Jonathan Scott

Jonathan_scott@AGILENT.COM

577-2504

ACKNOWLEDGEMENT

The author wishes to express his most sincere thanks to all those who have given assistance during the time spent working on this thesis.

In particular I wish to thank my supervisor Dr J Godfrey Lucas, who for his policy of letting his minions keep to themselves, has had to put up with long delays while innumerable tangents were followed to their ends, all the time maintaining encouraging comments and critical assistance. Perhaps pure research could ask little more.

Thanks go to members of the Air Navigation Group, especially Alan Young, Jim Bruderer and Mark Simmonds, all of whom were willing to lend hands, eyes and minds.

I would like to thank members of Electrical Engineering, and in particular Richard Small and David Skellern, to both of whom I oft went with problems, and from whom I always returned the wiser, and Mr Frank Butler, whose drawing skills were greatly appreciated.

I would also like to thank my friends who offered support beyond the call of mere acquaintance.

Finally I wish to acknowledge the help of the University of Sydney and the Department of Transport Australia, whose resources I used for the work - may they benefit from the labours.

DEDICATION

This thesis is dedicated to people whose skills inspire me:

Peter Single, the engineer, and a wise man;

Eric Mills, thinker, problem solver, perfectionist;

Godfrey Lucas, occasional idealist, eternal optimist;

Andrew Taubman, epitome of loyalty, perpetual ally;

Philippa Stone, philosopher, perceiver, mirror.

PREFACE

The original motivation which has led to this work was the desire for an instrument capable of capturing a singly occurring event at civil radar frequencies, and in particular at 1.35 GHz. It is intended that such a device be used to record the brief pulse returned by a radar target in a particular position and orientation, with the object of comparing this to that returned by a false target image (referred to as an 'Angel'). Such a comparison, it is hoped, will lead to a technique for separating the two, a problem of some concern to researchers in the area. It is clear, however, that an instrument with this capacity will be a general purpose tool of some power, with application in many other areas.

The work started as an investigation of currently available oscillographic equipment which might have been up to the task already, then it evolved into a feasibility study of the two very disparate approaches which appeared to offer a chance of fulfilling the requirements. The first approach, which offered higher ultimate performance with presently available resources once the technical hurdles were overcome, was eventually abandoned for several reasons outlined in Chapter 2. Published works of other researchers were seen, with hindsight, to tacitly confirm the author's conclusion that the technology in the relevant area is not up to the task. Microwave Transient Digitisation or MTD, as the instrument's function is boldly termed, is something which has not yet reached commercial development. This is partly due to its specialisation, partly its inherent complexity, and partly its youth as a technical possibility. Despite the specific use within a specialised field for which the instrument to be described is intended, and despite the fact that it is still a large and heavy instrument, tied to other instruments for support, the author has taken care to produce a machine, and to write this report, with an eye open to the future. It is envisaged that this class of instrument will eventually resemble an ordinary laboratory oscilloscope, in the same way that the sampling oscilloscopes of today resemble their continuous forbears, and digital oscilloscopes, which are still relatively new, their analog forbears.

At the outset, the extent to which this instrument would be dependant upon the theory and practice of Digital Signal Processing (DSP) was not envisaged. Proper use of the instrument in its present form demands some knowledge of the appropriate theory. While a device of this design can be configured to hide its more complex functions from the user, as sampling CROs have come more and more to do, there is a sacrifice of performance in this, which has been avoided.

Finally, this thesis is necessarily user's manual and hardware description manual for the 7912ADM MTD. The appropriate chapters are formatted after the fashion popular in the commercial field, for this reason. The reader is asked to recognise this requirement.

T A B L E   O F   C O N T E N T S

Appendices

1. <u>REVIEW</u>

1.1. Introduction

The Cathode Ray Oscilloscope, or CRO, in it's various manifestations is the most ubiquitous laboratory instrument available to the electronic engineer. It is the one instrument that is regarded as essential for many tests and measurements. The reasons for this are severalfold. Information is presented concisely in graphical form, making it fast and easy to assimilate; bandwidth is typically very wide for a single electronic instrument; the modern instrument is very capable and is performing tasks which were previously the domain of other instruments, such as counters, RMS voltmeters and spectrum analysers; finally, the control layout and operation of different units is sufficiently standardised to permit quick adaption from unit to unit.

It follows that a good deal of research and development effort is invested in improving specifications and diversifying the functions of oscillographic equipment. Over the last few years different types of oscilloscope, or different subsections of modular oscilloscope systems, have appeared. Often, rather than enhancing the specifications by pushing an existent design, a fundamentally differently designed instrument is created, but one which retains the familiar format on the outside. The engineer's familiarity with the layout and final effect of the instrument help it to gain acceptance.

The specification which is most sought after is bandwidth. The increasing speeds of new IC's and new technologies have laboratories striving for faster test instruments. Dependant upon which other facet of an instrument can be traded off, the differing approaches have different state-of-the-art bandwidths. A description of currently devised and proven technologies is presented later in this chapter, with a review of their achievements.

1.2. Thesis Considerations

Before proceeding to the review of oscillographic technologies, it is appropriate to discuss the aims of this thesis, to note the layout, and to thus assist the reader to choose in what order he should read the various parts.

The broad aims of the work undertaken are twofold. Initially, an instrument is required for an application demanding a bandwidth approaching 2 GHz, having as high a sensitivity as possible, and being capable of recording a single event, or a small number of events, for use in commercial radar systems.[23] Secondly it is desired to develop a general purpose transient recording system whose design contributes to the field of high frequency measurement and which will be useful in the area of modelling and testing of navigation systems. The aim of the thesis is to report the conclusions of the investigation.

Chapter 1 presents a brief summary of the range of systems available or reported at the time of writing. Chapter 2 describes the two approaches which offered promise of achieving the goals of the work and then reports that investigations eliminated one avenue. Chapter 3 describes the electronic hardware built. Chapters 4 and 5 report the performance measurement techniques and the signal processing algorithms adopted.

The instrument operating system (DIOS) user's manual is included for those who wish to become familiar with it. Appendix A contains a discussion of how an oscillographic system, particularly a digitising one, is measured and evaluated. It should be read if the reader wishes to familiarise himself with the terms involved. Appendix B details an analysis of discontinuities in transmission lines, relevant to conclusions stated in Chapter 2. Appendix C discusses a novel constructional technique devised for the work. Further appendices carry data sheets, mathematical derivations and system documentation which are tendered for completeness, but which it is not intended that the reader should cover, unless specifically interested.

1.3. Current Methods

Regular review papers [1,2,3] are available which summarise the diverse approaches into six categories. However, some of these are irrelevant to this application, while certain fields have effectively subdivided with the progress made. It is intended therefore to discuss each published or marketed genre separately.

1.3.1. Classical Oscilloscopes

The classically designed oscilloscope has all but reached the ideal performance asymptote as defined by the physics of its tube's mechanical layout.[4] This is in the low hundreds of megahertz at best. Figure 1.1 summarises the basic contents of an elementary classical oscilloscope. The electronic technology available today is sufficient to make the cathode ray tube with simple electrostatic deflection plates the sole limiting factor of the instrument's bandwidth. Improving the bandwidth by shortening the plates or further accelerating the electrons reduces sensitivity. If the plates are made closer together then the beam current must be reduced, reducing the maximum writing rate to the phosphor, and so on. With the maximum possible deflection applied, these constraints define an inescapable bandwidth/size-of-image limit. This limit is reflected in what is referred to as the tube's Figure of Merit. Improvements in phosphor technology may enhance this a little, but the physical limit essentially remains.

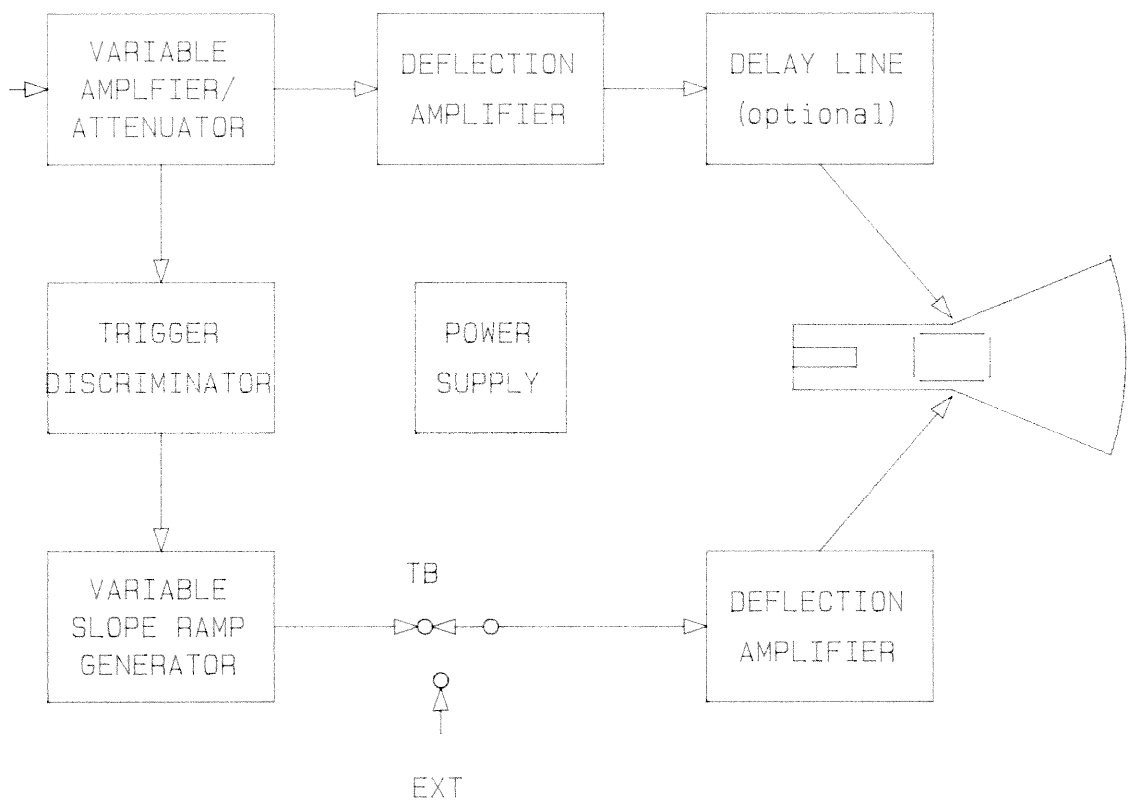

BLOCK DIAGRAM OF A CLASSICAL OSCILLOSCOPE

(a)

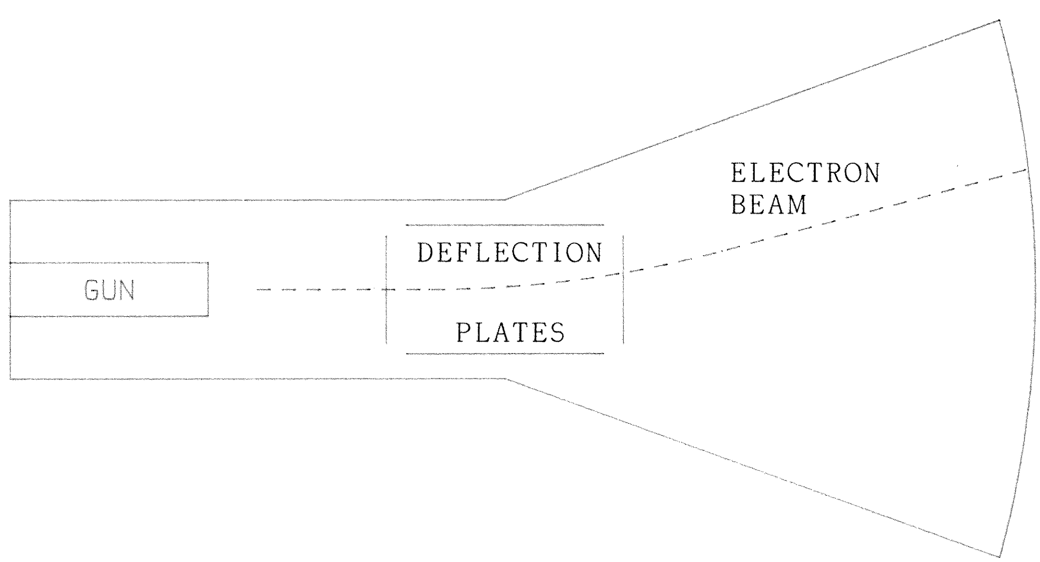

SIMPLE CATHODE RAY TUBE

(b)

Figure 1.1

1.3.2. Oscilloscopes with Travelling Wave Tubes

The first design innovation to enhance tube performance consists of replacing the plain plates with 'plates' which are configured to act as transmission lines with slow group velocities.[1,4] These are terminated at the phosphor end of their length, the deflection signal being injected at the gun end and travelling along the path of the beam at the electron velocity. Thus a signal "accompanies" the electrons it is to deflect for some distance without significant relative velocity. Several mechanical configurations are possible, but a helical structure has been found to be most suitable.[2,4] This type of tube is referred to as a Travelling Wave Tube, or TWT. TWTs will be familiar to microwave power amplifier designers.

Sensitivities of several volts per division are typical. The transmission line impedance seen at the input, and hence the match, depends upon the drive to the plates being push-pull (180° out of phase). The largest native tube bandwidth reported is 5 GHz[4] but this permits a screen of only one to two centimetres and requires several volts of drive per division of deflection. The number of effective bits of vertical resolution are not reported, but must be very low. A bandwidth of 1 GHz is the greatest reported for a system with amplification.[5] This is achieved by the use of filters designed to have a response characteristic which is the inverse of that obtained from the deflection amplifiers. Dynamic range is understandably low, but a sensitivity of 10 mV per division is available and the system is usable to at least 1.4 GHz. The noise performance of the amplifiers and the beam characteristics do not degrade the number of bits of effective vertical resolution below the usual 6 obtained from a standard 8 cm cathode ray tube (CRT).[Appendix A] As in classical oscilloscopes, a permanent record is available via film, but extremely sensitive film must be used for fast events, with its attendant difficulties.

1.3.3. Sampling Oscilloscopes

The widest bandwidths available in a time domain display instrument are those achieved by sampling techniques. Currently well documented room temperature designs achieve bandwidths above 18 GHz,[6] which approach the limit imposed by the semiconducting elements used as switches in the gating process.[7] A markedly different electronic approach employing the same type of diode elements achieves similar bandwidth in commercial applications.[8] Further discussion of sampler hardware will be found in Chapter 2.

The technique of operation implied by the title of sampling oscilloscope, however, achieves these specifications in exchange for certain other capacities. The main limitation is the inability to handle single events. Though the aperture time of the sampling gate may be very short (of the order of tens of picoseconds) it can typically only be opened and closed relatively infrequently. A sample rate of approximately 100 kHz would be an average goal. This delay necessarily occurs as a result of the time required to generate and convey the fast drive pulse to the switching elements, and to process the acquired charge to reveal the actual level at the gate input during the aperture time.

A satisfactory oscillographic instrument must have a bandwidth that is greater than or equal to that part of a system in which the observer is interested. This implies an absolute minimum number of discreet effective samples in an output waveform. A typical expected minimum for a general purpose instrument would be 100 or so, though fewer are sometimes technically sufficient. A general laboratory oscilloscope could, for comparison, typically resolve 0.1 division in 10, or have an equivalent of 100 samples.

A sampling oscilloscope with the above typical values will clearly require at least 1 ms to produce one output display, even though it may depict events in a time window of 1 ns. For this reason the techniques are referred to as 'equivalent time' techniques. Contrary to the implication in the usual description of the operation of sampling equipment, the utilisation of available cycles of the input waveform is very low. An example of time and frequency domain waveforms is given in figure 1.2. They exemplify not only the requirement of periodicity, but also the inefficiency and consequent slowness which is inherent in the sampling (equivalent time) process. In practice, the number of 'unused' waveforms (periods within which no sample is taken, and whose occurrence is thus not helpful) is often much greater than

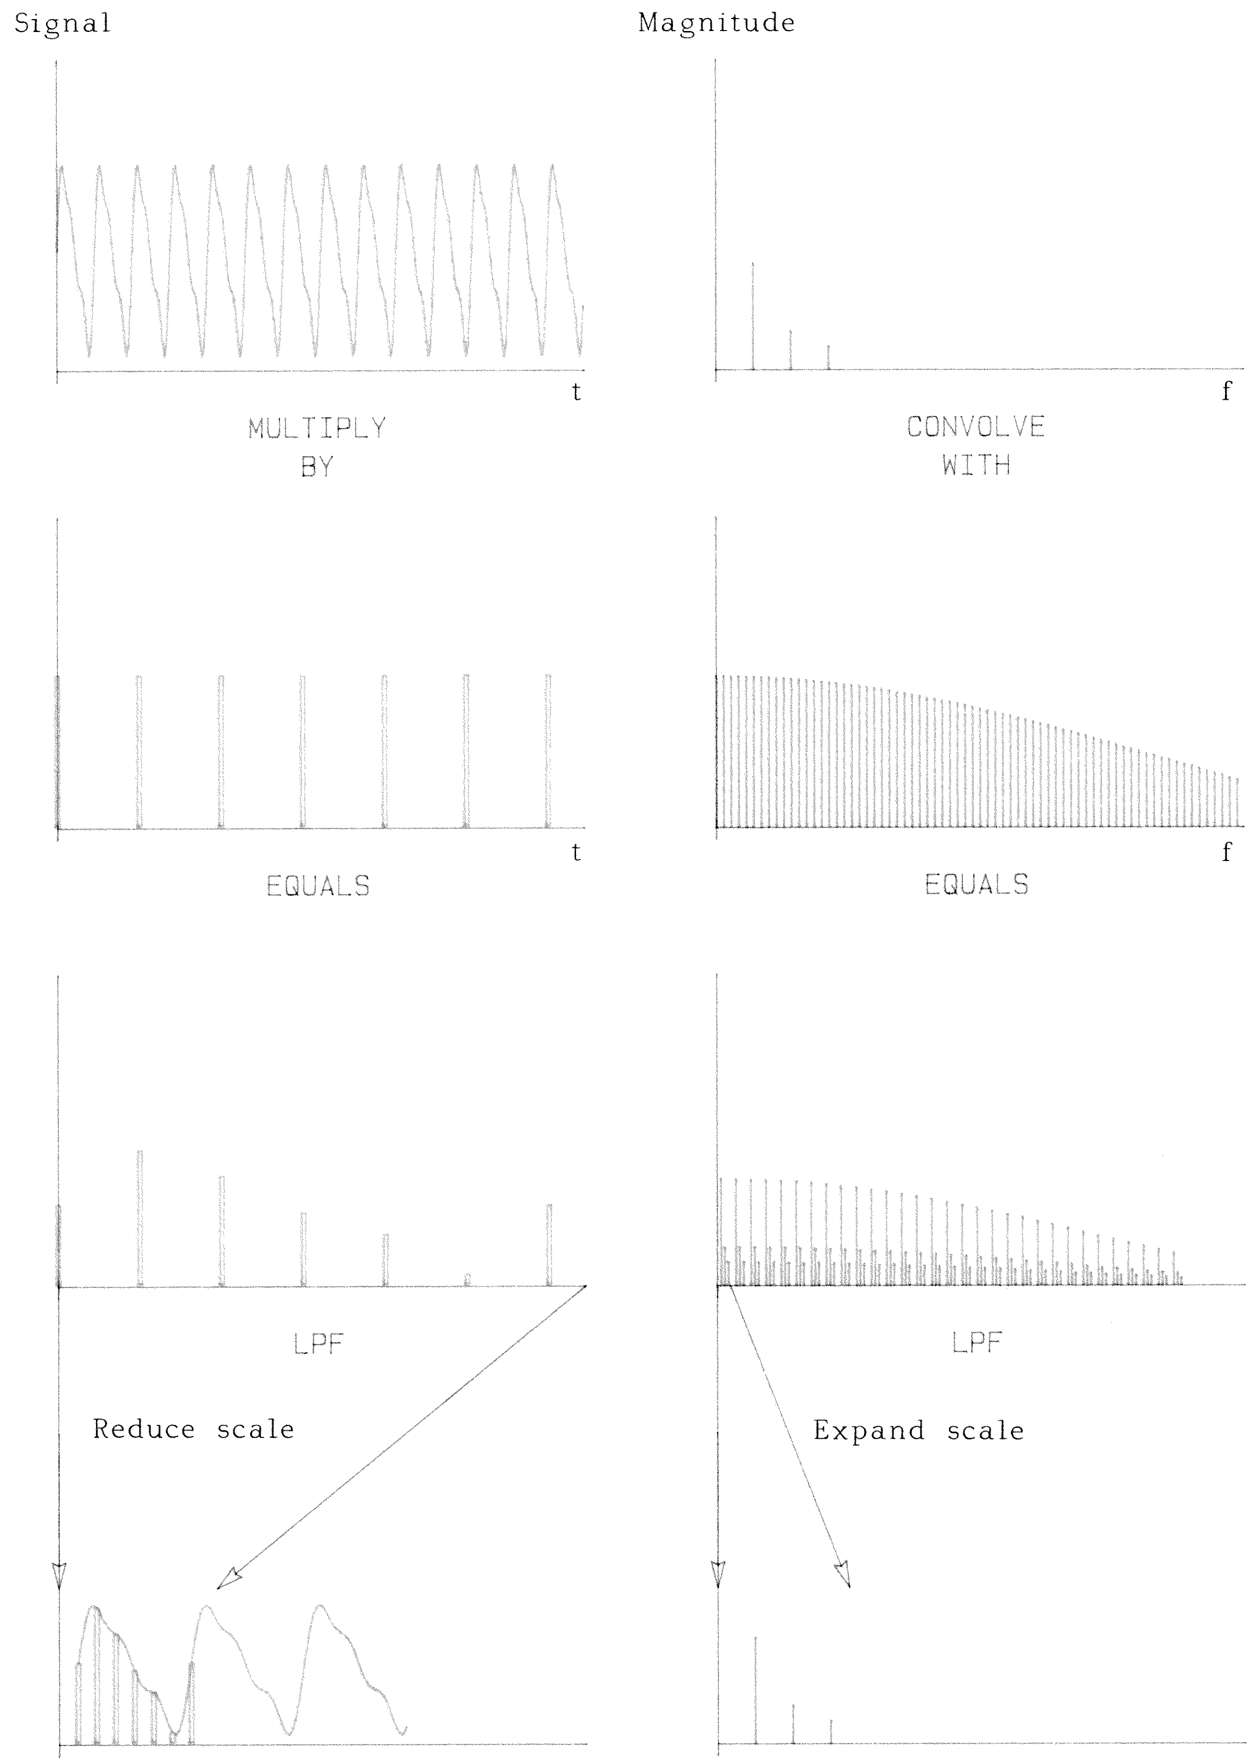
Signal
Magnitude
t
f
MULTIPLY
BY
CONVOLVE
WITH
t
f
EQUALS
EQUALS
LPF
LPF
Reduce scale
Expand scale

Figure 1.2

any figure fitting conveniently on one of these pages could show clearly.

In contrast to their speed, sampling systems offer excellent sensitivity in spite of the complete absence of amplification in the signal path. Full scale sensitivities of less than 20 mV are possible, with noise limited vertical resolution which approaches 5 bits at that sensitivity. In practice signal averaging is used to reduce the noise level, increasing again the inefficiency in the time domain.

1.3.4. High Repetition Rate Sampling Oscilloscopes

This class of instrument is characterised by a sampling system capable of acquiring, processing and storing the results in very quick succession. The primary advantage of collecting the data sequentially in this way is that only a single occurrence is required. Another advantage, however, is inherently found in such a system. The output is most conveniently held in digital form, in Random Access Memory (RAM) usually, because of its speed.

The system sample rate, and hence bandwidth, is determined by the speed of the sampler/converter and the electronics used to pass the results out to slower memory. State of the art logic circuits, which are not yet commercially available, can achieve frequencies above 3.5 GHz,[9] while 300 MHz is commercially available. This means that the latching and storage system is not usually the slower (limiting) element in the system. The fastest presently practicable sampler converters take 200 Megasamples per second[10] with at least 6 bit vertical resolution guaranteed above 20 MHz and 8 bits below this limit. This system does not employ semiconductor logic in the actual converter. Although semiconductor based flash converters are now available with conversion times approaching those of the above system, the latest systems employing them are not yet the fastest available.[11] Thus this approach is limited to a bandwidth somewhat below 100 MHz, and is not viable at microwave frequencies.[3]

1.3.5. Slow Sampling with Signal Replication

(or Low Repetition Rate Sampling with Signal Replication)

It is theoretically possible to replicate a signal, effectively making periodic a single event, by recirculating it in a closed electrical system. It can then be dealt with using techniques with relatively long processing times. The necessary low loss recirculators are only practicable at superconducting temperatures, considerably reducing the appeal of this approach.[12] In addition, while signal recirculation has been demonstrated,[13] criticism is available[2] which suggests that there has yet to be demonstrated the feasibility of recovering the original waveshape from the recirculated copies. Thus at this time the technique does not appear to offer promise for practical systems.

1.3.6. TWT Scan Conversion Oscilloscopes

From the techniques noted so far, the travelling wave deflection cathode ray tube offers the greatest bandwidth in a system which is capable of capturing transients. However, resolution tends to be low, and the output record present only dimly on standard (non-storage) phosphor. A scan conversion tube is one where the phosphor is replaced with a target offering better spatial resolution and/or improved sensitivity, and which allows transfer of the record to some further display system. The improvement in target characteristics permits the overall figure of merit of the tube to be vastly improved.

Two major techniques for fabrication of suitable targets appear in the literature. The first is based on microchannel electron multiplier plates, or simply channel plates.[14] A tube of this type has been produced with a bandwidth of 2.5 GHz, and another with up to 5 GHz bandwidth, though the latter is not yet reported as practicable.[15] As yet unrealised, a design is proposed for such a tube with a raw deflection bandwidth of 7 GHz.[16] The target consists basically of large numbers of minute electron multiplier tubes, typically on 50 micrometer spacings. These feed a phosphor display with an intensified (or 'amplified') image. Future designs are intended to employ a fibre optic pickup arrangement in place of the phosphor screen, permitting eventual direct conversion of the trace to digital form. This proposal has not yet been demonstrated however. The microchannel plate has the added advantage, when used approaching its saturation level, of reducing blooming and excessive contrast effects. This opens the possibility in the future of using charge coupled device digitising targets behind the plate, in place of fibre optic systems.[15,16] Charge Coupled Devices (CCDs) are currently unsuitable because of their excessive tendency to flood with electrons when locally overloaded.[84]

The second scan conversion design replaces the phosphor with a diode array, fabricated by well explored microelectronic manufacturing processes. The target captures the waveform as charge distribution on the semiconductor plate, which is sensed at relatively low speed by a second electron gun moving in a raster pattern. A tube with >2.5 GHz basic response is available.[17] Owing to the proprietary origin of this tube, little data is published regarding current developments, but it would seem safe to assume that research continues towards improvement of this specification. The target technology seems able to achieve better resolution than other methods, and the tube should thus at least achieve similar figures of merit. Research on the semiconductor target's sensitivity and resolution proceeds,[18] though again at a proprietary level, despite the lack of intention to market the product.[18,19]

Both approaches apply the input signal directly to the deflection plates, in order to achieve the stated bandwidths. This has avoided the need to construct amplifying elements with the required bandwidth and drive capability. The channel plate approach currently offers sensitivities of several volts for full deflection, though the number of bits of resolution for this level is not stated.[15] It promises sensitivities of better than this figure with future target designs. The plate characteristic impedance is nominally 50 Ohms. The diode target tube has a sensitivity of 8 volts peak for full deflection, offering at least 8 bits of resolution, dependant upon subsequent processing.[17] The input impedance is 180 Ohms. The drives must be push-pull, applied out of phase, with the stated sensitivities referring to each plate. A complex passive network may be added to remove the restriction of complementary drive. The higher impedance value of the second tube type can be utilised to advantage. A filter network has been realised which improves the bandwidth to 3.5 GHz for single ended drive, with the same voltage specification into 50 instead of 180 Ohms.[18]

The scan conversion technique thus offers good specification now, and promises improvement. A more detailed discussion of these tubes will be found in Chapter 2.

1.3.7. Multiple Sampling Gate Oscilloscopes

Because a sampling gate theoretically appears as an open circuit when not actually open to sample the line, such gates are frequently designed as feedthrough systems. This makes cascaded connection of many gates possible. An array of samplers placed along a transmission line, suitably synchronised to sample at the appropriate moment to capture different portions of the waveform on the line, forms a sampling oscilloscope capable of capturing a single occurrence. Such a system potentially offers both sensitivity and bandwidth which are comparable to the traditional sampling system without the restrictions of equivalent time operation. The obvious penalties paid in the tradeoff are complexity, the difficulty of synchronisation and the fixed timebase.

The highest reported bandwidth for such a system is 3 GHz,[20] where all sampling gates were driven by the one pulse generator system, removing synchronisation problems. The system obtained an absolute maximum of 40 samples, with a claimed random noise of approximately 1 mV in each sample. Some of the difficulties of such a system are already well documented.[20]

A bandwidth marginally inferior to this is reported for a system where synchronisation of the drivers is required.[21] In this case however, only 18 samples were obtained, reducing the synchronisation problem considerably, as half that number of gates were driven by each generator circuit. In both cases the samplers themselves resembled those used by sampling oscilloscopes.

A novel design has also been reported which employed an interrogate-pulse technique.[22] Here the samplers are passive, being driven by a pulse passed to them from the same line as that carrying the signal to be sampled, but travelling in the opposite direction. Although offering a number of other advantages, such as simplicity and low cost, this technique achieved only a low vertical resolution of approximately 4 bits, and a bandwidth of less than 1 GHz.

1.4. Summary

Two of the techniques discussed show potential for achieving transient capture with gigahertz bandwidth. They are those based on scan conversion cathode ray tubes, and those based around cascaded sampling gates. These will be examined in further detail in following sections.

2. <u>DESIGN OF A TRANSIENT DIGITISER</u>

2.0.1. Chapter Overview

Chapter 2 is divided into two major sections. It was noted in Chapter 1 that two techniques apparently offer a method of realising a Microwave Transient Digitiser (MTD) with the specifications sought. Each section deals with one of these techniques.

Section 1 (with Appendix B) deals with the design of a digitiser based on cascaded discrete sampling gates. It indicates that such a machine has been proven feasible, but with performance limitations.

Section 2 discusses scan conversion tube techniques. These show promise of improvement in the future. There are also currently available representatives with promising capabilities approaching those required in some instances. These are discussed in Section 3. The feasibility of modifying a current commercial 500 MHz system to meet design requirements is indicated.

Conclusions about the future of the techniques are drawn. The value to general research effort of the available alternatives is taken into account.

2.1. Multiple Sampling Gate approach

2.1.1. Overview

A simple block diagram of a Multiple Sampling Gate Transient Digitiser (MSG TD) is given in Figure 2.1. A signal introduced at the input of the transmission line propagates toward the termination. The sampling gates are arranged to open and close simultaneously. The period for which the gate is open is referred to as the aperture time, T_a. During this period the gate acquires a sample which is proportional to the average input voltage present. The resultant set of samples "photographs" the waveform present on the transmission line at the instant of sampling.

The intersampler spacing (in units of time) is effectively defined by the length and propagation velocity of the transmission line sections between the samplers, and is labelled T_d. For a single line system as shown, there are immediately two constraints on the bandwidth. The sampling theorem states-

$$F_{max} = 1/(2T_d) \qquad (2\text{-}1)$$

where F_{max} is the largest frequency resolvable without aliasing. Also the aperture function Sinc(pi.ft) requires-

$$F_{3dB} = (\mathrm{Sinc}^{-1}(2^{-.5}))/(pi.T_a) \qquad (2\text{-}2)$$

BLOCK DIAGRAM OF AN MSG TRANSIENT DIGITISER

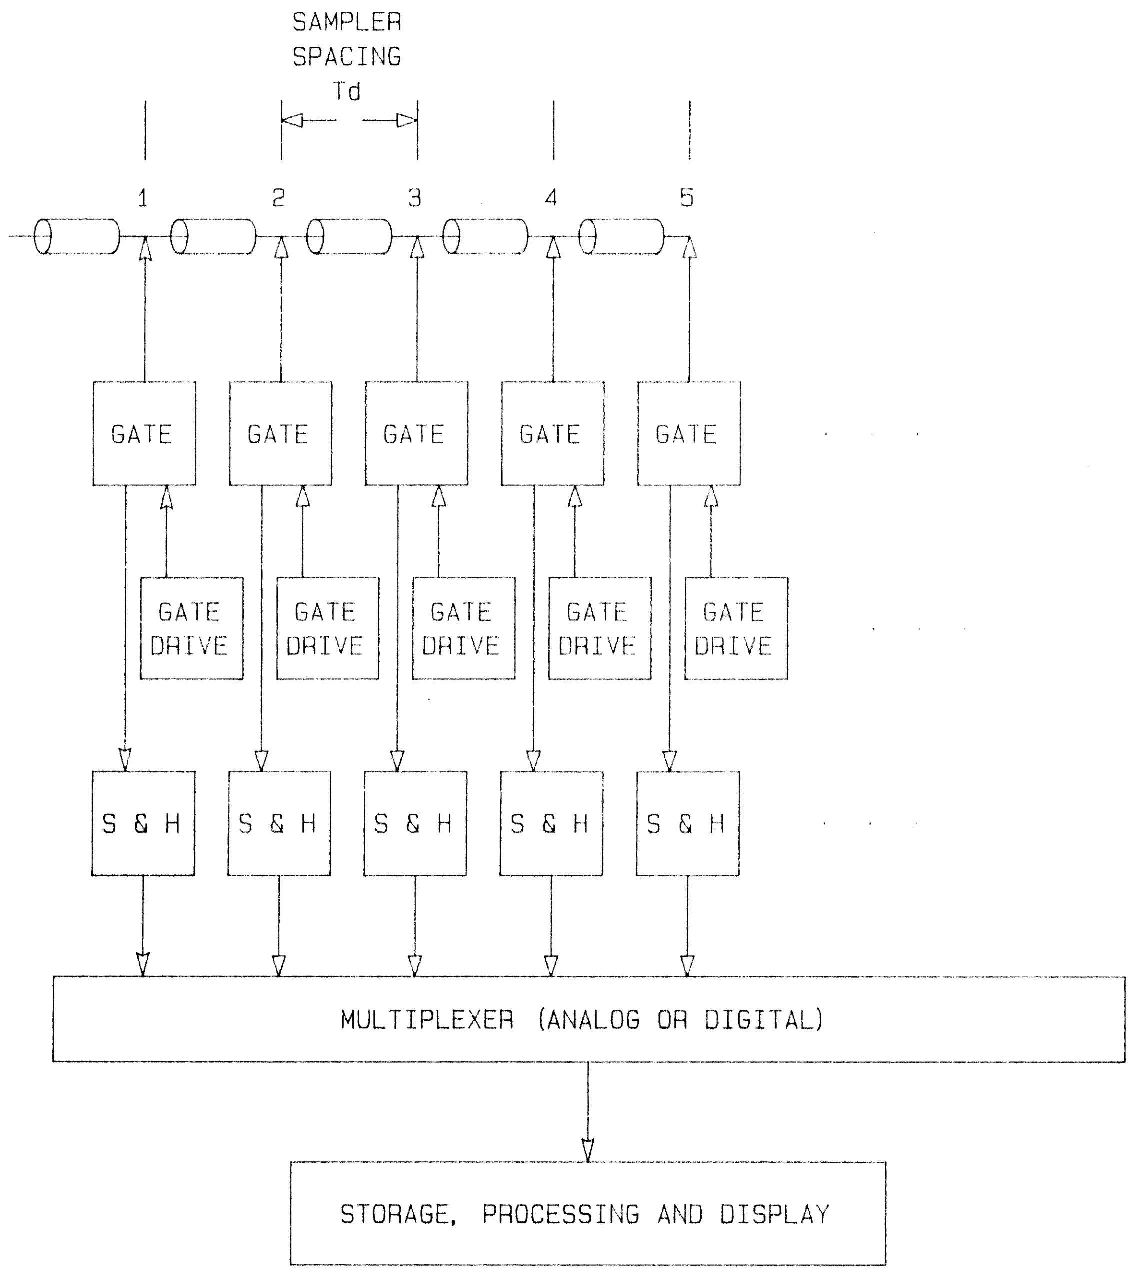

EXAMPLE WAVEFORM RECONSTRUCTED FROM SAMPLES TAKEN

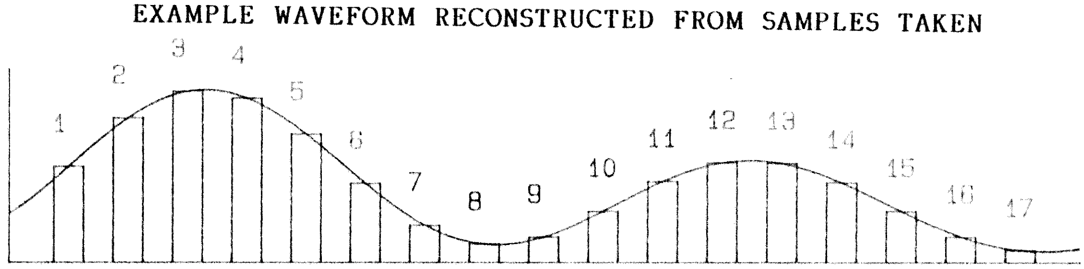

Figure 2.1

Placing these two frequency limits at the same value gives the relation

$$T_a = 0.866\ T_d \tag{2-3}$$

This implies that the samplers should be approximately as far apart in time as their apertures are wide; little is to be gained from closer spacing, and bandwidth is lost if wider spacing is used. (In practice the corruption signals to be discussed below prevent successful operation with gates more closely spaced.)

Disregarding the DC component, minimum resolvable frequency in a system with N gates will be

$$F_{min} = 1/((N-1).T_d) \tag{2-4}$$

For a practical system this implies a minimum number of samples. A general purpose system would be expected to achieve somewhere in the vicinity of 50:1 bandwidth, or 100 samples.

2.1.2. Sampling Gate Design

The basic equivalent circuit of a sample gate is shown in Figure 2.2a. When taking a sample the switch closes for the aperture time. The capacitor charges up to a fraction of the voltage on the sample line. The fraction of the actual voltage to which the capacitor charges depends upon the values of the capacitor, series resistor and the aperture time. In a practical

THE EQUIVALENT CIRCUIT OF AN ELEMENTARY SAMPLING GATE

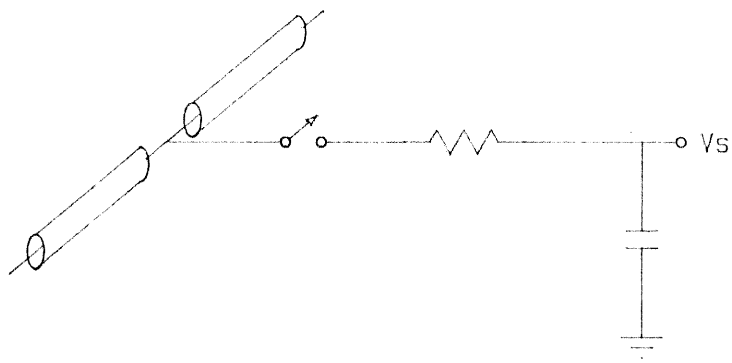

(a)

THE OUTLINE OF A PRACTICAL SAMPLING GATE ARRANGEMENT

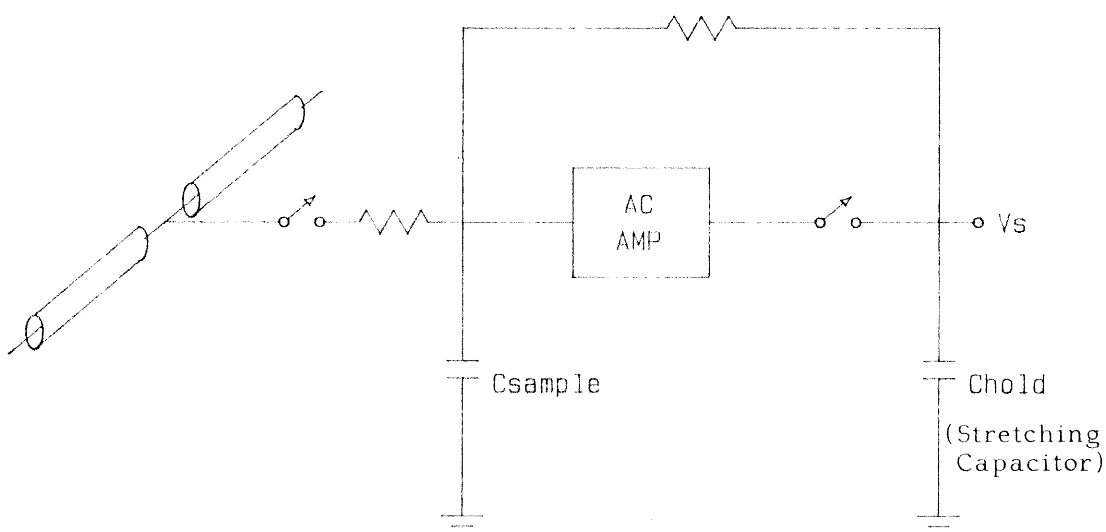

(b)

Figure 2.2

circuit with picosecond aperture time this fraction is typically 5%.[24] Here the sample line is shown as a feedthrough arrangement, though it can as well be terminated at the sample point.

A secondary effect of the switch momentarily closing is the corruption of the waveshape present on the sampling transmission line by the removal of energy to the capacitor. This is referred to as the introduction of sampling 'kickback'. This may equivalently be viewed as the introduction of an external signal which subsequently propagates away in both directions down the transmission line. Although less than 1 femtojoule may typically be removed, the kickback usually introduces a significant percentage error distortion in the amplitude.[20] The kickback signal limits the proximity of gates, since the corrupted voltage on the line must not be confused with the true signal by adjacent gates.

In a practical situation the charge trapped on the capacitor will leak away. Thus an amplifier is incorporated to facilitate the measurement of the charge magnitude. The circuit effectively operates by integrating the voltage on the capacitor as it is charged and subsequently discharged, and amplifying the result. The arrangement which is typically used is depicted in Figure 2.2b. An AC coupled amplifier is employed, eliminating DC offset and drift problems, which can swamp small signals. At the instant sampling commences, the output of the amplifier is connected to the second ('stretching') capacitor. It remains so connected for a duration longer than the aperture time, however, integrating the output of the AC amplifier. The integration effect arises from the finite output impedance of the amplifier. If the gain from the input to the stretching capacitor is correctly adjusted, the stretching capacitor will be charged up to exactly the value of the sampling line voltage, which was to be measured. This process depends upon the fact that all the time constants are invariant from sample to sample, of course. The feedback resistor then charges the sample capacitor up to this same level.

The advantage to this latter feeding back of the determined level to the sampling capacitor is found when subsequent samples taken by the same gate are either at the same voltage level or a nearby one. It is common in equivalent-time instruments to effect some signal averaging by having the sampling gate operate several times at the same instant within the period of the waveform to be sampled, averaging out random (non-systematic) errors. In this situation the subsequent operations of the sampling gate provide signals that serve to trim the result, correcting for small amplitude errors. This occurs because no signal will be amplified if there is no difference between the voltage on the line and the sampling capacitor voltage, since no current will flow across the switch during the aperture interval. However, a correction signal will be applied if there had been an error and a voltage difference thus existed between the poles of the sampling switch. This error trimming is referred to as 'null seeking', and the circuit arrangement is called a 'null seeking loop'. The loop gain is sometimes reduced when this mode is used, requiring several samples to seek the correct level, but reducing jitter and noise by effectively giving less weight to each sample.

