TRANSCRIPT of SOUND TRACK

from

TEKTRONIX TRAINING FILM:

SAMPLING OSCILLOSCOPES and TECHNIQUES

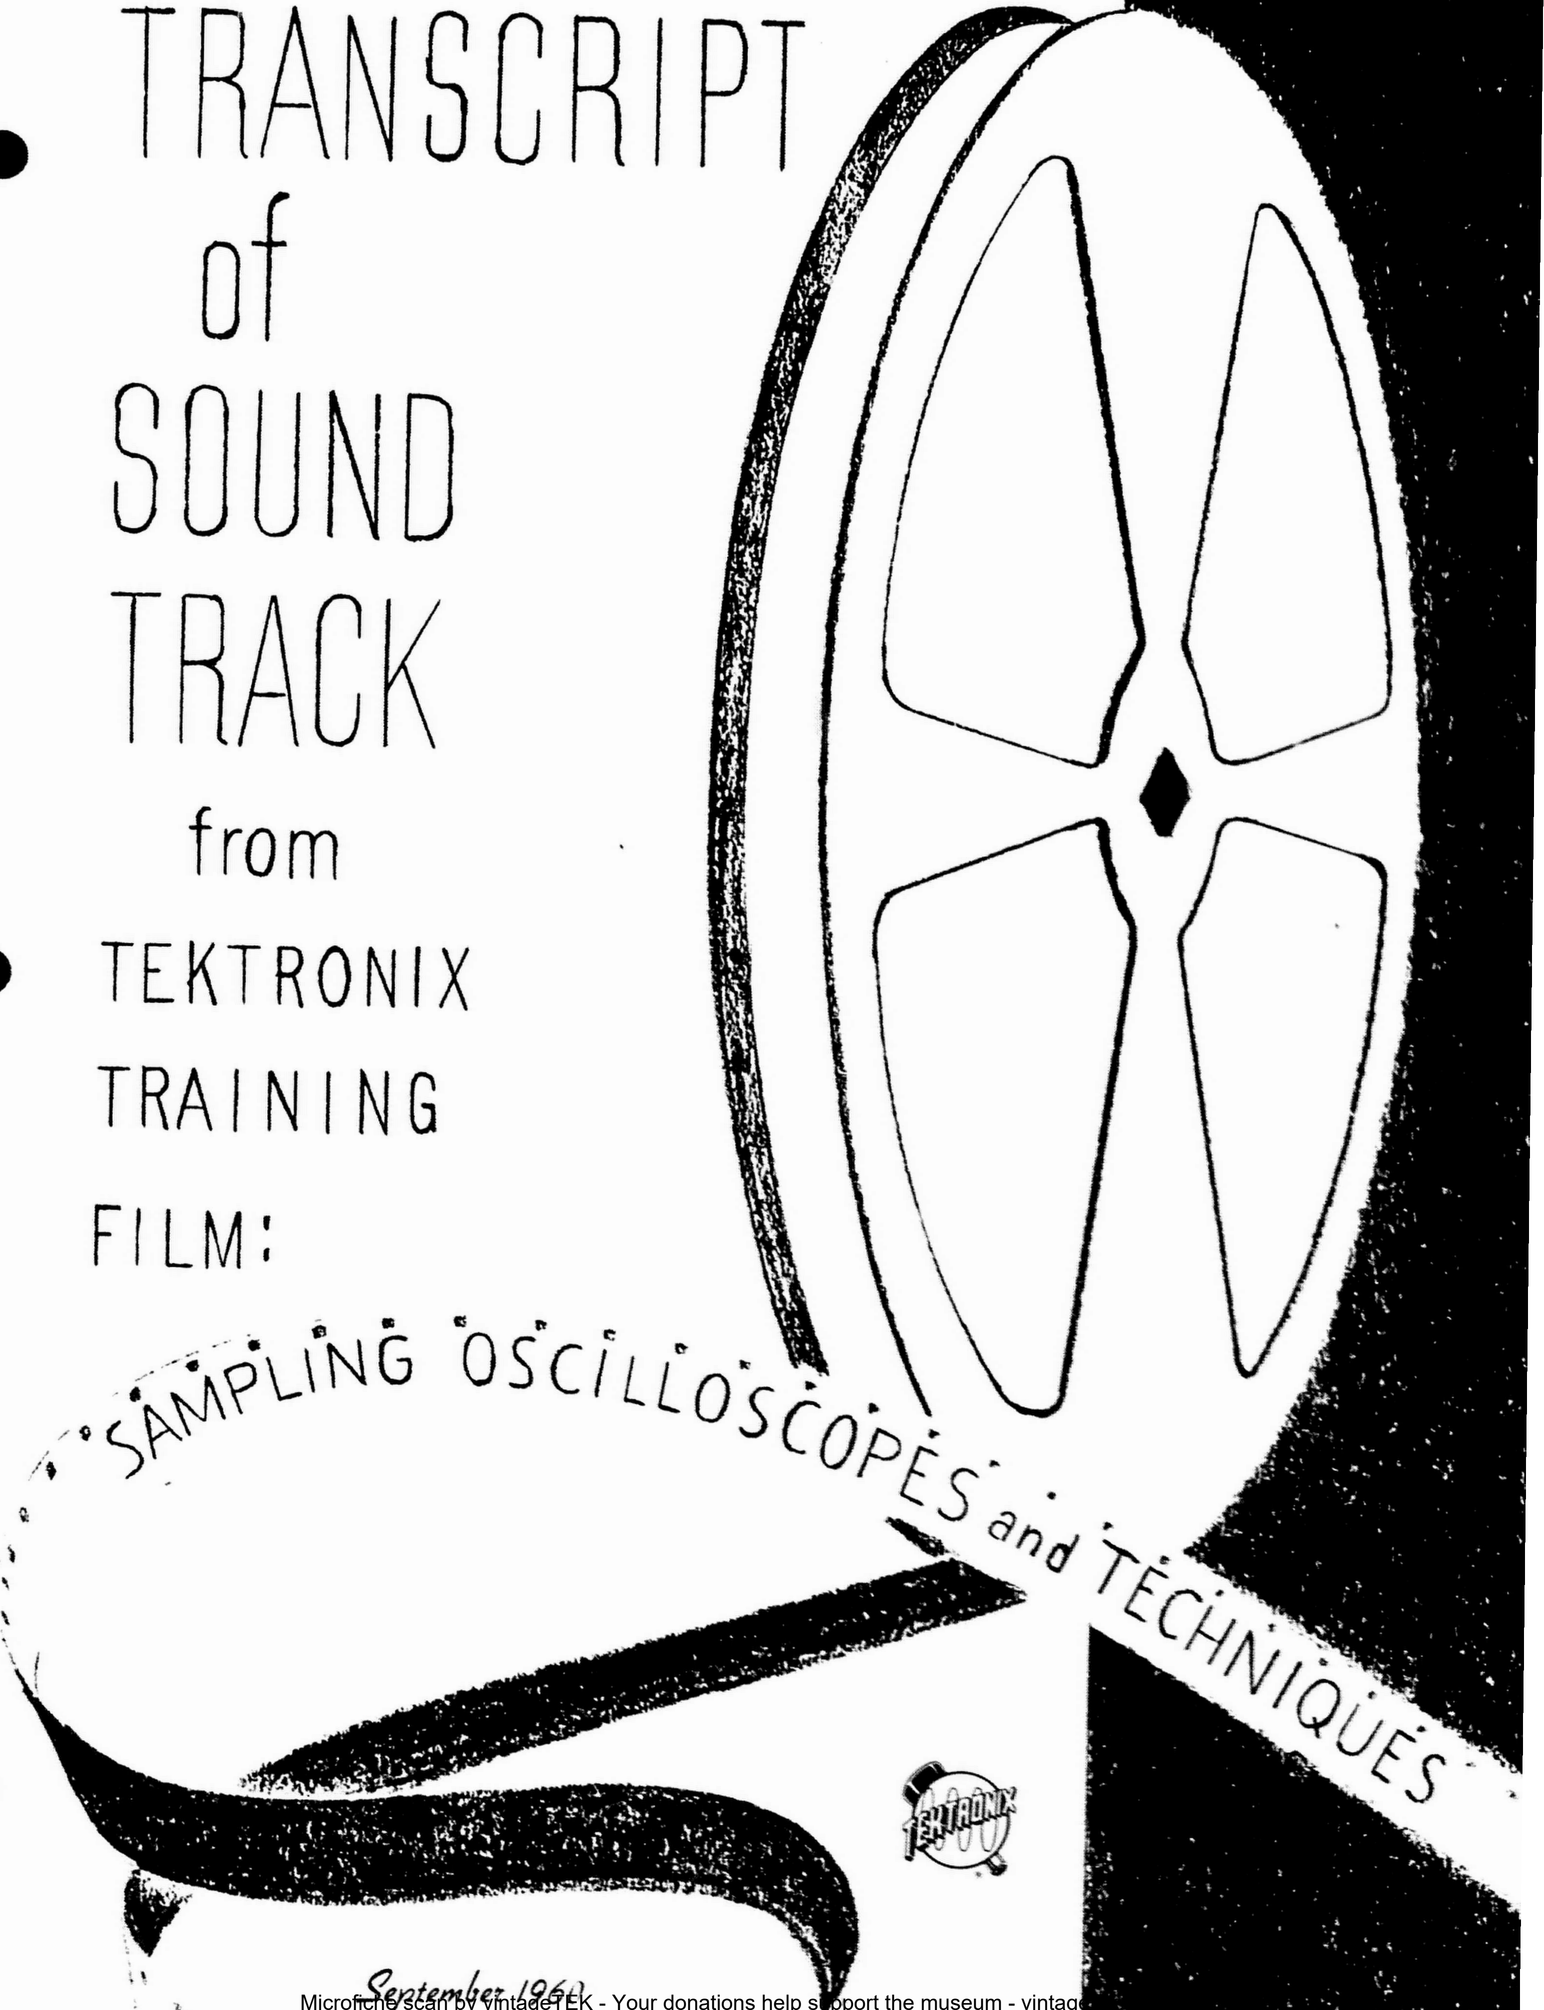

September 1960

SAMPLING OSCILLOSCOPES AND TECHNIQUES

by Norm Winningstad

Preface: In order to get information on the Type N Plug-In Unit to the field as rapidly as possible, a motion picture was made of an engineering meeting on the Type N Unit. The following is a reproduction of the sound track of that movie, slightly edited, corrected, and brought up to date. Because it is desired to reproduce the information as soon as possible, it tends to read like a lecture—which it was—and was altered only where necessary to avoid ambiguity. Photos were made, where necessary, of what was displayed on the crt.
C.N.W.

The Tektronix Sampling System is going to be explained today in a rather brief form, and we'll try to cover the points of the system which don't usually appear in detail in instruction manuals and go over some of the concepts of the sampling system, which will help in understanding sampling 'scopes in general. It is quite a change to go from conventional oscilloscope operation to sampling 'scope operation, and there are several traps that you can fall into, and we will try to expose these traps so that we can more easily understand the presentation in the sampling 'scope scheme.

