Cathode Ray Tubes:

Getting Down to Basics

FORWARD

This book was written to serve three major purposes:

1. To understand the interaction and interdependence between the CRT and electronic circuitry in measurement devices having CRTs. The cathode-ray tube (CRT), as the output or display section of oscilloscopes, graphics terminals and other measurement devices, requires understanding in isolating a malfunctioning electronic circuit, the design of circuitry that interfaces with the CRT is dependent upon the requirements of the CRT and before these circuits can be fully analyzed the requirements of the CRT must be known. The proper operation of the various controls and adjustments directly associated with the display requires an understanding of the probable effect upon the CRT.

2. To understand the basic theory or principles of CRT design and operation. In today's world of solid-state devices the principles of operation of vacuum devices is relatively unknown. This book is an attempt to give the engineer, technician or other reader a basic understanding of CRT operation.

3. To consolidate previous CRT-technical documents under one cover. Over the past 30 years there have been a number of CRT theory booklets and technical reports written by Tektronix covering CRT design and theory of operation; some were published and now are out-of-date and no longer in print, others were never published.

I wish to acknowledge those whose publications have been used in this book and whose efforts continue to advance the performance of cathode-ray tubes.

Vernon L. Isaac
CRT Product Assurance Manager
Tektronix, Inc.

For further information, contact: Tektronix, Inc., Corporate Offices, P.O. Box 500, Beaverton, OR 97077. Phone: (503) 627-7111; TLX: 151754; TWX: (910) 467-8708; Cable: TEKWSGT. Subsidiaries and distributors worldwide.

STORAGE TUBE BASICS

All storage tubes rely on the mechanism of secondary emission from a dielectric surface. In this chapter we shall study the principles of this mechanism.

The purpose of storage tubes is to record the movement (and in some cases the intensity) of an electron beam over a target area. In order to make the various parts of the target separately addressable, the target must be made of a dielectric material of such composition and construction that lateral leakage is kept to a minimum. The target should therefore be thought of as a collection of separate insulated points, and since it is insulated it can be described as "floating".

When a beam hits such a dielectric target, secondary emission can occur, due to the landing energy of the electrons, other electrons are knocked out of the target surface and collected by a nearby collector element. In order to do its job as a collector, its voltage has to be more positive than that of the target itself, but not so positive as to appreciable attract the primary beam electrons.

It stands to reason that when the electron beam lands with very little energy on the target it will knock out few, if any, secondary electrons. As the landing energy is increased, secondary emission will increase until, beyond a certain limit, the landing electrons hit the surface with such force and penetrate the material so deeply that more and more of the secondary electrons produced by the impact are trapped within the material instead of escaping. The landing energy of the electrons is determined solely by the potential difference between the cathode from which they originate and the target on which they land; it is not affected by intermediate accelerating and decelerating potentials.

If the landing speed is varied and the amount of secondary emission plotted, a curve with the general appearance shown in Figure 22-1 results. This is true of all dielectric materials, but the exact voltages at which the crossovers and peak occur depend on the characteristics of the dielectric material. The figures shown are typical for the dielectrics used in the tubes discussed in this text. (The slightly negative starting point of the curve is due to the thermal energy with which electrons are leaving the cathode.)

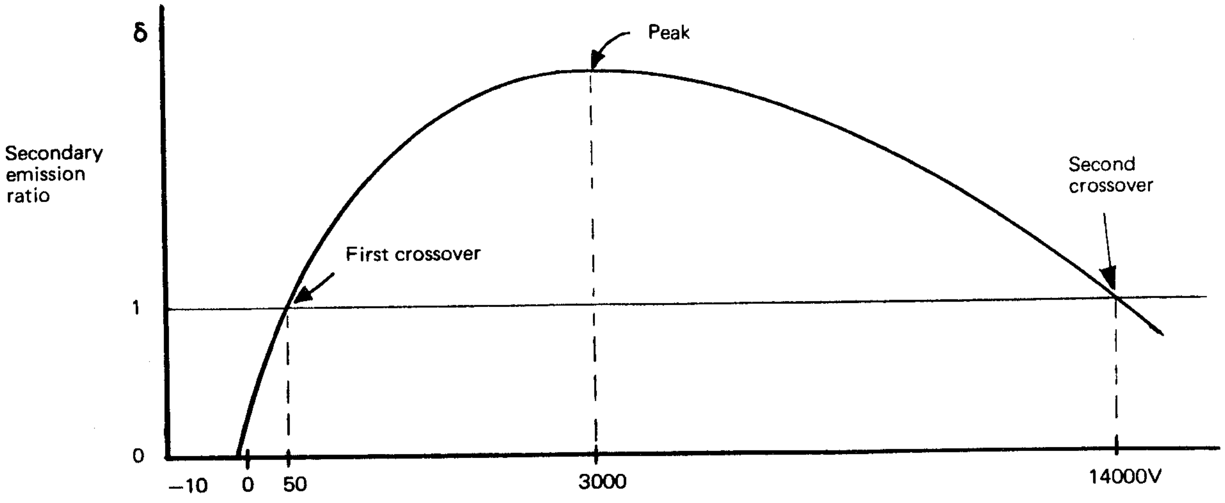

Figure 22-1

Secondary emission curves are generally shown in terms of the ratio δ of secondary emission to primary (or incident) beam. Such a presentation is valid and useful because, for a given target material and construction, the ratio does not change with the intensity of the primary beam, and plotting the curve in these terms brings out the essential points about secondary emission.

The line on the graph representing $\delta = 1$ is of great significance. Portions of the curve above it represent conditions under which the target loses more electrons by way of secondary emission than it gains, and conversely at points below this line the target gains more electrons from the primary stream than it loses. Since the target is in fact a dielectric, which is electrically floating, its surface voltage will drift up or down whenever there is an imbalance between the number of electrons landing and leaving. But before we go to the trouble of studying the effect of this voltage drift in detail we must look at two factors which will lead us to modify the shape of this curve.

First, we assumed that the collector was always slightly more positive than the target, so that any electrons liberated from the target would be attracted by it. But since the target is in fact an array of insulated and independent points, what constitutes the target? How could we measure it? And how could we make the collector always sit at a level slightly higher than the target? As a practical solution the nearby collector is simply held at a reasonable fixed positive voltage, typically 150 V. This will be sufficient to collect secondary electrons as long as the target voltage does not exceed +150 V. But if, for any reason, a point on the target does exceed +150 V the liberated electrons will tend to return to the target at the most positive element in the neighborhood. This does not in any way affect the basic secondary emission curve shown in Figure 22-1, but if our interest centers not so much on the electrons knocked out of the surface but on the net gain or loss to the target, then we have to redraw the curve at and above 150 V to show that at such voltages the target does not in fact lose any electrons because of secondary emission. The curve drops to a secondary emission ratio of zero. This can be seen in Figure 22-2.

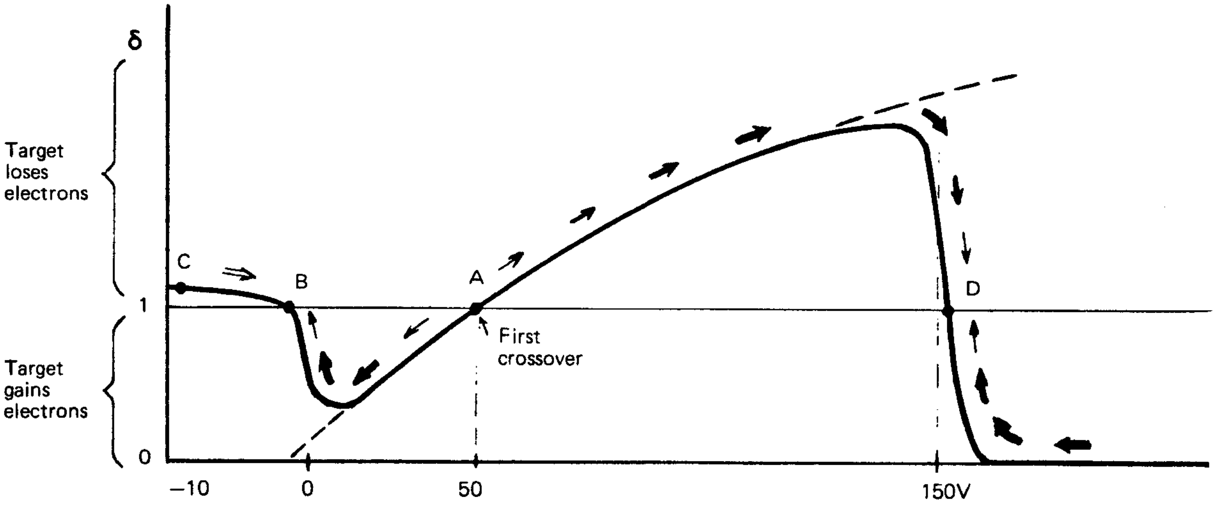

Target voltage (relative to cathode)

