlies UHD, then, is not a belief in some fundamental technological shift, but rather the conviction that digital technology is highly elastic, able to accommodate rising numbers of bits through many doubling cycles.

Of course, there is a familiar and very practical demonstration of the ability of technology to expand in precisely this way, namely, Moore's Law, which postulates that semiconductor technology can double transistors on a chip every couple of years. Over the course of the last 57 years, Moore's law has carried chip makers across 33 doubling cycles, from a single transistor in 1959 to over 8 billion transistors today, with no clear end yet in sight.

For digital video, the postulated elasticity of bit transmission must be founded on some combination of the abilities to contract the rising flood of bits, through improved digital compression technology, and to increase digital channel capacity, to accommodate however many bits cannot be eliminated.

5.3 Prospects for Better Data Compression

Clearly the best way to deal with the rising floodtide of bits created by UHD would be to simply eliminate it by improvements in data compression technology. In large part, this *is* the story of the digital SD to digital HD transition. The switch from MPEG-2 to the more aggressive set of compression tools in MPEG-4 successfully reduced the flood of additional bits created by HD to something closer to a trickle (where it did not eliminate it entirely).

To take just one data point, with MPEG-2 technology, CableLabs recommended minimum video streaming rates over cable of 3.75 Mbps for SD, but 15 Mbps for HD, i.e. a 400% increase to enable the new resolution standard. Using MPEG-4 technology, however, Netflix now streams full HD at a maximum rate of 5.8 Mbps, a relatively modest increase of just 55% over the 3.75 Mbps rate recommended for MPEG-2 SD. And even that modest increase can be drastically cut by using an "HD Lite" format, which Netflix streams at a minimum of just 2.35 Mbps. The bottom line here is simply that the overall impact on bitrates of the SD to HD transition was largely cancelled by better compression technology.

5.3.1 Historical Overview of Video Compression

There is, therefore, historical precedent for thinking that that a similar scenario may play out for the HD to UHD transition. It is worth pausing here to briefly review the history of digital video compression. Although useful digital data compression dates back to the very dawn of the computer era, 1951[9], digital video compression is much newer. The Moving

[9]1951 is both the year the first commercial computer, UNIVAC I, was sold to the U.S. Census Bureau, and the year an MIT doctoral student named David Huffman developed the first successful algorithm for generating prefix-free variable length codes (aka "Huffman" codes). The basic idea behind Huffman's variable length coding (VLC) scheme—namely, use the shortest binary codes for the data that appears most frequently—has since been refined, first as CAVLC (Context Adaptive Variable Length Coding), then as CABAC (Context Adaptive Binary Arithmetic Coding), but Huffman's key insight remains at the heart of lossless or "entropy" compression.

Pictures Expert Group (MPEG) was not formed until 1988, and did not issue its initial MPEG-1 standard (aimed at CIF-format video conferencing) until 1993. MPEG-2, the first standard aimed at broadcast-level video (digital "full SD" or 704 x 480/566 interlaced formats), appeared the following year.

Although MPEG-2 was successfully upgraded to handle new HD resolutions (both 720p and 1080i)—leading to abandonment of a proposed MPEG-3 standard for HD broadcasting—the resulting bitrates for high quality video tended to overwhelm transmission resources. As mentioned above, 15 Mbps was recommended as a minimum for digital cable, while the standard ATSC rate for over-the-air broadcasting, 19.4 Mbps, was set even higher. Lowering these MPEG-2 bitrates by reducing video quality tended to defeat the whole purpose of HD. Of course, this is not to say it was never done, as some cable providers scrambled to provide HD services by streaming MPEG-2 in the 10-14 Mbps range.

The saving grace for HD turned out to be the fact that it was slow to ramp. Although the HD standard itself dates to 1996, with the first HD broadcast occurring in 1998, HDTV penetration of U.S. households did not pass even the 10% mark for a decade (2007). This provided time for the MPEG committee to develop a robust set of more aggressive compression tools, organized under a new MPEG-4[10] standard. MPEG-4 is also known by acronym, as AVC (for Advanced Video Coding), as well as by its ITU-T designation, H.264. The timely arrival of MPEG-4 compression, with its ability to cut HD data rates back to near-SD levels, prevented HD programming from ever generating a widespread bandwidth bottleneck crisis for transmission of video signals, whether beamed to satellites, sent across cables, streamed on the internet, or broadcast over the air.

Are we now posed to repeat this historical success story one more time? Just as MPEG-4 compression appeared at the beginning of the HD era, so the first version of a new MPEG-5 standard was released in 2013, at the beginning of a new UHD era. Like MPEG-4, MPEG-5 is also known by acronym, as HEVC (for High Efficiency Video Coding), and by its ITU-T designation, H.265. The 10-year intervals between MPEG-2 (1994), MPEG-4 (2003, 2005), and MPEG-5 (2013) seem too precise to be just a coincidence. Presumably, as MPEG-4 arrived just in time to facilitate transition to the new HD resolution standard, so MPEG-5 is arriving just in time to facilitate transition to the new 4K UHD resolution standard.

Sadly, the short answer to the question, whether MPEG-5 can repeat the success of MPEG-4 in stemming the rising floodtide of bits caused by a new resolution standard, is "no**".**

5.3.2 Limitations of MPEG-5 (HEVC) Compression

Following the advice of the 1968 Jerry Lewis film, "Don't Raise the Bridge, Lower the River", MPEG-5 compression presumably offers a way to lessen the flood of bits created by shifting

[10]Since the MPEG-3 name had been assigned to the aborted attempt to develop a separate standard for HD, the naming of released MPEG standards skips from MPEG-2 to MPEG-4. The MPEG-4 standard was released in 2003, and received a set of major fidelity range extensions (frext) in 2005. These extensions included a new "High" profile that significantly improved compression rates over the original "Main" profile (e.g., by replacing CAVLC entropy coding with CABAC). Therefore, as HD programming started to spread widely after 2006, High profile was widely adopted by broadcasters as the preferred version of MPEG-4.

from 2K HD to 4K UHD resolution. Moreover, looking at the 10-year pattern followed by releases of new video compression standards, it might reasonably be expected that still more advanced MPEG-6 compression will arrive around 2023, presumably just in time to handle the next 10X increase in bits, created by shifting from 4K UHD to 8K UHD resolution.

Regrettably for this optimistic scenario, new compression standards are forced to follow a path of diminishing gains. The logic here is not hard to see. Unlike Moore's Law (mentioned earlier in this series as the prototype for seeming endless improvements in capacity), compression counts down rather than up. Counting up—in the case of Moore's Law, doubling transistor counts every couple of years—pushes against an elastic ceiling with no obvious limit to its expansion capabilities.

Alas, the logic of counting down is very different. Rather than an indefinitely elastic ceiling above, there is a hard floor below. Recall the two-fold strategy that underlies digital compression. First, eliminate everything redundant ("lossless" compression). Second, when everything redundant is gone, eliminate everything unimportant ("lossy" compression). When all the unimportant data (however defined) is gone, the limit of what can be done—without comprising quality—has been reached.

Moreover, the closer the approach to the fixed floor set by the pool of all and only significant data, the higher the resistance to further gains. Using relatively modest compute power, MPEG-2 takes 1500 Mbps of raw HD video down to 15 Mbps—a reduction of 99%. By throwing a lot more computational resources at the task of compression, MPEG-4 eliminates a further 67% of the remaining data, leaving a residue of just 5 Mbps.

By throwing an extravagant amount of computational energy at the task, MPEG-5 *might* be able to eliminate another 50% of what remains, reducing 5 Mbps to 2.5 Mbps. (At least, that is the goal set for MPEG-5, though current demonstrations of the technology generally settle for reductions of 20-30% rather than the postulated 50%.) MPEG-6, should it be developed, would have to push far harder still to eke out another 25-33%, going below 2 Mbps. And it is likely that no amount whatsoever of computational power could ever push HD video below 1.5 Mbps (setting the hard floor at 1 significant bit in every 1000).

Note the diminishing returns from additional compression. The great majority of what can be achieved by compression is achieved by MPEG-2, eliminating 99 out of every 100 bits. MPEG-4 manages a substantially additional gain by tossing out 2 of every 3 remaining bits. Realistically speaking, at this point, the bit stream has largely been wrung dry. Only 1 bit in 300 remains, and all of these cannot possibly be eliminated. Worse, in absolute terms, even if the 5 remaining Mbps went to zero, the savings would still be only half the 10 Mbps reduction realized in the previous step. The figure shown below illustrates this curve of sharply diminishing results for dramatically escalating efforts.

To be sure, the above numbers are merely illustrative, and lack both scientific authority and universal validity. But the basic point they make is inescapable. There is hard floor to what can be achieved with data compression, and each successive step taken in the direction of that floor will achieve less in the way of results while requiring more in the way of effort. In truth, the last really dramatic step possible with compression technology was MPEG-4.

MPEG-5, and any further video compression standards that may be developed, will not be dramatic leaps forward from MPEG-4, but rather increasingly modest gains.

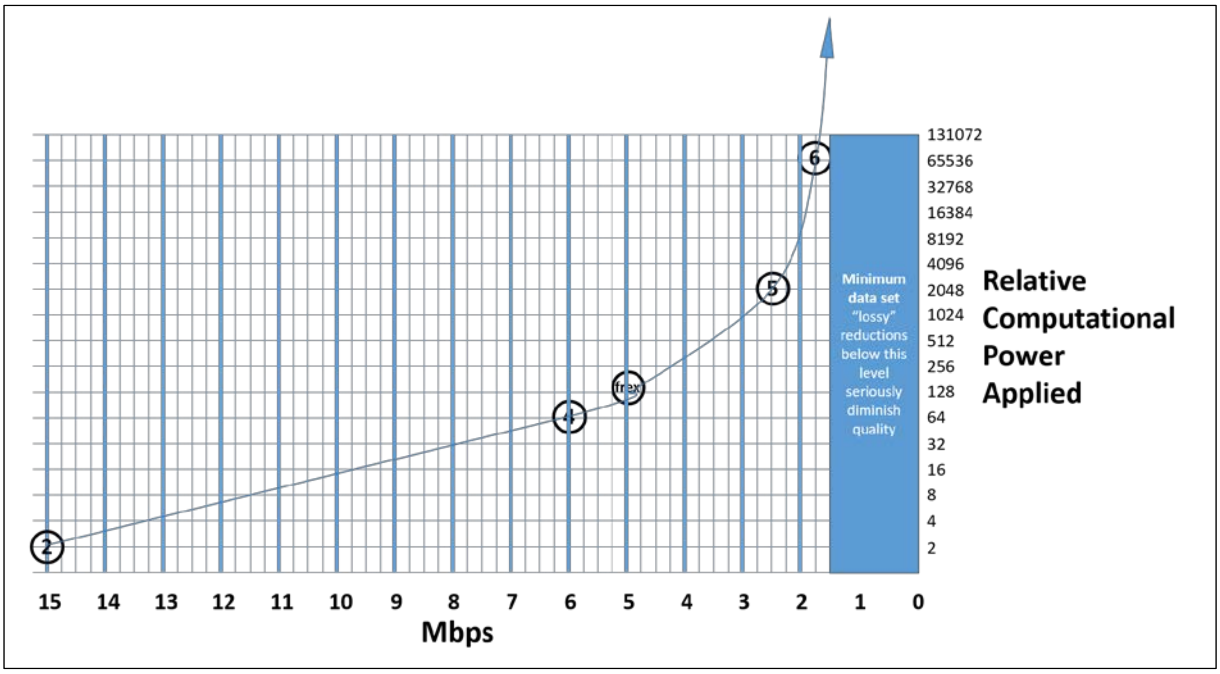

The compression curve for HD video, from MPEG-2 to a hypothetical MPEG-6. As compression approaches the floor set by the amount of non-redundant, non-trivial data (generously, if arbitrarily, fixed here at 1 bit in 1000), real gains decrease while the amount of effort required to make further progress rises exponentially. The assumption used in this graph is that successive MPEG standards appear at 10 year intervals, allowing time for 5 "Moore's Law" doublings in available computational power. This means each new standard has roughly 32X the power of the previous standard to throw at the compression problem, starting from an arbitrary level of 2 for MPEG-2.