The sampling switch is invariably realised by using very fast diodes. Figure 2.3 gives the circuit of the preferred 2-diode sampling gate, and block diagrams of its supporting circuitry.[6,7,20,24] The general mode of operation is relatively straightforward.

The bias supply keeps the sampling diodes reverse biassed when the gate is closed. Equal but inverted pulses are applied by the pulse shaping circuitry to the sampling capacitances, forward biassing the diodes for the sampling aperture interval. During the interval, charge flows around the loop formed by the pulse shaper, the sampling capacitances and the diodes. When zero voltage is present on the sampling line, the charges accumulated on the two sampling capacitances are equal but opposite, assuming that the components around the loop are matched. When the line, which appears as a generator of internal impedance equal to half the line characteristic impedance, offers a non-zero voltage, the charges differ. Such a difference in the charges on the sampling capacitors results in a non-zero voltage appearing at the pulse amplifier input.

This voltage is amplified as previously noted, capturing the sampled signal value before the charge leaks away through the

CIRCUIT OF 2-DIODE GATE WITH BLOCK DIAGRAM OF SUPPORT SYSTEMS FOR BIAS, DRIVE AND PROCESSING

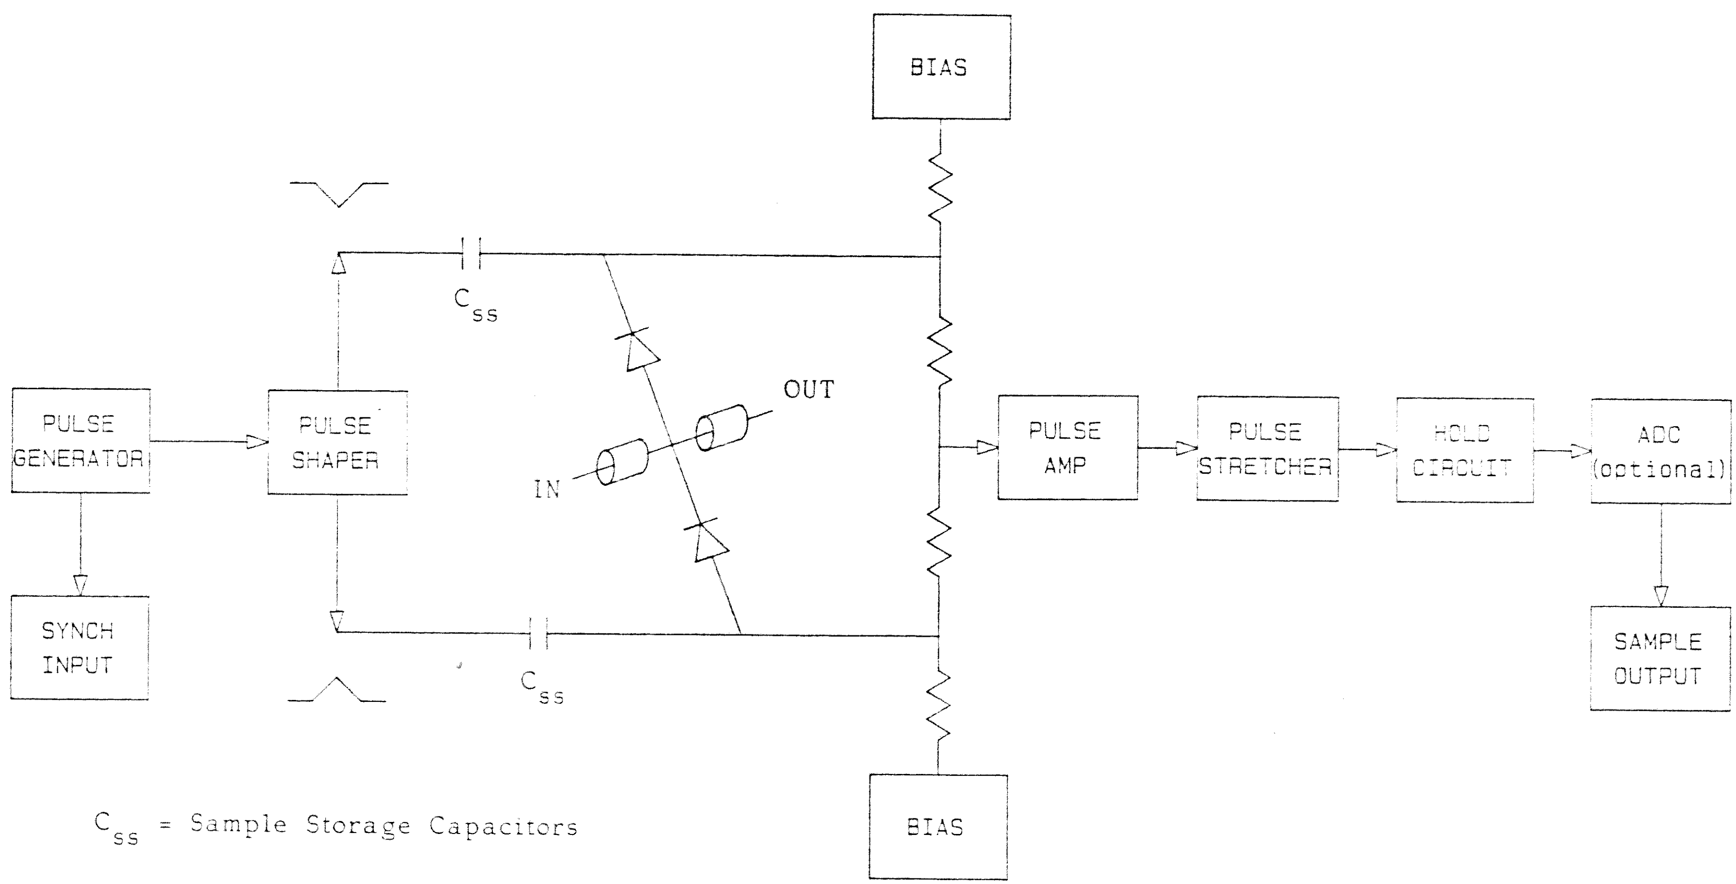

C_{ss} = Sample Storage Capacitors

Figure 2.3

bias circuitry. The DC feedback system then adjusts the bias voltages to prepare for the next operation of the gate.

It is shown in reference 7 that the gate equivalent circuit reduces to the small signal equivalent circuit reproduced in Figure 2.4. Making certain simplifying assumptions based on the realisable properties of the hot-carrier diodes employed as the switching elements, and assuming that the <u>package</u> capacitances (C_p) of these diodes can be masked out by an subtracting an equivalent capacitance from the distributed capacitance of the sampling line, it is shown in the same reference that the transfer function from line source voltage to junction voltage is of second order and depends upon:

C_j = junction capacitance of the sampling diode,

R_s = ohmic series resistance of the diode,

R_{sg} = pulse generator series resistance,

L_s = diode package series inductance,

L_2 = inductance between earth of the sampling line and the pulse generator earth connections,

R_j = dynamic impedance of the forward biassed diode,

C_s = sampling bridge capacitor value,

and ZO = characteristic impedance of the sampling line.

The transfer function is rather complicated, and so is also set out in Figure 2.4.

This function is found to place extremely stringent constraints on the diode parameters and the assembly in order to achieve wide

$$V_d = \frac{e_{in}\left[\frac{1}{LC}\right]}{s^2 + s\left[\frac{1}{CR_1} + \frac{R + Z_o/2}{L}\right] + \left[\frac{R + Z_o/2}{R_1 CL} + \frac{1}{LC}\right]}$$

$L_1 = L_s/2$
$R = R_s/2$
$R_1 = R_j/2$
$C = 2C_j$
$C_1 = 2C_s$
$C_2 = 2C_p$

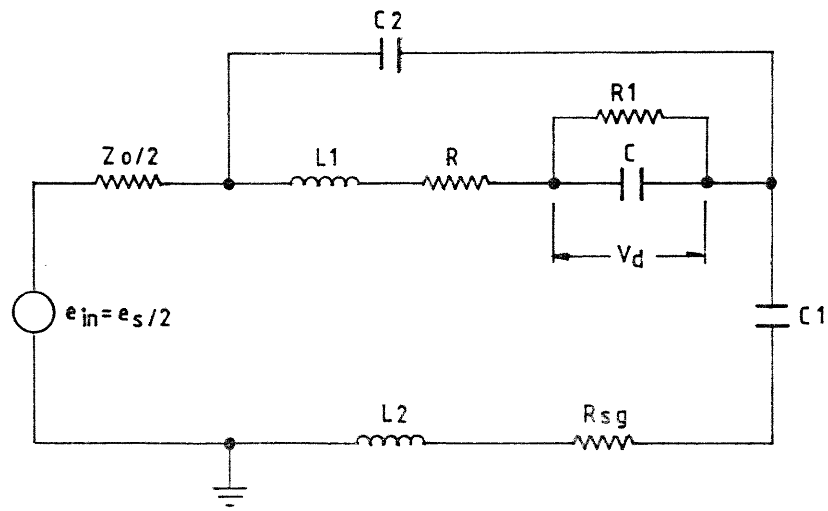

EQUIVALENT CIRCUIT OF 2-DIODE SAMPLING GATE AND ITS TRANSFER FUNCTION

FIGURE. 2.4

bandwidth with acceptable damping. Grove[7] employs an extremely ingenious, 3-dimensional, precisely machined assembly to effectively set $L_{sg} = 0$ by making the two transmission line earths the same point in space. Masking out the package capacitance (not the major capacitance in the circuit) at the cost of overshoot specification, the following diode parameters are consistent with the design bandwidth of 12.4 GHz for the sampler response:

$$C_j = 0.1 \text{ pF}$$
$$R_s = 20 \text{ Ohms}$$
$$L_p < 500 \text{ pH} \qquad (6)$$

Such values are barely obtainable, and the cost of diodes of this calibre is typically some hundreds of dollars.

In addition, not only the bandwidth, but also the sensitivity of the sampling gate in operation and the degree to which it approaches an ideal open circuit when not sampling, are affected by diode parameters and assembly considerations. The isolation of the sampler drive pulses from the signal being sampled depends upon the diodes as well as the resistors in the bridge being identical. If the isolation degrades a primarily level-dependant error is added to the resulting charge. In an instrument where there is no facility for intelligent post processing of the samples (that is a non-digitising system) this directly compromises the reliability of a sample value. As indicated in Appendix A, this does not present as significant a problem in the digitising instrument situation, where even non-linear corrections are simply made, once the error function is known.

The problem of the gate loading the line whilst in the isolated state is more serious. The reverse biassed diodes of the bridge appear as their junction capacitances in series with the drive and bias circuit, placed across the sampler transmission line. In addition there may be diode package capacitance to be taken into consideration. Some alleviation of this problem is to be found by removal of some of the distributed capacitance required to maintain the characteristic impedance of the line. However, the loading capacitances are not directly connected to ground, but to the pulse driver and bias circuits, so that direct subtraction is not a perfect solution. Even disregarding this, there is a second problem. Since the junction capacitances vary non-linearly with applied bias, it would never be possible for them to be exactly masked with a constant subtraction. Further, they are not usually specifically controlled in manufacture, so their values are not easily available for individual trimming. There is also a limit to the capacitance available for removal from any given transmission line before the distance involved becomes a sufficient fraction of a wavelength so that it cannot be regarded as subtracting from a lumped element. This whole problem is mentioned later, and dealt with more fully in Appendix B. As a typical example of manifestation, Grove's design[7] suffers from a reflection coefficient of magnitude approximately 0.3 at its full design bandwidth; indeed 0.1 pF of unmasked capacitance will represent a reflection coefficient of approximately 0.1 at 6 GHz.

An alternate scheme[8] for reducing the perturbation to the sample line by the sampling gate achieves a better SWR but does so at the cost of greatly extended kickback, requires a significantly more expensive hybrid arrangement of six diodes, and has a physical length much greater than the simpler two diode arrangement. Such a diode assembly alone can cost in excess of one thousand dollars, and presents problems in cascading (both because of the kickback, and the physical size).

An ideal gating pulse is square, with amplitude sufficient both to overcome the diode reverse bias, and force the diodes to conduct, and with width equal to the required aperture interval, T_a. A pulse narrow enough even to enable a bandwidth of 1 GHz cannot directly be developed with technology which is currently available. It is possible, however, to obtain step functions with small risetimes, and these can be converted to impulse-like functions by introducing them to shorted transmission line sections. This technique produces a pulse whose rising and falling edges have the shape of the original step and its mirror image.

The gating pulse is invariably generated by the use of a two step process. The first is the development of a triggered fast step by means of avalanche mode operation of a bipolar transistor.[25,26,27] The second is the shaping of that pulse by means of a Step Recovery Diode circuit.[6,20,24,28,29,30] The avalanche mode of operation employs the negative resistance region of a bipolar transistor to obtain the fastest possible risetime from a triggerable circuit, consistent with an amplitude in excess of the built-in voltage of hot carrier diodes. (Tunnel diode circuits can generate a sharper pulse but are limited in their amplitude to the 500mV or so of negative resistance region in their forward bias transfer curve.) The fastest reported avalanche pulse circuits achieve amplitudes of 10-15 volts with risetimes of 130ps.[27] These circuits use common transistors. Investigations of performance in avalanche switching circuits carried out, indicate that approximately 60 volts is the highest amplitude that can be expected without risetime falling below 1ns.[25] Although avalanche multiplication takes only some 10ps[26], the author's measurements suggest that the depletion capacitances of devices having sustaining voltages above this value are sufficient to restrict the dV/dT to values giving slower final performance.

If the risetime of the pulse produced is insufficient for the required gate aperture and sensitivity, or the maximum sensitivity possible is desired, the pulse is refined by a Step Recovery Diode (SRD) circuit. Several different arrangements are common[28] but they all rely on the same property of silicon PIN junctions. When forward biassed, charge accumulates in the depletion region. The SRD is distinguished by a very strong dependence of internal dynamic impedance upon this stored charge. When reverse biassed the impedance remains low until the charge is removed, whereupon it reverts to the normal back biassed high impedance.[30] This transition is very rapid, and can be used to develop the required fast edge.

The gating pulse finally used to turn the sampling diodes on and off is normally roughly triangular. This is because it is required to be so narrow that the flat top section is reduced to zero width, composing the pulse of the fastest moving part of the original step. The actual pulse shape does not directly define aperture time or thus affect bandwidth.[7] Provided the pulse is of sufficiently narrow construction that the aperture interval can be reliably fixed with bias adjustment, etc, the step risetime can exceed the aperture time with only reduction of sampling efficiency. This is illustrated in Figure 2.5. (Sampling efficiency is defined by the ratio of voltage accumulated on the sample gate capacitor to the voltage present on the line, and thus reflects the ultimate sensitivity of the sample gate. It is affected by the degree to which the diodes are forward biassed by the driving pulse.) This reduction comes about because the "triangular" drive pulse is able to turn on the diodes less hard, though for as short a time as, a narrower, rectangular pulse.

2.1.3. Sampling Gate Requirements for MSG TD

To date in the literature only one sampling gate design has been reported which differs in any fundamental way from the standard sampler designs developed for single gate applications.[22] This design has not been pursued here, and no further reference could

DEFINITION OF APERTURE INTERVAL

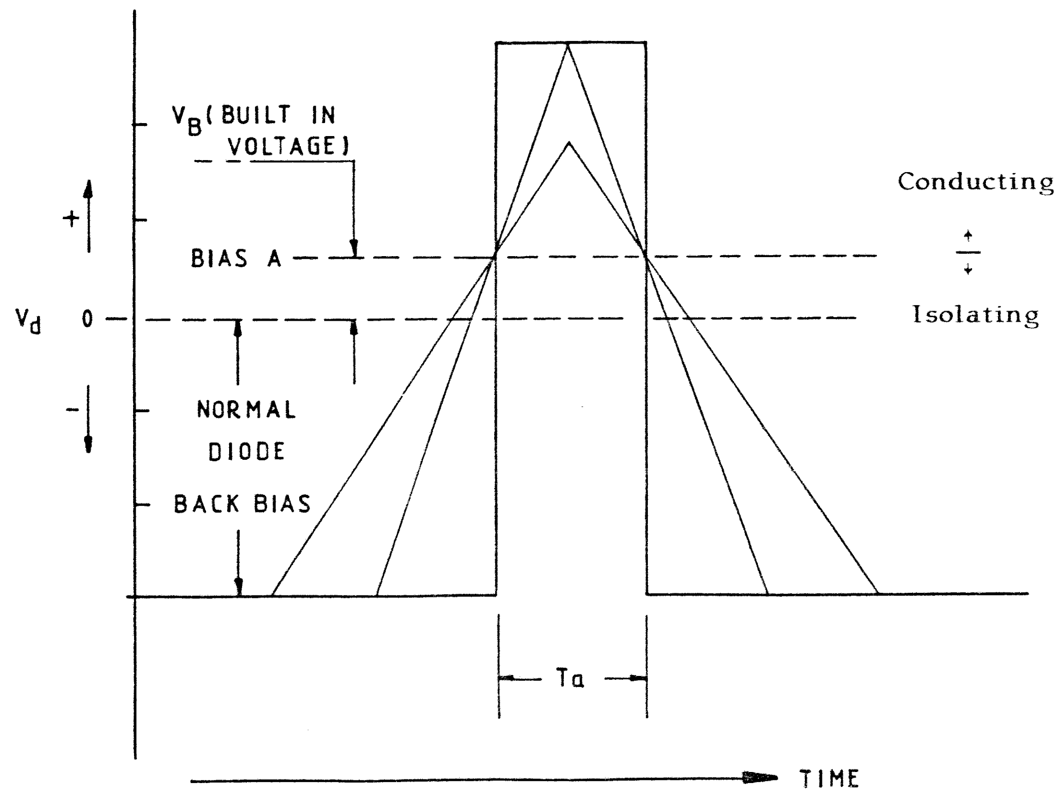

Note that the Aperture Interval, T_a, does not alter with the varying pulse shapes depicted.

FIGURE. 2.5

be found in the general literature. While reducing cost and complexity it offered a significantly reduced performance to cost ratio, leading to its apparent unpopularity.

Other designs reported have made little or no attempt to alter the basic unit from the form refined for single gate situations. Transient oscilloscopes utilising such gates have either not encountered the limitations of cascading these gates, or have tolerated them.[20,21] Problems with cascaded gates have to date not been serious because of the small number of actual gates employed. Three properties of a sampling assembly assume much greater import with the cascading of many units.

2.1.3.a Cost

A conventional sampling oscillography instrument typically costs from $10,000 to $20,000 and may employ one or at most two gates. Approximately $2,000 of the final price is assigned to each sampler along with its immediate support electronics. While this is a most satisfactory breakdown of costs in that situation it is clearly unacceptable to allot such an amount to the realisation of a component which is required in tens or hundreds per instrument. The most expensive single task instruments cost in the vicinity of $50,000. Making the conservative estimate that one fifth part of the expense represents components and materials, all parts in a Transient Digitiser should require a total of approximately $10,000, at the very most!

2.1.3.b Synchronisation

Referring again to Figure 2.1, it is evident that all the samplers must sample simultaneously in order to gather meaningful data. Failure of a gate to open and close at the same instant as other gates would give an equivalent of severe short term non-linearity in the timebase of a conventional CRO, even without taking into account the distortion effects produced by one sampler seeing the kickback of another. This is not a requirement of conventional sampling systems. The cyclic nature of the process renders it uncaring about delays between trigger and sample, provided the delay does not change within the period of one real time sweep, typically a few milliseconds. (This may be seen by considering the frequency domain analysis of the process given in Figure 1.2.)

The problem of kickback interference proves to be of minimal extent. The kickback of one sampler of the most modern type[7] has been measured by this author to produce 2% to 5% distortion when picked up directly by another.

2.1.3.c Discontinuity

In a single-sampler instrument the passive line discontinuity introduced by the gate when not sampling, though not desirable, is not of great consequence. Where feedthrough operation is not vital attenuators may be inserted to dampen reflections. Elsewhere the SWR must be tolerated, but a value even a few times worse than a single connector is rarely sufficient to perturb the signals to a degree which cannot be accepted. In a transient arrangement of cascaded samplers this is not the case. The signal must be fed through, and must pass all but one gate to reach the last in the line. Techniques to minimise the discontinuity will be dealt with, but it is shown in Appendix B that even the small reflection at each gate, achievable within the cost constraints, causes severe corruption after a number of such interferences.

The strength of the above performance constraints (defining an acceptable sampler for use in a cascading situation) are dependant upon the number to be cascaded as well as the individual sample requirements. Before considering the feasibility of realising the function in this manner it is important to know how many samples are to be taken.

Previous designs have not required the serial connection of more than twenty gates.[20,21,2,3] At most forty samples have been provided in the permanent record of any transient event captured by individual samplers.[20,21,22] In all the cases reported in the literature the designer sought to measure characteristics of a signal whose waveshape was approximately known beforehand. This permitted the use of few samples, as a characteristic is typically either local to some part of the signal graph or ultimately resolvable on lower frequency systems. As noted earlier, a general purpose instrument, or one where the waveform to be subsequently presented is unknown, should have in the vicinity of 100 samples.

Many transient digitising instruments are available for frequencies below a few megahertz. A feature incorporated in such devices is the ability to reconfigure the semiconductor memory, permitting the acquisition of few long records or many brief records. This flexibility gives designers an incentive to extend the window lengths both up, to permit long events to be characterised with good frequency resolution, and down, to offer many event stores with the available memory. In spite of this the lower record length cited is 128 samples. (Powers of two are preferred for the obvious convenience reasons.) All the referees implicitly conclude that this is a minimum desirable record size. The fifty (or sixty-four) to one bandwidth may seem unduly high initially. However, note that a basically square pulse or pulse-modulated waveform requires many harmonics to properly characterise it. Equivalently in the time domain, consider that such a waveform will need closely spaced samples where it is moving quickly or changing its envelope quickly, but will waste them in areas where little information is present. Certain recent digitisers which off-load their samples to memory as they are taken have the ability to change sample rate as the window is passed, but this facility is not possible with the fixed system proposed here.

The remainder of this section discusses the feasibility of adapting a sampler of the type presently employed in single sampler systems to the task of forming a cascaded array of samplers for transient digitisation.

2.1.4. Design Feasibility of an MTD Sampling Gate

The first constraint upon the sampler design will be the tenfold to thirtyfold reduction in cost which is imposed by the hundredfold increase in the their number. Even these figures greatly weight the expense of the instrument in favour of the samplers themselves, and puts the instrument in the highest class of systems available today. The component cost for each of the replicated sets of blocks of figure 2.1 must be kept to approximately $50.

2.1.4.a Fabrication Techniques

Current designs involve machined metal parts and exotic substrates in the sampler heads and their connections.[7,8,6,24] Such constructional techniques must be abandoned for two reasons. Firstly the cost of such methods is too great in the light of the above analysis. Secondly, because this approach is not practical for devices measuring more than a few centimetres in any dimension, connectors would have to be used between samplers or small groups of samplers, in order to achieve the required spacing between elements of the cascade. The required number of connectors would together contribute an unacceptable set of discontinuities, even with SMA or the latest K-band types.[49,appendix B] The only alternative constructional method is to use photographically printed circuit board (PCB). This technique offers easy component mounting, can be used to produce units of sufficient size, and can provide transmission lines of adequate quality.

Microstrip[31] and the more recent CoPlanar Waveguide[32] (CPW) are structures which permit the realisation of transmission lines on dielectric board. Microstrip requires continuous metallisation on the underside of the board to provide a groundplane, the signal travelling in a mode about a strip on the upper surface. Coplanar waveguide places both ground and active conductors on the same side of the supporting dielectric slab. This difference in configuration makes microstrip significantly less suitable than CPW for the connection of components in shunt or shunt-series with the guided signal, as is required in the sampling circuits reviewed. Microstrip does not permit easy realisation of high quality short circuits for the same reason of separation of the groundplane. CPW also boasts lower dispersion and loss.[34,35] Although considerably more work has been reported on microstrip, particularly with respect to discontinuities, giving a better design basis,[37,34,35] CPW is a more suitable technique for the application at hand. Microstrip has dimensions which are a strong function of the substrate thickness, whereas CPW does not. In practice, on suitable substrate, CPW pattern dimensions can vary with little effect upon line properties. In other words, provided the shape of the pattern is kept unaltered, the dimensions are purely relative. CPW has the property that it requires narrow slots to achieve a characteristic impedance of 50 Ohms on the more economical high grade substrates, with lower dielectric constants.[33,34,36] Normally a restriction, limiting the range of impedances available, this property works to advantage in this situation. The slots can be sized to accept beam lead components, the dimensional sensitivity to substrate thickness (and thus variation in it) is lower with lower dielectric constant, and a smaller absolute distance increase is required in the slot to subtract distributed capacitance.

2.1.4.b Semiconductors

Another major cost in conventional samplers is the semiconductors, mainly the switching and pulse shaping diodes. Designs cited earlier use proprietary devices packaged in a fashion specific to the environment used, and are available at great cost. The requirement that the diodes in the sampling gate bridge switch in sub-nanosecond intervals demands the use of hot carrier Schottky devices. For minimum package parasitics, and in order to be compatible with the printed circuit construction, it is necessary to use beam lead devices. Though intended for hybrid circuits as well as Microwave Integrated Circuits (MICs),[38] the author has developed a technique of bonding this package type directly to a PCB conductor. This involves bonding the devices directly using conductive adhesives.

(This and a variety of other techniques were investigated, including ultrasonic bonding, thermal and electrical welding and soldering. The conductive adhesive approach produced superior DC and RF performance with greater safety to the device in handling. Bonds did not deteriorate over a period of years. Connection imperfections are negligible compared to package parasitics and device parameters for the diodes tested and discussed in this work. The technique is now in common use here in the microwave laboratories of the University. The technique is outlined in Appendix C.)

A finite number of low cost, beam lead, hot carrier devices are available.[38,39] Hewlett-Packard-recommended sampling and gating type HP 5082-2837 beam lead diode is available for below $10 per unit. A version of this diode is used successfully to well above 3 GHz in the fastest Transient Recorder reported to date.[20] Since the bridge arrangement of a diode gate cannot distinguish between signal and mismatch in the arms of the bridge, some matching of these diodes would have to be undertaken at construction time. (It should be noted that in [20] **no** effort is made to check for either systematic or signal related error, which a mismatch would introduce. Only "reliable" (repeatable) performance is claimed. Although the error introduced would be largely systematic, and thus removable in a **digitising** instrument, the use of totally unselected units would invite signal related error problems.)

Unfortunately, these diodes have a zero bias junction capacitance alone of up to 2.0 pF. Even with the largest practical reverse bias, this figure cannot be reduced by a factor of two. (The back bias limit is set not by the breakdown voltage of this particular diode, but by what values it is practical to overcome with Step Recovery Diode pulse shapers.) The combined capacitance of two such devices (the minimum and preferred number) could not possibly be subtracted from a microstrip or CPW line. This can be easily shown by observing that no practical line has this value of distributed capacitance (to be subtracted) in 1/4 of a wavelength at 2 GHz.

Reduced junction capacitance can be obtained at the expense of reverse breakdown voltage. Although intended for signal mixing applications, Hewlett-Packard type HSCH-5311 is a more suitable diode. It offers a zero bias capacitance of below 0.2 pF, with a maximum reverse breakdown voltage of 4 Volts. It is available in batch matched sets for $17, or unmatched for $13, approximately. Although not manufacturer specified, a transit time of 100 ps has been typically found.[50] This implies that it will achieve similar switching performance to the 2837, even without taking the reduced capacitance into account. The total capacitance added by two such devices can feasibly be substantially masked by distributed capacitance removal from a printed transmission line.

Appendix B develops and analyses a best case example of the line discontinuity which can be expected from a gate using these devices, and discusses the effect which is to be expected upon a signal passing a cascade of such gates. The results will be assessed at the appropriate point below.

2.1.4.c Drive Synchronisation

The last special requirement for cascading is the synchronisation of sampling action. There are two methods by which this problem has already been approached. The first is to develop a single pulse of large amplitude. This is divided down sufficiently to drive all the samplers in a group without becoming too small to satisfy the requirements of overcoming the back bias and forward drop of the sampling diodes. This is primarily used in the systems described to date.[20,21,22] The relative timing of samplers is governed by cable introduced delays, fixed at the time of construction.

The second method involves the use of more than one pulse generator. An initial generator has a delay produced electronically. Subsequent generators are electronically delayed also, but their delays are available for trimming. The relative timings are set either with the aid of a suitable calibrating instrument (which may be an equivalent time unit) or a calibrating set of waveforms. Each pulse generator may drive several gates, set up as for the first method, with the groups requiring electronically defined relative timing. One instrument described to date uses this technique.[21]

One further method is available. This approach arises from the multistep technique used to develop the pulse in all sampling designs reported to date. The possibility exists of driving many step recovery diode stages from fewer avalanche drivers. While initially appealing as a further option, this can be eliminated from the possibilities. The main criticism is that avalanche drivers do not develop sufficiently more voltage output than is normally handled by SRD circuits.[25-28,38] In addition the avalanching transistor typically costs much less than either an SRD or other components in a sampler, so the technique does not offer much saving in cost.[27,38]

The single pulse generator approach has been demonstrated with as many as 40 gates.[20] However, this number would seem to tax the conventional pulse generator technology to the limit. Indeed, a total of less than forty gates has forced other workers to the second and more troublesome approach.[21] A requirement of 100 or more gates makes the approach untenable with conventional pulse sources. The available fast SRDs cannot process sufficiently high voltages or power levels with adequate speed.

The possibility of using a different pulse generator exists. A "mercury line" pulse generator can deliver very high voltages with sub-nanosecond risetimes. These devices consist primarily of a mercury wetted-contact switch built into a length of high voltage transmission line.[29] The switch closes without bounce, and the two sides of the circuit are instantly connected. One side of the circuit is a length of line charged to twice the required level. The second half is more of the same line which leads to the load. As one line discharges into the other a pulse with good risetime is transmitted to the load. This would solve the pulse generator problem with the original technique.

The problem associated with this method is that a large number of high quality transmission line dividers and cables would be required. This remains a possibility for the gate driver circuit, but is unwieldy, requires a large number of power dividers and requires that the pulse be sent through both the dividers and a considerable length of transmission line without significant degradation. In addition it introduces potentials of around a kilovolt, which demands bulky expensive transmission line cables.

The remaining alternative is to employ "distributed" pulse generators. The cost involved in this is not excessive, being substantially the cost of the SRD in each generator. A trigger pulse must be conveyed to each generator, but the line need not be of high quality as degradation will not affect the sampler operation directly. One advantage and two disadvantages are offered by this method.

It is elementary to provide a variable delay of the order of up to 1 ns by adjustment of the SRD idle bias.[28] This would allow some form of convenient adjustment to the relative timing of each sampler or each group of samplers on one board section. The converse problem is that there would be some drift associated with the relative generator timings. It is not reported in [21] or [22] whether drift was a problem with either the single or multiple generator systems in existence, so it is not possible for the author to assess whether this is an improvement or a liability. Calibration of the timings of the several generators would require a set of calibration waveforms, but could be controlled by the instrument software provided that the delay adjustments were programmable.

The second disadvantage centres around the placement of SRD pulse circuits on the same boards carrying the signal transmission line. The diode switching time is adversely affected by only small capacitances occurring across the SRD or to ground. In conventional designs[6,8] the pulse shaper diodes are separately mounted, using either low capacitance posts or special high quality transmission line assemblies. Both measurements and simple SPICE modelling suggests that sub-picofarad strays can easily compromise the risetime. This would necessitate some design care, possibly involving a special placement of the SRD circuits.

2.1.5. Summary

A summary of the attributes of the multiple sampling gate approach to transient digitisation is now presented. This will be compared to the alternative at the end of the chapter.

There is no doubt from previous experience reported in the literature and from measurements made by the author that the overall approach is feasible to some degree.

The major drawback is that the full potential of sampler bandwidth and accuracy will not be realised in a cascaded, single event design. It is shown in Appendix B that the elimination of line discontinuity is impossible within the constraints imposed, and that this will corrupt the waveshape to a significant extent.

The cost of MSGs will be large whatever method is used. Not only the fast components involved in the actual sampling must be used in quantity, but also the support amplifiers, and ADCs. If separate generators and programmable delays are used also, the cost increases further. The budget is stretched to the very limit.

A number of problems such as synchronisation calibration, stray capacitance elimination, and compensation for systematic error different in each sample, must be overcome. These problems are not expected to prove impossible, but must be taken into consideration as requiring further research.

The 'timebase' of the instrument will be fixed by the printed transmission line intersampler spacing. This will not easily be varied.

2.2. Scan Conversion Approach

2.2.1. Overview

As described earlier, a scan conversion instrument is based around a CRT whose target presents to the beam a small sensitive face, but makes available to the user a larger, amplified image. Much of the instrument is similar to the basic CRO with a travelling wave tube. However, it contains in addition whatever functions are necessary to process the trace and present it for convenient human viewing.

This section of Chapter 2 discusses the advantages and disadvantages of transient digitisation by means of scan conversion, (SC). Available scan converters are then listed, and the anticipated future design improvements are mentioned. The feasibility of designing an instrument satisfying the goals of this work is noted.

2.2.2. The Advantages of SC Techniques

There are three general advantages to Scan Conversion Transient Digitisation (SC TD), and two disadvantages.

The principle of scan conversion very effectively separates the high frequency signal handling sections from the low frequency ones. Because the SC Tube (SCT) accepts the high frequency signal and performs the temporary storage function, subsequent sections of any system using it need not be "aware" of the fast data rate. The wide bandwidth deflection structures which are the heart of the high frequency end of the SCT do not represent radical technology. They are already available and have demonstrated bandwidths of 2.5 GHz and above.[2,4,15,16,18,19] It is expected that these deflection systems will be improved in the future. It is thus to be expected that SCTs and instruments using them will become available with higher bandwidths,[14,15] as there is no concern that the system beyond the deflection mechanism might need to be enhanced to cope with greater speed. The availability of the required bandwidths and a promise of improvement form a strong advantage for any system designer.

The sweep rate, and thus the effective sampling rate, is as easily varied with an SC system as it is with a CRO. In addition, since the record in temporary storage is two-dimensional, greatly reducing the sweep rate results in the capture of the signal envelope, rather than a single valued data stream with aliasing. (It is true that cunning placement of sampling gates on discrete lines can allow a factor of two change in effective sample rate at the cost of reduced sensitivity and some extra hardware.[20] However, this technique is very limited in the extent to which it can alter sample rate, and can never capture envelope information.)

Target vulnerability to burning by an intense beam may place a lower limit on sweep speed in an SCT.[17,57] Despite this restriction, an SCT based instrument will typically offer from three to six or more orders of magnitude variation in effective sampling rate, while the alternative may offer one octave or two.

This fundamental flexibility in rate and information format constitutes a significant advantage, especially where the analysis of modulated signals is involved, because the envelope function of a signal contains information about the modulation function.

The targets of tubes offered to date have been channel-plate-and-phosphor or semiconductor based designs. These require considerable post-processing. Future designs will have charge coupled device or fibre optic output.[14,15] These may offer increased sensitivity, resolution or ease of processing. The vesting of the bandwidth reduction process in a single component with separate halves offers a kind of modularity to any system employing an SCT. Should any new techniques become available, much of the design base used in previous instruments (particularly the wide bandwidth part) would not be altered even in the face of a total redesign of the low frequency end. Of course, this 'modularity' operates in reverse, permitting update of the high frequency section while retaining the post-processing equipment. This is an advantage not foreseeable for MSG systems.

2.2.3. The Disadvantages of SC Techniques

An SCT must be fabricated by a large, well equipped laboratory. In the case of the semiconductor target tube, both an advanced semiconductor fabrication line and a CRT assembly facility are required. None of the component parts of either the beamforming, deflection or target systems are available from normal equipment sources. (This is partially a consequence of the tight proprietary nature of the field of instrumentation engineering.[58]) For this reason it is not practical to contemplate design of an SCT for a single application, and at best one must use what is offered in the marketplace. There are two manufacturers: Tektronix in the USA and Phillips in Europe. In practice this is not so severe limitation. In addition, the increase in interest in instruments with gigahertz bandwidths and transient capability is likely to ensure the availability of progressively more advanced SCTs in the future.[3,15,18]

The engineering disadvantage of the approach lies in the relative insensitivity of travelling wave deflection tubes. Several volts to tens of volts of drive (into transmission line impedances) is typically required to produce full deflection. All the systems reported to date either do not utilise the full bandwidth capability of the tube or apply the signal directly to the plates, without active amplification, tolerating the insensitivity.[2,3,4,16,17,19] A discussion of the current performance of various tubes follows.

2.2.4. Assessing SCT Performance

Three fundamental specifications may be taken to characterise a scan converting cathode ray tube to sufficient a degree to permit comparison between units.[4,15] These are bandwidth (BW), drive requirement for full deflection, and resolution. Information in the literature, particularly pertaining to the analog display versions, tends not to be presented in a form which unambiguously conveys each of these. In non-digitising tubes it is often necessary to estimate resolution. Bandwidth is normally expressed straightforwardly, though it is necessary to note that the response curve exhibits zeroes rather than the smoother (Gaussian) roll off associated with an amplifier. Drive requirement must take into account the base impedance, as well as the fact that the signal may need to be push-pull. In order to make a meaningful comparison between units, it is important to ensure that the specifications are expressed in appropriate terms.

The beam writing speed may also need to be taken into account. Where this is sufficient that it does not affect the bandwidth specification, it can be ignored. However, in many tubes it is a limiting factor, and will impose a "full deflection bandwidth" limit, much as the slew rate of an amplifier may introduce a full power bandwidth.[Appendix A]

2.2.5. Current SCT Performance

There are two scan converter tube designs currently available, and one which will be available in the near future. The first design employs a channel plate intensifier and a phosphor screen. There are several units in the series.[15] They have been developed in French laboratories. The series is soon to be expanded to include a digital output version. Although announced as having been demonstrated feasible in 1973[15], these digital output channel plate devices are more recently only being transferred from simulation to realisation stages.[16] The second design is offered by Tektronix in the USA. It has a silicon target, which reads out the stored image in serial analog form in response to a raster scanning of the target by a second electron gun within the tube.[17]

The following table summarises various scan converter tube performance figures, and includes a plain travelling wave deflection, phosphor target tube, believed to be close to the best obtainable without scan conversion assistance[15], for comparison:

Id No./Source	Target Type	BW	Drive/Z0	Resolution*
50D13BE/LEP	Best non-SCT	5G	100V/50	~5 bits
TMC1/LEP	Channel Plate	3.3G	10V/50	~6 bits
TMC2/LEP	Channel Plate	2.5G	1V/50	~5 bits
TMC3/LEP	Channel Plate	1.5G	.25V/50	~5 bits
T79/Tektronix#	Semiconductor	2.5G	8V/180	9/12 bits
-No Id-/LEP	Plate/Vidicon	7G	5V/50	~6(?) bits

*The Tektronix tube resolution is governed by the readout mechanism. Proprietary electronics developed by Lockheed are claimed to develop 3 extra bits.[18,19,68]
Resolution figures are estimated for phosphor output tubes.
#The T79 tube has full part number 154-0698-00.

The TMC2 and the T79 tubes offer optimum performance for an instrument requiring 2 GHz bandwidth. The Tektronix tube is more readily available, and has the Tektronix support electronics also available. The output is digital, which conveys several advantages apart from the simplicity of storage and display of the final result. These include the possibility of precise signal averaging when multiple occurrences are available, and ease of post processing for both linearisation or response compensation.

The TMC2 tube offers better sensitivity. It is also likely that it may be available in the future with digital output in some form.

Both of these tubes will be discussed in more detail in the following sections.

(It is the speculative opinion of the author that each of the designers of SCTs will eventually incorporate the innovation of the other, patents and commercial rivalry permitting. The Tektronix tube, as will be seen, would benefit greatly from the introduction of a channel plate amplification stage to overcome writing rate limitations and flooding or burn-out limitations. The French units will eventually be forced to place their digitising system within the tube, rather than relying on processing light energy derived from a phosphor, because this is likely to continue to limit resolution.)

2.2.6. The Tektronix SCT

The internal layout of the T79 tube is shown in figure 2.6. The salient parts are threefold. Firstly, the travelling wave electron gun, focusing and deflection system are found at one end of the assembly. Secondly, a target presenting a 2 cm square sensitive face to this gun is to be seen in the tube centre. Thirdly, a second gun is situated at the opposite end of the tube, facing the 'back' of the target. (In conventional microcircuit terms it is the front which is presented to the second gun.)

INTERNAL LAYOUT OF TEKTRONIX T79 SCAN CONVERSION TUBE

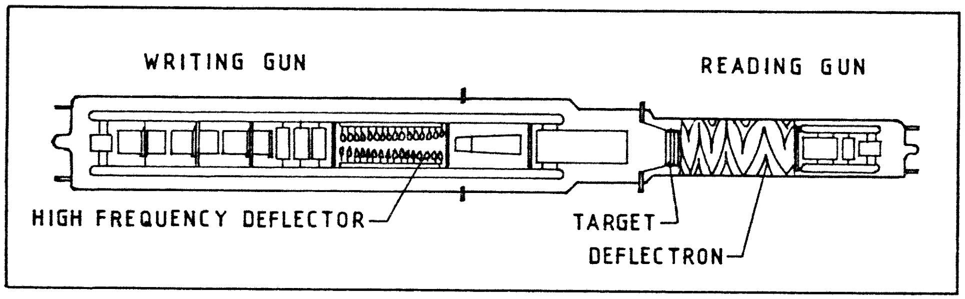

FIGURE. 2.6

The target is a semiconductor plate, upon which diodes have been fabricated into a square array, 512 diodes on a side. The principle of operation is relatively straightforward. The second gun scans the array, raster fashion, reverse biassing the diodes. The diodes retain their charge across the depletion layer. When a sweep is required, the TW gun is turned on, and its focussed spot traverses the target. Where the high energy electrons hit, they penetrate to the depletion region and allow some charge to leak across a diode. Thus immediately after a sweep of the fast gun, the trace trajectory is recorded in the charge distribution across the array. This may be "read out" by having the second gun scan the target in a known raster pattern, and observing how much current (the number of electrons, or charge) was required to recharge each diode in the array. Various processing algorithms may then be employed to recover either an image or a set of numbers for further processing.

