X-bit's Guide: Contemporary LCD Monitor Parameters and Characteristics

This guide is offering you a detailed description of various LCD monitor parameters, which are essential for comparative analysis, as well as the measuring techniques used for each of them. Besides, since the majority of LCD monitor parameters are determined by the type of the matrix the monitor is based on, we will also introduce to you the major four types of contemporary matrices. They are: TN+Film, S-IPS, MVA and PVA.

by **Oleg Artamonov**
10/26/2004 | 02:58 PM

I've been repeating the same idea through all my previous reviews of LCD monitors on our site: one of the main disadvantages of currently manufactured LCD monitors is that each particular model lacks versatility. I mean a good CRT monitor is always good whatever you're doing with it (working with text, processing photographs, playing games and so on), whereas an LCD monitor that's suitable for games won't suit for handling photos, and one with an excellent color reproduction would be a bad choice for dynamic games.

Formally, latest LCD monitor models all seem to have technical parameters that allow using them in any area as the manufacturers claim viewing angles of 160 degrees, a contrast ratio of 500:1, a true representation of all 16 millions of colors that you want to see, and the gap between the specs of different models is seemingly negligible – well, can the average human eye, not equipped with measurement tools, tell the difference between the 160 viewing degrees of a good TN+Film matrix and the 170 degrees of a PVA, MVA or IPS matrix? However, theory and practice are more diverse in practice than they are in theory, and if you placed two monitors next to each other – say, one with a TN+Film matrix and another with an IPS one – you would easily see their parameters to be different visually, even if you've never worked with LCD panels before.

I'm not driving to say that the manufacturers deliberately overstate the characteristics of their devices, thus confusing the user (well, sometimes they do, but not too often). I say that it's a matter of perspective or *what* the manufacturers mean by a specific parameter and *how* they measure it.

Speaking generally, any measurement of any value should start with a precise definition of the measurement method and conditions and the applicability limits of the result. Without that, the measurement result makes no sense and has no practical worth.

Regrettably, many reviewers and testers, following the latest trend to get objective data rather than subjective impressions, forget this simple rule and fall a prey to two common mistakes: they either erroneously regard the number they've got (I say "number", but not "result" because a number can only become a result after all the above-said items are complied with) as a real parameter of the LCD monitor, which it is in fact not, or put much emphasis on parameters of secondary importance, which but slightly affect the investigated characteristic (for example, "color reproduction" is a complex characteristic which cannot be described by a single parameter). The first mistake is often committed when some side factors, specific for the particular measurement method, come to the fore, obscuring the measured value. The second mistake is often the result of the equipment being not capable of measuring the significant parameters right, so the reviewer has to base his/her suppositions on the insignificant parameters.

Here are a couple of examples to prove my point. An attempt to measure a monitor's contrast ratio with the help of a digital camera would be totally wrong without accounting for the camera's matrix's own noise, the camera's gamma correction (which is performed when saving into any format, except RAW), the backlight noise and other related factors. An error of the second type is an attempt to compare the speeds of two LCD monitors by measuring their making black-to-white transitions. The result would be useless, even if everything is measured with high precision and correctly, just because it is the response time on gray-to-gray transitions that's more important here, rather than the transition between the two limiting states.

Thus, in order to compare LCD monitors by their specs or by numbers you've got through some tests, it is of paramount importance to understand what these numbers mean as well as by which method and under which conditions they were arrived at.

This article is intended as a comprehensive description of the main parameters of LCD monitors as well as of the typical methods the manufacturers use to measure them. Then, since the bulk of the LCD monitor's properties is determined by the matrix it is based on, and since there are only four types of the matrix now employed (TN+Film, S-IPS, MVA and PVA), I am going to describe the distinguishing features of each type of the matrix.

Pixel Response Time

Response time seems to be the most popular characteristic of any LCD monitor as it is often the first parameter the customer pays attention to when evaluating a model.

As you know, the state of a pixel in the LCD panel changes when the angle of the liquid crystals is changed under the influence of the applied electric field. But liquid crystals is rather viscous stuff, so it takes them quite an amount of time to turn by the necessary degree – up to a few or more milliseconds. The following picture illustrates this process (the X axis shows time in milliseconds, and the Y axis – the brightness of the pixel; the pixel is changing its state from fully closed to fully open):

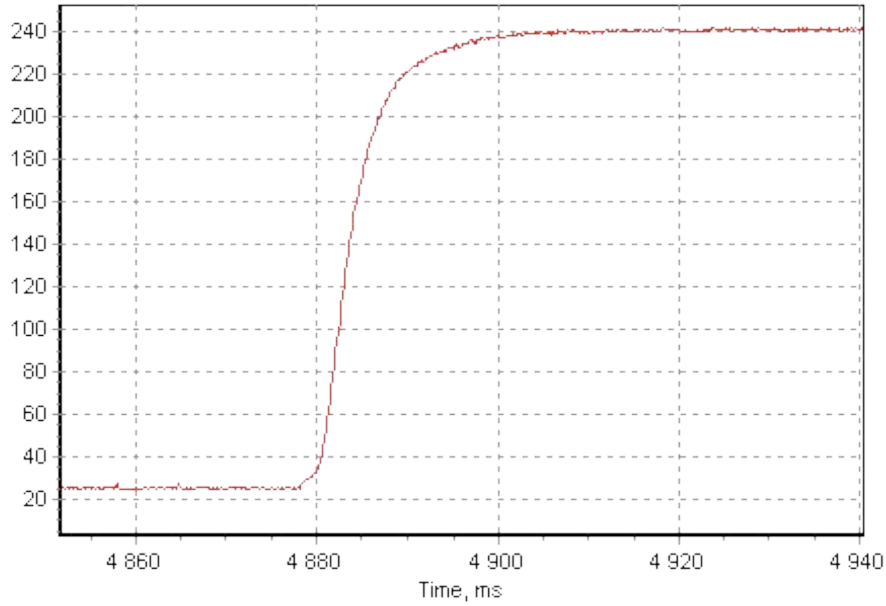

The manufacturers of matrices and monitors traditionally measure the response time as the total time it takes the pixel to change its state from black to white and to black again. More precisely, they measure the time it takes the pixel's brightness to grow from 10% to 90% and return to 10%. This convention – measuring within the 10%-90% range – is not any trick of the manufacturers'. It's rather a necessity since it is basically impossible to pinpoint the moment the pixel begins to shine or the moment it reaches 100% brightness due to noises present and limitations of the measuring equipment. So we should rather talk about the pixel's brightness being in or out of a certain interval, which is defined as 10% here.

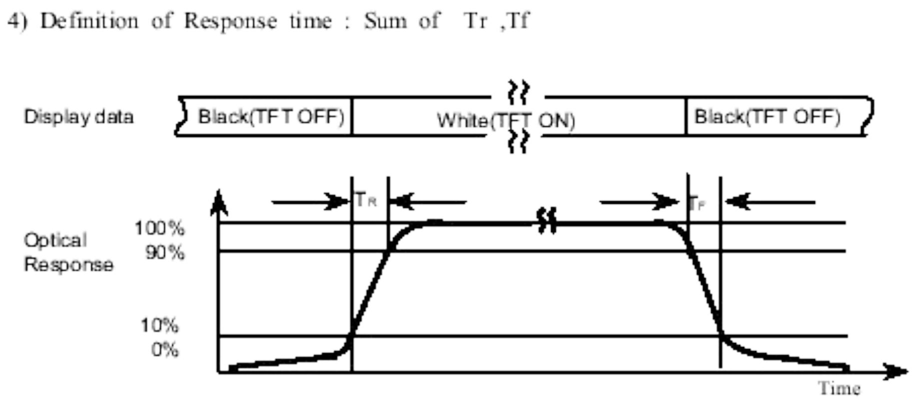

Unfortunately, this way of measurement gives rather a vague idea of how the monitor would work with dynamic graphics because the response time measured is the minimal response time of the matrix. Suppose we're interested in black-to-dark-gray (like in many "dark" games) rather than black-to-white transitions. The crystals have to turn around by a smaller degree in this case, but the speed of turning is proportional to the intensity of the applied electric field, and it is this field that determines the angle of turning – the smaller angle we need, the weaker electric field we must apply. Thus, we've got two contending trends – the angle of turning is getting smaller, but the speed of turning is getting smaller, too! In practice, the turning time (i.e. the monitor's response time) depends on the proportion of these trends. Measurements suggest that the response time is the smallest when the pixel's state (color) is transitioning from black to white. Other transitions take longer to be made, and the exact value depends on the type of the matrix (I'll describe this in detail later on, in the sections dedicated to each particular matrix type).

So, the monitor's specification says little as to how fast this monitor really is, since the dependence of the response time on the pixel’s initial and final states varies depending on the matrix types. We can’t even directly compare monitors with different matrices (for example, with TN+Film and PVA) relying just on the numbers provided by the manufacturer.

For a correct comparison we need a two-dimensional chart with a graph of the dependence of the response time on the pixel’s final state (i.e. the graph would show all possible black-to-gray transitions) or a three-dimensional graph (a surface) of the response time at all possible transitions, including those between two non-limiting states (i.e. between two shades of gray). Unfortunately, the manufacturers of matrices and monitors don’t usually offer this information.

As an example, the following graph shows the pixel rise time (Y axis) of a 25msec TN+Film matrix at transitions from black to shades of gray (X axis):

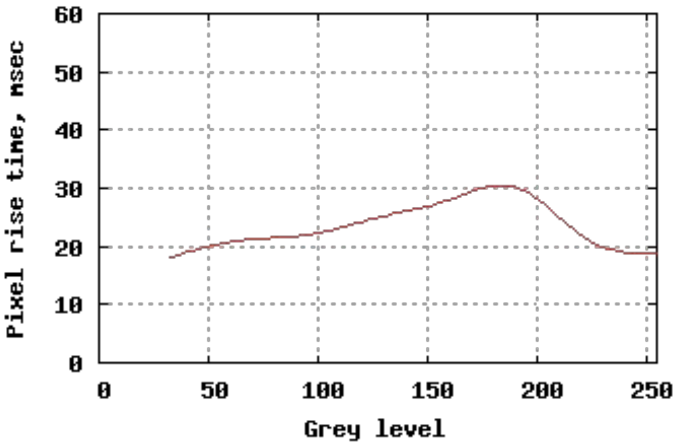

This peculiarity of LCD matrices will be first of all visible in dynamic games with a low-contrast image – the *ghosting* effect may be quite annoying although the formal responsiveness of the matrix may be specified to be low enough.

That said, even the time of a black-to-white transition is not a constant – it depends on the contrast setting of the monitor, and sometimes on its brightness setting, too. Generally speaking, the luminosity of a certain pixel is described as L = B + x*C, where B is a value that directly depends on the Brightness setting of the monitor, C is a value that depends on the Contrast setting, and x is the signal the pixel receives from the computer (x=0 corresponds to black color, while the maximum value of x – to white; I’m not counting gamma correction in, as I’m going to deal with it later on).

The level of contrast is controlled rather simply: the x signal from the graphics card doesn’t go directly to the matrix, but is first multiplied by the coefficient C. Thus, it is clear that white color that corresponds to the maximum angle of the crystals is only reached at the maximum contrast setting. If the contrast setting is below the maximum, the crystals will be turning around from “shut” to “almost open” and, due to the above-said reasons, this turning will take more time than what is specified by the manufacturer (who meant a “shut”-“fully open” transition).

So, here’s a general rule: reducing the contrast you increase the monitor’s response time.

