

Display Myths Shattered: How Monitor & HDTV Companies Cook Their Specs

By Dr. Raymond Soneira

Take everything you think you know about displays and throw it out the window. It's time for a clinic on what display specs really mean.



Vision is our most amazing and complex sense, so it's no surprise that display technology is so amazing and complicated. It's also no surprise that most consumers don't have a good understanding of how displays function, or the best way to select them, buy them, use them, and adjust them.

Not only are displays getting more complicated and harder to understand, but the competition between manufacturers has gotten so brutal that marketing gimmicks—ploys that exploit the average consumer's technical ignorance—are playing an increasing role in driving sales. The goal of this article is to point out and explain some of the most important myths, misconceptions, and misunderstandings about display technology. Much of what you're going to read is like the classic tale of *The Emperor's New Clothes*. What you've been told about the latest and greatest thing really isn't there, or better, or meaningful, or even visible.

In the following pages, I'm going to discuss user controls, contrast ratios, pixel response time, and color gamut. These topics comprise just a portion of what a savvy consumer needs to know. But for now, let's just start our journey with what should be the best question to ask before buying a new display: "What are the most important manufacturer specs to compare?" Unfortunately, the answer is none, because they're all exaggerated marketing specs rather than objective scientific specs. The only specs that are useful and meaningful are those in reviews that evaluate every display with the same consistent methodology.

Confusing Users with User Controls

One reason why most consumers don't understand their monitors and TVs is because some of the most important user controls have misleading and technically incorrect names. No wonder folks can't figure out how to adjust them. In fact, they misadjust them, and then usually just leave them misadjusted permanently. Here are some highlights—well, lowlights really—of inane user-control engineering.

When Brightness Isn't


On mobile displays with only a single user control, the control labeled “brightness” does in fact actually control the brightness of the image on the screen by increasing or decreasing the backlight intensity. However, on most monitors and TVs, the control labeled “brightness” does not control the brightness. It actually controls the signal-level setting for black on the display, which indirectly has a minor effect on brightness.

Contrast? Not So Much

The control labeled “contrast” has absolutely no effect on image contrast. It actually controls the brightness of the image, by increasing or decreasing the amplitude of the video signal. Monitors and TVs really should have a true contrast control, but the closest you'll find on some HDTVs is an obscure control labeled “gamma,” and I have yet to see one that works properly.

Controls of a Bygone Era

Even more shocking, today's digital monitors and HDTVs still have the same basic user controls that were found in the original analog NTSC color TVs from 1953: brightness, contrast, tint, and sharpness. These controls only made sense for analog signals on the old NTSC television system. Brightness controlled the CRT direct-current bias, contrast controlled the video amplifier gain, tint controlled the phase of the color subcarrier, and sharpness performed analog high-frequency peaking to compensate for the limited video bandwidth of the old vacuum tube amplifiers. Today, none of these controls are necessary for digital signals.



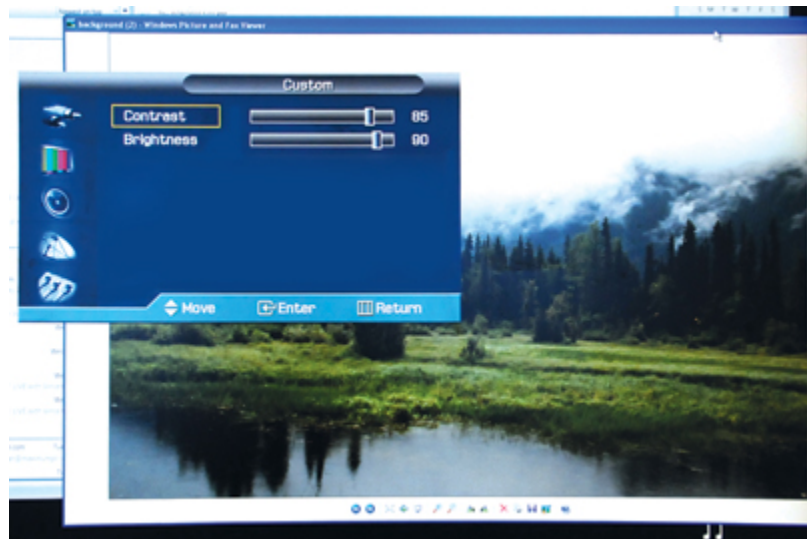
Rotary controls for a mid-century analog CRT. Those Contrast and Brightness controls are legit.