2.2.7. The TMC2 SCT

The internal layout of the TMC2 tube is shown in figure 2.7. Unlike the tektronix tube, output is by image directly projected onto phosphor, rather than as coded electric signals. The salient components are again threefold. Firstly there is a TW gun assembly, as in all SCTs. The second component, the sensitive target, is a microchannel plate. The fabrication and operation of this device is a fascinating discussion in itself for which the reader is referred to reference 14. It may be thought of as a bundle of microscopic tubes, whose packing is

INTERNAL LAYOUT OF PHILIPS TMC2 SCAN CONVERSION TUBE

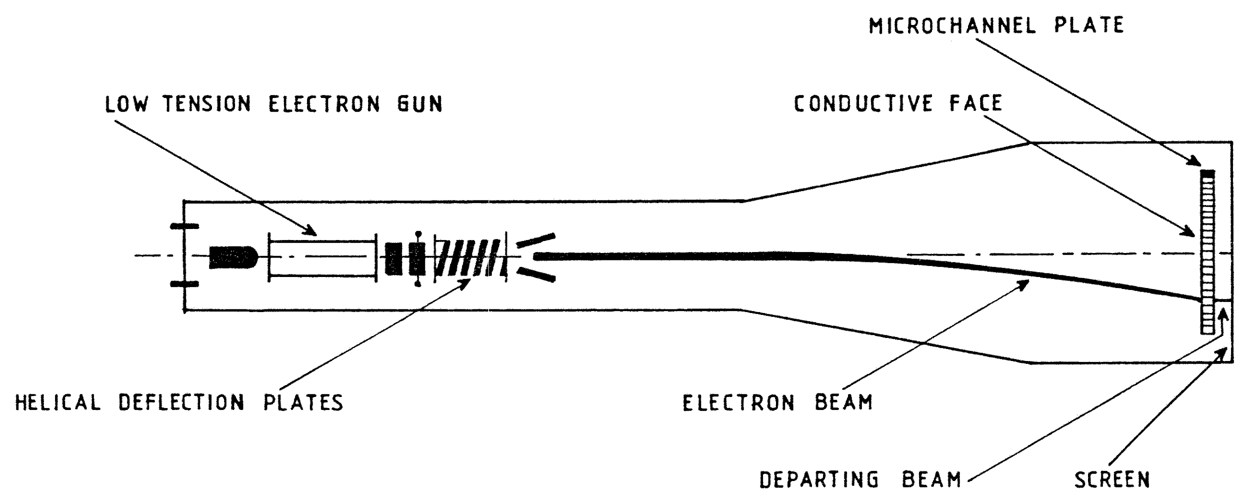

FIGURE. 2.7

very comparable to that of the diodes in the previous tube discussed. Each of these tubes is an electron multiplier, such that when an electron enters the input end, many are ejected from the output end. Thus the beam may be formed of very few electrons of relatively low energy, and this beam will be amplified and made to consist of many electrons when it emerges from the plate. The third (readout) component is merely a phosphor screen which responds to the electrons, directly converting their energy to visible light. The electrons may be deflected outward if the screen is required to be larger than the plate, by suitable electrostatic lenses.

The output is of course not directly available for further processing in this design. However, the minute electron multiplier tubes supply energy and can thus be made very sensitive. They can also be arranged to have a well defined saturation current. This confers two advantages. The sensitivity is such that even a single moderately energetic electron entering a tube can excite a visible flash. Thus the beam current can be quite low, for a given writing rate, and the beam composed of few electrons. Because the target can respond to few relatively un-energetic electrons, the deflection requirements are reduced. It is for this reason that the deflection requirements are much less stringent in the tubes utilising channel plates. The saturation current limitation has the effect of preventing blooming of the spot or excessive phosphor currents which can occur in other designs when the beam moves slowly and an excess of electrons is present at some points in the trace. Blooming is not only potentially harmful to the target, but may reduce resolution.

2.3. The 7912 Series Scan Conversion Digitisers

Tektronix has released a number of 7000 series compatible mainframes based on its tube. The currently available one is identified as the 7912AD.[17,67,57] This section will initially discuss the performance of this instrument. One significantly different instrument has been produced. It is available from the Lockheed company. This instrument will be described also.

2.3.1. The Tektronix 7912AD

2.3.1.a Hardware

An overall block diagram of the 7912AD is shown in Figure 2.8. This figure is simplified even beyond that given in reference 17. This is to suit this description, which will touch only on points important to this work, and will ignore some features of internal organisation.

The instrument may be regarded as having five sections: vertical deflection, horizontal synchronisation and deflection, readout, power supply and microprocessor control. Both the power supply system and the horizontal synchronisation and deflection equipment are substantially straightforward. The power supply is extensive, having to supply many subsystems with a variety of

SIMPLIFIED BLOCK DIAGRAM OF THE 7912AD

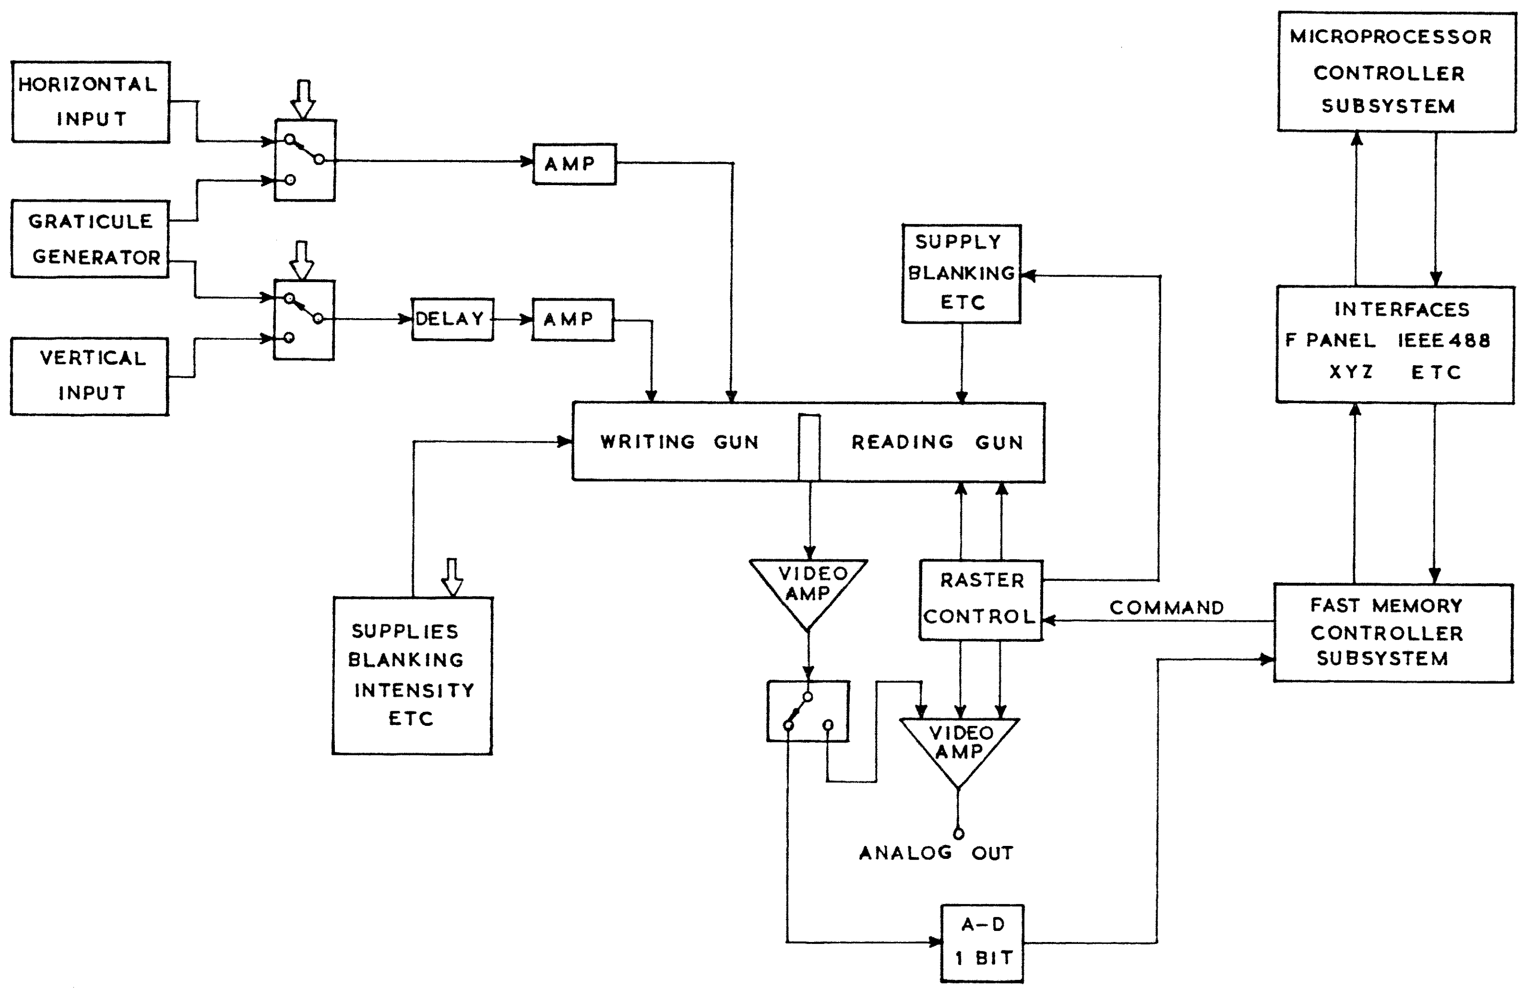

FIGURE. 2.8

voltages, but it is unremarkable in an engineering sense. The horizontal functions are provided via a standard Tektronix 7000 timebase plugin interface. The only additional component found in the 7912AD which is absent from other 7000 mainframes is an electronic switch in the plugin interface to allow the graticule generator to inject its calibrated signal. As this does not degrade its performance, it leaves the horizontal system in the same unremarkable position.

The vertical system will be discussed further in Chapter 3. It is standard 7000 series compatible, and is thus nominally limited to a bandwidth of 500 MHz.[67] (The instrument is supplied with a 250 MHz plugin, as this is the highest bandwidth available with programmability.) It is not recommended as convertible for use with the passive 1 GHz plugin (now discontinued), which requires for its installation the removal of all other vertical system components. The major nature of such a conversion coupled with the facts that it is not recommended, and can deliver sensitivities of only around 4 volts per division, have prevented experimental verification of the 1 GHz bandwidth. However, analysis of the design provides no reason to suspect that the conversion would not deliver some form of usable signal up to 1 GHz. Figure 2.9 shows the measured response of the 7912AD in various configurations. As can be seen 600 MHz is the greatest bandwidth obtained, and additional poles are evidently acting below 1 GHz. An OM335 amplifier response is included for comparison. The OM335 from Philips represents low cost technology of an older vintage than the 7912 series.

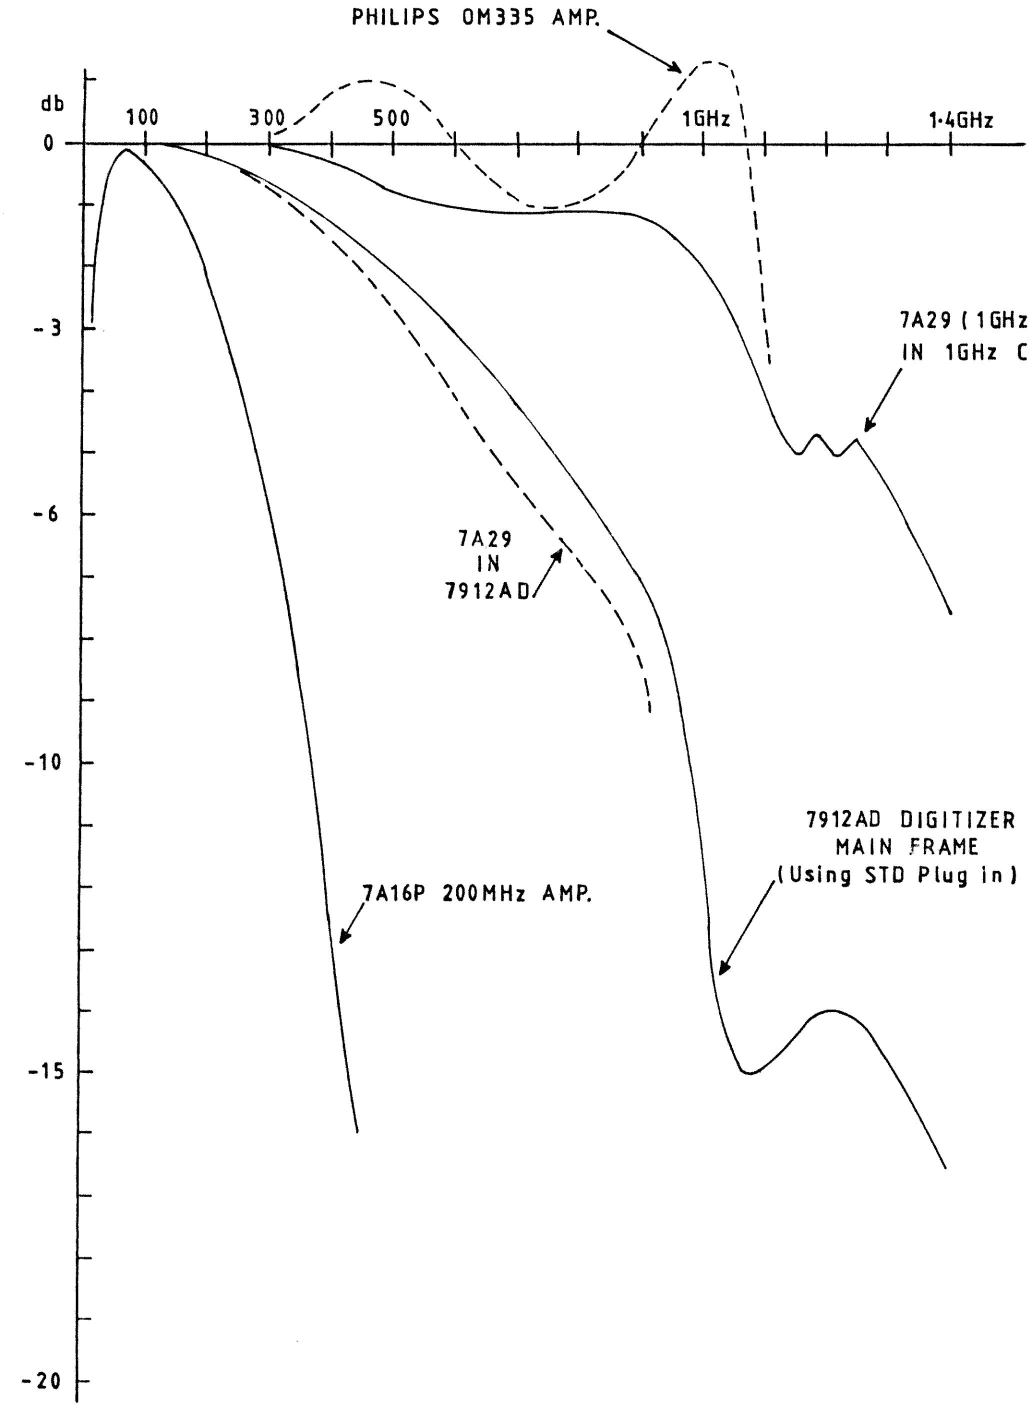
PHILIPS OM335 AMP.
db
100
300
500
1GHz
1·4GHz
0
-3
-6
-10
-15
-20
7A29 (1GHz AMP)
IN 1GHz CRO
7A29
IN
7912AD
7912AD DIGITIZER
MAIN FRAME
(Using STD Plug in)
7A16P 200MHz AMP.

FIGURE. 2.9

The offered bandwidths are seemingly well below that of which the tube is capable. The limitations are introduced in the vertical processing chain before the tube is reached. Figure 2.10, reproduced from [19], indicates the measured bandwidth variation with beam energy (accelerating voltage). Tektronix uses the tube at around 10 kV, sacrificing some bandwidth for the higher beam current and ultimate sensitivity. These are necessarily beneficial to the operation of their cruder image digitisation process and amplifiers, respectively.

As was indicated above, reading is affected by scanning the target with a beam operating in a raster fashion. As the reading beam passes a diode, any current flowing into the target plate is sensed. The current represents charge being stored in the depletion region of the diode. Where a diode remains charged there is little current flow, but where it is recharged by the beam, some current flows, indicating that the diode was 'hit' by energetic electrons from the writing gun. A sensitive video amplifier is used to amplify the current flow signal. This signal follows two paths.

The first path is analog. The raster signals are combined with the amplified video signal to form a standard composite video signal. This is fed to an external connector, to which a monitor may be connected. With the electronics set for continuous scanning, the monitor then displays a continuously updated image of the target. In essence, the instrument is acting like a

T7910 CHARACTERISTICS

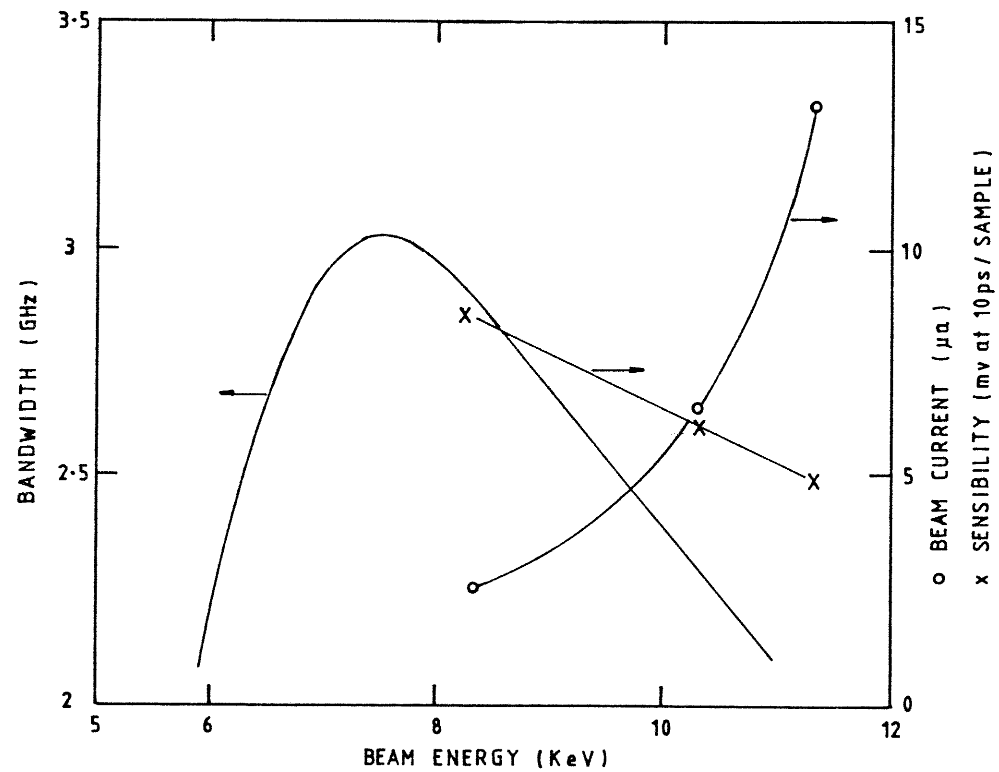

FIGURE. 2.10

conventional non-storage CRO. This mode is provided on the 7912 primarily for setting up by the operator. The setup procedure is greatly accelerated when the operator has a fast view of the effect of any adjustment. This is not possible in the mode where results are digitised and possibly further processed.

The second path is digital. Figure 2.8 shows the video signal fed to an Analog-to-Digital converter (A/D converter, or ADC). In the 7912AD this converter is in fact a 1 bit converter or simply a comparator. The ADC output is fed initially to a memory controller, which consists of a 2900 processor and fast cache memory, then finally to the 6800 system microprocessor.

Although very succinctly coded, the waveform memory effectively contains a bit map of the $(512)^2$ target positions, each one coded as having been hit, or not hit, by the writing beam. This information is available from the processor over the control bus, and is termed a 'RAW' waveform dump. This information may be viewed as encoding an array of pixels in an image of an oscilloscope screen. There is of course no grey scale - only on and off, or black and white.

In practice, what is mostly required from a digitiser is either a scalar array forming a single valued time function, or possibly a pair of such functions describing the envelope of the digitised waveform. Thus there arises the question of what to return from the image array when either of these functions is needed. In addition, it will be seen that there is a need to be able to interpret the returned array in the light of its method of derivation.

2.3.1.b Firmware

The complications arise from the fact that the beam path is much more than one pixel wide. At the minimum intensity required to make a positive, reliable impact on the target, the spot diameter is typically 4 pixels. When used at higher intensities, or when the beam moves slowly in one part of the trace, a width of 15 pixels is not uncommon. Where the absolute speed at which the spot traverses the target changes drastically, such as when depicting a step, the width changes mid-sweep. Should the current beam intensity set a maximum writing rate which is momentarily exceeded by the spot, the trace record can be lost entirely. A high repetition rate sampling converter or an array of samplers will inherently return a set of values whose interpretation is straightforward: The scan converter output processor, on the other hand, must handle the above complicating factors, distilling a result from the image data. (These effects will be elaborated upon in Chapter 4.)

The 7912AD contains some firmware to perform this "distillation". It also has firmware to affect signal averaging by point by point averaging of the single valued arrays derived from the raw data. The true resolution is often not that apparently returned by the firmware. No reference could be found to work done toward determining the actual resolution of the tektronix SCT in various situations, either in Tektronix' own documents or in the commonly available literature. The author has observed that the nominal specifications for the 7912AD[57,67] are far from telling the whole story. This is attributed to the digitising SCT technique, the data conversion mechanism and the interpretation methods employed.

The first step in data reduction performed by the firmware is the determination of so-called 'EDGE' array. This is a 512 element, double valued array which represents the upper and lower boundaries of the (hopefully) continuous "river" of pixels set by the writing beam in a sweep. Crude algorithms return empty elements in the array where trace width or the rate of change of width indicate some aberration or error. Empty elements are also returned when the trace is lost. Small islands of set pixels are eliminated as irrelevant. The arrays are, of necessity, 9 bit integers, corresponding to the 512 target matrix dimension. When a single valued array is required, as is usually the case in digital signal processing, the 'ATC' array is derived from EDGE. This is an Averaged-To-Centre version of EDGE, and is obtained by adding the corresponding elements of EDGE.

Signal averaging is accomplished simply by repeatedly obtaining the ATC array as described above, and then summing the arrays and dividing by the total number, again returning a ten bit array similar to a single ATC array. In practice the division is by a factor of two. The number of arrays averaged is forced to be a power of two in order to facilitate this easy integer division.

2.3.1.c Shortcomings

There are a number of shortcomings evident both in the hardware and firmware of this instrument. Firstly, the "1 bit" ADC does not even approach capturing all the information available from the target. This is perceived merely by comparing the binary and analog video displays. The analog trace is seen to have a "grey scale" associated with it, which is lost in the binary case. One upshot of this is that trace position is less certain. This will be exemplified by the description of operation of the Lockheed instrument in the next section. The second effect is the apparent reduction of maximum writing rate. As the speed of target traverse increases, the brilliance is of course reduced. Since it is only sufficient at the fastest sweep speed to be perceptible above the noise of the target on the analog display, it is often lost after the level comparison that serves as an ADC. Software can overcome the problem to some extent. It is nevertheless a notable problem.

Secondly, the EDGE array derivation is insensitive to trace breadth. It returns a value derived without awareness of the contents of horizontally adjacent vertical lines of pixels in the image. It does check trace width for variation, but can only reject all data in areas with gross aberrations. This can produce misleading data. (The discussion of the DIOS command ABP will elaborate on this problem: see Chapter 5.) In practice, the spot width is always at least several pixels. This produces a kind of "forced interpolation". This is beneficial to an **image**, but of course means that there are **not really 512 independent data supplied.** This explains the paradox which arises when one considers that the 7912AD claims to deliver 512 samples from a basically CRO-like image, when we know that the humble analog service CRO offers only about 100!

Thirdly, in deriving the ATC array, the errors of the EDGE process are compounded by the use of crude linear interpolation to fill gaps.

Fourthly, the Signal Averaging (SA) process has been found to fail entirely if either the instrument triggers on a glitch, which is common with very low repetition rate signals such as transient instruments encounter, or the trigger jitter is comparable in displacement to the width of trace features.

Lastly, and importantly, the bandwidth of the system does not realise that of which the basic tube is capable. This observation is expanded in Chapter 3.

2.3.2. The Lockheed R7912M

The Lockheed Palo Alto Research Laboratory has designed, under the guidance of Dr Ray Smith, a modified version of an early Tektronix instrument.[18] The R7912 (or a later version up to the 7912AD) is purchased directly from Tektronix, and modified in house.[19] For convenience, the unit will be referred to as the R7912M, (M for Modified).

2.3.2.a Hardware

The Tektronix original is purchased to provide the tube complete with its support circuitry. The extensive power supply and horizontal sections are satisfactory in their present forms. However, the vertical section, the readout circuits beyond the raster generator, and much of the control section concerned with processing the image is removed and discarded. It is replaced with a new passive vertical system, a new ADC subsystem, and a new data processor. The processing algorithms must of course be altered to cope with the new image coding, but the removal of the standard control hardware obsoletes the Tektronix firmware anyway.

The systems which are installed by Lockheed are proprietary, and documentation is not normally released. It is thus impossible to give precise details of their design here. However, much can be gleaned from a knowledge of the Tektronix instrument, from specifications and facts published in reports, and from information provided by Lockheed.

The vertical system which is installed is a passive system which possesses two salient properties. One, it converts single ended 50 ohm signals to push-pull 180 ohm ones. Two, it applies a boosting characteristic to the response above 2.5-3 GHz, extending the response to 3.5 GHz. (The accelerating voltage is lowered to extend bandwidth, because further modifications relax beam current requirements.) This requires a sacrifice in sensitivity, and must also compromise the smoothness of the roll-off above 3.5 GHz. It is not hard to accept that the transformer-filter mechanism required to achieve these things is well within the abilities of recent microwave technology.

As indicated above, the "1-bit ADC" of the 7912AD effectively ignored the intensity distribution information held in the analog signal obtained by the readout mechanism. The R7912M effectively replaces it with an ADC with much greater resolution.[68] (The number of bits is not revealed.) This allows the true trace position to be determined with greater accuracy. This determination is made by the software accompanying the unit and replacing the 7912AD firmware.

The memory controller must now handle many more bits of information. In addition, the coding scheme and trace processing algorithms must be different. These considerations demand the use of different control hardware.

2.3.2.b Software

Lockheed claim that 12 vertical bits of resolution can be obtained from the hardware. It is not known whether this definition of resolution includes noise, assumes averaging or relates to simply written traces. The trace determining algorithms are also not disclosed. It is not hard, however, to imagine processes similar to the Tektronix "standard" algorithms.

(It is only known with certainty that the ADC subsystem obtains the specifications of resolution quoted, by the means above, from reference [68]. The author guessed the mechanism in general outline, and was fortunate to have the suspicions confirmed during a number of exchanges with the chief design scientist of Lockheed MRC. The software for processing remains proprietary.)

It is known[68] that the trace centre determination involves some form of interpolation using the magnitudes of individual pixel samples, whose magnitude is known to more than one bit. Whether the interpolation is confined, as are the Tektronix originals, to vertical slices, or whether it is a more intelligent mechanism, is open to conjecture. The process may well mimic the algorithms the author has developed in the "1-bit" case. They could be quite complex, approaching the problem in a true image processing sense; true trace centre may be determined from the three dimensional map. This process is imaginable as determining the course of a set of mountain peaks on a contour map. The great resolution thus obtained compensates for the relative insensitivity obtained with the passive vertical stage: This is some 8 volts p-p for full deflection.

2.3.3. Further 7912 Instruments

The viability of installing new vertical signal processing circuits in existent 7912 instruments is proven by the Lockheed work. The possibility for software improvement of accuracy, resolution and functionalism exists. The fundamental bandwidth requirement from the Tektronix tube is available. There would seem to be a strong case for using a Tektronix instrument, much as Lockheed has done, to provide support circuits with little effort. These factors indicate the feasibility of achieving the goals of the work undertaken by designing a new 7912-based instrument.

2.4. Conclusions

Many factors must influence a decision to proceed with either an MSG or an SC design approach. Some of these are: The probable cost; the certainty of success; the bandwidth achievable for a given effort; the ultimate bandwidth of the technique, and the ease of subsequent improvement; the potential sensitivity; the likely resolution; the instrument's flexibility; the ease of operation of the finished product; the likelihood of future support; the effort required to bring the instrument to a useable level of operation.

Both approaches have predecessors whose work indicates that the goals set up here are within reach of the techniques. However, designers of MSG instruments express concern with the difficulties, while SC instrument sources express optimism in the future.

Both approaches have demonstrated the required bandwidths. In an SCT the bandwidth is related to deflection sensitivity and target sensitivity. In an MSG array the bandwidth is directly traded for sample number. The low numbers of samples taken in existent MSG designs prevents prediction of whether the competition between sample number and bandwidth will prove to be a limiting factor. The predictive modelling reported in Appendix B is not optimistic, indicating that a 2 GHz instrument will be operating near the limit. In addition, the bandwidth is not constant along a cascade, and there would be mismatch problems, which would greatly complicate software compensation of response.

Any improvement in the fundamental restrictions currently imposed could likely be installed by replacement of the tube in an SC device. However, an improvement at the component level in an MSG TD would prove difficult even if possible, and would require skilled, delicate work on each sampler.

Resolution and sensitivity of sampling gates used without null seeking or signal averaging is much lower than is seen in today's sampling instruments. Further, Appendix B indicates that at best the resolution will be seriously degraded by mismatch echoes. No work reported offers an analysis of either effect. The former is dependant upon the in-house matching of semiconductors not even intended for the job for which they would be used. Even if single laboratory sampler prototype performance could be achieved on a large scale, no information about degradation of match with time is available. Impressive claims made about the sensitivity of the SCT from Tektronix are available. These have not been refuted, though the resolution is probably compromised under certain signal conditions.

The fixed timebase and one dimensional store of an MSG design has reduced flexibility compared to the SCT approach. In addition, the basic "CRO-like" operation of an SCT digitiser and the "TV" mode for setting up make its use easy to those familiar with oscilloscopes. (This may seem a small point when discussing a specialised state-of-the-art instrument in the $50,000 plus category, but it has been shown time and again that the "user interface" affects an instrument's acceptance.[69])

Support for an SCT instrument is vital, because of the major facilities and engineering design base required for tube fabrication. Sampling gate components are not so critical. However, the support for SCTs is available, and enthusiasm in the technical press suggests that it will remain or improve. On the other hand, no recent MSG TD work is reported, and although the hardware proposed is not likely to become unavailable, advances are not anticipated.

Finally, the indirect costs of designing from the ground up, as would be the case for a sampling gate unit, would be very large in comparison to the conversion of a Tektronix instrument. The latter alternative provides support circuitry already.

The scan conversion approach appears to offer the most promise, now and in the future.

3. THE 7912ADM DIGITISER

This chapter contains three logical steps. It describes the Tektronix 7912AD vertical signal chain in more detail, focusing on its bandwidth potential. It proceeds to a discussion of the basic alternatives for enhancement. Finally it describes the hardware which has been installed in order to realise a machine towards the goals set up, and to enable the investigations carried out and described in this thesis.

3.1. The Tektronix 7912AD Hardware Configuration

3.1.1. The Vertical Amplifier Chain

As indicated in the last chapter, the 7912AD, as supplied by Tektronix, includes extensive limiting circuitry in the high frequency signal path. This prevents realisation of the full potential of the SCT. In the recommended configuration, supplied as standard by Tektronix, the bandwidth is specified as 250 MHz. Sacrificing programmability, the widest bandwidth offered by Tektronix is 1 GHz. This requires the removal of the graticule display facility, (the calibration mechanism) as well as the delay line. This option also involves removing all the active vertical components and replacing them with passive networks which fix the sensitivity at 4 volts per division, or 32 volts peak-to-peak for full scale deflection, as noted before. Retaining the vertical amplifiers, the fastest plug-in (the 7A29) offers a 500 MHz sinewave bandwidth, "usable" to 700MHz. The following table lists the critical components through which the signal passes in each case, and includes comments on the effect of each step as measured on one particular unit. Parentheses indicate that the component is absent in the passive (1 GHz) system option.

Component	Comment
Plug-in amplifier	Bandwidth specification 250 MHz or 1 GHz.
Plug-in connector	Introduces a serious mismatch. Bandwidth disturbed below 1 GHz.
(Graticule switcher)	Bandwidth limit below 1 GHz.
Delay lines	Response falls off as frequency$^{1/2}$, 1 dB at few hundred Megahertz.
(Plate amplifier)	Bandwidth limit of below 600 MHz.

It is clear that all these elements supplied in the vertical signal processing system are, to varying degrees, unacceptable in an instrument which is to achieve 2 GHz bandwidth.

3.1.2. The SCT

The Scan Converter Tube has a basic capability exceeding 2 GHz as noted in section 2.2. It has a drive requirement of 16 volts p-p applied out of phase to each plate. The plate presents a 180 ohm characteristic impedance. For correct focusing, a dc potential of 40 volts should be present on each plate. The tube has an externally attached termination for the travelling wave deflection plates which permits the dc potential of the termination to be set. This potentially permits operation down to dc. It also simplifies the bias of the driving stage.

3.2. 7912AD Redesign Proposals

3.2.1. The Lockheed Upgrade Philosophy

Lockheed engineers addressed the problem of obtaining optimum performance from this SCT. They accepted the deflection insensitivity of driving the tube passively, in order to obtain wide bandwidth without having to deal with active microwave components and circuitry. They then recovered sensitivity by improving their ability to determine the true beam centre of the trace. This required a major design effort which replaced the whole control system, and which, though proprietary in nature, is admitted to involve greatly improved resolution of the charge quantities on each target pixel.[19] Some criticisms may be directed against that approach, were it to be used at this time.

When the Lockheed design work was done, the instrument then offered by Tektronix had control hardware and firmware which was crude in comparison with the today's 7912AD. With its full remote instrument control abilities and inbuilt signal processing firmware the 7912AD represents a considerable advance. This means that the Lockheed modifications applied to the present 7912 would involve the loss of considerably more capacity than was originally the case.

Any engineer who has had to build a 12 bit ADC will be aware that the 72 dB dynamic range thus implied is very wide in engineering terms, and demands consideration of the problem of noise. This is especially true when dealing with gigahertz bandwidths. The 7912AD trace suffers noise disturbances in the least significant bit, even delivering only its nominal resolution - 9 bit integers. (The averaged-to-centre trace is represented by 10 bit integers. Each datum is derived from two 9 bit integers by plain addition, and as such should offer 10 bit resolution. Nevertheless, only 9 bits is claimed by the makers, and no mention of this consideration occurs in the manufacturer's literature.)

That there are noise sources internal to the 7912AD yet separate from quantisation noise can be inferred: Taking a Fourier transform of a typical narrow band signal acquired by a 7912AD, and ensuring that leakage effects are controlled, reveals a 'noise floor' in addition to the narrow part of the spectrum resulting from the signal. This noise floor contains components outside the bandwidth of the incoming signal, and the tube's theoretical deflection capability. These signals are still well below the Nyquist limit at the higher sweep speeds. If an ideal narrow band signal is ideally quantised and processed in the same fashion, the expected quantisation noise floor[73,74,101] is observed. Yet the floor developed mathematically is perceptibly below that obtained in practice. The conclusion must be that noise is added to the signal as it passes through the 7912AD.

Several mechanisms could account for these noise signals. A possible source is "beam inhomogeneities", introduced by random motion of electrons in the beam or gas atoms in the evacuated tube or signals on the accelerator supplies. Stray electrons, accelerated by tube potentials, often hit the target. These can alter the charge on a pixel below the discrimination level set by the readout comparator and alter the placement of beam edge. The discrimination level corresponds to different charges in different pixels, because of different sensitivity of pixels, induced by age and wear, etc. (This will be demonstrated in Chapter 4, in discussion of accuracy of analog readout.) More likely, noise in the sensitive video amplifiers of the reading system would result in apparent writing beam noise. (These amplifiers must operate in close proximity to 10 kV regulators, digital circuits and raster generators, and yet detect minute charge flows from the target.)

When the standard instrument has noise processes evident at the 9^{th} bit level, considerable difficulties would have to be overcome to obtain the 12 bit resolution claimed. It seems likely that the supply regulator and the video amplifier would have to be upgraded, in addition to the further circuitry required in the more complex approach.

It should also be noted that improvement of the resolution is not an equivalent substitute for improvement of sensitivity. In favour of the replacement of the entire ADC subsystem is the fact that the resolution of level is improved when dealing with signals approaching maximum deflection capacity. On the other hand, as peak level falls, so does resolution. Even in the absence of noise and with image magnification, the 12 bit approach is equivalent to the 9 bit approach when the input is only 8 times below maximum. In the Lockheed case, this occurs for an input of 1 volt p-p. Below this level, where the gain of a system with amplification would be increased and no resolution lost, the Lockheed approach loses ground. Further, consider that amplifiers capable of handling the full dynamic range of the instrument may need to be introduced into the signal path, as is not uncommon. They would require very good performance indeed in order not to compromise the resolution. The necessary dynamic range of 72 dB in a single wideband amplifier is a significant demand at gigahertz frequencies.

3.2.2. A New Modification Proposal

There is basically no **necessity** for the limitations of the components used in the Tektronix-engineered instrument. They are included in order to preserve the 7000-series compatibility for which the company is well known, and presumably in order to draw on the design base which already existed within the company. Likewise, the Lockheed approach sought to avoid microwave design, at the expense of increased complexity in the LF processing section. It must offer reduced flexibility because of the fixed deflection sensitivity, and the resolution may be reduced in portions of fast moving waveforms.

Many of the parts required to construct a vertical system with a 2 GHz bandwidth are available "off-the-shelf". This availability is not recent.[49,59,60] There is a strong base of design information (proprietary and public) for systems covering very wide bandwidths and stretching well into the microwave region.[60,61-65,76-78] With the approach of commercial Microwave Integrated Circuit (MIC) components using Gallium Arsenide (GaAs), the limits are expected to be improved further.[70,71] Already DC and 0.1 GHz to 12 GHz amplifiers are reported in the literature.[72,79] An active amplification system with 2 GHz bandwidth suitable for use with the 7912 series instruments is quite feasible with currently available Silicon bipolar technology.

The remainder of this thesis reports the implementation of an instrument, based on the existent Tektronix (Tek) mainframe, but with a new modification philosophy inspired by the ongoing improvement in the field of wideband amplifiers, the expectation of improvement in SCT technology, and the requirements for radar work outlined earlier. The vertical channel has been entirely replaced with a wideband amplifier cascade, using commercial components wherever available. The instrument provides a digitised output which is pre-processed by a set of algorithms which allow for some of the limitations of the overall response and provide user convenience. Optional post-processing algorithms are also available. The software implementing these algorithms (as different from the firmware ones provided in the 7912) reside in an external controller.

A number of considerations have influenced the decision to proceed in this way.

The modification of an existing (and available) instrument, as opposed to design from scratch, is appealing for reasons of expediency. The advantages of producing a machine tailored to an application are outweighed by the resource costs of largely duplicating work already completed to the degree of commercialisation by someone else. This decision is implicitly supported by the Lockheed philosophy.

The use of active amplifiers (with an SCT) will deliver good sensitivity. The potential sensitivity is in excess of that available by any other means, including discrete samplers. Currently achievable amplifier bandwidths are sufficient for the immediate application. In addition, the limits may be expected to improve in the future, as evidenced by reports cited above.

Vesting as much instrument function as possible in the software provides advantages. The ability to accept a degree of imperfection in the amplifier and tube response, and to adapt quickly to any change in the performance of the signal chain, will offer two potential advantages. It will make the instrument flexible in use, allowing removal or addition of components at short notice to accommodate various situations. It will also permit the system to update an amplifier and/or the tube without serious redesign. (Recall that both amplifier performance and SCT technology are improving.) The approach will thus allow the instrument to remain viable while accepting future hardware.

Inbuilt deconvolution algorithms, available in the post-processing software within the instrument, are a form of "self-calibration". They represent one of a number of techniques currently being investigated, some of which are already implemented in commercial equipment.[80] Such mechanisms allow rapid adaption to changes in the system, which may be accidental (random) or specifically introduced for particular applications.

3.3. 7912ADM Hardware Configuration

3.3.1. Components

As noted, commercial amplifiers are available with bandwidths of 2 GHz. The Watkins Johnson (WJ) company offers a range of hybrids with 0.01 to 2 GHz bandwidth. These have been used wherever possible in the design. The circuit discussed below utilises five WJ amplifiers, and thence requires only two full bandwidth circuits which are not available commercially. These have been built to meet the unusual and demanding functions required for the instrument.

The relatively modest requirement (in microwave terms) of a 2 GHz maximum frequency means that lumped components are practicable.

Both resistors and capacitors are available in surface mounting chip packages. Experience has shown that such packages have minor parasitic impedances even at 2 GHz.

Active components have been used in the modules constructed in house. These components are Hewlett Packard Bipolar Junction Transistors (BJTs). Full available component specifications for parts used are given in Appendix D.

3.3.2. Schematic and Signal Budget

Figure 3.1 shows the block diagram of the amplifier chain, with typical signal levels included. Standard SMA connectors have been used throughout the signal path, with the exception of the low-level front-panel interface (which is an N-type connector) and the SCT interface (which are flying leads). The N-type connector is used in many commercial systems for both power handling and physical robustness reasons. The SCT interface is defined by the connections presented at the glass envelope.

3.3.3. The Phase Splitter Module

Because the SCT drive must be push-pull, and the drive is provided by two separate final stages, these must be a phase splitter to derive the two antiphase signals if the instrument is to accept single ended inputs. Although the precise relative phases need not be held to better than a few degrees (in band), it has been found easier to achieve clean operation at a level

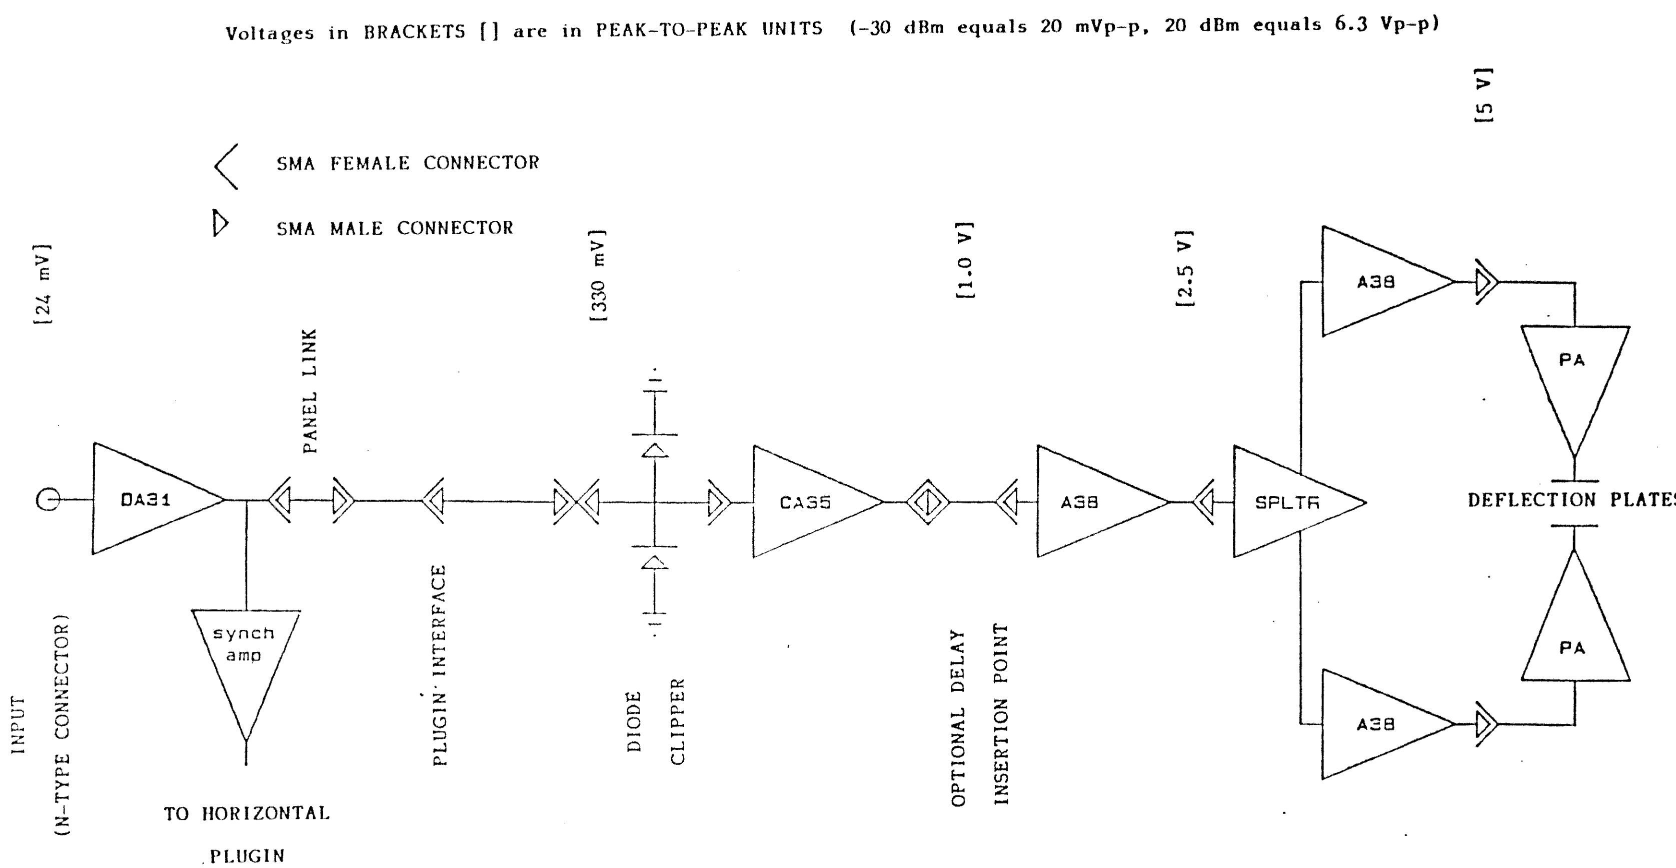

Vertical Amplification System Layout and Signal Budget

Figure 3.1

below the signal level presented to the final amplifiers. Small signal approximations are not seriously violated at the levels chosen, and circuit theory can be used to predict performance, simplifying the design procedure.