This is the first part of a 2-part presentation—we'll discuss the sampling schemes, the general block diagram of the Tektronix N Unit System, and go through the adjustment of the N Unit so that you will feel familiar with this. The second part, which will follow, will cover how we make use of the equipment to insert signals into circuits, take signals out of circuits, and how to arrange the timing of the different units. The sampling concepts that we wish to go into first are the thoughts of the tie-up between space and time. We are generally dealing in speeds that are in the order of 10^{-9} seconds, or a nanosecond. Light travels about one foot in that length of time, so it is a great help to think in terms of signals traveling in coaxial cables as having positions that are related to time. The usual oscilloscope presentation is a graph that essentially involves a plot of voltage against time. Now, what you didn't realize, perhaps, was that you are at a fixed position in space making an observation with variable time. You could just as well have taken a photograph of the signal inside the cable at a particular instant in time, freeze the action, make the energy stop in the cable, and measure the voltage as it exists in the cable. Then you would find there's somewhat of a difference between the two. In the case where we are using a conventional oscilloscope system, the leading edge appears first and then the trailing edge appears in our conventional diagram (figure 1a). But, if we looked, instead, at the position of voltage in a cable by freezing time and looking at many different positions, the order of the leading edge and the trailing edge would change and the direction of the propagation would be where the leading edge is, and the trailing edge would be further back from the direction of propagation (figure 1b.) On this basis then, we could get into a very simple view of the sampling concept by imagining that we have a coaxial cable which has holes drilled in the side and a very special type voltmeter involved. This voltmeter has the property that when it is told to, it can measure the instantaneous voltage on the cable at that particular position at that particular instant in time. So, if we tell all the voltmeters to read at once, and remember what they read, we have in effect taken a photograph of the pulse as it was traveling down the cable (figure 1c). We could then take readings from the voltmeters, plot the readings on a graph, and reconstruct the pulse by this method (figure 1d). Now, this would be what you would call a single transient sampling oscilloscope, because what we have done is with a single pulse traveling down the cable, we have taken readings that are samples of the pulse as it went by, and reconstructed these to produce the voltage as a function of space. Now, we had an overall distance where we had our voltmeter stations, and to convert this into time, we have to know that the space equivalent time is the distance divided by the velocity of propagation on the cable. Each position of the pulse in space is not only a position in space, but it is also a position in time, if you will, because the pulse travels with a velocity fixed by the materials used to insulate the conductors of the cable. Well, we can then, from position in space plot, reconstruct a diagam based on time, and

SAMPLING CONCEPTS

FIG. 1a

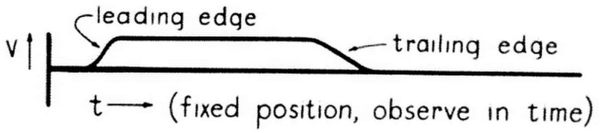

FIG. 1b

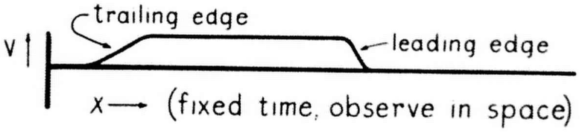

FIG. 1c

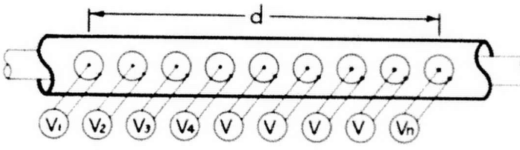

Equivalent Sweep time $= \frac{d}{v_p}$

Resolution $= \frac{d}{n v_p}$

FIG. 1d

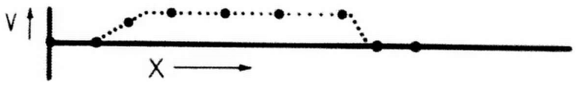

FIG. 1e

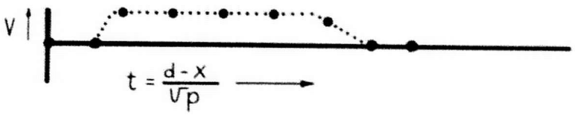

then we can have a display which is more nearly related what we are used to—the leading edge on the left side and the trailing edge on the right side (figure 1e.) Now, the important concept to notice here is that the total sweep time we will be displaying depends upon the total separation in space of our sampling stations, and the resolution of our plot is dependent upon the spacing between the individual stations. This, of course, shows why a single transient sampling 'scope in this form isn't practical. To change the equivalent sweep time we would have to change the physical spacing of the sampling units, and this is a difficult operation. So a single transient sampling oscilloscope in this form would not be a practical device for general-purpose use.

The next step is to notice that we can go to a repetitive sampling system—and in the repetitive system now, we are going to have only one voltmeter station on the side of a cable. Instead of examining the pulse at many positions in space, we will only examine at one position in space. But when we do this, we must examine the pulse at many positions in time, so we are illustrating here a pulse train (figure 2a). The pulses are varying in their periods; that

REPETITIVE SAMPLING CONCEPTS

Fig 2a Signal TRAIN

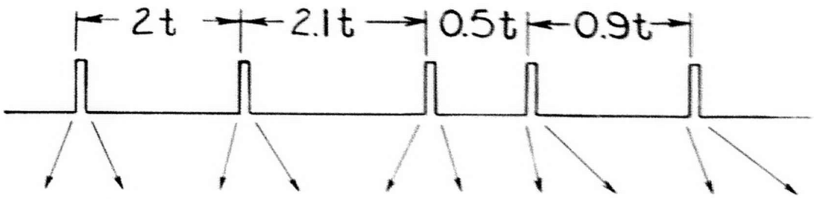

Signal Fig 2b Pulses

Successive Relative Sample Timing

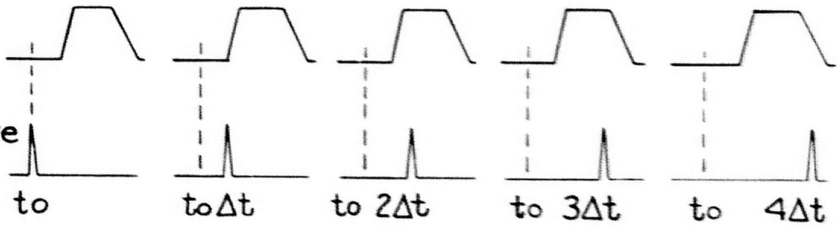

Fig 2c

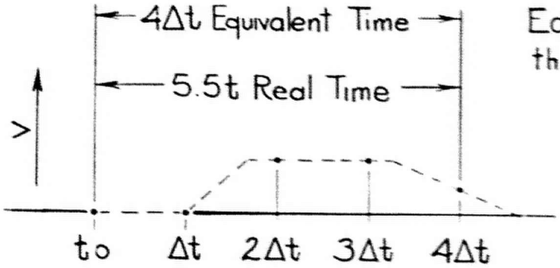

Equivalent time per dot is the slewing interval Δt

Real time per dot is the incoming signals repetition period is, the time separation between the pulses is not uniform, and it is not necessary to be uniform in this system. The reason for this is that every time we have an occurrence of a pulse, we derive a trigger from this particular pulse, and this trigger may be obtained at some time reference t_0. By making use of this reference time t_0 and delay cables, we can then have the pulse, expanded view here in time, appear with relation to this master timing pulse, however we wish, by adding or subtracting delay cables. Now, what we are going to do is have our voltmeter keyed on by this signal (figure 2b)—this timing signal—to take an instantaneous reading, and of course at the first time illustrated, it would read zero, because the value of the signal level at the occurrence of our keying pulse is zero. At the next repetition of the signal—two units of time later—we can then arrange our keying pulse, or voltmeter timing pulse, to occur at a later time with respect to the reference time by an amount delta t. The next time we can have it occur two of these time intervals later, the next time three time intervals delta t later, and so on. Now we can plot voltage as a function of time, based on the time intervals delta t that we moved over our timing pulse. With one position in space and with many repetitions in time, we can take samples of the signal, and reconstruct them to reproduce the pulse. Now as before, there's an equivalent sweep speed involved and equivalent time resolution involved. The time resolution in this case depends on the closeness of the samples in time or the size of delta t. The total equivalent time involved is just a matter of how many repetitions of the signal we sample before we stop increasing the delta t total and start over. Now we have two important factors here, and this is where most of the traps in the sampling system occur (that is in distinguishing between equivalent time and real time). The real time it takes to take all these samples and construct this figure (figure 2c) is the sum of the periods of all the signals that we had in the incoming signal train. The equivalent time (that is the reconstruction of the figure in equivalent time) is the time interval between the master reference pulse and the time we took our sample (that is some multiple of delta t).

Now to illustrate these factors, we have a sampling oscilloscope operating. What we are showing here (figure 3a) is a signal with many dots per display, and the signal recurrence rate is roughly 100 kc. If I turn the number of samples per display down to a smaller number, you can distinguish the individual samples being taken across the screen here (figure 3b)—and the significant difference is this—each one of these dots is about 10 microseconds apart in real time—so if we had 50 dots across the screen, it would take about 500 microseconds to make this display. If I increase the number of dots in the display, it takes longer time to go across the screen, because we have more dots—each one 10 microseconds apart—to fill up the screen. So the real time to get across the screen has changed by a factor of 10 when we go to 500 dots per display compared with 50 dots per display. However, the equivalent time remains the same.

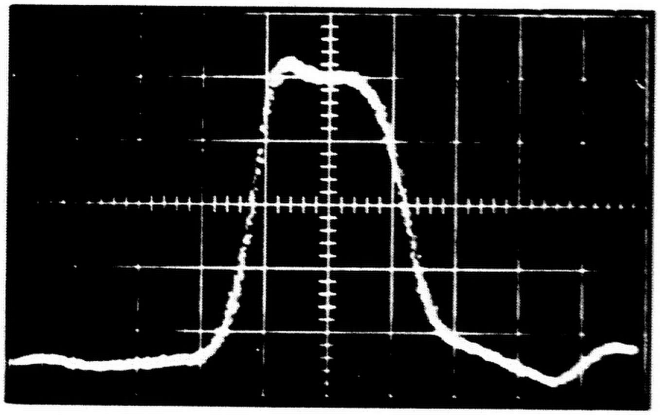

Figure 3a

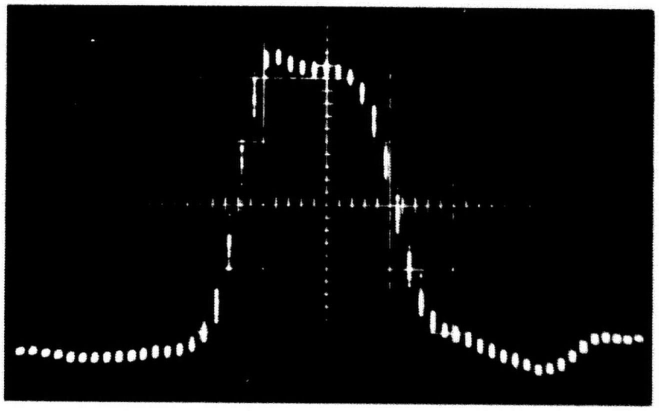

Figure 3b

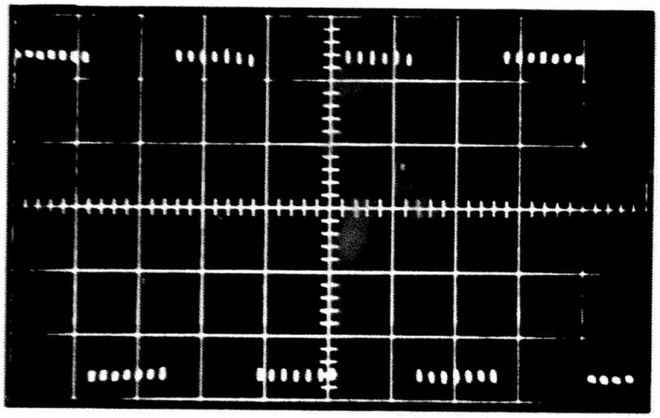

Figure 4a

You will notice that the risetime of the figure and the duration of the pulse stays the same regardless of whether we have a coarse trace or a fine trace—so the equivalent time is remaining the same regardless of the number of samples in a given display, but the real time across the screen changes. Now, to show you how you can get into a trap if you do not distinguish real time from equivalent time we have a calibrating scheme involving the standard calibrator on the oscilloscope. This calibrator is putting out a 1000-cycle per second square wave, and so what you are seeing now is switching between zero volts and the calibrator level (figure 4a). Now if I adjust the repetition rate of the signal generator just right, it looks as though we have extremely fast risetime square wave on here—but this is a false display in real time. If we change the number of samples per display, the display changes (figure 4b) so we know we are being tricked or hoodwinked with a false display. So the ability to distinguish between real time and equivalent time is quite simple—you merely change the number of dots in the display, and if the character of the signal changes, then you are seeing a false phenomenon in real time. But if the signal does not change but remains the same independent of the dots-per-display control, then you are seeing a phenomenon which is occurring in equivalent time and is legitimate.