Figure 22-2

The second modification of Figure 22-1 occurs at the opposite end of the curve. Once the target voltage is below that of the cathode, the primary stream of electrons will not land on it any more but will go straight to the collector. Again this does not affect the validity of Figure 22-1. It is a fact that if any electrons did land they would land with zero energy and would be incapable of knocking off secondary electrons. But if we are concerned with the balance sheet of the target we will interpret the situation differently and say that since the target neither receives nor loses electrons and since therefore, in this trivial sense, the gains and losses exactly balance, we are dealing with a δ of 1. The way in which the new curve deviates from Figure 22-1 is again shown in Figure 22-2.

Figure 22-2, then, is not a curve representing the secondary emission ratio but one which plots the secondary emission yield (in other words the net gain or loss of the target) against the landing voltage of the primary beam. For simplicity, and to conform with literature, we will continue to label the ordinate δ.

A small point still remains to be explained: why the curve to the left of point B is shown slightly above the $\delta = 1$ line. If no electrons can land on the target because it is more negative than the cathode from which they originate, one would have expected the curve to remain at $\delta = 1$, representing neither gain or loss. In fact, within the stream of particles coming towards the target are occasional positive ions, and these will be attracted by the negative target and land on it. Since a gain of positive ions is equivalent to a loss of electrons, it must be shown on the balance sheet the same way as a loss of electrons-in other words, as if the secondary emission ratio were greater than one.

We can now return to the study of the voltage drift. The target is a collection of separate addressable points. As long as electrons arrive and leave in unequal numbers, these points will move up or down in voltage. A net loss of electrons, and therefore a drift in a positive direction, happens whenever the target voltage is in regions where the curve is above $\delta = 1$, and a net gain and negative drift in regions where it is below $\delta = 1$. This is shown by the arrows in Figure 22-2. Therefore, as long as the beam continues to hit a given target area, that area will charge in the direction of the arrows. If you study these directions for a moment, you will see that they converge on two points, B and D. These are the only two points at which the target can stabilize. (The target cannot rest at A since the unexpected gain or loss of a single electron due to noise will bring it under the influence of one or the other divergent trend.) B and D are called, appropriately enough, the lower and upper stable points.

The speed of the voltage drift is obviously a function of the amount of discrepancy between landing and leaving electrons.

Whenever the curve approaches $\delta = 1$, the movement will slow down, whereas in regions of large gains or losses the voltage will change more quickly. In Figure 22-2 we tried to make this point by varying the thickness of the arrows. The region between B and C is a special case. Drift in that part is due to the landing of positive ions rather than electrons, and since these are fewer in ratio of perhaps 1 to a million, the drift from C towards B is measured in minutes, compared with tens or hundreds of microseconds on other parts of the curve.

We said that this drift towards the stable states occurs in any part of the target, as long as that part has an electron beam directed towards it. Therefore, if the whole target were to be flooded with a defocussed electron beam, all those portions of it whose surface voltage happened to set above A would move towards the upper stable point and the remainder towards the lower stable point, and under the influence of this floodbeam the target would be maintained at these points. This would give us a device capable of bistable storage of information in the form of a voltage pattern.

In order to be useful, we must of course have means of entering and deleting information-in other words, of writing and erasing-and we must make this voltage pattern visible. If the pattern becomes visible because of light emission from the storage tube itself, we speak of a direct-view storage tube. (The other method of making the pattern visible is by scan conversion: scanning the target with a reading beam which is then used to modulate some other light emitting device such as a TV picture monitor).

We will not discuss in this chapter how the pattern stored on the target is made visible in a direct-view storage tube. There are, as you know, two entirely different methods of doing this, which will be explained in later chapters. But with both methods it is convenient to use a higher target voltage for the written information and a lower one for the unwritten background. "Writing" therefore means lifting the target surface by means of a focused writing beam from a lower to a higher voltage-in the case of the bistable system, from the lower to the upper stable point.

How could this be done? Well, the target consists of a dielectric, and in order to increase the voltage of a given point on it we must cause that point to lose electrons. The only mechanism we have available is the one just studied: secondary emission. Writing beam electrons must arrive with enough energy to cause a secondary emission ratio of more than unity. The writing beam can only have so much energy if it originates from a cathode sitting at a considerably more negative voltage than the target. One could, in principle, stop the floodbeam, move its cathode sufficiently negative and focus it, then start writing on the target. Afterwards the flood condition could be re-established. (We shall see that in both types of storage tubes the floodbeam is the source of electrons which produce the visible stored display. In bistable tubes it also has the vital function of maintaining the written and unwritten parts of the target at their respective stable points as explained earlier.

In practice it is found simpler to use two separate guns in the same CRT envelope: a permanently defocussed floodgun, which maintains a floodbeam at all times, and a separate focussed writing gun, operating at a much more negative voltage, whose electron beam is controlled by a control grid in the normal manner.

When writing the target area, the writing beam action is initially opposed by the continuing floodbeam action. The target voltage will only move positive if the number of electrons lost due to greater-than-unity secondary emission of the writing beam exceeds the number of electrons gained from the lower velocity floodbeam. This will be considered in more detail. Once the target has moved above the first crossover point, the floodbeam will of course assist the writing beam in moving the target further positive.

Finally in this chapter, let us consider how this information could be erased after it is written. This involves moving all those areas of the target which are written back to the unwritten level. The target itself is, as we keep saying, floating. But the dielectric is in fact mounted on some kind of conducting surface, and if a negative pulse is applied to this surface, capacitive coupling will also move the target as a whole, negative by this amount. Once all points of it have been lowered below the first crossover, the continuing floodbeam will see to it that the target is then maintained at the lower stable point. This description of the erase process is only a preliminary one. The erase pulse is in fact more complex to take care of additional problems, and we shall look at these when discussing the two tube types in detail.

BISTABLE PHOSPHOR—TARGET TUBE CONSTRUCTION

We mentioned in the introduction that direct-view storage tubes fall neatly into two types, depending on the means adopted to make the stored pattern visible. These types are the phosphor-target tube and the transmission tube. In the first case the target dielectric is made of phosphor which will light up in the written areas and be looked at directly by the eye. In transmission tubes, on the other hand, the target forms a mesh which controls the flow of the CRT beam on its way to a conventional phosphor screen, acting much like the grid in an ordinary vacuum tube.

The earliest phosphor-target tubes had poor definition and an extremely dim display. This led designers to concentrate on transmission tubes. They on their part suffered from lack of reliability and were expensive to manufacture. Then Tektronix engineers returned to the phosphor-target idea and managed to refine it into a practical product. Their phosphor-target tubes represented a breakthrough in price and simplicity. They first entered the market in the type 564 oscilloscope in 1963 and are the subject of Tektronix patents.

The basic idea is simple enough: if phosphor is used as the dielectric in a bistable system such as we described, then the stream of flood electrons hitting the written areas with an impact speed equivalent to the 150 volts or so of the upper stable point will produce a light output, whereas the unwritten areas at the lower stable point receive no floodbeam electrons, or if they did the landing energy would be virtually zero, causing no light emission. This phosphor target will therefore continue to emit light from the written area as long as the floodbeam is present.