Brightness and contrast controls shouldn't be there because, for digital video, the black level is fixed at level 16, reference white at 235, and peak white at 255. Similarly, tint and phase have no real meaning for digital signals. Finally, the sharpness control isn't appropriate for digital displays because in a digital image there's no transmission degradation—the image is received exactly as it appeared at the source. Sharpening the image involves digitally processing the pixels, which leads to artifacts and noise unless it's done at resolutions much higher than the final displayed resolution, which, of course, isn't the case inside your monitor or HDTV.

Controls that Do Worse Than Nothing

Most monitor and HDTV user-menu options are actually unnecessary features added for marketing purposes—gimmicks to suggest the display has unique features that other models lack. Even worse, most of these options actually *decrease* image and picture quality.

In many cases, it's not even clear what these sham controls really do. The documentation seldom explains them, and I even know engineers from high-level manufacturers who don't know what the controls do, either. When I test TVs, I spend an inordinate amount of time using test patterns to figure out what the options and selections really do, and in most cases, turning off the fancy options leads to the best picture quality and accuracy.



Digital on-screen controls for a Samsung Syncmaster 242MP—that really have no business being labeled Contrast and Brightness.

The following is a list of useless (or near-useless) menu options and selections from three HDTVs sold by major brands: Black Corrector, Advanced CE, Clear White, Color Space, Live Color, DRC Mode, DRC Palette, Dynamic Contrast, xvYCC, Color Matrix, RGB Dynamic Range, Black Level, Gamma, White Balance, HDMI Black Level, Fresh Contrast, Fresh Color, Flesh Tone, Eye Care, Digital NR, DNle, Detail Enhancer, Edge Enhancer, Real Cinema, Cine Motion, Film Mode, Blue Only Mode.

Some of the terms sound impressive, but almost all of this is unnecessary puffery and jargon that confuses not only consumers but the pros, as well.

Contrast Ratio, Ad Absurdum

Both manufacturers and consumers are obsessed with contrast ratios. Because many people choose the model with the highest number, manufacturers have developed new contrast ratio specs to win at this game.

It's a sordid business, but deserves exposure, so let's jump in.

When Contrast Ratio Actually Matters

A careful, objective measurement of contrast ratio can be very revealing. After the display is accurately calibrated for optimum picture quality, the contrast ratio is determined by dividing the brightness of peak white by the brightness of black. In principle, the greater the ratio the better. Just be aware that contrast ratio is important only for low-ambient-light viewing, which is where black brightness values matter most. (In high-ambient-light settings, reflections off the screen abound, and they're all brighter than the display's own internal black.)

Even more to the point, a high contrast ratio really only matters when there is significant dark picture content, like you see in some movies. It's much less relevant for most TV shows because the picture seldom has much very dark content, and the image never really dips down to black except briefly between scenes. As for games, well, just consider which games really run a lot of "pure" black. If for some reason you're still playing in the dark, underground worlds of Doom 3, you need a high contrast ratio. But if the colorful world of Plants vs. Zombies is more your speed, you need not worry.

Keeping Up with the Joneses

Contrast-ratio specs are tremendously inflated. For the best LCDs, scientifically measured contrast ratios are actually between 1,500 and 2,000. But manufacturers almost never publish real contrast ratios anymore. You'll only find these true values in a small number of articles and publications. Yes, contrast-ratio values have been steadily improving over the years, but the year-to-year change is relatively small, which isn't good for marketing.

In their quest to quote ever-larger numbers, some manufacturers invented a completely meaningless spec called "dynamic contrast ratio," which is what's being prominently advertised now. Sometimes they don't even bother mentioning the "dynamic" part. Sadly, all manufacturers are now forced to play this game, as consumers wouldn't be interested in monitors and TVs that tout the true values. Meaningless contrast-ratio specs help substandard manufacturers by making their displays appear to be just as good as those from the best manufacturers, or even better, because the biggest liar wins. This not only hurts consumers, but it also hurts the better manufacturers because they're unable to publicize their superior technology.

Big but Not Big Enough

So what's really so "dynamic" about this bastardized contrast-ratio spec? It's really quite simple: When the display's video signal is entirely black or very close to black, the display's electronics go into a standby mode that significantly reduces the light output of

the unit. This much darker standby value is then used when computing the contrast ratio—instead of the real value of when a picture is actually present.