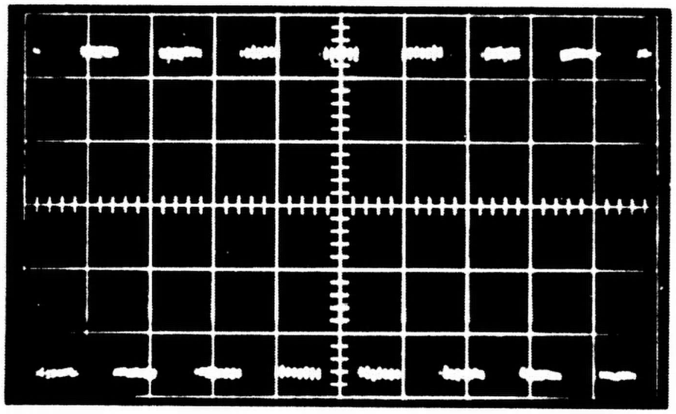

Figure 4b

Well, to hammer the point home a bit, we can reiterate that the signal period can vary in the sampling system provided the time interval that is smallest is in excess of 10 microseconds. This is because the maximum sampling rate is 100 kc—so if the smallest time interval that is involved is 10 microseconds or more, then the pulse spacing need not be uniform. On this basis then, the real time to make a display across the screen is dependent upon the total real time it takes the pulse train to produce enough pulses to make the display—so if you have 50 dots, it takes 50 signal periods to do it; if you have 500 dots, 500 signal periods—but the equivalent time on the screen is quite different. We very carefully arrange the circuits so that the time interval between dots is uniform in equivalent time and is uniform in space across the screen, even though the period is changing. Now, we do this so that the equivalent time between dots is a constant. This is arranged inside the circuitry to be so many picoseconds or so many nanoseconds between each dot, and when the figure is reconstructed then, we can use the dots as a counting method to measure risetime or measure time interval if we wish, because the dots are linearily spaced and have a direct equivalent time. We can detect a fraud in equivalent time by changing the number of dots in the display, because when we change the number of dots in the display we change the real time to make the display, and, therefore, we can detect false display. The other point to make is that you often need sufficient resolution; that is, enough dots to fully portray the figure—for example, (figure 2c), here we have shown some dots along the base line and a dot on top of the pulse, and we carefully show that the rise may occur sometime in between the dots, and we don't know exactly how we got from this point (delta t) to this point (2 delta t). What we need to do is take more dots in order to trail out the particular points. Now we can have the dots much closer than the risetime of the instrument. We can make the dots a few tens of picoseconds apart even though the risetime of the instrument is 0.6 nanoseconds. So we can get plenty of resolution, and this is important, because, for example, you could have a damped sine-wave train and if the dots came along at just the right time you might think you had a 100-megacycle damped train when it was really a 1000-megacycle damped train, because you had dots every 10 cycles. So it's necessary when using the Sampling System to be sure that you're observing an equivalent time phenomenon by changing the number of dots in the display and making sure that the display remains substantially the same.

We can now take a look at the block diagram of the N unit (figure 5), and as you can see, there are a few parts in here. The overall scheme of things involves bringing the signal and bringing a starting trigger in order to get the machine started, and then bringing out certain signals to operate the main oscilloscope. Now, we have arranged things to that the main oscilloscope does not

FIGURE 5

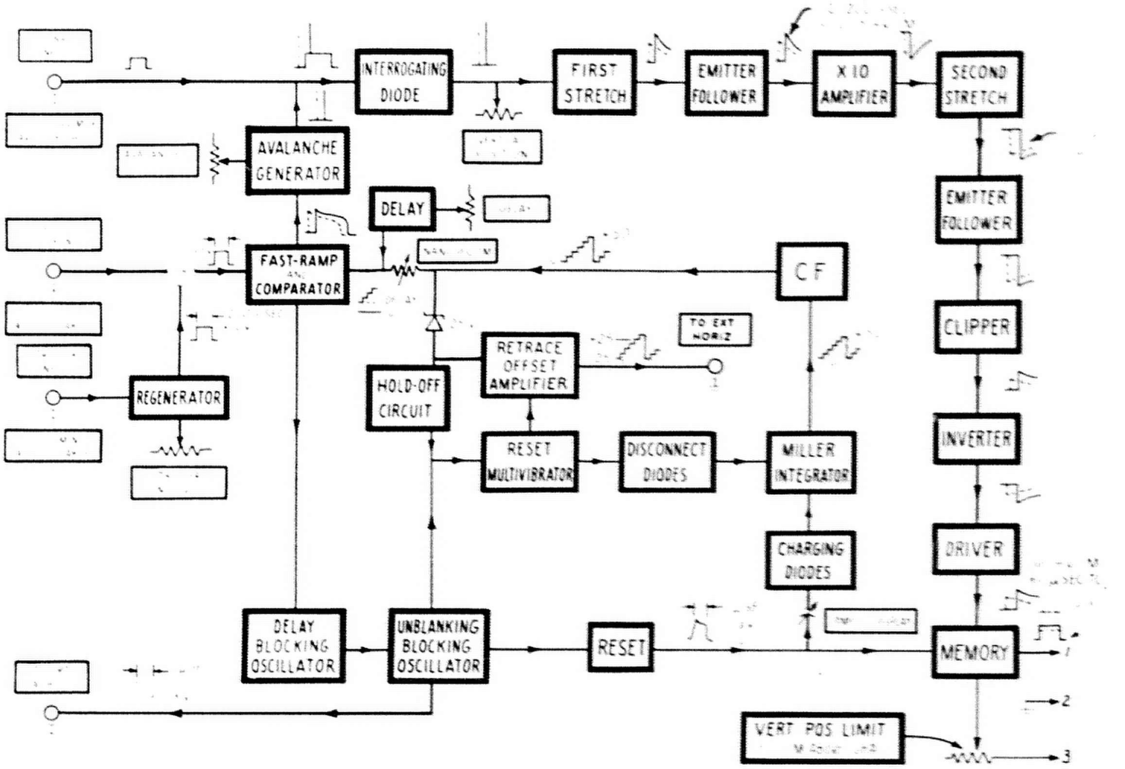

have to be modified. The N Unit will plug into any of our standard plug-in instruments and you will not have to modify the instrument. This means that we will have to externally patch over the horizontal sweep, since it is not a real time sweep anymore (it is an equivalent time sweep). We also must arrange our own unblanking signal to the unblank connections at the rear of the oscilloscope, because we did not wish to require that the instrument be internally modified to get the blanking inside. So we have essentially a triggering input system, a signal input system, and then an external lead to provide the proper sweep, and an external unblanking system, removing the necessity of modifying the oscilloscope.

Pulse Generation Channel

Well, essentially what must happen then in the system is that we must bring in a starting pulse the order of 40 to 45 nanoseconds ahead of the arrival of the signal. We'll discuss in the second half of this session how we arrange the timing of the signal to be in the order of 40 nanoseconds after the arrival of trigger input. But for now we'll just imagine that we have a trigger coming in, and inside the unit we have a trigger regenerator and a sensitivity control so that we can make a standardized signal from the relatively arbitrary trigger input. The standard size signal forms a starting gate on a system which is going to determine the internal timing of the system. This section of the unit (fast ramp and comparator) is going to make the small time intervals, delta t, at which we take our samples. Now, this part of the system operates comparatively simple; there are two signals coming into it of importance—one signal is, of course, the starting gate which permits a fast ramp to start, and the other is a hold-off signal. Now the hold-off signal is composed essentially of two parts—one is a manual delay control which allows you to change the time position of the signal, and the other is the signal which is proportional to the horizontal sweep which causes the spots to move uniformly to the right—each time we take a sample. So the hold-off signal coming in is a composite of the delay signal and a horizontal position signal—his combination keeps the comparator held off until the fast ramp catches up. When the fast ramp catches up, the comparator will then send a signal to start off an avalanche generator. The avalanche generator generates a very fast spike which has an effective half width of the order of a half nanosecond. This signal is mixed with the incoming signal to form a composite like I have shown here (just to the left of the interrogating diode). The time position of the avalanche spike, however, is made to vary successively as the incoming signals arrive.

Vertical Channel

This sum of the avalanche spike in the incoming signal is presented to what we call an interrogating diode. The interrogating diode is back biased by what corresponds to the vertical position control, and there's sufficient back bias so that the interrogating diode cannot conduct on signal amplitude alone. The avalanche spike is much larger than the signal, and is capable of overcoming the back bias on the diode, so essentially we transmit only the spike level above the setting determined by the vertical position control. So the output of the interrogating diode consists of a number of spikes proportional in amplitude to the signal when it was sampled. Now, if we sample ahead of the signal we get a small sample; if we sample when the signal is present, we would get a larger sample—if the signal is positive—and of course, if the signal is negative we would get a smaller sample. Well, these samples are comparatively narrow in time, and the big advantage of the sampling technique lies in the fact that these samples may be time stretched and amplified in narrow band width amplifiers, so the first thing we wish to do is to change the samples from relatively narrow time samples into samples which have a significant time duration. And the first stretch, which consists of a grounded base transistor, accomplishes this in the collector circuit by having a high impedance collector circuit with an emitter follower to remove the signal, and the time constant of the decay here is a few hundred nanoseconds rather than the ½ nanosecond wide samples. So we get an initial time stretch of many hundreds. Following the first time stretch, we amplify the signal from its level of the order of hundreds of millivolts up to the order of volts, so that we can perform another time-stretching operation with diodes, and preserve linearity of the system. By boosting the signal to the volt level we don't have to worry about the 300 millivolts of knee in the diodes. So now we have a signal which is time-stretched out to the order of 1 to 2 microseconds time constant, and the signal level is a few volts at this point. Well, as we stated before, the sampling pulse is considerably larger than the signal is; in fact, you could say that the signal is modulating the amplitude of the sampling pulse. So to have a linear system, we must restrict the percentage of modulation of the sampling pulse to the order of 3 to 10 %—or we won't have a linear system. This means that the spike size of the stretched signal height at this point consists mostly of avalanche pulse and only a little bit of signal. So, we have a clipper stage which is essentially a biased-off transistor, and it will conduct only the signals which are in excess of the cut-off level on the clipper. So now we have a signal which consists of a small pedestal, plus some signal that is useful. We further amplify the signals and put them into the final memory, which has a time constant of about 60 microseconds. Now at this point (terminal 1) we have a sufficiently large signal at the right impedance level to go into the main amplifier of the oscilloscope. At this point the time constant is quite long, and since we wish the system to operate with signals as close together as 10 microseconds, if we had a 60 microsecond time constant, it is obvious it is not going to reset itself in sufficient time to be ready for the next pulse, so we must have a reset signal coming into the memory, so that after about 2 microseconds of time we reset the memory so that it will be ready for the next pulse, which could come as soon as 10 microseconds. The other input to the oscilloscope—the other half of the push-pull input (terminal 3)—goes to a vertical position limit control which essentially adds a dc voltage to the other input of the oscilloscope to balance up the system so that the dc level inside the main oscilloscope will come out properly and we can make full use of the dynamic range of the driver amplifier. We will go through the adjustment of this part of the system in a little bit.