But there are of course problems. First, the phosphor must be suitable as a dielectric, which means it must offer a high secondary emission ratio and possess good insulating properties. The most efficient phosphor, P31, does not have these qualities; we generally use a modified form of P1 with about half the efficiency of P31, which means a dimmer trace. Furthermore, with the phosphor target at about 150 volts, compared with the more usual several kilovolts, the trace brightness is again appreciable reduced. Nevertheless, under subdued lighting conditions it is still a usable display.

Then there was the problem of poor definition in early tubes. This was traced to the lateral spreading of the written area after the passage of the writing beam, due probably to inadequate lateral insulating properties of the phosphor. Our solution was to deposit the phosphor either as a pattern of finely spaced dots or to lay it down as a layer of randomly arranged semicontinuous particles with the aim of preventing lateral leakage between areas.

The particles, whether regular dots or of random shape, must of course form a pattern so fine that the width of the focussed writing beam will cover several of these target elements. In this way the limits of definition are dictated by the fineness of the writing beam only.

You will remember that bistable target operation depends on a nearby collector to collect secondary electrons. This discontinuous nature of the phosphor deposition allows us to use the conducting film on which the target is deposited (Storage Target Backplate in Figure 23-1) as a collector. This film is so extremely thin that it appears to be transparent, so that light emitted from the phosphor can be seen through it by the observer. Secondary electrons knocked off the target will therefore be attracted through the gaps in the phosphor to the higher potential collector.

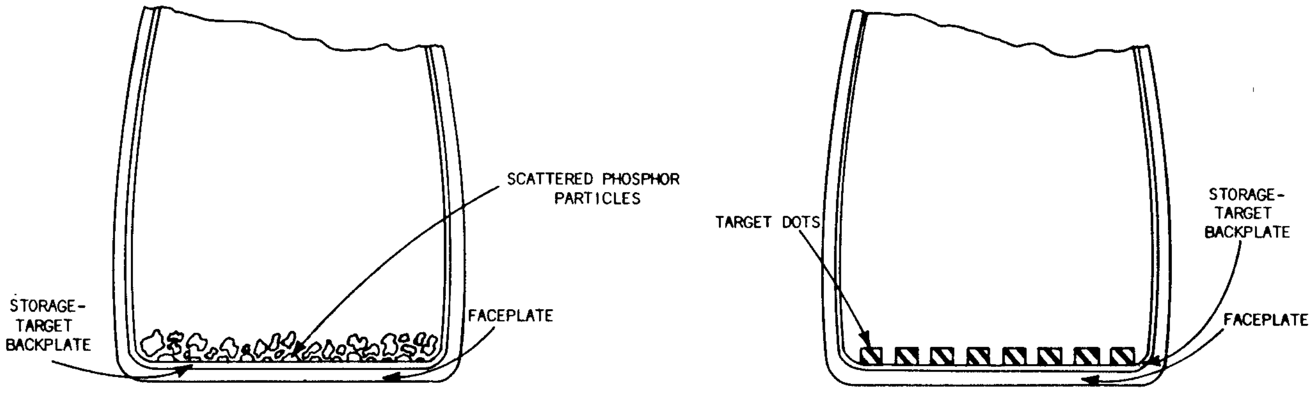
Figure 3A

Figure 3B

Figure 23-1

Perhaps we should consider briefly why the primary stream of flood electrons does not itself go directly through these gaps to the collector, thus defeating the whole purpose of the arrangement. The reason is that the flood electrons arrive with a fair amount of kinetic energy and are not easily diverted at the last moment to the minute gaps between phosphor particles. By contrast, secondary emission electrons have much lower energy and therefore move at much slower speed, which makes them more maneuverable. It is, incidentally, the difference between the high energy of the landing electrons and the lower energy of the secondaries which gets converted into heat and light emission from the phosphor.

Let's pause at this point to summarize briefly what we have learned of phosphor-target storage tube construction and operation. These tubes have a target composed of phosphor which can be written, i.e. lifted by a writing beam to a higher potential and will then attract electrons from a floodbeam whose landing energy is partly converted to light and partly used to dislodge secondary electrons. The secondary electrons find their way through gaps in the target to the storage target backplate which acts as a collector. The floodbeam is therefore used in the first place to make the written areas visible, but it also has the effect of shifting the target from whatever voltage it may have been left at by the writing beam or erase pulse to the upper or lower stable point, making this storage tube a bistable one. The beam originates from a floodgun and is deliberately dispersed to cover the whole target area. The writing beam comes from the writing gun which is so negative with respect to the target that when the writing beam lands it causes much secondary emission, thus lifting the target voltage. The writing beam is intensity-controlled, focussed and deflected in the normal manner.

With this basic picture in mind we must now go a little more deeply into the problems of target construction, since this will considerable increase our understanding of storage tube behavior. These are problems which are of concern at the design stage, but also have important effects on operating characteristics.

A suitable target material must be chosen. Then it must be decided whether to deposit particles according to figure 3A or 3B. In the first several years that Tektronix manufactured bistable storage tubes the semicontinuous method shown if figure 3A was predominantly used. More recently the target deposition process shown in figure 3B has become the most widely used technique. Having decided which material to use and which target deposition method to use we can then begin the process of optimizing the processes or materials to obtain the performance desired.

As target thickness increases a number of factors are affected in a beneficial way. Luminance increases fairly linearly, since the presence of additional material (and the higher operating voltage that this permits) generates additional light. Resolution increases rapidly at first: when the target is only molecules thick, wide spaces exist between particles and these fill in as thickness increases. Predictably, once a certain thickness has been reached, the increase in resolution levels off. But perhaps the most significant improvement resulting from greater target thickness is the increase in contrast, as shown in Figure 23-2.

Against this catalog of benefits resulting from increased target thickness we must set one factor which, after reaching a peak, decreases again. This is the stable range of collector operating voltages, and to understand what it is, and why it is so important to us that we sacrifice a great deal of contrast to it, we must consider one aspect of this storage tube which has been ignored to this point: the possibility of leakage from the target surface to the backplate on which it is deposited.

In unwritten screen areas there is in fact some leakage through the target, to the collector, which sits at a high positive voltage, lifting the phosphor surface above the lower stable point and causing a slight amount of light emission because of the increased landing energy of the floodbeam. It might seem to contradict basic theory that the target can rest at a point above the lower stable point, since the secondary emission ratio is then less than unity and it ought to gain electrons, but this effect is balanced by the migration of electrons through the target to the collector.

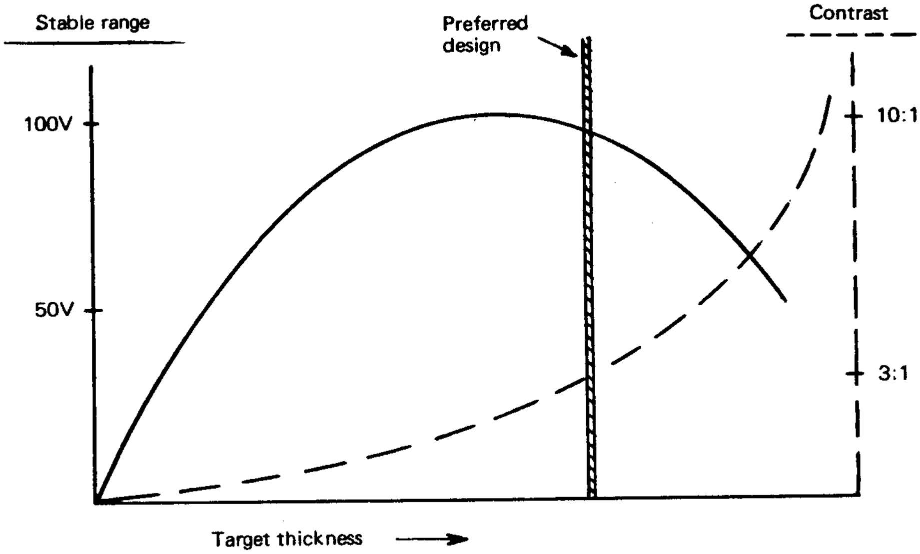

Figure 23-2

The amount of leakage will vary from point to point across the screen, since the phosphor layer is randomly semicontinuous, but some leakage will be observed almost everywhere and we can surmise that the rest potential might look something like the solid line RP in Figure 23-3. This will cause light emission to vary across the screen in a correlated manner. From normal viewing distances these variations average out and we simply observe an average background light level, corresponding to the average rest potential (ARP).