Obviously, this trick doesn't change the true black or true contrast ratio for any picture that's not all black, so it's meaningless for picture quality. The primary reason for measuring the spec this way is that published contrast ratios can now go from about 1,500 up to, well, infinity. In 2008, many TVs were advertised with a "dynamic contrast" in the range of 15,000 to 35,000. Now, in 2010, some go into the millions and beyond. There's no real improvement, of course. It's just the same trick with a bigger exaggeration.

When Infinite Means Nothing

At my local Walmart in Amherst, New Hampshire, the Sony KDL-52EX700, an LED-based TV, is listed by Walmart as having an "infinite contrast ratio" on its information label. First of all, "dynamic" was left off—it should say "infinite dynamic contrast ratio." This is then technically correct because the LEDs turn off when an all-black image is present. This results in a division by zero, and produces the infinite result. But this is also nonsense because the LEDs need to be on whenever an actual picture is present!

Labeling like this is intentionally misleading to consumers. Walmart should set an example for other retailers and refuse to show misleading manufacturer specs to its shoppers.

Response Times: How Fast Is Fast Enough?

All displays show artifacts of one sort or another when their screen images change rapidly. It's most easily detected with moving objects, or when the entire screen moves due to camera panning. In many cases, it's not the fault of the display. Rather, it arises somewhere in the signal path from the source, which can be caused by camera blur, interlaced scanning, MPEG compression artifacts, poor video processing, insufficient bandwidth, or insufficient CPU speed in the case of games. Further confusing matters, artifacts can occur for different reasons with CRT, LCD, LCoS, plasma, and DLP technologies. It can even occur with OLEDs, if switching speeds aren't sufficiently high.

But when people discuss motion artifacts, they are generally talking about LCD response time. And not surprisingly, the manufacturers' published specs for response time have become one of the major deciding issues for many consumers. As a result, in the last five years or so, manufacturers have somehow pushed response time numbers from 25ms (milliseconds) to an essentially untenable 1ms.

So what, if anything, do these specs really mean?

Behind the Basics of Blur

Motion blur arises when the liquid crystal—the active element within an LCD—is unable to change its orientation and transmission rapidly enough when the picture changes from one frame to the next. Because the standard video rate is 60 frames per second, a

pixel is expected to fully update its light-transmission opacity within 16.7 milliseconds (that is, in one 60th of a second). If it takes any longer than that, the image will show some degree of lag, which appears as a trailing smear or blur whenever there is motion.

LCD motion blur is generally evaluated with an industry-standard spec called response time that (supposedly) measures the time that it takes for a pixel to go from black to peak-intensity white, and then back to black again. However, most picture transitions involve much smaller and more subtle shades of gray-to-gray transitions, which can take longer to complete.


But it gets even more complicated than that because every pixel is actually made up of independent red, green, and blue sub-pixels that have their own separate intensities, frame-to-frame transitions, and times. The upshot is that visual blur within a detailed, moving picture is a fairly complex and nebulous phenomenon.

Motion Blur: Visual Proof

Motion blur is one of the most visually tangible display problems—the evidence speaks for itself in screenshots and photography, both of which can illustrate the relationship between response time and motion blur. In this article, I've included high-speed screenshots of moving DisplayMate test patterns, as well as moving test photos taken of a top-of-the-line, 120Hz HDTV (shot with a Nikon DSLR using a fast shutter speed of 1/160 of a second). These images were taken in 2008, but the results wouldn't be much different today.

The published response time for this HDTV model is 8ms. Since this corresponds to a double transition (from black to peak white, and then back to black again), the single transition time (from black to white, or from white to black) should therefore be about 4ms.

But is the pixel response time really that fast? To find out, I ran DisplayMate tests in which black and white squares move across the screen at measured speeds. In the examples here, one photo shows the squares racing across at 1,093 pixels in a single second. The second photo shows the squares moving about 50 percent faster, covering 1,609 pixels in a single second. The white tips seen on the edges of the ghost images are artifacts resulting from electronic overdrive processing that's being used to try to improve the response time by exaggerating transitions.



The screen on top was grabbed from a motion speed test of 1,093 pixels per second, with the black and white squares indicating real-world pixel response times of 60ms and 40ms, respectively. The screen on the bottom shows the same lengthy response times during a faster test of 1,609 pixels per second.

As you can see from my screenshots—each a brief snapshot in time with a shutter speed of 1/160 of a second, which is less than the refresh rate—it's possible to make out at least eight individual screen refresh cycles on this 120Hz display. Indeed, in the screenshots, each square is shifted from the other by 1/120 of a second, which is 8ms, and those ghosted squares indicate that the older images haven't yet dissipated. The upshot is that you're looking at a true response time of about 65ms. In fact, a response time of much less than the 8ms refresh rate would be needed for there to be no visible blur. Obviously, 65ms blur in the screen shots doesn't jibe with the manufacturer's single-transition response time spec of 4ms.