5.3.3 4K UHD Compressed Bitrates

Given this logic of diminishing returns, what is the bottom line for 4K UHD video? As we have seen, the increase in bits required by new UHD resolution standards is a compounding problem, composed of three independent terms: more frames per second, more pixels per frame, and more bits per pixel. When fully realized, this equation multiplies out to an order of magnitude increase for 4K UHD over 2K HD, raising raw data rates from 1.5 to 15 Gbps.

Now apply the above rules for successive compression levels to a raw 15 Gbps UHD data stream. Eliminating 99 of every 100 bits with MPEG-2 technology leaves a staggering 150 Mbps. Eliminating 2 of every 3 remaining bits with MPEG-4 reduces this to a still hefty 50 Mbps. MPEG-5—depending on how close it gets to the target of eliminating 1 of every 2 bits left by MPEG-4—will be able to further thin the UHD stream down to 25-35 Mbps.

5.3.4 Modifying UHD Requirements to Lower Bitrates

Looking at the limited headroom left to new compression technologies after MPEG-4, it does not appear that better compression will be of great help in lowering the 10X flood of bits created by new 4K UHD resolutions, let alone the 100X flood needed for 8K UHD. Over time, MPEG-5 may reach its goal of eliminating 1 out of every 2 bits left by MPEG-4 compression. However, it seems a safe bet that no amount of additional compression will ever halve the MPEG-5 bitrate again.

Fortunately, there is a simpler way to reduce bandwidth requirements – just cut back on the number of bits generated by new resolution standards. The only truly fixed part of the 4K UHD standard is its quadrupled pixel count. By definition, 4K doubles the W x H dimensions of 2K HD to an array of 3840 x 2160 pixels. So much is inescapable.

In addition, there is substantial pressure to adopt the new HDR10 Media Profile, requiring an upgrade from 24 to 30-bit pixels – a further 25% increase in the data required. Multiplying out both changes yields a total increase of 4 x 1.25 = 5X the number of HD bits. The other half of the 10X increase for "full" 4K UHD comes from the requirement to run at double the 30 fps rate of HD. But, since movies have survived for over a century at 24 fps (more or less) without (much) complaint, this part of the new 4K standard can probably be ignored, assuming both the higher pixel count and heightened color range are implemented.

Let us consider, then, a reduced 4K UHD standard, beamed or streamed at 30 fps, requiring a mere 5X increase over 2K HD bit rates. This version of 4K raises the HD raw data rate of 1.5 Gbps to (5 x 1.5) = 7.5 Gbps. Running this 4K figure through the compression ratios set out in the previous part of this series means MPEG-2 (100:1) reduces it to 75 Mbps, MPEG-4 (3:1) to 25 Mbps, and MPEG-5 (2:1 to 3:2) to somewhere in the 12.5 to 16.7 Mbps range.

Looking at these figures, let's use 18 Mbps as a safe target data rate for 4K UHD. 4K might be done in less, but a reasonably high quality version (excluding the doubled frame rate) should be doable within this envelope. As for 8K UHD, let's reduce its required bit rate, to the maximum extent possible, by implementing only its 4X increase in pixel count. This yields a quality MPEG-5 data rate of 4 x 18 = 72 Mbps. And, by applying a theoretical MPEG-6 level of compression, able to use 2023 levels of compute power to generate (say) a further compression of 4:3, this figure might come down to the neighborhood of 55 Mbps.

Thus, by shaving off parts of the new UHD standards, then applying maximum compression pressure to what remains, we have lowered the river of UHD bits as far as seems feasible. The rest of the gain needed to support UHD will have to come from raising the bridge.

5.4 Prospects for More Bandwidth

What are prospects for getting, first, 18 Mbps, then 55 Mbps, of bandwidth to a UHD TV? And by when? We will consider this question in two parts: wirelessly and wired. For wireless transmission, we will look at over-the-air broadcasting. For wired transmission, we will look at cabled Internet delivery.

5.4.1 Over-Air Broadcasting

As mentioned earlier, the upgrade from SD to HD resolutions was made possible by shifting TV broadcasts from analog to digital technology. The original digital standard, ATSC 1.0 (from the Advanced Television Systems Committee), was finalized in 1996, the same year the HD standard itself was approved. For purposes of the present discussion, there are two critical points about ATSC 1.0: an 8-VSB RF modulation scheme, and MPEG-2 data compression.

5.4.1.1 ATSC 1.0

8-VSB is an 8-level Vestigial SideBand scheme that, given a 6 MHz RF channel, is able to transmit 1 of 8 possible 3-bit digital codes (000 to 111) 10.76 million times a second. A little multiplication (3 x 10.76) shows this generates a raw data rate of about 32 Mbps. However, due to overhead[11], the usable amount of data from this scheme is only about 19.4 Mbps.

The selection of MPEG-2, as the technology for video data compression, is certainly unsurprising, since it was the only MPEG compression standard available for broadcast TV in 1996. Which is to say, every built-in ATSC 1.0 RF tuner produced over the past 20 years, from 1996 to 2016 (now), has included a companion MPEG-2 decoder.

Using current ATSC 1.0 technology, then, over-air stations are able to broadcast HD programs in MPEG-2 at up to 19.4 Mbps. While this exceeds the minimum 15 Mbps bitrate recommended for MPEG-2 HD by CableLabs, it provides little headroom for any higher resolution standard. In particular, an MPEG-2 4K UHD signal would require something like 75 Mbps (far in excess of even the 32 Mbps raw data rate available).

Clearly, what over-air 4K programming needs is a new version of ATSC technology. Preferably, the new standard will upgrade both of the critical bandwidth features listed above: first, a new RF modulation scheme, to maximize the number of bits that can be pumped over a 6 MHz RF channel[12]; and second, a new compression standard, to minimize the number of bits that need to be pumped.

5.4.1.2 ATSC 3.0

The good news is that this new standard, ATSC 3.0[13], is already far along in committee. A final version should be approved sometime in 2017. To be sure, judging by the history of ATSC 1.0—approved in 1996, but not fully implemented until 2011—the prospect of ATSC 3.0 coming to a TV set near you is still some distance off. Setting that issue aside for the moment, however, how much bandwidth improvement does ATSC 3.0 provide?

---

[11]The bulk of the hefty 40% transmission overhead penalty in ATSC 1.0 consists of FEC (Forward Error Correction) codes, which allow receivers to check the bits they receive, determine if any were somehow flipped in transmission, and, if so, correct them.

[12]Of course, the one thing ATSC 3.0 can't do is increase the 6 MHz RF channel allocated to broadcasters. For better or worse, the radio frequency spectrum is a limited natural resource—a fixed pie that, with the rapid growth of wireless devices of all descriptions, is now beset on all sides by demands for larger slices.

Indeed, as mentioned earlier in this series, with the "2GHz relocation", broadcasters actually surrendered RF spectrum to feed Sprint/Nextel's growing demand for PCS bandwidth. In the face of increasing competition for limited MHz, doubtless the best case scenario for broadcasters is just to hang onto the 6 MHz over-air channel they now have. As a matter of practical reality, keeping all of this channel for their own use may require over-air broadcasters to upgrade from badly outdated 1996 ATSC 1.0 technology sooner rather than later, simply to show that no one else can put their slice of RF spectrum to better use.

[13]ATSC 2.0, like MPEG-3, was a standard begun but soon overtaken by events, and eventually abandoned without issue.

First, ATSC 3.0 replaces 8-VSB with a more aggressive scheme for bit multiplexing: OFDM (Orthogonal Frequency-Division Multiplexing). This almost triples the nominal bit rate ceiling, from 32 Mbps to 90 Mbps. Of course, overhead will eat up some of this nominal 90 Mbps bitrate. But even a 40% overhead tax on 90 Mbps would still leave broadcasters with a real data rate of around 55 Mbps.

Second, ATSC 3.0 replaces the 1994 MPEG-2 compression standard with the latest 2013 MPEG-5 compression standard. The good news is that this change makes not only a 30 fps 4K UHD stream possible within the old 8-VSB transmission envelope (at 18 Mbps), but easily enables 60 fps 4K UHD within the new OFDM envelope (at 36 Mbps).

The bad news is that, unless real OFDM bitrates can be gotten up into the 70 Mbps range (not at all clear), ATSC 3.0 is unlikely to handle even a minimal version of a future 8K UHD upgrade. The good news is that 8K UHD remains a problem for the future rather than any sort of near-term concern.

Regrettably, much the same can be said of ATSC 3.0 itself. Not only is the standard not yet final as of the end of 2016, but serious practical issues render any timeline for its future deployment highly uncertain.

Chief among the practical obstacles ATSC 3.0 must overcome is its incompatibility with ATSC 1.0. On the transmission side, this means over-air broadcasters will need to make a substantial investment in new modulation, compression, and (to some extent) transmission, editing, control and other equipment before they can deploy ATSC 3.0. On the receiving side, this means none of the 300M or so TV sets sold in the US since the digital transition began (or many yet to be sold) will be able to receive an ATSC 3.0 signal. In other words, at whatever future cut-over date is set for the transition to ATSC 3.0, every set that worked up to that moment with 1.0 will immediately go dark—unless, by that time, all TVs have dual mode capability, able to run with either sort of signal (8-VSB MPEG-2 or OFDM MPEG-5).

The bottom line here is that, while ATSC 3.0 provides a clear technical solution for broadcasting 4K UHD over air, it is less clear that it will handle even the minimal demands of a potential future 8K UHD upgrade. Moreover, the incompatibility of ATSC 3.0 with ATSC 1.0 leaves not just its timeline, but even the practicality of deploying it, shrouded in mystery.

5.4.2 Internet Streaming

Unlike over-air bandwidth, which changes in a stepped way as broadcasters switch from one standard to the next, Internet capacity changes in a gradual and fairly regular way as broadband providers make ongoing investments in new and upgraded infrastructure, in response to a steadily rising demand for broadband services. As a result, the Internet has more capacity this year than it did last year, and will have still more next year—and every following year, for some indeterminate (but presumably lengthy) period.

This increase in average capacity is reflected in the periodic rise of advertised broadband speeds promoted by the various carriers. Currently, for example, Comcast's least expensive plan promises to deliver an "Internet Download Speed" of "up to 25 Mbps". Or, for a higher monthly charge, customers can purchase a download speed of “up to 50 Mbps”. Hence, it appears OTT delivery is already past the 18Mbps needed for minimal MPEG-5 4K UHD video, even in Comcast’s lowest service tier, and well past that point in their higher service tiers. So, does 4K/8K UHD video actually pose any problems for OTT Internet broadband delivery? And, if so, what are they?

5.4.2.1 Internet Speeds and Their Rate of Change

The fact that Internet broadband speeds are advertised as “up to” some number (for example, 25 or 50Mbps) is a bad start on knowing what Internet speeds really are. The number listed by a carrier typically is their best-case scenario. But, in a world where the only sure thing often seems to be Murphy’s Law[14], this “marketing” number is unlikely ever to be realized. Formerly, the rule of thumb was simply to divide the listed number by 2, in hopes of getting something like half the advertised rate.

However, the most recent annual report from the FCC (“Measuring Broadband America”, December, 2015) indicates broadband speeds are up significantly in the past few years. According to the FCC, broadband speeds nearly doubled from 2014 to 2015, with 90% of the customers of the top carriers receiving at least 95% of their advertised rate.

Akamai’s “State of the Internet” report for Q2 2016 is also upbeat, if more restrained. Over the past year, it reports peak connections speeds in the US have increased by about one-third, while average connection speeds are up better than one-quarter (28%).

The actual numbers reported by Akamai, though, are relatively modest: a current average peak speed of nearly 70Mbps, but an average actual speed of just over 15Mbps. Going by average speed (since peak speed, by definition, is both rare and fleeting) this seems to indicate that, whatever “up to” tiers broadband providers currently promote, in practice, for the majority of users, at the current rate of improvement, a real speed of even 25Mbps is still at least two years distant.

