REPRINT

Plasma Displays

By Richard Stevens

Following his excellent talk and demonstration at the Kingswood Warren open day last year, Richard Stevens from BBC Research and Development provides a layman's guide to the technologies behind these increasingly popular large flat screen displays and some personal views as to what he is likely to be buying when the time comes to replace that TV in the corner of the living room.

This article is intended as a quick primer on plasma displays, providing useful background information for those who are not experts in this field. It is a deliberately simplified layman's guide, intended only as a rule of thumb aid to understanding the principles involved, so not all the details will necessarily apply to any one display device. The information has been derived from my own work on these devices and from the published literature, and I should warn readers that developments in this field are happening so quickly that published information can rapidly become out of date.

What is a PDP?

A plasma display panel is a thin sandwich of two sheets of optical grade glass. The meat of the sandwich is exceedingly thin, and filled with nothing much more than a rarified Neon and Xenon mixture, much the same as the content of a fluorescent tube. It is quite different from a Cathode Ray Tube or a Liquid Crystal or a Field Emissive Display. Some of the image forming electronics for PDP, LCD and FED can be the same, but this is more because of the geometry of cell structure being akin to a memory map than to any similarity in the function of the display.

Figure 1 shows a diagram of a single cell. You notice that there is no gun structure, no focusing or aiming, no beam to deflect, in fact almost no controls at all. The cell can be either on or off. That is it. The plasma display is simply a flash gun.

There are just three electrodes in a generic AC plasma display. One lies beneath the phosphor, the other two are in a pair on the other sheet of glass and covered by an insulating dielectric layer and a MgO layer. The space is filled with the Xenon and Neon. (10% : 90% ratio at about 600 torr.) To prime the cell ready for excitation a small charge is trapped on the dielectric, ionizing some of the gas.

Originally published in the BKSTS journal *Image Technology*, March 2002, p. 24. Richard Stevens is with BBC Research and Development. This article is reprinted with permission of the BKSTS, for which SMPTE would like to express thanks.

Excitation is performed by passing an AC current through this ionization and so generating a plasma which has an ultra violet discharge. The UV stimulates the phosphor to produce the visible light. Once the excitation (or sustain) cycle has been completed, the cell needs to be "erased," so the charge is drawn back towards the conductor beneath the phosphor.

Figure 2 shows the sequence of an AC plasma burst exciting the phosphor at the back of the cell. Notice that the surface of the phosphor that you view the image on is also the surface being excited, so only a relatively low energy is required compared with when one excites the back of a CRT phosphor with a high-speed electron beam. The web site http://www.sni.net/siglo/examples.htm contains more examples in similar vein, including a co-planar cell system, which is the more common one in use these days.

Display Structure

Now the first thing we need to make a picture is some structure. Structure is defined by the physical layout of the cells in rows and columns, as in Fig. 3. It is all rather simple,

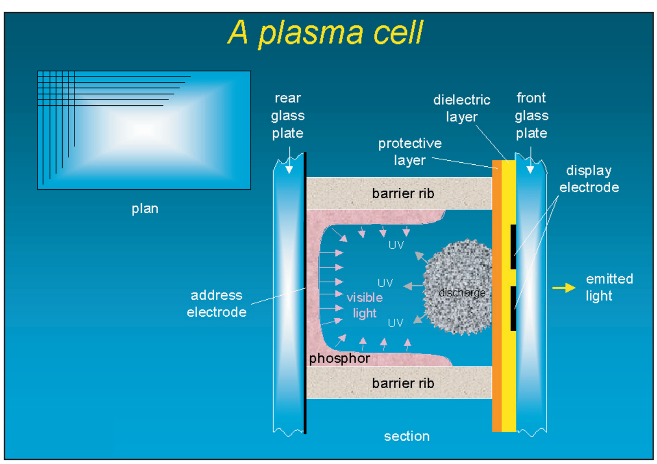

Figure 1. A plasma cell; section, and plan.

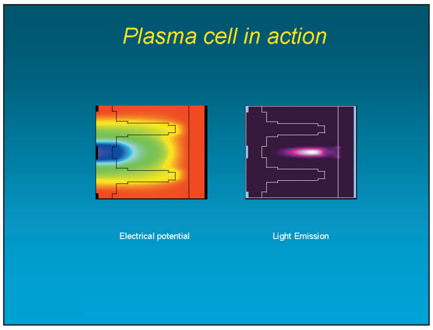

Figure 2. The plasma cell in action.