There are a number of general techniques which may be applied to obtain wide bandwidth operation. These fall into five groups:[61] Reflective matching, Lossy matching, Feedback, Distributed amplification and Active matching. More than one technique may be applied in any one amplifier. Reflective matching is the familiar process of designing the circuit with reactive elements (realised as lumped components or transmission lines) using Smith Charts. Lossy matching is a similar process, but one where resistive elements are also used. Feedback is the general use of any kind of component, but where signal feedback is applied from output to input of the active device(s). Distributed amplification refers to the use of multiple devices cascaded without separate interstage matching, but with overall matching (by some means) of the cascaded group. Active matching is a relatively recent use of active devices (usually FETs) to provide matching without gain, for the central amplifying device.

In general the selection of one or more of these processes for any given application is not straightforward. It is common, however, to use computer optimisation of circuit components. The optimisation processes are basically iterative numerical ones of considerable complexity. (This is the basis of programmes such as "Compact" and now "Super Compact".[79]) Indeed, multi-octave performance is not readily obtained using most of the matching techniques without computer adjustment of element values.

The lack of such software at the time of design, and the availability of relatively straightforward feedback configurations which satisfy the requirements for the circuits designed (without optimisation of a profound numerical type) dictated the use of feedback techniques.

At the levels indicated in Figure 3.1, a configuration strongly resembling the elementary common-collector and common-emitter phase splitter has been found to work satisfactorily. The emitter load (which should equal the collector load) acts as a shunt-series feedback element. Gain at both outputs is similar, except that the phase is inverted at the collector output. The bandwidth is very wide, despite imperfectly matched load, and device capacitances. Response extends beyond 2 GHz. The circuit is given in Figure 3.2A, and the board layout is given in Figure 3.2B. The layout diagram is stylised in order that the locations and identity of components should be clear. Also note that the board includes the two hybrids that follow the phase splitter in the signal chain. These are mounted below the substrate, embedded in machined recesses in the brass mounting plate. Thus it acts as a heat sink as well.

The circuit is fabricated in micro-stripline (microstrip) technology. This has been used because the WJ hybrids are designed for mounting in this medium. Series elements are

CIRCUIT OF THE PHASE SPLITTER AND FOLLOWING AMPLIFIERS

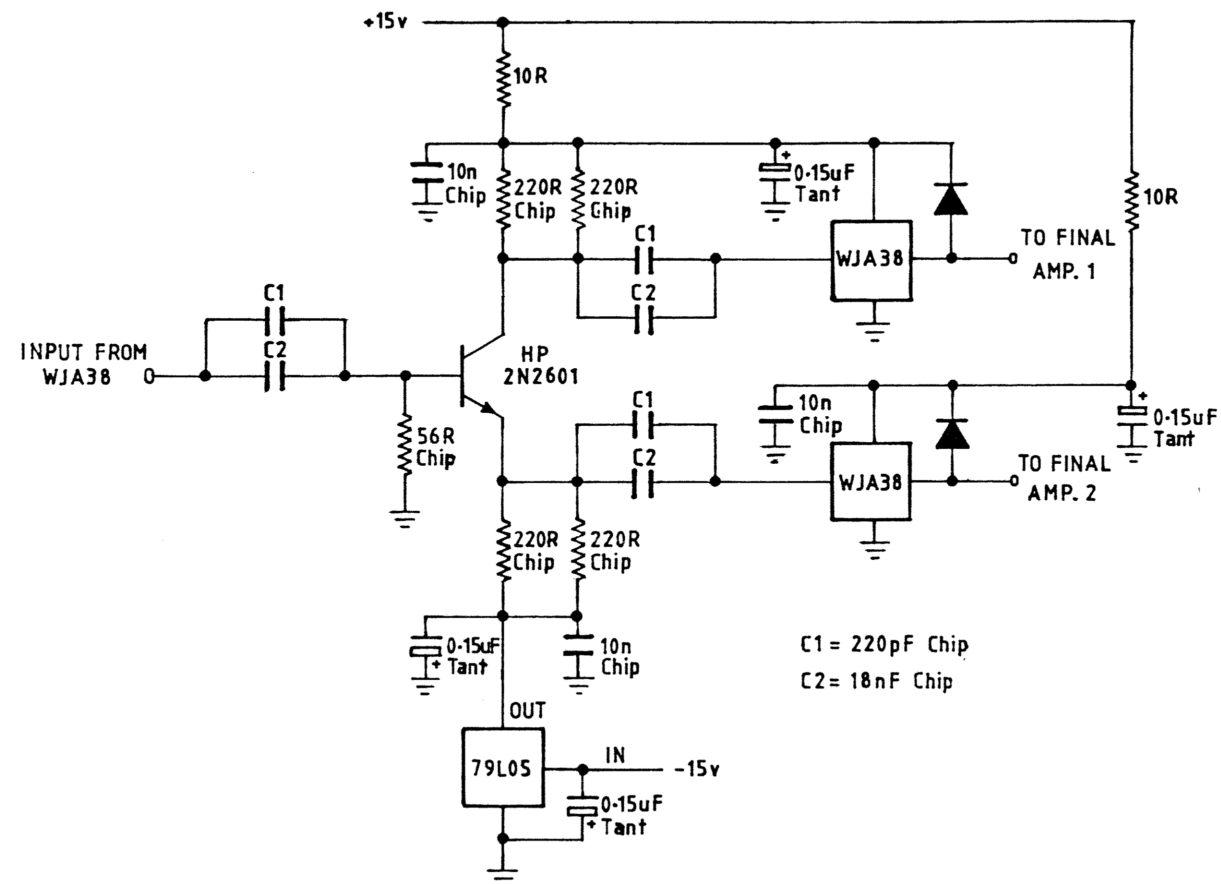

FIGURE.3.2A

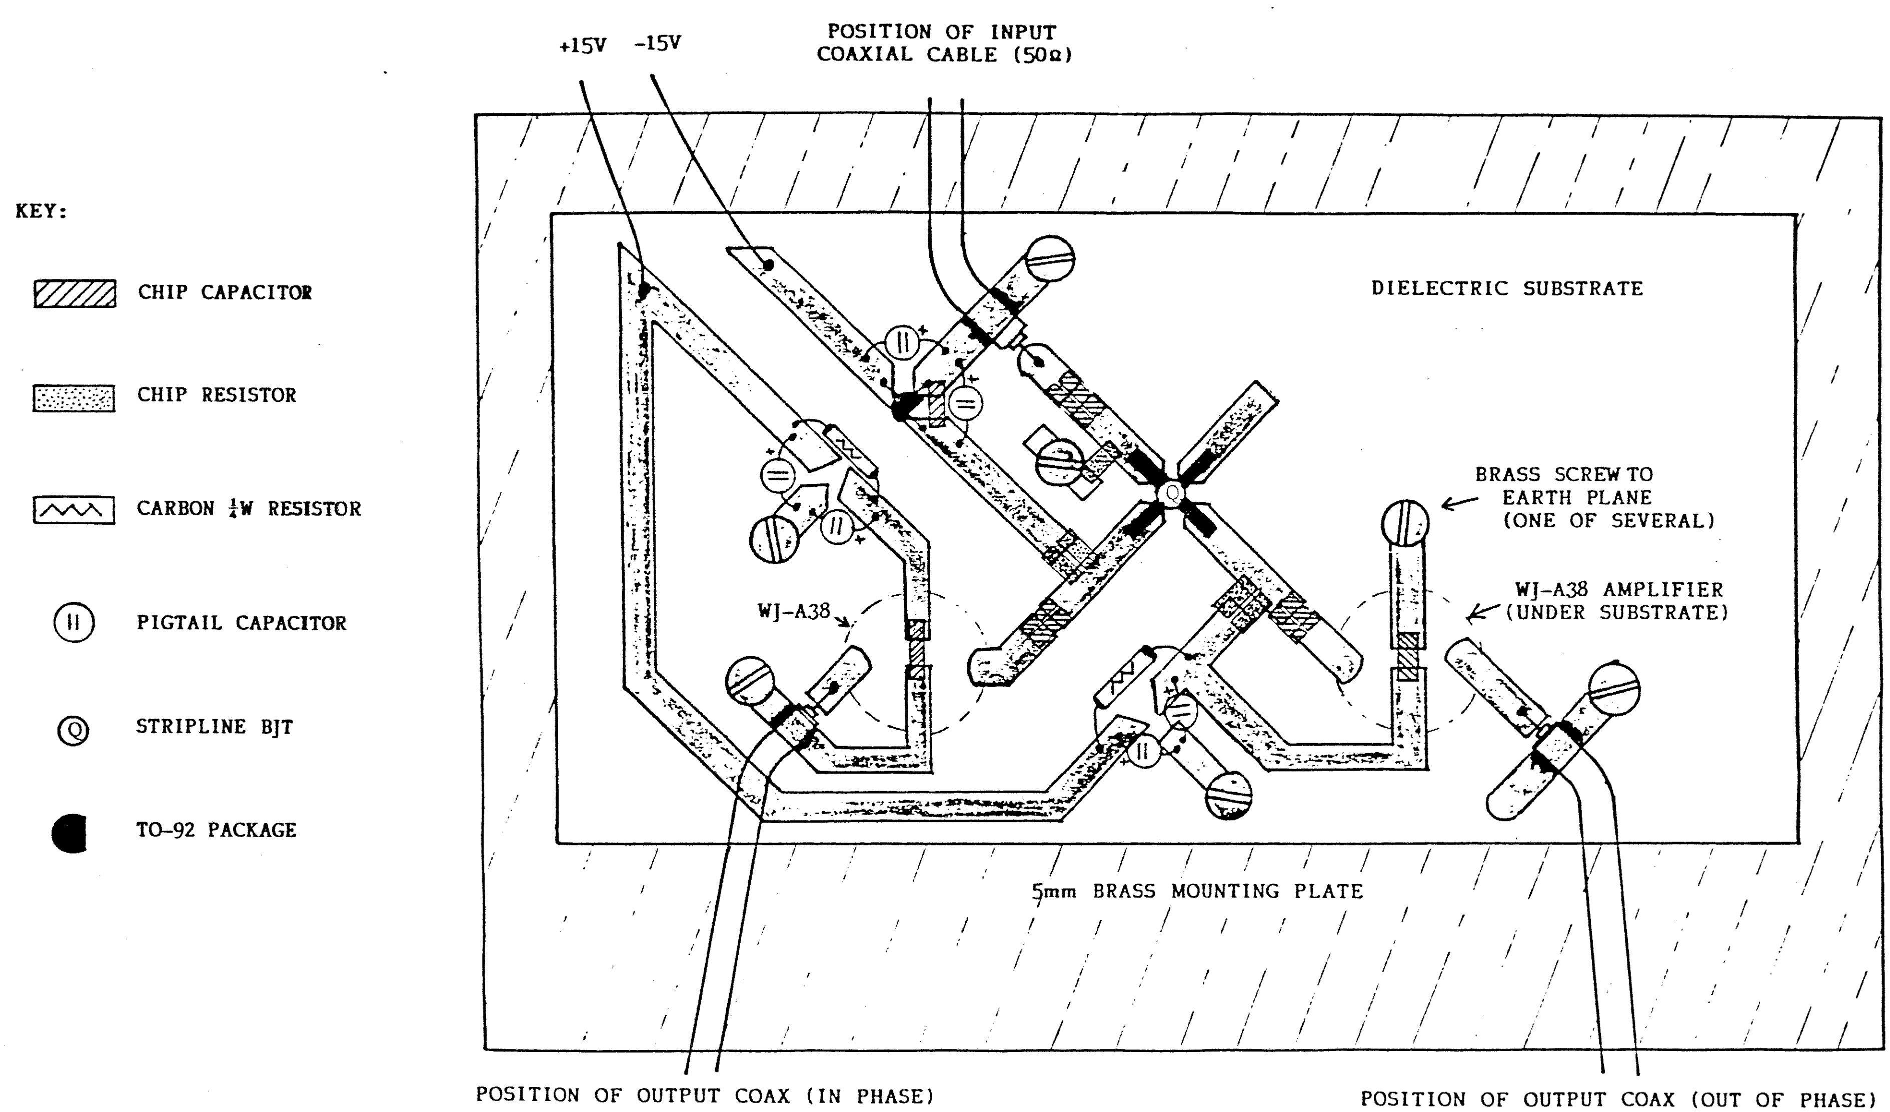

LAYOUT OF THE PHASE SPLITTER BOARD: NOTE THAT THE DIAGRAM IS STYLISED FOR CLARITY, AND INCLUDES TWO WJ AMPLIFIERS

FIGURE 3.2B

readily incorporated, but shunt elements must be provided with a grounding point which protrudes through the dielectric substrate. This is clearly depicted in the layout diagram. (Brass countersunk screws have been used.)

3.3.4. The Final Amplifier Module

The final amplifier is required to deliver a large signal output of 16 volts peak-to-peak into the SCT transmission line plates. Because suitable active elements have a maximum working voltage of 24 volts (V_{ce}) and the plates must be at approximately 40 volts, the whole circuit operates between levels of +24 and +48 volts. The collectors are biassed to +40 volts with respect to system ground, which is equivalent to +16 volts with respect to local ground. The special supply levels are provided by a power supply mounted near, and servicing, the final amplifier.

Once again, modest gain requirements permit the use of simple feedback arrangements. The amplifying element is a common-emitter circuit with emitter degenerating elements defining the gain. The input match is improved by shunt-series passive RC elements. The circuit is AC coupled to permit the operation with elevated local ground. A relatively complex biassing circuit surrounds each amplifier's active element, setting the bias and hence the target trace position. The bias circuit inputs are provided by a position setting circuit connected to a control on the plug-in facia. (See the next section.)

The circuit is fabricated in CPW. This permits the simple surface mounting of both shunt and series components. The SCT plates are accessible via a pair of pins on the top of the tube. A pair of short flying leads jump from the tube to the bottom of the final amplifier board. The connection offers a minimal mismatch because there is no need to pass a groundplane with the coplanar configuration. The Final Amplifier circuit is given in Figure 3.3A, and the board layout in Figure 3.3B.

Some comment on the layout is in order. The coplanar type of transmission line has facilitated neat and precise fabrication of the circuit, but was found to require specific care. The gold plated "bridges" visible in the figure are necessary to ensure the continuity of the ground connection for the high frequency signals wherever there is a side branch disrupting the ground side of a slot.

The capacitor pairs are used to provide satisfactory "lumped element performance" when effecting decoupling of the DC components. The larger value chip capacitors have sufficient inductance to become disruptive before 2 GHz. The second, smaller valued capacitor extends the utility to beyond 2 GHz.

3.3.5. The Power Supply and Position Circuits

The circuit of the power supply for the final amplifier is given in Figure 3.4. The supply provides +24 and +48 volts. The +24 volt line can source or sink current, as may be required

CIRCUIT OF THE FINAL AMPLIFIER (ONE HALF SHOWN)

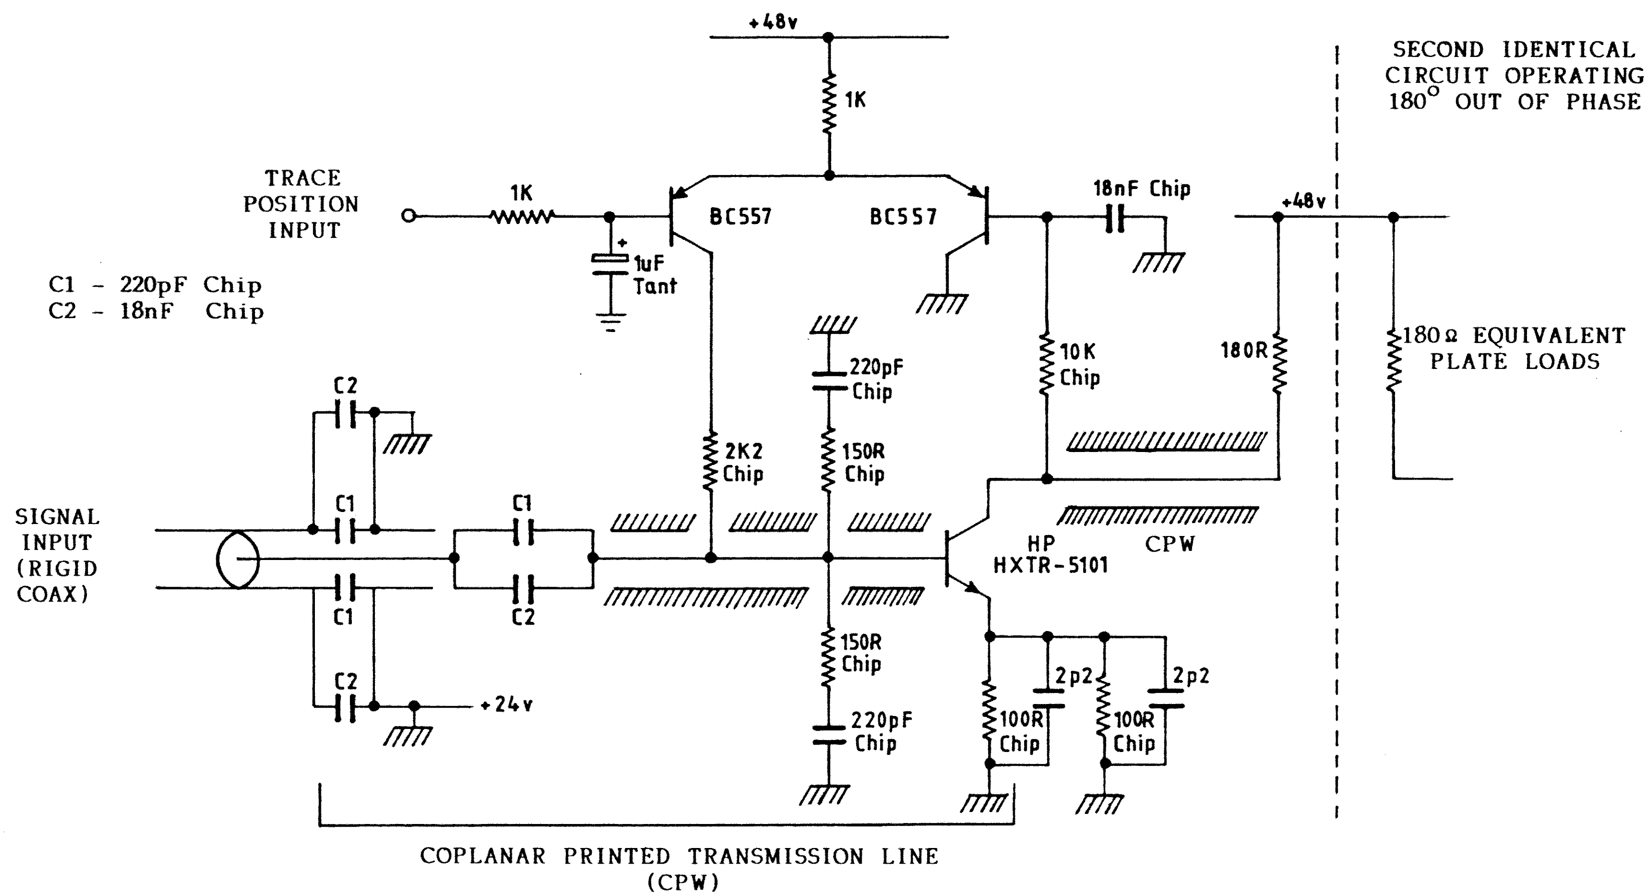

FIGURE. 3.3A

LAYOUT DIAGRAM OF THE FINAL AMPLIFIERS (BOTH SHOWN)

KEY: C1 = 220pF
C2 = 18 nF

= Ground Signal Bridge

= HP transistor (c=collector, b=base)

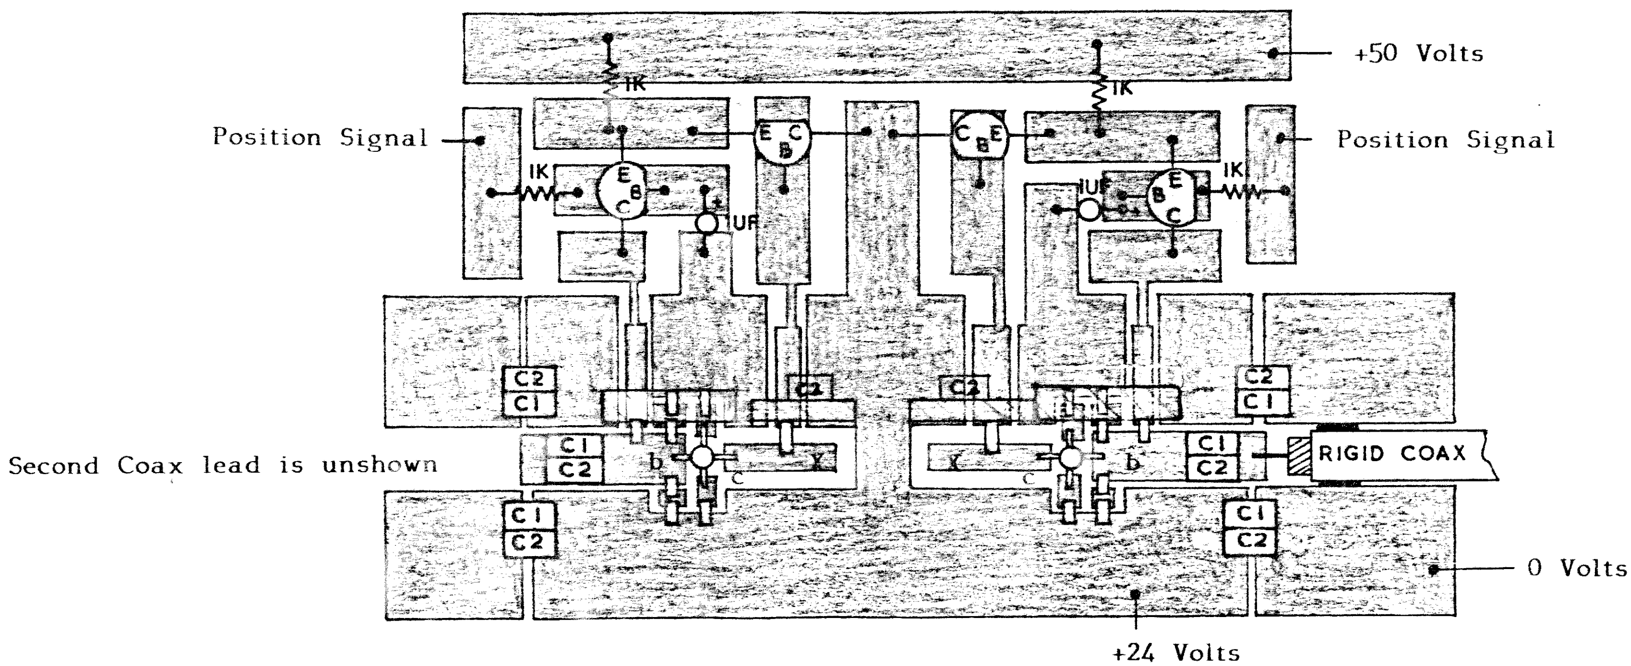

Flying leads to the SCT envelope attach at rear at points marked 'X'.

Chip components in the vicinity of main BJTs are un-labelled for clarity. They are readily identified by comparison with the associated schematic circuit diagram (a).

FIGURE. 3.3b

POWER SUPPLY FOR 7912ADM FINAL AMPLIFIERS

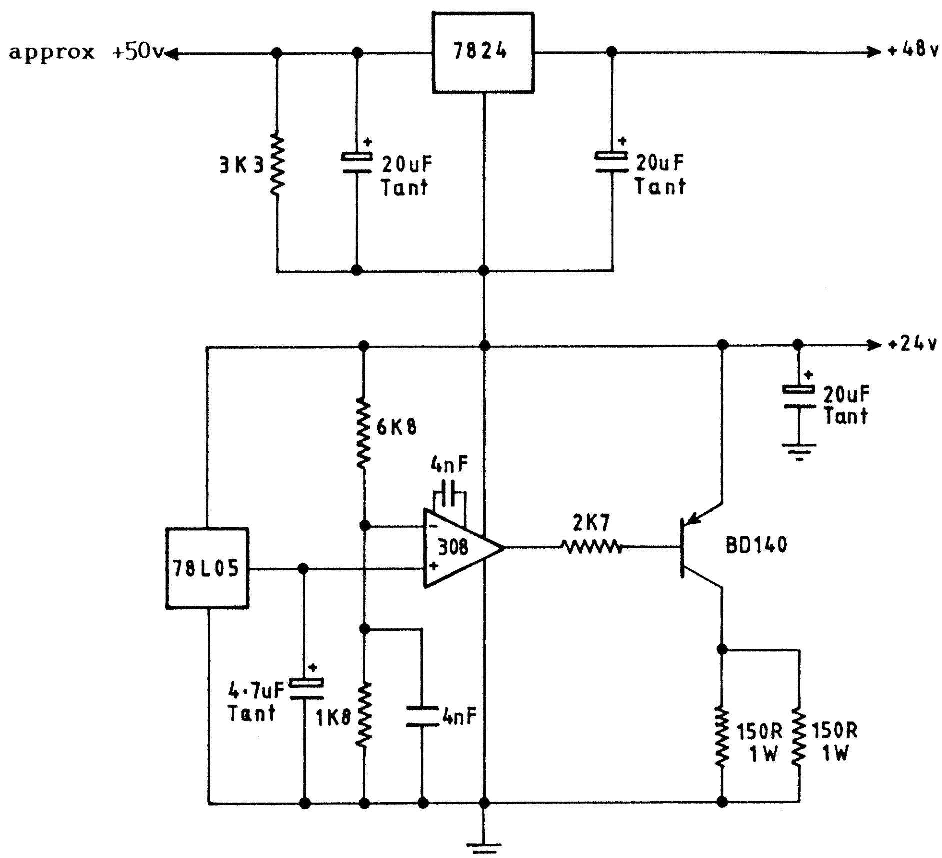

FIGURE. 3.4

because the final amplifier operates between the two rails. The upper 24 volt step is defined by a three terminal regulator, while the lower one is set by a circuit based on an operational amplifier, a driver stage and load dissipation resistors. This enables the sinking of current.

The position control circuit is required to provide two voltages. These must average +40 volts (+16 on the local ground of the final amplifier). The voltages range from +32 to +48, under the control of a potentiometer mounted on the control facia of the vertical input plug-in. Figure 3.5A shows the circuit, and figure 3.5B the layout. The board is mounted behind the plug-in interface carrying the control facia, in the instrument's left hand side bay, adjacent to the first hybrid amplifier module.

TRACE POSITION BIAS CIRCUITRY

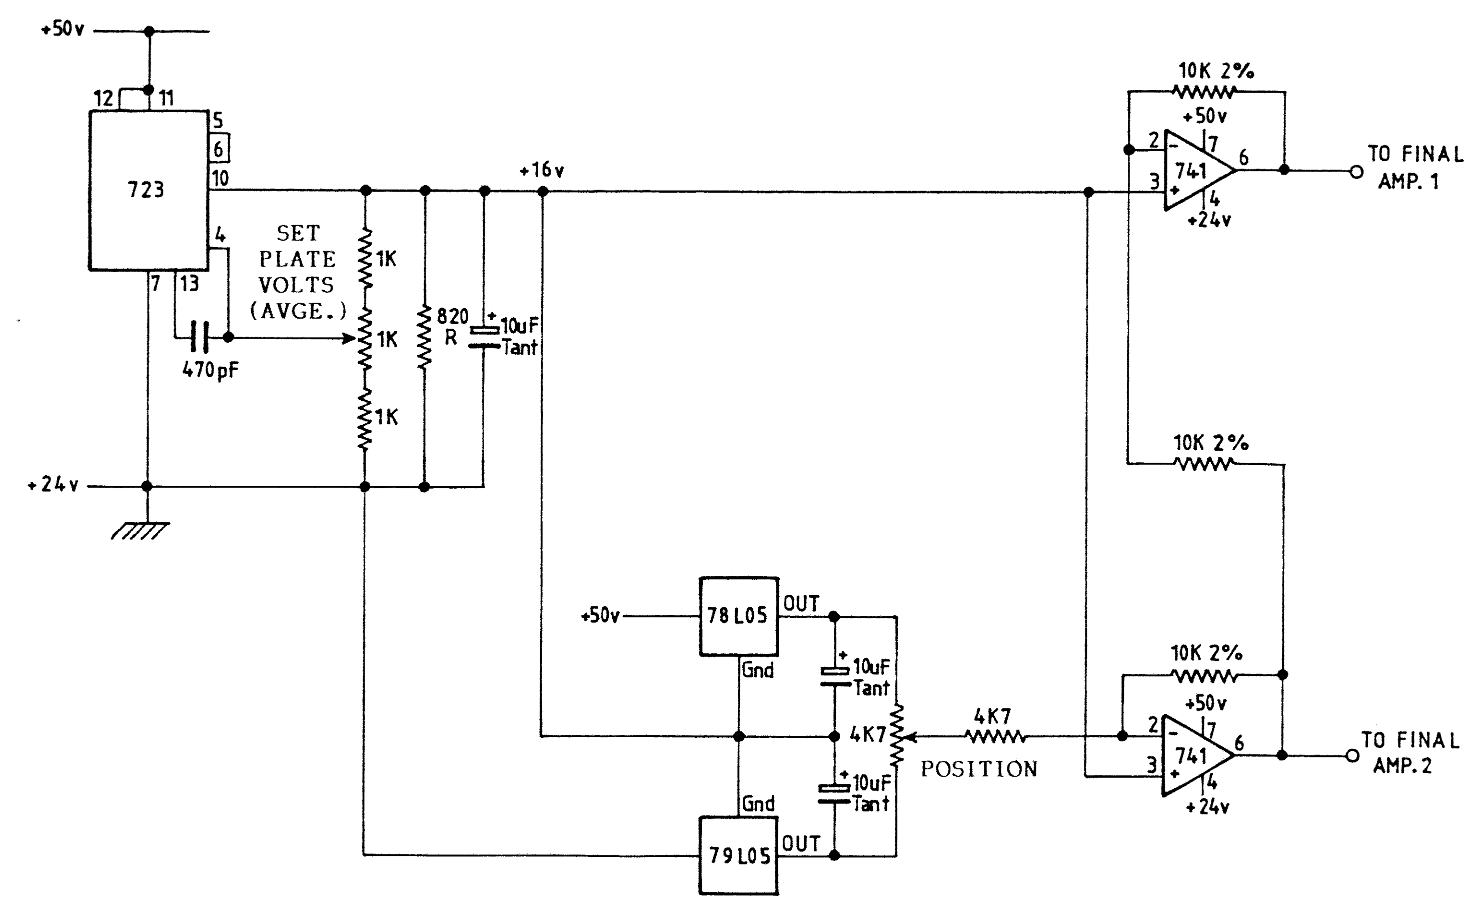

FIGURE. 3.5A

LAYOUT OF TRACE POSITION BIAS CIRCUIT BOARD

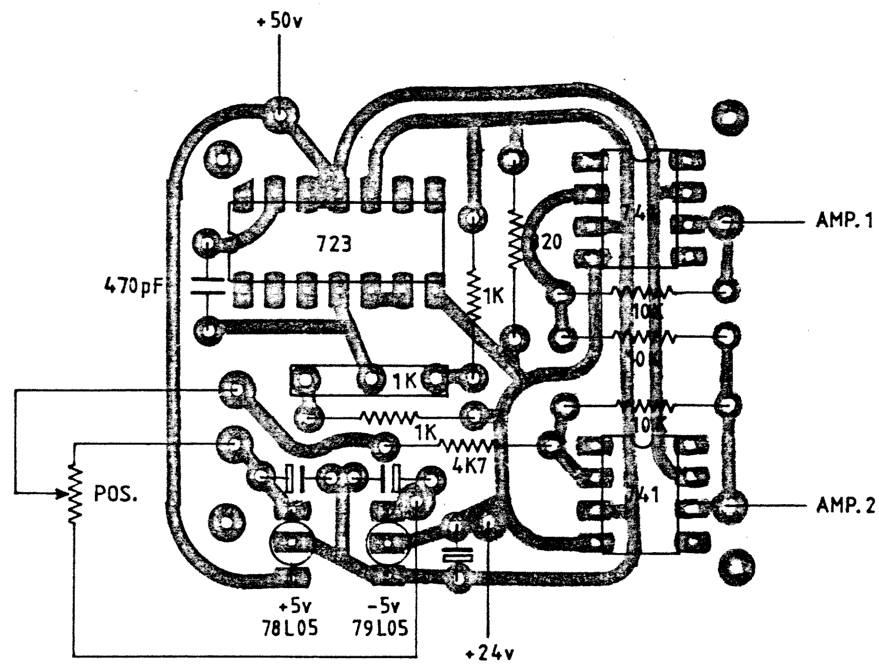

FIGURE. 3.5b

4. FREQUENCY PERFORMANCE ANALYSIS

As will be indicated in this chapter, there are four ways that the response of the 7912 can be assessed. Three of these are variations of the basic swept frequency measurement familiar to all electrical engineers. The variation arises because of the peculiarities of the Scan Converter Tube in the instrument. The fourth is a time domain type of measurement.

4.1. Swept Input Methods

The traditional and most direct method of determining the (magnitude) response of a component or system is to measure the envelope of the output response to a swept sinewave input. It is important to be aware of the fundamental limitations of such a measurement technique.

4.1.1. Limitations

There are three factors potentially limiting the accuracy of a swept signal envelope response measurement. These are:

1) The properties of the signal source which affect how well it represents an ideal swept sinewave function. These include the purity of the signal source, the accuracy with which its envelope is known or held constant, and its frequency stability.

2) The accuracy and resolution of the detector mechanism which is used to determine the envelope of the output.

3) The inherent fidelity with which the true frequency domain response of the system is reflected in the envelope of the time response to the finite duration input function.

The first and second limitations can be generally understood and estimated without serious difficulty. The last requires care during measurement. All are discussed below.

4.1.1.a The Signal Source

Any limits imposed by the signal source are independent of the remainder of the measurement system. They can thus be noted here, and a suitable signal source selected in the light of the desired performance at the time of a measurement.

The accuracy of the frequency of the signal source clearly places a bound on the accuracy of the result. In addition, the short term stability (jitter) may provide a limit on the resolvable feature size. If the source is jittering across a bandwidth delta-w, then the detected output can at best be taken as an average of the true response H(w) across the bandwidth delta-w. (Naturally, if there is bias in the jitter it would be reflected in the average. This is not usually the case.) If the detector can respond in the bandwidth of the effective frequency modulation of the jittering signal, and the output is sampled for long enough to permit observation of the output variation, then the presence of the jitter will be detected. In general, sources are available with performance well beyond the point where these problems become significant.

The purity of the signal is also a consideration. Whereas jitter represents the occupancy of a finite bandwidth centred about the average frequency, the purity of the signal represents the total amount of signal (expressed as additive voltage) from all other frequencies present. Care should be taken when selecting a source, since some modern sweepers with wide range have total spurious responses only 25dB below the fundamental. While signal averaging deals with the non-harmonic signals from such generators, this cannot be guaranteed to be the case in general. Again, however, suitable sources are available with specifications such as to remove this worry completely.

Accuracy of amplitude is probably the largest source of error. More advanced generators employ detectors in feedback loops to stabilise the output. The detectors used are comparable in accuracy to the best obtainable separately, and so the precision of a source amplitude function is as good as is measurable by the usual means. This is in the small fractions of a dB, and so is again not of concern. It is necessary to allow for this in the quotation of final accuracy, however.

4.1.1.b The Detector

There are three properties of a detector which may affect the result when determining an overall response. The first is the accuracy (mainly linearity) of the detection mechanism. The second is resolution. The last is speed of response to a signal whose amplitude is to be determined.

In measurements made with the 7912, the target is the detector. The limitations encountered will vary depending upon whether it is used in an envelope-capture mode, or in the normal trace capture mode with a single frequency. In either case, pixel position is available with more resolution than is generally required in a response measurement. In other words, since a response is rarely determined with more than a fraction of a dB of both accuracy and resolution at microwave frequencies, the target resolution (and deflection accuracy, etc) will not be of concern.

Where a single frequency sinewave is digitised, speed will not be of concern, of course. The detection is performed with an image of a sinewave, and so effectively 'responds' in one cycle. The situation is more complicated in the case of envelope mode, and the discussion of this is postponed to the appropriate section, (4.1.3).

4.1.1.c System Response

Even with a perfect signal source and an ideal detector, the response curve obtained does not identically represent that response which would be obtained from the network under test, to an infinite duration sinewave signal at each frequency. This arises from the interaction of two limitations: The network bandwidth is finite; also the exciting signal is of finite duration, and thus infinite bandwidth. A simple example of a practical and obvious case serves to aid the realisation of the limitation.

Consider the case where a high Q (narrow bandwidth) circuit is being observed. If the observer sweeps the signal source past the resonance frequency too quickly, only a small deflection, or no noticeable deflection at all, is perceived at the detector output. This might be intuitively expected by recalling that the response of such a circuit to a sinewave burst with rectangular envelope is a pulse of greatly smoothed envelope. This effect is a limitation found in analog spectrum analysers, which have a limiting scan speed for any particular required resolution.

Any response graph will differ to some extent from the ideal graph, because of this interaction of the bandwidth of the system under test and the spectrum of the test signal. The obtained function will be a filtered version of the ideal result. Because the bandwidth of the system forcing function (the swept input signal) is a function of the speed of sweep,[83] it can be arbitrarily compressed by reducing the rate of sweeping. A manual sweep measurement thus guards against distortion induced by this filtering. An observer typically chooses a portion of the emerging graph where the waveform contains higher frequency components, that is, where it moves faster. Here, he will sweep at a rate, then at a lower rate; if the results are identical, then the filtering effect is indistinguishable to his available resolution, and he is sweeping at an acceptable rate.

The filtering effect experienced will be a function of the system whose bandwidth is being sought. It is thus not possible a priori to calculate it. Where an automatic measurement or one with fixed forcing function is used, this limitation may (and indeed must) be estimated by some form of calculation. Such is carried out in Appendix E, where there will also be found a discussion of the results. Reference will be made to the estimate therein when appropriate.

4.1.2. Discrete Digitisation Method

The most rigorous method of obtaining the response consists of instructing the 7912 to take a single sweep digitisation at each of a large number of separate single frequencies. The timebase can be set to ensure that a suitable number of cycles are written to the target. This method possesses four advantages: The response envelope can be sampled at an arbitrary set of points and a detailed result is thus available, with the possibility of improved sampling at fast moving features of the curve. The output is immediately in digitised form. Detection is performed in software, giving optimum detector performance. The detector resolution is the full resolution available from the SCT in operation with a signal spaced on the target so as to obtain optimum resolution.

The major problem with this method is that it is inordinately slow to execute on a 7912 digitiser in the digitise mode. This is partly because of the delay in transferring the digitised result to the controller. It is also a result of the need to guard against spikes and other anomalies polluting occasional results. This requires one or a combination of: visual inspection of digitisations, software checking of digitisations and signal averaging.

(The delay is approximately 20 seconds per double valued array with the system used. This is extended to several minutes if the detailed Actual Beam Path (ABP) algorithm, as implemented, is invoked, to maximise fidelity and resolution. It is also more than multiplied by any signal averaging factor employed, if the Intelligent Signal Averager is used. These times would be reduced if a fast controller running compiled code were to be used. Refer to the ABP, EDG and ISA routine descriptions in the DIOS Manual (section X) for more details of the tasks.)

A typical run of the kind described above has taken around 4 hours. This is without continuous visual inspection or the use of ABP or other resolution enhancement routines, but with crude detection and rejection of abnormal samples. Estimates can be made of the time required to improve confidence in the result. As the confidence level itself is not readily assessed, these are not certain, but it is not difficult to arrive at figures of days.

A further problem with this method is that sine waves approaching full deflection on the highest timebase speeds exceed the maximum writing rate at their zero crossings. This does not prevent the measurement, but complicates the algorithms used to determine confidence in the reliability of each result.

Recalling from section 4.1.1 that instrument response is generally not resolved to more than a fraction of a dB, the need for resolution to more than 3 or 4 bits is unnecessary. Even under these circumstances it is a slow procedure, and the risk of undetected accidental disturbance of the measurements because of the lengthy duration is disquieting.

4.1.3. Single Trace Method and Results

The scalar response of a microwave component is often displayed on a simple CRO using a detector and a sweeper with a sweep rate of the order of several Hertz. The response of the 7912 can be encoded into a single write to the target using it in such an envelope mode directly. Here the target captures a two dimensional pattern, its height in any column relating to the magnitude of the deflection at the time (proportional to frequency) the spot traversed that column. The apparent elegance and efficiency of this method, which delivers (potentially digitised) results quickly, is slightly superficial.

The tube target susceptibility to burn-out has an effect. The slowest permitted sweep speed for the 7912 is 1 ms/division, or 10 ms for the full window. This **very fortunately** corresponds to the speed with which a state-of-the-art sweeper is capable of covering 2GHz. This rate translates to a sweep of 200GHz/second. This is potentially satisfactory, in the case of a substantially uniform wideband response up to 2GHz,[Appendix E] since the resolution would be first limited by the number of effective samples available. Even with all 512 samples (in 2GHz) the limit would be about 4MHz.

This potential resolution of 4MHz is not delivered. Since the target must act as the detector, the same problems which reduce resolution in the capture of certain waveshapes limit performance here. The spot dimension is usually about 4 or 5 pixels, even with best focus and optimal low beam current. Thus the "smoothing" effect of the size of the spot used to draw the image only allows the perception of features of width comparable to or exceeding 4 pixels, or 20MHz. Where the spot is larger, the potential for error is greater. This is the "forced oversampling" noted in Chapter 3. Here it limits resolution in the frequency domain, since the envelope capture depicts response magnitude vs frequency rather than voltage vs time.

