

Displays Of A Different Striped

**NEW IMAGE-MANAGEMENT
TECHNIQUES PROVIDE BRIGHTNESS
FOR HALF THE POWER, EASING
THE BRUTAL TRADEOFF IN VIDEO-
RICH GADGETS** By Joel Pollack


Today's flat-panel displays provide bright, crisp, and vivid images—and they use plenty of power while doing it. It's a tradeoff that hardly mattered when we rarely watched movies, played games, or surfed the Web on anything other than furniture-size monitors. But that power consumption is a serious engineering constraint today, when more and more of us are getting our visual data on the go, from cellphones, video iPods, and game players like Sony's PSP. And as serious as the constraint is now, it will soon become downright intolerable as engineers strive to wring far more vivid visual information out of the next generation of portables than can be displayed by anything now on the market.

Fortunately, remarkable power savings—as much as 50 percent—can be achieved by simply redesigning the display to provide no more information than the eye can absorb and the brain can digest. This strategy is called biomimetic, because it deliberately mimics a living system.


Biomimicry has long been used in audio. For many years microphones, amplifiers, and speakers have been designed in sizes and frequency ranges that match the human auditory system. Similarly, the telephone

Primary Colors


High-resolution color displays add to the media-rich features in next-generation cellphones, driving power consumption way up and operating life way down.




TRADEOFFS: Most flat-panel displays today, with their 1:1:1 ratio of red, green, and blue elements, compromise on power consumption and resolution. These elements, or subpixels, are arranged in either a conventional stripe pattern [left] or a delta pattern [right]. Because light from the blue subpixels does little to help the eye resolve images, most of it goes to waste.



EASY ON THE EYES: The Bayer pattern [left] saves power by including extra subpixels in the color green, which the human visual system resolves with particular efficiency. The PenTile Matrix [right] yields even better results by changing the size and arrangement of the pixels.



WHITE MAKES BRIGHT: Replacing one of the extra green subpixels of a Bayer pattern with a white subpixel [left] provides better color balance and improves brightness. A further modification in the PenTile pattern [right] pushes performance up another rung.



TWO DO THE JOB OF THREE: A PenTile red-green-blue-white pattern [left] requires fewer columns than a red-green-blue stripe pattern [right], providing both economic and performance advantages.

system crams calls through limited carrying capacity by editing the frequencies down to a limited bandwidth. Such compression techniques have also been applied to audio and video software, as seen in the Moving Picture Experts Group (MPEG) and other algorithms. Now it is time to apply biomimicry to displays.

It all begins with the retina, the part of the eye that converts photons into electrochemical signals that are interpreted by the brain as images. The retina's most discerning photosensitive elements, or photoreceptors, are the cones. Except in a color-blind person, the human eye has three kinds of cones, each having a different type of protein, called a photopigment.

One kind of photopigment is specialized to sense photons in the reddish-yellow band of wavelengths, the second in the greenish-yellow band, and the third in the blue one. Because the typical eye has about 30 red- and green-sensitive cones for every blue one, almost all the work of resolving an image—its luminance, edges, and other structural detail—is done with output from the red and green cones, which also detect color, of course. The blue cones detect *only* color [see sidebar, "The Human Visual System at a Glance"].

Yet despite the great preponderance of red and green cones in our eyes, most flat-panel displays produced today, like just about every color TV tube produced in the past half century, have a 1:1:1 ratio of red, green, and blue color elements. These elements, called subpixels, are arranged in either a stripe or a delta pattern [see illustration, "Tradeoffs"]. Because the blue subpixels do almost nothing to help the eye resolve images, most of it goes to waste.

Over the years, researchers have come up with ways to minimize the waste. In the 1970s, Bryce Bayer, of Eastman Kodak Co., came up with the Bayer pattern, with a 1:2:1 ratio of red, green, and blue subpixels, with the green subpixels linked diagonally, as in a checkerboard [see illustration, "Easy on the Eyes"]. That pattern was used by General Electric Co., in Fairfield, Conn., in avionics displays in the late 1980s. It made the displays somewhat more efficient, but it had problems. One was that the ratio of colors in the pixels gave the screens a distinct greenish cast. Another was that the scheme could not, for a given density of pixels, render imagery with the highest possible level of detail [see sidebar, "A Sharper Image"].