Pulse Generation Channel Continued

Now, the question is then, "How do we get the signal that resets the memory, and how do we get the signal which is going to provide the fast ramp comparator system with its variable hold-off?" We take a signal from the comparator when it goes off, and we send this into a delaying blocking oscillator. This is simply a unit which produces a 1 microsecond wide pulse, and we use the trailing edge of the delaying blocking oscillator to provide a 1 microsecond delay. This 1 microsecond delay serves the purpose of allowing the whole vertical system to get going; to get a signal into the memory, to get the memory capacitor charged up, to send it into the main 'scope, to go through the delay line in the main 'scope, to get the signal on the deflection plates, so that we have plenty of time for all this to be accomplished. So when the trailing edge of this one microsecond wide pulse comes along, we fire off the unblanking blocking oscillator. This also provides a 1 microsecond wide pulse, and this 1 microsecond wide pulse goes to unblank the cathode-ray tube. The signal level at this point is relatively low—it suits transistors—it's about 15 volts high, and there is a transformer on the end of this cable to step up the signal to the order of 45 volts to unblank the cathode-ray tube. In addition to that, the unblanking blocking oscillator sends a signal into the reset stage, and the trailing edge of the 1 microsecond wide unblank pulse sets the reset stage into operation, and the output of the reset stage, of course, resets the memory, after 2 microseconds (the delay associated with the delay blocking oscillator, and the unblanking blocking oscillator). In addition to that, the reset stage provides a signal to the samples per display capacitor. The samples per display capacitor takes the standardized signal coming from the reset stage and puts a definite charge on the capacitor. Of course, the amount of charge is then determined, since this voltage is fixed, by the size of capacitor you choose.

Horizontal Channel

The units of charge are sent into a Miller integrator, so each time we send in a lump of charge to the Miller integrator, the Miller integrator will advance 1 step in its output. So everytime we reset the vertical memory system we automatically deliver a definite unit of charge to the Miller integrator which determines the definite step size. So if we use a large capacitor we will have big steps and we would have a very few number of dots in our display, because we quickly move across the CRT. When the staircase gets up a certain height in volts, the system automatically resets. So if we have large steps we get a few dots in our display, by using small steps we get many dots in our display, to cover the 50-volt interval that the Miller integrator goes through. The cathode follower output of the Miller integrator, is capable of driving the resistors which determines the amount of signal sent to the fast ramp and comparator. By sending a relatively small staircase into the comparator, the fast ramp doesn't have to go very far before it overcomes the hold-off signal represented by the staircase being sent in from the Miller integrator. On the other hand, if we send the large signal down here, it takes a long time for the fast ramp to catch up. The fast ramp goes through a uniform sweep of zero to 200 nanoseconds each time, and the size of the step, then, is going to determine how long it takes, or how much time we'll have between each step to, in effect, determine our equivalent time per dot. The internal staircase goes from zero to 50 volts, and to make it convenient to have a magnifier system in our unit, we have a 25-volt Zener Diode which makes the staircase go from −25 to +25 volts in our horizontal output, thus by changing the gain of the horizontal amplifier by a factor of 10, we can provide a 10-time sweep magnifier, because the size of the steps will appear to be 10 times larger. The retrace offset amplifier has the function of stopping us from seeing the retrace dots while we're resetting the Miller integrator. The Miller integrator capacitor is relatively large, so that we can have stable operation with real time between samples of the order of a 50th of a second. The Miller integrator must have very small drift even though the time—the real time—between dots is as small as a 50th of a second. For that reason it takes quite a bit of time (in the order of 100

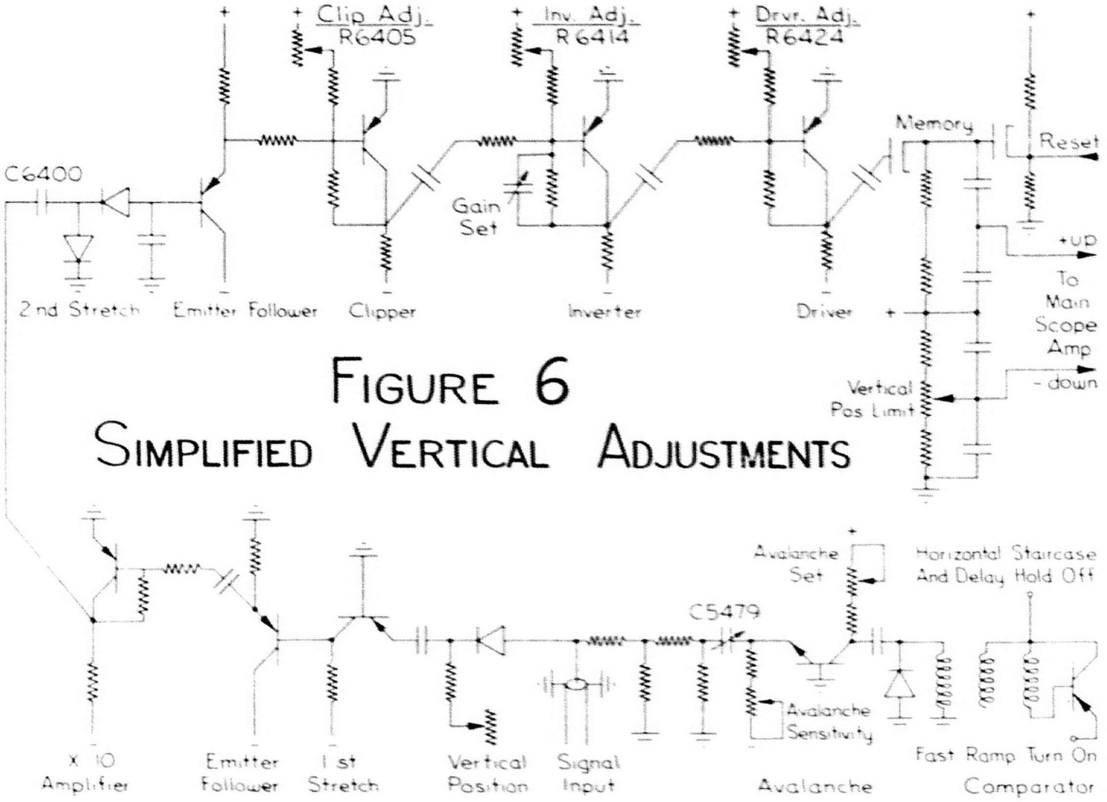

FIGURE 6
SIMPLIFIED VERTICAL ADJUSTMENTS

to 150 microseconds) to reset the Miller capacitor. This means we could have 10 to 15 dots occurring during the retrace time, and to prevent these from appearing, the retrace offset amplifier picks up the trace from 25 volts, moves it to about 30 volts, and holds it there until the retrace has been accomplished in the Miller circuit. This means the trace moves off the screen to the right side and then pops back later to be ready to go. So, to make this operation occur, and to reset the Miller Capacitor, we have the reset multivibrator acting in the sweep circuit. We take a sample of the −25 to +25 staircase, send it into a circuit which operates essentially a Schmitt trigger circuit. When we get to the peak of the Schmitt trigger, the Schmitt goes off, sending a signal to the retrace offset amplifier to pull the sweep over and tells the Miller integrator to reset. In this way we have our internal staircase resetting.

Vertical Adjustments

We are now in a position to go through the adjustments on the N Unit, and we can do this best by starting at the back end of the N Unit Vertical Channel (figure 6), and we will first set up the driver amplifier. The driver amplifier is required to have almost its total swing available to provide the memory with a large enough signal to operate properly. So we will be setting the dc level in the driver stage by biasing it to cut-off, and then moving one volt into conduction. Then when we finish that adjustment we will go to the inverter amplifier, and since it is required to go from a conducting condition to a cut-off condition and we want to make the most dynamic range available, we will adjust the bias on the inverter so that the transistor is de-crushed and then we'll back off a volt or two. And we can judge the quality of the transistor by being sure that the voltage on the collector at that time is not more than 5 volts or so. Then when we go to set up the clipper adjustment; that is when we will get into the problem of setting the clipper correctly, with relations to the main amplifier, so that the dynamic range of the whole system will be fully utilized. Once we have these adjustments, the dc settings in the main system will be completed, and then we can move on to setting up the avalanche circuits and checking that the avalanche spike is the right height to properly drive the dc settings that we have previously arranged. The easiest way to do this is to take the side off the main 'scope. When you do this, the N Unit is exposed, and all the adjustments that you have to make are available from this side. There are no adjustments on the back side or underneath; everything is accessible. We usually use a Type 540 Oscilloscope to make these adjustments.

We can conveniently use external triggering on the test oscilloscope by picking up the signal from the triggering generator. The vertical input can then be arranged to look at the driver output signal. By starting with the trace zeroed at the top of the test screen, and 5 v/cm at the probe tip, we can see then the normal operating level and then the height of the stretched signal going into the driver stage. By taking a screwdriver to the driver adjustment we can move the trace up and down and we can see when the driver stage goes into cut-off, because the signal starts disappearing and the trace stays put. Set the driver one or two volts out of cut-off, then we are sure that the driver stage will have a maximum linear swing. The inverter stage is similarily adjusted, except we can see that the transistor is just short of saturation and it is drawn out of conduction. We can crush the transistor by increasing the base drive (INV adjust); then we can see that we are saturated, and we set the inverter out of saturation by a volt or so. Now when we do this, this will allow the signal on the screen to have a dc level which is independent of repetition rate. The signal should stay substantially at the same level on the screen regardless of rep rate. We are running now at about 1000 cycles; we can go up to 10,000 cycles—and the base line does not shift. The inverter stage adjustment is very important for this point. We can see then that we have accomplished the first two adjustments—the setting of the dc levels on these units.

Now, in order to set the clipper and the vertical position limit we will want to be able to have a trace on the screen, and by hook-or-crook we can get this trace. One convenient way is to use the calibrator, and I happen to have a trace going, so we might as well stay with this—and with the trace on the screen then we will attempt to find out what the dynamic limits of the vertical amplifier system are. We are going to do this by using the clipper control to position the trace upwards—and use the vertical position limit control to keep the trace on screen downwards—what will eventually happen is we will run out of the ability to position upwards with the clipper control, because the memory system will run into its dynamic limit, formed by a reset hold-off diode. When we overcome the back bias on the reset diode the memory system will be shorted out by this diode, and we will not be able to further increase the trace on the screen. At that point we have determined the saturation limit of the vertical amplifier. We can then take the clipper adjustment and back off about 2 centimeters from the saturation point and then arrange the saturation point to be effectively 2 centimeters above the graticule by making the non-saturated level be at the top.

The first thing we notice here (figure 7 a) is that the trace is relatively unsaturated on the screen, and if I take the clipper control and move the trace upwards, and move it down with the vertical position limit we can see that we have reached the saturation condition (figure 7 b). Now, by backing down about 2 centimeters from the saturated condition, we now have a trace level that's well out of saturation (figure 7 c). We arrange that with the vertical position limit to be at the top of the screen. Now we move the trace with the clipper control back down to normal base line (figure 7 d); from now on all the positioning is done with the vertical positioning control.