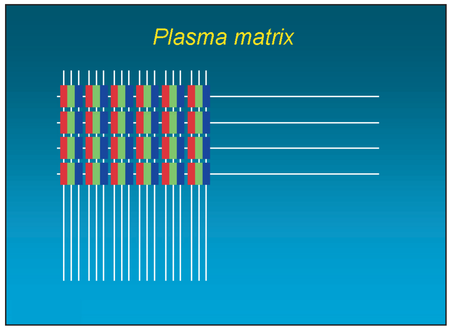

Figure 3. Plasma matrix showing pixels and rows.

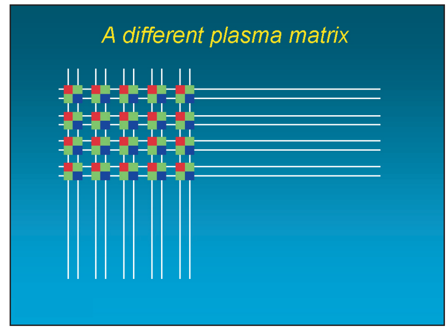

Figure 4. A different plasma matrix of pixels and rows.

spatial resolution is just a matter of mathematics; how many pixels would you like in your row, Sir? and how many rows? There are no focus or spot shape problems, no scan geometry correction required. We could tailor the display to suit the sample structure of the picture source; Mr. Nyquist stays happy, and all the scaling chip manufacturers would wring their hands in despair.

In theory a 720 pixel by 576 row display would be ideal for Europe, and one manufacturer is even making an interlaced display, (but 1024 by 1024, not 720 by 576). We need go no further if we forget HDTV and computer monitors.

Practical Constraints

The main difficulty is in making cells that are small enough, and that is mostly because of the current requirements on the excitation conductors, which in turn have a dependency on the work function of the phosphor. Small cells have a low work function, and so need higher currents to compensate for this, but, alas, small cells have to use thinner more resistive conductors. This means that major streaking problems can arise because of current starvation, and significant heat is generated within the panel, needing noisy fans to cool it.

So the answer is to make bigger cells, with thicker conductors, in bigger panels, to be viewed from further away, and then you won't hear the fans, maybe! One mm square for an RGB color triad pixel seems to be a commonly used size.

Figure 4 shows a different plasma matrix of pixels and rows. I can see some interesting possibilities for variations in color rendition and improving resolution. We might not have to have the cells in RGB RGB sequence all the time.

We might emulate a single chip CCD color camera, with RG on one row and GB on the row below, useful for television with YUV sources and giving more than a 50% improvement in horizontal resolution, no loss in vertical luminance resolution, plus a reduction in aliasing and edge coloring. This would not, however, be so good for a computer display.

A useful analogy is to compare a CRT and PDP in the same way as in the past we might have compared tube cameras with CCD cameras; the newer technology eliminates a whole range of problems inherent with the old, but brings along its own baggage as well.

So how do we turn on or off individual cells in a large array of cells in order to display a picture?

As Fig. 5 shows, there is an addressing period during which each individual cell on the screen may have the charge placed on the insulator. This is followed by a common excitation period, during which (for an AC plasma) the current is forced through all the primed cells, exciting the plasma and so the relevant phosphor. And of course preceding each address period there has to be an erase sequence which removes any space charges that may have been set for the last subfield. But this only means that the cells can be on, or off.

Dynamic Range

The second thing we need to make a picture is some dynamic range. We need grayscale, we need lots of differ-

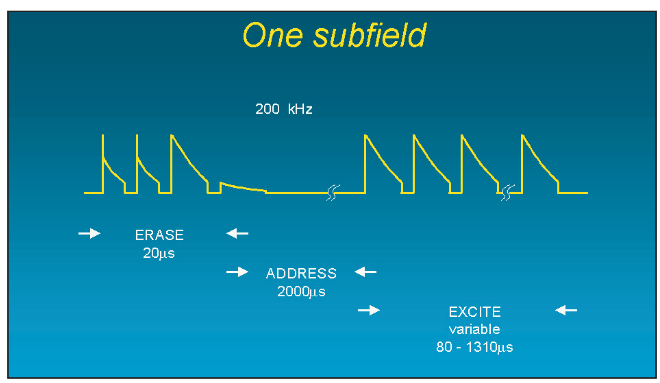

Figure 5. Waveform showing one subfield.