But, even supposing an average actual broadband speed gain of about 25% a year for the foreseeable future (rounding down Akamai’s reported 28% average gain for the last year), the result is still encouraging for broadband 4K UHD delivery. While the current average speed of 15Mbps does not quite reach the 18Mbps threshold set above for minimal 4K UHD delivery, in another year, with another 25% gain, it should be past that mark. And, by 2020, for most broadband users, 18Mbps should appear as nothing more than a steadily receding point in time’s rear-view mirror.

Moreover, since the current 15Mbps average must be composed of numbers that fall as much above that mark as below it, some broadband subscribers must be enjoying 4K UHD speeds right now. So, a minority audience is already in place for 4K UHD over broadband. This group will only grow larger year-by-year until, by 2020, it constitutes a substantial majority of viewers.

[14]Succinctly stated, Murphy’s Law is the principle that, “If anything can go wrong, it will”. Which is to say, in the real world, it is generally advisable to plan for the worst case scenario rather than the best case scenario.

To be sure, there are many qualifications to this optimistic conclusion about the practicality of delivering 4K UHD over the Internet, both now and in the near future.

5.4.2.2 Practical Internet Data Delivery

Projecting forward the 2016 Akami number of 15Mbps, growing at a rate of about 25% a year, brings average speeds up to 18Mbps in another year (the minimum threshold we set for quality 4K UHD), and promises to double that rate by 2020. If we simply stopped the discussion of broadband video with these numbers, the prospect of streaming 4K UHD over the Internet, if not immediately, at least in the next few years, could hardly look brighter.

But, of course, any claim about Internet speeds is a perfect illustration of the maxim that there are lies, damn lies, and statistics. An average connection speed of 18Mbps in 2017, or even an average speed of 36Mbps in 2020, may turn out to be little or no bandwidth at all in any given case.

The most obvious qualification about Internet connection speeds is that these rates are not specific to either a person or a device, but rather the connection bandwidth for a *subscriber*, which is typically a whole household of people and devices. As of the most recent Census (2010), the average U.S. household is 2.58 people and (according to Nielsen) contains about 3 TV sets. Let's assume these figures remain roughly accurate today.

Then, if all 2.58 people in an average household gather around just one-third of their TVs, and every other broadband appliance in the house is off, that TV may, in fact, enjoy all 18Mbps that will be that subscriber's average share of the Internet in 2017. But, if they split up into 3 parties, and each 0.86 person takes 1 of the 3 available TVs for his or her own use, then that 18Mbps subscriber bandwidth divides likewise into 3 streams that average 6Mbps each. Obviously, this leaves no one in the household with anything approaching 4K UHD bandwidth. And, even this one-third share of subscriber bandwidth assumes no other data-hungry appliance—smartphone, tablet, computer, game machine, etc.—is active at the same time, demanding its own share of bandwidth.

Nor does Internet sharing stop at the walls of a household. You may be alone at home with only one TV on. But if, when you attempt to stream a 4K UHD version of "House of Cards", many of your fellow subscribers also start streaming with similar ambitions, while, at the same time, other subscribers on your local feed are attempting to download TB files, the result is likely to be, at that moment in time, that no one on this shared feed will be able to get their fair share of "average" bandwidth.

Moreover, even if there are no local bandwidth issues, either in your household on your shared feed, there is still the risk of catastrophic failure elsewhere on the Internet. For example, back on the ranch, your OTT video server of choice may be overwhelmed by a sudden surge in demand or (what is much the same thing) a "denial of service" attack. Or your video stream—somewhere between issuing intact from a server and arriving, still intact, at your TV—might fall victim to any of the numerous other ills that can afflict Internet traffic (e.g., failing switches, lost packets, unresponsive name servers).

All of these public network problems—shared usage, bottlenecks, overloads, failing servers, bad switches, cyberattacks, etc.—concern what we might term the physical reality of a global data network. But subscribers also must contend with another class of issues, namely, data throttling and/or data caps imposed by the service providers themselves. Yes, the same providers that promise speeds of "up to" so-many Mbps in their ads.

5.4.2.3 Economic Limits on Data Streaming

The issue of data caps and data throttling by service providers are not so much real physical problems with Internet data delivery as economic limits placed on these data services by vendors to manage their costs and provide all subscribers with a fairly uniform level of service. For example, some carriers are currently rolling out plans that carry a 1 TB monthly data cap, arguing this is a generous limit.

Of course, like the claim that a basketball player is "short, the term "generous" is relative to context. Compared to a cellular plan with a 10 GB monthly limit, a limit of 1 TB a month is generous. Compared to streaming an uncompressed 6K UHD file, it is 10 minutes.

5.4.2.4 Streaming 4K UHD Over the Internet

For 4K UHD video, 1 TB/month is somewhere in the middle. A streaming rate of 18 Mbps is 18 MB every 8 seconds, 135 MB in a minute, 8.1 GB in an hour. The 1 TB cap, then, allows about 123.5 hours of minimal 4K UHD video/month. According to Nielsen, the average person watches 5 hours of TV a day. Assuming, then, that 1 TB plan data is not being frittered away on other uses (emailing photos, web browsing, downloading music, etc.), a 1TB cap allows 1 person in a household to watch 4K UHD video pretty much all they want—for the first 25 days of each month.

5.5 Over-Air vs. Over-Top Delivery

Comparing Over-Air Broadcasting (OAB) with Over-The-Top (OTT) Internet broadband delivery reveals contrasting scenarios of good and bad news.

5.5.1 Advantages and Disadvantages of OAB

For OAB, the good news is the availability of a dedicated 6 MHz channel. The bad news is that's all there is and (in the face of escalating demand for RF bandwidth) all there is ever likely to be. Of course, as ATSC 3.0 proves, improvements to technology over time make it possible to do a lot more with the same fixed resources. However, the rest of the bad news for the long-term future of OAB is that all technological improvements are subject to a law of diminishing returns.

For video compression (as discussed earlier in this series), after everything redundant, irrelevant, and insignificant has been discarded, a hard floor is reached where all and only significant data remains. Past this point, every additional bit removed by continued compression must come at the price of degraded quality. Much the same can be said for data multiplexing over a fixed-width channel. The width of the pipe and the speed of light provide hard constraints on how much data can be moved through it in a given unit of time; and, somewhere before that limit is reached, the cost of cramming additional bits down the pipe must outweigh any benefit conferred by the incremental gain in available data.

ATSC 3.0 is not yet at the ultimate limit of possible improvements to compression or multiplexing, but it is far closer to those limits than was ATSC 1.0. It is certainly close enough to what is possible that no significant further advances are now visible, even on a relatively distant horizon. ATSC 3.0 will have to carry over-air broadcasting to and, very likely, well past 2050. Which is to say, if some form of 8K UHD can be squeezed into the ATSC 3.0 envelope, over-air broadcasting will be able to upgrade from 4K to 8K UHD. Or else, not.

In summary, 1996 ATSC 1.0 technology was sized to deliver HD video, and has to be upgraded to ATSC 3.0 to accommodate the much higher bandwidth demands of new 4K UHD. The major issues with this transition are the incompatibility of 3.0 with 1.0 (and the consequent haziness of the timeline for 3.0 deployment), and the still vague upper limits of the 3.0 technology.

These qualifications notwithstanding, it seems likely that within the next 5 years or so, ATSC 3.0 will provide a reliable over-air mechanism sized to handle 4K UHD. The future of 8K UHD is less certain, but it is at least possible that it, too, can be handled within the envelope of improving Internet and ATSC technology, at some still more distant time.

But even if this technical optimism about ATSC 3.0 is well founded, there remain serious practical considerations for broadcasters, many of whom are still in the process of fully upgrading from SD to HD. To what extent will 4K UHD require ripping out all their lately acquired HD infrastructure? Enthusiasm for this prospect, likely low to begin with, is certain to be depressed still further by the reflection that the reward for successfully negotiating the 2K-to-4K transition may be nothing more than the opportunity to do it all over again, to support a potential 4K-to-8K transition.

5.5.2 Advantages and Disadvantages of OTT

Turning to OTT broadband delivery, the good news is that, although broadband is subject to the same limits on compression technology that apply everywhere else, there is no hard constraint on broadband pipe size. The size of a wired pipe can always be increased, e.g., by shifting from copper to fiber optic cabling, or by adding more fibers to optical cables.

For broadband, the bad news is that the Internet is not a dedicated resource, but one that is shared, and (by definition of www) shared world-wide. With the Internet, the real constraint on bandwidth is not technology but economics. Practically speaking, a shared resource cannot be built out to the prohibitively expensive standard of maximum possible demand. Rather, it must be sized to some more affordable standard, like “peak” loading.

Peak loading is simply an estimate of worst case (highest) actual demand. Statistically speaking, even under the most favorable conditions, this sort of calculation is certain to be wrong some of the time and, under less favorable conditions, may be wrong most of the time. But, regardless of how good or bad peak calculations are, they are only applicable in constrained circumstances. Thus, Amazon or Netflix may understand the demand for their specific services in a specific geography, and design capacity accordingly. But the Internet as a whole is too large and various for meaningful peak loading estimates.

Instead, we must make do with average available bandwidth, i.e., the bandwidth we actually have after accounting for both real physical issues and imposed economic limits. Some of the time, perhaps even most of the time, the Internet supplies us with all the bandwidth we really need (if not all we might desire). By Murphy's Law, however, just at what always seems the most inopportune time, it becomes maddeningly slow.

Indeed, if there were any major shift toward UHD video on the Internet—on the part not merely of major Video on Demand (VoD) OTT vendors like Netflix, Amazon, and Hulu, but by all the commerical sites in the business of streaming video, and also by corporate sites, by personal sites, and so on and on—the Internet would soon stop working for anyone. The same would happen if a substantial fraction of the TV audience that now receives service either free over the air (OTA, 17%) or via traditional paid cable/satellite subscription (75%), were suddenly to join the still small minority of mostly young Internet "cable cutters" (8%).

To repeat, the Internet is a perfect illustration of the claim that there are lies, damn lies, and statistics. Statistics about Internet usage are just that, statistics. Your numbers can and will vary widely, not only from year to year as the Internet grows, but from day to day, hour to hour, and sometimes minute to minute.

Still, there is a qualifiedly optimistic moral here. Under some circumstances, for some of the people, some of the time, the Internet either now is or, within a few years, will be capable of delivering a minimal 4K UHD stream. Always assuming, that is, that no sudden surge of demand, major shift in viewing preferences, or act of global terrorism intervenes to shut the whole thing down for everybody pretty much all of the time.

6 Is 16K SUHD Coming … Someday?

Is that, then, the likely future of resolution technology: a continuing series of upgrades to and past 8K UHD over future decades, as better technology makes it possible, not only to do more with each passing year, but (at a 25% rate of annual improvement) an order of magnitude more in every new decade? Frankly, it seems unlikely.

We began with the observation that SD was the only TV resolution standard in North America from 1941, when the original NTSC standard was adopted, until 1996, when the HDTV standard was finalized.[15] However, although 55 years elapsed between the initial SD

[15]The PAL standard, which emerged in Europe in the 1960s and was widely adopted outside North America, did not change the fundamental SD bandwidth equation established by NTSC. While PAL has 20% higher resolution than NTSC (576 vs. 480 visible scan lines), NTSC has a 20% higher frame rate than PAL (30 vs 25 fps). From the standpoint of data requirements, these two alternatives—moving bigger PAL data sets less often vs. moving smaller NTSC data sets more often—essentially cancel, leaving the amount of data moved per second in PAL and NTSC formats much the same.

resolution standard and a second, higher bandwidth, HD resolution standard, the third and fourth higher bandwidth standards, 4K and 8K UHD, were both approved in 2012.

This history shows a dramatically shortening lifetime for new, higher resolution standards: 55 years (SD to HD) declining to 16 years (HD to UHD). The natural question raised by this progression, assuming 4K gets us though the second decade of the 21st century and 8K can outlast the 2020s, is whether some new 16K/32K super-ultra-high definition (SUHD) standard is destined to appear, say, before the end of the 2030s. And so on.