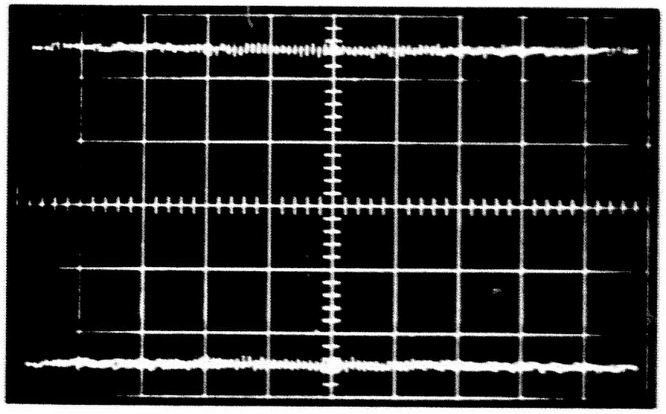

Figure 7a

We have set up the complete unit so that the widest dynamic range that is available is present. When you

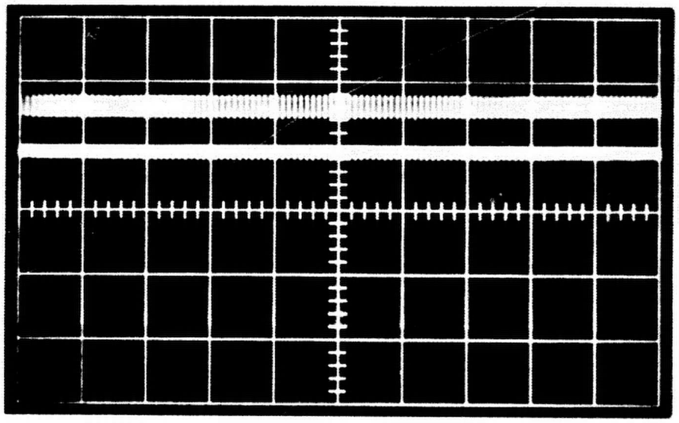

Figure 7b

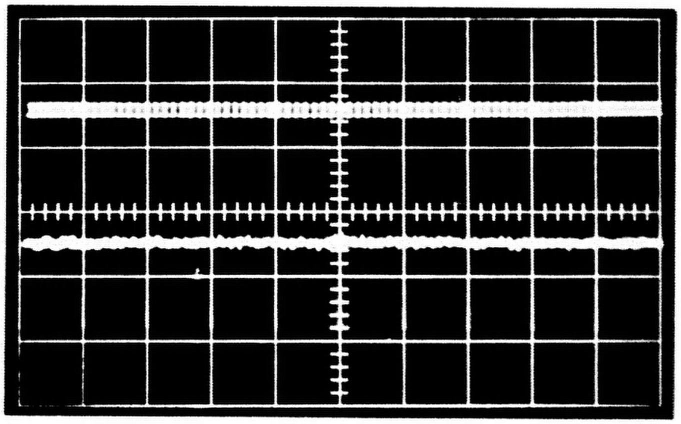

Figure 7c

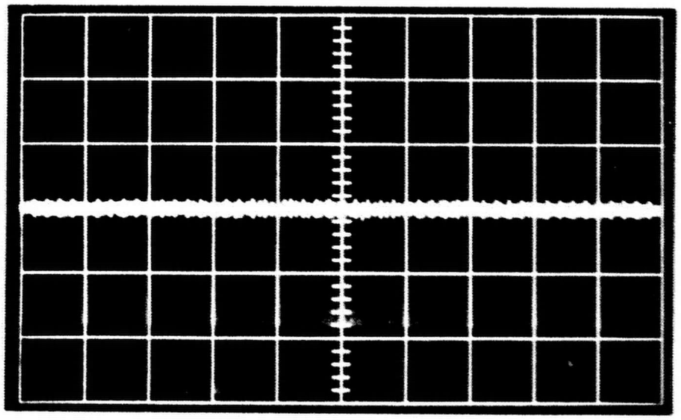

Figure 7d

change the N Unit from one Tektronix oscilloscope to the other, due to the fact that the main amplifier may be slightly differently balanced than the main amplifier that was used when you set the clipper control, it will be necessary to reset the vertical position limit, and this is done with the front-panel vertical position control, when you are simply changing the N Unit from one 'scope to the other. The procedure is the same—we turn the vertical position control to get the picture to go up; we turn the vertical position limit to make the picture come down, until you reach saturation. When you hit saturation, you back off 2 centimeters, put that at the top of the screen with the vertical position limit and from then on the vertical position control acts normally. So this vertical position limit control is extremely important to remember when you're plugging an N Unit into a main oscilloscope for the first time. If you neglect that you'll be very apt to be in trouble with linearity. You will have crushing of the pulse. The Vertical position control should end up between 11 and 1 o'clock; if not, readjustment of the vertical channel and the avalanche control may be necessary.

We have completed the adjustment of the dc levels and we have the widest dynamic range available—and so now we have to adjust the avalanche transistor to get it operating properly. The easiest way to arrange this, as starting point, is to set the trimmer capacitor (C5479) so that the avalanche spike could get through the memory diode with the vertical positioning control centered (12 o'clock). We produce, through the amplifier, a stretched signal at the input to the second stretcher that's about 1.8 volts high.

Avalanche Setting

The trick now is that when you have a pulse displayed on the screen when you go to adjust the avalanche set, is to make the pulse appear as far to the right and up as you can without getting excessive noise. The avalanche set adjustment is available on the outside, and is mainly useful as a temperature compensating adjustment. The avalanche sensitivity, that's internal, is the one which will usually get you into an unstable condition. What you can do is look at the effect on the trace and see when minimum noise occurs and when you are safest away from a free-running condition. This will vary a little bit from one avalanche transistor to another, as to which control is most sensitive—but by turning the two you can quickly observe when the signal becomes unstable—when the noise gets out of hand. We'll sometimes find that some avalanche transistors have two stable modes—one giving a very slow risetime and a very large sensitive operation; the other giving a fast risetime and a less sensitive operation. Of course, it's the fastest risetime mode that we want. Not all transistors exhibit this, however. Once we have done this step of checking that the avalanche transistor is operating reasonably properly we can then take a look to see that we have the standard size signal of 1.8 volts at the input to the second stretch. The test oscilloscope should show this. The trimmer capacitor will vary the height of it very nicely. When that adjustment has been made, the next step is to check the sweep timing, and the sweep timing can be done with our timing standard arrangement.

Sweep Timing

We have a damped train available for each sweep speed, and we can check easily since we have one cycle in the damped train per centimeter for each sweep speed. By checking that the sweep speed is right on, or how close each sweep speed is, we can then tell how to set the sweep speed adjustment control. The sweep speed adjustment is a single one for all sweep ranges, so that when we change from one main unit to another we can very quickly see that the sweep speed remains proper. So, if we find that some sweep speeds are 2% high and other sweep speeds are 1% high, we can trim them up to split the difference. In this way we can check if our timing is right, and we can then check the risetime and linearity of the N Unit. The setting of the avalanche trimmer capacitor is going to be a compromise between risetime and linearity.

As you set a smaller value of capacity in the trimmer, you of course get a narrower spike fed to the interrogate diodes, and, therefore, our risetime improves. But, as we do this the size of the spike sent through is smaller, and, therefore, we prejudice the linearity of the instrument, due to excessive modulation of the interrogation spike. So we must make sure that the adjustment of the system is proper by checking that first of all we have the right gain in the system, and in checking the right gain in the system we check that the sampling efficiency is correct and that the sampling diode in the following amplifiers are working properly. Generally, too-wide an avalanche spike will result in too slow a risetime and too much gain if everything else is normal. So the way we check this is fairly simple. We can check the gain by utilizing our calibrator adapter and the standard calibrator on the Tektronix plug-in type 'scope, making sure that we have the right number of centimeters deflection for the setting of the calibrator. Typically, for example, we will adjust this for 5 centimeters on a 50-millivolt signal, establishing that we have 10 millivolts per centimeter (figure 7 a). We can then cut down on the size of the signal so that we have a 1 centimeter size signal and position it up and down the screen to make sure that the 1 centimeter of deflection stays 1 centimeter within a millimeter as we go across the screen. The next step, since we have timed out the sweep, is to measure the risetime of the instrument with the Type 110 pulser, which has a sufficiently fast output pulse to measure the risetime. This procedure is a little bit lengthy, and I think it is best done individually by a Tektronix demonstrator rather than by myself here, because it would be quite hard to see the dots on the screen to establish the risetime (figure 8). So rather than go through that we'll spend the time on the application section that follows. Knowing the risetime and the gain, one can then decide on what to do with the avalanche capacitor. If in order to make risetime, you need less than 1.6 v at the input to the second stretcher, linearity will probably suffer. This indicates insufficient avalanche pulse amplitude.

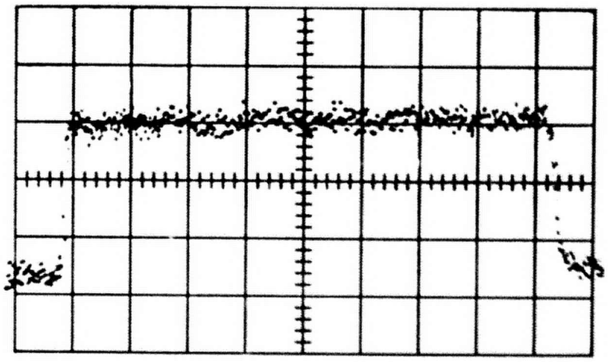

500 SAMPLES PER DISPLAY
Figure 8

TEKTRONIX SCENE II

The second half of this talk on the Tektronix Sampling Sytem will include the general arrangements to obtain the proper time relation between the incoming signal and the triggering signals, and then we'll talk a bit about how to get signals from our signal generators into circuits and how to take the signals out of the circuits, and present them to the N Unit. In order to appreciate what we are accomplishing here, we can take a look briefly at how the N Unit obtains its variable time.

Time Slewing

A regenerated trigger consists essentially of a pulse somewhat in excess of 200 nanoseconds long and about 10 volts high which is used to start the fast ramp circuit and comparator circuit in the unit. Now the time relations are approximately as shown here (figure 9). After the regenerated trigger comes along, there's a time delay before the fast ramp starts—and it's a little non-linear at first—so we

FIGURE 9 SIMPLIFIED SLEWING

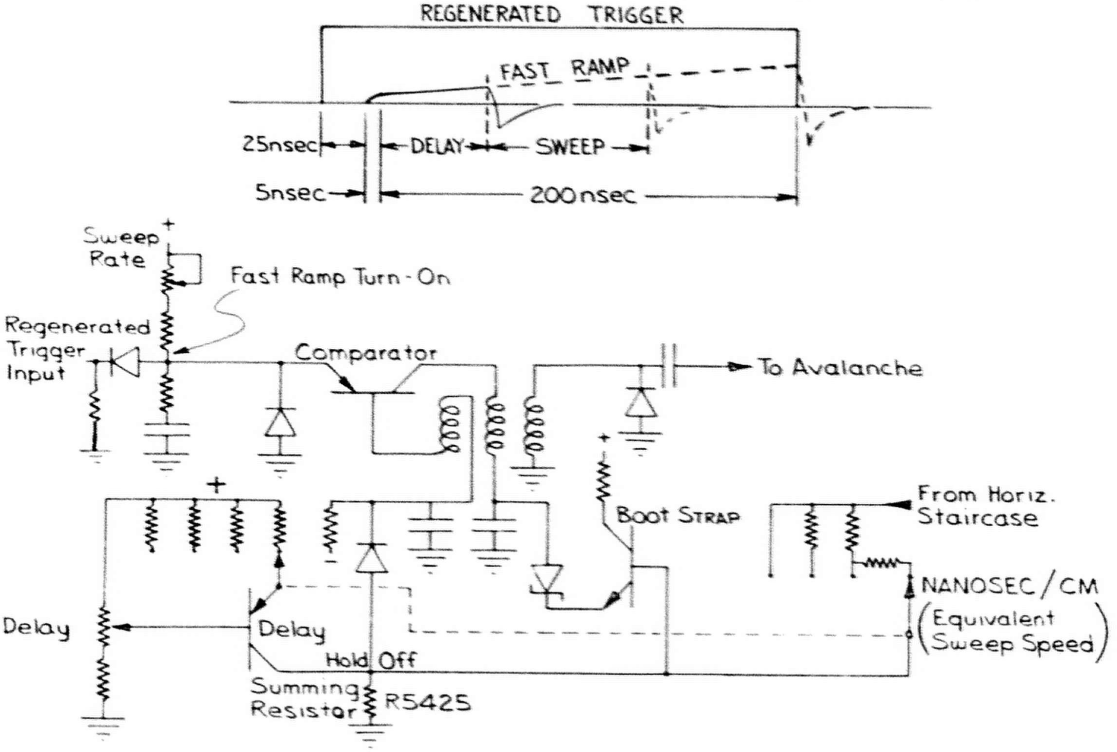

don't make use of this beginning time. Then, depending upon how much delay someone wishes to dial into the system manually, the sweep will start at some point on this fast ramp. We have a hold-off signal coming into a comparator, and a fast ramp coming into the comparator, and when we have coincidence (the fast ramp catches up with the hold-off signal), the comparator will fire off and we would see on an oscilloscope the time of firing on the fast ramp move progressively to the limit of the sweep and then start at the beginning again and move back each time we make a complete display across the screen. We can vary the delay time so that the position of the sweeping interval moves throughout this 200 nanosecond time slot. So if we use a very fast sweep—for example, 1 nanosecond per centimeter—the total sweep length would only be 10 nanoseconds, and this 10 nanosecond time interval could be moved anywhere within the 200 nanosecond time interval. If we use the smaller sweep—10 nanoseconds per centimeter—we have 100 nanoseconds of time which we can move back and forth in the 200 nanosecond period. So, with this system—regardless of what sweep speed range you're on—you always have a 200-nanosecond "Time Window" in which you can observe things. So you can set the trailing edge of a pulse a little under 200 nanoseconds long on the fastest sweep speed, and you can measure fall time with the same resolution that you can measure risetime. The way this is accomplished in a little detail—we have a summing resistor down here (R 5425), which is going to take a current from the manual relay system and add to it a current from the horizontal staircase that we were talking about, to make the total hold-off signal. Now, the delay signal for a given sweep across the screen is fixed, but of course, the staircase is what makes the sweep go across the screen, and therefore, we make the hold-off signal proportional to the horizontal determining voltage. As we choose resistors in the horizontal staircase input (nanosec/cm switch), we can force the system to advance at different rates; that is, the hold-off voltage will build up from the value determined by the delay setting to some other value, in steps, according to the resistors selected in the horizontal staircase. This hold-off signal is introduced into the base circuit of the transistor, and the emitter-circuit has the turn-on signal.

