HP Business PCs
Displays

July 2003

Technical Brief

Evaluating Performance Specifications for LCD Monitors

Contents

Abstract

The intent of this paper is to help the user understand how LCD monitor specifications distinguish monitor products and to what degree specifications indicate actual image quality. Information is presented on the three primary monitor features:

- brightness
- contrast
- response time

Introduction

Performance specifications of LCD monitors have been attracting a great deal of interest lately. In particular, brightness, contrast, and response time specs are drawing a lot of attention as more and more users are starting to understand how these affect their viewing experience. An examination of specifications must, however, go well beyond a search for "bigger is better."

Brightness and contrast specs are a good place to start. These seem to be the first numbers the potential monitor buyer looks at these days, and they can be valuable indicators of product performance. However, there is a point where "enough is enough," and further increases in the numbers no longer necessarily signify the better choice.

This paper will help the user understand

- which LCD monitor specifications are important
- what to look for in LCD monitor specifications
- when to compare monitors side-by-side

Brightness

Brightness—often listed by the more technically correct term, *luminance*[1]—is simply a measure of how much light is being produced by the display under certain conditions. There's a bit more to the actual measurement than this, but what most people call "brightness" is just what you think it is—a higher value here means that those areas of the display that are lit (the white background of a word processing document, for example) will appear brighter to the eye.

But how much is enough? And more importantly, how much is too much? Before the LCD monitor became a significant force in the market, monitor luminance specs were pretty much all the same: somewhere in the range of about 90 to 120 cd/m^2 (read "candelas per square meter," and sometimes referred to by an older term, *nits*)[2]. This similarity in specs came about because it just wasn't practical to achieve higher light output in a high-resolution CRT. In contrast, the CRT used in the typical television set, which is a relatively low-resolution sort of tube, would commonly produce something like 300–400cd/m^2—but keep in mind, a TV is usually viewed from a much greater distance than a desktop monitor, and often in an overall brighter environment.

Luminance is a measurement of how much light is emitted by the display in absolute terms. **But it's also important to judge a display by how bright it *appears to be* under lighting conditions similar to those that will exist in the environment in which the display is expected to be used.** For example, a monitor or TV that appears very bright in a dark room might appear dim—even to the point of being unreadable—if viewed in direct sunlight. In other words, what we're usually concerned about is how bright the monitor appears relative to its surroundings, not what the absolute luminance measurement might be.

---

[1] The real difference here is that "brightness" is a subjective perception, while "luminance" is a formally-defined objective measurement.

[2] An even older unit, the *foot-Lambert* (ft.-L), may also occasionally be encountered. 1 foot-Lambert is about 3.43 cd/m^2.

The LCD monitor operates on a completely different principle than its CRT-based cousins; instead of producing light itself, an LCD panel actually is just controlling the amount of light passed through from an always-on "backlight" (figure 1). And if you want a given LCD to appear brighter, all you have to do is give it a brighter backlight. This has enabled LCD-based desktop monitors to claim some very high luminance specs compared with their CRT predecessors. Values in the range of 250–300 cd/m^2 are today quite common, and many of the most recent products spec sheets will give even higher luminance figures. However, are higher "brightness" specs really needed, and what else should we be looking for here?