Basically, what we have done at our company, Clairvoyante Inc., in Cupertino, Calif., is to rectify these shortcomings of the Bayer pattern. First, our pattern of subpixels, called the PenTile Matrix, addresses the Bayer's color imbalance by going easy on the green. In addition, our pattern is rotated with respect to that of the Bayer by 45 degrees; this rotation has the effect of mapping the conventionally square orientation of the incoming pixel data to the high-luminance green subpixels on a one-to-one basis. We also resize the subpixels with respect to each other, making the green ones a lot smaller, a process that renders information with higher resolution.

Another seemingly slight tweak enhances the display's performance immensely. By adding a white (or clear) subpixel to form a red-green-blue-white pixel, we dispense with one out of four color filters. We thereby boost efficiency: color filters absorb wavelengths from the backlight, so they sap energy from a display. We call this scheme the PenTile RGBW pattern.

We weren't the first to come up with an RGBW display; past experimenters did it by simply swapping one of the two green subpixels for a white one in every pixel of a Bayer display. What we did differently was to squeeze the subpixels into long rectangles.

Even so, there is still an inherent inefficiency in using four subpixels instead of three. To get around the problem, we found ways to render a pixel with an average of just two subpixels—two-thirds as many as in the conventional RGB pattern. We do

PREVIOUS PAGE: DISPLAY PHOTOGRAPHY: MARK RICHARDS; ZEBRA IMAGE: MATTHIAS CLAMER

The Human Visual System at a Glance


Different kinds of information are processed by different photosensitive cells in the retina of your eye, traveling to the brain over different nerve pathways. The brain combines signals from these visual channels to interpret the image.

Light passing through the cornea is focused by the lens onto the retina, passing first through fine blood vessels, then through neural wiring to reach the photoreceptors. A layer of neurons on top of these photoreceptors then processes, encodes, and ships off the resulting data to the brain via the optic nerve. The optic nerve terminates in the brain's lateral geniculate nucleus, a kind of switchboard that directs data elsewhere into the brain for further processing.

The photoreceptors fall into two broad classes—rods for night vision and cones for high-resolution color vision. The rods provide high gain at the cost of low resolution, in part because each rod is connected to many cells in the optic nerve. The cones resolve fine details, because each one is connected through an intermediate layer of image-processing nerves to a single optic nerve cell.

For the sensors to work properly, the light hitting them must vary continually, so as to refresh the signal. The eye performs such refreshing, even on static images, by making slight intermittent movements, called saccades.

Half the cones are in the fovea, a 2-millimeter area in the macula. It spans 2 degrees of the visual field; of these cones, half fall in the very center, spanning



ning 0.5 degree. Cones come in varieties sensitive to long, medium, and short wavelengths, corresponding roughly to red, green, and blue, respectively. There are roughly 30 times as many red- and green-sensitive cones as blue-sensitive cones.

The red and green cones perceive the range of wavelengths between 500 and 700 nanometers on the visible-light spectrum. The human eye tends to focus on yellow (580 nm), because that wavelength gives the best overall image quality, exciting the green and red cones nearly equally. Far-red and far-blue wavelengths provide little resolution.

Output from the cones undergoes preprocessing in neural layers in the

retina, right above the cones. The neural layers encode an image into a luminance channel, which distinguishes black from white, and two chrominance channels, one managing red versus green, the other, yellow versus blue.

These preprocessing neural layers generate the luminance channel by combining inputs from the red cones and the green cones and effectively ignoring the blue cones. They generate the red/green chroma channel by subtracting green-cone input from red-cone input. The yellow/blue chroma channel subtracts the blue-cone inputs from those of the red and green cones.