In short, is there any natural end to this progression of increasing resolutions? Or, with steadily improving technology, do pixels multiply forever, from millions a second with SD, to billions a second with HD/UHD, and ultimately to trillions, quadrillions, and more a second with future standards?

The clear answer to this question, I think, is “no”. This is not because of any physical limit in the ability to transmit bits in a second. The real point is simply, that, with digital technology, it is never about what is *possible*, because it is always possible to add more bits to any number of bits. Rather, it is about how many bits are *enough* for a given purpose. With digital technology, you simply have to quit when continuing is pointless. Enough is enough.[16]

6.1 4K Enough Resolution for a Living Room TV

The real question, then, is how much resolution is enough resolution?[17] For a TV screen, the shortest route to answering that question comes from considering screen size relative to the powers of human visual resolution. It seems reasonable that the latter capability has to be at the heart of any discussion of how many pixels are enough for TV viewing.

To take up the subject of resolution first, the key fact about human visual resolution is that it drops off in a linear way with viewing distance. Earlier in this series, we mentioned a “retina display” as any display with a resolution of 300 ppi or higher, since that is the point at which pixels become too small to be individually distinguished by the human eye. The critical qualifier left out of that earlier discussion was viewing distance. For 300 ppi, the

---

[16] For a deeper discussion of this point, see Appendix 1: Binary Color Coding.

[17]To take film restoration as one example of practical digital limits, according to John Lowry of Lowry Digital (the studio responsible for transferring classic films like North by Northwest, Gone with the Wind, Citizen Kane, Star Wars, Indiana Jones, and James Bond to digital format for DVD distribution), 4K is enough to “capture everything on the film—everything”, while higher resolutions are “pointless”. Since the amount of digital resolution needed to match traditional analog film resolution is a subject where opinions differ, an extended version of Lowry’s quoted statement follows. It should be noted that Lowry is using the terms “2K” and “4K” as these are defined by the DCI (Digital Cinema Initiative) standards—2048 x 1080 and 4096 x 2160, respectively. While the number of lines are the same, the DCI standards have roughly 7% more pixels per line than the HD/UHD standards.

“*If I’m going to restore a film, my objective is to capture everything that is on that negative, which probably has a limit somewhere in the 3- to 4K range…. If you scan it at … 2K, there’s all kinds of information on that film that you just haven’t got. If you scan at 4K, it captures everything on the film—everything…. We create a digital master … that is just as good as the original camera negative in terms of resolution and grain structure…. In all the measurements I’ve done, I’ve yet to see much information on a film right up there at the 4K level—it usually rolls off between 3 and 4K. We’ve experimented at 6K, but, frankly, it’s pointless on a standard 35mm film frame.*” (quoted in “Creating the Video Future” by Josef Krebs, Sound & Vision, November 2004, pp 110-112.)

assumption is normal reading distance, or about 12 inches—about as close as most people want to get to anything they need to visually scrutinize.

But TVs, at least in sizes typical for a living room, are not meant for close-up viewing. Average TV viewing distance in a US living room is around 9 feet. The fact that visual resolution drops off in a linear way with increasing distance means simply that, at 9 feet, we can discriminate about 1/9th as many pixels per inch as at 1 foot. Using 300 ppi as the limit of resolution at 1 foot, then 300/9 = 33 ppi is the limit of our normal TV viewing resolution at 9 feet.[18]

Armed with this figure and the knowledge than a full HD screen is an array of 1080 lines of 1920 pixels each, a little math will let us pick out an appropriate HD TV set size, where "appropriate" means the biggest screen with no fewer than 33 ppi. Buying a smaller screen would, of course, increase resolution (ppi), but the tradeoff is a bad one, since none of the gain in resolution is noticeable from a distance of 9 feet, while the loss of viewing area with a smaller screen will noticeably diminish the overall immersion of the viewing experience.

So, what is the right size when choosing an HDTV for a living room with a typical 9 foot viewing distance? The short answer is a 65" set. A TV screen with a 65" diagonal measurement has a width of about 57" or a pixel density of 1920/57 = about 34ppi, or a shade more than we can actually discriminate at 9 feet. This screen will show us every detail we could possible make out.[19]

Assuming that viewing distances cannot readily be changed when purchasing a new TV, the effect of upscaling from HD to 4K UHD (keeping the maximally useful resolution figure fixed) is simply to double the appropriate screen width. Thus, at a viewing distance of 9 feet, a 4K UHD format screen with the same 34 ppi resolution as a 65" HDTV would measure 130" diagonally, and be about 113" wide by 64" high.

---

[18]This rule of thumb (300 divided by feet) scales amazingly well. At 300 feet (100 yards, the length of a football field), the limit of human resolution is a square about 1 inch on a side.

[19]A 70" screen, with a width of about 61", would render a bit less detail (1920/61 = about 31 ppi) than the 33 ppi we could, in principle, discriminate at a distance of 9 feet. Though, in reality, there is no point in being too precise about any of the numbers regarding human vision, ppi values, screen sizes, or seating distances.

Human vision, in fact, varies widely between persons (all the way from extreme near-sightedness to extreme far sightedness). It also changes considerably over time for individuals, and is affected by a variety of widely variable environmental factors, including lighting and atmospheric conditions like haze or smoke. As a result, the ability to discriminate two dots at distance may be very different not only for different people, but for the same person at different times, and for everyone at the same time as environmental factors fluctuate.

As for seating distances, simple changes in position (e.g., sitting up vs. leaning forward or back) can easily change the associated viewing distance by a foot or more.

The bottom line is that all, except where otherwise indicated, all the base numbers in this paper should be regarded as illustrative or typical, since absolute precision about human capabilities, etc., is generally unobtainable. Likewise, unless otherwise indicated, computed figures should be regarded as approximate, since many have been worked out on the basis of starting values (often arbitrarily plucked for purposes of illustration from a range of possibilities) that are manipulated by rounded and simplified calculations. For this reason, they are unlikely to be exactly the same as the equivalent numbers found elsewhere (though, hopefully, they remain similar enough to other, independently derived figures, to be regarded as at least approximately correct).

Given the challenge of placing a screen more than 9 feet wide and 5 feet high in a typical US living room can be overcome, if the nearly 5-foot wide 65” HDTV did not provide a sufficiently immersive viewing experience, it might reasonably be hoped that a 130” 4K UHDTV, with a screen approaching 10 feet wide, would do so. And, even if it does not, doubling screen size again, with an 8K UHDTV, does not seem generally feasible—since it would require a living room with 11 foot ceilings and an unbroken wall 19 feet long.

Thus, if the question is: “How many pixels are enough for a large screen TV designed for the average US living room?”, the answer is 4K. 8K appears to move us from the realm where “enough is enough” into the realm of truly, serious overkill.

6.2 What About 8K (and Higher) Resolutions?

8K resolution allows an “optimal” viewing distance of 9 feet in front of a screen 11 feet high and 19 feet wide. These sorts of H x W dimension surpass the merely “immersive”, on the way to a truly overwhelming experience for home consumption. In addition, they seem impractical in all but exceptional circumstances, based on simple room dimensions of ceiling height and wall width—independent of any new display technologies, e.g., “wallpaper” OLED, that would allow future home TVs of any size to be unrolled directly onto a wall.

Thus, if 8K UHD falls beyond the limit of what seems feasible for home viewing, it seems safe to assume the madness of constantly escalating resolution standards must end there. We should not have to worry about 16K SUHD (15360 x 8640) TVs, with their demands for 22 foot ceilings and 40 foot walls to support a “typical” 9 foot viewing distance. Even shrinking a 16K screen by moving closer and closer to it, reducing pixel size/spacing until a “retina” viewing distance of 1 foot and its associated 300 ppi resolution is achieved, does not help (enough). It still leaves a “tablet” screen of enormous size, over 4 feet wide and nearly 2 and a half feet high, far too big for any sort of practical handheld use.

Thus, I conclude, we are safe from 16K or higher screen resolutions and, in all but niche markets, safe from any real need for 8K UHD as well. Which is not to say it is impossible for screen manufacturers and others to attempt to generate a widely-perceived need for 8K UHD. Indeed, we headed the Table of Contents for this paper with an earnest quote in favor of upscaling from 4K to 8K UHD as rapidly as possible.

Assuming, then, the future holds a rising tide of opinion in favor of 8K resolution, as the market for 4K TV sets becomes saturated—of which the quote displayed at the head of this paper is merely the harbinger—it is important to keep in mind the fact that the basic difference (ignoring color depth and frame rate) between HD and 4K, and between 4K and 8K, is simply a 4X increase in the number of pixels a screen can display: twice as many lines of twice as many pixels.

If screen size is held constant, the only way to accommodate this change is to half the distance both between rows of pixels and between pixels in each row, cutting “recommended viewing distance” for the TV set in half. If pixel spacing and its associated “recommended viewing distance” is held constant, the only way to accommodate this change is to make TV screens twice as wide and twice as high. Either way, the screen looms 4X larger in your visual field, generating a more immersive experience.

The simplest way to summarize this 4X "loom" factor of higher resolutions is by a decreasing series of fixed ratios between screen width[20] and viewing distance for each of the resolution standards. For HD, the fixed ratio for ideal viewing distance is 1.8X the screen width. For example, relating this value to the above discussion of a 65" HD screen as the ideal size for a "typical" US living room (one with a viewing distance of 9 feet), 9 feet is about 1.8X the roughly 57" width of a 65" HDTV screen. Obviously, making the screen half the size and viewing it from half the distance, or any other proportional adjustment in screen size and viewing distance, would not affect this ratio. So, the ratio of 1.8X screen width as the ideal viewing distance for HDTV holds constant for all HDTV screen sizes.

Since 4K supports the same resolution at the same distance with a screen twice as wide, the shift from HD to 4K simply halves the HD ratio number, i.e., the ideal viewing ratio number decreases to 0.9X the width of a 4K UHD screen. Continuing on from 4K to 8K would halve this ratio number again, meaning the ideal view point distance for 8K UHD is just 0.45X the screen width. This set of decreasing ratios is illustrated in the figure below.

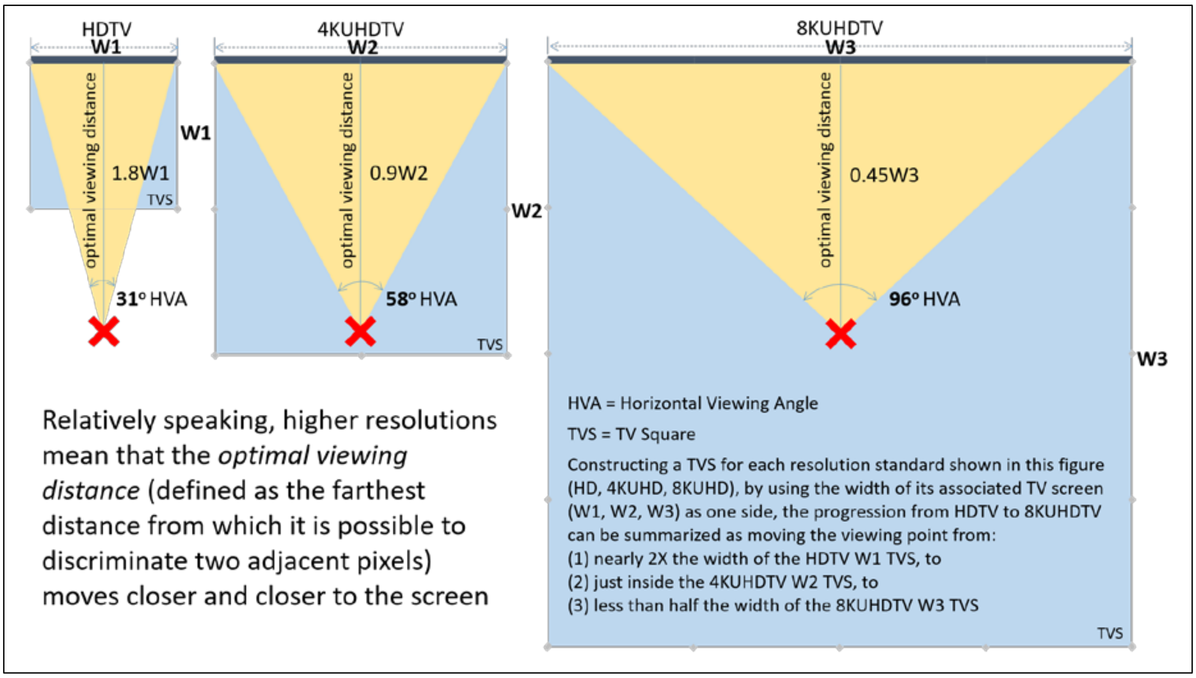

For reference, commercial cinemas built to THX specifications have a minimum horizontal viewing angle of 36 degrees (from the last row of seats) and a maximum viewing angle of 62 degrees (from the first row of seats). The "sweet spot" for theater viewing falls in the middle rows, at a horizontal viewing angle somewhere around 45-50 degrees.