The Brightness setting is somewhat better in this respect – the majority of LCD monitors control the brightness of the screen by changing the luminosity of the backlight lamps, so this setting has no relation to the matrix and doesn’t affect the response time in any way. Monitors with matrix-controlled brightness exist, though. For example, Sony’s products have an independent “Backlight” setting for adjusting the luminosity of the backlight lamps and a “Brightness” setting for the matrix. If you use the latter, the response time will be changing (measurements show that the response time may grow up considerably at low Brightness settings).

The *asymmetry* of the response time should be taken into consideration, too. I mean the ratio of the pixel rise and pixel fall times. For example, if we take two monitors with a total response time of 30 milliseconds, but one has 25/5msec pixel rise/fall times (typical for TN matrices) and another has 15/15msec times (this is typical for MVA and PVA matrices), they will display moving objects in a different manner. Particularly, thin black lines moving on a white background will seem thinner than should be on the screen of the first monitor.

The same lines will keep their thickness on the screen of the second monitor, although becoming a little lighter. The latter effect is better on the eyes, so an MVA matrix would seem faster than a TN one when scrolling text, their full response times being equal. This is another confirmation of my point that it's wrong to compare matrices of different types by their total response time only – you must also know the pixel rise to pixel fall ratio, at least.

Users are often asking if it is possible to measure the response time without any special equipment. Alas, it is not. The maximum you can do is to launch a dynamic game and evaluate the image subjectively ("it suits me" or "it doesn't suit me"). Users who are trying to estimate the response time by means of special tests that usually use a white square running on a black background (for example, Passmark Monitor Test) make one mistake, at least. The blurriness of the box is only indicative of black-to-white-to-black transitions, but it is not a crucial factor contributing to the monitor's responsiveness in the majority of cases, as I explained above. Moreover, users often carry over their experience with CRT monitors to LCD ones, evaluating the response time by the trail behind the moving square.

Due to the specifics of the CRT technology (the pixel lights up very fast, but produces an exponential graph when fading out), the box will have sharp edges and a barely visible and rather long trail (the "tail" of the exponential function that's describing the pixel's fading-out). LCD monitors won't necessarily have this trail, since they are usually described with quite another function, without the long tail. Some people even claim that modern LCD monitors have surpassed CRT ones, basing on this criterion. But they should take a look at the front and rear edges of the box – it is them that are the true indication of the matrix's responsiveness. The figure below illustrates the typical white square moving on a black background from left to right: the picture taken on a CRT monitor is above (sharp edges, but a visible trail), the one taken on a typical LCD monitor is below (no visible trail, but fuzzy edges):

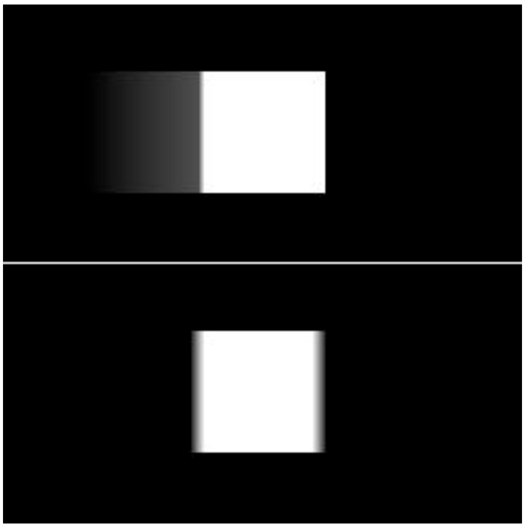

Viewing Angles

Yet another traditional problem of the LCD monitor technology is the angle you are permitted to look at the screen at. While the image on a CRT screen looks practically the same whatever your line of sight is, a small deflection from the perpendicular leads to a dramatic degeneration of contrast and color reproduction with many LCD matrices.

Meanwhile, all the manufacturers are now declaring seemingly acceptable viewing angles – the majority of monitors have 160° viewing angles both ways – horizontal and vertical. The problem is – like with the response time – how these angles are measured.

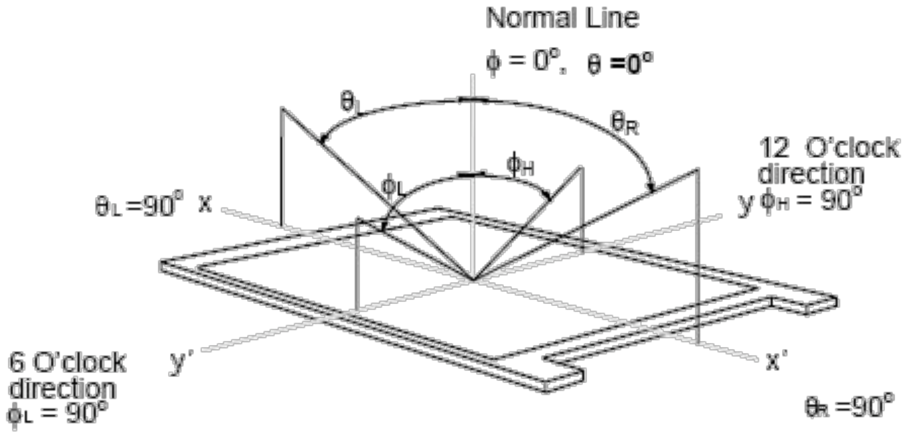

According to the current standards, the matrix manufacturers define the viewing angle as the angle relative to the normal dropped to the center of the matrix – when viewed at this angle, the contrast ratio in the center of the matrix degenerates to 10:1. This definition makes room for some pitfalls an unsuspecting user may get into.

First, it is known that the image becomes apparently distorted when the contrast ratio degenerates just in a few times, i.e. to about 100:1. In other words, the method the manufacturers employ is rather lax, and this is enough to be skeptical about the specified viewing angles, since you'll most likely notice the deflection from the ideal picture at much smaller angles. Furthermore, some manufacturers specify viewing angles for a contrast ratio of 5:1 rather than 10:1, thus making the angles of an inexpensive TN+Film matrix wider like 160/160 degrees instead of 150/140. Of course, this "modernization" gives nothing to the user – the matrix remains the same, although the specifications seem to say that the manufacturer of the monitor has started using new matrices with wider viewing angles, and only a small text at the bottom of the package tells you that it is only the measurement method that has actually changed.

Second, the contrast ratio is measured in the center of the screen, while the user sees the sides of the screen at a different angle than the center. For example, the next snapshot depicts a Greenwood LC521FT monitor and the camera is looking at it from below, at a small angle:

You are wrong if you think the monitor’s screen background is a gradient from black at the top to gray at the bottom. The monitor is actually displaying an absolutely even gray background of the color RGB:{128;128;128}, while this striking discrepancy in the brightness between the top and bottom of the screen is due to the insufficiently wide vertical viewing angle. However, the brightness in the center of the screen is much closer to the ideal than at the top of the screen (which seems almost purely black), and the standard method of measuring the viewing angles will give us a high enough contrast ratio as not to consider a vertical angle of 25 degrees as a limit (that’s the angle at which the camera is looking at the screen).

Three, the same snapshot illustrates another peculiarity of the angles the manufacturers specify. They usually declare a total angle to both sides from the normal (for the vertical angle – the above and below angles are summed up), while for this monitor (and for all TN+Film matrices for that matter) the above angle is much bigger than the below angle, and there’s a different artifact there – when you are looking at the screen from above, the bottom of the image first fades out and then, at a bigger angle, becomes inverted (white acquires a characteristic bluish tincture and becomes darker than light-gray shades). As a result, we have a big vertical viewing angle in the specification, but the top of the onscreen image becomes dark on even the slightest deflection of the monitor’s screen backwards.

Four, there’s the same story with the viewing angles as with the response time which is measured on black-to-white and white-to-black transitions only. The manufacturer specifies the contrast you see when your line of sight is perpendicular to the screen plane and says at which angles this contrast ratio drops down to 10:1, but you don’t know anything about how this contrast ratio is changing within these two points. The graph below shows the dependence of the contrast ratio on the angle of view for two different matrices (this is a purely theoretical example, not a result of any measurements – for illustrative purposes only):

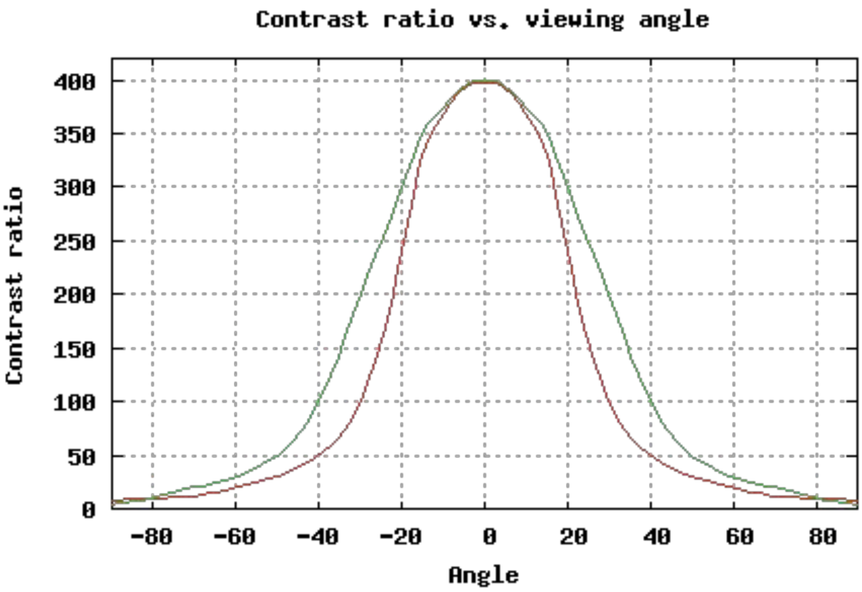

As you see, matrices with such curves will have identical specifications: the maximum contrast ratio is 400:1, the viewing angles (as measured by the contrast drop to 10:1) equals 160 degrees (80 degrees into either side). However, if you look at the matrices at an angle of 40 degrees, one will have twice more contrast than the other; in other words, the user of one matrix will have wider viewing angles than the other, although the specs of the two matrices are perfectly identical.

Five, it is only the reduction of the contrast ratio that's considered when measuring the viewing angles – distortions of color reproduction are not counted in. For example, here's a snapshot of the Greenwood LC521FT monitor with a purely white background onscreen:

As you see, white color is darker when viewed from a side, but it also acquires a strong yellow-brown coloring. Sometimes such a change of the color may be more visible than a change in the contrast, but the manufacturer doesn't count that in.