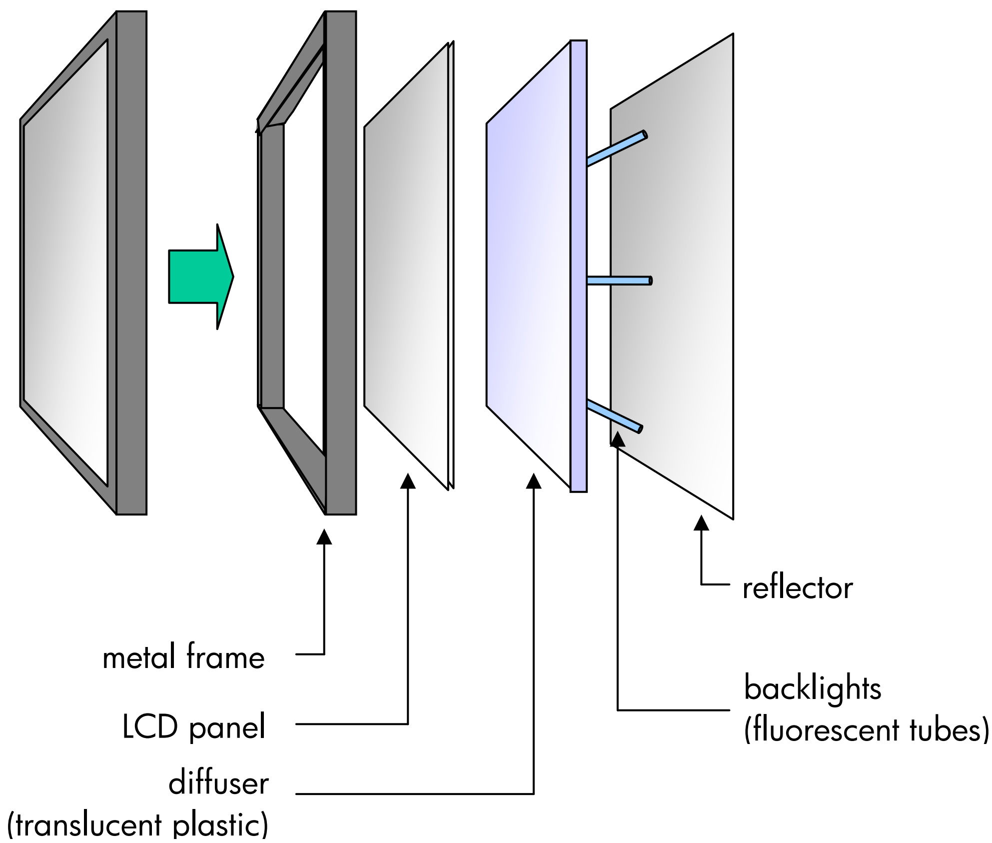

Figure 1—LCD Monitor Components

The LCD panel itself does not produce light, but instead acts as a "valve" for light produced by a backlight assembly, which is always on. Obviously, *if* all else is equal, a higher luminance specification should indicate a more capable, or at least a more flexible, product—meaning one that could still provide an adequately bright image over a wider range of ambient lighting conditions. But all else is rarely equal. A brighter display can be achieved by adding a more powerful backlight or a more sophisticated LCD design. Higher luminance may mean a more limited viewing angle. Adding more power into the backlight system can negate one of the major advantages of the LCD over the CRT, better energy efficiency.

And mere "brightness" isn't all there is to a good viewing experience—various factors besides viewing angle determine the color performance of the display, the response time, and contrast. Some brighter panels can actually have poorer response time.

Beyond a certain point, a higher luminance spec is simply not important. In a typical office environment, **most users will find a display luminance of between 150 and 250 cd/m^2 about right**, given good performance in other areas. (For further information, see the various international ergonomic standards dealing with displays, such as ISO 9241 Part 3, or the recent TCO '03 requirements for CRT and LCD monitors.) A call for a "brighter" display may actually reflect a need of better contrast or a wider color gamut. A display capable of providing 400 or 500 cd/m^2 may actually appear too bright in the office, and need to have its backlight brightness turned down for comfortable viewing.

When comparing monitors, consider the viewing conditions at your home or office. Showrooms can be much brighter or dimmer and the viewing distance can be quite different.

Contrast

LCD monitor makers claim ever-higher values in their contrast specs. A couple of years ago, 200:1 or 300:1 contrast specifications were common and considered quite acceptable; claims today of 400:1, 500:1, or even higher are not uncommon. But what exactly does the contrast specification tell you, and what is a reasonable figure?

Contrast, or more correctly *contrast ratio*, is just an indication of the difference in luminance or "brightness" between the brightest and darkest areas of the display—or more simply put, it tells how much brighter "white" is than "black." For instance, if white regions of the image are measured as having a luminance of 300 cd/m^2, while the same or adjacent areas only produce 1 cd/m^2 when "black," we would say that the display has a contrast ratio of 300:1 ("three hundred to one"). A *truly* black region should emit or reflect no light at all, but this is clearly not going to be the case with any current display under normal viewing conditions. Even when fully "off," *some* light is very likely going to be leaking through at any given pixel location—and more importantly, there isn't a display made that will reflect *none* of the light striking it from outside light sources. So what a contrast ratio spec really comes down to, once the luminance of "white" is set (which is what the luminance spec is telling you), is just "how black is black?"

From this, it's easy to see how one could get the idea that a bigger number for the contrast ratio spec is always better. However, such a simplistic approach ignores one very basic fact—we're talking about something that's meant to be viewed by human beings, and the capabilities of human vision are most definitely limited. Contrast ratio is very closely related to the question of usable *dynamic range*, meaning the range of luminance levels over which differences in brightness can actually be seen. Human eyes are capable of adapting to a very wide range of absolute luminance, but at any given moment can only detect differences over a range of not much more than a few hundred to one. If a light source exceeds this range on the high end, it's going to look "white" no matter how bright it really is—and more importantly for this discussion, a light source that is dimmer than the low end of the range covered by your vision will appear "black" (1,2).