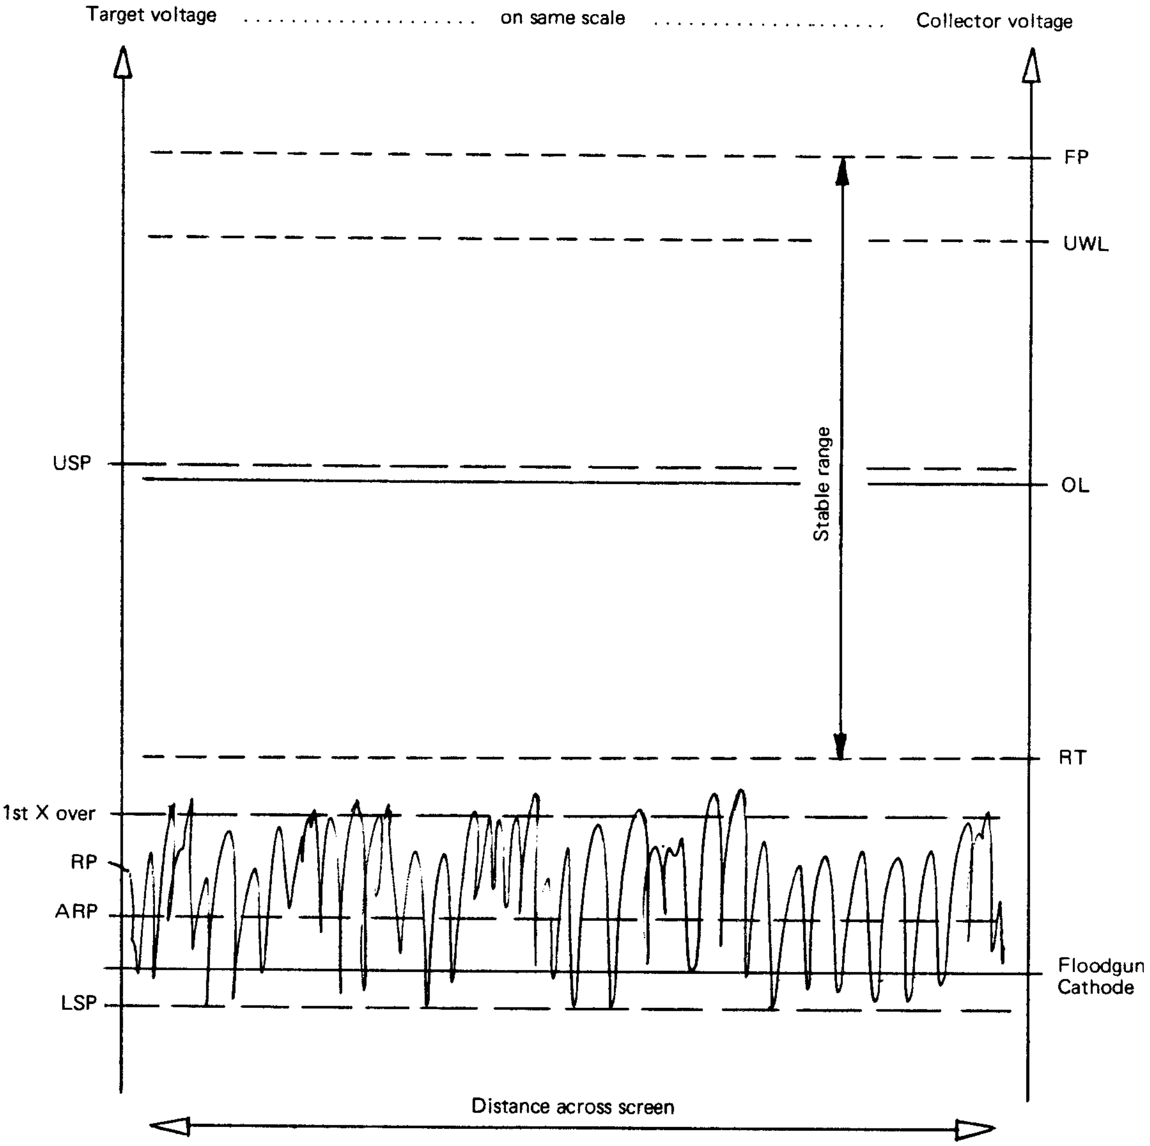

Figure 23-3

The solid line RP is no more than an artist's impression, but given such wide variations across the target, some points will inevitably exceed the first crossover level, and these points will therefore move to the upper stable point, a process which is often called "fading positive". Being individual, randomly distributed bright dots on a microscopic scale we can again see only their contribution to the average background light level.

Although on theoretical grounds one might wish to exclude these written dots from the calculation of the average rest potential, in practice this is not possible. The average rest potential is a purely theoretical value which cannot be measured directly since the target is floating. We assess the average rest potential on the basis of average light emission, and when making such light measurements we are bound to include the written dots as well as those in various unwritten states.

The full picture, then, is that dot by dot across the screen the rest potential varies in a random manner, causing a corresponding slight light output, with the exception that all those dots which happen to exceed the first crossover level will fade positive and emit the written light level. Only the average of all these light contributions can be perceived on a macroscopic scale, and from this average light level we can deduce the average rest potential.

The situation is illustrated in Figure 23-3, in a purely qualitative way, for the condition where the collector voltage is set to a typical operating level. Naturally, as the collector voltage is varied up and down, the amount of leakage also varies and the RP curve will shift up and down to some extent.

If we set the collector to increasingly positive levels, a point will be reached where spreading of the written trace occurs because areas adjacent to it are so near the crossover that capacitive effects or local dielectric breakdown are sufficient to make them fade positive. This collector voltage is known as the upper writing limit (UWL). At some still higher level, so much of the RP curve lies so near the first crossover that the whole screen will sponteneously fade positive. Both the collector voltage levels are shown in Figure 23-3.

Turning now to the consequences of decreasing the collector voltage below the operating level, we must recall that the upper stable point of the target always occurs at a voltage in the vicinity of the collector voltage, since it is the failure of the collector to collect which causes the abrupt drop in the target "balance sheet" curve of Figure 22-2. Now if the collector is lowered to the vicinity of the first crossover voltage, this will result in a curve as shown in Figure 23-4, and it is clear that under these conditions there is only one stable point, the lower stable point. The floodbeam will return all target areas to the lower stable point; written information is no longer retained. This collector voltage is therefore called the retention threshold.

Now we can define the stable range: it is the range of collector operating voltages between retention threshold and fade-positive. It is this stable range which is affected by the thickness of the target in the manner shown in Figure 23-2. In itself it will not concern us operationally, since we would be unwise to operate the collector near either of these extreme limits. But a large stable range will obviously provide a greater operating margin for the collector voltage. This margin is important to us for several reasons:

1. Setting the collector voltage operationally to the center of this range is a subjective procedure which will yield a certain spread from operator to operator;
2. In many instances the CRT heater is unregulated, and varying line voltages can cause performance changes (this is not as much of a concern in newer instrument designs as it was in earlier instruments);
3. Storage CRTs are subject to aging effects which might, if the operating margin is too small, require frequent recalibrations.

4. Even with the best manufacturing techniques there is usually some non-uniformity across the target, calling for different optimum collector voltage settings, and in the presence of a large operating margin the choice of a suitable compromise setting is much easier.

For all these reasons we consider a large stable range so important that we sacrifice much contrast to obtain it, as suggested by Figure 23-2.