Six and the last, the manufacturers declare only vertical and horizontal viewing angles, whereas you may want to look at the screen from the right and above, for example. The following chart shows the dependence of the contrast ratio on both viewing angles (courtesy of Fujitsu):

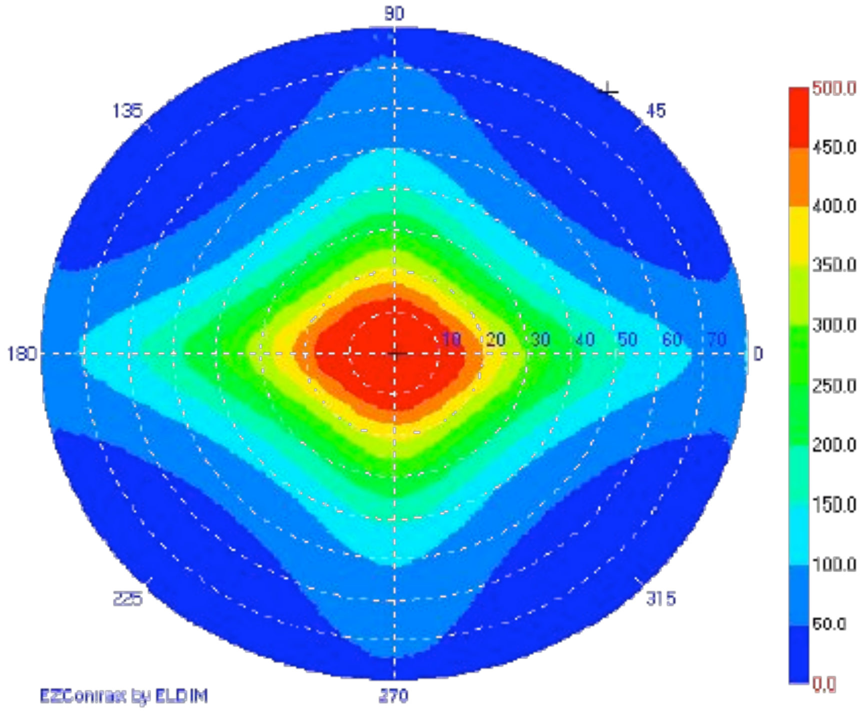

Thus, it turns out that vertical and horizontal viewing angles are the maximum ones, and they are put into the specs, while "diagonal" viewing angles are much smaller.

So, the item in an LCD monitor's specification called "viewing angles" does not actually say much about what you'll see on the screen of this monitor. Moreover, it involves so many reservations and depends on the particular matrix type that an extensive investigation is required to objectively evaluate the real viewing angles, so the only practical way for the user is the following: do not trust blindly to the specification, but examine the monitor models you're interested in with your own eyes.

Brightness and Contrast

Strictly speaking, it's not correct to talk about brightness and contrast of an LCD monitor, because the manufacturers often just copy the specification of the matrix the monitor is based on. However, while the monitor's electronics doesn't affect such parameters as response time and viewing angles, it is not so with brightness and contrast.

Let's first clarify the meanings of the terms: brightness is the luminosity of white color (i.e. the matrix receives the maximum signal) in the center of the screen, and contrast (or contrast ratio) is the luminosity ratio of white to black, also in the center of the screen.

The contrast-related problems are natural for LCD matrices due to the very principle of their operation. Unlike with the absolute majority of electronic display devices (CRT, electroluminescent and LED displays, OLED and so on), the LCD matrix is a passive rather than active element. It can't radiate light; it can only modulate the light that is passing through it. There's a backlight unit behind the LCD matrix, while the matrix only controls its own opacity, weakening the light to a certain extent. It achieves this by turning the polarization plane: the liquid crystals are placed between two co-directional polarizers.

So if the light hasn't changed its polarization plane between them, it goes through the second polarizer without any losses. But if the polarization plane has been turned by the liquid crystals, the second polarizer will stop the light stream and the corresponding cell will look black. But due to various reasons (the polarizers are not perfect, and the crystals cannot be perfectly positioned, too) it is impossible to stop all the light, so some portion of the backlight will pass through the matrix, highlighting slightly the black color of the monitor.

As I mentioned above, these measurements are performed by the matrix rather than monitor manufacturer, on a special platform where the matrix is attached to a test-signal source and the backlight lamps are powered with an electric current of a definite value. Thus, they arrive at some etalon values. In a real monitor, however, the electronics adds to the results, too. The electronics differ from the laboratory's signal generator and is also adjusted by the user to some extent who can control brightness, contrast, color temperature and other parameters of the device.

Thus, the real parameters of a monitor often don't comply with its specs. For example, if the monitor's electronics "highlight" black color a little (of course, it is an obvious defect, but quite a widespread one among some inexpensive models), the real contrast ratio proves much lower than specified.

Moreover, a contrast ratio of 500..700:1 as declared by many matrix manufactures for their products is still far from perfect, notwithstanding the seemingly high number. In fact, the monitor with such a contrast ratio cannot provide the really deep black color – black is going to look dark-gray in a dimly lit environment. If the real contrast ratio is only 200..300:1, it is easy to notice the backlight lamps to shine through the black color.

Some people try to justify the LCD monitor manufacturers saying that an extra-high contrast ratio strains the eyes. This is not true at all. You just can't have a "too low" level of black. Ideally, black should be not just low, but equal to zero – this would mean that the monitor is reproducing the pure black color, without any tricks. Of course, the specified contrast ratio would be very high in this case (the screen surface is not an ideal black body and will reflect external light, but I'm speaking about the specification now, which is measured without any external lighting).

There's one more myth: some say that the manufacturers are increasing the specified contrast ratio of their matrices by increasing the brightness of white color, keeping the intensity of black color on the same level. Thus, the specified contrast ratio grows up while the effective one remains the same since the user works with a contrast ratio that's comfortable for him/her rather than at the maximum possible. It's clear from the operational principle of LCD matrices that their brightness can only be increased by intensifying the backlight.

Given L is the backlight intensity, the level of white color is then L_w=L*n_w, where n_w is the pass-through coefficient of the open pixel (it is slightly below 1, since some portion of the backlight is still lost on its passing through the crystals and polarizers). The level of black is then calculated as follows: L_b=L*n_b, where n_b is the pass-though coefficient of the closed pixel (it's slightly above zero). The contrast ratio C is thus described by the formula: C= L*n_w/ L*n_b=n_w/n_b. The pass-through coefficients of the closed and open pixels depend on the characteristics of the matrix only, but not on the backlight intensity, so the specified contrast ratio of the matrix is in no way dependent on the backlight, but is only determined by the matrix's own properties. Thus, the manufacturer cannot better the specs of a matrix by increasing its brightness – this common opinion is not grounded at all.

Sometimes it is argued that the external lighting should also be taken into account. For example, a normal daylight in the room does contribute to the level of black. In this case, the "visual" contrast ratio is $C_{vis}=(L*n_w+L_{ext})/(L*n_b+L_{ext})$, where L_{ext} is the external lighting. As follows from the formula, higher L values lead to higher C_{vis} values. But I want to repeat once again that I'm talking about the specified contrast ratio of the matrix as measured by the manufacturer without any external lighting.

Besides the fact that the contrast ratio of the matrix is measured on a special test platform, rather than in the finished monitor (i.e. without counting in the specifics of the electronics of this monitor), the user is also free to tweak the brightness and contrast settings, thus affecting various parameters of the image. The way the settings affect the image depends on the realization of the controls in the particular monitor model.

First, it's not quite correct to say that the user changes brightness and contrast of the monitor with Brightness and Contrast settings, respectively, since a question arises – what brightness is adjusted and at the expense of what the contrast is changed. As I mentioned above, the brightness L of a pixel can be ideally described by the formula: L = B + x*C, where B is a value proportional to the Brightness setting of the monitor, C is a value proportional to the Contrast setting of the monitor and x is the signal the monitor receives from the computer (x=0 for black color and x=maximum for white color). It follows then that the Contrast setting is adjusting the brightness of white color (to be exact, of all shades of grays, but black color remains intact), while the Brightness setting affects both black and white colors at once.

In most monitors the Brightness control is realized through changing the intensity of the backlighting – this is logical enough. The employed cold-cathode fluorescent lamps allow doing this in two ways: by adjusting the discharge current in the lamp (this provides rather a narrow adjustment range since the charge loses stability after the current is far from the nominal) or by a pulse-width modulation of the power of the lamp at a rather small frequency (it's small considering the physics of the lamp's charge, but big enough for the eye not to notice it; this frequency is usually from 200 to 500Hz in practice). Pulse-width modulation is a widespread method of controlling voltages and currents. Its key point is such: depending on the desired voltage, the width of the applied impulses is controlled, while their frequency and amplitude remain the same. The average voltage is proportional to this width. The control process is illustrated by the oscillograms below:

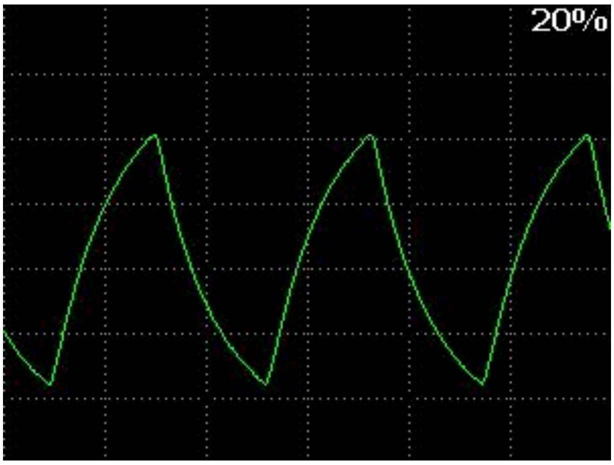

This signal was taken directly from the monitor's screen with the help of a photo-sensor, not from the lamps' power circuits, so the impulses had been already smoothed out by the afterglow of the lamps' phosphors, and it's clearly visible in the animated picture how the average brightness grows up. The distances between the peaks remain the same as the brightness is changing, so we evidently deal with nothing else but pulse-width modulation.

Brightness can also be controlled with the matrix – when the user chooses a higher brightness setting, the monitor adds a definite constant to the signal sent to the matrix. Alas, the contrast ratio degenerates with this method – the backlight lamps are always working at the power necessary to provide the maximum possible brightness, so when working at a small brightness, even if the added constant equals zero, the monitor will produce a higher level of black than a model that controls brightness with the backlight lamps. Suppose the brightness of black is described by the formula L_b=L*n_b, where L is the brightness of backlighting, and n_b is the pass-through coefficient of the closed pixel.

Suppose two monitors have the same matrix with a maximum brightness of 250 candelas per sq. m (and, accordingly, the same n_b coefficients), but we want to have a screen brightness of 100 candelas per sq. m. In this case, the monitor with the backlight-controlled brightness reduces the value L in 2.5 times compared to the maximum, but L remains the same in the monitor that controls the brightness through the matrix. Clearly, the level of black (L_b) is going to be 2.5 times lower in the monitor that controls the brightness with the backlight lamps.

Besides that, as I mentioned in the previous section, controlling brightness with the help of the matrix has a negative affect on the response time characteristic. These effects can be clearly seen in Sony's monitors that allow controlling the screen brightness either with the matrix (the "Brightness" parameter in the screen menu) or with the backlight lamps (the "Backlight" parameter).

But what screen brightness should you choose? It all depends on the task and the external lighting. The screen brightness should be from 70 to 130 candelas per sq. meter for working with text, while 200 and more candelas per sq. meter may be comfortable for games and movies. In contrast to LCD monitors, CRT ones typically have an operational brightness of 90..100 candelas per sq. m (models produced in the last two years feature extra brightness modes, made available by the technologies for high-precision focusing of the ray, but they are still suitable only for movies and games; LCD panels, with their ideal image sharpness at any brightness, have long surpassed them), but the contrast ratio of a good CRT monitor easily exceeds 1000:1, an unattainable peak for the majority of LCD monitors.

Irregular backlighting may also be seen with LCD monitors, mostly with low-contrast matrices. This is often perceived as light or dark stripes and spots (light spots may sometimes correspond to the locations of the backlight lamps), or as light stripes near the edge of the matrix – they occur if the matrix is too pinched in the case during the assembly process (I mean the metallic case of the unit, rather than the plastic bezel of the monitor, which serves decorative purposes only).