The fast ramp circuit essentially has 10 mils coming up through a resistor and a diode towards a 225-volt power supply. The 10 mils coming up through the diode will determine a certain voltage on the emitter which is relatively small. When the regenerated trigger comes along, it is a plus signal, and this turns off a diode. When the diode turns off, the ten mils can no longer come through the diode, so it must come out of the fast ramp capacitor. The resistor in series with the capacitor makes a small step occur at the beginning to turn on the comparator abruptly and reduce the minimum delay needed in the system. After the small step occurs, then the capacitor starts charging up, and it only charges 2 volts out of 200, so we have a quite linear fast ramp occurring at the comparator. So as the comparator emitter heads up towards the base hold-off voltage, eventually there will be a coincidence and the comparator will fire. To guarantee that the comparator fires in the same way every time, we make the collector potential be boot strapped, according to the base-return voltage on the hold-off bus, so as the hold-off voltage goes from the order of zero towards 2 volts, the collector supply will follow the base advance—so when the comparator goes off the collector has the same potential with respect to the base regardless of where its firing on the fast ramp. This assures us a uniform driving trigger to the avalanche circuit, and hence, a uniform timing in that circuit. This is the reason that we need the order of a 40 to 45 nanosecond time delay between the arrival of the signal and the arrival of the regenerated trigger. We must wait internally for something like 30 to 35 nanoseconds for internal delay times to get to a linear part of our fast ramp. Then we

FIGURE 10 GENERAL SCHEMES

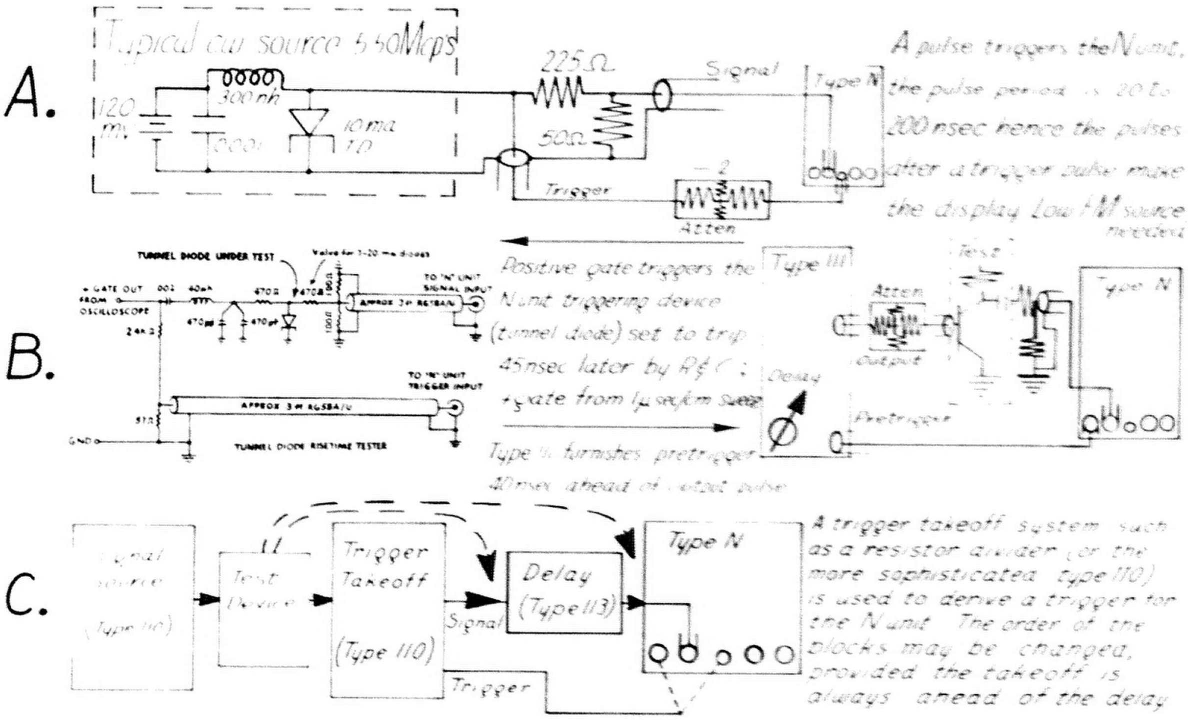

A pulse triggers the N unit, the pulse period is 20 to 200 nsec hence the pulses after a trigger pulse make the display. Low FM source needed

Positive gate triggers the N unit triggering device (tunnel diode) set to trip 45nsec later by R & C; + gate from 1μsec/cm sweep

Type 111 furnishes pretrigger 40nsec ahead of output pulse

A trigger takeoff system such as a resistor divider (or the more sophisticated type 110) is used to derive a trigger for the N unit. The order of the blocks may be changed, provided the takeoff is always ahead of the delay like to have when we're triggering on a fast pulse—a field of view ahead of the pulse in the order of 5 to 10 nanoseconds, and so, in addition to the built-in delay, we recommend an external delay that's the order of 5 to 10 nanoseconds more, so that you can see a bit ahead of the pulse before you apply it.

Time Delay

Now the general way that you go about providing the time delay for the N Unit can be expressed something like this—there are 3 fundamental ways to do it—one way is to use a signal that has a high repetition rate and a uniform repetition rate, so that you can trigger on one pulse and observe the following pulse in the display. For example, if we had signals with repetition rates which are between 5 and 50 megacycles per second, the time interval between the pulse would be between 20 nanoseconds to 200 nanoseconds in the pulse train. Well, we can trigger the N unit and count-down, so it turns out, since we have a time field of 200 nanoseconds, we can observe the pulses. If we trigger on the first pulse and look at a later one in the train, we can then provide the time delay by the repetition period between pulses. We have shown this figure 10, example "A", a free running tunnel diode relaxation oscillator.

We have about a ten-mil tunnel diode setting here, and it's operating as a relaxation oscillator at about 20 or 30 megacycles. This means we will have a period in the order of 30 or 40 nanoseconds between pulses, and if we take the output of the tunnel diode and send it into the trigger system of the N Unit, we can count down from the 30 or 40 megacycles to the 100 kilocycle sampling rate that the N Unit operates on. The signal is also divided down in the form of a resistor probe to reduce the signal level down to the order of millivolts for the use of the N Unit. Now this particular divider forms a 10-1 divider so that the presentation on the screen will be about a 100 millivolt per centimeter. By doing this, then, we should be able to observe directly on the N Unit without the need of an external delay line, without the need of external pulse generators, the risetime of the tunnel diode when it is operating directly into the N Unit. We have on the screen now the tunnel diode (figure 11a). The signal being divided down and sent into a vertical input to provide a display of a 100 millivolts per centimeter. The triggering signal is sent around into the N Unit trigger, and by triggering on the signal, we can show it's triggered, because if we turn off the supply the trace disappears. We count down from the nominal 40 megacycles that the unit is running at and display the picture. The sweep speed is about 10 nanoseconds per centimeter, so we can see that the pulse interval is 20 nanoseconds—so the repetition rate is about 50 megacycles. If we wish to observe the risetime in detail we can open up the sweep speed and take a look at it (figure 11 b) or if we wish we can make use of the magnifier and see the sweep speed on that basis (figure 11 c). So either by magnifying or increasing the sweep rate we can observe the switching time in detail of the tunnel diode just with the N Unit and no delay line or other device. We are in effect obtaining our necessary delay by triggering on a pulse which is off the screen to the left and observing those which follow many nanoseconds later. And by making use of our delay knobs, we can observe over a relatively wide 200 nanoseconds time range, so that even if the pulse repetition rate was as low as almost 5 megacycles, we could still observe the leading edge of the next pulse. We can do a similar trick by making use of delaying the signal coming to the tunnel diode. We have another tunnel diode scheme, which we'll show you schematically in a moment (figure 10 b), and we are going to make use of the positive gate on the oscilloscope to provide the power for running a tunnel diode, and in this case we are going to integrate the plus gate so that it arrives late at the tunnel diode. This scheme then allows us to observe the switching time of the tunnel diode and I'll show you circuit-wise in a moment what it's like, and we can observe the rise (figure 12) by going to the fast sweep speed, or we can observe more of the pulse on the slower sweep speed—and again we don't make use of a delay line or other device—it is the internal power supplied from the plus gate of the main oscilloscope which provides a trigger for the N Unit. Then

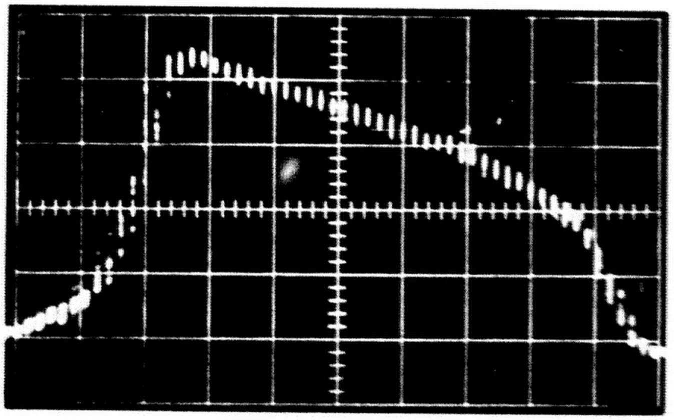

Figure 11a

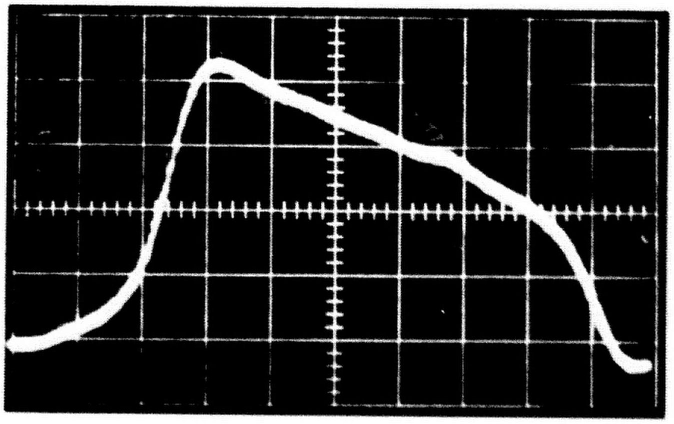

Figure 11b

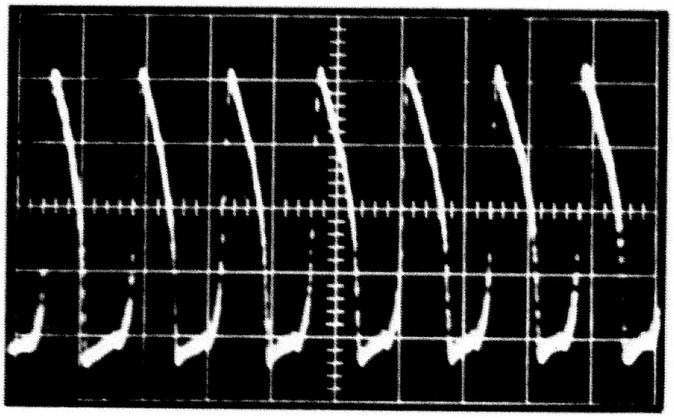

Figure 11c

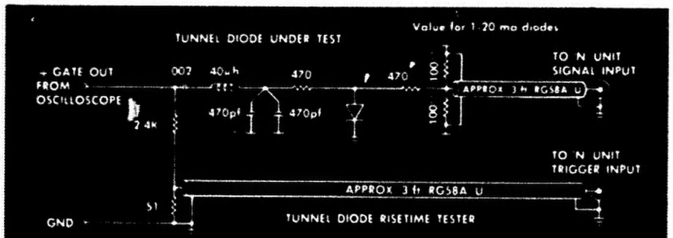

A convenient low-cost method of testing tunnel (Esaki) diodes with nanosecond switching speeds is shown above. A Tektronix Plug-In Oscilloscope provides both the current ramp source for the tunnel diode and the pretrigger for the Type N Unit. The N Unit is set up in the usual way — however, the oscilloscope main sweep generator is allowed to free run at 1 μsec cm. The + GATE OUT not only triggers the N Unit but also provides a delayed current ramp with a low rate of change—which allows the tunnel diode to switch at essentially its own rate.