Providing an acceptable strong image is written, the actual speed of response of the target is not going to be a limitation. However, where there is potential image overwriting, because of the spot dimension, it must be realised that there will be an effect similar to a slow detector response. Rapid (and hence closely occurring) variations may be washed out of the stored image by the overwriting. A narrow notch might be filled by overlap of the spot writing its steep sides close to each other.

In order to write an envelope to the target, the intensity must be set at a level which represents a considerable **total** current in the target. A limitation is in fact the target burn-out security circuit, which shuts down the beam if excessive current is detected (protecting the target from damage). The sensing circuit does not know the distribution of the beam across the target, and so must be pessimistic in its operation. Just below the limit, the overall image is faint, and the charge collects at the envelope edges, since, of course, a sinewave spends most of its time at the extremes of the function.

The bloom of the trace blurrs the edge, because the spot dwells at the important envelope edges. The true vertical position of the trace is confused, and there can be no use of the Actual Beam Path (ABP) routine, or other algorithms which require the spot path edges. The obtained envelope edges are perceived as having been displaced outwards, and the displacement will be a function of the instantaneous path width. The envelope path width affects current density arriving at the target areas being exposed, assuming constant beam current. These errors should be systematic, and thus could be removed. Empirically determined functions could be developed to provide correction, given intensity (beam current) and sweep speed data. This is not considered worthwhile.

The limitations discussed above are readily discernable in the sweep images of Figure 4.1. The two images are of single sweeps written to the target by an HP8350B Sweeper. The sweep is linear, from 10 Megahertz to 2.3 GHz. Amplitude markers are imposed at 10 Megahertz, 1 GHz and 2 GHz. A signal at -11 dBm (approximately 180 mV p-p or 22 mV/division) is fed in at the mainframe link. The conditions for the two images are identical except for the intensity (beam current) setting of the 7912 and any compensation effects of the wide latitude film used in the camera. The film and processing prevents the perception of differential brightness in the white parts of the two images. The brown halo visible surrounding the high intensity sections is an aberration of the camera. A Celtic video photography system was used with Agfa C-41 wide exposure tolerance film. The pictures are taken of the analog video output of the 7912. Virtually instantaneous output of target pattern is available from the 7912 by means of the analog video output circuit described in Chapter 2. This basically transfers the continuously written target charge distribution to an external monitor.

It is interesting to observe that the centres of the envelopes are not successfully written in either case, in spite of the wide latitude film. Next, note that the blooming in the image taken at higher current is sufficient to completely conceal the amplitude marker at 2 GHz, barely visible in the one taken at a lower setting. The charge hurled at the target is sufficient at the edges to 'flood over' and fill the troughs corresponding to the markers. Just vertically above and below the marker there is

SINGLE SWEEP SCALAR RESPONSE OF 7912ADM (version 3.00)

Markers at 0.01, 1.0 and 2.0 GHz

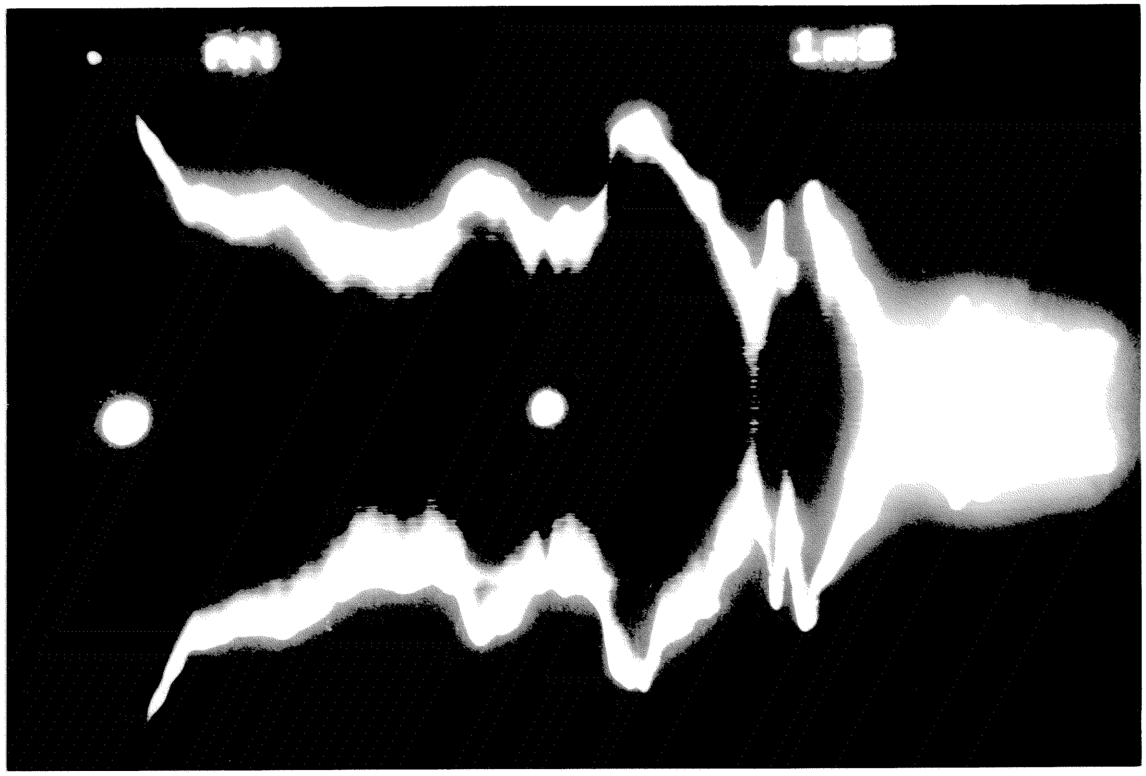

Lower Beam Current

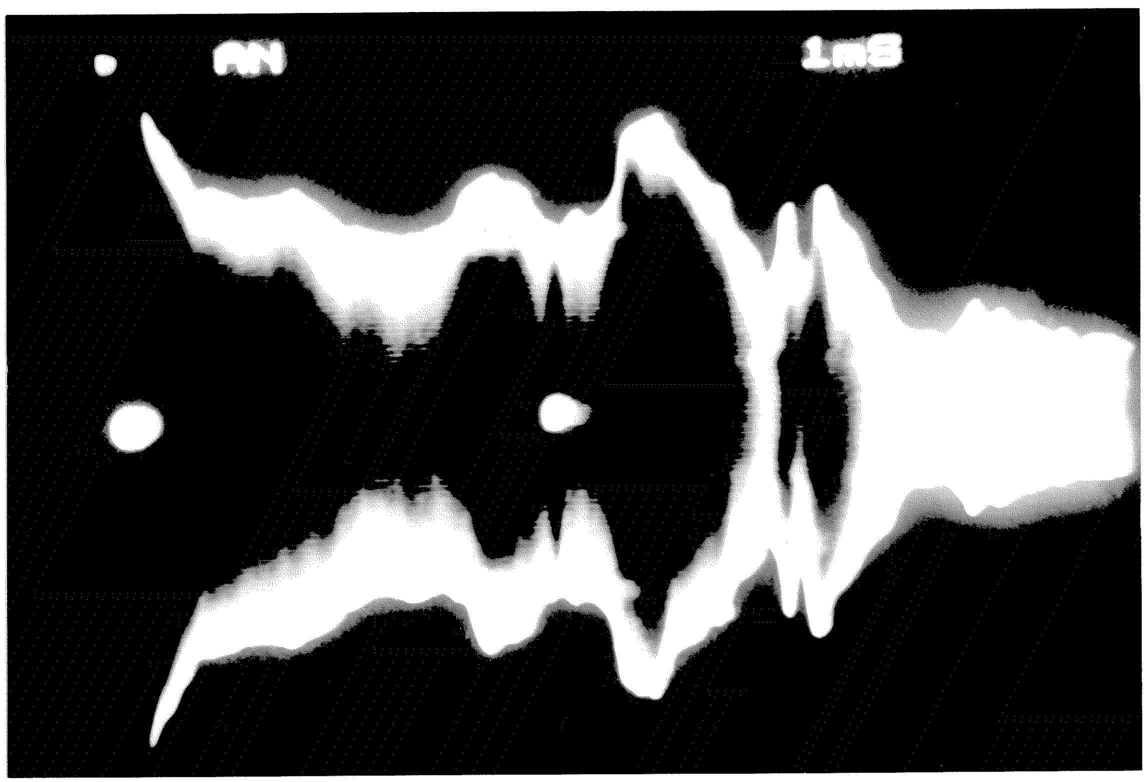

Higher Beam Current

FIGURE 4.1

a small section of the target which has not saturated. Further above and below, it is filled by the increased bloom of the trace, because the spot has spent more time at the extremes. Conversely, note that near the upper peak at 1.3 GHz on the lower current image there is a section of the trace which is almost lost completely because of the rate of change of deflection. Observe that between the 1.65 GHz double peaks the trough depth is not certain. (The depth will be noted shortly, and comparison drawn in section 4.1.5.)

The differences in the top and bottom trace edges are attributed to the differing sensitivity of patches of the target. Figure 4.2 depicts a sinewave written at low current. A square dark patch is discernable in the background. The tube used was not new when made available for the work here. It had been used for an indeterminate period in a configuration which caused scanning only of the centre one quarter of the total target area. The use (or abuse) of the tube has apparently had a detrimental effect upon the sensitivity of these pixels (diodes). There is in fact a small area (some tens of pixels) which either never functioned, or has been destroyed. Its proximity to the centre of the target may be coincidental, or it may have resulted from an excessive stream of undeflected electrons. During the course of the work Tektronix provided a ROM revision which causes the controller to turn off the beam current if the machine is left unattended with triggered sweeps for a period of several minutes. This step was most likely prompted both by target and cathode wear.

ANALOG DISPLAY OF A SINGLE SWEEP AT HIGH SPEED

EXPOSING TARGET IRREGULARITIES

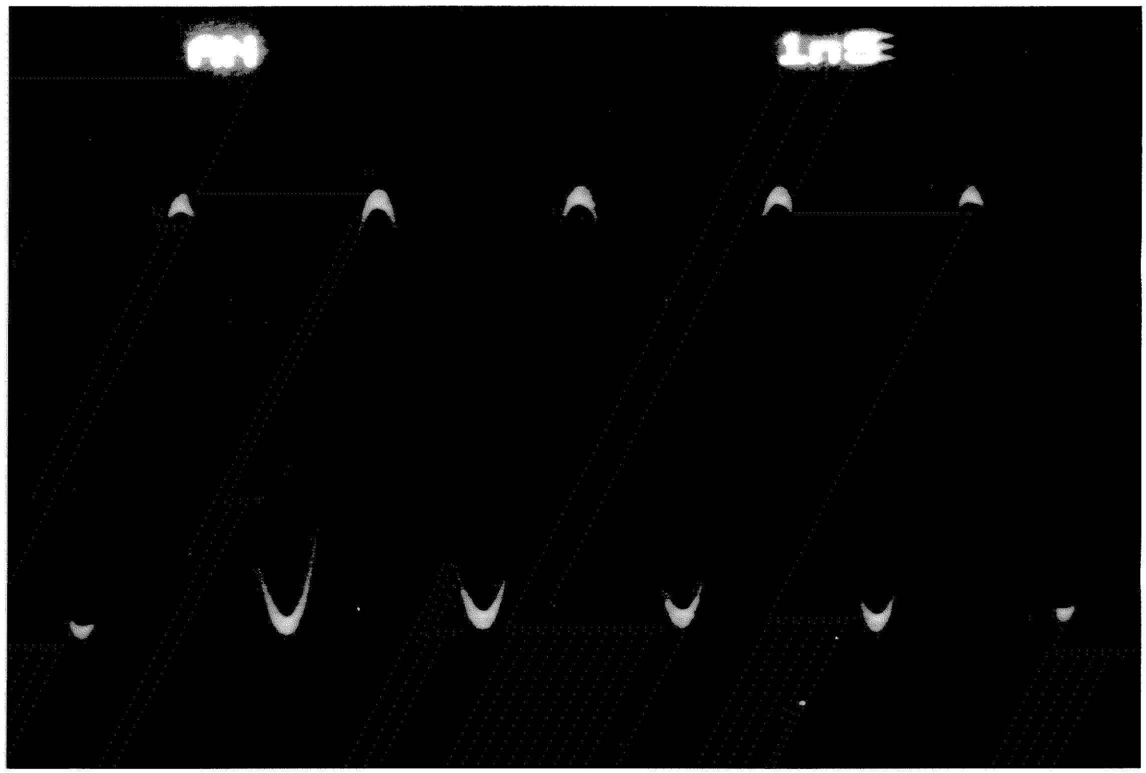

FIGURE 4.2

Figure 4.3 presents a response graph, obtained similarly to figure 4.1, for the case of the signal injected before the low noise amplifier and trigger pickoff circuits which are installed in the plugin. Three observations may be made comparing the result with the plugin amplifier to that obtained without same. The response above 2 GHz falls off rather more dramatically. This is a direct result of the DA-31. Its multistage feedback construction and relatively high gain for a single unit might lead one to expect a fast decline in response above the highest specification frequency. The overall magnitude response shows additional minor "ripples". These can be attributed to interstage mismatches arising from the spacing of the DA-31 from the next stage in the mainframe. This will be analysed further in the discussion of the results. Finally, below and to the right of the central marker, the target defect has been made apparent by the surrounding illumination.

(To give a better concept of the formation of the images, one can imagine a contour map of charge versus XY-position on the target having a very gentle undulating form. In the case of an envelope waveform, the charge distribution of the map can be regarded as having been built up by a gentle to and fro hosing of the charge buckets (pixels) by the beam. Charge collects at each pass of the beam. If the beam current is high, one pass of the beam over a pixel may register sufficiently to be read out. This is the normal case in a single simple waveform digitisation. To acquire an envelope waveform, the beam current is set much lower. The

SINGLE SWEEP SCALAR RESPONSE OF 7912ADM

INCLUDING PLUGIN PRE-AMPLIFIER (version 3.00)

Markers at 0.01, 1.0 and 2.0 GHz

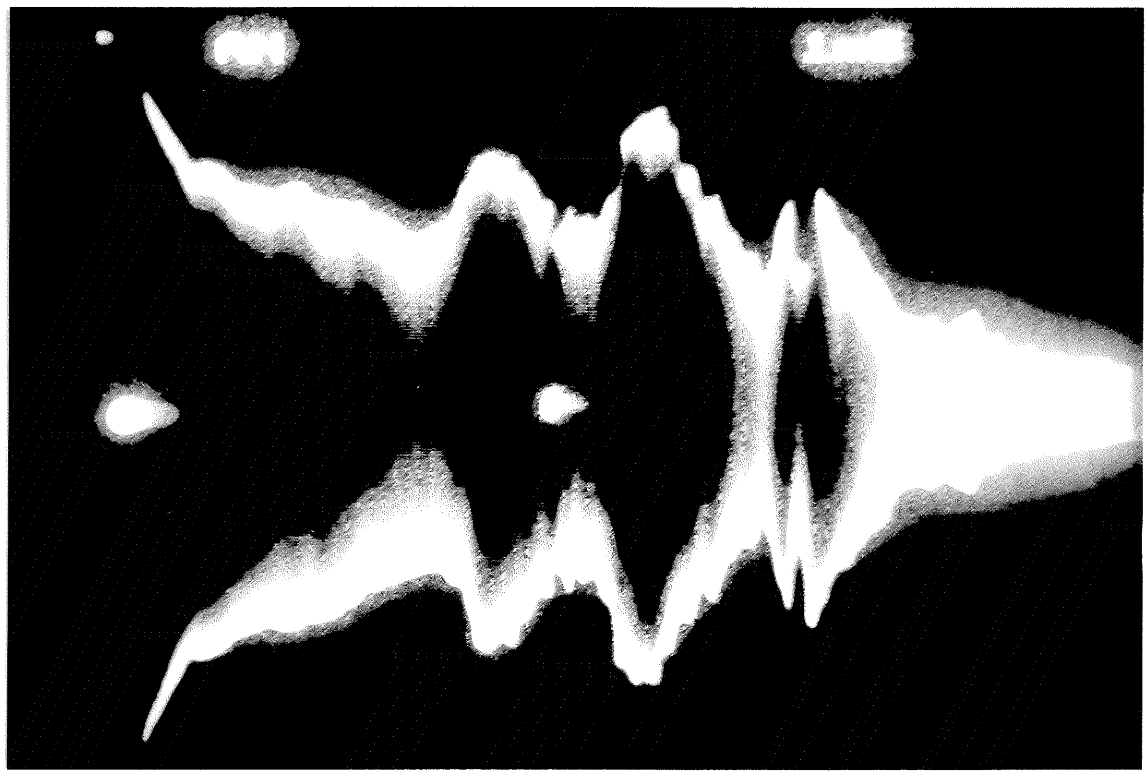

FIGURE 4.3

multiple passes build up charge in the diodes representing each pixel. The charges accumulated at each pixel site correspond to hills on the contour map. When the beam is confined closer to the centre, the hills will be that much higher, since the same charge is delivered to less area. The current is the same, but the current density at the target increases. The hills so formed will naturally "fall" down, becoming lower and wider, which corresponds to the charge buckets "overflowing", resulting in charge spilling into adjacent pixels. This description fits the case of an overexposed CCD imaging device.[84] In reality there is a causal difference, because the 7912 target operates 'subtractively', hits depleting charge rather than accumulating it. Nevertheless the effect in the 7912 appears similar, though less serious, which is a prime motive for the preference of its technology over that of CCDs. The flooding effect probably arises from spreading of the beam. The variation of charge distribution over the "exposed" area of target is clearly discernable in the analog images of envelope waveforms given in figures 4.1-4.3.)

4.1.4. Analog Display Method and Result

The penalty in using the analog signal fed to a monitor is that the result is not digitised or stored. The facility is intended to aid in setting up the instrument using repetitive, or at least relatively frequent, events, and so has no need of permanent storage. However, with some drawbacks, it does allow manual discrete frequency response measurement.

The resolution is comparable to a simple conventional CRO; about 5 bits. The output is obtained by plotting manually-read peak-to-peak amplitude results, similar to that of Figure 4.2 above. This approach permits very slow sweeping, with data intelligently recorded when some feature is encountered. On the whole, beyond these limitations, it is a much more confidence-inspiring technique, because the operator is close to the process, and can adjust the frequency as precisely as required, limited only by his control of the generator.

The response of the 7912ADM measured by this method is shown in Figure 4.4. For comparison purposes, the results from the Single Sweep images of Figure 4.1 are plotted on the same diagram. In transferring the results from the photographs, the most favourable image was used for each point plotted.

4.1.5. Discussion of Results

The degree to which the response rises and falls, sometimes quite rapidly, is initially striking. There is signal available to well beyond 2 GHz, but up to that frequency as much as 13 dB between troughs and peaks is observed. These variations can be attributed to interstage reflections. The typical input and output VSWRs of the WJ amplifiers used range from 1.5:1 to 1.7:1; worst case values for typical operating temperature ranges are 2.0:1. (These figures represent a voltage reflection coefficient ranging from 0.20 to 0.33.) It is easy to attribute

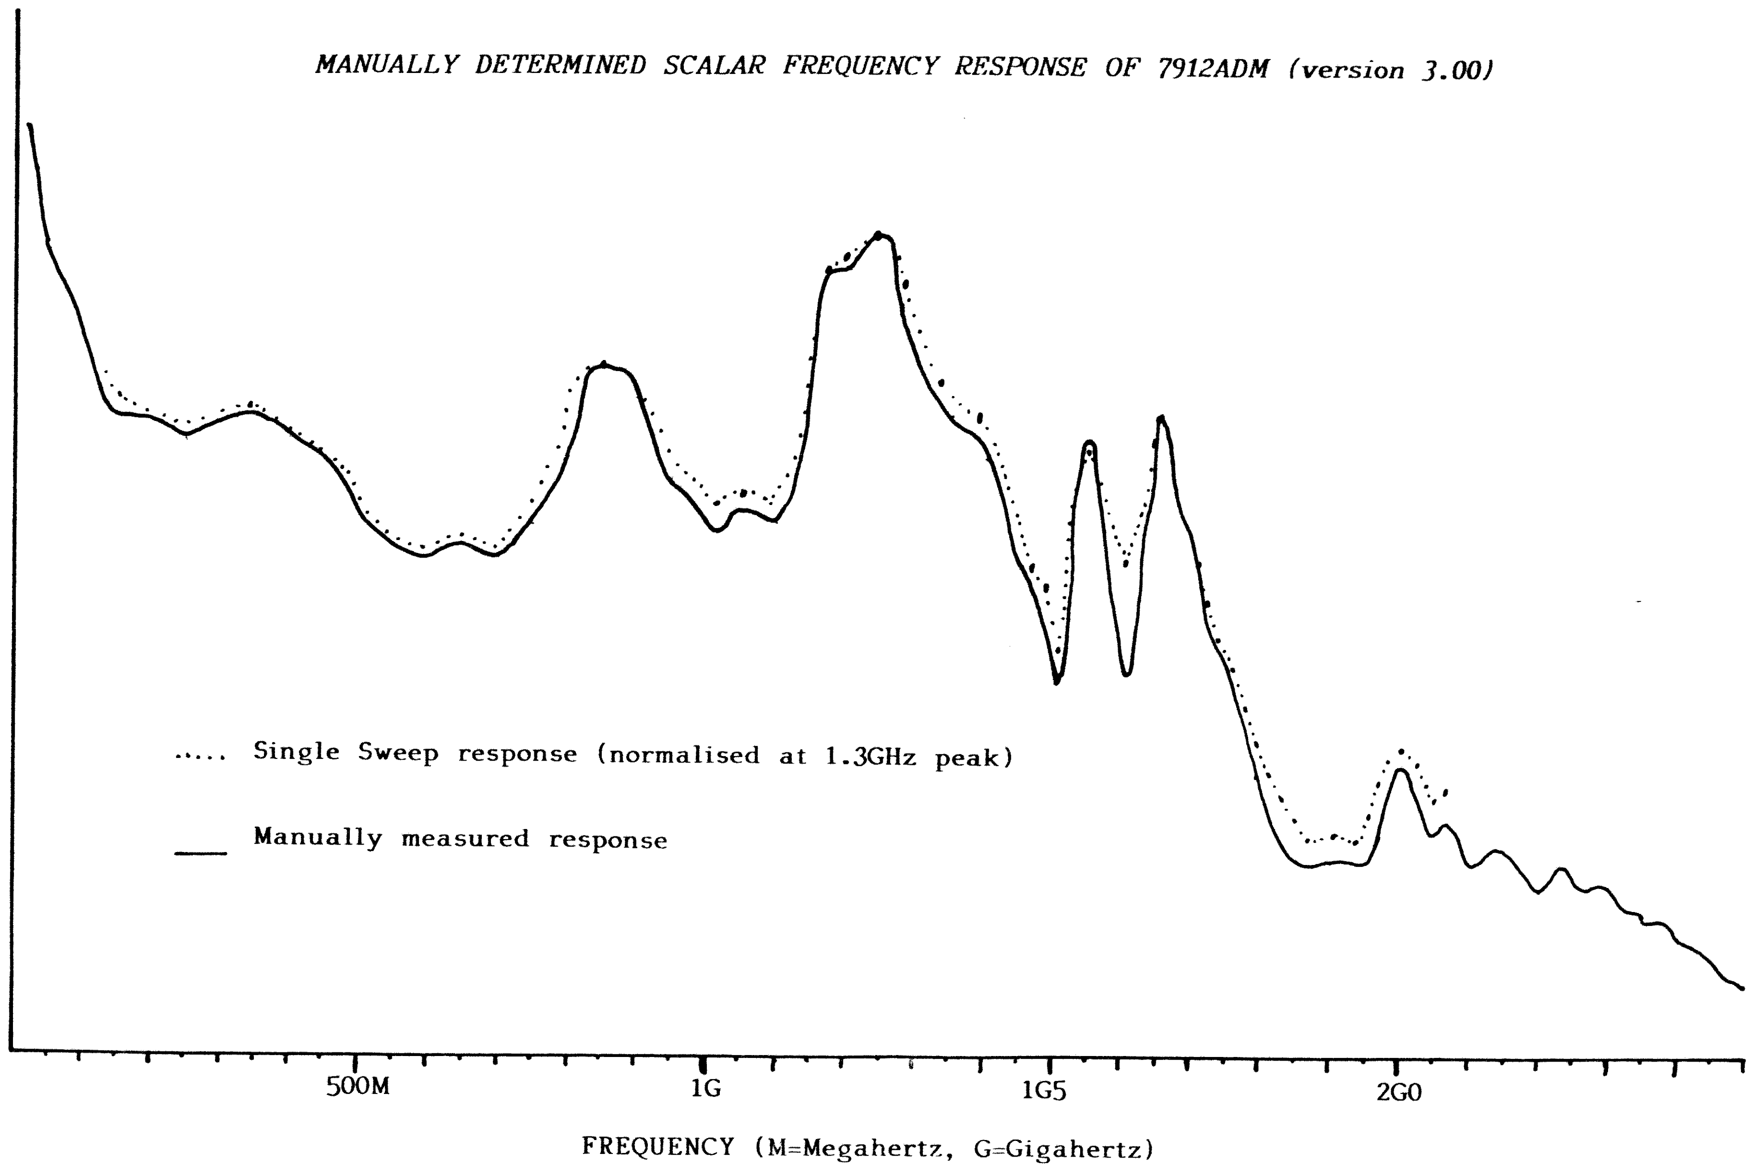
MANUALLY DETERMINED SCALAR FREQUENCY RESPONSE OF 7912ADM (version 3.00)
..... Single Sweep response (normalised at 1.3GHz peak)
Manually measured response
500M
1G
1G5
2G0
FREQUENCY (M=Megahertz, G=Gigahertz)

FIGURE 4.4

13 dB of variation to the combination of interstage interactions and normal roll-off approaching 2 GHz. This theory derives support immediately from two observations. Firstly, the magnitude variation is cyclic. The periods can be observed and the frequencies of repetition are found to correspond to quarter wavelength transmission lines of lengths very close to the interstage coupling transmission lines used. The major effect (of approximately 420 MHz) is found to correspond to the lines from the Final Amplifier to its WJ-A38 drivers. The Final Amplifier input SWR, though empirically optimised by selection of the matching arm components, is of the order of 2.5:1. This is a most acceptable figure given the available design tools and the uncommonly high levels at which it is asked to work. Secondly, the theoretical roll-off for the Final Amplifier is 2.1 GHz, operating into ideal 180 Ohm loads. Although not directly measurable with available tools, parasitics associated with the flying lead connections imposed by the 7912 SCT are believed to account for a premature fall in its response. The pictures suggest an underlying fall in amplitude above 1.5 GHz. Sub-picofarad lumped capacitances only are required to account for this order of effect.

Discussion of the implications of such a response will be taken up at the end of this chapter in section 4.2.6.

The differences between the results of the two methods of measurement are exemplified in Figure 4.4. The photographs of single sweeps covering the band expose the limitations of the technique. The author has high confidence in the accuracy of the manual method. The blooming of the bright images of Figure 4.1 leads to concealment of trough depths, as well as errors in estimating the relative magnitudes of peaks. Apart from the inability to show sharp features, the blooming appears to introduce the expected systematic error in the form of increased trace thickness at lower deflections.

4.2. Digitised Pulse Method

An ideal Dirac "delta" function (the impulse of unit area and zero width) contains all frequencies in equal amounts. It is well known that the response of a system to that impulse is simply the transform of the frequency response.[82] In fact, any wideband pulse may be used to characterise a system, provided the pulse has frequency components of adequate amplitude spanning the frequency range required for the characterisation. In addition, since the phases of the input components are known, and since they all pass the component under test simultaneously, the phase response can be determined as well as the magnitude response.

The idea of measurement of system response by the reaction to a wideband input[85] is used to advantage in many branches of engineering, more usually in low frequency applications. Two significant advantages of this technique make it appealing for the measurement of response of a 7912 digitiser.

Firstly, acquisition is rapid. It requires a single array of samples upon which to operate. These may be obtained with a single bus transfer, or if the external Intelligent Signal Averaging (ISA) routine is used, one transfer per sample averaged. In comparison with either the manual or multiple digitisation methods, this is a very favourable characteristic. (The Tektronix signal averaging routine, accessed in DIOS by the SAx command, would be quicker than the ISA routine. Unfortunately, it tends not to be able to cope with the signal frequencies required. This is because of the inherent timing jitter in triggering. Thus averaging should not be carried out within the 7912. Refer to the descriptions of SAx and ISA.)

Secondly, phase data is available. It might be thought at first that this information is redundant, since it is well known that the Hilbert transform relates the phase and amplitude functions for minimum phase systems.[86] However, the deviation from minimum phase is not, in general, easily determined. With the extensive and rather empirical design techniques becoming increasingly utilised in microwave design, it is not easy to control this property in hybrid amplifiers. (Coarse graphs of phase deviation for typical devices are sometimes tendered with specifications.[Appendix D]) In addition, the imperfect matching of amplifiers to transmission line impedances can create non-minimal response because of reflections. The problem of distortion of wideband signals resulting from mismatch has arisen already in the discussion of MSG techniques[Appendix B].

In order to achieve frequency response compensation (deconvolution), as it is intended be applied when using a 7912ADM with its controller, the phase data is required. It is thus noted that the pulse method alone can handle non-minimum phase vertical deflection networks and still provide complex response data for correction algorithms.

4.2.1. Selection of the Pulsed Exciting Function

Several practical considerations favour the use of a step function as the stimulating pulse.

Whereas an impulse settles back to its beginning level, the step function settles at or near the peak stimulating level. This has the advantage that it is more likely that it will be evident if any stages have been driven out of their linear regions. It is of course possible to have driven an interim section into overload due to overshoot in an earlier stage, but this is considerably more likely to be detected with a step than an impulse. The impulse function, on the other hand, can readily conceal such a condition, since it settles back to the quiescent level quickly.

Development of practical high quality voltage step functions is considerably easier, and the subject better researched, than for impulse signals. Indeed, the samplers discussed in Chapter 2 settle for poor quality impulse signals (having low level echoes which these systems can tolerate). These signals are developed from step functions, which are much more readily fabricated.[28,24-27]

Considerable work has been done on the processing, especially the frequency analysis calculation, of step-like functions[87,88,90]. The various approaches have been shown to be essentially equivalent, removing the concern of selecting a processing technique from among those available.[89]

4.2.2. A Practical Step Signal

As indicated in the last section, the techniques for development of step voltage signals are well explored.[24-28] Step signals with equivalent bandwidths[29] more than an order of magnitude above 2GHz are regularly employed, particularly in fast sampling systems. In many applications, however, the precise shape of the pulse (and hence its precise spectral resemblance of the ideal case) is not of special concern.

In sampling oscillographic systems such as Time Domain Reflectometers (TDRs) there is a concern to achieve a clean step response with minimum overshoot. The response of such systems as a function of exciting pulse overshoot has been determined in some detail.[91] It has been shown that the system overshoot is lower than the pulse overshoot, when sampler and pulse have comparable bandwidths. The study assumed overshoot of duration comparable to the risetime, as is most usually observed. It is also shown that the perceived overshoot rapidly diminishes as the pulse bandwidth exceeds that of the sampling system. This is merely confirmation of the instinctive observation that most of the energy in the brief perturbation after the step is contained in frequencies about the equivalent bandwidth of the step, as determined from its risetime. This is noted more formally in Appendix E.

A 25 ps risetime TDR exciting pulse can be used as the test signal. The pulse has effective bandwidth at least seven times above that of the anticipated system. It is engineered for minimum overshoot in a 12.4GHz system. The overshoot can be estimated from the results of [91]. The step amplitude of 250 mV minimum (and 270 mV measured on three systems) is more than sufficient. As such, it represents probably the best available step source for this situation.

The step as measured by an adjacent Tektronix S-6 sampling head[8] is reproduced in Figure 4.5. It is presented in DIOS hardcopy format, using a utility programme written by the author for the 7S12 TDR with custom hardware. This utility (DOSCRO) digitises, averages and filters the 7S12 output and stores the result in RCW format, for later processing by DIOS.

Figure 4.5 is of interest because it shows overshoot and apparent underdamped response. The effect is in fact similar to that responsible for the undulation of the response in Figure 4.1. Rather than actually being present on the 25 ps step, it is induced by the reflection coefficient of the S-6, which was noted to be high in Chapter 2. The TDR system is in fact not specified

JOBNAME DATE **WARNINGS** SCALE FACTORS/INFO

IDENTIFICATION TEXT TIME (PROCESSING HISTORY)

instrument version

Hardcopy output template

TDR Step function as viewed by S-6 sampling head

TEKFREQ 25 7 1983 0 interpolations N/A 500e-12S

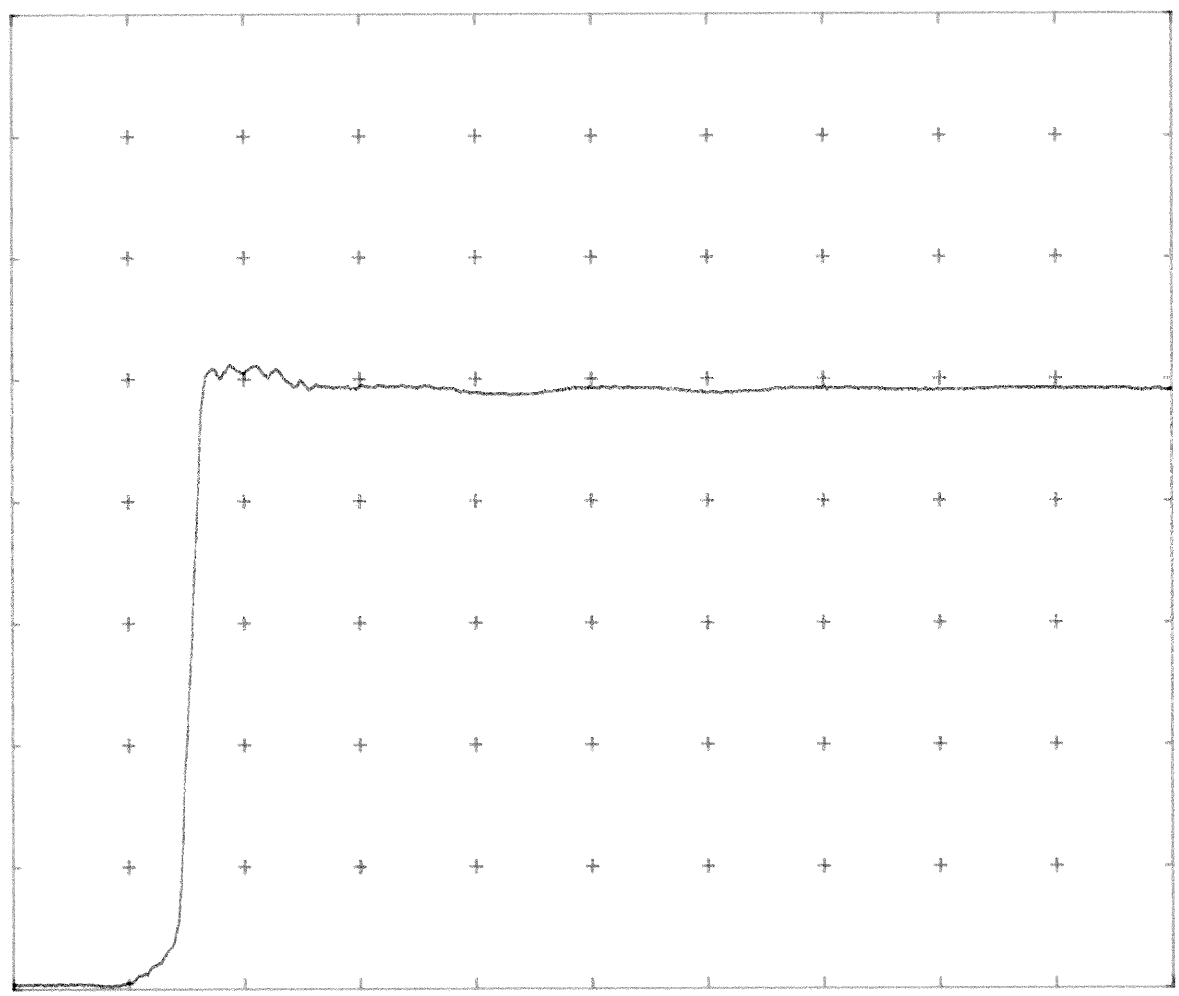

DOSCRO output

version 2.49

FIGURE 4.5

as meeting its quoted accuracy until viewing a transmission line which is some metres distant from the source/sampler pair. This limitation is often overlooked by engineers not familiar with the problems of mismatch in transmission line systems. The aberrations of the step would be virtually invisible in an ideal 2 GHz instrument.

4.2.3. Ideal Performance

The Fourier spectrum of a step function[82] has a spectrum whose magnitude varies inversely with frequency, neglecting the DC term. The question arises as to what resolution will be available from the process of determining the system response by analysis of the digitiser's response to the practical step.

Ideally, the step will be adequately sampled, and will settle to its final value within one half of the least significant bit of vertical resolution within the available time window. Moreover, the step response will be successfully recorded to the limits of the capability of the sampling mechanism. In such a case, the resolution available at any frequency will be defined by the difference between the noise floor and the magnitude of the component at that frequency in the discrete Fourier transform of the exciting function. The noise floor will be defined ultimately only by quantisation.

The continuous Fourier transform of an ideal unit step (ignoring the DC term) is

$$F(w) = 1/jw \qquad (4\text{-}1)$$

Transferring to the discrete case, with N samples in a unit time window, components up to a normalised frequency of N/2-1 will be theoretically resolvable. The n^{th} component will have relative magnitude equal to 1/n, or:

$$\text{Magnitude w.r.t. fundamental} = -20\log(n) \qquad (4\text{-}2)$$

The noise floor on a sample resolved to 9 bits is approximately -54 dB with respect to full scale.[74] (That is, the dynamic range is 54 dB.) When an N point transform is taken, the dynamic range improves by a factor of:[102]

$$20\log(N^{0.5}) = 10\log(N) \qquad (4\text{-}3)$$

which is 27 dB for 512 samples. Hence the noise floor will be approximately 81 dB below a full scale sinewave signal, **in an ideal system with specifications equal to those <u>quoted</u> for the 7912AD.** (The specifcations of the 7912 indicate that 9 bits of "readout" are obtained.)

Thus for any frequency, w, the n^{th} multiple of the reciprocal of the sample window T, the signal to noise ratio in the case of a digitised ideal step of full amplitude will be given by combining

(2) and (3) to obtain:

$$SNR \doteq (81 - 20\log(n))\ \text{dB} \qquad (4-4)$$

For a given SNR there will be a certain probable error present in the determined magnitude of the relevant component. The signal will probably be reported as having an amplitude which differs from its true value by less than the amplitude of the noise signal. Because of the statistical nature of the error there is no fixed rule as to what is "acceptable".

For this case, the 2 GHz component will need to be sufficiently above the noise to allow the response to be determined with "acceptable" certainty at that frequency. With the ideal situation assumed, and a sweep speed of 5 ns/division, this component will be statistically centred 41 dB above the noise.

4.2.4. Factors Limiting Performance

It is no surprise that the theoretical performance is not even approached. A number of practical factors conspire to worsen the resolution of the system. Some of these factors are ameliorable, while others are intractable. These must be accepted by the engineer using an SCT of this type, and the amplifiers which are used to produce a suitably sensitive instrument.

4.2.4.a Magnitude of Step Signal

It has been assumed that the step will be digitised to the **full scale** capacity of the digitiser. In practice some allowance must be made for possible overshoot and fluctuation in magnitude, and in addition, the attenuation available may be in discrete steps. This is because attenuators with the necessary response are usually barrel attenuators. Resolution is typically limited to the nearest 3 dB step, or possibly 1 dB if suitable sets of attenuators are available. In low frequency situations, an allowance (for staying well below FSD throughout a system) of about 70% is typical,[44,45] which is approximately 3 dB. Thus it would seem reasonable to allow for this figure in the practical evaluation.

4.2.4.b Electronic Circuit Noise

The active (and to a degree the passive) components used in developing the step, and amplifying it for application to the deflection plates, will contribute noise.

The 7912ADM is provided with a maximum sensitivity input, offering some 3 mV per division of sensitivity. Reference to the specifications of the front-end amplifier used gives the estimation of its contributed noise as slightly less than 200 microvolts peak-to-peak referred to the input. This is equivalent to a noise signal seven bits below FSD (8 divisions). The 7912ADM is also provided with an input at a higher level.

This offers a sensitivity of about 40 mV p-p, with equivalent noise below the 9 bits of claimed vertical resolution.

Fortunately, random noise effects are reduced by averaging with the use of the ISA routine. When including the low noise plugin amplifier in the measurement, averaging is advisable.

4.2.4.c Forced Oversampling

The earlier allowance of some 27 dB dynamic range improvement made the assumption that there were 512 **independent** samples provided for Fourier transformation. In practice, oversampling has been forced on the SCT target by the problem of spot size compared to pixel size.

It is difficult to estimate the degree of loss of improvement. It is not justified to say that the oversampling itself reduces the independence of the samples, and thus that the improvement will only be that corresponding to the 100 or so samples necessary to describe the available frequency components. However, there is no doubt that the samples are not completely independant. The same image-resolution problem concerned with spot size and pixel size that necessitates oversampling, also ensures that information is shared between columns of pixels. Thus, the degradation will be between zero and 7 dB. An allowance of some 6 dB is considered appropriate.

4.2.4.d Step Edge Resolution

Most AD conversion mechanisms seem to lose resolution as they approach their speed limits.[10,19,43-46] This is the case in the 7612 CRT-based converter[10] which resembles the 7912 SCT. Although there is no specific admission of this, there is much _implicit_ admission in the literature surrounding the 7912.

It is noted in Chapter 5 that much of the information about higher frequency content in a waveform with sharp transitions is contained in the transition regions. These are particularly demanding as far as the SCT recording mechanism is concerned. In the transition, the rate of change of writing speed and thus spot size is greatest, and in the overshoot and settling region, the spot motion is slower, wider, and confined to smaller oscillatory movements. These two conditions (changing spot size and rapid oscillatory behaviour) are most conducive to concealing the true beam path. In addition, a transition is best recorded with high intensity, while precise tracking of the beam in small oscillations is optimally recorded with moderate beam current. There must be a trade-off in the actual current used. This is discussed in detail in section 4.1. Suffice to say that the limitations of blooming at transition ends and loss of spot intensity due to beam current ("slew") limiting in the fast moving centre combine to reduce the certainty of the samples where they are most significant. It is not easy to obtain an estimate of the severity of the reduction.

Test procedures applied to high repetition rate sampling oscilloscopes to determine the number of effective bits obtained when dealing with fast moving signals are mostly not applicable to the 7912. The sinewave fitting method of determining the number of effective bits, selected for use in DIOS, gives the best estimate of the performance. It gives an estimate of the combined deterioration resulting from the effect of blooming and amplifier non-linearity (discussed next) without writing rate limitations, on sine waves. No better technique has been discovered.

This test estimates the 7912ADM resolution at 4.3 bits. Although it is substantially certain that 9 bits of resolution are available to fix the level of a flat section of a waveform, this figure provides a more credible report of the resolution in fast moving parts of the trace.