[20]Note the relevant screen dimension here is width, not the commonly quoted diagonal measurement. For 16:9 HD/UHD screen ratios, width can be estimated by subtracting 13% from the diagonal dimension. 13% is a bit awkward for direct mental calculation, though it can be factored into 10% plus a short third of that number. For example, the width of a 65" screen, measured diagonally, can be guesstimated by subtracting 6.5 inches (10% of 65), plus another 2 inches (the down round of 1/3 of 6.5) = 8.5 inches, from 65. This answer, 65 – 8.5 = 56.5 inches is close enough to the true value of 56.65241 inches as to make no practical difference.

Note the jump from HDTV to 4KUHDTV takes us all the way from 31°, just outside the lowest recommended 36° "back row" angle for a commercial theater, to 58°, just under the highest recommended 62° "front row" angle for a commercial theater. Not to mention, noticeably closer to the screen than the 45-50° viewing "sweet spot". Again, the moral seems to be that 8KUHDTV, with its massive overwhelming 96° horizontal viewing angle, is a bridge too far for home use.[21]

7 Summary

Since this is a long paper that covers a wide range of topics, let's conclude by summarizing the major points made.

1. The primary purpose of a digital resolution standard is to fix pixel count, specified as a W x H array of rows and columns.
    1. HD is a pixel array 1920 x 1080, or roughly 2K x 1K. Multiplied out, this is about 2 million pixels
    2. The UHD standard has two parts.
        i. The first part, known as "4K", doubles the HD pixel array in both dimensions, to roughly 4K x 2K, or 8 million pixels.
        ii. The second part, known as "8K", doubles the "4K" pixel array in both dimensions, to roughly 8K x 4K, or 32 million pixels.
2. A digital resolution standard also specifies the number of bits that compose the data for a pixel, and the video frame rate.
    1. Quoted bit numbers (e.g., "8-bit", "10-bit", etc.) are generally for "channel length", or the number of bits used to encode each of the three primary colors that compose a pixel (Red-Green-Blue, or RGB). Channel numbers have to be multiplied by three to give the number of bits per pixel, e.g., "8-bit" HD pixels are 3 x 8 = 24 bits long.
    2. Video frame rates are specified in frames per second or fps.
    3. By multiplying (frames per second) times (pixels per frame) times (bits per pixel), a resolution standard specifies a raw data rate in bits per second. The raw data rate for HD is about 1.5 billion bits/second. The raw data rate for 4K UHD is about 15 billion bits/second. The raw data rate for 8K UHD approaches 150 (~144) billion bits/second.

[21]The extravagance of 8KUHDTV can be approached from the other end, by looking at the limits of human vision. While no horizontal viewing angle of a flat surface can reach 180 degrees (at least, as long as there is any separation of viewpoint from surface viewed), the human visual field is actually slightly wider than a half-circle, at about 190 horizontal degrees. The binocular field, however, defined as the range of points that can be seen by both eyes at the same time, is substantially narrower, about 120 degrees.

The remaining 35 degrees on each side of our central binocular field constitutes our peripheral fields of vision. By definition, since objects in our peripheral fields can be viewed by only one eye, peripheral vision lacks stereoscopic perception and its associated sense of depth.

At 96 degrees, 8KUHDTV underlaps our entire range of binocular vision by only a 12 degree sliver on each side. A potential future 16KSUHDTV, with an optimal viewing ratio of just 0.225 screen width (half the 0.45 ratio of 8K), would provide a truly extreme horizontal viewing angle of about 132 degrees, actually exceeding our binocular field and overlapping into our peripheral vision by about 6 degrees on each side.

3. Everything about pixels beyond the W x H array, their length (color coding scheme), and the rate at which they succeed each other falls outside a resolution specification. Specifically, pixel size, and even how bright and vivid a particular color code appears on a screen, is a function of the display used to render pixels.
4. Any assessment of the new UHD 4K (and 8K) resolution standard must be done in two parts.
    1. One part assesses the impact of the new standard on consumers. More specifically, what enhancement does UHD bring to the viewing experience, and what are the physical and financial challenges of acquiring new UHD-capable sets?
    2. The other part assesses the impact of the new standard on businesses involved in the creation, acquisition, and distribution of television programs. More specifically, what technical and financial challenges are posed for these businesses by a shift from today's SD/HD landscape to a new landscape that also encompasses UHD?
5. The impact of the new UHD standard on consumers—the improved viewing experience—is primarily borne by the quadrupling of pixel count, from an array of about 2M pixels for HD to an array of 8M pixels for 4K, then (potentially) to an array of 32M for 8K.
6. Although vital for determining the data rates required for businesses to implement new 4K and 8K standards, from the standpoint of a viewer, channel length and frame rate can be thought of as the invisible (or, if you prefer, technical) parts of a resolution standard. To consider channel length first:
    1. The 8-bit "true" color standard used for HD, which produces 256 shades of each primary color and 16,777,216 24-bit mixed RGB colors, is sufficient to divide the human color space below the threshold of just noticeable differences. In fact, the 8-bit standard produces 2, 3, 4, or more distinct 24-bit codes for each discernible mixed RGB color.
    2. Shifting to a 10-bit "deep" color standard simply multiplies this overabundance of codes by 64. Thus, for example, where 8-bit color generated (say) 3 codes for what, to the human eye, appears as the same RGB mixed color, the effect of 10-bit color is to generate 192 separate codes for that same color. Thus, what is deeper about "deep" color is not the color itself (which, to the eye, remains exactly as it appeared long before the advent of digital color coding) but just the number of binary codes that can generate indistinguishable versions of that color.
7. An exception should be noted to the argument that "deep" color coding (primary color channels of 10 or more bits), considered from the standpoint of the viewer experience, simply multiplies codes for the same limited set of colors that have been distinguishable by the human eye since the beginning of humanity. In fact, 10-bit channels/30-bit pixels are required to supporting the new HDR10 Media Profile standard for High Dynamic Range or HDR colors.
    1. Unlike "deep" color, HDR color does provide a visible enhancement to the viewing experience, by more accurately reproducing the human eye's ability to perceive different colors at the same time, in scenes with high contrast between lighter and darker areas. It is, therefore, likely that HDR color will become a "must have" consumer feature of new 4K sets.
8. As for frame rate, at any speed past about 15 frames per second (fps), the eye blends images together to produce the illusion of smooth motion. Although higher frame rates can help smooth out very fast motion, the plain truth is that motion pictures have survived for a century on a rate of 24 fps without significant complaint. Hence, advances over the 1080i 30 fps rate are likely to prove of only minor viewing significance to most consumers.

9. To summarize the consumer side, then, the primary issue for the enhanced UHD viewing experience is the 4X multiplication of pixels (higher resolution) offered by UHD, followed by the ability of the new standard to support 25% bigger pixels, to provide HDR color. Frame rate is not likely to be a major consumer concern. That, at least, seems to be the moral of the HD shift (since, given a choice, most people seem to opt for higher resolution 30 fps 1080i video over lower resolution 60 fps 720p video).
10. We will return to the subject of possible limits to ever-escalating UHD resolutions on the demand (consumer) side. For now, however, let us turn, instead, to the supply side. What are the practical issues confronting businesses involved with the creation, acquisition, and distribution of video, in making UHD programs available to a mass TV audience?
    1. Despite the increasing proliferation of 4K sets (from 5% of households in 2015, to 15% of households in 2016, projected to jump to near 50% by 2020), there is still very little 4K content of any kind heading into 2017 (including zero regular network/local programs). Does this slow takeoff for 4K content indicate some fundamental problem in providing 4K programs to a mass audience?
11. In general, for businesses involved with the creation, acquisition, and distribution of TV programs, the critical issue is not the enhanced UHD viewing experience itself, but only the extent to which this perceived enhancement helps create consumer demand for 4K programs. Given demand, the primary business issue is simply the feasible cost of providing 4K programs.
    1. This cost burden falls unevenly on businesses, ranging from companies that see UHD as pure opportunity, to companies that see it as pure loss.
12. On both the business and consumer side, the SD to HD transition was enabled by a fundamental technological shift, from analog to digital video. Now that the digital transition is over, however, no similar technical shift exists to enable the transition from HD to UHD.
13. Thus, the difference between HD programs and UHD programs is just the difference between some bits and an order of magnitude more bits. But how this flood of new UHD bits affects costs depends on where a business stands in the chain of creating, transmitting, and receiving bits.
14. To begin at the receiving end of the content chain, for TV set makers—once the challenge of manufacturing screens able to display 8M 30-bit HDR pixels at reasonable prices has been overcome—4K is all opportunity. UHD creates a whole new market for display sales, as buyers seek to replace their old HD sets with new UHD sets.
    1. By the same argument, in future, display makers are likely to welcome an upgrade from 4K to 8K sets, since this again would open a new market for sales, replacing (what by then will be) an increasingly static replacement market for 4K sets.
    2. Indeed, as long as improving technology continues to enable manufacturing still higher resolution displays, set manufacturers are motivated to seek an endless cycle of upgrades, from older lower to newer higher resolution standards.
15. At the other end of the content chain, although the cost of creating programs does rise substantially with 4K content (since existing HD cameras, editing equipment, etc. must all be replaced with more expensive 4K products), these added costs are not decisive. In an overall budget for TV production (including salaries, location costs, marketing expenses, etc.), the added costs of generating video in the more bit-intensive 4K format are doubtlessly a minor line item. Here, the critical question is: Does producing a program in 4K increase audience share, relative to producing the same program in HD? If so, the (relatively minor) extra cost likely will be regarded as money very well spent.