The DisplayMate tests clearly demonstrate that the TV's real LCD response time is considerably longer than its published spec would indicate.

But What about Moving Photos?

It must be stated that DisplayMate test patterns are very sensitive to imaging effects—this is by design. Photographic images, on the other hand, typically consist of a very complex and varied admixture of blended picture elements. With so much going on in

an image, motion blur is easily obscured and lost within the complex, variegated imagery of a typical photograph. In particular, photographs of real-world content lack uniform backgrounds, and uniform backgrounds—like those in my DisplayMate tests—make it easier to see the motion blur trails. For this reason, we would expect moving photographs to show much less visible blur than what's demonstrated with test patterns.

To wit: If you look carefully at the magnified marching band images below, you can see a total of at least six refresh cycles in the second photograph. This corresponds to a real-world pixel response time of 50ms. But the motion blur is still much less noticeable than what we see in DisplayMate's punishing test patterns.



Witness two photos shot by Lauren Soneira. On top, a static image—thus no motion blur. On the bottom, the image is moving at 1,018 pixels per second, and you can see six refresh cycles indicating a 50ms response time.


Photographs are static images and moving them across a screen is quite different from live video, where images are a complex and varied mixture of continually blending

picture components that are themselves constantly changing in both time and position. With all this screen activity going on, we would expect to detect much less motion blur in live video than with either of the moving static photographs or test patterns.

And Now for the Tests You've All Been Waiting For

To evaluate motion blur and artifacts in live video with lots of high-motion picture content, we set up a side-by-side comparison shoot-out with 11 HDTVs, and had both consumers and experts evaluate them. The top-of-the-line LCDs had 120Hz screen refresh rates, the top-of-the-line LCD had strobed LED backlighting, and the other units had standard 60Hz screen refresh rates. Two of the units were plasma displays, and one was a pro-grade CRT studio monitor. The goal was to determine the degree to which this varied technology affected visible motion blur.

All of the HDTVs were fed identical, simultaneous digital video using an all-digital High Definition Tivo and a Blu-ray player. They were all compared side-by-side in the configuration as shown in the photo. The content included both daytime and nighttime sporting events, TV shows, and movies, all with lots of action. If any viewer thought he or she detected motion blur on any HDTV, we would repeatedly press the eight-second Tivo backup button and watch the sequence over and over again on all of the units until we fully understood exactly what was happening on each display. We did the same thing with the Blu-ray player and its content.



Behold our 11-display shootout—with the lights on. Photo by Dieter Michel.

The conclusions from all participants were consistent across the board, and will likely surprise most consumers: There was essentially no visually detectable motion blur on any of the LCD HDTVs in any of the video content we assembled.

When people thought they saw motion blur, with only a handful of minor exceptions, the blur was either in the source video or a temporary visual illusion that disappeared when the segments in question were reviewed. Unlike what we empirically identified in moving test patterns and moving photographs, the eye is unable to detect the blur in live video because the images are much more dynamic and complex—and undoubtedly because of the way the brain processes and extracts essential information from visual images.

So, Is Blurring Even an Issue for Videos, Movies, and Games?

For all of the tests—the DisplayMate test patterns, the moving photos, and the live video—we found that there was no visually detectable difference in motion blur for the mid- to top-of-the-line LCD HDTVs. This regardless of their claimed pixel response times, 60Hz or 120Hz refresh rates, strobed LED backlighting, or motion-enhancement processing. If you find this surprising then just re-read the classic tale of *The Emperor's New Clothes*.

The underlying reason why higher refresh rates don't mitigate blurring is that the true pixel response times of displays are considerably longer than the 60Hz video frame

rate, so it doesn't matter whether the screen refresh rate is 60Hz or 120Hz, or whether the LED backlights are strobed off during the frame updating. Similarly, adjusting the electronic processing enhancements that some models offer—controls that are supposed to reduce motion blur—only served to introduce objectionable contours, edges, and other artifacts onto moving objects without reducing the overall motion blur.

So that's the story on video. What significance do these results have for PC gamers?


First, while motion blur isn't generally noticeable with live video, it's more likely to be seen by gamers who intently focus on particular moving objects. For this reason, the blur illustrated above with test patterns and test photos applies.