So, the conclusion to this section actually repeats what I said at the end of the preceding sections: we can compare two monitors on matrices of the same type by their specified contrast ratios (the higher specified contrast ratio will most probably result in a higher real contrast ratio), but we can't compare monitors on matrices of different types. Then, we also cannot say anything about the absolute (not relative, as in terms of "better-worse") contrast ratio basing only on the numbers provided by the monitor manufacturer.

Color Reproduction

The manufacturers usually utter one number when talking about the reproduction of colors on the screen of an LCD monitor. The number is either 16.2 or 16.7 millions of colors. However, there's a trap here, too. Many matrices nowadays – and almost all inexpensive ones – cannot reproduce more than 262 thousands of colors, as represented by 18 bits or 6 bits per each of the three basic colors.

The image on an 18-bit matrix looks depressingly bad without any additional measures. Such a matrix is in fact only suitable for office work and maybe for games. That's why the manufacturers realize the so-called Frame Rate Control (FRC), a method of emulation of the missing colors when the color of a pixel is changed slightly with every frame. For example, the monitor has to output the color RGB:{154;154;154}, and the matrix doesn't physically support it, but it supports the two neighboring colors, i.e. RGB{152;152;152} and RGB{156;156;156}. If we were outputting these two colors alternately with the frequency of the refresh rate, the similarity of these colors and the inertia of the human eye (which doesn't perceive flickering at a frequency of 60 hertz) as well as of the matrix itself (which is "smoothing" the moment when the colors are being switched) would give us what our eyes would perceive as some in-between color, i.e. the required RGB:{154;154;154}.

Of course, this is still an emulation, which cannot match the true "true color" reproduction, so such monitors are often specified to output 16.2 millions of colors – if you see this number in the specs, know that you are dealing with an 18-bit matrix. Unfortunately, the text "16.7 millions of colors" doesn't mean anything – many manufacturers write this in the specs of their 18-bit matrices, too.

More complex FRC mechanisms are also used in practice, accompanying the more traditional dithering (when the necessary color is formed by several neighboring pixels of slightly different colors), i.e. changing not just the color of a single pixel each frame, but of a group of four pixels, for example. This allows for a more precise reproduction of colors the matrix doesn't support, but the hitch remains the same – such matrices can hardly be considered "full-color" ones.

So, the quality of color reproduction of such matrices is largely determined by the quality of the FRC mechanism. There are two common problems: first, there appear transverse stripes in smooth color gradients. Sometimes it looks as if there are no FRC at all. This problem, however, mostly belongs to the first generation of "fast" matrices, although some stripiness can sometimes be seen in a modern model. The second problem is that the FRC algorithms may fail on some specific pictures (for example, a one-pixel grid, especially if it is combined with a smooth gradient), resulting in flicker – this flicker may be so strong as to make it impossible to work with the monitor. The latter effect can only be seen in the most inexpensive models, though.

We should also remember that the quality of FRC (and the side effects thereof) may depend on the brightness and contrast settings of the monitor (if brightness is regulated by the matrix, not by the backlight lamps) – in this case, the image may be flickering at certain settings only. Overall, flickering usually occurs with rather specific images, not preventing you from working comfortably with the monitor.

Besides the matrix color depth, there's another property that affects the color reproduction of an LCD monitor. It is *gamma compensation*. When talking about brightness and contrast above, I simplified the matter saying that the correlation between the input signal and the pixel's brightness is linear (L = B+x*C), but it is not actually such. This is a power dependence and it looks like this: $L=B+x^{gamma}*C)$, where *gamma* is some constant.

Gamma compensation can be said to have originated and to exist due to historical rather than technical reasons. In fact, the natural dependence between the input and output signal of the cathode-ray tube is close to a power dependence with an exponent of 2.5. Operating systems for the PC platform used to have no color management systems at all at first, so gamma=2.5 is traditionally considered the standard for the Wintel platform.

The Apple Macintosh, however, was traditionally employed for printing, processing photographs, performing color correction and so on, and the value of gamma was slightly adjusted on that platform – reduced to 1.8. Of course, for the user to see the original colors of the image on the screen, it should be first processed with the function $i=I^{l/gamma}$, where i is the resulting brightness, L is the original brightness of the picture and gamma is the gamma of the system this picture is being processed for.

Thus, the picture viewed by the user will be described by the following formula: $L=B+(I^{l/gamma})^{gamma}*C=B+I*C$. That is, the user will see the original i, but corrected with respect to the contrast C and brightness B of the monitor. Since the value of gamma varied between the platforms, pictures had to be compensated in a different way, and a picture prepared on a Mac would look too dark on the PC and vice versa – a PC-oriented image would be too light on the Mac. That's why about a decade ago Microsoft, Hewlett-Packard et al. came up with the sRGB standard, "A Standard Default Color Space for the Internet", which set the value of gamma to 2.2 (more precisely, the gamma curve in the sRGB specification is a combination of two independent functions, but with gamma=2.5 it can be described with an acceptable degree of precision with a single function).

Thus, sRGB-compliant images look equally well (or, as skeptics put it, equally bad) on Macs and on older PCs (with gamma=2.5). Right now the sRGB specification is de jure as well as a de facto standard, and modern monitors are generally calibrated for gamma=2.5.

You may be wondering why this gamma compensation is necessary from the technical point of view. The proponents of compensation say that it allows increasing the precision of reproduction of darks (of course, at the expense of lights): the human eye has a logarithmic characteristic of sensitivity, i.e. it more easily perceives a change of dark tone than a change of light tone of the same value. Thus, we can improve the precision of darks at the expense of lights. But a theoretical calculation says that gamma being equal to 2.2, a precision equivalent to 9-bit encoding is only achieved for 7% of the darkest color tones, and equivalent to 10-bit encoding for 3% (there's no sense talking about 11-bit precision of reproduction of darks, since such colors don't practically differ from pure black); meanwhile the color precision degenerates for 75% of lights.

This can be compared to the losses inflicted by saving an image in a middle-quality JPEG format (JPEG also brings in geometrical artifacts, but this is irrelevant for my topic). Looks like everything's all right, yes? We've improved the precision of darks at the expense of lights – we've got what we want?

Well, not exactly so. First of all, images are not of perfect quality themselves – being restricted by the capabilities of the camera, scanner or any other device you've made them with. The precision of darks is first of all determined by the noise in the camera's CCD or CMOS array (there can be numerous reasons for the noise – photon shot noise, read noise, dark current of the matrix and so on). So, the signal-to-noise ratio, even of high-quality cameras with cooled matrices employed for scientific purposes (in astronomy, spectroscopy, microbiology and so on), is 60..65dB (to reach that, at least two-step Peltier-based cooling is employed, complemented with active air cooling of the heatsinks; the resulting temperature of the CCD array being from -10°C to -40°C). This corresponds to about 10-bit precision (1 bit = 6.2dB). Ordinary cameras, up to professional ones, provide a SNR of 40..50dB at best, which corresponds to a precision of 7-8 bits. So, what's the sense of additional precision bits if the junior bit only contains the matrix's noise at the standard 8-bit precision?

Moreover, gamma compensation reduces the color-reproduction precision by itself – due to truncation errors as well as during subsequent processing of the compensated images. These distortions are most noticeable in dark areas of the image, although gamma compensation is intended to render them with a higher precision than a linear display. Anyway, we're not unlikely to see the industry to abandon gamma compensation in the near future. Too much equipment is designed to work with it.

But let us return to LCD monitors. I mentioned above that CRT monitors have an exponential display characteristic; with LCD ones it is close to S-shaped. In other words, to achieve the necessary exponential dependence on an LCD monitor, we need a correction table that translates the natural dependence into the required. Thus, the monitor's color-reproduction properties will also depend on the profile the manufacturer flashes into the firmware. The manufacturers find a kind of compromise between calibration of each particular monitor (that's unpractical with respect to the speed of the production line as calibration takes about a quarter of an hour or more) and a single-time calibration at the start of production of a new model (this is unacceptable from the point of view of quality – matrices from different batches may have different characteristics). But even a batch-specific calibration is no guarantee of the best result. The calibration curves for an Acer AL1715 monitor are shown in the picture below as an example (red, blue and green lines are the calculated curves for gamma=2.2; measured curves for the appropriate colors are given in black):

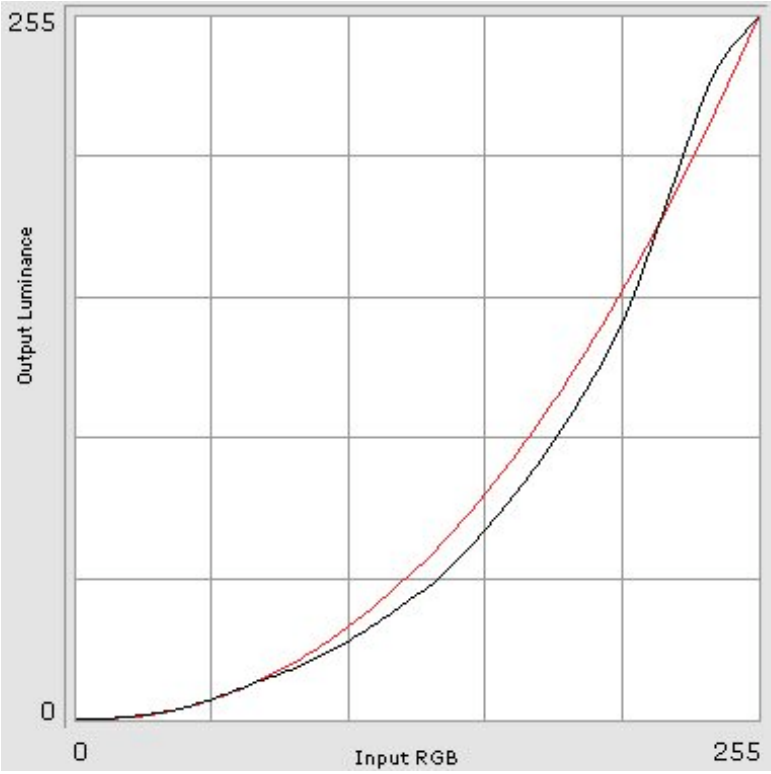

The graph reveals that the “native” S-shaped characteristic of the matrix is not fully compensated: the brightness swoops down in the middle of the range, and the lights, on the contrary, are displayed brighter than they should be. This deviation is rather small, and the home user won’t even notice it, but there can be worse cases:

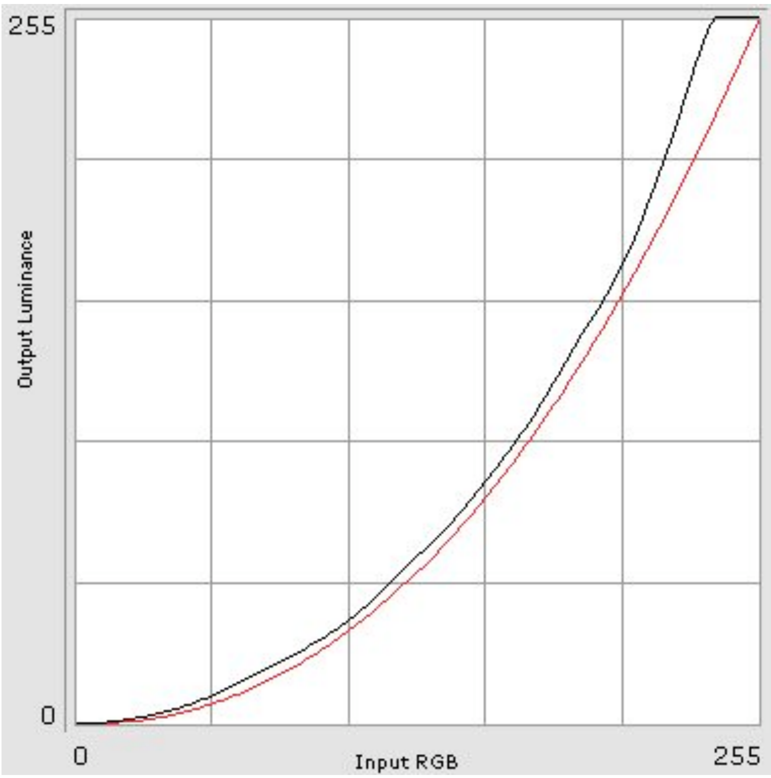

This graph is drawn for an e-Yama 17JN1S monitor. As you see, the manufacturer overdid it – the colors are saturated earlier than the rightmost point of the graph. In practice this means that, for example, the color RGB:{224;224;224} will be displayed as pure white rather than light-gray. In other words, the monitor doesn’t distinguish between some light tones, reproducing them all as white color. A similar thing may happen to dark tones when the monitor displays a dark gray as black. Sometimes you can even witness a paradoxical situation when, for example, the color RGB:{5;5;5} is darker than the pure black RGB:{0;0;0}. The quality of the color reproduction setup also depends on the brightness and contrast settings.