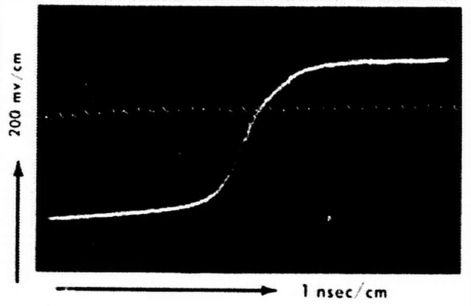

Typical waveform of gallium arsenide tunnel diode in Tektronix Tunnel Diode Risetime Tester (part number 013-029). Price $50.

Figure 12

by taking the plus gate and integrating it we provide a time delay to cause the tunnel diode to trigger about 40 nanoseconds later, so that we can properly display in time relaxation the fast rise of the tunnel diode. This lets us very quickly and very simply observe the characteristics of tunnel diodes.

The scheme which we just saw on the face of the oscilloscope made use of the integrating action of an RC combination to delay the signal coming into the tunnel diode so that it would trigger on a slow current ramp and permit its triggering point to occur roughly 40 nanoseconds after the N Unit was triggered. This scheme of taking the trigger directly into the N Unit and then providing a time delay in the signal going to the device to be tested is a very powerful one because it requires so little equipment. The size of the resistor can be easily selected to make the current sent to the tunnel diode reach its peak at the proper time interval (the value is not critical, since the N Unit has a 200 nanosecond time "Window"). By changing this resistor you can accommodate tunnel diodes of different current ranges (there is an improved version now available Tek No. 013-029). The divider shown going to the N Unit allows a relatively high impedance point at the tunnel diode, and so the slow current ramp permits observing the tunnel diode's true switching speed—it's not speeded up significantly by sweeping current that comes into firing. This allows us to see the true speed. Now a very similar approach is used in the Type 111, where we internally delay the output signal (Right side, figure 10 b). The Type 111 is a device which utilizes the principles we have spoken of above, in that the Type 111 generates a triggering signal for the N Unit and then internally delays the generation of an output signal. The output signal can be delayed with respect to the pretrigger signal by the order of 30 to 250 nanoseconds. This allows us then to choose the time at which the output step signal arrives at a device under test, and hence, the time that the signal from the device under test arrives at the N Unit. This method of a pretrigger pulse generator is a very powerful one in that again we do not need a delay line; we have only a relatively inexpensive pulse generator and the N Unit.

The factors involved here are that the Type 111 must have a very small time jitter between the production of the output pulse and the pretrigger, just as we hope the tunnel diode triggers very uniformly with respect to the current going into it, so the display will have a small amount of time jitter. The Type 111 uses essentially the same delay circuit that the N Unit does to produce its slewing sampling pulse—and so the time jitter approximately doubles when you use the Type 111 over the Type N. The total time jitter will be not more than about 100 picoseconds (10^{-10} seconds) in time jitter when using both of these units in cascade.

The Type 111 was being used to generate the pulse that we saw earlier (figure 3 a). Right now the Type 111 is being used in making this pulse (figure 3 a), and we can turn the delay knob on the Type 111 or on the N Unit to put the pulse wherever we want it on the screen. The normal procedure would be to set the time delay on the Type N to minimum and then turn the Type 111's time delay knob to present the pulse at the extreme right on a 1μsec/cm sweep (from then on, the delay on the Type N is used). The pulse length can be varied and the Type 111 has about ½ nanosecond risetime. The advantage of the Type 111 being we don't need the delay cable—we can present a pulse of high or low repetition rate as we wish by making use of the variable rep rate feature of the Type 111. The main limitation on the Type 111 is that it does not give as clean a pulse, nor as fast a pulse, as from the Type 110.

If the Type 111 System is not practical, because you need a very clean pulse, or if the signal source has an extremely high repetition rate over the triggering capability of the N Unit—which is 50 megacycles—you may wish to use still another system; this third system would be what you would use if the device under test wasn't able to furnish a trigger of its own, or if you could not arrange to get a sample of the signal coming into the device under test. For some reason or another we'll assume that we can only operate on the signal coming out of the device under test. You could then make a resistance divider which would send some of the signal energy to the N Unit input and some of the signal energy over to trigger the N Unit. This would be perfectly satisfactory if the amount of energy which you can divide off for triggering purposes is large enough to trigger the N Unit. In order to do this the N Unit requires that its trigger input be approximately + ½ to 2 volts. If the signal that is coming from the device under test is too small to furnish that size of trigger, or if the repetition rate is above 50 megacycles, the Type 110 will need to be used (figure 10 c). In any case, whether you use a resistive energy divider or the Type 110, you will wish to produce a time delay between the signal and the trigger as they arrive at the N Unit, so a delay line such as the Type 113 is very useful here. With an ordinary simple resistive-type divider the delay line should be at least 45 nanoseconds long. When used with the Type 110, due to the internal delay of the Type 110 producing a regenerated trigger, we need about a 60-nanosecond delay line, and of course, the Type 113 has 60-nanoseconds of delay, so this is adequate for both cases. The trigger take-off system attenuates the signal voltage only 2% between input and output, and from this triggers a relatively large signal that we can use directly in the regenerated trigger input. The relative order of the blocks in the diagram (figure 10 c) can be changed, just so you have the trigger take-off ahead of the delay; for example, we could in principle put the trigger take-off between the signal source and the device under test, if that was reasonable. Or, we could put the device under test after the delay line if we wanted to. The order is not particularly important just so that the trigger take-off is ahead of the delay, and in turn the delay is between the trigger take-off and the signal path coming into the N Unit. If we do this then, we can take advantage of the very fast-rise pulse generator in the Type 110, if we wish, to drive the device under test, or we can instead make use of the high-frequency count-down ability of the 110. What we will show on the screen next is the pulse from the Type 110 signal generator as displayed through the delay cable on the N Unit, and then after that we'll show a 100-megacycle count-down signal.

The signal that's now on the screen (figure 8) shows the fast-rise mercury pulse. The mercury relay's output is at approximately 720 cycles per second. This is fast enough that we can obtain a picture on the screen, we can position it with the delay control, and as fast as we can reasonably turn the delay control, the picture follows. This is the advantage of having a high repetition rate mercury relay system. The display rate, of course, decreases as we increase the number of dots in the display, as the real time to make a display gets longer with a larger number of dots. This is the basic disadvantage of the very fast mercury relay is its relatively limited repetition rate compared to the Type 111. But with this system we can trigger on signals as small as—well, we are seeing a 20-millivolt signal on the screen now (figure 13) accomplishing the triggering—so with a 20-millivolt signal we can very nicely trigger, or of course, on larger signals, we can also trigger. In this case we very definitely need the Type 113 delay cable to provide the proper timing between the trigger coming into the N Unit and the signal coming from the trigger take-off. The other advantage of the Type 110 is its ability to count down from a high-frequency.

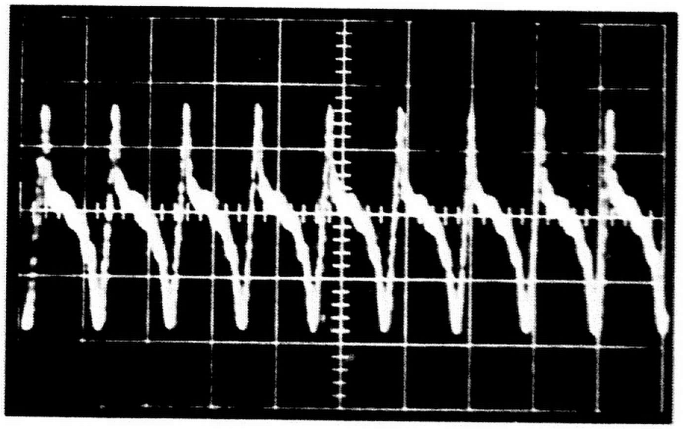

Figure 13

We are at 10 nanoseconds per centimeter, and we're locked on now to a 50-megacycle signal. This is a genuine triggering operation, because as we remove the signal, we should also be able to make the picture remove. We'll run the frequency up to 100 megacycles now, and lock in on 100 megacycles.

Locked in on 100 megacycles with 10 nanoseconds per centimeter (figure 15), of course, we have one cycle per centimeter and we can open up the sweep rate, and look in at another sweep speed, if we wish—so we can examine the pulse train in detail. Well, what we have done now, is shown a few of the general schemes for firing off the N Unit. We have the very simple case of not requiring any delay cables or any trigger take-offs or anything—it's all built in as part of the system (figure 10 a) and we have the relatively complicated case where the signal cannot furnish a pretrigger—we must derive our own trigger—and synchronize on this (figure 10 c). We have shown that we can do this with signals the order of 20 millivolts (figure 14) or as high as 100 megacycles and count down on them (figure 13). The next trick then is to find out how we introduce signals into a device under test and how we get signals out of a device under test.

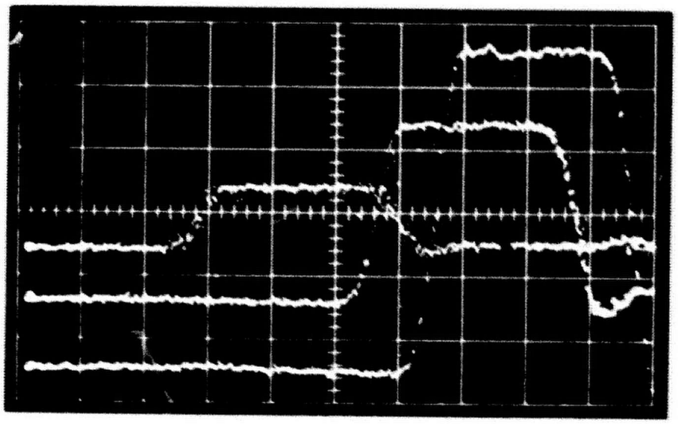

Figure 14

Introducing Signals into Circuits

This diagram will show some of the points here (figure 15 a). First of all, the basis of the mercury relay type generator operation involves the mercury switch and a transmission line of 50 ohms impedance being charged up to some dc voltage like "V". When the mercury switch closes it very quickly dumps the charged transmission line into a load (we'll assume a matched load). Well, the output voltage which is obtained is ½ of the voltage to which the transmission line was charged, so if we charge up the transmission line to 100 volts and the switching closes, a 50-volt pulse goes towards the load, a 50-volt pulse goes in the reverse direction. When the leading pulse arrives at the load, the load voltage remains constant until the reverse direction pulse (which is opposite polarity), terminates the pulse. The pulse length has a duration which is twice the length of the charge line, since the trailing edge travels down the line, reflects, and comes back out a double transit time late. The pulse amplitude is ½ of the voltage to which we charged the charge line. Now the trick is this—it's voltage across the switch that counts, so if the load had a dc voltage associated with it, even though you charged up the charge line to 100 volts you wouldn't get a 50-volt pulse out unless the voltage on the load was zero. If the load had a 10-volt level on it