16. After content creators come OTT (Over-The-Top) companies, who make video content available for Internet streaming. Although, this split, between video creation and OTT video distribution, is not as sharp as previously, since major OTT companies (Netflix, Amazon, Hulu) are now also creating their own original content.
    1. Netflix, in particular, has actively promoted UHD as the future of video content, even producing the 2016 season of their own high-profile "House of Cards" program in an upscale 6K format. While this shift to UHD doubtless means that Netflix incurs higher costs, not only to produce their own original programs in UHD, but also to acquire more expensive UHD content from other sources, again, these added costs are doubtless written off as money well spent, in Netflix's campaign to establish itself as an industry leader in video programming.
17. Of course, to garner any marketing or other benefits from 4K video, OTT companies also must make 4K content available to their audiences. The strength of the OTT position, however, is that they do not actually have to deliver 4K content to anyone, ever.
    1. OTT responsibility ends with making the content they provide available in a variety of streaming formats, ranging from high-resolution, high-bitrate 4K versions (if available), to low-resolution, low-bitrate mobile formats. Which of these versions actually gets streamed to any customer is strictly a function of (1) the type of receiving device the customer has, and (2) the amount of Internet bandwidth available between sender and receiver when streaming occurs—two issues over which an OTT company has no control.
18. The major cost problem posed by 4K for businesses, then, does not lie at either end of the content chain, with either video production/acquisition or video reception. Yes, costs do go up at both ends—UHD sets are more expensive to make than HD sets, UHD programs are more expensive to make and acquire than HD programs—but the benefits for set makers, content creators, and OTT providers appear to easily outweigh any expenses that may be added to their operations by shifting from HD to 4K.
19. If UHD poses a real cost challenge on the supply side of video, therefore, this challenge must lie in the middle, with the distribution of 4K content to 4K receivers. The true barrier to an industry-wide transition from HD to UHD programming is the herculean task of transmitting an order of magnitude more bits across available wired and wireless channels in real time.
20. A simple way to appreciate the size of this order-of-magnitude task is to consider its potential impact on the Internet. According to Sandvine's **2016 Global Internet Phenomena** report, video already accounts for the lion's share (over 70%) of all peak down-streaming traffic (with Netflix accounting for an astonishing half of the video share). ***Shifting just 15% of current down-streaming video traffic to a 10X higher rate requires more downstream bandwidth than the entire Internet now provides for all its various purposes***.
21. In general, there are two ways to cope with the 10X rise in bits per second generated by the 4K UHD format that threatens to overwhelm available digital transmission channels.
    1. First, eliminate as many of the additional bits as possible. This may be done both by employing better compression technology, and by implementing only some parts of the specification (in a UHD "lite" format).
    2. When all the disposable bits have been eliminated (by one or the other of these means), the only remaining solution is to increase transmission bandwidths sufficiently to enable delivery of however many bits still remain.
22. **Better Compression:** Regrettably, potential help from this source is not only sharply limited but steadily diminishing until, finally, it vanishes altogether.

1. For video compression, by far the biggest step was the step from no compression (with analog technology) to MPEG-2 compression (with the initial shift to digital video). In truth, MPEG-2 compression is sufficient to eliminate almost all the bits in a raw video stream, leaving perhaps 1 out of every 100. All remaining compression steps deal only with the 1% of data left after MPEG-2 compression.
2. The next most effective step is the shift from MPEG-2 to MPEG-4 compression, which can eliminate perhaps 2 out of every 3 bits left by MPEG-2 compression.
3. MPEG-5 compression (the proposed standard for 4K transmission), will eliminate, at most, just half the bits left by MPEG-4 compression.
4. The gain from any further compression standard, e.g., some future MPEG-6 technology, would be even smaller, eliminating perhaps only 1 bit in every 3 or 4 left behind by MPEG-5 compression.

23. It is not hard to see the logic behind this series of diminishing gains.
    1. Obviously, the goal of data compression is not the trivially easy (and pointless) task of eliminating all bits, but rather the much more difficult task of identifying and eliminating, first, any redundant bits, and second, any insignificant bits. When this has been accomplished, what remains is just a residual hard core of important bits. The only way to make any further reductions is to begin eliminating the important bits, too. For video, this means visibly degrading image quality.
    2. Non-degrading compression gains, therefore, are limited to throwing out what is either redundant or insignificant. For any bit pool, the number of bits in this disposable group may be larger or smaller, but it is always limited to some fraction of all the bits in the pool.
    3. Thus, the difference between less vs. more powerful compression technologies is simply a difference between less vs. more effective targeting of the disposable bits in a bit pool. The more effective the compression technology used, the fewer disposable bits remain in the bit pool. The fewer the disposable bits that remain, the smaller the possible gain from applying still better compression.
    4. In the end, better compression must fail as a useful data reduction strategy. Either no disposable bits at all remain to be eliminated, or the disposable bits that do remain are too few and elusive to repay the effort needed to eliminate them.
24. An alternative (and far simpler) way of eliminating the bit flood created by UHD is to not implement the less visible but more expensive parts of the standard, e.g., for 4K, the demand to double the frame rate from 30 to 60 fps. Continuing to transmit 4K at 30 fps would, at a stroke, cut the raw data rate increase for 4K in half.
25. Counting everything—including the proposed ability of MPEG-5 to cut MPEG-4 data rates in half, and ignoring any increase for doubled frame rates—it seems that quality 4K frames with 30-bit HDR pixels could be transmitted at a data rate conservatively set at 18Mbps. Even lower rates may be feasible. But data rates approaching MPEG-4 HD (5Mbps) are out of the question for 4K. Accommodating the minimum new data demands of a 4X increase in pixel count and a 25% increase in pixel length requires substantially higher data rates, if not 5X more than MPEG-4 HD (25 Mbps), then 3-4X more (15-20 Mbps) at a minimum.
26. What are the prospects for an increase in transmission bandwidths to at least 18 Mbps? We will look at this in two parts, first, for wireless Over Air Broadcasting (OAB), then for wired Internet streaming, or Over The Top (OTT) broadband delivery.
27. For OAB, the changes made to improve video compression (over MPEG-2) and data multiplexing (over 8-VSB) in the new ATSC 3.0 standard should provide ample bandwidth,

not just for 18Mpbs 30 fps 4K video, but for full 60 fps 4K transmission at 36Mbps or higher. It is possible, but less certain, that it will also prove sufficient for 8K transmission (another near 10X jump in raw data rates over 4K), though probably only in a limited form that does not embrace 12-bit color and frame rates of 120 per second.

    1. The deployment of ATSC 3.0, however, remains shrouded in mystery. In part, this is simply because the standard is not yet complete (as of the start of 2017). In part, this is because ATSC 3.0 is incompatible with ATSC 1.0—an awkward fact that makes it unclear exactly how or when a transition from 1.0 to 3.0 can be implemented.
    2. Although it is not clear when ATSC 3.0 will deploy, it is clear that ATSC 1.0 is entirely inadequate to support 4K data rates. Which is to say, until 3.0 deploys, OAB will be limited to SD and HD transmission only.

28. For Internet delivery of 4K content, as mentioned above, the mere existence of 4K receivers, and 4K content queued for delivery to those receivers by OTT providers, is inadequate to trigger 4K delivery. Given the availability of a 4K receiver, the remaining missing parameter is the amount of bandwidth available on the Internet.
    1. Looking at current figures for average Internet bandwidth (15 Mpbs, at the end of 2016), it seems likely that a minority of customers already have data rates adequate to support 4K.
    2. At the current growth rate in Internet bandwidth (about 25% a year), this group should grow rapidly. By 2022, it should be possible to accommodate much, if not all of the streaming that takes place in HD today, in a 4X more expansive UHD format, on an Internet that provides three times the capacity of the 2016 Internet.
29. I conclude the technical barriers to delivering 4K video streams in real time, while substantial, are not insurmountable, either for OAB or Internet streaming.
30. The barrier to 4K OTT video is not so much financial as temporal. The Internet continues to be the beneficiary of substantial annual upgrade investments which, over the next five years, should be sufficient to enable 4K streaming for a large part of the OTT 4K audience.
31. The financial barriers for OAB broadcasters are more significant, which will play out as a temporal delay in OAB 4K. The ATSC 3.0 standard, needed to enable 4K OAB broadcasting, will be complete this year. But it will not be widely (let alone fully) implemented for some years to come for financial reasons. First, it requires broadcasters to invest heavily in new equipment to support 4K (on top of an ongoing investment in HD equipment). Second, since 3.0 is incompatible with current 1.0 technology, a 3.0 transition puts current OAB customers at risk. As a result of these financial barriers, it is impossible to accurately project a timeframe for an OAB transition to 4K, though the generalization that it is likely to be later rather than sooner is probably safe. If 5 years is long enough for this transition, then OAB 4K may be reaching a majority of its potential audience about the same time that Internet growth is enabling a majority of the OTT streaming audience to receive 4K.
32. If 4K content is created, 4K receivers are common, and transmission paths can support 4K data streams, does that mean 4K is inevitable? As technical progress on the supply side continues to improve through the 2020s into the 2030s, will 4K then be followed by 8K? And, with continued technical progress through the 2040s into the 2050s, will 8K be succeeded by 16K? And so on, until by 2100 the world is at some presently inconceivable 128K resolution standard with data rates of 1.5 quadrillion bps?
33. In short, suppose there is no technical barrier to still higher bit rates on the supply side, or at least none that time and human ingenuity cannot surmount. What about potential limits on the demand side? After all, for any product, demand is the key economic issue. No demand,

and supply of any product will vanish (however easy it may be to produce). High demand, and—through one means or another—any product that can be provided, will be (however difficult it may be to produce).

34. On the demand side, there are two major issues for viewers. One is the quality of the viewing experience itself. The other is the cost of the viewing experience. These two issues are tightly interrelated. If the cost penalty of a better experience is sufficiently high, the mass market will opt for the lower quality experience, however great the quality difference. If the cost penalty is sufficiently low, the mass the market will opt for the higher quality experience, however small the quality difference.
    1. It bears mentioning that, in any area as subjective as viewing preferences, it is not necessary for differences in viewing quality to be objectively measurable. What counts for the mass market is only the widespread perception that devices of type A are superior to type B, for buying patterns to favor type A (that is, in the absence of any significant cost advantage for B).
35. As mentioned earlier, any significant cost advantage for HD over 4K is rapidly vanishing, if it has not entirely disappeared. Removing the economic barrier throws the demand issue entirely onto the question of perceived video quality. What are the quality advantages of 4K over HD? Also, as mentioned earlier, for a viewer, these reduce primarily to 4X higher resolution and 10-bit HDR color. Better color is a noticeable advantage. What is the advantage of higher resolution?
36. As long as the market remains fixed on any given resolution standard, as display screens get bigger, the pixels in its fixed W x H array must get bigger in proportion to cover the screen. A screen twice the width and height of another screen (with 4X the area of the other screen) must have pixels that are twice as wide and high as the other screen (that are 4X the size of pixels on the other screen). The practical effect of the fact that bigger displays require proportionally bigger pixels is a proportional increase in "recommended viewing distance".
37. The primary effect of introducing a new resolution standard is to unlocks this requirement to grow and shrink pixels, as screens grow and shrink. For screens of the same size, a new resolution standard shrinks pixels. For pixels of the same size, a new resolution standard expands screens. Specifically, for 4K vs. HD, a 4K screen of the same size as an HD screen, will have pixels ¼ the size of HD pixels. A 4K screen with pixels the same size as pixels on an HD screen, with be twice the width and height of the HD screen.
    1. For example, a 40" HD screen will become an 80" 4K UHD screen with pixels of the same size, or remain a 40" screen with pixels ¼ the size of HD pixels.
    2. The transition from 4K to 8K would have exactly the same effect on 4K pixel/screen size. An 80" 4K UHD screen would become a 160" 8K UHD screen with pixels of the same size, or stay at 80" with pixels ¼ the size of 4K pixels.
38. The effect of this change to bigger screens with pixels of the same size, or screens of the same size with smaller pixels, can be summarized as a change in the immersive experience of watching TV, which in turn reduces to a change in the optimal viewing distance for a screen. This change in optimal viewing distance can be precisely analyzed with a basic understanding of human visual resolution and a little pixel math. Specifically, we need to:
    1. Fix a base value for human resolution, in terms of ppi and distance
    2. Determine how to calculate pixels per inch (ppi) for any display screen
    3. Determine how to scale ppi resolution numbers with distance
    4. Calculate an "optimal viewing distance" (ovd) for TVs of all sizes
    5. Establish a fixed viewing ratio for each resolution standard: HD, 4K, 8K