A properly set-up monitor reproduces the entire color range in a rather wide range of user-defined settings (for example, for the white color brightness from 50 to 150 candelas per sq. m), but some models can only work normally at the default settings. If you move the brightness or contrast controls just slightly off the ideal position, you lose either dark (when increasing the brightness) or light (when increasing the contrast) tones.

You can also notice that the curves of different colors do not coincide on the diagram – some go lower, so go higher… This leads to deviations of the image tonality (in other words – of the color temperature) from the desired, and as the difference between the curves of the basic colors may vary throughout the dynamic range, the deviation of the color temperature will also vary depending on whether we're outputting light or dark tones.

This makes it impossible to correct the tonality precisely with the monitor's standard settings – i.e. with the independent RGB controls – as having achieved the necessary color balance at one part of the range (for light grays, for example), you will find you've worsened this balance at another part of the range (dark grays). The only way out of this predicament is calibration of the monitor with a hardware calibrator that creates an ICC profile basing on all the peculiarities of the monitor's color reproduction.

Well, you can try to create a profile manually, as special software for that purpose comes enclosed with many monitors, but manual calibration isn't always a success. Sometimes the color reproduction of the monitor is originally so badly set up that even a hardware calibrator is useless.

I said "color temperature" above, so let's consider this characteristic in more detail. The color temperature determines the tonality of the onscreen image. The lower the temperature – the warmer the colors are (that's the way we're all perceiving colors – the spectrum of hotter bodies is perceived as cold). The need of such temperature arises as the human eye doesn't have any universal white color, which would always be white. The eye adjusts itself to a specific range depending on the environment.

You can see this in the following experiment: take a cell phone with white highlighting of the screen and put it on a sheet of white paper. Then take a look at it at normal daylight and at lamplight. In the first case, the phone screen would seem white or even slightly yellow; in the second case it will have a bluish tone, since the eye will base its "balance of white" on the sheet of paper and the color of the sheet of paper is determined by the spectrum of the light source. The spectrum of lamplight is bluish, of daylight – yellowish. In the same way the tonality of the white color on the screen of a monitor will vary slightly depending on the external lighting, only in a slighter degree, since the monitor's screen is bigger than the phone's and the eye is adjusting itself to the screen, too.

That's why it is recommended to set a color temperature that would produce pure white color without any additional tones – under the given external lighting.

Color temperature is measured in Kelvins (K) and equals the temperature of the ideal black body radiating the same spectrum. There are three commonly used values. Printing and photography often involves 5500K temperature (this value was introduced by folks from Kodak under the name of "daylight". There is a joke that in the real world this value is the color temperature of midday sunlight near the offices of this company). For computer-processed images a color temperature 6500K is used (it corresponds to bright sunlight in a clear sky, while a slightly clouded sky has a color temperature of 6500-7000K). Sometimes, but quite rarely, the color temperature is set to 9300K giving out a slight bluish hue (an analog from nature – it is the color temperature of a thin shadow on a bright day).

Talking about artificial light sources, the ordinary bulb has a color temperature of about 2000K, powerful studio incandescent lamps and daylight lamps with "warm" phosphors – about 3000K, daylight lamps with "cold" phosphors – 4000K.

The aforementioned sRGB standard recommends a color temperature of 6500K (D65), and this is the default temperature of many computer monitors. But still, the color temperature of LCD monitors has some specific features. First, as I mentioned above, color temperature may vary considerably between different tones of gray. On a well-calibrated monitor the temperatures of white and 50% gray differ by a score or two of degrees, but for the majority of monitors this difference is as high as 500K and more (up to a thousand degrees for high temperatures). It means that if white looks really white, light-gray will look bluish (it is the temperature of gray that's usually higher, but there are exceptions, too).

Second, CRT monitors allow adjusting the color temperature rather smoothly along the entire range, from 5000K to 9300K, with a stepping of 50..100K, but LCD monitors usually have three or four temperature values for the user to select among them. For a smooth adjustment of the temperature, you can only use the independent controls for the three basic colors (R, G and B), which is rather unhandy and requires much experience to achieve an appropriate result.

Moreover, the temperature of a CRT monitor is usually changed very neatly, while LCD monitors may produce some artifacts at different temperatures. For example, at a reduced temperature, the screen gets a strong pinkish or greenish hue. On a temperature increase, gray color may become so bluish that the calibrator just can't measure its temperature…

One more parameter that slightly affects the quality of color reproduction is the so-called gamut range. It is a known fact that the human eye perceives light with wavelengths from 390nm to 760nm, different wavelengths being different colors from violet to red. Display devices, however, reproduce a much narrower range of colors.

This is most obvious from the so-called CIE color diagram. The CIE color space describes a particular color tone with three coordinates: two determine the color and the third determines the brightness. The diagram has only two (color) coordinates and outlines the color space perceived by the human eye (the borderline of this space corresponds to pure colors, from a monochromic light source, while the internal area is represented with colors of a more complex spectrum):

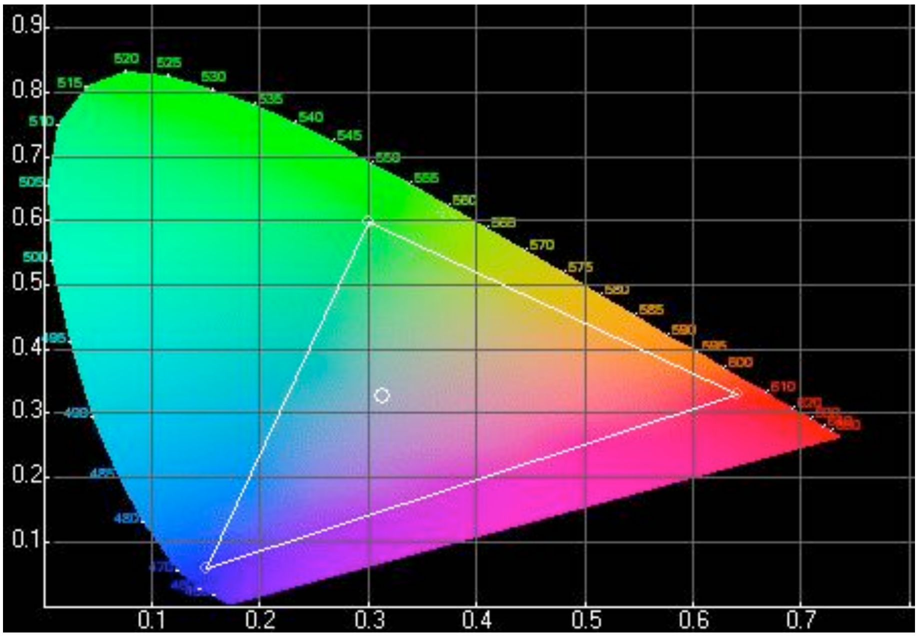

The white triangle on this diagram bounds the area that corresponds to the color range of sRGB-compliant devices and the white dot in the middle of the triangle is the white color point at a color temperature of 6500K. As you see, the range of sRGB colors is small compared to the range the human eye can perceive, so many colors are left off when the image is created (for example, an sRGB monitor cannot reproduce a single pure color at all) – they are just replaced with the colors that lie on the edges of the triangle. Sometimes this leads to noticeable visual artifacts, of course.

Besides sRGB, there are other color spaces that describe much wider color ranges. First of all, I mean AdobeRGB and NTSC as they are more often referred to in monitor-related discussions. While the sRGB color space embraces only 35% of the colors perceived by the human eye, AdobeRGB covers 50.6% of them, and NTSC – 54.2%. Well, almost all monitors today comply with the sRGB standard, and since the gamut range is determined physically by the characteristics of the backlight lamps and the color filters of LCD monitors and phosphors of the CRT monitors, there is a small variation between different models (even, between CRT and LCD ones) when it comes to color reproduction. Other factors influence it much more.

Exceptions to this rule are some professional-class CRT monitors. Not long ago, NEC-Mitsubishi introduced a 22" RDF225WG model with a gamut range of 97.6% of the AdobeRGB space thanks to a better green phosphor (as the above-shown diagram suggests, the green color is the most problematic). In the next year LCD monitors are expected with a gamut range wider than that of the AdobeRGB standard, although by 1% only. Unlike in the CRT model from NEC-Mitsubishi, the LCD models will achieve it by using white LED highlighting instead of traditional mercury-based cold-cathode lamps. The lamps have an irregular spectrum, consisting of several peaks and stripes, while light-emitting diodes produce a more even light, which also fits well into the passbands of the matrix's light filters, improving the image.

Right now such monitors are under development. Of course, they will be oriented on professional applications, with appropriate prices. LED highlighting will come to home monitors eventually, if LCD matrices haven't given way to OLED matrices (which don't require backlighting at all) by that time. Well, I wouldn't pay much attention to that. The fact that the first CRT monitor with a gamut range close to AdobeRGB was only released in the last spring, while the color correction industry is a few decades old, indicates that gamut range is not the determining color reproduction factor.

As you see, a good reproduction of color is a hard and complex task. When talking about contrast, viewing angles or response time, I could say that the value declared by the manufacturers described but some of the characteristics, but here the text "16.7 millions of colors" means nothing at all with respect to color reproduction.

Safety

One topic that is often discussed when comparing CRT and LCD displays is that of safety. Although it is not directly connected to image quality, I'll dwell upon it, too. And I'll be talking about how dangerous CRT monitors are rather than about the safety of the LCD technology.

First of all, the monitor can be dangerous through radiation. I heard complaints that CRT monitors emit various types of radiation, from alpha particles to gamma radiation. Let's browse through them one by one.

1) Alpha radiation is a stream of helium-4 nuclei. The cathode-ray tube has no helium to start with, so this radiation can only be the result of some nuclear reactions, which is absurd.

2) Next goes beta radiation, which is a stream of electrons. The cathode-ray tube does have a stream of electrons accelerated to energies of about 25,000 electron-volts (since the operational voltage of the kinescope is about 25 kilovolts), but there is one centimeter of glass between the electron stream and the user – no electron can break through it.