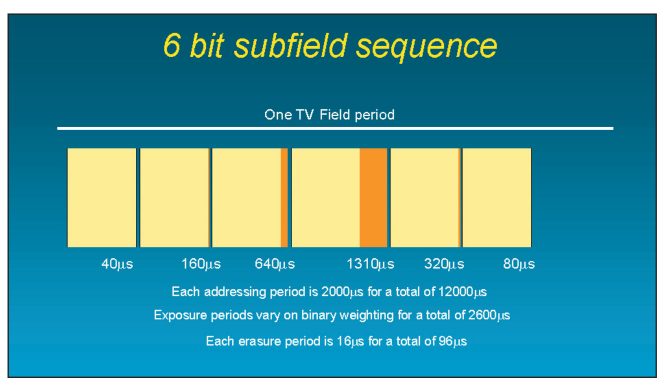

Figure 6. A 6-bit subfield sequence.

ent colors. We do not want visible quantizing levels, or posterization.

This is the very awkward bit. The plasma screens are nowhere near as good as a CRT on dynamic range. In fact the PDP can be On, or it can be Off. A good CRT might have a DC contrast ratio as great as 1000 to 1 with any value you want being possible between these limits, although one normally only expects an AC contrast of better than about 50 to 1. On a PDP the dynamic range, or brightness variation, is determined simply by changing the perceived duration of the flash.

Trickery

So we have to resort to some form of trickery; this can work to our advantage if done correctly, but unfortunately there are economic constraints. The usual way of varying the perceived brightness is to vary the number of AC excitation pulses applied to the cell. On a display with maybe a million cells it is not practical to control the times for which they are on individually and simultaneously.

Accordingly the usual technique is to break up the display time into subfields, for example on a binary weighting scheme. So in a six subfield system like the one shown in Fig. 6, the least significant subfield is on for 1/32nd of duration of the most significant one.

Measured Results

Figures 7 and 8 show the results of some measurements I made a couple of years ago on a six subfield system.

The problem with the six subfield method is that the light output from the display follows a linear law.

Unfortunately all our current video sources are expected to be displayed on a CRT, which has gamma. So there is a discrepancy in the voltage to light transfer characteristic. This is one of the biggest single problems associated with plasma panels.

Another of the problems with this kind of scheme is that the subfields may not be on a true binary sequence. The subfields are quantized by the AC excitation cycles, and so the grayscale might not be monotonic. A reasonable 8-bit panel is typically only accurate to 6.5 bits or even worse, as you can see in Fig. 9, which was derived from measurements I made on a six bit system. Here the brightness variation is controlled by the total number of excitation cycles, but the total is unlikely to be properly divisible into the binary scheme.

Brightness Variations

On some displays, as the total light output, or more accurately the power consumption, rises, there comes a point when a protection scheme comes in and the number of excitation cycles per subfield is reduced to protect the system. The monotonicity fails again, but differently! So if correction had been applied for the poor monotonicity in one way, now it might fail the other way but worse; for want of a better phrase I guess it could be called a differential gain effect. So I have found it is best not to attempt to correct for the monotonicity errors.

It gets worse still. I mentioned before that the display is a linear device where the number in equates to the brightness level out. All our video source material is expected to be applied to a CRT which has a Gamma power law, where the number in equates to a lot more out as the numbers rise—Fig. 10 shows the ideal gamma law curves, and Fig. 11 provides a more detailed view, showing just the first 64 levels.

Broadcast video is 10-bit accurate and this gives very good results on a CRT, but when translated to the linear domain we need close to 16 bits. The displays are generally only 8 bit (although some of the more recent ones are 10 or 12 bit, depending on refresh rate). Errors due to this restriction show up as poor resolution in the dark scenes and poor definition of subtle colors, in other words posterization. We have the cartoon or comic book look.

Flicker Problems, Color Fringing

Things can get even worse! Having eight flash guns going off one after the other at maximum intensity for varying lengths of time 50 times a second, or worse still—sometimes not, produces an objectionable amount of flicker, and the

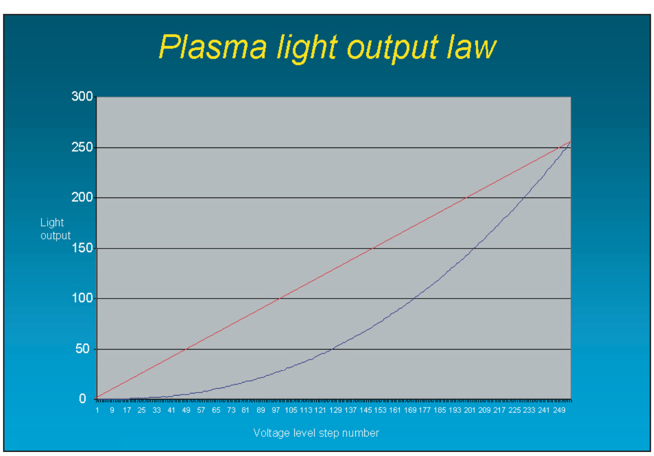

Figure 7. Plasma light output law: Linear curve and Gamma curve.