When contrast was first mentioned as a significant factor in connection with Figure 23-2 you may have been puzzled since it is normally taken for granted in oscilloscopes that unwritten areas of the screen are practically black and the contrast therefore practically infinite. The discussion of the average rest potential will have explained why, on phosphor-target storage tubes, the contrast is quite limited. But although Figure 23-2 shows a typical contrast ratio of only 3:1, some improvement can in fact be expected after a few hundred operating hours. The reason is that much background light is contributed by those dots which have faded positive, and as these phosphor dots operate continually at full light output they will be the first to age and eventually turn dark, leaving the unwritten part of the screen much darker. On most tubes the contrast ratio will reach 20:1 after about 500 hours. Contrast ratios in the order of 100:1 are possible after several thousand hours of operation.

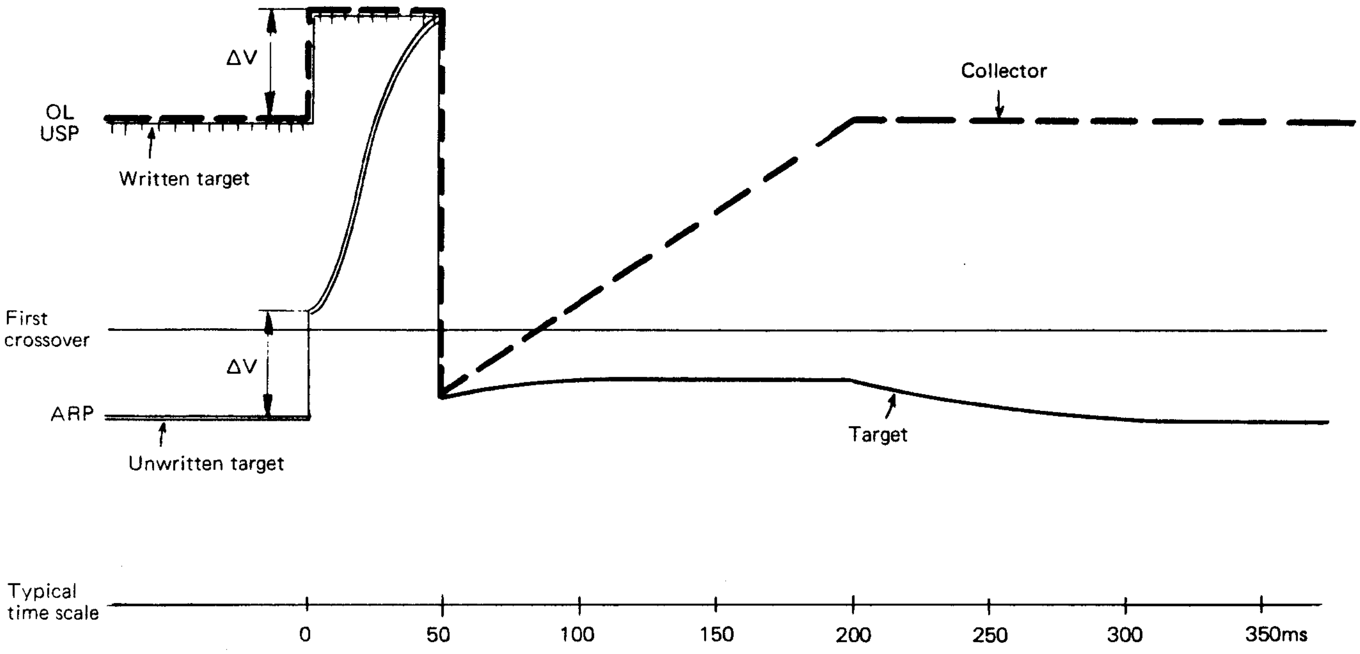

Figure 23-4

OPERATING CHARACTERISTICS OF THE PHOSPHOR-TRAGET TUBE

One of the main limitations of a storage tube is its inability to store traces if the beam is moving too fast-if it exceeds the maximum writing speed. The bulk of this chapter will be concerned with the definition of writing speed, what factors influence it and how it can be improved. Then we shall return to the topic of erasing and see in detail how it is done.

In a bistable tube, writing, as we know, is the process of raising the voltage of those points on the target which are scanned by the writing beam above the first crossover despite the continuing attempts of the floodbeam to return them to the rest potential. (Once the critical first crossover level has been passed, the floodbeam will carry them to the written level even without any further contribution from the writing beam.) The effect of the floodbeam is to add a given number of electrons to unit target area in unit time. But this number depends on the secondary emission ratio and is highest where the "balance sheet" curve departs most from the $\delta = 1$ level, trailing off to zero as the first crossover is approached. Since we can neither measure the secondary emission in an actual CRT, nor even be sure from what rest potential the target must be lifted, it is impossible to quantify the demands made on the writing beam if it is to achieve storage.

But the effect of the writing beam itself is also far from straightforward. Consider first the situation of a stationary beam. Even though it is focussed, the spatial distribution of beam intensity follows the normal Gaussian distribution curve shown in Figure 24-1. At the point on the target where it peaks, the beam density per unit target area is greatest, hence the number of secondary electrons lost in unit time is highest.

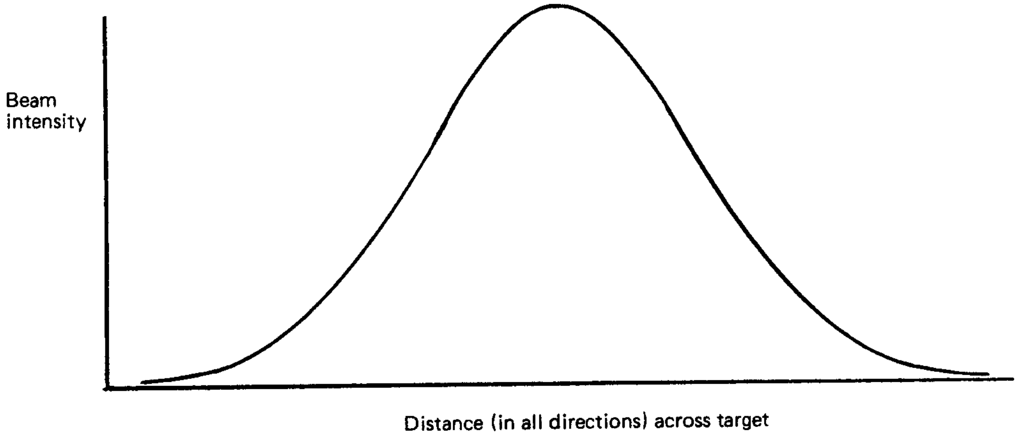

Figure 24-1

If this number exceeds the number gained from the floodbeam action the target will begin to charge up. However, the charging process takes time and relies on the continuing presence of the writing beam if it is to reach a successful conclusion, namely that the target voltage passes the first crossover. With greater beam density, the disparity between electron loss due to writing beam and gain due to floodbeam increases and a shorter beam dwell time is enough to achieve storage.

Away from the center of the writing beam, since the beam intensity decreases, the number of electrons lost per unit time by the target will also decrease. As long as it is still greater than the gains made from floodbeam action, the target will still move positive, but it will require a longer beam dwell time to reach a successful conclusion.

Let us review the picture given in the last three paragraphs, and assume for simplicity that the target rest potential is at point B of Figure 22-2. To achieve storage, the requirement is that the center of the writing beam (where its intensity is greatest) should cause the target to lose more electrons per unit time than it gains from the floodbeam, and that the writing beam should dwell long enough at that spot to cause the resulting positive target drift to reach the first crossover. We can instinctively feel that something like the product of dwell time and beam intensity is significant here, but there is a certain minimum intensity below which no amount of dwell time will achieve storage because the target gains more electrons from the floodbeam than it loses from the writing beam. it would be misleading to try to quantify this complicated situation in a formula, but we will refer to the dwell time-intensity product in this loose sense later in the text.

One last consideration, if we start with the minimum dwell time and beam intensity which will just achieve storage at the beam center, and then increase either factor, areas away from the center of the beam will also manage to reach the first crossover. As dwell time or intensity are increased we therefore obtain a stored dot of increasing diameter.

A stationary beam such as we have discussed is typical of information display devices. For instance, a computer readout monitor might display alpha-numeric characters, each of which is made up of selected written dots in a 7 X 9 dot matrix. Figure 24-2 shows an example.