Second, don't pay much attention to a manufacturer's response time specs because they are so different from the real response time and motion blur that we have demonstrated here.

Third, while 120Hz refresh rate monitors and HDTVs don't inherently improve on motion blur over the 60Hz models, they are generally equipped with better performing panels and electronics, so they may still produce superior image and picture quality. And if you're a movie buff, the 120Hz units should offer better motion interpolation from the 24 frames per second used in all movies shot on film. The 60Hz models need 3:2 pull-down, which produces judder, but most people seldom notice it.

Fourth, be aware that the latest 240Hz displays don't offer any real picture-quality performance improvements, and are just a marketing gimmick taken to an absurd level.

Color Gamut or Marketing Gambit?



Color gamut, which is the range of colors that a display can produce, is undoubtedly the most misunderstood and exploited spec—precisely because it's natural to believe that the range should be as large as possible. While that's true for most specs (even when they're exaggerated), it's definitely not the case for the color gamut.

The color gamut that you want on all of your PC monitors, laptops, HDTVs, and even smartphones is the same color gamut that was used when the content you're viewing was created. If a different gamut is employed, you'll see different colors than you're supposed to see.

Virtually all consumer content is created using industry standards that specify the exact color gamut to be used. For computers and digital cameras it's sRGB. For digital HDTVs, it's called ITU-R BT.709 (often referred to as Rec.709). Fortunately, both of these standards specify the same exact gamut. Yes, there are other color gamuts for specialized applications (more about that later), but sRGB and Rec.709 cover virtually all consumer content, and that is the color gamut you want on all of your displays. The

color gamut in these standards specifies the exact color coordinates for the three red, green, and blue primary colors, which are used to produce all color mixtures on screen.

Now that you're versed on what color gamut is, let's share what it isn't—as illustrated by six examples of manufacturer misinformation.

Bigger Isn't Better

One common misconception frequently exploited by manufacturers is that a wider color gamut indicates a better display that produces more realistic colors. This is absolutely wrong. A larger gamut will simply make all the screen colors for standard production content appear more saturated than they ought to appear. Indeed, displays claiming more than 100 percent of the standard color gamut simply can't show colors that aren't in the original source image. Expanded gamuts are just gimmicks that make consumers think they're getting something better.

The Perils of Recalibration

If you do get a display with a larger color gamut, it's necessary to reduce the gamut back to the sRGB/Rec.709 standard values by adjusting color saturation via a user control. Unfortunately, if the display isn't calibrated at the factory to match the standard color gamut, it's unlikely you'll be able to visually adjust it properly yourself. This kind of adjustment typically requires professional calibration using instrumentation.

NTSC? Never!

The often-quoted NTSC Color Gamut is from 1953. It's also obsolete and irrelevant. Computers, digital cameras, and HDTVs use the sRGB/Rec.709 color spaces, and specs should refer to them instead of NTSC. As explained above, values greater than 100 percent of the standard color gamut aren't desirable—unless you like punchy, unrealistic, oversaturated colors.

Adobe RGB or Not to Be?

As stated above, there are specialized color gamuts for specialized applications, and some of these are larger than the sRGB/Rec.709. Adobe RGB, one of the more common ones, is used by imaging professionals and you'll find it as an option on some digital cameras and scanners. Just be aware that if you use the Adobe gamut, you will also need a display that produces the Adobe gamut, and only a small fraction of consumer displays can do this. If you display an image produced with an Adobe gamut on a monitor with a standard sRGB/Rec.709 gamut, the colors will be incorrect and oversaturated.

How the Eyes Play Tricks

Adobe RGB is a larger gamut than sRGB/Rec.709, but be aware that for most applications, gamut size doesn't matter very much. The further out you go in color space, the less frequent the colors appear in nature, so the human eye doesn't notice that they're not quite right except in rare circumstances (like when viewing a full-screen rendering of a very red tulip). When faced with a gamut beyond their rendering range,

displays simply wind up reproducing the closest most saturated color they can under the circumstances.

Bit-Depth Misconceptions

Manufacturers will also dupe consumers by advertising useless and misleading specs about the number of screen colors produced by their displays. Screen color counts have absolutely nothing to do with display's color gamut, though manufacturers will attempt to tie them together. In reality, a display's maximum number of colors is a function of the total number of intensity-level combinations that the device can produce.