FIGURE 15

INTRODUCING SIGNALS INTO CIRCUITS

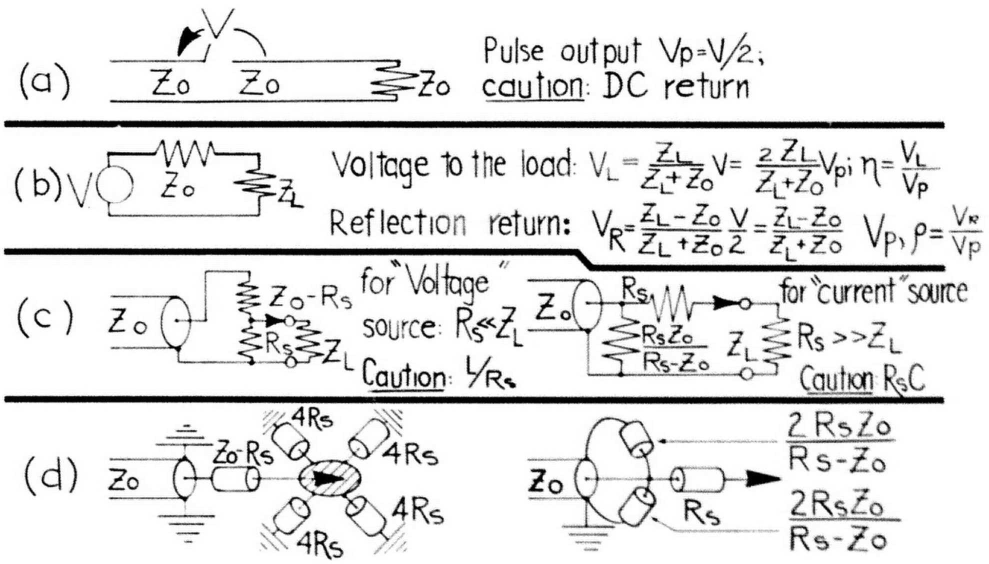

that would either add or subtract— depending upon the polarity—to the charging supply and alter the voltage across the switch, and hence, alter the size of the pulse coming out. So, one caution to remember when using the mercury relay type pulse generator is that you must have a dc return in the output circuit.

If there was a capacitor in series with the load, removing the dc return through the load, the capacitor would charge up to the supply voltage, and so your pulse train would decrease in amplitude and finally disappear. When using this type of a generator you must be certain to have some kind of a dc return, which will allow the charge on the output to drop to nearly zero within the repetition period of the relay. So this means if the load has a dc return time constant of the order of a couple of hundred microseconds, it is sufficiently fast. Now, the next question is, "what happens if you send the pulse down the cable and the load is not equal to the characteristic impedance of the cable?" Well, in that case you'll get a reflection at the load, and here we are going to summarize what happens (figure 15 b).

When you propagate a pulse down the cable, if you left the cable open-circuited at the end, you would get twice the amount of voltage you expected, and of course, if short-circuited, you would get zero volts. Well, this gives an equivalent circuit, which is a generator of voltage "V", which is the same as the voltage across the switch in the original case, having an internal impedance Z_o and a load impedance of your choosing Z_L. So the voltage will divide up across the load exactly according to the voltage divider that you would expect. So the voltage to the load, V_L, is equal to the ratio formed by the voltage division here, times the source voltage, and this is the same thing as saying it is two times this fraction, times the pulse voltage that you are transmitting down the transmission line originally. So, for example, if we had a 100-volt power supply working into a couple of 50-ohm cables, the pulse coming down the cable would be 50 volts. If we left Z_L as being open or infinite, then you would have two times the pulse voltage because the positive 100% reflection adds to the 100% incidental signal; or in other words, the voltage would double on the end of the line. If the load impedance was zero then of course, we would have zero volts output. The transmission coefficient is the ratio of the voltage to the load, compared to pulse voltage sent down the cable. It would be equal to two times the divider ratio that forms this network (figure 15b). Now, there will also be a reflection return from the load, and the reflection that's returned is equal to the difference between the load impedance and the transmission line impedance, divided by the sum of the impedances times the voltage propagated down the cable. The reflection coefficient is easily defined as the ratio of the voltage reflected to the voltage sent down the cable in the form of the pulse originally. So, if we had an infinite impedance for the load, or in other words, an open circuit, then this ratio would cancel out to unity (figure 15b), and so we would have a full pulse voltage reflected back down the cable. On the other hand, if the load was zero ohms, then we would have negative reflection coefficient, and we would have a complete reflection of the pulse back down the cable. So, depending upon whether we want it or not, we can get a reflection coefficient from −1 to +1, and we can get a transmission from 2 to 0 depending upon the relation of the load to the characteristic impedance on the line. Now, usually a reflection from the load is permissible, provided the length of the transmission line between the load and the open end of the charge line is long enough to prevent the reflection return from coming back at a time that would be troublesome in your display. Another way, of course, is to arrange things so that when you introduce the signal to the load, you compensate for the fact that the load is not correct, and in that way prevent a reflection. We have shown this (figure 14c) by showing a case where perhaps you want a voltage source to drive your load rather than a match source. What we have done essentially is to say that if the voltage source is supposed to be small compared to the load, we put a resistor in series in such a way that the generator cable thinks it is terminated in 50 ohms, but the load sees a very low impedance source, and so we have the complete case accounted for here we have prevented the reflection from going back to the generator, and at the same time, we have presented the load with a low impedance source. Now, the caution here is that if this source

FIGURE 16 OBTAINING SIGNALS FROM CIRCUITS

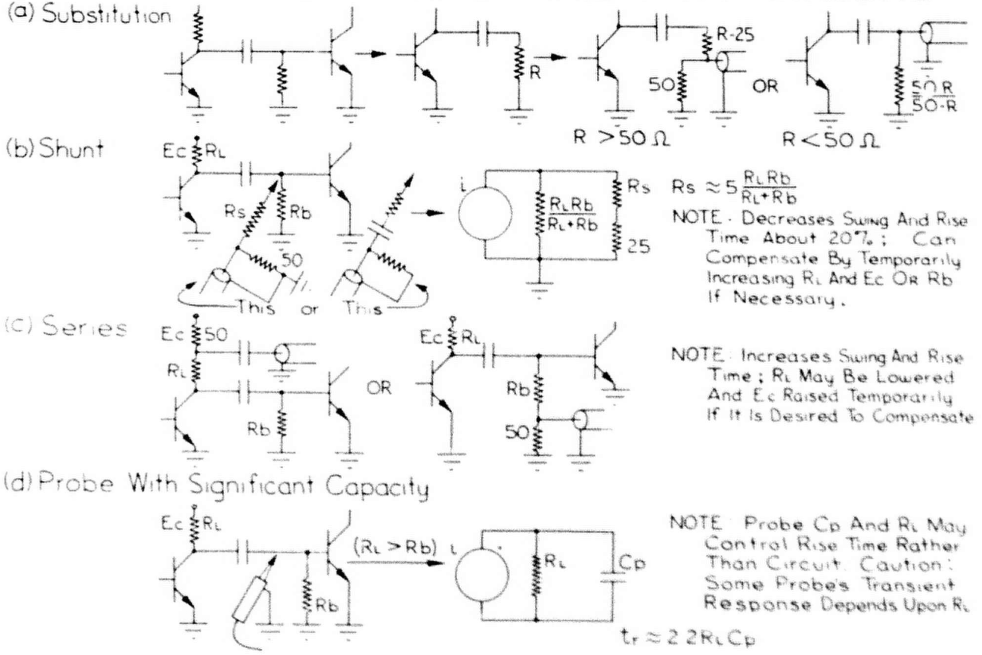

$R_s \approx 5\frac{R_L R_b}{R_L + R_b}$

NOTE - Decreases Swing And Rise Time About 20%; Can Compensate By Temporarily Increasing R_L And Ec Or Rb If Necessary.

NOTE - Increases Swing And Rise Time; R_L May Be Lowered And Ec Raised Temporarily If It Is Desired To Compensate

NOTE - Probe Cp And R_L May Control Rise Time Rather Than Circuit. Caution: Some Probes Transient Response Depends Upon R_L

$t_r \approx 2.2 R_L C_p$

impedance is going to be low, the inductance associated with that resistance becomes a predominating factor, and the L over R time constant may prevent you from building up the current through the small resistor fast enough, resulting in a spiked-up leading edge. So, we have shown in figure 15d a physical configuration which allows you to obtain a low L over R ratio by paralleling up, for example, 4 resistors to form the load resistance we want. Thus one can get a fairly low-inductance broad-area place to connect on your load. A similar case is shown in Figure 15c to make a "current" source, where we are using a large resistor to drive the load. We put in a parallel resistor to make the cable think it is terminated in 50 ohms, but still provide the load with a "current" source. The caution here is that the distributed capacity along this resistor must be taken into consideration, and values much beyond a 1000 ohms become quite difficult if you wish to preserve less than a nanosecond risetime.

Obtaining Signals from Circuits

To get signals out of the circuit, we have 4 main ways—first by substituting in place of the existing load in the circuit under test (for which we can draw an equivalent resistance as shown in figure 16a). We can then substitute by disconnecting the load part of the circuit, and substitute our cable with appropriate resistors. Thus we can get a signal into the N Unit; again we try to match things up, so that the circuit thinks it still sees its normal load resistance, "R", but the N Unit is properly matched into its cable. A second method is to use a resistive probe to go in shunt with the existing circuit. You may or may not require a coupling capacitor to prevent dc trouble as shown in figure 16b. Now, when you do this, it's convenient to keep the equivalent circuit in mind—the collector load resistor and the base-leak resistor form a parallel combination which is loaded by our probes. We would like the probe impedance to be roughly at least 5 times this parallel combination, so that the loading would be 20% or less. Now, the nice part of this is that if the loading is 20% from dc to very high frequencies, and, therefore, all that happens is the signal swing's capability is slightly altered, and the risetime is also slightly altered by the lowering of the impedance level of the system (you can compensate for this temporarily, if you want, by increasing the load resistance and the supply voltage). But the point is, you nominally load the circuit uniformly at all frequencies with this probe, and you can get your results. You can also put the N Unit in series with the load resistors that exist already in the unit. The convenient part of this system is that you can then provide a permanent fixture terminated in a 50-ohm resistor so that the circuit operates normally. Then when you want to see what's happening, you remove the 50-ohm termination, and plug in the 50-ohm cable going to the N Unit so that it's always ready to be monitored without disturbing the circuit.

The last case (figure 16d) is where we have a probe with a significant input capacity, such as a cathode-follower, or other probe that might have an input capacity of 3 to 4 picofarads. Three picofarads at 500 megacycles is about 100 ohms, so at 500 megacycles a probe of this sort would load the circuit down with 100 ohms, but of course, at dc it wouldn't be loading it hardly at all. So with this type of probe you will find that it's mainly useful with slower risetimes, because the equivalent circuit shows that the collector load, which may be a few hundred ohms, and the input capacity, which is the order of 3 or 4 picofarads, may form a time constant which limits the risetime at the circuit to a few nanoseconds before you even consider what the circuit is doing. So you have destroyed the fast risetime of the sampling system by adding this extra capacity instead of a uniform resistive loading to the circuit. A further precaution is that some probes have a variable transient response depending upon what this load resistance is—low values may allow the probe to ring, whereas high values will result in an over-damped system. The capacitive probe also grossly alters the current distribution at high frequency—in fact, it may allow tunnel diodes to oscillate.

The physical configuration of resistive probes must, of course, follow what we showed in the previous drawing (figure 15), to minimize the inductance and stray capacity. But it's quite a simple thing to do, and a little practicing will quickly convince you that a resistive probe is far better operationally than a capacitive probe.