4.2.4.e Amplifier Non-linearity

The non-linear transfer function from input to scan converter, resulting from the amplifiers used, produces non-linear distortion. This must also be taken into account. The mechanism introduces additional frequency components, or alters existent ones. The severity of the effect is reflected in the number of effective bits obtained. This is taken into account in the allowance used in the last section. The method of measurement is outlined in Chapter 5.

4.2.4.f Conclusion

Taking the above factors into account, the estimate of the ability of the method to reach 2 GHz can be revised. The factors are:

Dynamic range	= 53 dB
Deflection allowance	= -3 dB
"Oversampling" allowance	= -6 dB

thus

Usable dynamic range	= 44 dB

For 5 ns/division sweep, the 2 GHz component is reduced by 40 dB with respect to the fundamental. Hence with a sweep of 5 ns/division it might now be expected that the 2 GHz component be no more than 4 dB above the noise. At 2 ns/division, all other factors being equal, a margin of 12 dB might be expected.

4.2.5. The Measured Step Response

Figure 4.6 is the hardcopy output from DIOS of the step response of the 7912 (version 3.00), obtained using the PLT command. The significance of symbols and labels may be explained by reference to the PLT command in the description of DIOS. The timebase speed is 2 ns/division, or 20 ns for the whole window. The image shown is the result of averaging 32 occurrences.

By way of comparison, figure 4.7 depicts the same waveform as that shown in figure 4.6 with a second waveform superimposed.

STEPV3   16 12 1985 -10 interpolations   NONE   +2.E-09S

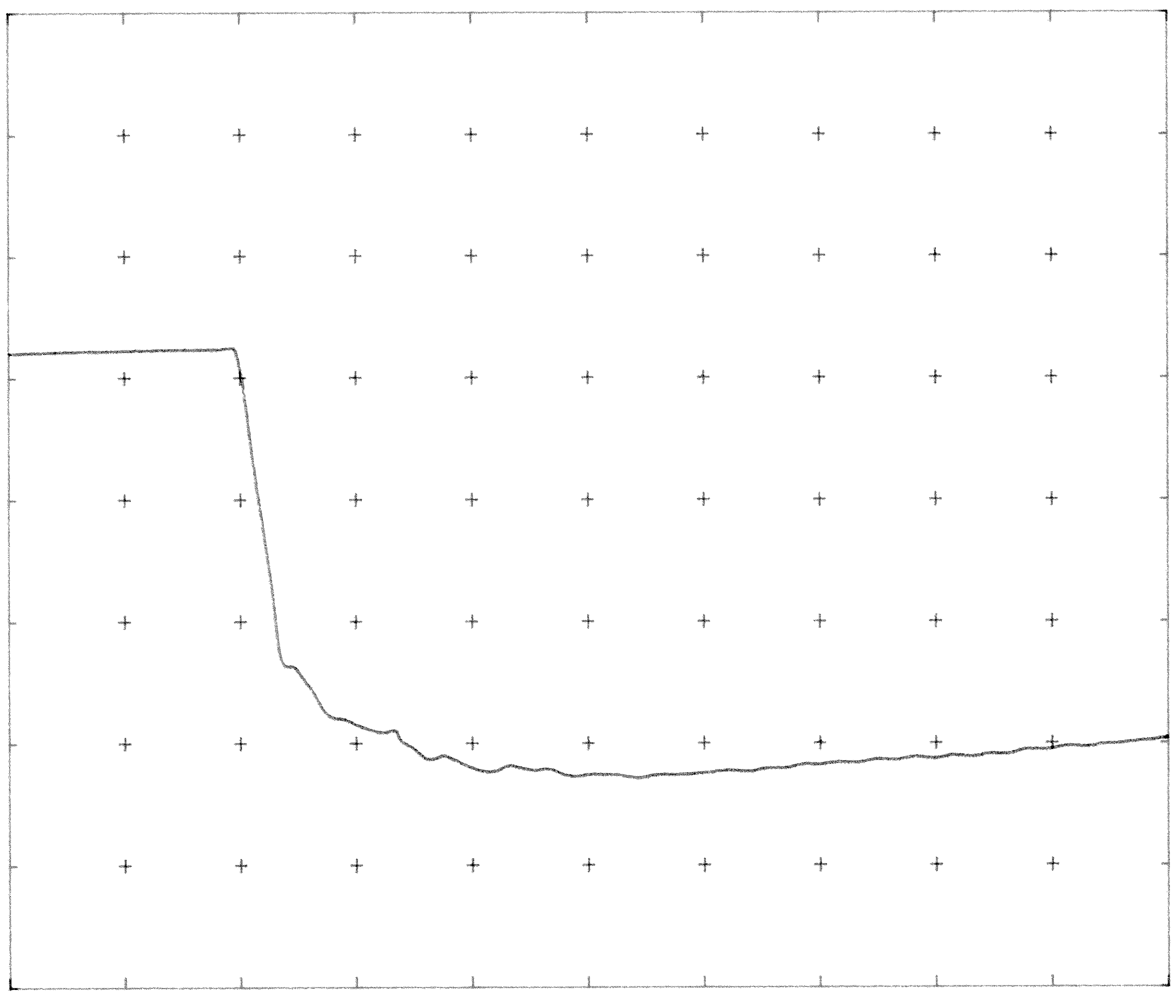

SASW-intelligent signal averaged step response (32 averaged)   14.40hrs

version 3.00

FIGURE 4.6

Averaged and unaveraged responses compared

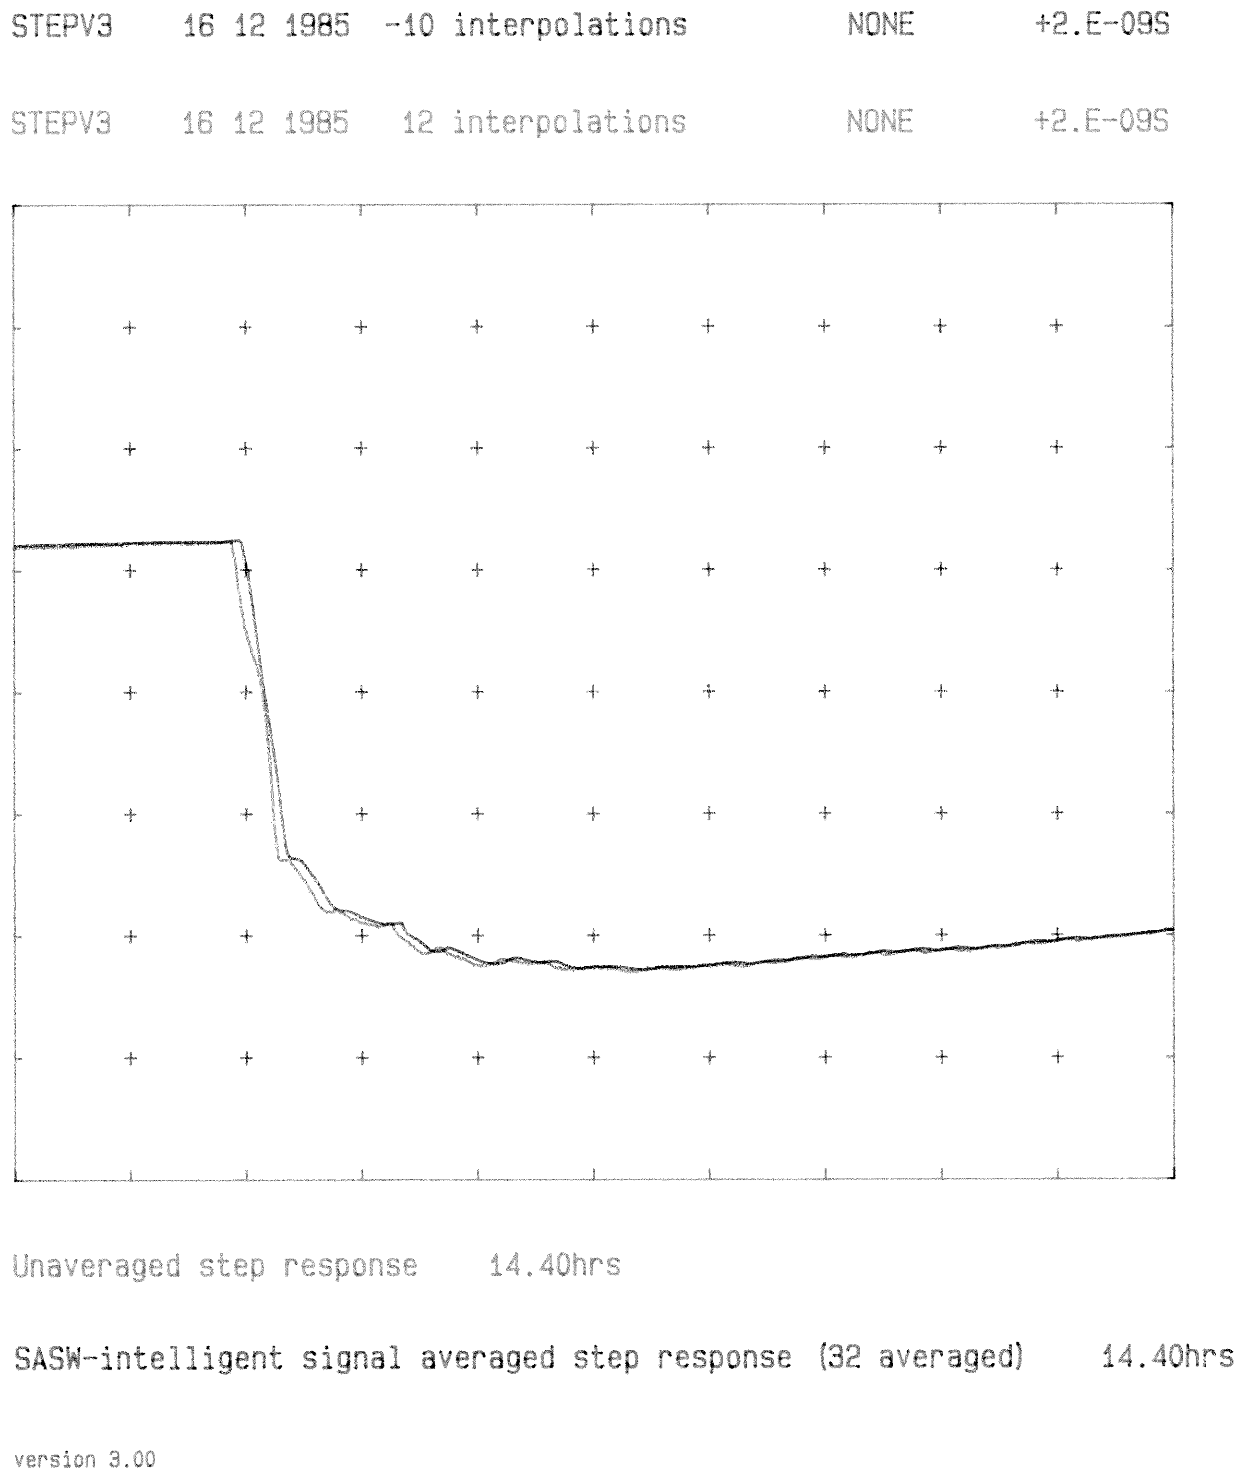

FIGURE 4.7

The second waveform is an unaveraged single step response. Two features are of interest. Firstly, it is displaced in time from the averaged result. This is a manifestation of trigger jitter. The second point to note is that the second trace has small noise perturbations. These are most noticeable in the two divisions after the main edge, where the traces do not overlap. (The clarity of traces is enhanced in originals over copies, since they are of course plotted in suitable colours.)

Figure 4.8 is a plot of an "edge" data set for a trace taken of the step response at maximum intensity. The plot is rather confusing in the region of the step. There the "width" of portions of the beam path is quite large, since it is of course a vertical width on a piece of the path where the beam is substantially travelling vertically. This figure serves to give the reader some idea of the difficulty in obtaining a single value waveform from the available data.

4.2.6. Analysis and Comments

The response is anything but "ideal", both in the frequency domain, as seen in the results in section 4.1, or in the time domain, as seen in section 4.2. The amplitude response vs frequency varies from maximum to minimum by as much as 13 dB across the 2 GHz band. The step response has significant disturbances on the leading edge.

The step response of Figure 4.6 does not exhibit a shape which

A plot of the EDGE arrays of a digitised waveform

STEPV3 16 12 1985 NONE +2.E-09S

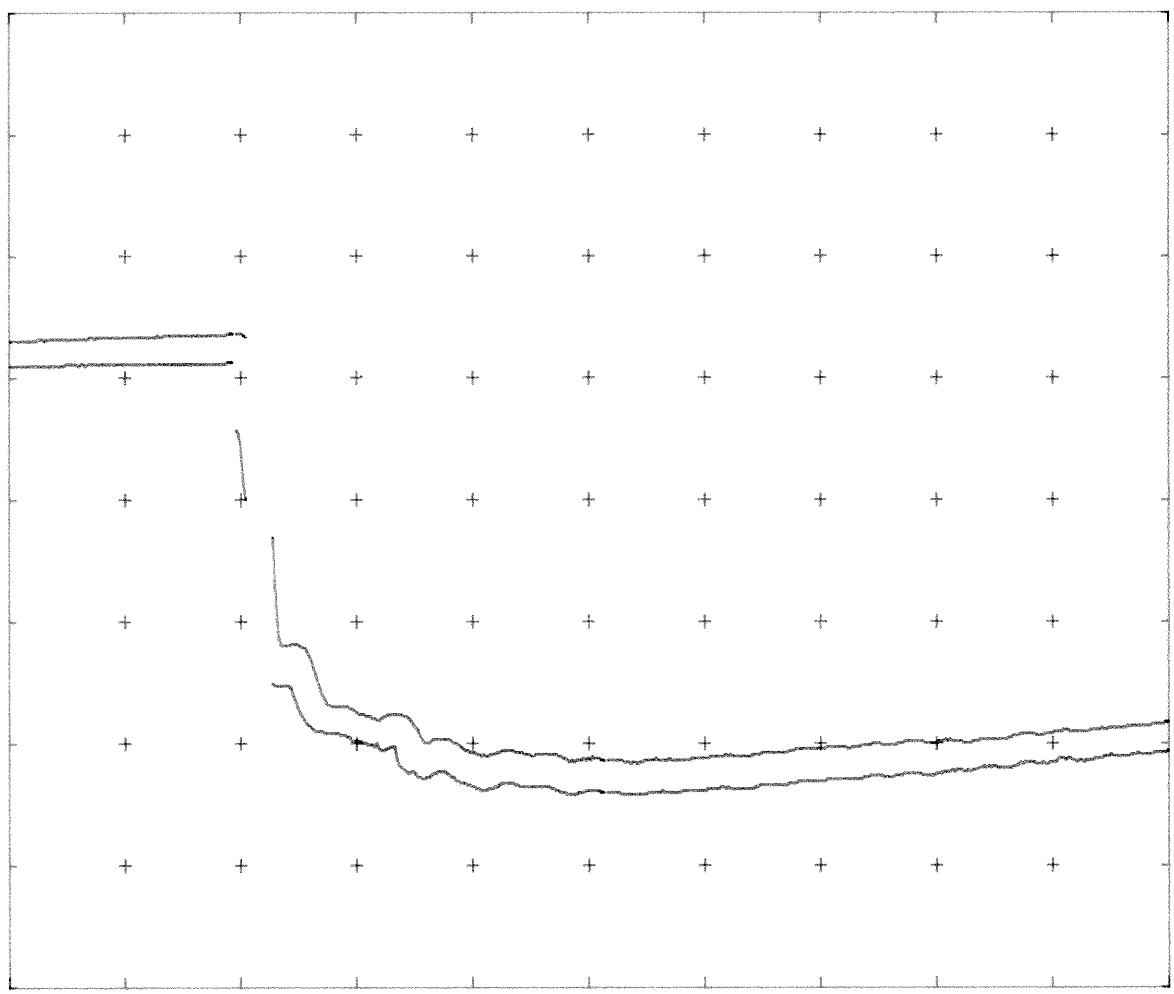

Unaveraged Step Response (Max Intensity) 15.01hrs

version 3.00

FIGURE 4.8

experience might lead one to expect, assuming that the response of the instrument is the result of simple poles and zeroes. The "overshoot" is not that which might be the product of a normal sharp roll-off. (The classic models, such as "second-order-dominant", give rather sinusoidal or Gibb's Phenomenon-like perturbations.) The disturbances appearing here are those which would be expected from internal reflections, and are characterised by delayed signals of vaguely step-like shape. The shape is vague, because the reflection coefficient is a function of frequency, and a differing amount of each component is reflected. Considering Figure 4.6, the main step response appears at the beginning of the second division marker. It has a magnitude of about 2.5 divisions, and a risetime of about 400 ps. It finishes quite sharply, and with the beginning of a flat top, visible momentarily at about 2.4 divisions (2.4d). From 2.45d to about 2.9d appears a small somewhat distorted replica of the main step. Another such replica is discernible with the eye running roughly from 3.3d to 3.7d. Undulations further along, for instance from 4d to 5d, have lost the sharpness of feature which is evident in the first ones.

The delays involved are approximately 1.1 ns and 2.8 ns. The first corresponds to the same length of line as is found between the WJ-A38 drivers and the Final Amplifiers, and which was previously attributed with the 420 MHz spacing of peaks in the scalar response. These results then agree. The 2.8 ns event, which appears less clearly, approximately corresponds with two observations. Firstly, it is not far removed from twice the

1.1 ns event, and secondly, it corresponds to a 180 MHz undulation in the response obtained by the single trace method. (This second periodicity in the scalar figures is much less strongly evident, because of the interaction of several competing periodic functions, but can be taken as manifesting itself not only by some peaks it alone causes, but also by its effects on coincident peaks.) The delay corresponds to the length of line which occurs both between the plugin interface and the WJ-CA35, and the optional delay line insertion point and the first WJ-A38. Thus a multiplicity of smaller potential causes may be taken to account for the blurring of the second replica.

Small sinusoidal undulations appear in the 7d to 9d portion of the step response. These have a frequency of 1.5 GHz. As noted from the scalar result it would seem that the response begins to fall at about this level. These are then attributed to the action of the poles responsible for the ultimate loss of high frequency performance. This closely parallels the usual overshoot phenomenon experienced in classical "second-order-dominant" systems.

The Final Amplifier is designed to deliver slightly more than the signal required for full deflection. In order to deliver this it requires the full minimum guaranteed output of the driver stages. It is this point in the signal chain where the available "headroom" is at a minimum. At frequencies corresponding to the troughs in the response above about 800 MHz, clipping becomes evident at a deflection exceeding 6 divisions. This effect may also be attributed to mismatch at the Final Amplifier input, since at those frequencies where there is subtraction the drivers would be asked to deliver higher output than expected.

It should also be noted that the tail of the response is not flat. This is to be expected from the LF response, which is designed to extend to only 10 MHz. (In fact it falls off below 5 MHz.) The time constant is some hundreds of nanoseconds.

Referring to Figure 4.3, response variations of relatively short periodicity are observed in addition to those evident before. These are attributable to the length of line connecting the DA-31 amplifier in the plugin to the CA-35 amplifier in the mainframe.

Unfortunately, because of the echo-polluted nature of the step responses obtained, it is not possible to simply apply a transform technique and expect the true frequency response to be revealed.

This is exemplified by considering the information returned by a single complex sample of a DFT acting on the step response of Figure 4.6. This sample derives its value from every real input sample. The initial step will contribute information, as will the delayed and distorted replicas. This sample is a single complex number, and can thus represent only a single magnitude and a single phase value. Ideally it alone conveys information about the frequency it represents in the spectrum of the signal. The echoes have arisen because of an arbitrary delay acting on some of the energy originally time coincident with the main, unreflected step. The information about this delay is lost when a single value is derived, yet the delay must affect the result. Hence the transform cannot contain sufficient information to completely describe the system's action on the input signal, nor will it solely represent the linear response of the system, disregarding the introduced echoes.

Only if allowance for the echoes can separately be made, can the frequency response of the system be derived from the step response. Homomorphic signal processing, introduced and discussed in the next chapter, offers the only possibility of achieving this.

Despite this failure, useful information may be inferred from the risetime of the step, which is visible ahead of the mismatch corruptions. The risetime is approximately 250 ps, which corresponds to a "3 dB bandwidth" of 1.4 GHz[29]. There may be some minor reduction of the risetime (smearing of the sharpness of the image) as a result of the signal averaging and trace extraction, and the sharpness of the actual roll-off in the octave above 2 GHz will reduce the risetime compared to that expected for a single or double pole response. This suggests that the bandwidth is indeed in the region of 1.5 GHz.

4.2.6.a Performance in Perspective

This final section explains some reasons for the apparently abnormal response of this instrument. In fact the result obtained is that which is to be expected from commercial state of the art equipment. It is the initially confusing nature of plain signals, such as a step, which probably prevents companies such as Tektronix from producing instruments with this bandwidth. Only a very large scale concerted effort of research and development, producing customised circuits, would improve upon the situation.

The results may be put in perspective by considering both the performance of lower frequency instruments and of amplifiers working up to similar frequencies. Appendix A noted that a low frequency oscilloscope will typically have a magnitude accuracy specification of the order of 2%. It will also have a bandwidth limitation, usually specified by the minus 3 dB point. The accuracy specification is intended to apply to the flat mid-band portion of the response curve, which is usually very wide. The 2% error permitted represents less than 0.2 dB.

A typical state-of-the-art high frequency amplifier with multi-octave bandwidth, as might be offered by companies such as Watkins-Johnson or Avantek[60] or as reported in the literature[61] will offer a "gain flatness" of $\pm$ 0.5 dB to $\pm$ 1.5 dB in a matched 50 Ohm system. (These figures represent errors of 6% to 20%.) Less tightly specified circuits offer considerably wider limits. In addition, input matching performance is variable over the specified bandwidth. Worst case SWRs of above 2.0:1 are typical.[Appendix D] (This represents a reflection coefficient of one third.) The cascading of several devices of this order of performance produces a response which is likely to fluctuate several dB over the specified band, and which will exhibit reflections of the order of magnitude of 0.33^2, or roughly 10%.

Two factors explain the differences in performance. Wideband gigahertz amplifiers represent the state-of-the-art in wideband signal processing technology. Whereas the wideband amplifiers used in oscilloscopes use only a simple degree of response compensation in order to achieve their bandwidths,[48] the multi-octave gigahertz amplifiers in the current literature use one of several complicated approaches to render an otherwise very variable response tolerably flat.[61] Obtaining a response that is "flat" to better than 1 dB or so is unreasonably difficult over the bandwidths addressed by microwave amplifiers. The amplifying elements themselves are induced to provide level response only by the introduction of numerous compensating elements into a feedback network. Allowance must be made for the parasitics which are inevitably present in components of finite physical dimensions, because these parasitics are significant at gigahertz frequencies. The first explanation of the different response performance is simply that it is impossible or uneconomical to produce a flat response at frequencies where physically unavoidable parasitics are significant without feedback or other compensating mechanisms to offset degradation.

The second reason for the performance difference between microwave wideband circuits and oscillographic instrument circuits is that there has been less motive to refine microwave circuits. The offered gain flatness is most adequate for current applications of these devices. Precise oscillographic instruments, where good match and impeccable phase response is needed for immediate optical fidelity, have not yet begun to be used at those frequencies. It is only with the advance in wideband systems stretching to those frequencies that the need is being felt. Gigabit data rate circuits may alter this circumstance in the future, however.

5. POST-PROCESSING ALGORITHMS

5.1. Introduction

This chapter introduces the various techniques which have been investigated for numerically processing the data acquired by the 7912ADM, and comments upon their effectiveness and application. Before proceeding to a discussion of these signal processing algorithms, the categories into which they fall are delineated.

5.1.1. Trace Determination Algorithms

Trace determination (or extraction) algorithms are basically processing methods which derive an array of data from the image which is written to the SCT target. They may be regarded as trying to overcome some of the inherent limitations of the SCT itself, and attempting to convert the data into a conventional and "friendly" form.

A number of such algorithms have been developed by the author. Some are original, answering a need of the application, while others supplant unsuitable routines already available in the 7912AD firmware. They are a form of digital image processing. They are invoked by calling the appropriate routines in the Digitising Instrument Operating System (DIOS). Each is described in detail under its own heading in the separate DIOS command manual included with the appendices. The trace extraction and related commands which are original or of special interest are

ABP, ENV, ISA, and ITC; descriptions of these commands are included in this chapter. Some information provided here is duplicated in the operator's command manual.

5.1.2. Correction or Compensation Algorithms

These are digital signal processing procedures which attempt to reverse degradation and distortion of the signal which has occurred in the analog signal processing chain of the instrument. They operate on a single valued (ATC) data array. They can operate without a priori knowledge of the signal. (To require otherwise would clearly be unacceptable in a general purpose system.) However, for such routines to return data which results in improved measurements with best efficiency, the operator must appreciate their effect and select and invoke them appropriately. This is particularly the case with windowing before Fourier transformation, data conditioning before cepstral analysis, and echo signal cancelling after cepstral analysis. Routines in this category use a knowledge of the system itself (learned from "calibration" of the routines with the system).

Compensation algorithms which have been implemented fall into three groups: Time domain algorithms; Frequency domain algorithms, which include routines utilising the discrete Fourier transform (DFT or FFT) and effect deconvolution[73,82,94]; and "Quefrency" domain algorithms, which are routines utilising homomorphic signal processing (Cepstral processing)[96]. Relevant theory is dealt with in this chapter. The related commands are ACY, ACN and DEB in the first group, CAL, DCW, FFT, HAN, IFT, LHN, LRx, LSx, RHN, RRx, RSx, SPM, SPT and TWx in the second group, and CEP, ICP, PAC, RSE and SEP in the third. The interested reader is referred to the operator's manual for specific descriptions of these commands.

5.2. Trace Determination Algorithms

As originally noted in the discussion of the SCT in section 2.3, the information obtained after a digitisation in a 7912 is a (coded) array of one-bit numbers in a 512-square matrix. This must be converted to a single-valued, or occasionally double-valued, data array. The physical processes involved in the writing operation can make this conversion difficult, or produce low confidence in the result. The algorithms below have been used to assist in this process.

5.2.1. Interpolation and Trace Centering

The Interpolation to Trace Centre routine replaces the 7912 ATC routine. It executes an algorithm similar but slightly superior to the one found in the digitiser firmware. It takes in the double-valued data array called an "EDGE" array. (This contains two single-valued functions which represent the upper and lower bounds of the river of pixels "hit" by the writing beam - the edges of the beam where they can unambiguously be determined. See the EDGe command in the operators manual or reference 17.) A description of the algorithm follows.

The routine executes five steps. Firstly it searches the two input functions for missing data. Any horizontal address missing either an upper or lower limit or with limits reversed or out of range is flagged as absent. Secondly, the simple sum of valid pairs is formed. (These two steps are usually executed conveniently in one move.) If all points are invalid, the routine stops.

Thirdly, the routine looks to see if the end samples are valid. If they are not, the endmost valid sample is copied to the end, in both directions. This differs from the 7912 routine, which can leave invalid samples (as different from zeroes) in the data. This may affect further routines. The copy count is included when determining the maximum number of iterations, though no interpolation is undertaken.

Fourthly all gaps are filled with interpolated sample points. The slopes at the two ends of the intervals are determined, and these are then used to fill the gap by **cubic curve fitting.** This contrasts with the 7912 routine which uses linear interpolation only, which affects the waveform more adversely. In practice the cubically interpolated waveform appears more natural, whereas the linear one is visibly discontinuous when more than 3 or so samples have been interpolated. A cubic interpolation of over ten steps often goes unnoticed. This is especially true when interpolating over a step which has been lost. (This is the most common location of a momentary beam loss, since the writing rate is changing most rapidly there.) The cubic approach has been found to be misleading, however, when the waveform is noisy or patchy near the end-points of the gap, as the two end-points are used to determine a slope at the point of beam loss, and this is corrupted by the final point being out of place.

Note that if the gap has only one endpoint available at either end (that is there are two gaps separated by a small valid area), the cubic fit must be aborted in favour of a linear one. If the gap is only one point wide, linear interpolation is used directly. If the single endpoint was the last or first valid point, the slope is automatically set to 0.

The method of determining the endslopes could have been made more complex. It has been observed that increasing the order of the fit beyond cubic opens the possibility for erratic fits (as is a well known phenomena in curve fitting to an arbitrary order on the same arbitrary number of points). However, keeping the main fit at cubic, the endslopes can be determined by taking a weighted average of the slopes suggested by three or more points at the gap end. Such efforts have been tried, and have been found not to be rewarding.

The largest gap size is accounted for in the determination of the largest number of points calculated in a single interpolation.

5.2.2. Actual Beam Path location

This command instructs the controller to determine the Actual Beam Path from edge data. Whereas ITC determines each of the samples without reference to information in the adjacent sample positions, except in the case of a missing sample, ABP is a two dimensionally conscious routine. It also acts on edge arrays.

Because ABP makes maximum but safe use of the data available in edge arrays, the extraction of the edge arrays from the raw data should be carried out with lax spot-width and spot-delta-width constraints. This will ensure that little useful data is eliminated in that step.

ABP takes the form of a preprocessing routine which falls through to the ITC routine. It modifies the edge data in each location to reflect information gathered from nearby locations, then allows ITC to convert the edge data to a single valued trace, interpolating and extending as required to force a valid sample value at each horizontal location.

Two mechanisms which act in a Scan Conversion Tube combine to produce potentially misleading images. ABP is an attempt to compensate to some degree for these mechanisms. The first problem arises from the variation of spot size with writing rate. When the speed of the beam across the target rapidly reduces, as is the case at the end of a step transition, the spot may widen from only a few pixels (or no written pixels at all in places) to tens of pixels. As the spot widens it can flood electrons backwards into columns of pixels over which it has already passed, and into which it has yet to proceed, polluting the datum associated with each of those columns. The second problem is a similar corruption of data in nearby columns caused by the spot retracing its path before it has moved one spot-width further along the target. This may result in the trailing part of the spot overwriting those pixels which its leading edge wrote moments before, or writing pixels previously not hit. By way of example, this occurs where a sinewave oscillation is written with too low a sweep speed to separate the adjacent peaks of the wave.

In both these cases, the output of an ATC command can be very misleading. The decaying sinewave oscillation may be portrayed with severely reduced amplitude, rejected altogether by an EDG algorithm with too strict width or delta-width constraints, or converted into a flat region with irregular perturbations. The end of a step transition can acquire or lose fine detail in the "smudge" generated by the sudden spot width increase. The risk of these happenings is mentioned in the 7912 manuals, but little is provided by way of a solution.

ABP operates by a process of circle fitting inside the confines of the lines which define the trace edges. The call results in three passes being made over the data arrays. In the first pass, an iterative honing process is used to determine the centre and radius of the largest circle which can be fitted over written pixels in the image, its centre remaining on a fixed vertical line. This process is repeated for each vertical column of pixels, the circle being centred in the middle of that column. (The assumption is made that the beam has been adjusted to produce a circular spot on the target. This is part of the normal 7912 adjustment procedure.) In the current version, the centre position is resolved to the nearest 0.25 vertical units, where one unit is equal to one pixel. This value is chosen for convenience.

The second pass rejects the values of any column all of whose written pixels are covered by circles of other columns. (It is necessary to resolve to sub-pixel levels to prevent numerical rounding causing rejection of otherwise valid columns.) Acceptable columns are modified to leave the data in edge form, but in a fashion that will allow the normal ITC routine to retain the improved resolution which the circle fitting technique yields. That is, the numbers are stored with fractional parts.

The third pass consists simply of falling down to the ITC routine. That routine cannot tell that the data has been processed, since it handles rejected data (the data is made negative, as is the Tektronix convention) and operates in floating point rather than in integer arithmetic.

(Two observations arising from experience with the problems addressed by ABP need to be made here. Initially, it is significant that the image aberrations will not be avoidable with any scan conversion system. This was not taken into account before embarking on the project. They may be resolved (as has mostly been the case in the author's experience) by suitable alteration of timebase speed and triggering conditions, but this is only practical when the event being observed is "reasonably recurrent", that is it is not a true single or very rare event, which is just the type of occurrence with which scan conversion alone is equipped to deal.

A second observation is that the author has not been able to devise an image processing technique which will deliver results which can be shown to be optimal. That the ABP routine is an advance over the plain ATC method is sometimes evident to a user's eye: the single valued waveshape returned by ATC alone lacks portions which an expansion of the waveform exposes. The use of ABP will at least warn the user of the uncertainty by rejecting the suspect portion (and giving an interpolation warning), and may resolve it all or in part. It has been tested by the author on data numerically concocted by simulating the spot width variation. In this rather artificial case, ABP reveals reduced image distortion. However, since the simulation assumes all and only the aberration mechanisms ABP tries to deal with, it might be expected to do just that. ABP also often rejects pieces of the waveshape which might have been resolved by some superior approach. The lack of theoretical support for the circle-fitting algorithm must bother an engineer using it. Nevertheless, no solution for this dilemma has been found.)

5.2.3. Intelligent Signal Averaging

ISA, the improved signal averaging routine incorporated in DIOS, supersedes the Tektronix firmware routine SA. The mnemonic is used because the routine exercises simple judgement in the determination of the signal average.

ISA requires that there be a single valued waveform (ATC or better, an ABP waveform) in the controller for enhancement. In addition, it also requires an integral iteration count to be supplied.

ISA repeatedly calls SSW, EDG and ITC. This command sequence loads the computer with a new waveform data set. When the latest single valued waveform has been determined, it is compared to the original or partially averaged one. The latest waveform is shifted (up to plus and minus one division) relative to the other, and the minimum mean square error between aligned samples is sought. (The mechanism is similar to the DEB routine described above, but is not free to scale the waveforms.) If this error is acceptable, the waveshapes are deemed to be sufficiently identical and optimally aligned, and the last one is included in the running average. If the error is excessive, the routine reports that the waveshapes are uncorrelatable, and the last captured one is rejected.

There are two situations where this method will allow signal averaging but where the original 7912 firmware one will not.

This method controls trigger jitter. Each waveform is shifted (and extended) as required to compensate in software for the hardware uncertainty of the trigger. The triggering circuits of the 7912 have been found to be almost completely useless for signal averaging by the normal method with sweep speeds in excess of 10nS per division. Secondly, false triggerings cannot pollute the averaging process. Particularly in the very noisy environment where signal averaging is required, false triggering or triggering on events other than the one it is required to analyse is prevalent. A waveform which does not resemble the one used for enhancement will be rejected, so an automatic, correct-triggering filter is implemented after the event, but before the averaging has taken place.

Certain algorithm parameters have been selected for optimum operation of the process. The 'jitter scan width' of one division has been found to be more than adequate, though a little costly in terms of computation time, because of the number of iterations during each of which a mean square error between data sets must be found. The need to make the jitter scan wider with increasing timebase speed has not been felt to be sufficiently important, but could be easily incorporated if necessary. It would reduce the delay in some circumstances.

The mean square error limit is not as critical as might be expected. Basically, waveforms either can be shifted to make a good match, or they are corrupted in some way and will have very large mean square errors. A value corresponding to an error of five pixels in each vertical searched is satisfactory.

There is no need to sample every vertical in calculating the error. In theory only every second needs to be sampled if the waveform is oversampled by a factor of two, etc. In practice sampling every five irrespective of the oversampling ratio gives a satisfactory compromise of speed of execution and seems not to be too infrequent at any level.

ISA continues searching for the required number of acceptable waveforms (given by the iteration count) and so will loop indefinitely if the events stop or the signal is lost, etc. (Provision for manual abortion of the search is provided in DIOS.)

5.2.4. Envelope Acquisition

The Envelope (ENV) routine causes an envelope waveshape to be determined from the raw coded data. The routine returns a double valued set of data similar to EDGe data.

The main difference between ENV and EDG determination algorithms is that EDG tends to reject data points which imply a beamwidth which is impossible. A trace deliberately written with a slow timebase is not normally available except in coded format. ENV will return an effective "edge" data set; that is, a pair of continuous (where there is any data recorded at all) single valued arrays.

5.3. Time Domain Correction

Practical amplifiers exhibit non-linear transfer functions. Similarly to the case of compensation of the spectrum of a signal for spectral distortion, correction may be applied where the transfer function is known. The technique of "pre-distortion" is a familiar analog example used with transconductance multipliers implemented with BJTs.[51] When the signal is stored as numbers in a computer it may be "post-undistorted" in software. In addition the correction function may be derived from the result of processing a known signal under software control. This allows the system to be quickly adapted to change.

For such a mechanism to be implemented with any degree of ease, the form of the signal used as a reference must be carefully selected. It must be readily available and must have a lower level of distortion than that which it is desired to realise in the end product signal. Fortunately the problem of sensing distortions in a digitising instrument has been addressed in recent literature.[42-46]

A sinewave signal is observed to be the most suitable for the task[45]. This signal excites at a single point in the spectrum, which obviates the kind of problems addressed with the other correction signals and processes discussed in the literature cited above. Sine waves of good harmonic purity are not difficult to generate. The technique for digitiser testing which is presented in [45] employs deliberate aliasing to make the test more exacting. Although not used with intentional aliasing, because of the impossibility of doing this with an SCT instrument, the technique is found to give a good measure of the number of effective bits available.

The number of effective bits of accuracy delivered by the conversion process is determined by fitting a mathematically very accurate sinewave curve against that obtained by the digitiser. The mathematically derived curve is optimised in four parameters - frequency, amplitude, phase and offset - to give minimum mean square difference between the actual and theoretical data sets. The mean square error gives a single figure of merit. This figure is returned by the Determine Effective Bits (DEB) function in DIOS.

In the determination process described above, a theoretical and the measured curve are simultaneously available. The difference at each value is thus also available. These differences plotted against expected value (the theoretical curve) represent the error function introduced by the non-linearity of the digitiser deflection system.

A correction function of any form might be fitted to these data. In practice the distortion has been found to be mild with the circuits employed. A linear correction is elementary to determine and of quite acceptable utility, and has been used for convenience.

(There is no reason why higher orders could not be used should electronics in the signal chain demand this increased attention. However, as anyone who has blindly used matrix arithmetic to obtain a pure polynomial fit to a set of points will realise, there is a strong tendency for such fittings to produce a curve which does fit all provided data, but which swings wildly to enormous extremes between fixed data. Some care must be exercised in selecting the order of the model, and in determining the coefficients. An intelligent model can be visualised by considering the effects which give rise to the distortion. Transistors exhibit an exponential transfer function. This gives rise primarily to second and third order harmonic distortions. The final amplifier operates push pull; as in the case of a differential pair, this will tend to cancel the even harmonics. A good model to fit for approximating the transfer function would reflect this tendency, and might be cubic. It would need to use sufficient points typically from 12 to as much as 100, which is the number of independent data which it is estimated that the digitiser delivers.)

DIOS performs the linear fit to the error function whenever DEB is executed. The parameters are stored for future use.

When the function is activated, data read in over the bus are corrected according to the stored data. The determined linearity is approximately 4.3 bits without correction. This is increased to almost 5 bits with correction.

5.4. Frequency Domain Compensation

Where the distortions are those introduced by passing a wide band signal through a system whose magnitude response is not smooth, and the response of the system is known, it is theoretically possible to reverse the distortions. The process is referred to as deconvolution, because it is the reverse of the effect of convolving the system impulse response with the desired original signal. It ideally reveals a waveform which more accurately represents the original signal.

Having the linearly corrupted waveshape, it may be corrected for imperfect system response by moving to the frequency domain, and adjusting appropriately the sample values by dividing each by the complex magnitude of the response at that frequency. If a frequency domain parameter is sought it is then available; for a time domain property to be observed, an inverse transformation may be applied. The basic scheme is illustrated in Figure 5.1. The transformations are necessary, employing the equivalence of convolution in one domain to multiplication in the other, because there is no known direct inverse convolution function.[94]

It is informative to briefly note the effects of the phase part of the response of a system on certain time domain parameters. Where a measurement seeks to determine the presence or magnitude of a frequency component, a phase deviation may not be significant. However, properties such as risetime and overshoot (ringing) of a step transition are commonly sought in the time

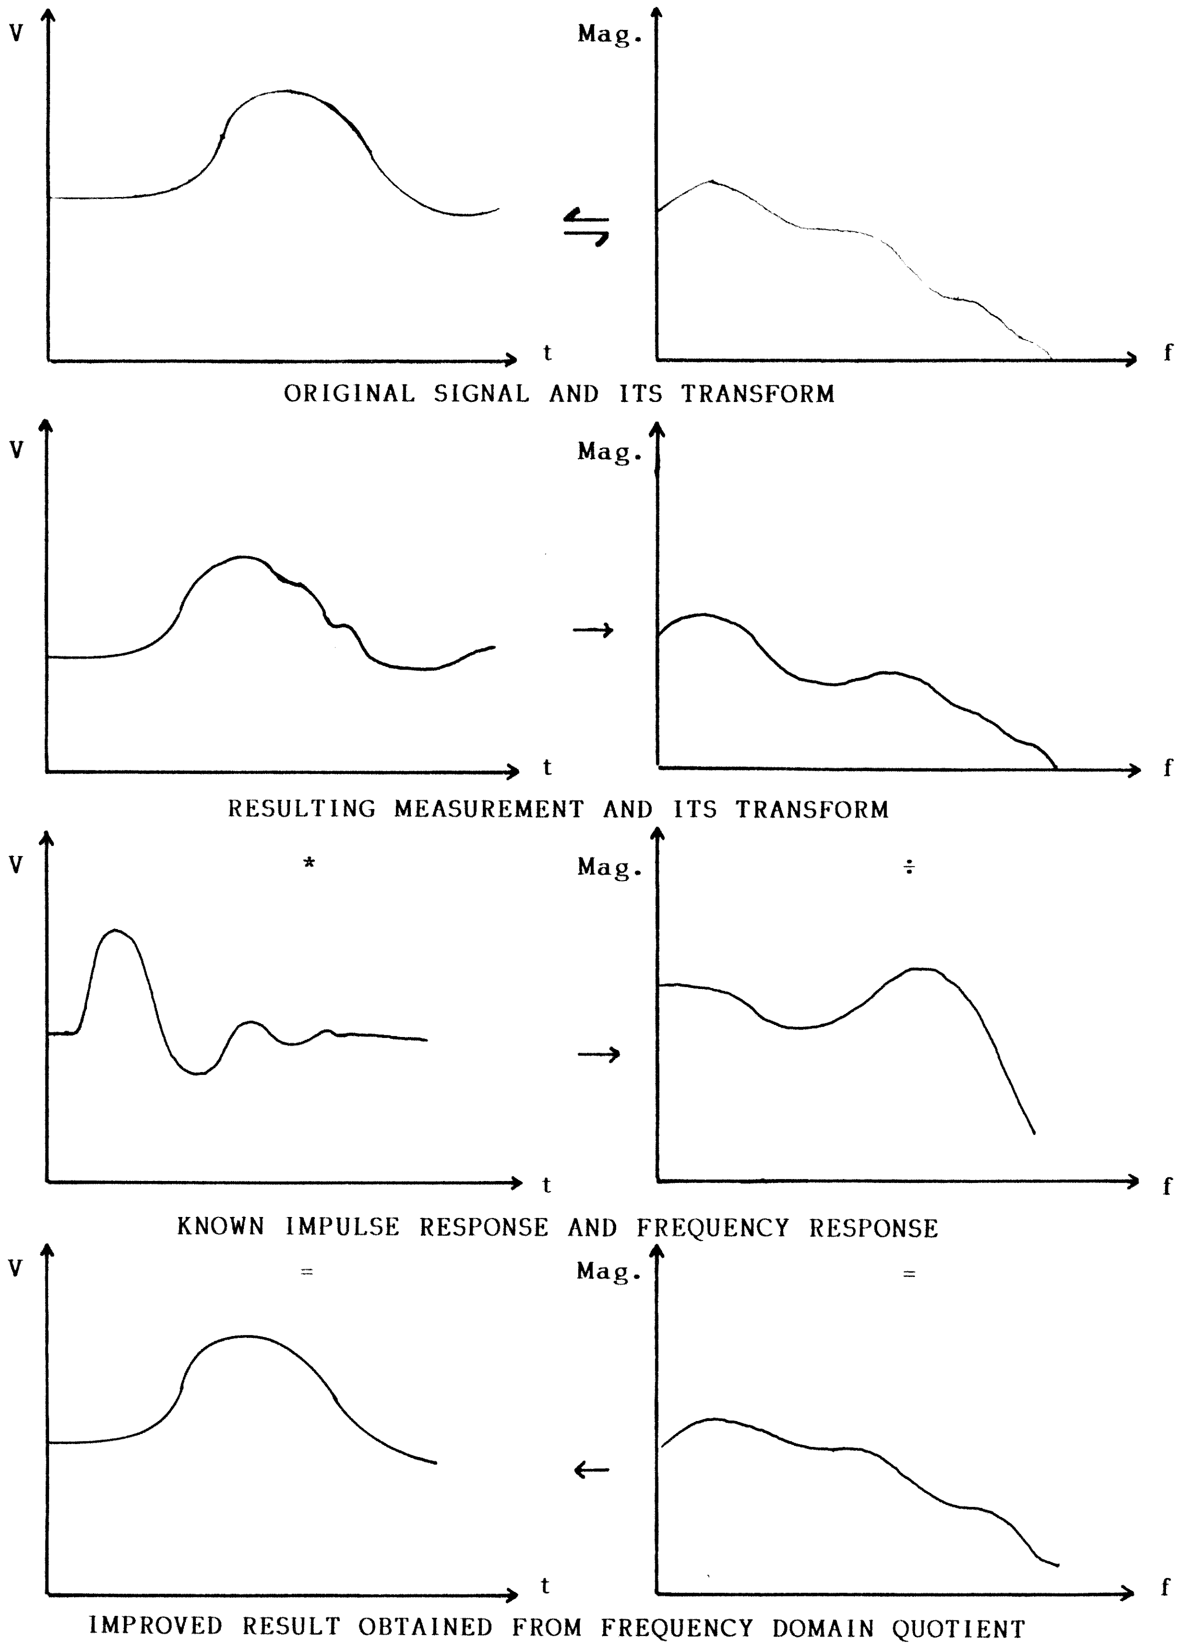

FIGURE 5.1: EXAMPLE OF DECONVOLUTION

domain. These are quite sensitive to the relative phases of components adding to form the sum. The adherence to linear phase in a system used for such time domain measurements assumes importance. This fact, and the consequential desire to be aware of, and to correct for, phase disturbances, makes it necessary to have the complex response of the system.