3) Hitting the phosphor, and being stopped by it, the electrons give only part of their energy to it (the phosphor shines due to this energy, by the way), while the remaining energy goes to the so-called bremsstrahlung (deceleration radiation), the radiation emitted by electrons slowed down in matter. The spectrum of the deceleration radiation stretches from zero to the maximum energy of the decelerating particles, i.e. its maximum energy may be 25 keV, which corresponds to very soft X rays (the term *soft radiation* is applied to a radiation with quantum energy up to a hundred thousand of electron-volts due to its relatively low penetrating property; for comparison, modern X-ray photography works with energies up to 150 keV).

4) Next to the X-ray radiation goes gamma radiation with quantum energies of tens of megaelectron-volts. It is created during nuclear fission reactions or as deceleration radiation when particles of energies of tens of MeV are being slowed down. Evidently, there are no nuclear reactions going on inside the monitor case, and there are no million-volt voltages in it, so it cannot emit gamma radiation.

So, the only type of radiation that may worry you is 3) soft X-ray radiation with an energy of about 25,000 electron-volts. To oppose it, the front glass of each cathode-ray tube (deceleration radiation is always directed forward, in the direction of movement of the particles it is created by) has lead and other metals (with lead alone, the glass eventually becomes blurry), which effectively stop this type of radiation. Thus, the X-ray radiation of the monitor originally has a very soft spectrum and after it has passed through the lead-doped glass of the tube it doesn't exceed the natural background radiation.

The second thing with respect to safety of the monitor, besides radiations is as follows: old monitors could be uncomfortable – not really dangerous – due to a powerful electrostatic field at the front side of the tube, but all monitors of the MPR-II standard (not mentioning the series of TCO standards) have a conductive grounded coating on the tube, which reduces the electrostatic field to acceptable values.

Thirdly, the deflecting system of the CRT monitor produces rather a powerful electromagnetic field. This system, however, is located at the root of the tube, i.e. rather far from the user. Moreover, it is covered with a protective metallic screen in good modern monitors – in other words, you can only hurt yourself by sitting all days long with your head resting on the rear or side panels of the monitor. Nothing is going to happen to you if you're sitting in front of the screen, as normal people do.

Thus, it seems like the modern CRT monitor has nothing to hurt your health with. All the ascribed radiations are either non-existent or dumped to the safe level. By the way, the users usually complain of headaches, eye strains, and worsening eyesight. But I guess if a person received such a dose of radiation as to feel a headache or eye-strain, he'd better think about making its will, rather than about replacing the monitor.

So, the discomfort from the monitor is due to the bad quality of the image rather than to any radiation. The most common problems are bad focusing, bad convergence, or a graphics card that's outputting a "fuzzy" image on the screen. It is the insufficiently sharp onscreen picture that makes your eyes sore. Sometimes, the monitor is just set up incorrectly: the contrast is at the maximum, while the brightness control is at zero, or vice versa (the brightness is too high, while the contrast is low). Working in full darkness is a strain for your eyes, too…

That said, the main advantage of LCD monitors is not the lack of radiations (there are no electron beams or deflection systems or high voltages at all, save for the power voltage of the backlight lamps), but rather the lack of such notions as "convergence" or "focusing" in the LCD world – LCD monitors are always outputting an ideally sharp image, save for some inexpensive models that have problems with setting up for the analog signal from the graphics card. This, however, cannot eliminate the problem of setting up the brightness and contrast correctly: excess brightness or insufficient contrast of an LCD monitor will hurt your eyes much like those of a CRT monitor do.

This is the end of the theoretical part of this article. Let's now go over to practice – to the various types of LCD matrices available in the market and their typical characteristics.

TN+Film Matrices

TN matrices are the oldest type of the LCD matrix, tracing its origin to the times of passive matrices. They have acquired the word Film into their name early – this additional film improves the viewing angles. Nowadays all matrices of this type have this film, so it is not necessary to mention it – speaking about modern matrices we can use terms "TN" and "TN+Film" interchangeably.

TN matrices never had the very best parameters. Their poor color reproduction was their main drawback – it was so specific that you had to get used to it even in office applications, not to mention processing photographs. That's why some people claimed an imminent replacement of the TN technology with other matrix types, first IPS, then MVA, but reality proved different.

The name of the technology – Twisted Nematic – comes from the way of organizing liquid crystals in the panel: when voltage is applied, the crystals curl up into a helix whose axis is perpendicular to the panel's surface. Regrettably, the shape of the helix is somewhat irregular (the borderline crystals are not exactly parallel to the surface, but are rather at a small angle to it), and the optical properties of the helix vary greatly depending on the line of sight – along its axis or at an angle. The first drawback prevents TN matrices from having a good contrast ratio, while the second drawback – from having wide viewing angles.

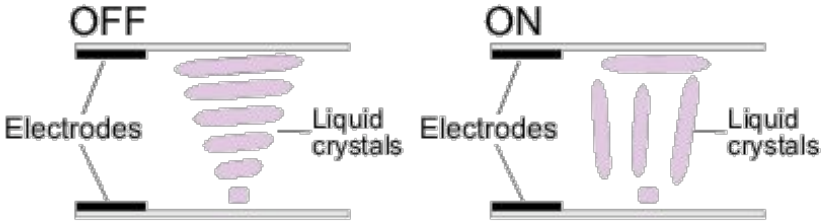

The TN technology got a second wind on the arrivals of matrices with a response time of 16 milliseconds. They were the only matrices at that time that you could specify such a low response time for, making a solid foundation for the marketing departments to build upon – they all started touting "unbelievably rapid matrices". It's always better to emphasize one parameter in an advertising campaign – one that the user intuitively understands (or thinks that understands) – and write it in bold big letters on the product's package. Craig Barrett once put this idea down in a very laconic form, talking about the successful sales of Intel's processors: "They buy the megahertz". The clock rate was an intuitively comprehensible parameter for people who bought central processors. They thought it determined the processor's speed, and it took AMD a lot of time and effort just to shake this common notion. Analogously, the response time became (or was made by the marketing departments of the manufacturing and selling companies) a parameter determining the quality of the matrix for the end customers.

Besides that, the TN technology is the cheapest of the available LCD matrix production technologies, so monitors with TN matrices could be sold at a lower price than their competitors. This killer combination – low price and the intuitive parameter – had a devastating effect on all other matrix types. Drawing the parallel with central processors further – imagine Intel starting to sell its many-gigahertz processors at much lower prices than the competitive products from AMD. AMD wouldn't stand this competition long, I guess.

But this situation became reality two years ago in the LCD monitor market. The manufacturers rolled out their TN matrices and TN-based monitors which cost less than the competitors on IPS and MVA matrices and were "better" than them (better in one respect only – in the response time). As a result, all 17" LCD monitors currently available, besides just a few models like Samsung's PVA-based monitors and Iiyama's H430S with an S-IPS matrix, are built on TN matrices. The TN technology is poised to conquer the market of 19" monitors, too. Until now, this market has only been safeguarded by a total lack of large TN matrices.

Regrettably, the response time is not the true indication of quality. Actually, the response time of 16 milliseconds was obtained with the help of a trick, made possible by the measurement method. As you remember, the response time is only measured on the matrix's switching from black to white and vice versa, but now take a look at the following graphs, which show the time it takes to make a transition from black to a shade of gray:

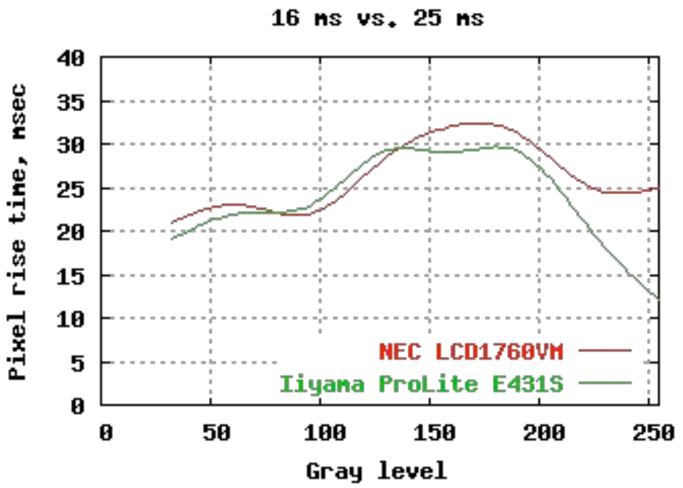

This chart displays the results of two monitors: NEC LCD1760VM with a response time of 25msec, and Iiyama ProLite E431S with a response time of 16msec. You see that the two graphs nearly coincide, save for the black-to-white transition where the 16msec matrix suddenly has a great advantage. This is not a unique feature of the particular monitor model – all TN matrices with a response time below 25msec draw a similarly-shaped graph. Of course, they have a slightly lower response time on black-to-gray transitions, too. The maximum response time of 12msec matrices is no more than 25 milliseconds, but I think if the 25msec matrices were developing further, they would reach the same result, save for the black-white transitions.

Samsung even achieved a black-to-gray response time of less than 20msec in the SyncMaster 710T, but this is a single case where the maximum response time closely complies with the specification. In all the other matrices tested in our laboratory, the advantage of fast matrices on black-gray transitions is small (2-3 milliseconds compared to previous-generation matrices). In other words, you are going to find that a 16msec matrix is not 1.5 times faster than a 25msec one, and a 12msec matrix is not 1.33 times faster than a 16msec one – the difference is smaller.

Anyway, the steady improvement of the response time, although not as big as you might guess from the numbers the manufacturers are specifying, is a good thing overall. Right now, 25msec matrices are completely ousted out of the market, and 16msec models are ruling it (I mean the TN+Film technology only here). And 12msec matrices are coming up, followed by 8msec ones. But still, even with this response time, LCD matrices have a long way to CRT monitors. For the ghosting effect in moving images to vanish completely, a response time of 4msec is necessary, and in the whole range of colors rather than on black-to-white transitions only.

Above, I was expressing my complaints about the marketing folk from the manufacturing companies emphasizing one matrix parameter – its response time – silencing the rest of them. Now let me correct this situation.

First, the matrix has viewing angles. The problem of the early 16msec matrices was in awfully small viewing angles, which prevented the user from working normally. Even sitting motionlessly and looking at the center of the screen perpendicularly you could see that the top of the screen was much darker than its bottom, and colors were yellowish at the sides. By the way, this thing – a strong darkening of the image when it is viewed from below – unmistakably distinguishes a TN matrix from the other types, which have no such defect.

Of course, some improvements have been made since. The horizontal viewing angles are now wide enough for you to sit before the screen with your friend without complaining at the “impure” white color; the vertical angles aren’t provoking any big discomfort too, although the vertical irregularity of the screen brightness is still visible even in the best samples. Unfortunately, the manufacturers of monitors on TN matrices, trying to match the competitive matrix types in this parameter at least on paper now often declare the viewing angles measured by the fall of the contrast ratio to 5:1 rather than to 10:1. This way TN matrices acquired 160-degree angles into their specs, without any real improvements. I want to warn you once again and remind you the above-described method of measuring the viewing angles. A true specified angle of 140 degrees doesn’t mean that its defects are only visible “when you are looking at it from under the desk” or “dance before it when working” as some users think, because *visible* distortions of the image appear much sooner before you reach those specified angles. The number “140 degrees” means you’ll see *strong* image distortion when looking at the screen at such an angle. For example, the vertical irregularity of brightness is visible in a TN matrix whatever position you choose before the monitor, so if this matters to you, a monitor on a TN matrix would be the worst choice you could make.