What this says is that there's definitely a visible difference between displays offering, say, a 50:1 contrast ratio on the one hand and 150:1 on the other, but relatively little difference between displays of 350:1 and 450:1. And somewhere above this point, further increases in the measured contrast ratio become essentially meaningless in terms of continued visible improvement in the image.

A big part of the reason for this is that desktop monitors aren't normally used in otherwise totally dark rooms—and yet that's the sort of conditions under which these contrast ratio measurements are obtained. They measure only the maximum contrast ratio that the display is capable of under the most ideal conditions (ideal, at least, for measuring contrast ratio!). But if you're *not* using your PC in the depths of a coal mine at midnight, it's pretty certain that ambient light sources are reducing the actual contrast ratio delivered by *any* monitor to something far less than the spec-sheet number. A specification such as "450:1 contrast ratio" is obtained by looking only at the light actually coming out of the panel in the white and black states; such a measurement typically will completely ignore the light that might be reflected by the panel in normal use (figure 2).

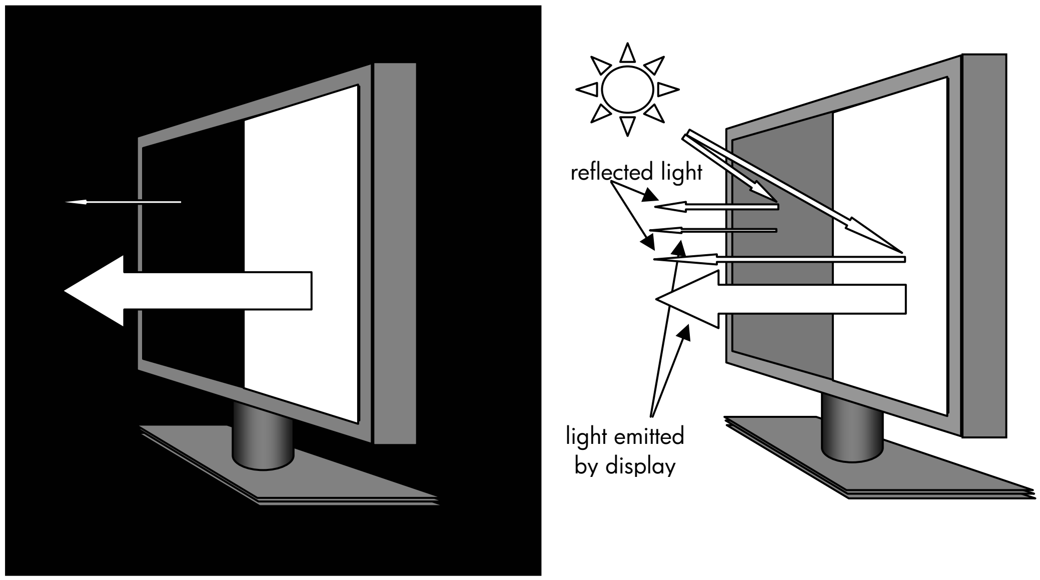

Figure 2—Maximum Contrast vs. Actual Contrast

Very high contrast ratio values may be measured in a dark room, when comparing only the light actually emitted by the "white" and "black" areas of the display. Under normal lighting, however, reflected light from outside sources will often reduce the contrast actually delivered by the display.

Under typical office lighting conditions, it is very difficult to obtain a true contrast ratio in the image *as seen by the user* of much more than 100:1 or so—*regardless* of the value of the "dark room," spec-sheet contrast ratio given.

Response Time

The last of the "big three" specs is response time. This is an area in which LCD monitors, until fairly recently, have been at a disadvantage relative to their CRT-based competition. CRT displays can respond essentially immediately, at least as compared to the frame rates encountered in normal TV and PC video. Typically, these are somewhere in the neighborhood of 50–100 Hz, meaning that each frame (or field, in interlaced systems such as television) takes about 10–20 milliseconds (ms) to complete. We might assume, then, that any display *should* respond to changes rapidly enough to have everything properly "written" on the screen in this time or less. (Otherwise, a given image wouldn't be completely displayed and stable before the next one came along!)