5.4.1. Use of Windows

Although theoretically straightforward, the practical process of transforming an arbitrary waveform from one domain to the other is not so simple. The fundamental pitfalls have been well documented.[92,93] The most important here is referred to as leakage. It is conveniently viewed as pollution of the results of a discrete transformation arising from the implicit assumption that the signal being analysed repeats with the period of the data set. If the signal truly does have a period equal to the interval over which it has been sampled, no problem arises. However, this is not generally the case. If the signal finishes at a level different to that at which it started, the assumed-replicated signal has a step where it wraps around. This step contributes components to the transform which are not present in the actual signal.

The usual method of dealing with this possibility is to multiply the data set by a "window" function. This tapers the waveform so that no sharp discontinuity arises when it is repeated end to end. The multiplication in the time domain corresponds to a convolution in the frequency domain. The transformed result is "smeared" by convolution with the transform of the window used. A great range of windows is in use, and their transforms and effects are readily summarised.[95]

In certain circumstances, the leakage can be analytically removed.[101] The circumstances, however, tend to be rather specific. The requirements in [101] are essentially a priori assumption that the signal consisted of discrete pure tones. From that assumption it follows that the precise nature of the tones can be distilled from the leakage altered transform data. Such methods are not nearly sufficiently general for use here.

No single window suits all applications. The window for a particular waveform is selected on the basis of precisely what is required in the result, and how the information is distributed within the original time window. Factors which may be considered include, in the frequency domain, the amplitude accuracy of samples, the ability to resolve closely spaced peaks, and the desire to resolve smaller peaks away from large ones. In the time domain, such factors as whether critical transitions of the waveform are located near the interval ends are important.

The above notwithstanding, instruments usually provide a small finite selection of useful windows from which the operator selects in the light of his requirements. Some windows prove to be more generally useful than others. The Hanning and Tukey window types have been found most satisfactory, and are built into DIOS. In addition, another window type may be calculated and stored in a disk file for immediate use. This defaults to the Dolph-Chebychev window, which exhibits minimal bandwidth at the expense of there being no sidelobe falloff. The sidelobe level is selected to be suitable for use with the fixed dynamic range of the Scan Converter output. Other windows can be developed[100] which offer maximal amplitude accuracy. The 'flat top' window used in some Hewlett-Packard instruments is specifically designed for use in FFT based spectrum analysers.

(It is observed that in general, these instruments have means by which the bandwidth can be almost arbitrarily narrowed, in exchange for, say, speed of processing. In this case, a window which gives optimal amplitude accuracy makes sense, since speed can be traded for bandwidth, to make up for the shortcomings of the window in this area. For the case here, where the sample number is fixed, this window is of considerably less appeal, as the resolution in frequency is compromised in comparison to windows of general appeal. This is the reason that it is mentioned only as an alternative.)

The user will in general have to apply a window before taking the FFT of a data array. It is not the intention here to delve into a discussion of window selection as there are good texts available on the process. However, most neglect to adequately point out the sensitivity of the result to the position of sharp transitions in the time domain record. Where there is a step or pulse in the record, it is found that most of the information is contained in that part of the waveshape. If the pulse or step is near the edge of the record, a window carelessly applied to the samples can attenuate the important part of the waveshape needlessly. The Tukey window is particularly useful in that it can be specified to cover wider or narrower portions of the sample interval.

In addition to the Tukey window, DIOS provides the ability to window only one side of the data. This forces the other extreme of the record to converge to the level of the first, in case the information-carrying part of the signal is at the first end of the record. Commands for rotating and shifting the data in the record to allow more central placing of certain features are also provided.

(Such functions are not available on more "conventional" digitisers, because movements of the data samples within the record destroy the trigger-relative timing of events. This is of less concern at higher frequencies, since the whole process of triggering becomes more uncertain. When working in microwave systems, separate triggering is often employed, which makes the instant at which a record started rather arbitrary anyway.)

5.4.2. Response Smoothing (Deconvolution)

DIOS can store a record of the measured response of the 7912ADM. When it is desired to refine a set of samples as far as possible, the following procedure is used to produce that result which would have been obtained if a system with flatter response had been used. This process of deconvolution thus appears to the user as a built-in mechanism which is "flattening" the effective response of his instrument.

Initially, the waveform is viewed. If it begins and ends at the same (or a similar) level, windowing may not be necessary. (This might be checked using a cursor provided by the CUR function, or by rotating data. Once the ends of the record have been viewed adjacent to each other, the record can be rotated back.) A minor discrepancy can be dealt with by using a Tukey window which acts only on the very ends of the record, or a one-sided Hanning "half-window". The latter is especially useful for dealing with step waveforms where the step is closer to one end of the record.

If there are significant numbers of adjacent interpolated points, or the waveform is oversampled, it may be packed down (using PAC, for instance), before being windowed. This ploy is equivalent to extending the record to twice its length, allowing the window transformed to cover a longer time interval, regarded from the time scale of the original samples. Since the very nature of the scan converter produces oversampling, this technique can be quite useful. It is vital when transforming to the quefrency domain, as will be noted below with the discussion of that technique.

After suitable preparation, the FFT is taken. Then the compensation is effected. Each sample in the frequency domain is moved in magnitude and phase in the direction required to restore it to that value it would have had if the machine had been ideal. The adjustment required to achieve this is readily derived from a knowledge of the instrument's response, previously determined, and corresponds to the division of the signal spectrum by the instrument's response. A correction array is stored, and interpolation is used where the samples do not coincide.

Several practical matters must be considered in association with the compensation (or deconvolution) process as implemented for the 7912ADM.

The compensation is effected by the command SPM in DIOS. The method of interpolation used is classical Lagrangean interpolation. Although 4^{th} order is probably adequate for the interpolation, 6^{th} order has been used. The order was selected empirically. The suitability of the method is assessed in the light of its digital signal processing interpretation.[103] Improvement by the techniques of linear programming suggested in [103] has not been seen as justified because it is very unlikely that the interpolation will involve functions sampled anywhere near their Nyquist rate. This circumstance reduces the advantages of the improved techniques.

It is of course necessary to assume that the system response is not changing more rapidly with frequency than can be represented by the samples obtained from the FFT of the waveform to be processed. Correction could be meaningless in this case. It is associated with the ideas of "resolvable feature size" discussed in Chapter 4. Figure 5.2 exemplifies the situation. It is equivalent to saying that the time window is assumed to be of sufficient length, or that the frequency domain response function is sampled at an adequate rate. This can only be guaranteed if the waveform "settles" within the interval sampled, is not significantly erased by later application of shaped windows, and the interval can be extended arbitrarily. The idea of "settling" is not precisely a correct description. It refers to the effect only with respect to a pulse or step situation; nevertheless, it serves to give some feel for the actual situation.

In the vicinity of zeroes of the system response, there may be considerable error (noise) introduced. This is merely a reflection of the fact that one would be trying to divide by zero at some point in the array quotient taken for compensation. (If the system does not pass a particular frequency, all data about that frequency will be lost from the signal as it passes through the system.) In general, this situation should be avoided in a general purpose instrument. However, a similar effect arises where the response of the system is measured with less certainty, i.e. towards the higher frequencies. (Recall the discussion in Chapter 4 of the pulse response determination method and the usable measurement range of the technique used.) Recently, work has been done on refining the deconvolution process with the aim of minimising the noise in the vicinity of response zeroes.[94,104] However, these (frequency domain iterative) techniques basically consist of interpolation processes, and thus require data on both sides of the patch of small magnitude

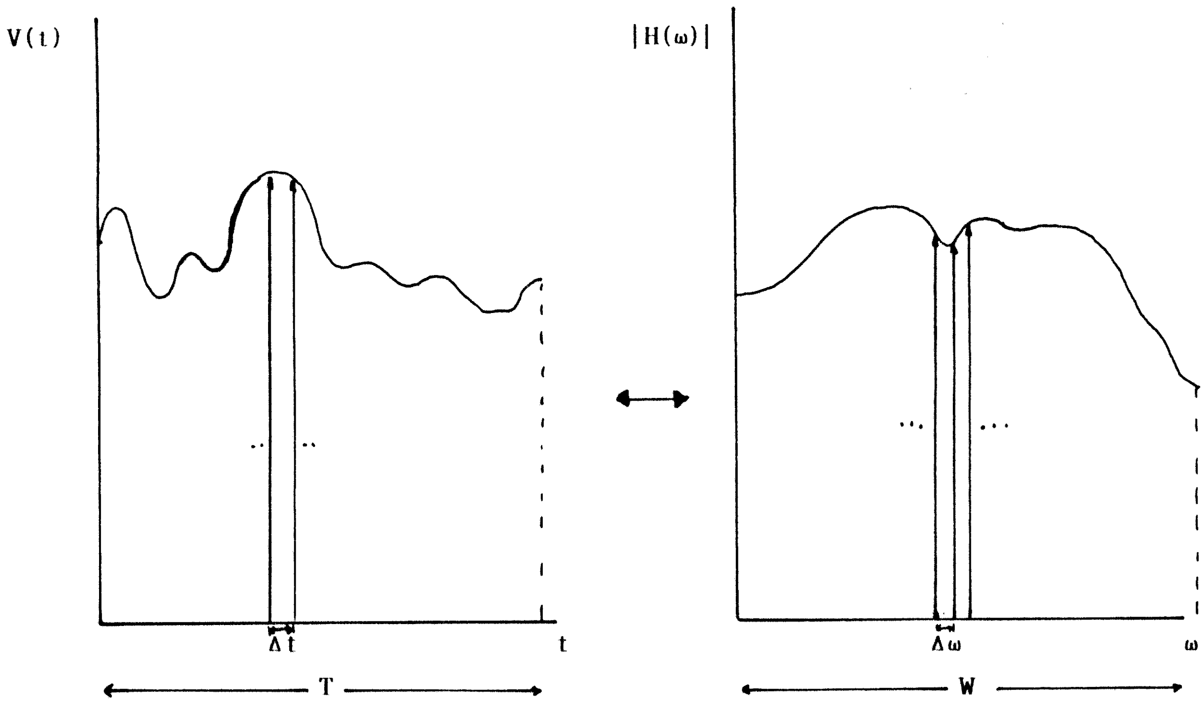

|H(ω)| must be sampled at its Nyquist rate.

Figure 5.2

samples which are giving rise to the "numerical noise". They are thus of little utility in improving the poorer performance at the upper band limit encountered here: Neither can they recover information lost if a frequency is notched out of a signal's spectrum. It is concluded that the plain division technique will return as sound a result as is possible in the present situation.

Deconvolution by means of the time domain technique of "cleaning" is an alternate possibility. This operates by iterative subtraction of one of the convolving functions from a convolved result, leading to an estimate of the second function involved in the convolution. There is no reason to expect that this very lengthy numerical approach will be any less susceptible to errors, working in the time domain, than the previous method, working in the frequency domain. Both approaches require that the function to be reversed be purely linear. The corruption by addition of delayed signals, to be discussed in the next section, represents contradiction of this requirement.

5.5. Quefrency Domain Correction

Homomorphic signal processing is a technique which enjoys popularity in the fields of speech processing and seismic signal analysis.[81] First proposed in 1963, it is perhaps more commonly referred to as "cepstral processing".[96] The cepstrum of a signal is defined as the inverse Fourier transform of the complex logarithm of the Fourier transform of the signal. The inverse is defined similarly, with the complex exponential replacing the logarithm. Techniques for calculating it are now quite advanced.[97,98,99] These will be discussed in more detail shortly.

The utility of the cepstrum lies in the following property: If a signal is composed of an original part to which a delayed, attenuated (or amplified) replica of the original is added, its cepstrum will contain an added delta function whose amplitude and position indicate the delay and relative amplitude of the echo. If the delta function is removed (and the original value of the cepstrum somehow returned to its place) returning to the time domain will reveal the original less the echo - the additive copy is effectively removed. In practice the delta function is either subtracted if its characteristics are known, or the value interpolated from neighbours. The removal is remarkably efficient given plenty of samples. (In certain circumstances the echo can be removed in the time rather than the quefrency domain; that is, without removal of the delta function and reverse transformation, but rather by plain computation on the original, knowing the delay and size of the echo from the cepstral analysis. This is not in general possible, however.)

The application to speech processing is obvious. In seismic processing the interest lies not in recovering the original, but in the causes of the echoes, which are discontinuities in the Earth's crust. These are described to some degree in position and type by knowing the delay and relative amplitude of the echoes. Thus it is the delta functions, not the remaining signal, which contains the information sought.

5.5.1. Application to Microwave Signals

In a microwave system, there is considerable effort and cost expended to reduce mismatches in transmission lines. A mismatch of course introduces an "echo" to the signal, albeit usually travelling in the opposite direction. A cascade of sampling gates, as noted in Chapter 2, introduced (frequency dependant) mismatches to the signal at the gate itself, and by virtue of the side transmission lines used for generating pulses and the sampling diode capacitances, produced echoes which travelled in both directions away from the gate. In narrow band systems where reactive matching is common, these reflections are deliberately induced to produce signals which add to the present continuous signal for some desired effect. (This is usually the eventual cancellation of other reflections to give a match. Such systems are not usually viewed as introducers of echoes, and so the statement seems strange.)

Where a wideband signal is exposed to a situation of mismatch, it will be corrupted by echoes of itself. It appears that the cepstrum might be useful in detecting and remedying this. There does not appear to be any report of the application anywhere in the literature. This is probably attributable to the scarcity of digitised microwave signals in laboratories around the world. The application to the 7912ADM is obvious in the light of the results of the response measurements reported in Chapter 4. However, there are problems with the application which become more significant in microwave situations than in those encountered elsewhere.

In the determination of the (complex) cepstrum of a signal, the complex logarithm of the Fourier transform must be found. The complex logarithm is defined as:

$$\log(Z) = \log(\text{magnitude of } Z) + j(\text{phase of } Z) \qquad (5.1)$$

Here, the phase of Z is not the phase modulo 2.Pi, but the actual "unwrapped" phase. The determination of this unwrapped phase is a numerically burdensome task which has thwarted researchers for some time, and continues so to do.[96-99] The problems arise because a division by each term in the first Fourier transform is necessary, and the successful calculation at each stage depends upon the result of the preceding stage. When one term is zero or near zero, the determination is impossible.[99] (The original paper called the unwrapping task "saphe cracking", a name which has not stuck despite the very apt anagram/pun.)

A zero term in the transform is sufficient to completely prevent phase unwrapping. A term which is not zero because of noise (numerical or otherwise) is sufficient to cause the determination of an <u>incorrect</u> unwrapped phase. Thus, there is associated with any cepstral determination some level of confidence, set by the constraints in the numerical integration carried out by the unwrapping algorithm. An unwrapped phase is by no means certainly correct.

Any oversampled waveform should have zeroes in its transform. This is simply because the FFT guarantees samples beyond the band limit, if the waveform is sampled in the time domain beyond the Nyquist frequency. The very nature of scan converter tubes ensures oversampling, as has been noted previously. Untouched, this would eliminate cepstral processing as a possibility. For this reason, in the absence of noise, it would be invariably necessary to "pack" a time domain record down, to force it to approach the Nyquist sampling limit. The PAC (pack) command in DIOS is provided for this function. In practice, it is often necessary to PACk a waveform in order to perform cepstral analysis. Noise, however, sometimes permits the determination without PACking, as it is of course indistinguishable from signal to the routine.

Further, microwave signals are often narrow band in nature, such as a modulated carrier. These signals will have no determinable cepstrum, since their FFT should have zero magnitude samples at all frequencies outside the signals band limits. A single sinewave signal will also have no cepstrum.

Even when all components are present in a signal, the cepstral determination programme requires significant time to develop the transform, since it must execute at least 2 discrete Fourier transforms as well as the numerical integrations for phase determination at each step.

For echo detection, the cepstrum must be such as to betray the delta functions which signify echoes. The delta functions may become blurred if the echoes are not independent of frequency, because the echoes are not pure replicas of the original. They also become difficult to certainly identify if the cepstrum is far from "smooth", or the delta functions are so small as to "hide" in the features of the remainder of the cepstrum.

Yet one further hindrance to the development of a mismatch removal algorithm of any generality is the range of echo situations which arise in microwave hardware. In speech processing the echoes (even if frequency dependant) are usually singly occurring delayed replicas. In connected transmission lines the energy can "bounce about" with less loss, because the signal does not disperse as does sound, even in a reverberant environment. (Researchers find that very reverberant locations can present problems to the process, though it is uncertain whether this is due to the multiplicity of echoes.) In the example case of a shorted stub fed from a matched source, the echoes occur as a series of delta functions with fixed ratios of amplitude and delay increments. (This is easily seen by considering the signal which propagates up and down in the stub, seeing a match at neither end.) The signal which must be removed from the cepstrum so that the inverse returns the uncorrupted signal, is completely characterised by the delay, initial amplitude, and amplitude ratio between spikes. This holds, however, only for the situation described: Another layout will have another set of additive functions in the cepstrum.

The DIOS functions SEP, for setting the salient parameters of an echo caused by a stub to a wideband signal, and RSE, for removing the echo known to have been caused, are specific to the situation. Another situation will demand another correction strategy. It is pointless trying to build up a library of such functions because of the potential diversity.

To initially investigate the utility of cepstral processing, deliberately corrupted signals have been produced and digitised, and appropriate cepstral programmes included in DIOS. An earlier version of the final amplifiers was modified to include shorted coplanar transmission line stubs in filter networks situated at the inputs of the final amplifiers. The stubs introduced a series of frequency-dependant reflections, each of opposite sign and reduced magnitude to the last; the magnitudes at any frequency and the delays are readily calculable for comparison purposes.

The initial reflection will have a delay of twice the stub length, because energy is diverted into the stub and returned from the shorted end. The use of coplanar line ensures that the position of the short circuit will be accurately known, and indeed also easily alterable by means of a sliding shorting bar. Subsequent echoes will arise because the energy reflected from the shorted terminations will be partially reflected once again into the stub, and again returned with double the first delay. The magnitude of each reflection at any frequency may be calculated knowing filter component values and the characteristic impedances of the transmission lines. The ratio of the magnitude of each reflection to the last will remain constant.

Figure 5.3 shows the signal obtained with a very nearly frequency independent reflection. With hindsight it is easy to pick the presence of the first couple of reflections by eye. Figure 5.4 shows the same signal PACked and suitably windowed. The delta functions corresponding to the first reflections are visible in the cepstrum and the result of interpolating past these and reverse transforming is shown in Figure 5.5. This figure shows a step by no means ideal, but there is great improvement. The spacing of the peaks of the high frequency ripples after the step suggest that all the reflective energy has not been removed from the signal, but it is not possible to say conclusively that these perturbations can be attributed to mismatch. Unfortunately, the technique rapidly becomes unable to cope when the reflections are reduced for any significant part of the frequency spectrum; that is, when the stubs are made to reflect energy unequally across the bandwidth of the system. In this case, the delta functions expected in the cepstrum become indistinguishable from the rest of the data.

TEKFREQ   18 7 1983   12 interpolations   ^NONE   +2.E-09S

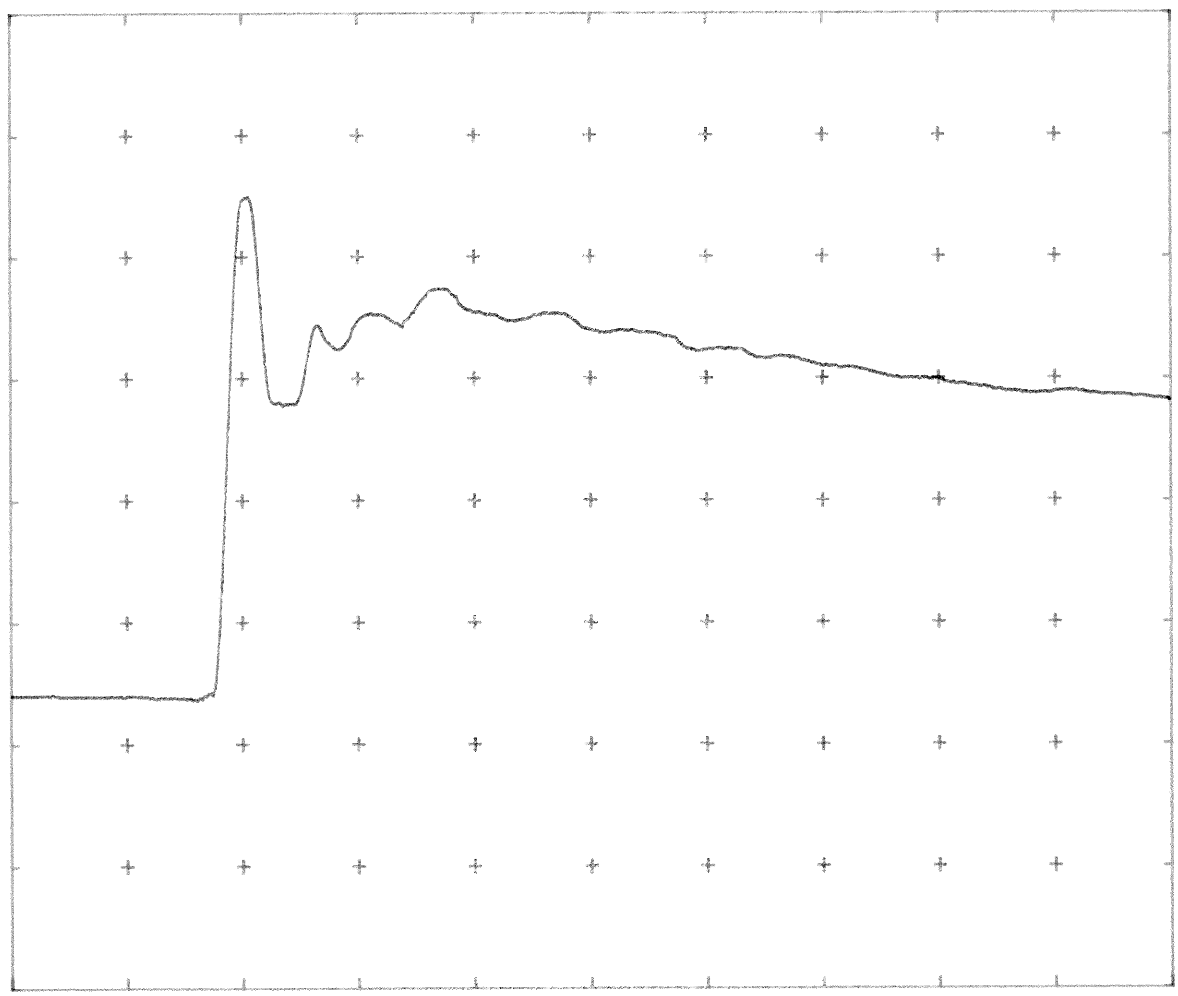

ABPoutput for wfm ACTU   12.04hrs

version 2.49

Figure 5.3

TEKFREQ 18 7 1983 6 interpolations none 4.00e-9S

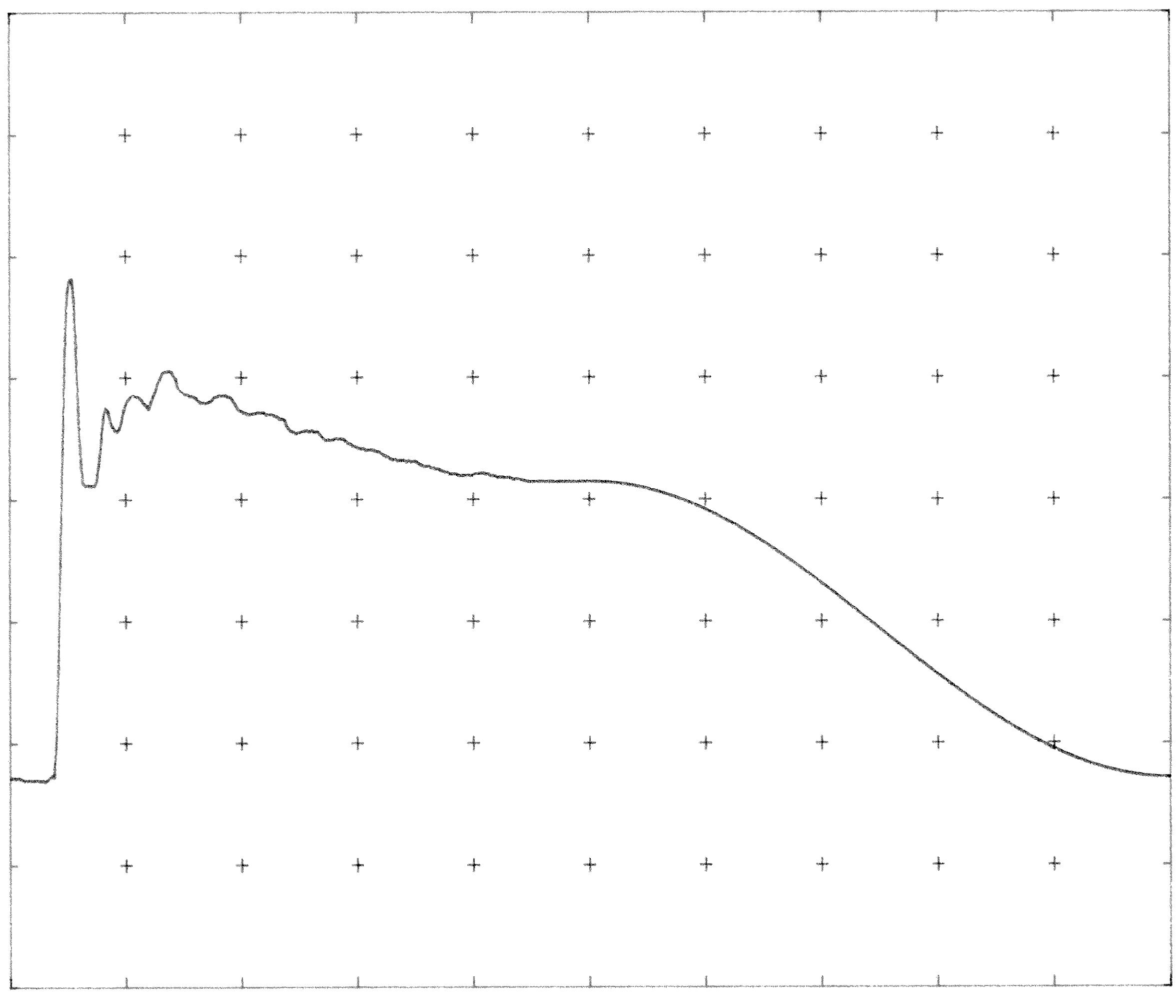

PACd waveform ABP of Step 12.04hrs (H/2)

version 3.00

Figure 5.4

TEKFREQ   18 7 1983   6 interpolations   none   4.00e-9S

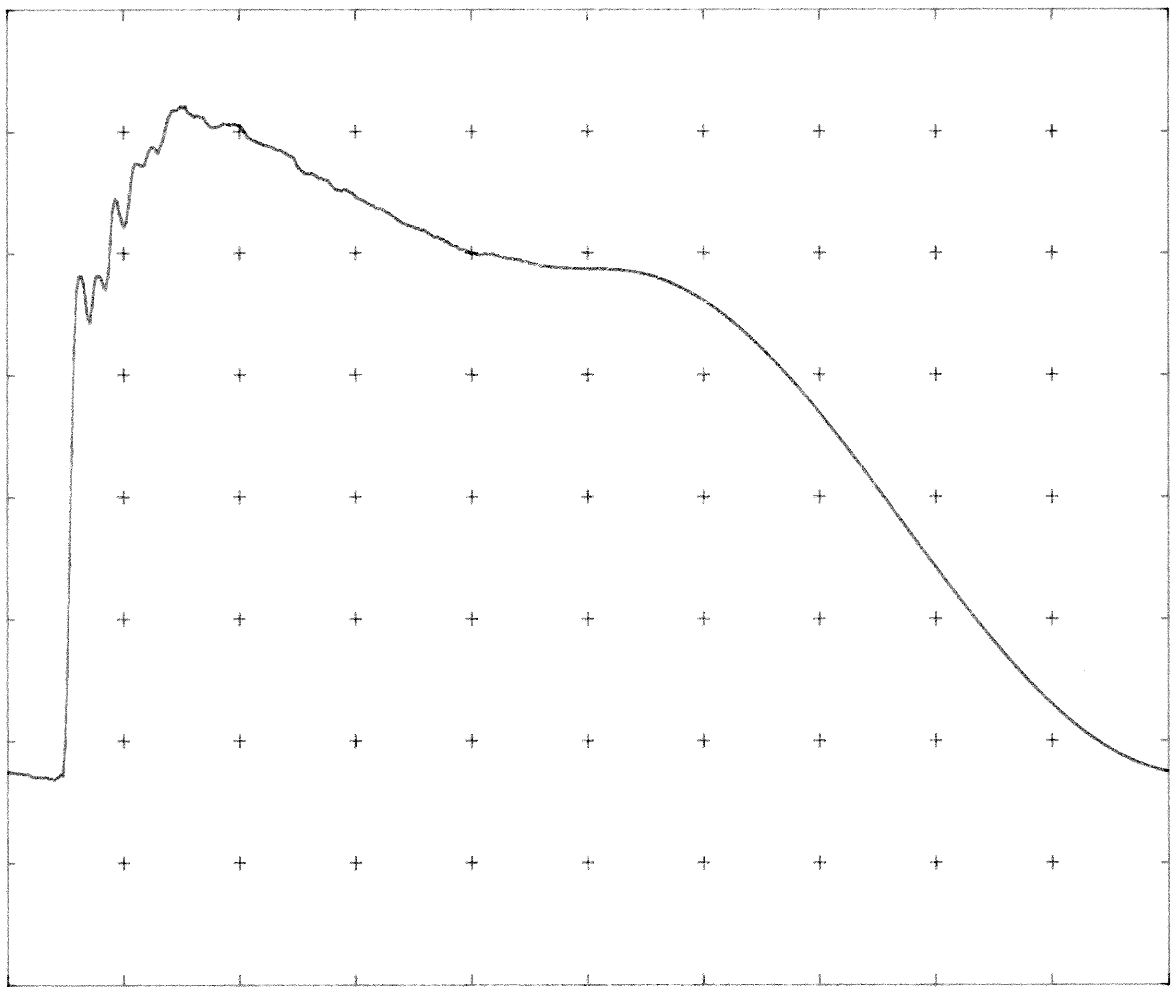

CEPW   12.04hrs (H/2) (M)

version 3.00

Figure 5.5

5.5.2. Conclusion

In order to determine the (mismatch free) frequency response of the 7912ADM, it was noted in Chapter 4 that it would be necessary to first characterise and remove the echoes in the step test signal. This has proved to be substantially impossible. Figure 5.6 shows the cepstrum of the (averaged) step signal encountered in Chapter 4. No delta functions are readily discernible. Figure 5.7 shows the cepstrum of an unaveraged example of the same signal. The cepstrum of the averaged signal of Figure 5.6 is plotted superimposed upon the cepstrum of a single event for comparison in Figure 5.8. (The originals produced by DIOS are in colours, enhancing the usefulness of the comparison.)

The comparisons are not enlightening. Either the cepstra are erroneously determined because of some failure in the numerical phase unwrapping function, or more likely the "delta" functions are concealed in the cepstra shape. Even with the knowledge of the location of at least one of the functions which might be expected to betray the echoes, it is not possible to isolate or eradicate it. The additive impulse signals are believed to be concealed largely because they have been made less sharp by the frequency dependant nature of the echoes, which do not appear to have the full low frequency content.

The unfortunate conclusion to be drawn is that the process has limited use. In critical situations, such as the test step case

STEPV3   16 12 1985   Cepstrum

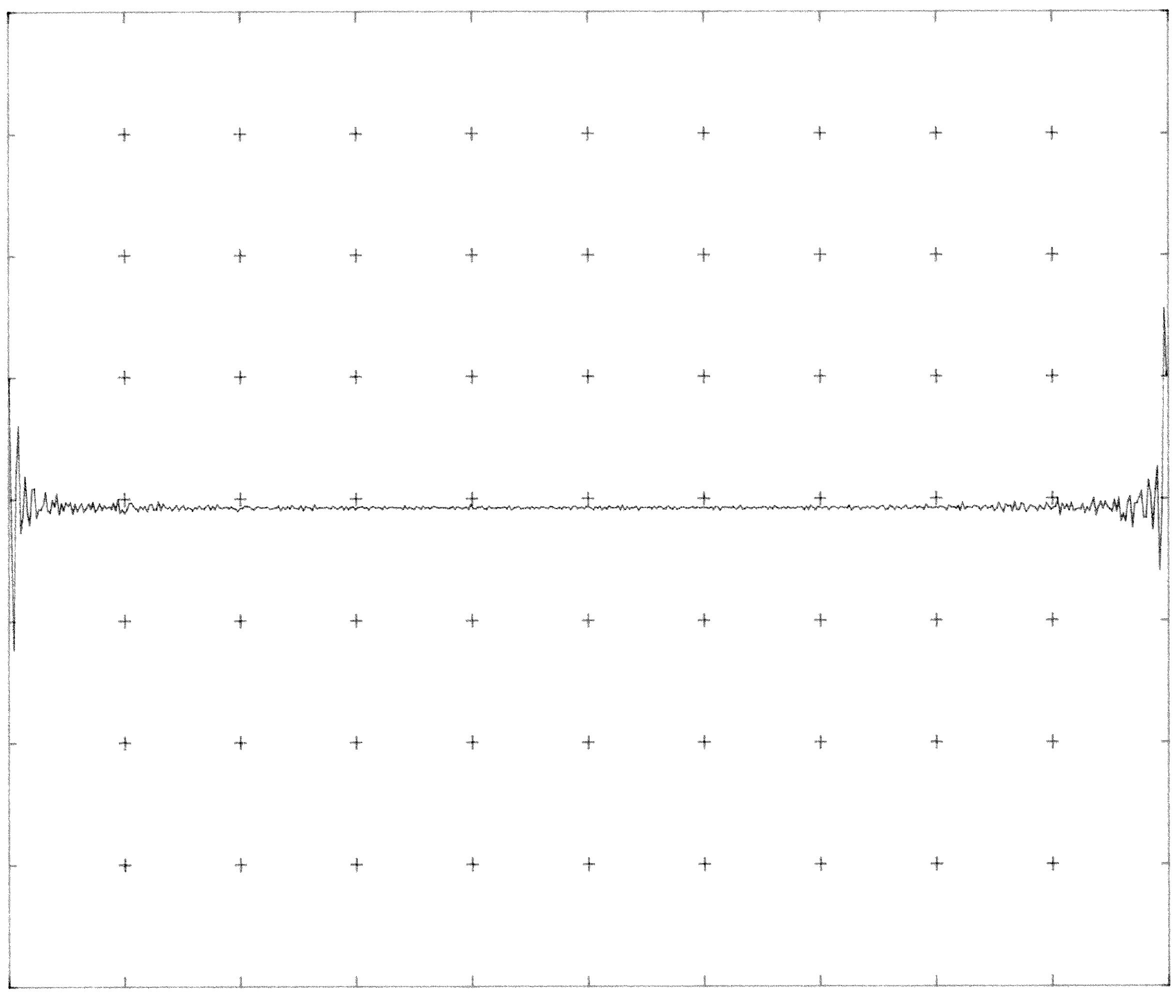

SASW-intelligent signal averaged step response (32 averaged)   14.40hrs (H/2)

version 3.00

Figure 5.6

STEPV3   16 12 1985   Cepstrum

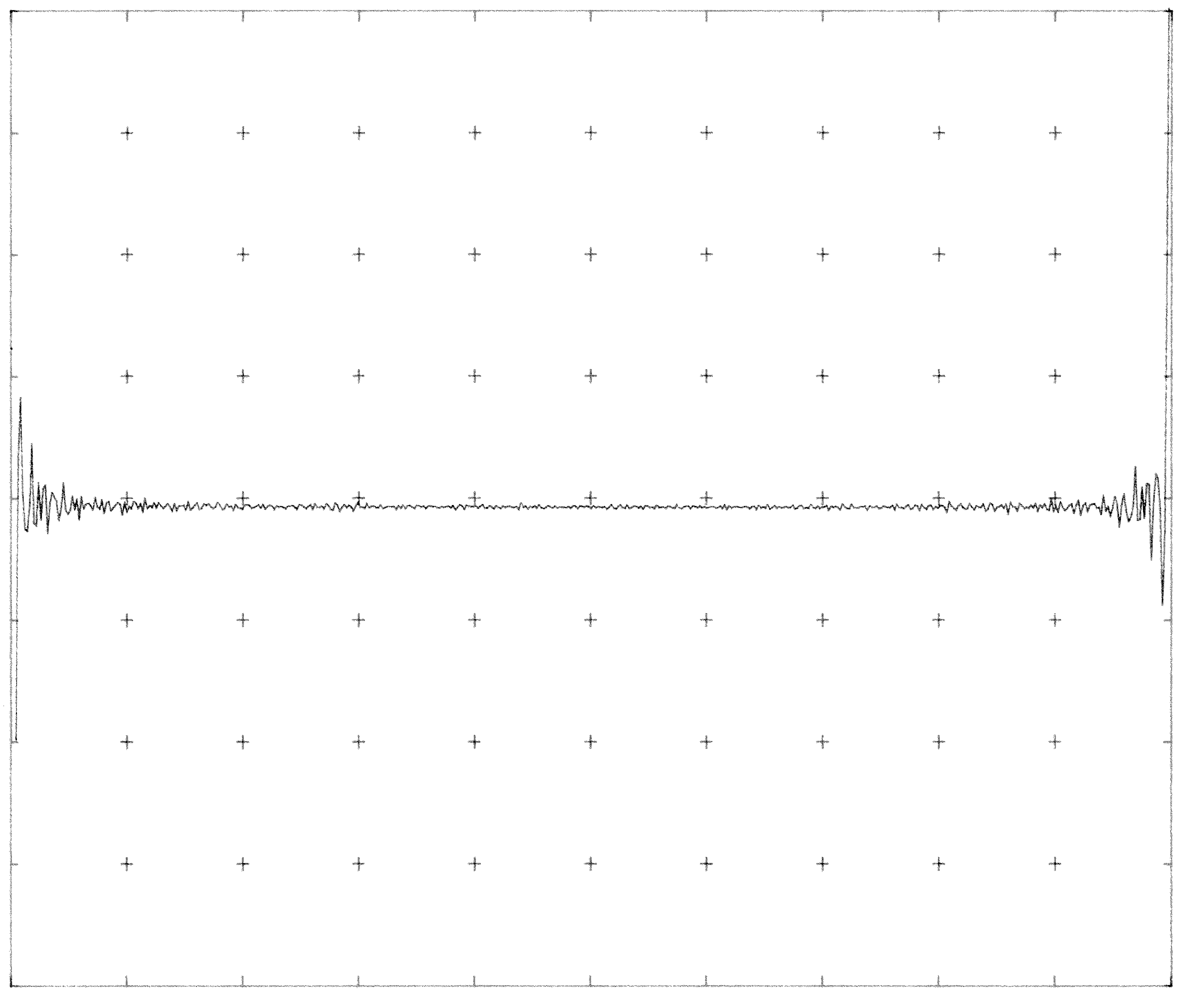

Unaveraged step response   14.40hrs (H/2)

version 3.00

Figure 5.7

STEPV3 16 12 1985 Cepstrum

STEPV3 16 12 1985 Cepstrum

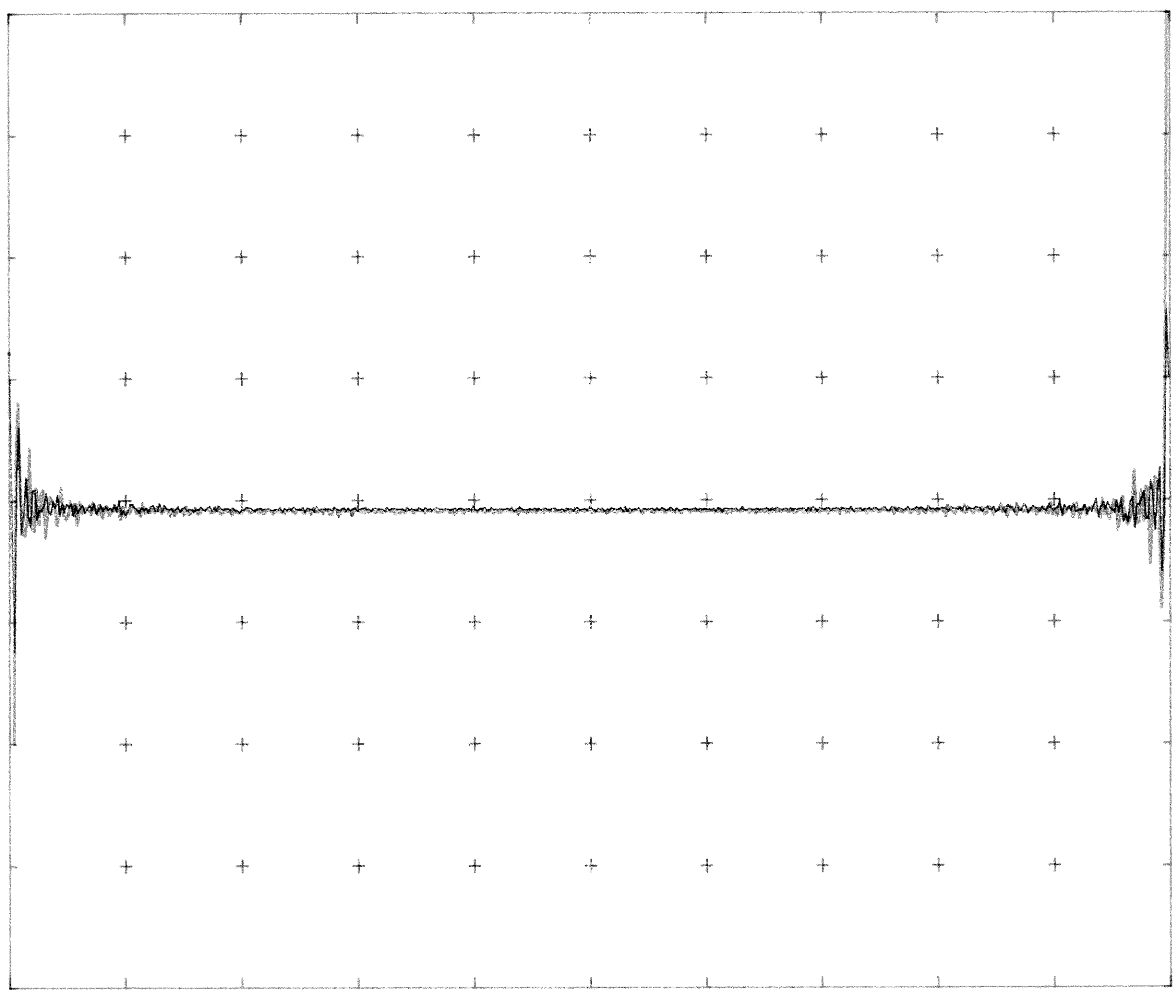

Unaveraged step response 14.40hrs (H/2)

SASW-intelligent signal averaged step response (32 averaged) 14.40hrs (H/

version 3.00

Figure 5.8

above, the technique has failed to resolve the echo signal parameters of delay and magnitude. Where the echo situation can be deduced using a suitable test signal, because it is simple or calculable, routines can be set up which will apply the correction to the cepstrum of an arbitrary signal which has passed the same circuitry. However, it is by no means certain that the cepstrum of each new signal acquired will be determinable, and thus the correction effected. In addition, the removal may be incomplete, and the degree of success is not readily estimable. This is because the cepstral analysis techniques do not allow one to decide whether remaining corruption is the result of linear or non-linear processes, because they are **insufficiently sensitive to the presence of small scale, frequency dependant echoes.**

The inability to remove the mismatch corruptions, and their significance at the frequencies involved, prevent the really useful application of deconvolution techniques. Simple "boosting" of parts of the spectrum is possible, but any detailed compensation requires the echo-free response, which cannot be determined.