Let's do the math. Standard 24-bit color has eight bits per primary color, and eight bits generate 256 intensity levels. Because there are three primary colors, the number of possible color combinations is 256^3 —16.8 million colors. Now, if a manufacturer uses 12-bit color processing internally within the same display, there are (in theory) 4,096 intensity levels and 68.7 billion possible colors. Sounds impressive, yes, but the display's color gamut remains the same as before and the additional number of colors doesn't mean anything visually.

Still not convinced? First, remember that essentially all consumer content is 24-bit color. Thus, the source images have only 16.8 million colors, and the display can't "invent" intensities and color combinations that don't exist in the original.

Second, true onscreen 24-bit color does a good job of meeting the human eye's color and brightness discrimination abilities.


Third, be aware of the real reason why additional processing bits are necessary. Onscreen intensity levels are not supposed to be linear. Rather, they should follow a standard gamma curve with a nonlinear mathematical 2.2 power-law exponent (meaning the screen brightness for any sub-pixel varies as $s^{2.2}$, where s is the input signal intensity level 0-255). The extra processing bits are necessary just to get the display to produce the gamma curve accurately on screen.



Maybe Sharp Should've Consulted Mr. Spock Instead


By now, you've surely seen ads for Sharp's Quattron four-color technology. George "Mr. Sulu" Takei dons a lab coat, and fawns over Sharp's introduction of a yellow primary-color sub-pixel to the traditional three-sub-pixel, RGB primary-color arrangement. According to Sharp, this results in "expanding the color gamut and faithfully rendering nearly all colors that can be discerned with the unaided human eye, especially golden yellow."

If you have read this far, you already know that Quattron is just another shameful marketing gimmick. HDTV television and movie content is produced and color-balanced on three-color displays that are accurately calibrated to Rec.709. Sharp's fourth primary color is yellow, and there isn't anything for it to do because yellow is already being accurately reproduced with mixtures of the existing red and green primaries. More importantly, a Quattron display can't show colors that aren't in the original three-color source image. So what good is it? None, unless you like to see over-exaggerated yellows.



But could it be that existing consumer HDTVs are unable to reproduce the standard sRGB/Rec.709 color gamut, so Sharp's fourth primary color actually has something useful to do? We decided to find out.

Colors and color gamuts can be accurately measured and then plotted in a chromaticity diagram to compare values to the standard. What follows is from a 2008 article where I used a spectroradiometer to measure the color gamut of HDTVs in the DisplayMate lab. To the right, are the results for a Sony consumer LCD HDTV. The black triangle is the Rec.709 standard and the red dots are the measured values for the red, green, and blue primary colors of the Sony display. Notice that the Sony measurements all fall exactly where they should on the triangle vertices. It's perfect! In short, this Sony HDTV accurately shows exactly the same colors seen by, say, the director at a TV studio. *Ipsa facto*, Sharp's fourth color is absolutely superfluous and can only decrease picture quality and accuracy! Undoubtedly, part of the Quattron's "Yellow Push" is being produced with simple video processing. Some people have been impressed watching the Sharp demos on the Quattron, but manufacturers demos are always fine-tuned to get a maximum wow response, so be careful before jumping to any conclusions about how it will perform displaying content at home.



See that narrow strip of gamut real estate between the top of the black triangle and the inside of the white line? That's where Quattron technology would have to live.

Note that in our figure, the outer white curve represents the limits of human vision. While the Rec.709 standard is much smaller, it's important to note that the colors between the black triangle and white curve aren't common in nature. Yes, a display can only reproduce the colors that lie inside of the polygon formed by its primary colors, but because yellow falls between the red and green primaries, Sharp's yellow primary would need to lie somewhere *outside* of the red and green leg of the color triangle. But there isn't much room between the Rec.709 triangle and the human vision curve, is

there? For this reason, it's difficult to see why a yellow primary sub-pixel is needed unless Sharp isn't able to put its red and green primaries where they belong.

Sharp shows its Quattron color gamut in some promotional material by using an old (x,y) distorted CIE Diagram from the year 1931 because it makes its extended color gamut look much larger than it really is. Our figure is the (u'v') 1976 Uniform CIE Diagram and shows the color gamut accurately.

---Dr. Raymond Soneira is President of DisplayMate Technologies Corporation of Amherst, New Hampshire, which produces video calibration, evaluation, and diagnostic products for consumers, technicians, and manufacturers. A research scientist with a career that spans physics, computer science, and television system design, Soneira was a Long-Term Member of the Institute for Advanced Study in Princeton, and a Principal Investigator in the Computer Systems Research Laboratory at AT&T Bell Laboratories.