If we want to specify the minimum requirement to achieve storage, we generally adjust the beam intensity to the maximum value before defocussing occurs, and then measure and quote the necessary dwell time. This specification is called "dot writing time", and a typical value is 5 microseconds or less. You will appreciate that the beam intensity setting is a somewhat subjective adjustment.

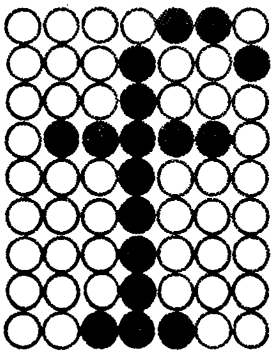

Lower-case f in a dot matrix

Figure 24-2

In oscilloscopes, the beam is normally moving and we must now study this new situation. If a given spot on the target lies in the path of this beam, then as the beam approaches, its intensity will increase in a manner which corresponds to the slopes of the beam current distribution curve. It will reach a peak when the beam is centered on the spot, and then decrease in a similar manner. But whether storage will take place depends on the same considerations which we enumerated previously, i.e. whether the maximum beam intensity is great enough and the dwell time is long enough. In this situation quantitative analysis is even more futile. Specifications are verified by again selecting the highest beam intensity before defocussing occurs, and increasing the beam velocity until the beam moves so fast that there is insufficient dwell time for storage to occur. This specification is called "writing speed" and is typically, for phosphor-target tubes 0.1 cm/μs.

For the moving beam we can say that if the dwell time is made longer by moving the beam more slowly, areas to the side of the central path of the beam will receive a sufficient dwell time-intensity product to become written. As the beam is slowed down we therefore get a progressively wider stored trace.

At the end of this discussion we hope that you will have an instinctive feeling for the principal factors affecting dot writing time and writing speed. We will now consider in what way the writing speed, and also the brightness and contrast of the stored display, are affected by the collector operating voltage.

The published specifications assume that the collector operating level (OL) is set normally, let us say to the center of the operating stable range in Figure 23-3. As we increase the collector voltage, leakage increases, the average rest potential increases, and consequently the target rests nearer to the first crossover. This means that a lesser dwell time-intensity product will suffice to achieve storage; holding the intensity constant we can increase the beam velocity and still store. The writing speed specification has been improved. But the improvement is not spectacular and the change of collector voltage has other side effects which are more important and which we will look at shortly.

If the collector voltage is decreased the opposite effect takes place. The ARP drops and the writing beam must linger longer to achieve writing. In fact, for a specified beam velocity, if the collector voltage is decreased sufficiently, a level will be reached at which the dwell time-intensity product is no longer enough to achieve writing. This collector voltage limit is called "writing threshold" (WT). Unlike all other collector voltage limits (FP, UWL, RT), this one is not a limit due to basic construction features of the tube; it is dependent on the beam velocity which we specified.

For such a specified velocity, the writing threshold represents the lower limit of the collector voltage operating margin to which we referred earlier. Neither can we operate successfully above the upper writing limit since trace spreading occurs. This defines the collector operating range and is shown in Figure 24-3. A writing speed specification is only realistic if it puts the writing threshold in approximately the position shown in Figure 24-3, giving a usefully large operating range (typically operating range is $>$10 V).

Now to the other effects of departing from the normal collector operating level. We said that as the collector voltage is raised, the ARP goes up. Therefore the light level of the unwritten area will increase. But also, since the upper stable point follows the collector voltage up, the brightness of the written trace increases. The converse is true when the collector voltage is decreased. We must now consider whether, on the average, these effects produce traces with more or less contrast, and whether, if we have a choice, it is important to get the maximum possible contrast or the maximum possible absolute light output. (Contrast as defined here means the brightness ratio of written to unwritten areas.)

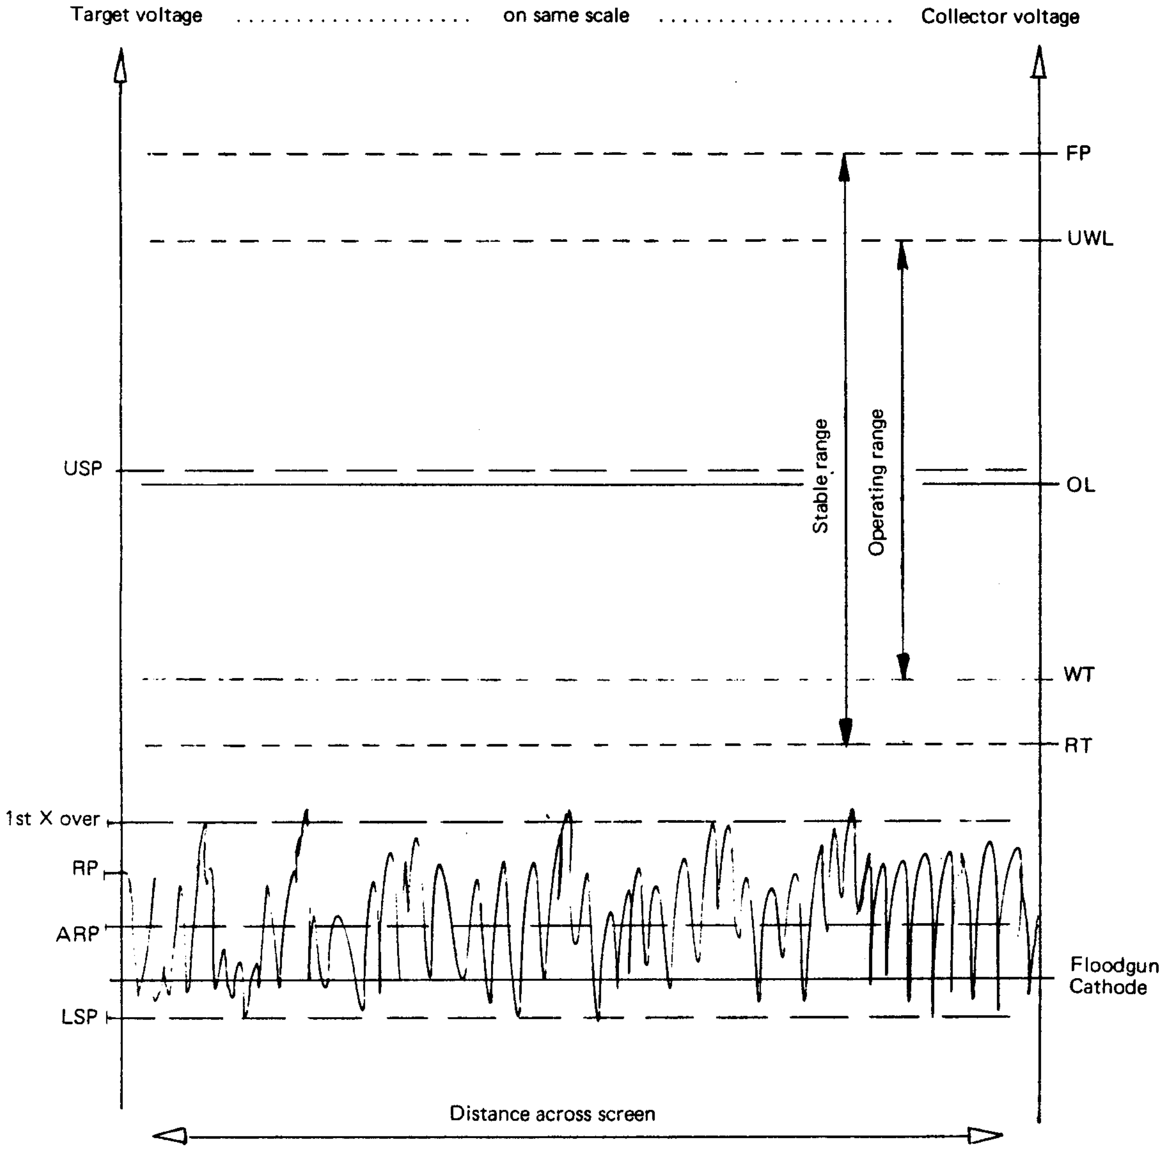

Figure 24-3

The simplest way to study this matter is by way of a numerical example. The figures in table 1 have been chosen for convenient mental arithmetic, but they are typical for the actual performance of phosphor-target tubes. The set of figures headed "in total darkness" shows the actual light output of the CRT. As you can see, the brightness of the unwritten areas increases more rapidly with increased collector voltage than the brightness of the written trace, so the contrast gets poorer. In the next two sets of figures, 6 and 100 candela /m^2 respectively have been added to the CRT light output figures to show what happens when these amounts of ambient light are reflected from the CRT. The contrast decreases, but it decreases least if the CRT light output is high, because the ambient light cannot swamp out the tube light as easily.