The CRT, as noted, is a very fast response sort of display. It operates by scanning a beam of electrons across a chemically coated screen, which lights up wherever and whenever the beam strikes it. It does this essentially immediately—within microseconds of the beam first striking the screen at any given point, that part of the screen is emitting light. Unfortunately, it also responds almost equally quickly in the other direction; once the beam moves on and is no longer actually driving that spot on the screen, it very rapidly stops producing light. It doesn't stop immediately—the chemicals that coat the screen are created to exhibit *persistence*, which means that light continues to be emitted for some time after the drive source (in this case, the electron beam) is removed. (This is also why your TV screen can often be seen to be faintly glowing for a short time after it's turned off in a dark room.) What seems to be a stable and complete image on a CRT screen, though, really isn't. What's *really* being produced by the CRT is a very bright and very rapidly moving point of light (and a rapidly fading and much fainter version of the image in its trail), which is integrated by your eye (via what's called the *persistence of vision*) into the appearance of a steady image. The image is actually anything but steady, and this is what causes "flicker" to be such a problem in CRTs. Your eyes aren't fast enough to see the bright, moving dot of light across the screen, but you *can* perceive the overall unsteadiness of the light produced. "Really quick response," therefore, may or may not be a good thing (3).

LCDs obviously work quite differently than the CRT. In an "active-matrix" or "TFT" type of LCD, which is the kind found in almost all LCD monitors, a transistor that is built into every pixel (actually, every individual red, green, and blue *subpixel* that makes up each complete pixel) acts to hold that location in whatever state is desired, essentially until new information is "written" into that location. The video information is sent to each pixel or subpixel, the liquid-crystal material is set to pass the correct amount of light, and then everything pretty much stays steady until that location has to change. The problem here is that the LC material doesn't respond immediately. After the proper video level is "captured" by the transistor at that location, the liquid crystal starts to change to the desired state, but it takes a relatively long time to get there—up to 50 or 60 ms in some older types, much longer than the duration of the displayed frame. Once it's stable, everything is fine, but this relatively slow response has meant that LCDs until recently had some visible problems with showing motion. Rapidly moving objects in the image—even a cursor or scrolling text—would experience smearing and "ghosting" as the LCD could not keep up with the rapid changes. LCD monitors had none of the flicker that plagued CRTs, but they just didn't do a very good job with fast motion.

This situation is now very quickly improving. New developments in LCD design and drive techniques have brought the response times down well under 20 ms; in the best of the current panels, down to the 12–16 ms range, equal or better than the fastest standard video frame time. And since we're talking only about how quickly the LCD responds—nothing about changing the way it stays stable once it *has* responded—we get both very good motion-video performance *and* no flicker problems.

Unfortunately, **faster is not always better**. While it is, in general, true that lower response times mean a "faster" display and, therefore, one that will perform better in terms of how well it shows motion, there are some things to be aware of when comparing products in this area. First, the "response time" specification most commonly published is the measurement of how quickly the display changes from black (completely off) to white (completely on) *and back again*. In other words, it's the sum of the "turn-on" and "turn-off" times for as big a change as can be made in the light output. You do not know which is faster, turn-on or turn-off (and if they're not reasonably close to the same, some odd effects can happen – imagine how the display would look if it took a significantly longer time to turn a given pixel *on* than to turn an adjacent pixel *off*), and you also can't directly compare this sort of number to one which only represents the response time in one direction or the other.

Further, while it might seem that the full black-to-white (or white-to-black) transition should always be the worst case, that's not always true. In some LCD types, transitions between two intermediate "gray" levels can actually take considerably longer than going from full black to full white. This can make for some very annoying problems, especially in terms of color. Moving objects on displays showing this sort of behavior can change color considerably *while in motion*. (Imagine what happens, for instance, if the red and blue parts of the image are changing over the full range—and so are changing very quickly—but the green is only changing a little, and therefore much more slowly.)

Finally, just as with the contrast and brightness specs, response time should not be viewed as a "more (or less, in this case) is always better" sort of thing. At some point, further improvements in the display's response time may be far less important than improvements in other areas. Even worse, the absolute fastest response time (or highest contrast, or brightness, or…) may mean poorer performance in other areas. The same design that gives the ultimate in speed might be costing you in brightness, color accuracy, power consumption, etc.

Summary

The bottom line is that you should not carry any spec comparison too far. Obviously unacceptable candidates can be eliminated, but, within a fairly broad range, the final decision should be made based on the actual performance of the display as *you* see it, preferably under viewing conditions (and with the sort of images) that resemble those you'll most often be using.

References

1. Rose, A (1948). The sensitivity performance of the human eye on an absolute scale. Journal of the Optical Society of America 38, pp. 196-208.
2. Color Science: Concepts and Methods, Quantitative Data and Formulae, 2nd Edition. G. Wyszecki and W. S. Stiles, publ. 1982, John Wiley & Sons.
3. Farrell, J.E., et al (1987). Predicting flicker thresholds for video display terminals. Proceedings of the Society for Information Display, Vol. 28, No. 4, pp. 449-453.

Copyright and Trademarks

July 2003

342218-001