—J.P.

it by using software algorithms to create, in effect, virtual pixels. Basically, the algorithm fools the eye. It defines an edge of an object in an image with the red, green, and white subpixels, and adds the requisite dash of blue off to the side, on the ground that the eye cannot discern the exact location of the blue bits, anyway. Such tricks provide very crisp images with good color.

Another trick enhances the color on a single pixel indirectly. The goal is to minimize the number of color subpixels needed to display an image by getting each one to work as hard as possible in resolving the image. And there is a lot of freedom in doing so. For instance, to enhance a red pixel on a gray background, you can add a dash of white—increasing the luminance of the red—and also turn down the surrounding blue and green. The eye perceives this reduction of the blue and green as an enhancement of the red.

The bottom line is that brightness and color can be conveyed in more than one combination of red, green, blue, and white. Orange, for instance, will look the same to a human eye whether it comes from a single, pure wavelength at 600 nanometers or from the combination of two or more wavelengths from the red and yellow bands of the spectrum.

In tests, our PenTile Matrix pattern did well on many kinds of image files, including video, computer-generated graphics, and still and moving images [see illustration, “Best by Test”]. The PenTile technology, however, works well only in high-resolution formats. If the resolution is too low—below about 185 dots per inch for near-range devices like cellphones—it can produce artifacts, such as texture in the background. The matrix did especially well displaying the higher-performance versions of the Joint Photographic Experts Group (JPEG) and MPEG-2 compression formats. A version of MPEG-2 is what is used to compress video data so that a whole movie and more can fit on a DVD. Briefly, MPEG-2 has several versions; the kind used on DVDs is known as 4:2:0, which indicates the ratio of the samples used to convey the moving image's brightness (the 4) and the color (the 2 and the 0). A superior form of MPEG-2 is 4:2:2, because it allows for more detailed color sampling.

Our PenTile Matrix did particularly well with this superior form of MPEG-2. That was encouraging to us, because MPEG 4:2:2 is the compression format used for high-definition video, which is growing in popularity with the proliferation of big-screen TVs

A Sharper Image

Any system that dices a picture into a matrix of pixels risks the distortion known as aliasing. A simple example is the jagged rendering of certain characters seen on a computer display. More interesting distortions, such as the moiré pattern, may emerge when one matrix is used to sample data on another, as happens when one looks at a screen door through another screen door. Yet although the human eye filters the world through a matrix, it rarely produces aliasing, and that is why engineers try to copy its powers.

The trick is to find out how the eye ignores extraneous information without overburdening its sampling mechanism. The world's television standards have been designed according to this biomimetic strategy.

Analog color encoding schemes in both the NTSC color TV and the JPEG/MPEG standards resemble the eye in that they break an image into a luminance channel, which conveys brightness, and



two chrominance channels, which convey color. Because both of the chrominance channels are defined by their difference from a third color, it is possible to deduce that third color, which is needed to show a full-color image. The only difference

from the human visual system lies in what is inferred. The JPEG/MPEG technologies provide red and blue channels and infer green; the human visual system provides red versus green and yellow versus blue and infers luminance. (NTSC is the National Television System Committee, JPEG is the Joint Photographic Experts Group, and MPEG is the Moving Picture Experts Group.)

In modern digital cameras, the image forms on a charge-coupled device that contains an array of photosites, each of which converts light into electricity. Each photosite—regardless of the color of the color filter array above it—maps to a conventional, full-color, red-green-blue pixel. Algorithms sample the array of photosites to produce a signal that fits a different array—that of the display. The success of the translation depends on the sampling rate.

Sampling a signal can be compared to a person running in the dark over a road with regularly spaced bumps. If he takes small steps, he will detect the shape of the bumps, but if he takes longer steps, he will miss the bumps altogether or misconstrue their form. The general rule is that signal components more than half the sampling frequency will produce aliasing; therefore, the algorithm that assembles and interpolates pixels will work only so long as the channel signals are each less than half the original sampling frequency for their color.