   6. Relate a standard's viewing ratio to its horizontal viewing angle, as a basic measure of immersion.

39. In fact, human resolution varies widely from person to person, time to time, and with circumstances. Therefore, the base value assumed here—based on a "retina" screen of 300ppi viewed at a distance of 1 foot—can be thought of as a "best case" number, that doubtlessly assumes better vision than most people enjoy.
40. That is to say (for people with excellent vision), when viewing a screen from a foot away, fewer than 300 ppi waste powers of perception, since more can be discriminated at that distance than can be represented by available pixels. Conversely, more than 300 ppi waste pixels, since unaided human vision is not powerful enough to see the extra detail provided by the added pixels.
41. Since each resolution standard fixes the number of pixels/row, determining ppi for any screen is as easy as dividing the number of pixels supported by the resolution standard it implements, by the width of the screen, measured in inches. Thus, for a 60" wide 'full HD' screen with 1920 pixels/row, its ppi value is 1920/60 = 32. Reimplemented in 4K, the same width screen would have 3840/60 = 64ppi; while, done in 8K, its ppi value would be 128 (7680/60).
    1. More simply, since 4K is double the row length of HD, for the same width screen, 4K doubles HD pixel density, while 8K doubles 4K and quadruples HD pixel density.
42. A simple rule of thumb for scaling human visual resolution with distance is to convert the base "retina" screen values of 300ppi at a distance of 1 foot to a ratio, 300/1, and simply change the number of feet in the denominator. Thus, at 10 feet, those with excellent vision can discriminate no more and no fewer than about 300/10 = 30ppi. At 100 feet, the limit of our resolution is about 3 ppi (300/100); and at 300 feet, the limit of our resolution is a single pixel about 1 inch on a side (300/300).
43. The "Recommended Viewing Distance" parameter offered by TV set manufacturers can be thought of as a way of compensating for pixel size. Ideally, pixels will be seen as color points rather than as distinct rectangular color blocks. Fortunately, any pixel block, regardless of size, can be reduced to a visual point by the simple expedient of viewing it from an appropriate distance. Thus, as screens (pixels) get bigger, recommended viewing distances also increase, to maintain the appearance of pixel points.
44. For our purposes, we will redefine "recommended viewing distance" as "optimal viewing distance" (ovd), meaning the furthest distance from which the full resolution of the screen still can be discerned, i.e., the distance at which all pixels remain discernible as individual points. From any further away, the screen will have details (pixels) that are impossible to discern. From any closer, the pixels will start to appear as blocks, rather than as points, creating artifacts (like "jaggies") that degrade the visual experience.
    1. For calculation convenience, we will reduce pixel size to the single dimension of width or ppi. Of course, for HD and UHD screens, height will scale with width in the ratio of 16:9 (W:H).
45. For any TV screen, its optimal viewing distance in feet can be determined by dividing the 1 foot value of 300ppi by the ppi value of the given screen. Thus 300/ppi = ovd. For example, for a 5-foot wide (60") HDTV screen, with a ppi value of 32, its ovd is 300/32 = 9.375 feet. Or, the denominator and quotient of this equation can be reversed (300/distance = ppi), to determine the optimal ppi value for a given viewing distance, e.g., 300/9.375 = 32.

   1. Thus, if your seating distance in a room is 6 feet, you want a TV with a ppi value of 300/6 = 50ppi. For HDTV, this means a set with a screen width of about 38 inches (1920/50 = 38.4).
   2. Note that ovd is relative to pixel size (ppi), not to screen size. Any screen with 32 ppi is optimally viewed from a distance of 9.375 feet, regardless of how wide or narrow that screen may be. This means, for example—assuming your seating distance does not change—when upgrading from an HD to a 4K TV set, optimally speaking, you want a 4K set that is twice the width of your HD set.

46. For pixels of the same width (screens of the same ppi), for any resolution standard, the ratio of ovd to screen width (sw) is fixed. Consider a 60” 32ppi HDTV screen, optimally viewed from a distance of 9.375 feet. This ratio: ovd/sw = 9.375/5 = 1.875 is the same for all HDTV sets. Thus, a 32-inch wide HDTV is optimally viewed from a distance of 60” (1.875 x 32 = 60)
    1. Running the calculation for a 32-inch wide HDTV the long way, 1920/32 = 60ppi. 300/60 = 5 feet. Or consider a 10 foot (120”) wide HDTV screen. By the 1.875 screen width ratio, this set is optimally viewed from 18.75 feet away. Working the calculation, the long way, 1920/120 = 16ppi, 300/16 = 18.75 feet.
    2. In short, the 1.875 ratio of ovd/sw is invariant for HD screens. In general terms, HDTV is optimally viewed from a distance of just under 2X the width of its screens (or from 1.875X its width, in precise terms).
47. Since ovd is fixed by pixel size (ppi), the ovd/sw ratios for 4K and 8K screens can easily be computed from the ovd ratio for HD. Specifically, the 4K ratio is ½ the HD ratio, or 1.875/2 = 0.9375 screen width. Similarly, the 8K ratio is ½ the 4K ratio, or 0.9375/2 = 0.46875 screen width. Rounding off, the optimal viewing distance for HD is just under 2X screen width, for 4K is just under 1X screen width, and for 8K is just under half the screen width.
    1. Working the 4K ratio calculation formally, for an HD screen of 1w (width unit), a 4K screen with the same ppi will be 2w. Since ovd does not change for the same ppi, for a HD ovd of 1d (distance unit), the general HD ovd/sw ratio of 1d/1w becomes 1d/2w for 4K, or simply 1/2.
48. The diminishing fixed ratio of sw to ovd for each successive resolution standard also fixes a growing horizontal viewing angle (hva) at the optimal distance. Specifically, the hva for HD is 31 degrees, for 4K is 58 degrees, and for 8K is 96 degrees.
49. The recommended horizontal viewing angles in a movie theater range from a minimum of 36 degrees in the back row, to a maximum of 62 degrees in the front row. Comparing these recommendations to the hva computed for TV resolution standards above, while the HDTV hva of 31 degrees falls below the minimum theater angle of 36 degrees, the 4K hva of 58 degrees falls near the upper 62-degree limit of a theater experience. The extreme 96-degree hva of 8K is probably closer than most people want to get to a screen for an optimal viewing experience (just as most people do not want to sit in the front row of movie theater, and still less if it were moved to just half its current screen distance).
50. Thus, the conclusion of this analysis of resolution standards in terms of horizontal viewing angle or screen immersion is that, while HDTV falls below minimum theater standards, 4K will reproduce a near-front row theater viewing experience. Combined with the 10-bit HDR color advantage of 4K over HD 8-bit color, 4K is a significant viewing enhancement over HD.
51. 8K, however, when viewed objectively, looks like an unnecessary step for the home viewer. Doubling 4K screen width without changing viewing distance is likely to make the viewing experience, if anything, too immersive for comfort. Conversely, maintaining a comfortable screen width for the available viewing distance with 8K will make the pixels too small to be

discerned, i.e., is simply a waste of resolution. Also, 12-bit color, assuming it is the same as HDR 10-bit color except with 64X more color codes, is just a further waste of bits.

52. This conclusion, that 4K is a sensible home standard but 8K and higher standards are not, can be reinforced by considering the size problem for living-room screens. caused by doubling the width and height of successive screens. If we begin with a 60" diagonal HDTV set, for an average living room with a 9-foot viewing distance, the 4K set for this room ballons to 120" diagonal measurement. This is a screen over 9 feet wide and 5 feet high. Locating a screen this size in a living room is doubtless challenging, but doable in many (if not all) cases.
53. But, if we assume 8K means doubling size again, to a 240" diagonal set with a screen about 19 feet wide and 11 feet high, we doubtless pass beyond the boundary of the possible. Assuming the 4K screen dimensions do not exceed available space in most average living rooms, it seems certain that these 8K dimensions will, in all but the most exceptional cases.
54. Based on this analysis, it seems that 4K will bring a true theater experience into the home, with visibly better immersion and color. In this respect, it is both a useful and desirable upgrade of the HD standard. Also, while certainly very large, 4K screen sizes do not seem impossible large, especially if a room is planned around the 4K screen size from the beginning, as its dominant feature (e.g., in a "home theater" room).
55. However, based both on its excessive horizontal viewing angle and outsized screen dimensions, a further upgrade to 8K seems to provide more than the average consumer either wants or can possibly use. Since a future 16K standard would redouble all the 8K values, 16K (or higher) seems to be a potential standard past consideration.
56. For broadcasters, the technical issue raised by 4K is how to move the 10X more data required by this standard, relative to HD, to receiving TVs in real time. Of course, the raw 4K data rate of 15 billion bits/second can be reduced both by compression and by implementing "lite" forms of the standard, with lowered requirements. But even the most advanced form of compression currently available, MPEG-5 or HEVC, is limited in its ability to prune unneeded bits, and there probably no acceptable way to get around the need to quadruple HD pixel count, and shift from 24-bit pixels to 30-bit pixels, to support HDR color.
57. Consequently, the bottom line for 4K data transmission is likely to be
58. at the Bigger channel numbers (more bits/pixel) allow finer subdivisions (more distinct shades) of a primary color. An 8-bit channel provides binary codes for 256 separate shades, while a 10-bit channel number codes 1024 separate shades.
    1. Human Vision (HV) can perceive no more than about 150 shades of a color. Hence, 8-bit "true" color primary channels divide a color into more shades than we can see.
        i. This means that, with the 256 divisions enabled by 8-bit channels, adjacent shades of the same color are imperceptibly different from each other (i.e., their difference falls below the threshold of just noticeable differences). Thus, for example, a red color bar built up from 256 strips, shading from pure red (255 0 0) on one end to pure black (0 0 0) on the other, will not appear to the eye as composed of distinct strips, but as a continuous gradient, shading by imperceptible degrees from red to black.
    2. Channel numbers bigger than 8 bits simply code more indistinguishable gradations of a shade. Thus, while 10-bit channels code 4 primary color shades for every 1 shade coded by 8-bit channels, the additional 3 color shades are, to the human eye, indistinguishable both from each other and from the original 8-bit shade.

3. The number of “mixed” RGB colors that characterize individual pixels can be determined by multiplying the number of shades for each primary color.
   i. For example, 16-bit “high” color pixels use 5 bits to encode shades for red and blue, and 6 bits to encode shades for green. A 5-bit channel provides codes for 32 shades, a 6-bit channel for 64 shades. An 8-bit “true” color channel provides codes for 256 primary color shades. A 10-bit “deep” color channel provides codes for 1,024 primary color shades. Hence:
      - 3 5/6-bit “high” color channels provide 32 x 64 x 32 = 65,536 mixed 16-bit RGB colors.
      - 3 8-bit “true” color channels provide 256 x 256 x 256 = 16,777,216 mixed 24-bit RGB colors.
      - 3 10-bit “deep” color channels provide 1024 x 1024 x 1024 = 1,073,741,824 mixed 30-bit RGB colors.
   ii. The number of mixed colors can be proliferated endlessly, by the expedient of making the channel number bigger and bigger (“deeper” and “deeper”). Thus 12-bit channels (4,096 primary shades) provide 64 times more mixed colors than 10-bit channels, or 68,719,476,736 mixed 36-bit RGB colors. 20-bit channels (1,048,576 primary shades) provide over a quintillion mixed 60-bit RGB colors. 30-bit channels (1,073,741,824 primary shades) provide over an octillion mixed 90-bit RGB colors. And so on and on.
4. The question, then, cannot be ‘How many primary/mixed colors are possible?’ because, however many there are, it is always possible to have more. The question, instead, must be ‘How many primary/mixed colors are enough?’ This question, however, is always relative to some purpose.
5. For purposes of human vision, i.e., generating all the mixed RGB colors the human eye can perceive, 8-bit “true” color primary channels are enough. There is no precise answer to the question, ‘How many different colors can the human eye perceive?’ but the general answer is probably somewhere in the low millions. Even for the most perceptive, it is certainly well below 10 million.
   i. Which is to say, while the 2,097,152 mixed RGB shades generated by 7-bit primary color channels (with their ability to code only 128 color shades) may fall short of human capabilities, the 16,777,216 mixed RGB shades generated by 8-bit primary color channels exceed the ability of even the most discriminating of eyes to detect differences by many millions of shades. In other words, for every mixed RGB color the eye can perceive, 8-bit “true” color coding will (on average) provide 2, 3, 4, or more different codes.
   ii. 10-bit “deep” color coding, with its billion mixed RGB colors, does not change the ability of the human eye to perceive, at most, a few million mixed RGB colors. Rather, it multiplies codes for what, to the human eye, appears the same color by a factor of 64. For example, if 8-bit color coding provided 3 different codes for what appears to be the same mixed RGB color shade, with 10-bit color coding there will be 3 x 64 = 192 different codes for that same shade.
6. The value of “deep” color, then, has nothing to do with the viewing experience. Whether a specific pixel on the screen is represented by one of 3 possible distinct 24-bit codes or one of 192 possible distinct 30-bit codes is an imperceptible

technical fact about it. As such, this fact about the pixel cannot possibly matter to the viewer or affect their experience in any way.