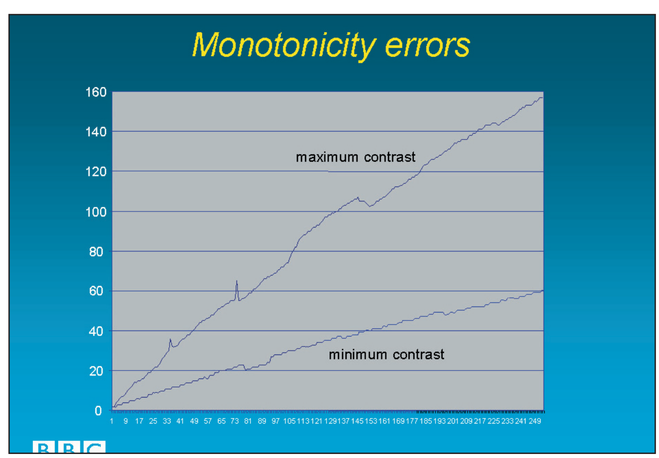

Figure 8. Monotonicity errors.

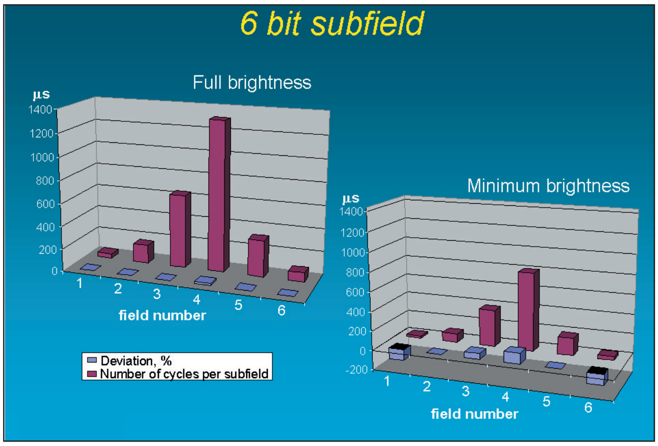

Figure 9. 6-bit subfield—effects of brightness variations.

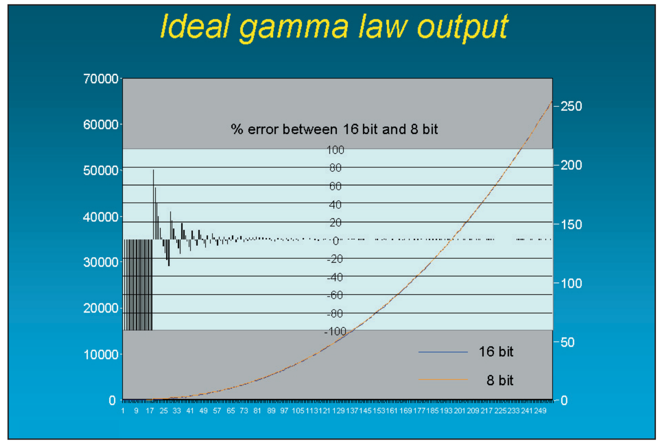

Figure 10. Ideal gamma law output from 16 bits, and from 8 bits and % error.

motion portrayal is awful, with heavily colored fringes. This is because the flashes are at different points on the timeline, but the eye expects the timeline to have a smooth and even tempo, so will attribute a spatial error to the scene even though it is actually a temporal error. Hence the colored fringes visible on some systems, as indicated in Fig. 12.

Reducing the Number of Subfields

Fortunately there are some tricks that can be used to get around most of these problems. One of them, reducing the number of subfields, kills several birds at once, and it is a standard feature on most manufacturers' products. At first sight you might think that the image would be degraded further, since the most significant six of the eight different subfields are reduced to two identical sets of three, with the remaining two filling in the gaps. Figure 13 shows the amended subfield sequence of a 6-bit system.