Collector Voltage	in total darkness			with 6 cd/m^2 ambient light			with 100 cd/m^2 ambient light		
	written trace cd/m^2	unwritten areas cd/m^2	Contrast	written trace cd/m^2	unwritten areas cd/m^2	Contrast	written trace cd/m^2	unwritten areas cd/m^2	Contrast
100 V	24	4	6 : 1	30	10	3.0 : 1	124	104	1.2 : 1
140 V	50	10	5 : 1	56	16	3.5 : 1	150	110	1.4 : 1
180 V	80	20	4 : 1	86	26	3.3 : 1	180	120	1.5 : 1

Table 1

Which is preferable? To see the trace at all, we need contrast and the more we have the better. But it turns out that for different ambient lighting conditions different collector voltages will give the best contrast, so no hard and fast rule is possible. Photography, of course, takes place in total darkness as the camera shuts out all ambient light, and would therefore benefit from a low collector voltage.

It has already been said that the improvement in writing speed which can be achieved with higher collector voltage is only marginal. There are two other techniques, however, which are capable of increasing the writing speed by a factor of 10 or more. These will now be discussed.

To understand how they work, we must first visualize what happens when the beam moves faster than the maximum writing speed and fails to store. In such a case, the dwell time-intensity product is not enough to raise the target voltage above the first crossover, and as soon as the writing beam has passed, the floodbeam begins the destructive process of moving the target back to its rest potential. Nevertheless, the writing beam did raise the target above the rest potential. The secret of the two techniques is to make use of this charge pattern before the floodbeam destroys it.

The first technique is useful on repetitive sweeps, and is called the "integrate" mode. By stopping the floodbeam altogether, the destructive process can be halted. Any charges laid down by the writing beam will remain on the target, if not indefinitely, at any rate for minutes. If the signal is repetitive, successive beam sweeps will scan the same target areas and will add to the charge patten. This is a cumulative process which must eventually lead to the point where the written target areas cross the first crossover point. If the floodbeam is then restored it will move these areas to the written state and the trace can be seen.

For a given beam velocity it will take a given number of beam sweeps to build up this charge. We have already seen that these factors are not amenable to numerical analysis. In practice it is a matter of trial and error to find out how long the integrate mode has to be applied before a trace gets stored. There is unfortunately no way of knowing, while the mode operates, whether storage will be achieved, since without the floodbeam we cannot see a stored trace. When integration is stopped, and by the time we have inspected the display and perhaps decided we didn't integrate long enough, the portions of the trace which failed to write will have been moved back to rest potential. Any second attempt at integration will therefore be starting again from square one.

But imagine now that we wish to store a single transient, some unique event, at a speed exceeding the normal writing speed. Since we cannot repeat the event, the integration technique is useless. Yet even that one sweep did leave some charge behind. The second technique, called "enhance" mode, again attempts to salvage the situation. A positive pulse is applied to the collector, of such amplitude that capacitive coupling will lift the whole target by just the amount needed to bring the written area above the first crossover. Figure 24-4 makes this clear. The floodbeam will then immediately set to work separating the written and unwritten potential further. We maintain the positive pulse long enough to ensure that at its end the written areas do not drop back below the first crossover.

The curvatures recall the fact that the floodbeam is most effective at voltages where the secondary emission ratio departs most from unity, and floodbeam action slows down as a δ of 1 is approached.

Figure 24-4 also makes the point that immediately after the beam passage the floodbeam starts removing the laid-down charge. The enhance pulse must therefore be applied as soon as possible-in other words, as soon as the sweep is completed. But on slow sweeps, say 5 μs/div. or slower, even this may be too late. The enhance pulse will only rescue the later portions of the trace while those near the beginning of the sweep will already have been partly or wholly destroyed by the floodbeam.

Nevertheless, if enhancing were that simple one would have to ask why the technique is not made a permanent feature of fast sweeping storage, giving at a stroke a tenfold improvement in writing speed. But Figure 24-4 is oversimplified in an important aspect. The average rest potential is a fictitious level, and the actual target rests over a broad range of levels. When the writing beam adds a charge to this, the written areas, too, will end up over a broad range of levels. There will therefore be no one correct amplitude of enhance pulse which can raise all the written, and none of the unwritten areas above the first crossover.

In fact, the smaller the charge left behind by the writing beam, the more likely it will be that even with optimum enhance pulse amplitude some written parts will remain unstored, and some unwritten parts will become stored. The exact amplitude then becomes a matter of experimentation until the user subjectively feels that he has achieved the best compromise, making for the clearest display.