7. Rather, “deep” color is of value in other contexts, as shown by the following two very different cases.
    i. Color film captures thousands more colors than the human eye can distinguish. Thus, if the goal is to preserve all the information present in old color pictures, a scanner with 16-bit color channels (“48-bit color”) is desirable.
    ii. Video editing generally results in information loss, which may include collapsing different pixel codes to a single code. In this case, the more codes there are for the same color, the less likely it is that collapsing codes will result in any visible degradation of the image.
8. The subject of “deep” color (primary color channels of 10 or more bits that subdivide each just noticeably different color into a set of imperceptibly different colors) is sometimes confused with a different color topic, known as HDR or High Dynamic Range color, since the HDR10 Media Profile standard also uses 10-bit channels/30-bit pixels. HDR10, however, is not used to define hundreds of codes for the same color (or, what is the same thing, hundreds of imperceptibly different colors) but, rather, to display a wider range of colors across the same image.

59.

Appendix 1:
Binary Color Coding

Although the long strings of 1s and 0s that comprise binary numbers can seem quite daunting on first encounter, understanding binary numbering is really very easy. The basic rule is just that every bit added to a binary number doubles the number of possible combinations supported. This can be seen most readily by starting at the beginning, with 1 bit, which has only 2 possible values (0, 1). Adding a second bit allows 4 possible values (00, 01, 10, 11). And so on: 3 bits have 8 possible values (000, 001, 010, 011, 100, 101, 110, 111), 4 bits have 16 possible values, 5 bits 32 possible values, etc. By the time you reach the 8-bit values used in "true color" RGB encoding, this doubling algorithm has passed by 64 (6 bits) and 128 (7 bits) to reach 256 possible combinations of 1s and 0s between 00000000 and 11111111.

The doubling rule itself is most readily understood by the fact that adding a bit simply allows us to write all the numbers of the previous set twice over, the first time tacking a 0 on to the front of all the previous numbers, the second time tacking on a 1 (compare the 1-bit vs. 2-bit and 2-bit vs. 3-bit values shown above).

8 bits is regarded as "true" color since it is the first channel value safely past the outer limits of human color perception. That is to say, if you build up a color bar out of 256 strips, each with an adjacent shade of, for example, red—running from a strip of pure red on one end to a strip of pure black (no color) on the other—this color bar will not appear to the eye as 256 distinct stripes, but rather as a single continuous gradient, shading from red to black by insensible steps. Which is to say, when a color is divided into as many as 256 distinct steps, we have moved below the threshold of noticeable differences between adjacent steps—in other words, no one can tell shade 1 from shade 2, shade 2 from shade 3, and so on down the row of 256 shades.

Since there are 256 8-bit binary numbers (possible combinations of 1s and 0s between 00000000 and 11111111), an 8-bit channel provides encoding for 256 distinct shades each of Red, Green, and Blue, or 256 x 256 x 256 = 16,777,216 "mixed" colors. Since this is more shades of a primary color than the eye can distinguish (human visual discrimination tops out somewhere around 150), 8-bit channel encoding is also called "true" color—a term introduced to distinguish 24-bit pixels from earlier standards that supported fewer mixed colors, e.g., the 5/6/5 RGB 16-bit pixels of "high" color (providing a greatly reduced, if still large, palette of 65,536 mixed colors.)

In fact, for most people, the same would be true of a color bar built up from 128 strips (7-bit channels), but the very keenest eyes under ideal conditions might be able to distinguish very faint stripes in this bar. So, 7-bit color channels are not quite past the limits of human perception. But 8-bit channels with their associated 256 shades are well past the outer limits of vision.

No one knows for certain exactly where the outer limits of human perception fall, which in any case will vary widely by individual. However, a ballpark number for best case human vision is probably somewhere around 150 primary color shades, with most falling well short of that number. But if we take 150 as an outer limit, then it might be possible for some to distinguish 150 x 150 x 150 = 3,375,000 RGB mixed colors. Since 7-bits produce only 128 x 128 x 128 = 2,097,152 mixed RGB colors, 7-bit color channels fall short of the "true" color mark, understood as encompassing every possible RGB mixed color anyone might ever distinguish under the most favorable circumstances. But 8-bit color channels, which multiply the number of 7-bit RGB mixed colors by 8 (= 2 x 2 x 2), are easily sufficient to include not only all the colors anyone might ever be able to distinguish under any circumstances, but many millions more besides that no one can tell apart from their neighbors.

But, wait, if 8-bit "true" color channels are sufficient to generate millions of more colors than anyone can perceive, it is only natural to wonder what is the advantage of a new 10-bit "deep" color standard?

Before answering that question, let's consider the binary logic of 10 bits. This much, at least, is straightforward. Since each new bit added to a binary number doubles the number of combinations previously possible, the move from 8- to 10-bit color first doubles the number of possible primary 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, to repeat, what is the point of moving from 8-bit "true" color channels to 4X more expansive 10-bit "deep" color channels? Isn't this just adding over a thousand million more mixed colors no one will ever see?

There are two answers to this question. Regrettably, the first answer is that the marketing of "deep" color often relies on a simple "more is better" argument. If 8 bit "true color" is good, then 10 bit "deep" color 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 of the 256 8-bit shades into 4 separate sub-shades, all of them indistinguishable to the human eye. This point is worth dwelling on because there is a hidden truth lurking here, just below the surface.

In fact, the term "deep" color is a generic term, embracing 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? Or 1024-bit color? Indeed, why should the ambition for more color bits ever end?

The lurking hidden truth 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 over 20 years ago, there has been no clamor for yet another upgrade to 128 bits, and no such demand seems to be on 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 insanity.

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 eight. Add as many colors as you like past the 8-bit “true” color RBG 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 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 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 a family photo album, it’s a good idea to shop for a 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. But that reason is ***not*** about showing new mixed RGB colors that no eye will ever see.

Rather, it has to do with a color technology known as High Dynamic Range (HDR). HDR color is not about adding a billion or so colors we can't possibly see, but about overcoming a limitation of our image capture technologies to generate pictures that more closely resemble what we can see.

We have already mentioned one way in which artificial capture technologies may exceed the capabilities of the human eye, namely, in their ability to detect wavelength differences within the visible spectrum too subtle to be noticed by the roughly 6 million color-sensitive cones in our eyes. Thus, while the human eye cannot profitably exceed 8-bit "true" color channels (since its ability to discriminate distinct shades of a color tops out somewhere around 150), our superhumanly sensitive cameras and scanners might need to discriminate considerably more than 256 shades per primary color channel.

However, there is also an important way in which the human eye is generally superior to artificial capture mechanisms. The 120 million or so highly sensitive rods in our eyes, which detect light levels, respond to a much wider range of luminosities than our light-recording mechanisms can capture. This issue is familiar to anyone who has selected a camera exposure value ("f-stop" and shutter speed) for a scene with strong light to dark contrasts. The eye, with its naturally high dynamic range, can make out details at both ends of this light-to-dark spectrum. Relatively speaking, however, our light-recording mechanisms typically have low dynamic ranges.

In practice, this means that, setting a camera at one end of its exposure scale, results in darker parts of the image taking on enough contrast to distinguish features, while lighter parts of the image wash out. Conversely, setting the exposure at the other end of the scale, results in lighter parts taking on enough contrast to make out features, while darker parts black out. Whereas, the compromise of setting the exposure in the middle, results in feature loss in both the darkest and lightest parts of the image.

HDR color solves this problem, in effect, by taking three different exposures of every scene, one optimized for the lighter parts, one optimized for the darker parts, and one optimized for the middle parts. The best parts of each of the three resulting images are then combined into a single image, that more nearly represents the high dynamic range the eye perceives when viewing the scene.

The problem for classic 8-bit "true" color coding is not merely that 8-bit channels are inadequate for capturing HDR images, but high dynamic range images cannot be reproduced on displays using just 24 bits/pixel. The recent (August, 2015) Consumer Electronics Association HDR10 Media Profile standard, for HDR compatible displays, uses 10-bit channels (as the "HDR10" name suggests). Of course, 10-bit channels require displays to process 30 bits of color information for pixels, which is to say, to make use of HDR capabilities, content providers must transmit 30-bits/pixel of color information to displays.

Appendix 2:
Adaptive Bitrate Streaming

To make sure streaming content gets though to their customers, streaming companies take advantage of something known as Adaptive Bitrate Streaming (ABS). ABS leverages the fact that, unlike traditional broadcast technologies, where data flows only one way from transmitter to receiver, the Internet is bidirectional. An Internet provider can not only send data to a customer, they can receive data back. Two-way communication opens up a whole new realm of interactive possibilities, including ABS.

The bidirectional communication involved in ABS consists of sending data to a target address and, in real time, receiving back information about available bandwidth at the target destination. Thus, the usual protocol on receiving a video request is to begin sending a relatively low-bitrate stream to the target address. For this reason, when you first start up an on-demand video, the picture is often not very good.

As video delivery starts, however, the receiving device reports back to the sender on its available bandwidth. Based on that feedback, the sender then adjusts the bandwidth of the stream being sent to the target. Thus, if more bandwidth is available, the sender switches to a higher bitrate feed, as appropriate, and the softer, lower-bitrate picture originally sent suddenly improves. Or, if even the initial, lower bandwidth feed exceeds the available capacity, the feed starts at a still lower bitrate. Or, should available bandwidth be below the lowest usable level for the target display, a “try again later” message is sent.

A standard ABS ladder, with full HD as its top rung, might consist of 10 steps, ranging, say, from a high resolution of 1920 x 1080 at 5800Kbps to a low resolution of 320 x 240 at 235Kbps. Using this ladder, a client that fails to meet the high requirement of 5800Kbps might still get a (softer) full HD picture, at a lower bandwidth of 4300Kbps. If the client device fails even that bandwidth test, the picture will shift from “full HD” to 720p HD (1280 x 720) at 3000Kbps. A softer version of this format might be available at, say, 2350Kbps. Below 2350Kbps, the picture reduces again, to a 720 x 480 SD format, streamed at 1750Kbps. And so on, all the way down to the bottom rung of the ladder, the small screen 320 x 240 mobile format streamed at 235Kbps.

In short, you may have a “full HD” TV, but if all don’t have at least 2350Kbps of streaming bandwidth available when you’re watching House of Cards, you won’t be seeing it in HD. And if you don’t have at least 4300Kbps available, you will be seeing it in 720p HD rather than a full HD 1920 x 1080 format.

This sort of ABS ladder is the secret behind Netflix’s cheerful promotion of 4K content. Yes, they will make House of Cards available in a 4K UHD format. All it takes is adding more UHD rungs to the top of their ABS ladder. But they do not say you, in your particular circumstances, will receive House of Cards in 4K. In fact, they do not guarantee any streaming customer will ever receive the program in this format. In reality, the version customers receive depends entirely on the sort of receiver they have, and the fluctuating bandwidth available at that receiver while the program is playing.

Thus, if you have a UHD receiver hooked up to a pipe with (say) 18Mbps available, Netflix surely will be happy to send you House of Cards in a 4K 3840 x 2160 format. Of course, this will not be the "full 4K" version of the program (with 10-bit HDR color channels streamed at a rate of 60fps, which certainly would require far more bandwidth). But, at least, it will be 4K resolution. However, if no one has 18Mbps available at 9PM on a Sunday, then no one watching House of Cards at that time will get the UHD version of the program. Yes, many, perhaps all can still see the program. But it will be shown in a lower resolution, lower bitrate version, as determined by the particular circumstances of each customer.

To repeat, OTT providers just provide the content—in every format, up to and including UHD—but they are not responsible for which format actually gets delivered to you. That is determined by fluctuating circumstances, over which they have no control, across an Internet they never built.