Then, the contrast ratio of TN matrices is not ideal, either. Although many manufactures declare a contrast ratio of about 500:1, the real contrast ratio seldom exceeds 300:1, some rare samples getting as high as 400:1. In practice this means that you can’t get a good black color on a TN-matrix monitor, and the black background on the screen will look highlighted if you’re in a dim room (for example, when you’re watching a movie). I should note, however, that the contrast ratio of matrices depends heavily on the manufacturer. For example, the latest matrices from Samsung have a standard contrast ratio of 300…400:1, while matrices from Chunghwa Picture Tubes (CPT) often have a deplorable contrast, so monitors on them can only be recommended as inexpensive office devices.

By the way, another drawback of TN matrices is that if a thin-film transistor crashes, there appears a bright dot on the screen, as pixels in TN matrices let light pass through in their inactive state. Such bright dots are more annoying than black pixels, especially if you’re intending to use the monitor at home, i.e. in the evening, for watching movies or playing games.

Next, the color reproduction of matrices of this type is far from perfect, too. All fast matrices are 18-bit, i.e. they employ Frame Rate Control to display all 16.2 millions of colors, but the colors of TN matrices are anyway bad – they are faded, unimpressive, rather far from natural ones, which makes TN matrices a bad choice for working with colors even amateurishly.

Thus, the small response time turns to be not just the main, but the only advantage of TN matrices. All of their other parameters are rather average. Monitors with matrices of that type suit for playing games and watching movies and for routine office work, but professionals may want to consider other types of the matrix. Unfortunately, this means limiting yourself with monitors with a diagonal of 19” and bigger, since the majority of 17” models have TN+Film matrices now.

IPS Matrices

The IPS technology was developed by Hitachi in 1996 to solve the two plagues of TN-matrices: small viewing angles and low-quality color reproduction. The name – In-Plane Switching – comes from the crystals in the cells of the IPS panel lying always in the same plane and being always parallel to the panel’s plane (if we don’t take into account the minor interference from the electrodes). When voltage is applied to a cell, the crystals of that cell all make a 90-degrees turn. By the way, an IPS panel lets the backlight pass through in its active state and shutters it in its passive state (when no voltage is applied), so if a thin-film transistor crashes, the corresponding pixel will always remain black, unlike with TN matrices.

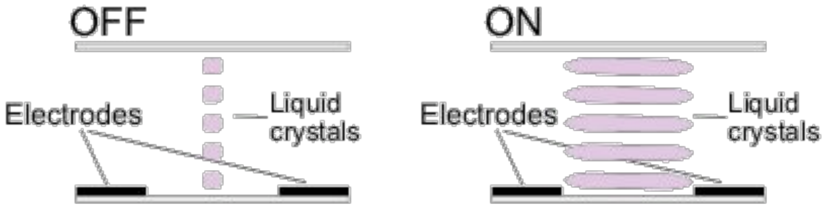

The figure above shows that IPS matrices differ from TN ones not only in the structure of the crystals, but also in the placement of the electrodes – both electrodes are on one wafer and take more space than electrodes of TN matrices. This leads to a lower contrast and brightness of the matrix.

The original IPS technology became a foundation for several improvements: Super-IPS (S-IPS), Dual Domain IPS (DD-IPS), and Advanced Coplanar Electrode (ACE). The latter two technologies belong to IBM (DD-IPS) and Samsung (ACE) and are in fact unavailable in shops. The manufacture of ACE panels is halted, while DD-IPS panels are coming from IDTech, the joint venture of IBM and Chi Mei Optoelectronics – these expensive models with high resolutions occupy their own niche, which but slightly overlaps with the common consumer market. NEC is also manufacturing IPS panels under such brands as A-SFT, A-AFT, SA-SFT and SA-AFT, but they are in fact nothing more than variations and further developments of the S-IPS technology.

S-IPS panels have gained the widest recognition, mostly due to the efforts of another joint venture LG.Philips LCD, which is outputting rather inexpensive and high-quality 19” and 20” matrices. The price – you can buy a 19” LG L1910S monitor on an S-IPS panel of the latest generation for slightly more than $600 – is an important achievement since IPS matrices have long been the costliest, and this fact impeded their development greatly.

Besides the high price, the response time was among the serious drawbacks of the IPS technology – first panels were as slow as 60msec on the “official” black-to-white-to-back transitions (and even slower on gray-to-gray ones!). Fortunately, the engineers dragged the full response time down to 25 milliseconds lately, and this total is equally divided between pixel rise and pixel fall times. Moreover, the response time doesn’t greatly grow up on black-to-gray transitions compared to the specification, so modern S-IPS matrices can challenge TN ones in this parameter.

The following diagram compares the pixel rise time of the 16msec TN+Film matrix of a NEC LCD1760NX monitor and the 25msec S-IPS matrix of an LG Flatron L1910S. As you see, the graphs are much like each other:

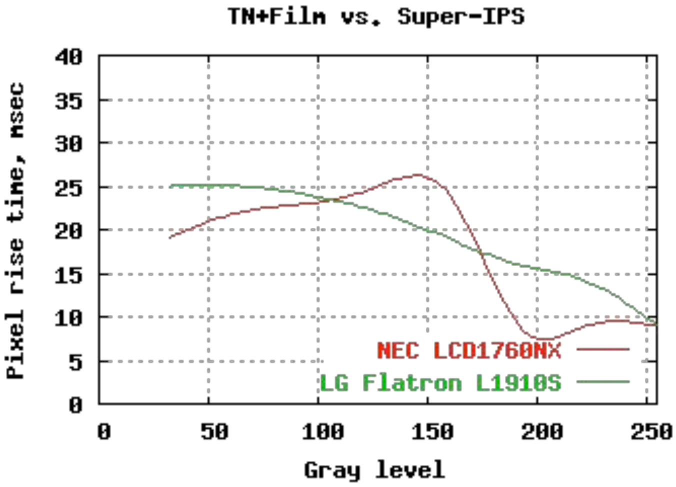

As for advantages, the IPS technology has always been better than TN+Film in terms of color reproduction and viewing angles. In fact, S-IPS matrices leave no chance to other LCD technologies in the color-reproduction quality. They have soft and pleasant colors, which are natural and close to high-quality CRT monitors. That's why all LCD monitors for professional work with color are based on S-IPS matrices, starting from relatively inexpensive to hi-end models of the Eizo ColorEdge series with integrated tools for custom hardware color-calibration.

The viewing angles are a treat after TN matrices: you can't notice any distortions of the image, sitting in front of an IPS matrix. There's only one specific defect – when you're looking at the screen from a side, black color acquires a characteristic violet hue (by the way, this defect allows telling an IPS matrix from any other), but the manufacturers are improving on this. In most cases, this is an insignificant defect anyway.

The only real problem of the S-IPS technology is the low contrast ratio (about 200:1, like that of an average TN+Film matrix). In means you see a dark gray instead of pure black. That's not noticeable at daylight, but if you're working in a dimly lit room, you may be disappointed at the highlighting of the black color (coupled with the characteristic violet hue when you're viewing the screen from a side).

Alas, but due to the reasons explained in the previous section IPS matrices are fully driven off the market of 17" monitors (save for the Iiyama H430S model, whose high response time only makes it suitable for work with static images), so the users who are not satisfied with the low image quality of TN+Film matrices but willing to have a low response time have to consider 19" models. Fortunately, you can choose here as quite a number of popular models are based on S-IPS matrices: LG Flatron L1910S and L1910B, NEC MultiSync LCD1960NXi (don't confuse it with the LCD1960NX model, which is based on a matrix of another type), Philips Brilliance 190B5 and many others.

As concerns their possible applications, monitors with S-IPS matrices are the only wise choice for serious work with color. Besides that, these matrices seem to be more balanced than others with their excellent viewing angles and low enough response time, so they will suit people who are choosing a monitor for games, movies and the Internet. TN+Film matrices that have recently entered the 19" monitor market, have a better response time but worse viewing angles (140 degrees only), and thus can't be called a good choice for a big-diagonal monitor.

MVA Matrices

The MVA technology (Multidomain Vertical Alignment) was developed by Fujitsu in 1998 as a compromise between TN+Film and IPS technologies. On the one hand, MVA provided a full response time of 25 milliseconds (that was impossible with IPS then and not easily achievable with TN), and on the other hand, MVA matrices have viewing angles of 160..170 degrees, and thus can compete with IPS in that parameter. Besides that, MVA provides for a much higher contrast ratio than TN or IPS.

MVA’s precursor, the single-domain VA technology, had been developed by Fujitsu two years earlier. Small viewing angles were its main disadvantage. Take a look at the figure below – a half-open pixel is on the right. If you look at such a pixel from above, everything will be all right: the crystals will be at an angle of 45 degrees to your line of sight and the pixel will look gray as it should. But if you look at it from the right, you’ll see the same crystals under a right angle, which is going to appear as white color. A look from the left would mean your line of sight being parallel to the crystals – you’ll see black color. Thus, VA matrices didn’t just have small viewing angles – the specific effect from a high viewing angle also depended on which side you have deflected your head to.

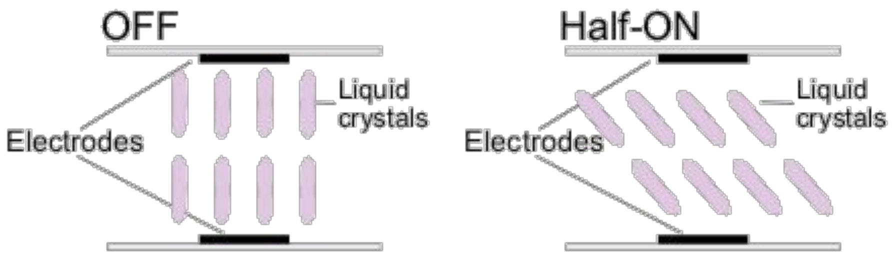

This problem was solved by dividing each pixel into domains which worked synchronously. Crystals in the domains are oriented differently, so if one domain lets light pass through, the neighboring domain will have the crystals at an angle and will shutter the light (of course, save for the display of white color, in which case all the crystals are placed almost in parallel to the matrix plane). Like with IPS matrices, an inactive pixel doesn’t pass light through, so defective pixels look as black dots on MVA matrices.