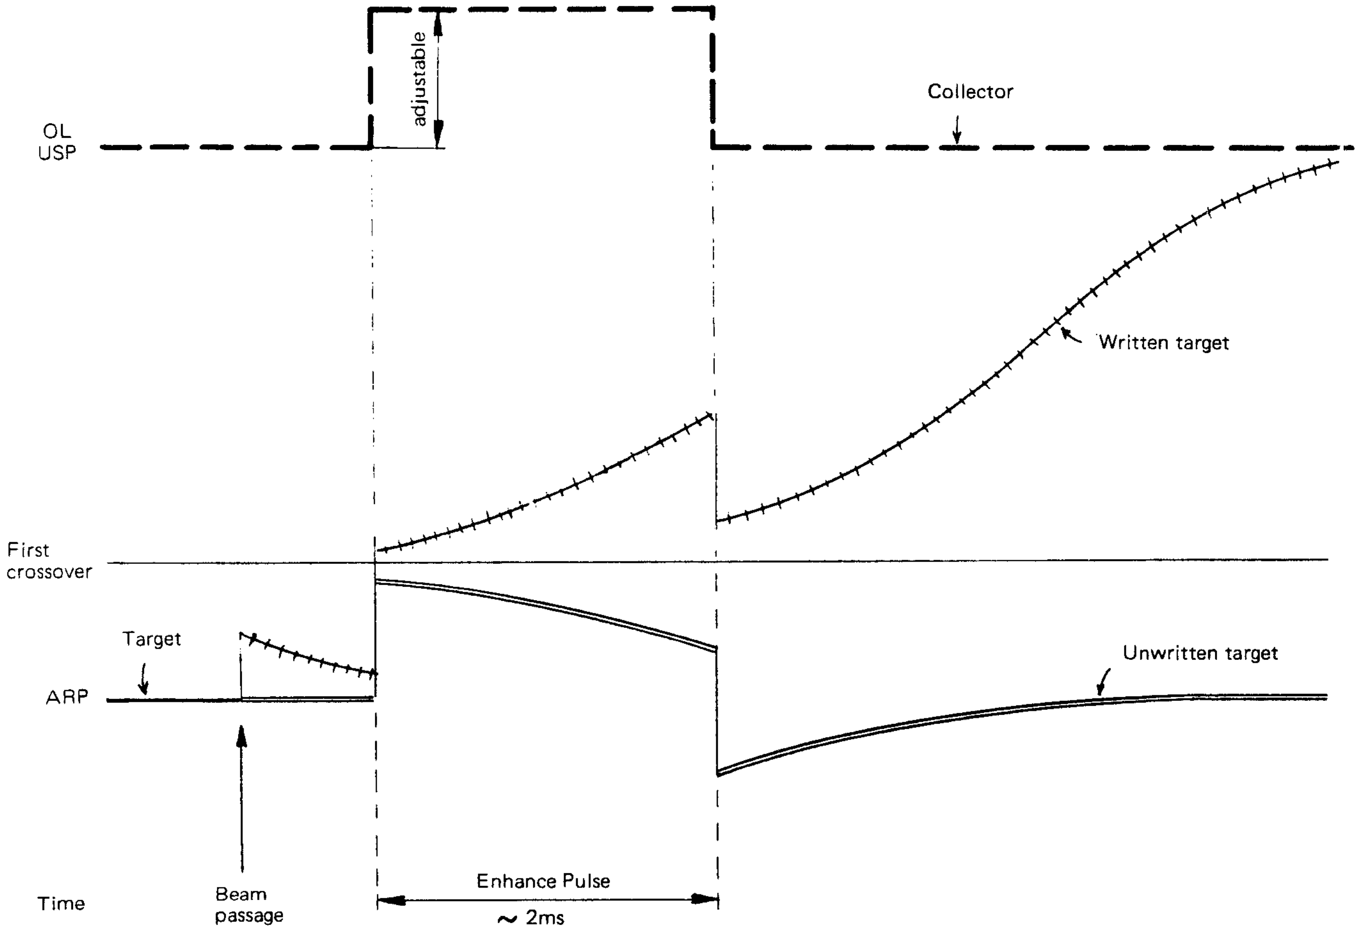

Figure 24-4

When we said that the enhance technique allowed a tenfold increase in writing speed, this was meant as a guideline only. In any given situation it depends on the kind of compromise the user finds acceptable. By contrast, the integrate technique really has no upper speed limit; it just depends on whether you can afford enough time to integrate long enough to accumulate enough charges to reach the first crossover. In cases where the signal repetition rate is 1 Hz or so and the required sweep speed is very fast, this can become a question of operator patience.

The next topic in this section is the erase process used in phosphor-target tubes. Basically, the erase pulse is a negative pulse applied to the collector, which capacitively moves the whole target negative. The aim is to move the written portions from the upper stable point to below the first crossover, after which the floodbeam can complete the erasure. But there are two problems. The first arises from the fact that sooner or later we will have to return the collector back to its normal operating level, and if we do this too fast we will capacitively move the target back up. This is true even if the negative pulse was long enough to give the floodbeam a chance to stabilize the target at the rest potential, because the voltage separating rest potential and first crossover is much smaller than that between first crossover and operating level through which the collector must move. The solution is to make the trailing edge of the erase pulse so slow that any capacitive coupling effects on the target can be countered by floodbeam action.

The other problem with erasing is that when small written areas are surrounded by large unwritten areas, and the target is capacitively lowered, the unwritten areas will move to a potential which is so greatly negative that the floodbeam is totally repelled from the target. The small written areas are in effect then shielded from the floodbeam and not returned to rest potential. At the end of the erase pulse they can easily become written again. Since small written areas among large unwritten ones are typical in normal storage tube use, this cannot be tolerated. Shielding effect of this kind can be avoided if the whole target is first written and then the erase pulse applied. So the erase pulse proper is preceded by a so-called fade-positive pulse large enough to lift the unwritten areas above the first crossover (V in Figure 24-5). Figure 24-5 shows the complete sequence.

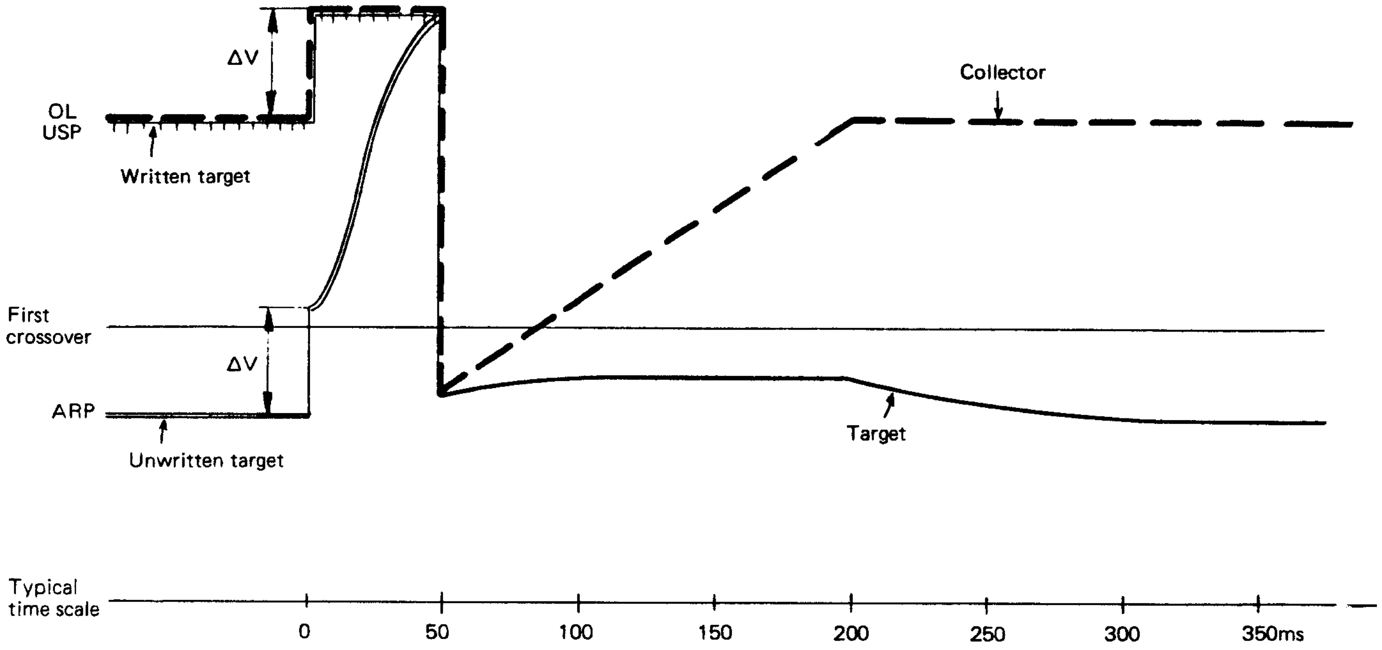

Figure 24-5

It sometimes happens that after the prolonged stored display of a trace the target resists erasure. During such periods the charge pattern becomes buried deep inside the target, rather than just on its surface. This is not to be confused with a phosphor burn in a conventional tube, where the phosphor is permanently damaged due to excesive heating of the phosphor during the passage of an intense writing beam. But though the phosphor is undamaged, unless the buried charge can be removed, an irritating residual image will remain. In the first place, to avoid buried charges, most instruction manuals give a time limit for storage, typically 1 hour. But if a buried charge does appear, and cannot be cleared by repeated erasures, the whole target may be written and left in this state for about 10 minutes, after which erasure should be successful.

The stored time limit just mentioned serves another purpose. As we said earlier, dots which sit at fade-positive for prolonged periods will obviously be the first ones to burn the phosphor. Parts of the screen which are frequently written will become dimmer and the tube will eventually fail for writing speed or poor appearance. One should avoid displaying the same waveform in the same position day after day. Second, it would be prudent for this reason to limit viewing time to no longer than is necessary.

An alternative solution is to reduce the floodbeam. This will result in a dimmer display and reduce the aging process somewhat but may still be sufficiently bright to be useful. Some oscilloscopes have a storage brightness control with which the floodbeam can be adjusted between 100% and 10%. (At the lower end, the floodbeam is so weak that it allows the target to accumulate charges from successive sweeps as in the "integrate" mode, provided the sweeps follow one another at intervals not much longer than 1 ms.) On information display monitors, viewing time at full brightness is often limited to 90 seconds, after which the display automatically goes into a low-brightness standby condition until the operator requests another viewing period. This condition is called the "hold" mode.

This completes our description of the storage characteristics of the phosphor target tube. But a word is in order about using it in the non-store mode. To stop the storage effect we simply have to set the collector below retention threshold. The tube then behaves like a conventional non-storage tube. No matter how high the writing beam charges the target, in a matter of milliseconds-before the eye can see it-the floodbeam returns it to the rest potential. With the collector below RT, leakage will be very small and the average rest potential so low that the screen is completely dark.