Chromatic aliasing appears as color fringes around the edges of objects. It was a problem in the earliest algorithms that attempted to sample not pixels but subpixels—the elements of which pixels are made. Those early subpixel rendering algorithms used filters based on decimation, a process of dropping information in order to fit high-resolution data into a lower-resolution format. Decimation works moderately well in displaying coarse-grained images in slow video, but when it tries to cram detailed, fast-moving video into a coarse display—say, by displaying NTSC TV resolution on a 240-by-320-pixel color mosaic—it can produce chromatic aliasing. Decimation is commonly used in viewfinders, where the low quality doesn't matter.

The next step in the development of the art of subpixel rendering was taken by IBM Corp. display engineers in 1988, when they reduced chromatic aliasing with a filter that represented a group of pixel values as a weighted average. In a later development, Microsoft Corp., in Redmond, Wash., designed its ClearType font-rendering technology to minimize aliasing with a system of filters that accounted for the slight spatial displacement of all red subpixels vis-à-vis the green ones and the green ones vis-à-vis the blue ones.

—J.P.

and the imminent arrival of high-definition video players, such as those offering the competing Blu-ray and HD DVD formats.

Vertical black-and-white lines show off the advantages of the PenTile RGBW [see illustration, "Two Do the Job of Three"]. While the conventional pattern must turn three adjacent columns on and three other columns off, our display renders the same detail by turning just two on and two off. Because this scheme requires just two-thirds as many subpixels as the conventional one, fewer transistors and drive lines are needed to control the subpixels, and less of the display's area needs to be obscured by opaque elements. In other words, the aperture ratio increases, letting more light through. This improvement, together with the white subpixel, provides about twice the brightness for a given draw of power. The savings in manufacturing costs more than balances any increase occasioned by the addition of a fourth, clear, color filter.

Engineers can use the gains to save power or to rev up the brightness. In a cellphone with a 2.8-inch display, you can get a luxurious 350-candela-per-square-meter level of brightness for 475 milliwatts. That's a power level that would give you a meager 175 cd/m² if you used a conventional pattern. If, however, you are content with the meager brightness, you can make do with 233 mW—and with a backlight having half as many light-emitting diodes.

In anticipation of the expected demand for optimized displays, several manufacturers have already begun incorporating these power-efficient subpixel patterns and rendering algorithms. So far, the following LCD companies have publicly demonstrated PenTile display technology: AU Optronics, BOE Hydis, CPT, LG Innotek, Samsung, and Wintek. Silicon Works and Tomato LSI have also developed chips.

Clairvoyante hopes to see a PenTile display in a commercial product by the end of this year.

The art of designing products to conform to the needs of the body has been dignified by a name, ergonomics. The expansion of information technology into new domains means that engineers must now learn to make products that conform to the needs of the mind and the senses, as well.

ABOUT THE AUTHOR

JOEL POLLACK is president and chief executive officer of Clairvoyante Inc., a display-architecture company in Cupertino, Calif. He has also worked in a variety of research and management jobs in the display businesses of Sharp, Tektronix, and Xerox.

TO PROBE FURTHER

The advantages of biomimetic displays are discussed in "Reducing Pixel Count Without Reducing Image Quality," by C.H.B. Elliott, *Information Display*, December 1999, Vol. 15, pp. 22–25.

Check out the blog by Greg Hitchcock and Bert Keely, inventors of Microsoft's ClearType, at <http://blogs.msdn.com/fontblog/default.aspx?p=2>.

For a discussion on the total list of processing operations required to view a pixel, see "What Is a Pixel?" by J.F. Blinn, in *IEEE Computer Graphics and Applications*, September–October 2005, Vol. 25, no. 5, pp. 82–87.

More information on subpixel patterns is available in "Subpixel Rendering on Nonstriped Colour Matrix Displays," by D.S. Messing et al. in *Proceedings of the 2003 International Conference on Image Processing*, Vol. 2, pp. 949–52.