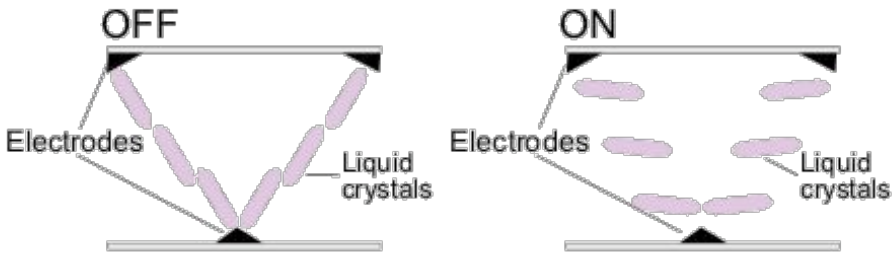

For a few years analysts were forecasting a bright future and a big portion of the market for MVA matrices: TN matrices should have been ousted into the low-end market sector as they were originally cheaper than MVA ones, while expensive S-IPS matrices should have reigned in the top-end sector, leaving the market’s mainstream to the MVA technology. Those forecasts never came true, though. Besides the above-described effect from the onslaught of cheap 16msec TN matrices, MVA matrices were rather too expensive for their responsiveness. No, I’m not mistaken: despite all the claims about an excellent (by the standards of that time) response time of 25 milliseconds, MVA matrices proved to be among the slowest. The trick is in the measurement method, like it nearly always is with the response time characteristic. Here’s a graph:

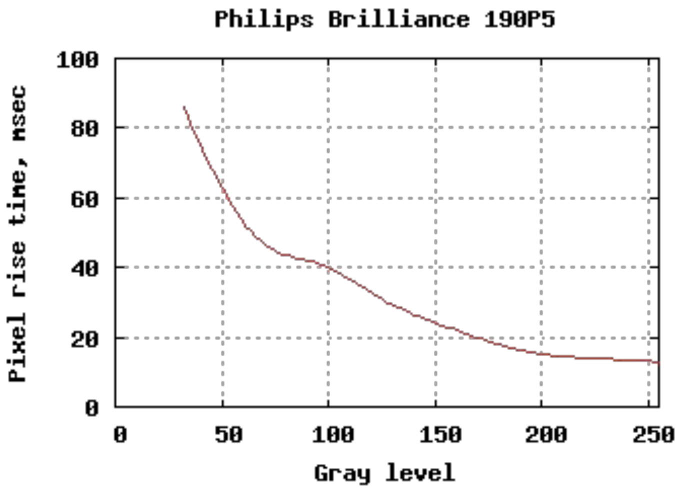

This depressing picture is true for every MVA matrix – the response time grows dramatically when there's a smaller difference between the pixel's initial and final states. Thus, such matrices are practically unsuitable for dynamic games, i.e. for home use. Of course, "suitability" is a subjective category, and some people may be quite satisfied with the image produced by an MVA matrix, but they are objectively slower than TN as well as IPS matrices anyway.

The manufacturers were promising 16-millisecond MVA matrices to compete with fast TN+Film ones, but this claim was just employing the unawareness of the users about their ways of measuring the response time. The inexperienced user thinks that the given full response time value reflects the matrix's responsiveness, so "25msec MVA" is bad, while "16msec MVA" is good. In practice, of course, the shape of the pixel rise time graph remains the same – the curve just goes a little lower with 16msec matrices. Of course, it's good the time to switch a pixel from black to dark-gray has diminished from 90msec to 80msec, but the latter number is still too high to compete successfully with matrices of other types. Thus, the transition from 25msec to 16msec MVA matrices is first of all good for people who use the monitor to work with text or line drawings, since the new matrices blur text less when scrolling. Fans of dynamic games had better prefer a monitor with a 25msec S-IPS matrix rather than with a 16msec MVA one.

The color-reproduction properties of the MVA technology proved to be deficient, too. Such panels give you juicy and bright colors, but due to the peculiarities of the domain technology many subtle color tones (dark tones in the first head) are lost when you are looking at the screen strictly perpendicularly. When you deflect your eye of sight just a little, the colors are all here again. The panel manufacturers sometimes claim a wider gamut range, but as I mentioned earlier, these are rather qualities of the color filters and backlighting rather than of the matrix. Thus, MVA matrices are somewhere between IPS and TN technologies as concerns color reproduction. On the one hand, they are much better than TN matrices in this respect, but on the other hand the above-described shortcoming prevents them from challenging IPS matrices.

Of course, the MVA technology has some undisputable advantages, too. Their contrast ratio may be considered as such, but… Well, when this technology was being promoted into the market – when a contrast ratio of 300:1 was considered an achievement for an LCD monitor – high contrast ratio was touted as an advantage of MVA matrices. Since then, however, TN-matrices have made a wide step forward and the tables have almost turned upon the MVA technology. Moreover, MVA matrices, originally developed by Fujitsu, are now manufactured by several companies of varying manufacturing ability. While modern Premium MVA matrices from Fujitsu and AU Optronics have a real contrast ratio of 400..600:1, products from Chi Mei Optoelectronics (CMO) rarely boast a contrast ratio better than 200:1, which is worse compared to modern TN matrices from major companies like LG.Philips or Samsung. Thus, the fact that a monitor is based on a MVA matrix does not guarantee that its contrast is going to be at a proper level.

It’s all right with the viewing angles of MVA matrices. Like with the IPS type, the declared numbers are the “real” viewing angles. In other words, it’s hard to notice any irregularities when you’re sitting in front of the monitor. The image retains its contrast and color even at big angles of sight (contrary to TN+Film matrices, for example, with their white color changing into yellow if you view the screen from a side). I should also note that the vertical viewing angles of MVA matrices are no worse than horizontal ones.

As you see, the MVA technology is rather ambiguous. MAX matrices seem to belong to text processing and line drawing applications where their excellent viewing angles and high contrast ratio (considering the above-said things about the different manufacturers and production dates) will come in handy, while the color reproduction and the black-to-gray response time won’t play a big role. MVA-based monitors will also suit people who don’t play dynamic games.

The speed these matrices provide is enough to watch movies or play strategy games (and other genres, not critical to the reaction speed), while their deep black color (thanks to the high contrast ratio) will please people who often use their computer in the evening or at night. If you, on the contrary, need a monitor to work with color or to play fast games, then S-IPS matrices should be your choice, irrespective of what the manufacturers of MVA matrices say. Unfortunately, like with S-IPS matrices, the MVA technology has completely left the market of 17” monitors, so you can only buy a 19” MVA monitor. But there are PVA matrices, too, which will be discussed in the next section.

PVA Matrices

The PVA technology (Patterned Vertical Alignment) was developed by Samsung as an alternative to MVA. Well, it was not the only case when Samsung chose the same development strategy – there was the ACE technology, practically analogous to the ordinary IPS. However, I can’t say PVA is a copy of MVA developed by Samsung only to avoid paying licensing fees to Fujitsu. You are going to see below that the parameters and the development ways of PVA and MVA are so different that PVA can be truly regarded as an independent technology.

The liquid crystals in a PVA matrix have the same structure as in a MVA one – domains with varying orientation of the crystals allow keeping the same color, almost irrespective of the user’s line of sight. In fact, the viewing angles (as traditionally measured by the reduction of the contrast ratio to 10:1) are limited not by the matrix, but rather by the plastic framing around the screen.

Alas, there’s the same problem with PVA matrices as with MVA ones – their response time grows catastrophically when there’s a smaller difference between the initial and final states of the pixel.

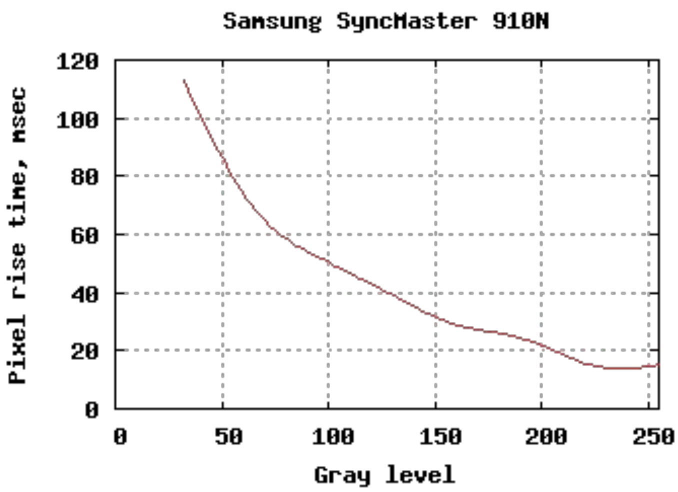

No so long ago Samsung released the PVA-based SyncMaster 193P model with a full response time of 20msec, but it's like with the 16msec MVA matrices – the matrix is really faster than its predecessors, but this improvement is negligible considering the above-illustrated correlation between the pixel's initial and final states.

Color reproduction is not perfect, too, like with MVA matrices: when you are looking straight at the screen, the matrix "loses" some shades, which return after you deflect your line of sight from the perpendicular a little.

The contrast ratio parameter is really good with the PVA technology, though. First, PVA matrices are manufactured by Samsung alone, so there can't be any variation in quality between different manufacturers. Second, Samsung is actively working to improve the contrast ratio and with some results already: monitors with PVA matrices (they mostly come from Samsung, too) typically have a contrast ratio of 600..800:1. Latest models – SyncMaster 910N and 910T – boasted a contrast ratio of 1000:1 and higher in our recent tests (the calibrator I use just couldn't measure the level of black on the 910T model, so the contrast ratio was kind of "infinite"). Generally speaking, PVA matrices are the only matrix type today for which the declared contrast values are true (sometimes the real characteristic is even better than specified). In fact, only they can show you the really deep black color.

Overall, PVA matrices could be said to be an improved version of MVA. Without any new defects, save for those already present in the MVA technology, PVA matrices feature a better contrast ratio and a much more predictable production quality due to their being manufactured on the facilities of one manufacturer only.

Conclusion

Summarizing the article in a single phrase, I'd once again say, as I often did in my LCD monitor reviews: there is no all-purpose LCD monitor and none is going to emerge anytime soon. While a high-quality CRT monitor can be successfully employed for playing games, working with text and processing photos, in the liquid crystals realm each matrix type is more appropriate for each of the enumerated applications.

All the types of modern matrices have high declared parameters, but it is the employed measurement methods that allow the manufacturers to present their produce in the best light, while the numerous conventions and simplifications introduced into the measurement process often play the crucial part – the excellent example is the response time of MVA and PVA matrices.

All modern matrices can be divided into three types: TN, IPS and *VA. TN-matrices have the smallest response time, but can't boast big viewing angles or high contrast ratio or good color reproduction, which makes them in fact fitting for games and simple office work only. A monitor on a TN matrix would be the worst choice for serious work. You can easily tell a TN matrix from any other type by a strong darkening of the image when it is viewed from below, even at a small angle.

Monitors on IPS matrices seem to be the closest to my ideal of versatility and could aspire to be called the best if it were not for two problems: low contrast ratio (no better than with TN matrices) and a characteristic violet hue that appears in black color when the screen is viewed from a side (by the way, it is this violet hue that visually differentiates IPS matrices from *VA ones). On the other hand, monitors on IPS matrices have very good viewing angles and feature an excellent color reproduction, thus they are the only sensible choice for working with photographs among all LCD monitors. Considering that the response time of the latest IPS matrices is approaching that of TN matrices, thus allowing to play dynamic games, monitors on IPS matrices will be a good choice for home use.

MVA and PVA matrices boast an excellent contrast ratio and viewing angles, but they are not very responsive: the response time degenerates quickly as the difference between the initial and final states of the pixel is decreasing. Thus, such monitors suit badly for playing games. They also have some problems with color reproduction – they are worse than IPS matrices in this respect, so MVA and PVA matrices both are unsuitable for working with color. On the other hand, thanks to the high contrast ratio, such monitors will be an excellent choice for working with text, line drawings, and will make a good home device, if you don't need a high-speed matrix. Choosing between PVA and MVA, it's better to go PVA as they have a much better contrast ratio and repeatability of quality from model to model. Moreover, if you're targeting a 17" model, there'd be no choice at all since 17" MVA-based LCD monitors are not manufactured anymore. If you go MVA, pay attention to the level of black color (you have to turn the monitor on in a dim room for that, though), since matrices from different companies are greatly different in quality, not always providing a really high contrast ratio. MVA and PVA matrices can be easily spotted among the others by the lack of any artifacts when you take a look at them from a side – no violet hue in black areas and no darkening when viewed from below.