This has the effect of doubling the subjective flicker rate, in much the same way as a film projector shows the same frame twice or three times to reduce flicker. The image is still only refreshed at the original rate. The second big advantage is that the color fringing on motion is greatly reduced, so much so that you only can see it if you are looking for it with an expectant eye. A third benefit, strangely enough, is that the monotonicity gets better, because with fewer subfields it is easier to make the sums right. In this instance cheaters prosper! The big compromise made is that there is some static dithering to spread the dynamic range spatially, providing an increase in fine grain noise, somewhat akin to a half toning system on a laser printer. It is still only an 8-bit result, but can look as good as 16-bit linear on appropriate displays which use some of the techniques I have developed, although pictures can still look posterized on some of today's standard displays. I have demonstrated pictures from a 1250-line interlaced high-definition source, down-converted to Rec. 656/601 625 lines interlaced, which were then de-interlaced and scaled to fit on the panel as 480P with 852 pixels a line, and they were judged to be excellent by many critical viewers. There are no analog bottlenecks restricting these signals in the 625-line domain.

Plasma Panels in the Home?

How likely is it that we shall be using plasma panels for our home TV displays in the near future? I think that it is very likely, with a few constraints. Firstly, viewers should be realistic about the viewing distance.

It has been proven many times that a natural viewing distance for displays of this size is about 5 times the picture height, or about 2 1/2 m for a 42 in. 16 x 9 panel. This may be serendipitous, but I would like to think it is by design, since the resolution limit of the human eye at 2 1/2 m is

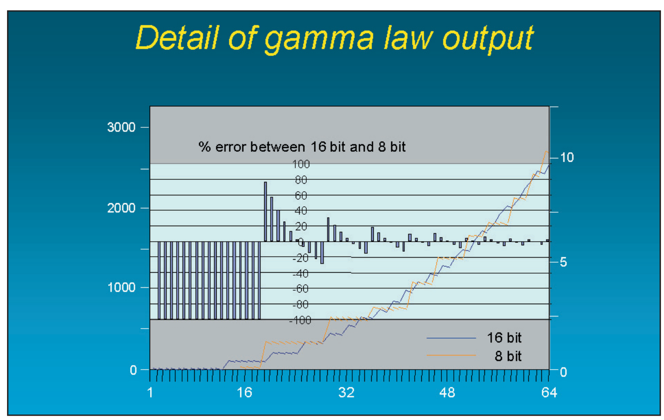

Figure 11. Detail of gamma law output from 16 bits, from 8 bits, and % error, but just the first 64 levels.

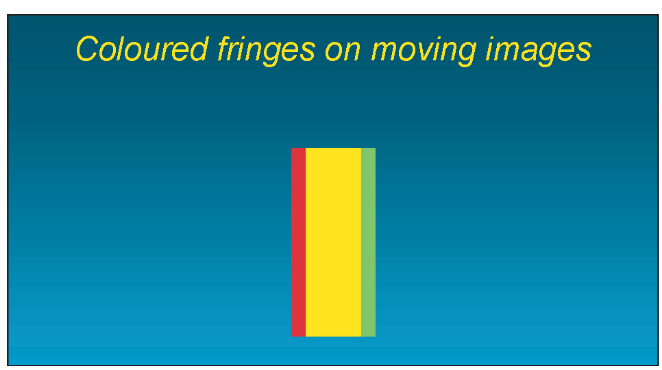

Figure 12. Colored fringes on moving images.

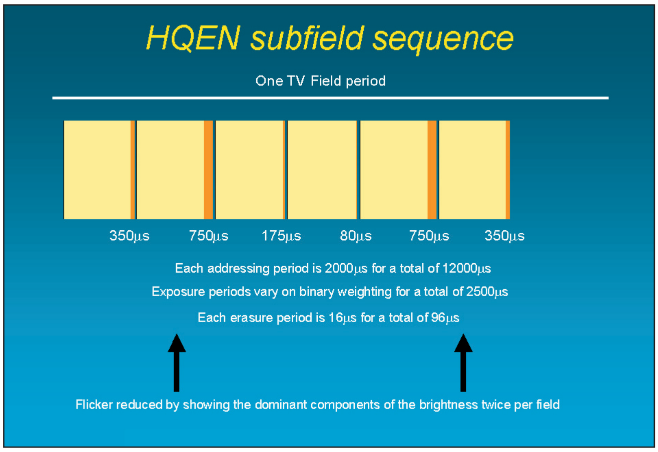

Figure 13. 6-bit HQEN subfield sequence.

Summary of pros and cons

	DMD projection	CRT	PDP	LCD	FED
geometry	**excellent**	average	**excellent**	excellent	**excellent**
focus (detail size)	**good**	variable	**good**	excellent	**excellent**
picture placement	**excellent**	variable	**excellent**	excellent	**excellent**
picture size	**big 60" to>360"**	medium 2" to 38"	**larger 21" to 60"**	small 1" to 28"	**medium?**
chromaticity	**excellent**	excellent	**good**	poor	**unknown**
flicker	**good**	average	**better**	good	**excellent**
greyscale	**average**	excellent	**average**	poor	**unknown**
viewing angle	**average**	good	**excellent**	poor to OK	**good**
weight	**medium**	heavy	**medium**	light	**medium**
cost per square inch	**high to medium**	low to medium	**medium to low**	high to medium	**unknown**
box	**bulky**	bulky	**slim**	small	**slim**
spatial resolution	**medium**	high	**medium**	high	**medium**
contrast ratio	**good**	good	**good**	poor to average	**good**
noise	**big fans**	silent	**fans likely**	silent	**silent**
domestic receivers	**yes**	yes	**yes**	yes	**not yet**
future developments	**yes**	unlikely	**yes**	yes	**probably**
will I buy one? (2 to 3 years time)	**no**	no	**likely**	unlikely	**maybe**

Figure 14. Summery of pros and cons.

about 1mm. (Remember that the most commonly used pixel size in PDPs is 1mm.)

A 42-in. PDP is large compared with the CRT displays currently used in most homes, but being seen from a distance of about 2 m or more should not prove difficult in many living rooms. At this distance most people will not notice the pixel structure, and the picture size is no longer intimidating. Instead the "Home Cinema" effect comes to life and one can get used to it very quickly indeed!

HDTV—Not Yet

Curiously enough, when I compared 1250-line high-definition images on a 38 in. HD CRT monitor and the same images down-converted to 625-line Rec. 601 on the 42 in. PDP, at this distance of 2 m I preferred the PDP, since it had more contrast and appeared to be just as sharp. Of course, as I walked up closer to the displays the HD CRT display continued to get better and better, and the PDP was just too many pixels and not enough picture. So we do not really need HD just yet, not until we go to bigger displays, say 50 to 80 in. or more.

For those of you who have looked critically at the plasma displays currently on the market, it is important to note that current generation PDPs are optimized as computer display devices, and they have a black mask around each pixel to make the text look sharper. This is a bad thing for a TV display as it contributes to high frequency aliasing by putting a lot of energy into the system at the sample rate; it breaks the visual integration of the picture. I feel that when the market expands there will be displays optimized for TV, but that can also handle computer input, and different displays optimized for computers, but which can handle TV input. Some of the engineering sample prototypes that I have examined can produce extremely good TV pictures in spite of the lack of dynamic range, but the black mask on the production models makes the pictures so much poorer with the high energy 8 MHz stimulating all the edge of band aliases. It was almost like the difference between a good shadow-mask CRT and the early 30AX PIL CRTs, where you had to choose between sharp but dim pictures or bright pictures seen through a picket fence.

Comparing Display Technologies

It is always interesting to compare the various possible display devices, when considered for domestic viewing, and Fig. 14 attempts to show the pros and cons of each. It is important to note that there are some new developments on the addressable-matrix LCDs, especially to do with better contrast ratio and grayscale, but the costs are high and the probable target market will be workstations. There are also some new developments coming along using DMDs in much higher contrast ratio rear projection systems. Putting my theoretical "money where my mouth is," you will see that I have indicated my preference for my future household TV display, and I shan't be buying a large CRT!

Sound and Pictures

A digital display of digital transmissions of digital pictures, with the full surround sound experience is the target. Removing all the restrictive quality bottlenecks of PAL and analog transmitters, could signal as big a change as that from Vinyl to CDs.

What Does the Future Hold?

As better programs become available on digital transmission, and DVDs, the requirements for a higher resolution display become more pressing. I don't want anything less than a PDP with direct digital interface to a set top box with a direct digital output, or to a DVD player with a direct digital interface.

I see a bright future for video separates. The display device now has a life in excess of ten years, but the transmission formats are in a state of flux. I will probably never again buy an integrated receiver as my main display device at home, but expect to change between better (and cheaper) set top boxes over the years.