6. SUMMARY

6.1. This Dissertation

The search to develop a functional transient digitisation instrument with usable performance reaching 2 GHz has been described. The dissertation commenced by considering all presently identified techniques which have been put forward as means to achieve the overall goal. All but two were eliminated and those two were investigated further.

After considerable investigation, the first approach was abandoned in response to technical constraints. Work done indicates that the approach is unlikely to be cost effective, or indeed superior at any cost, in the forseeable future. The second promising alternative was then followed, and a machine produced. The instrument is far from ideal, being limited in the short term by the resources which were available for its construction. Its performance and utility are summarised below. Nevertheless, two important achievements are to be noted.

Firstly, ongoing improvement in the technologies used, as reported in the literature, will soon make more advanced instruments feasible. Additionally, the need for such instruments grows with the advances in the semiconductor and communications fields. Much of the groundwork has been laid for further investigation, and for the development of further generations of instrument, in response to the anticipated improvements and needs in associated technologies.

Secondly, a very usable machine has resulted. Although much of the work was undertaken with the aim of developing a machine of general and simple application, this aim has not been achieved. The transient digitiser remains a specialised laboratory instrument. Its cannot be straightforwardly used by an unskilled operator; he must well understand the workings of the device and be specifically aware of machine limitations in order to guarantee any reliable extraction and interpretation of results.

6.2. The Instrument

Conventional oscilloscope specifications cannot apply to the 7912ADM. Neither can conventional operating procedures or display interpretations be adopted. This will remain the case as long as the state of the art continues to advance at the present pace.

The 7912ADM system is capable of acquiring and processing waveforms with frequency components from below 5 MHz to above 2.0 GHz. It has an accuracy, when calibrated against a suitable standard, of at least 4.3 bits and up to 5 bits, and a resolution of at least 9 bits, at full scale deflection. It has a vertical sensitivity of better than 25 mV for full deflection. It has full-target deflection capability in the absence of mismatch, and half scale deflection capability under most other signal conditions.

The processed waveforms suffer from mismatch distortions. General attempts to eliminate these by numerical processing have proven unsuccessful, and this failure has prevented the operation of numerical deconvolution strategies. Nevertheless, a human operator can interpret what his eye sees, recognising much of the distorting contributions, and gleaning useful information from a waveform in spite of the additive corruptions. For instance, the risetime and monotonicity of a transition may be determined provided it is less than the delay of the first echo, simply by making the measurement upon the first edge encountered in the data record, realising that subsequent perturbations are accidentally introduced replicas.

This degree of utility **is** worthwhile. The width of a very infrequently occurring laser pulse has been determined using the 7912ADM, and the result is tendered as an example. Figure 6.1 below depicts the waveform acquired when a fast vacuum photo-diode is used with the digitiser. The hump pulse occurring in response to the laser hitting the diode is clearly discernable; the shape is largely determined by shortcomings in the diode and its associated equipment, rather than in the 7912.

In the near future, scan conversion instruments undoubtedly represent the best technology available. It is perhaps a simple observation of engineering history repeating itself to note that radical methods such as electro-optical sampling techniques[105] which are just being used in the most advanced research

LASER   30 10 1984   1 interpolations   NONE   +2.E-09S

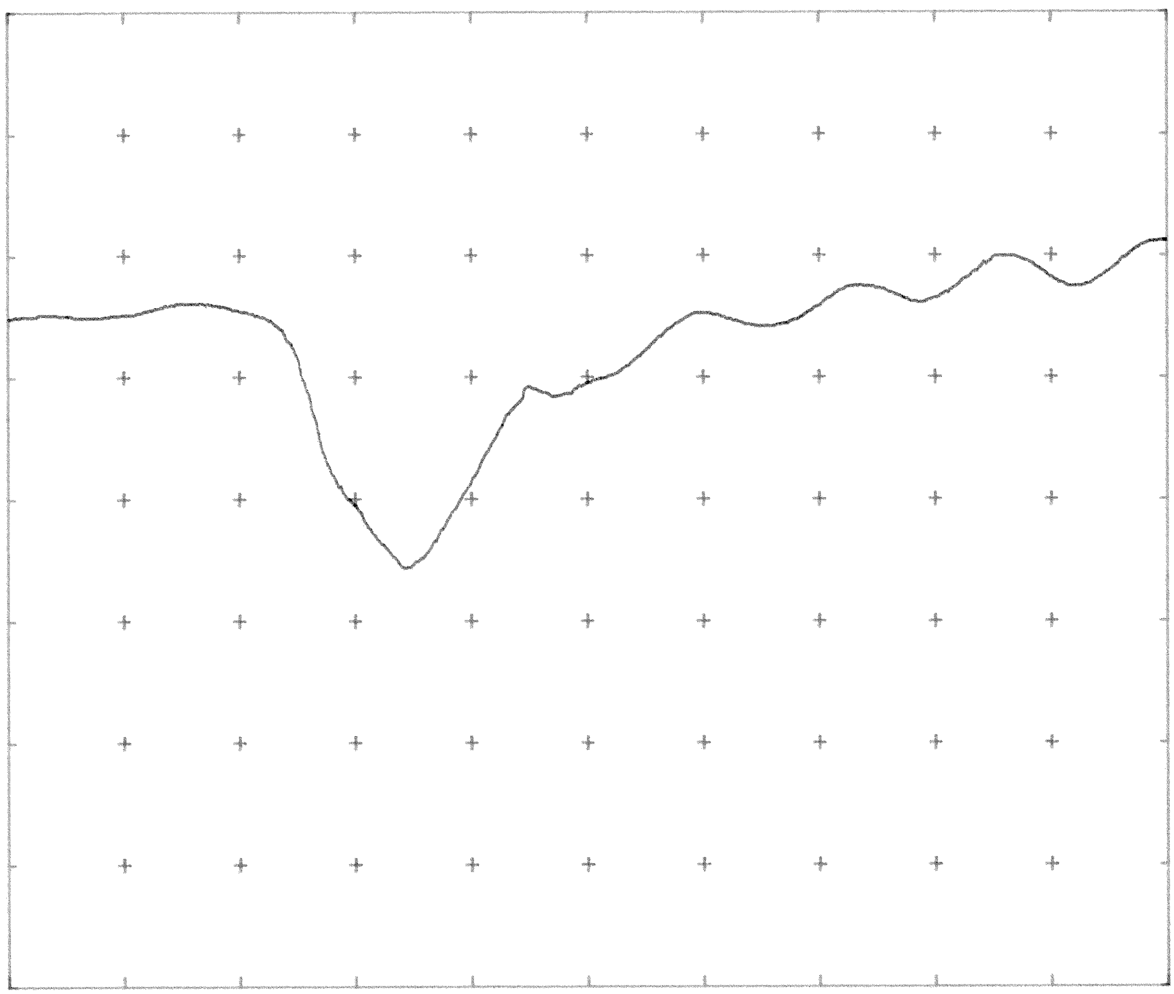

Laser Pulse   12.01hrs

version 2.49

establishments for repetitive signal measurements, may sweep the conventional techniques away after all.

X. THE INSTRUMENT OPERATING ENVIRONMENT: DIOS

X.1. Overview

The 7912AD was not designed as a stand-alone instrument, but as a peripheral of an automated measurement system. Although it has certain front panel controls and outputs for a local display, these are provided to facilitate rapid initial setup only, and do not permit the machine's full capabilities to be utilised by any means. This is more so for the 7912ADM, for whose full operation considerable processing is used. A full software operating system should:

1) permit friendly manual control of the digitiser;

2) automate control of the digitiser where possible, and handle all low level protocols, etc;

3) store and display control setting files;

4) handle data file storage, display and output;

5) process waveform files in predetermined ways using the correction and compensation algorithms;

6) automatically document data files with time, date, and keep a processing history, etc.

This document describes the format of the Digitising Instrument's Operating System (DIOS). It describes the environment it presents, and lists and explains all the commands provided. Thus it serves both as a user's manual for DIOS, and as a description of the processing and other special features which it has been found facilitate and ease the use of the instrument.

X.2. DIOS Organisation

The versions of DIOS written by the author are in HPL, a powerful but line oriented proprietary language developed by Hewlett-Packard for their 9825 computer. DIOS runs on a 9825T. HP still support HPL on their series 200 machines, and DIOS could be modified for direct use with these. However, it is apparent from a number of machine features that the 200 series machines are not primarily intended for HPL, and HP are no longer supporting the language. The last version of DIOS is version 3.0, and it has been so broken into sections as to facilitate both easy editing and minimum effort upgrade to any language supporting IEEE-488 instrument interface bus standards. The HP implementation of this is called HPIB, while the Tektronix version is called GPIB. A listing of DIOS is available, along with all relevant utilities, from the University of Sydney, Air Navigation Group, Department of Electrical Engineering. Owing to length as well as the observations above, it is not included in this thesis.

Modelled vaguely after CPM, DIOS consists of a kernel and a series of discretely executable blocks. Some blocks are internal to the kernel section beyond the key processing routines. These are said to be "innate". Some blocks are resident in the memory, as permitted by available memory space. These are referred to as "resident". Others reside only on mass storage, being loaded if required and deleted when they relinquish control to the next higher level. These are "external", or "disk resident". All functions are called by a 3-letter (or 2-letter and 1-number) mnemonic. The three-character requirement for all commands means that no entry terminators are required at the main command entry level. Online help is provided.

Data entry within executing routines may require terminators as appropriate for the situation.

Resident and External commands can be converted to the other type (that is external to resident or vice versa) simply. The conversion is effected merely writing residents to a disk file of that name and deleting them from RUNTEK, or by chaining (concatenating) the diskfile to RUNTEK. The kernel will search memory first, and look at primary mass storage if the routine is not found. An error is only noted if the kernel is unable to locate the routine anywhere. (While this search has been implemented at the HPL level, HP BASIC supports this at the machine operating system level.)

The blocks of DIOS may be individually added, translated or deleted when editing the operating system. A missing block merely causes the command interpreter to return an 'unknown mnemonic' message. This manual should provide sufficient conceptual information that a programmer can duplicate DIOS on any machine, provided he is familiar with its own internal workings. The entities listed below are assumed to exist for proper operation of the system. Entities which are marked with an asterisk may be temporarily absent, with the consequent inexecutability of routines that need them, or with substitution of a warning message for the missing information.

The 9825 system was provided with a tape drive mass storage backup integral with the processor, which will act as main mass storage, providing all facilities except execution of disk resident commands. Thus loss of the disk drive makes the machine more portable, but removes disk resident DIOS commands and slows execution considerably. (Thus it is usual for external commands to be those not needed in a remote data collection situation, while the necessary commands are resident.)

X.2.1. DIOS System Entities

Host Processor

Keyboard

Printer*

One Line of Text Display (the System line)

Graphic display*

Plotter*

Real Time Clock* (Provided by an HP2240A)

7912AD or 7912ADM*

On/Off control interface for 7912* (HP2240A)

Secondary Mass Storage (Disk Drive)*

X.3. Using DIOS

When RUNTEK is first run, DIOS will boot and run the INItialisation routine. (See INI below.) Once operating, the wake-up message "Ready" is displayed, along with the current system time if the real time clock is on line. The time is displayed as a negative number if the clock is not on line. The number gives the last valid time. The kernel keeps a record of the number of times the digitiser has been initialised. This represents a history of usage. This and other data associated with the machine (such as a record of target defects and information used to monitor the digitiser operation) are kept in an associated mass storage file.

The (necessary) system display line is used for all system messages, error reports, warnings, etc., as well as input and activity display.

Commands are entered by typing their 3 mnemonic characters. Where more than one command is required to be executed in sequence, the mnemonics may be typed in directly one after another and they will stack up and be executed in turn. Two properties of this system should be noted. An unknown mnemonic will be ignored, a warning issued (on the printer if it is activated - see the printer control commands) and execution continued. Thus a sequence can be executed without a step if it is typed in wrongly or a command is absent. Also certain routines will halt to await specific input from the keyboard.

When this occurs, keyboard input is switched from the command level to the routine executing. The routines which do this are of course only ones which would be executed when some form of manual override is desired by the operator. In these circumstances the operator is highly unlikely to be stacking commands, and the problem should not arise. The system time is displayed, with a ready message, when idle.

The continuous execution of the stack may be halted at any point in the input stream by the presence of a semicolon in the data stream. Thus a command may be followed by a semicolon, and execution will stop when the semicolon is reached. The semicolon may not break a mnemonic, that is occur in the 2^{nd} or 3^{rd} character position, or it will be interpreted as part of that mnemonic. Extra semicolons are ignored. Execution resumes whenever a carriage return (<cr>) is pressed, but will halt if at the appropriate point if another semicolon is subsequently entered. (Thus a <cr> effectively purges all semicolons present in the input stack at the time of entry.)

Editing of the input stream is achieved only by removing the end character from the stream, using the erase key. One command, KSW, is executed out of order in the stream. That is, it is executed at once if it is present anywhere in the stream. This command aborts the execution of the whole stack. (It also escapes the user from any 7912 operations pending - see the KSW entry below.)

The input stream is not case sensitive.

Pressing the space bar at any time displays the unexecuted commands in the stack on the system display line. It is not entered as data in the stream. Whenever a command is taken from the input stack and processed for execution, its mnemonic alone is displayed. (This informs the operator what the machine is currently doing.) Returning to viewing the command stack requires that the space bar be pressed (or an executable character entered) each time the stack is reduced in the natural progress of execution of stacked commands.

X.4. DIOS Commands

X.4.1. ABP

This command instructs the controller to determine the Actual Beam Path from edge data. Whereas ATC or ITC determines each of the samples without reference to information in the adjacent sample positions, except in the case of a missing sample, ABP is a two dimensionally conscious routine.

ABP will report a fatal error if there is no edge data in the controller.

Because ABP makes maximum but safe use of the data available in edge arrays, the extraction of the edge arrays from the raw data should be carried out with lax spot-width and spot-delta-width constraints. This will ensure that little useful data is eliminated in that step.

ABP takes the form of a preprocessing routine which falls through to the ITC routine. It modifies the edge data in each location to reflect information gathered from nearby locations, then allows ITC to convert the edge data to a single valued trace, interpolating and extending as required to force a valid sample value at each horizontal location. (See ITC.)

Two mechanisms which act in a Scan Conversion Tube combine to produce potentially misleading images. ABP is an attempt to compensate to some degree for these mechanisms. The first problem arises from the variation of spot size with writing rate. When the beam's linear speed rapidly reduces, as is the case at the end of a step transition, the spot may widen from only a few pixels (or no written pixels at all in places) to tens of pixels. As the spot widens it can flood electrons backwards into columns of pixels over which it has already passed, and into which it has yet to proceed, polluting the datum associated with each of those columns. The second problem is a similar corruption of data in nearby columns caused by the spot retracing its path before it has moved one spotwidth further along the target. This may result in the trailing part of the spot overwriting those pixels which its leading edge wrote moments before, or writing pixels previously not hit. By way of example, this occurs where a sinewave oscillation is written with too low a sweep speed to separate the adjacent peaks of the wave.

In both these cases, the output of an ATC command can be very misleading. The decaying sinewave oscillation may be portrayed with severely reduced amplitude, rejected altogether by an EDG algorithm with too strict width or delta-width constraints, or converted into a flat region with irregular perturbations. The end of a step transition can acquire or lose fine detail in the "smudge" generated by the sudden spot width increase. The risk of these happenings is mentioned in the 7912 manuals, but little is provided by way of a solution.

ABP operates by a process of circle fitting inside the confines of the lines which define the trace edges. The call results in three passes being made over the data arrays. In the first pass, an iterative honing process is used to determine the centre and radius of the largest circle which can be fitted over written pixels in the image, its centre remaining on a fixed vertical line. This process is repeated for each vertical column of pixels, the circle being centred in the middle of that column. (The assumption is made that the beam has been adjusted to produce a circular spot on the target. This is part of the normal 7912 adjustment procedure.) In the current version, the centre position is resolved to the nearest 0.25 vertical units, where one unit is equal to one pixel.

The second pass rejects the values of any column all of whose written pixels are covered by circles of other columns. (It is necessary to resolve to sub-pixel levels to prevent numerical rounding causing rejection of otherwise valid columns.) Acceptable columns are modified to leave the data in edge form, but in a fashion that will allow the normal ITC routine to retain the improved resolution which the circle fitting technique yields. That is, the numbers are stored with fractional parts.

The third pass consists simply of falling down to the ITC routine. That routine cannot tell that the data has been processed, since it handles rejected data (the data is made negative, as is the Tektronix convention) and operates in floating point rather than in integer arithmetic.

(Two observations arising from experience with the problems addressed by ABP need to be made here. Initially, it is significant that the image aberrations will not be avoidable with any scan conversion system. This was not taken into account before embarking on the project. They may be resolved (as has mostly been the case in the author's experience) by suitable alteration of timebase speed and triggering conditions, but this is only practical when the event being observed is "reasonably recurrent", that is it is not a true single or very rare event, which is just the type of occurrence with which scan conversion alone is equipped to deal.

A second observation is that the author has not been able to devise an image processing technique which will deliver results which can be <u>shown to be optimal</u>. That the ABP routine is an advance over the plain ATC method is sometimes evident to a user's eye: the single valued waveshape returned by ATC alone lacks portions which an expansion of the waveform exposes. The use of ABP will at least warn the user of the uncertainty by rejecting the suspect portion (and giving an interpolation warning), and may resolve it all or in part. It has been tested by the author on data numerically concocted by simulating the spotwidth variation. In this rather artificial case, it reveals reduced image distortion. However, since the simulation assumes all and only the aberration mechanisms ABP tries to deal with, it might be expected to do just that. ABP also often rejects pieces of the waveshape which might have been resolved by some superior approach. The lack of theoretical support for the circle-fitting algorithm must bother an engineer using it. Nevertheless, no solution for this dilemma has been found.)

X.4.2. ACN and ACY

These functions turn on (ACY) and off (ACN) the Automatic Correction facility.

When activated, waveform data loaded from the 7912ADM is corrected before storage for non-linearities known to be introduced by the deflection circuitry. (Note that the DEB function must have been successfully executed to obtain the correction calibration data before use of the facility has any meaning.)

The facility defaults to the inactive (off) state.

Correction could in extreme cases result in data being corrected to values greater than FSD or below zero. Minus numbers are not allowed, because they are the signal used by Tektronix to flag absent data. Hence "clipping" can occur on negative excursions purely in the correction stage. In general this is a very rare occurrence. It is also unusual for the operator to be aware of the fact that the clipping occurs in the correction rather than the digitisation operation.

X.4.3. ALT

ALT displays the alternate (non-current) waveform store on the graphics display device. (See DSP.)

The routine operates by toggling the domain identification flag, calling DSP, and only then returning control to the next higher level.

X.4.4. ATC

This routine causes the main data memory to contain a single-valued 512-sample time-domain waveshape. The method by which this is determined depends upon the conditions present at the beginning of execution.

If an EDGe waveform is currently in the memory, control is passed to the Intelligent Trace Centre (ITC) routine. This routine provides a slightly improved version of the ATC routine provided as firmware in the digitiser.

If there is not an edge data set available in memory, and the digitiser is on line, it is requested to determine its atc array, and then to transmit this array via the bus. (Errors which may arise as a result of the state of the digitiser are handled by the automatic error routine, which would abort ATC if required.)

The routine also requests and stores the largest number of interpolations which occurred together at any point in the routine, and the current scaling data, which includes any horizontal and vertical multipliers.

For a description of the 7912's internal atc routine, see the ITC command description.

X.4.5. CAL

The CAL routine computes a new frequency response correction array given a single valued waveform array of a digitised step.

The step must be presented in a positive going sense. The routine will use either the current store contents, or a nominated file from mass storage. The file should not have been altered by any post-processing algorithms of the corrective sort. (See Chapter 5.)

The routine will accept a nominated maximum frequency to be covered by the correction array, or will use the default of 3 GHz. The value is stored in the "TEKfrq" file along with the correction array.

The routine determines the position of the step within the window, and brings it to the window edge, so as to reliably reference the phase of components. (The time delay would correspond to a linear phase shift in the correction array, which would hence be introduced in all further calculations. This would have the effect of time shifting all waveforms which were subsequently corrected.)

The routine operates by taking first order differences of the array terms, which approximates differentiation. Because of the forced oversampling inherent in the acquisition process, the approximation is sound. The differentiation converts the data array into an impulse response from a step response.

The input data array should of course have as large an amplitude as possible, consistent with not overloading the circuitry. If the input array peak-to-peak value is below half maximum, or if there is not an oversampling factor of at least four between the requested maximum frequency and the maximum frequency available from the Fourier transform, the routine aborts and flags an error. The error always prints a message, even if POF is active.

Sixth order Lagrangean interpolation is used wherever there is a discrepancy between the values available from the transformation and the values required for the correction array.

The corrector array is tapered to a value of unity (no correction) at the nominated maximum frequency, using a Tukey window operating above 2 GHz. It is also damped in that manner for frequencies below 'midband'. The midband region is determined from the first few samples above DC, and the array forced to a unity effect at DC, with a Tukey window, starting at the lowest midband sample. (This is typically the second sample, as response is usually flat to 10MHz.)

The routine displays the correction array and writes it to "TEKfrq" when it finishes.

X.4.6. CAT

CAT displays the catalog of control setting files and waveform files on the primary mass storage device. Other file types are suppressed.

CAT will not display the files on the secondary mass storage device if the primary one is on line. File access is limited to the DSC statement in that case.

The catalog is printed if PON is active. It is printed if the secondary mass storage device is the only one on line even if PON is not active. An error results if CAT is executed, the primary mass storage device is off line, and the printer reports out of paper.

CAT loads all files sequentially from secondary mass storage, and leaves the last one stored there (chronologically) in memory.

X.4.7. CEP

CEP attempts to calculate the Cepstrum of the current waveform.

CEP requires a single valued time domain waveform to be present. An error is flagged if this is not the case. CEP may require the waveform to be packed. (See PAC.) Because of the uncertainty of devising the unwrapped phase of the log-transform, progress reports will be printed to assist further steps if PON is active. The complementary function is ICP.

For a full description of the use and execution of cepstral processing (homomorphic signal processing) see Chapter 5. The algorithms used by CEP are those defined in the literature and are similar to those available on the IEEE standard DSP package. Full references are given in the descriptive text.

CEP is used mainly to detect and to assist in the correction of "echoes", or mismatch phenomena present in the digitised output of the 7912. These may arise from its own amplifiers, or from external transmission line circuitry.

X.4.8. CUR

The CUR command places a cursor on the graphics display which will coincide with the trace of the current waveform.

It requires a single valued trace, and assumes a time (or quefrency) domain display. A warning message is issued if these conditions are not met, but no error is flagged, and execution does not stop.

The cursor follows the trace, and is associated with time and magnitude numeric displays on the screen. The time co-ordinate is given in seconds in engineering notation. The amplitude co-ordinate is given in divisions. The cursor is moved and the delta-time between two cursor positions obtained with the numeric keypad keys. The system display is loaded with a prompt to indicate how to move the cursor and how to obtain delta-time readings. Time readings are printed if PON is active.

The routine is terminated (leaving the current cursor(s) in place on the display) by pressing any alpha key. Thus the entry of any new DIOS command causes CUR to return control to the calling level. Note that the command assumes that a DSP command has previously been given. A cursor with no trace will be obtained if the screen was previously clear, or a cursor whose cross hairs do not coincide with the trace will result if the current store has been updated since the last DSP command. This necessary to handle graphics devices whose contents are not accessible.

X.4.9. DCW

DCW causes DIOS to load an array from mass storage file "tekDCW" and to use this array as a window for single valued time domain data in the controller.

DCW requires the primary mass storage device to be on line. It will flag an error if the current main data store is not a single valued time domain (ATC type) waveform.

The default window represented by the array is a Dolph-Chebyshev. This window offers minimum main lobe width for a given side lobe maximum level, although other parameters of the window are by no means optimal. It has been selected because the window array is not simply calculated. A utility programme is provided for loading the data file. It is set up to calculate the Dolph-Chebyshev function with -60dB relative sidelobe amplitude.

Other windows can be introduced for use by the operating system simply by loading them into the data file accessed by this command. Since the number of data is the same as the number of samples stored in a waveform data (wd) file, it is elementary to convert a data file into a window file and vice versa.

X.4.10. DEB

The DEB routine determines the digitiser's number of Effective Bits, in accordance with the usual definition, given in reference 42. This definition of 'effective bits' includes in one specification the effects of both differential and integral non-linearity, quantisation errors, noise, etc.

The routine requires the current waveform to be the result of digitising a high purity sinewave. (Note that the result will include the noise and distortion present on the original signal, and so will not reflect the digitiser's performance if the EB value returned is merely that of the incoming signal. Phase jitter and non harmonic components present in sweep generator signals, particularly the newer SRD/YIG source types, generally seem to **preclude their use** in this application.)

The sinewave input must contain more than two but less than 15 cycles. The number of cycles need not be an integer. The waveform should not have been windowed as this is a form of distortion. The waveform must exceed half full amplitude. If any of these conditions are not met, the routine will flag an error, halting operation.

The routine operates by fitting an ideal sinewave to the data samples. The frequency, phase, amplitude and dc-shift of the ideal sinewave are adjusted to obtain the minimum Mean Square Error between the two data sets. The adjustment is an iterative, four part procedure. Each of the above four variables is adjusted in turn. The adjustment steps are adapted (both up and down as required) after each cycle of adjustments. When the mean square error is reduced on two consecutive passes by less than an amount corresponding to 1% of 1 effective bit, the routine terminates. This has been found to produce stable results upon which no useful improvement can be made by extensive iteration.

The routine prints its results when completed even if POF is active, as a record of the determination.

When the determination is complete, a best linear fit is made on the error function. These parameters are stored for use by the correction function. (See ACY and ACN. Correction theory is presented in Chapter 5.)

X.4.11. DMC

DMC allows the Definition of a Macro Command. A macro command is a set of up to 20 single legal DIOS commands which are invoked by the use of a single command. The execution command is EMC.

The macro defaults at powerup to being the commands required to perform a frequency compensation on a single valued time domain waveform.

The mnemonic DMC is used as both the beginning and end delimiter of the definition. All keystrokes entered between the two occurrences of DMC (correctly separated by a multiple of three characters) are stored, and entered into the input stream in place of the mnemonic EMC, whenever it is encountered. The first occurrence of DMC halts the execution of the stack. The second occurrence removes all entries between the occurrences, and the two DMC commands, and resumes execution by returning control from the DMC routine to the kernel.

X.4.12. DOC

This command displays a short form documentation of DIOS on the graphics output device. The documentation consists mainly of indicating the use of all global variables, so that the contents can be anticipated when viewing any routine, or the correct information accessed by a new routine being written. Re-usable local and global variables are also noted.

The documentation file is stored on the primary mass storage device as a series of commentary lines. These are displayed line by line, without alteration.

DOC halts execution of the command stack. Keyboard input is redirected, and a key press is required to proceed from screen to screen, or to return keyboard input to the command stack once all screens have been viewed.

X.4.13. DSC

DSC transfers the machine usage file DATEK and collected waveform files from the secondary mass storage device to the primary one. The complementary function is TAP.

DSC clearly requires that both primary and secondary mass storage devices be on line. A fatal error is flagged if this is not the case.

The transfer of any file which would overwrite an existent file by virtue of identical names is not carried out. Because the secondary media is not lost, these files may be recovered by either using a new primary media, renaming the primary media file or deleting it. The former of these methods is not recommended for the reasons concerning the machine usage file DATEK, as described below, and in the NAM command description.

The purpose of differentiating secondary and primary mass storage devices stems from the availability of both compact inboard (tape) storage and superior outboard (disk) storage mechanisms in the HP9825. It is intended that the digitiser be usable for a specific task (once prepared) with only the minimum of equipment. The transfer of files to the secondary and portable storage system would be undertaken just before switching to the portable mode of operation and going out into the field. Upon returning the acquired data files (acting as a control token) would be transferred back to primary storage.

The routines relating to the two systems could easily be altered to reflect small (microfloppy) and large (host system) mass storage facilities, if say, an Integral computer was substituted for the 9825 in later versions. If desired, the restrictions on use of certain (external) commands can be relaxed with improved secondary storage capability. Nevertheless, it is recommended that the distinction be kept (requiring the use of DSC to recover data files from secondary media) in order to ensure that the usage file DATEK remains associated with the digitiser, by the effective passing of a 'token' (the data files sought) between storage levels.

X.4.14. DSP

This command displays the contents of the current data memory on the graphics device.

If a superimpose request is pending, previous displays are not erased, and the current display is shifted to prevent overwriting of characters.

Two data stores exist within the controller memory, each capable of holding a single or double valued set of 512 points. The first always holds time domain data when it is present, and is set to be the current store whenever it is loaded from the digitiser or a mass storage device. When a transform (for example to the frequency or quefrency domains) is performed, the destination data store is designated as the current store. An internal flag indicates what type of data is current, and then the associated store (a particular type of data is always associated with a specific store) is taken as the current store. The contents of the alternate store can be inferred from the contents of the current store. (For instance the alternate store will contain time domain data if the current one contains frequency transformed data, because the frequency domain data is derived immediately from the time domain data.)

Time domain data is displayed with scale, identification text, time and date of capture, jobname and peak to peak amplitude. The number of interpolations used is displayed. A warning is displayed if excessive interpolations have been required. The processing which the data has received, such as windowing and/or spectrum modification, is also indicated. This serves to provide a history of the data.

The routine displays a connected set of dots in the case of ATC data, since all samples are forced if any are valid in this case. Edge data are displayed as dots, connected when they are adjacent and in the same edge half. RAW data is displayed as unconnected dots. ENV data is displayed as vertical lines running from the maximum to the minimum values.

Current intensity is also displayed, as a warning in case the instrument is in use. This data is not stored or plotted associated with the data displayed, nor is it displayed if the digitiser is off line. (When first viewing a waveform captured, this gives the correct intensity. This is most normally the case when the digitiser is on line.)

Frequency domain magnitude data are displayed with a logarithmic scale. 0 dB represents the amplitude of a full scale single frequency input. The display is expanded if the maximum resolvable frequency exceeds 5 GHz, until the abscissa spans DC to just 5 GHz. A marker shows the mathematically ideal noise floor. The frequency step (resolution) is displayed.

Quefrency domain data are displayed in the same frame as the time domain data, with a flag indicating its actual form.

X.4.15. EDG

The EDG routine requests the 7912 to determine and transfer its edge data waveform, according to currently set parameters. The edge data represent a double valued set of horizontal samples which give the upper and lower pixels of the target which were detected as written in each pixel column. (These should **not** be confused with the **actual upper and lower limits of the beam path** as suggested by the 7912 firmware routines and their associated documentation. A more full description may be found with the Actual Beam Path (ABP) routine.) The routine does **not** automatically determine an edge waveform from a RAW data set if it is in memory, analogously to the ATC command determining a single valued function from edge data if already present. (See ENV.)

EDG responds with an error if the digitiser is not on line.

X.4.16. EMC

EMC allows the Execution of the Macro Command. A macro command is a set of up to 20 single legal DIOS commands which are invoked by the use of a single command. The definition command is DMC.

The macro defaults at powerup to being the commands required to perform a frequency compensation on a single valued time domain waveform.

Whenever EMC is encountered in the input stack, the previously defined macro command sequence is inserted into the input stream in place of the EMC mnemonic. An example might be the keystroke sequence PONPTMCEPPTMPOF, which will cause the activation of the progress printing facility, print the time, execute a cepstrum calculation, print the time and de-activate the printer, whenever EMC is typed. This would be useful for finding out how long a cepstrum calculation was taking, is it was being executed often. The mechanism is basically a time saving facility.

X.4.17. ENV

The ENV command causes an envelope waveshape to be determined from a raw data set. The command flags an error if the controller does not contain raw data. The routine returns a double valued set of data similar to EDGe data.

The resulting functions are treated as beam edge information by all routines except DSP, but are flagged as envelope data.

The main difference between ENV and EDG determination algorithms is that EDG will reject data points which imply a beamwidth which is impossible without the beam overwriting its own track. Thus ENV will return an effective "edge" data set if there is no overwriting evidence. ENV is a more general determination.

X.4.18. ESV

The ESV (Envelope to Single Value) command performs a conversion from envelope data to a single valued waveform. The command flags an error if the controller does not contain envelope data. The routine returns a single valued set of data indistinguishable from ATC data.

The routine effectively performs a detection function in conjunction with the ENV routine. The envelope shape of the input waveform is determined at each point by returning the difference between the two values associated with each element of the array.

The result is treated as a normal single valued ATC-type data set in order to permit the application of processing routines which require data of that form. Since the timebase setting is usually such as to make the application of any response compensation algorithms ineffectual, because compensation is not applied at the lower band edge, it should be noted that such application is meaningless.

X.4.19. FFT

FFT causes the controller to calculate the Discrete Fourier Transform of the current waveform. The routine requires that the current waveform be single valued, and in the time domain. The Cooley-Tukey fast algorithm is used. FFT reports an error if these conditions are not met.

In the HPL implementation, all data are stored with 12 digits of resolution. Calculation is carried out with guard digits. Underflow and overflow situations are handled without reporting an error by substitution of the smallest non-zero, and largest finite, values allowed by the machine. When the routine completes, a warning is issued if a default value has been substituted. Because the occurrence is not typically associated with a significant inaccuracy in the results, the warning is not printed even if the PON condition exists. Normally, noise processes in the digitiser prevent the need for use of defaults.

No scaling of the data is performed in FFT. Thus the resulting samples are artificially large. The IFT routine scales the reverse transformed samples to return precisely the same values (within the numerical rounding limits defined by the processor operating system). (See IFT.)

Transformation to the frequency domain does not erase the time domain samples.

X.4.20. HAN

This command causes a Hanning window to be imposed over the time domain data.

If the current waveform is not a single valued time domain waveform the routine will report a fatal error.

The "Hanning" window used is a second order one, defined earlier as a raised cosine:

Weighting Function = 1+cos(X)

Note that most references reserve the term Hanning (or Hann'd) for the cosine squared member of the family. This is the one used in this command. Any even order member of the family can be obtained simply by repetition of the command the required number of times. It would be simple to adjust the routine to deliver the first order member of the family, allowing any order to be obtained by repetition, but the second order is the most commonly preferred one, and thus this method is used because a single call defaults to the most common requirement.

The function consists of calls to two other routines, LHN and RHN, described in their proper places within this manual. The effect is identical, but the separate routines allow some flexibility as they are accessible directly by themselves. Each affects only its own half (Left or Right) of the window.

X.4.21. ICP

ICP calculates the inverse Cepstrum of the current waveform. The complementary function is CEP.

ICP requires there to be a cepstrum waveform in the current store. An error is flagged if this is not the case. Unlike CEP there is no uncertainty as to whether the inverse cepstrum can be numerically determined at all.

Where the echo parameters cannot be reliably determined from the cepstrum, such as is often the case because the echoes are "muddled" in the form of the signal, the echoes or mismatch corruptions may be largely removed by simply attenuating the cepstral delta (impulse) functions* and then taking the inverse cepstrum. This is most often the case when the echoes are frequency dependant, such as is the case with a capacitive or inductive connector fault in a transmission line system.

* This is possible using the FFT, SPT and IFT routines; these and other "time domain" DIOS routines accept quefrency domain input for this explicit purpose.

X.4.22. IDT

This command allows entry of the IDentification Text which is stored with a file, and which is used in file naming. The first four characters of an identification text are used in the mass storage file name, and thus must comply with system requirements for filenames. (See STW.)

Up to 64 characters of user text are allowed. All but the first four are not checked, so any keyboard character is legal. In the HPL system, spaces are not allowed in the filename, and so are forbidden in checked text.

The time of capture of a waveform is stored as ASCII text at the end of the IDT string. When the text is changed, the time remains unchanged. Whenever a waveform is moved from the 7912 to the controller, the time is appended or updated in the IDT string. The intention is that the text can end its description with the characters "at", and the time of recording will be handled by the system.

X.4.23. IFT

The IFT routine calculates the inverse Discrete Fourier Transform of a frequency domain waveform.

The IFT calculation routine is substantially identical to the FFT routine, with the exception that scaling is performed on the results so that the same magnitudes as were originally present in the time domain would be returned, if the frequency domain samples are not altered.

Transfer from the frequency to the time domain does not erase the frequency domain samples.

X.4.24. INI

INI causes a full initialisation and self test of the digitiser.

INI requires that the digitiser be regarded as off. If this is not the case a fatal error will be flagged. Note that the digitiser may be on (and will need to be switched on manually if the Measurement and Control Processor is not on line).

The routine executes a self test of the digitiser, acquires the target defect array, checks to see that it has not changed and warns the operator if it has (indicating some possible hardware problem or wear and tear), enquires the date, checks for a programmable vertical plugin, and then commences the warm-up calibration delay. Once this has expired, or if it is cancelled according to instructions in the system display, the last digitiser control settings are recalled and the TIM routine entered. This last step loads the correct time into the clock.

Note that the INI command does **not** initialise the RUNTEK programme. This is handled by routines not available to the user from DIOS command level. See the Appendices.

X.4.25. INT

INT fine tunes the intensity setting of the digitiser. It is intended for use only in unattended operation situations.

INT operates by examining 1) the total number of 'hit' pixels on the target, 2) the number of columns with no hit pixels, and 3) the maximum number of hit pixels in any column. The intensity is adjusted to produce a scan with approximately 4-10 hit pixels per column; however, the intensity is weighted heavily against any column having more than 15 pixels hit, and will not permit more than 30 to be hit. It is also weighted to increase the intensity to minimise unwritten columns.

Usage notes: The weightings are rather arbitrary. It appears that they are not very critical. The intensity will home in on a satisfactory value in most cases. It should be noted, however, that the routine can be 'fooled' by a badly selected timebase setting, which results in overwriting of trace parts. In this case the INT routine can leave large unwritten parts, because the overwritten sections force the intensity down. It should also be noted that the routine repeatedly calls SSW, ans so should not be relied upon in situations where trigger events are rare.

X.4.26. INV

INV inverts the current waveform. INV is equivalent to inverting the polarity of the incoming waveform with a hardware control. The inversion is noted in the displayed trace history information.

INV requires the current waveform to be a single valued time domain waveform. An error is flagged if this is not the case.

X.4.27. ISA

ISA invokes the improved signal averaging routine incorporated inside DIOS. It supersedes the Tektronix firmware routine SA. The mnemonic is an acronym for intelligent signal averaging, because the routine exercises simple judgement in the determination of the signal average.

ISA requires that the digitiser be on line, and that there be an ATC (or better an ABP) waveform in the controller for enhancement. A fatal error is reported if either condition is not met.

ISA requires an integral iteration count to be supplied. This is taken from the normal input stream, so that operation is not halted when ISA is called. The integer is normally a single digit, but as any integer is valid, there must be a way of indicating that a longer integer is being entered. This is effected by halting input execution if necessary with a semicolon, and then resuming when the number is keyed by entering a <cr>. Because of data transfer delays, large iteration counts will cause long execution times.

ISA repeatedly calls SSW, EDG and ITC. When the latest single valued waveform has been determined, it is compared to the original or partially averaged one. The latest waveform is progressively shifted up to plus and minus one division relative to the other, and the minimum mean square error between aligned samples is sought. If this error is acceptable, the waveshapes are deemed to be sufficiently identical and optimally aligned, and the last one is included in the running average. If the error is excessive, the routine reports that the waveshapes are uncorrelatable, and the last captured one is rejected.

There are two situations where this method will allow signal averaging but where the 7912 firmware one will not. This method controls trigger jitter. Each waveform is shifted (and extended) as required to compensate in software for the hardware uncertainty of the trigger. The triggering circuits of the 7912 have been found to be almost completely useless for signal averaging by the normal method with sweep speeds in excess of 10nS per division. Secondly, false triggerings cannot pollute the averaging process. Particularly in the very noisy environment where signal averaging is required, false triggering or triggering on events other than the one it is required to analyse is prevalent. A waveform which does not resemble the one used for enhancement will be rejected, so an automatic, correct-triggering filter is implemented after the event, but before the averaging has taken place.

Certain algorithm parameters have been selected for optimum operation of the process. The 'jitter scan width' of one division has been found to be more than adequate, though a little costly in terms of computation time. The need to make the jitter scan wider with increasing timebase speed has not been felt to be sufficiently important, but could be easily incorporated if necessary. It would reduce the wait sometimes.

The mean square error limit is not as critical as might be expected. Basically, waveforms either can be shifted to make a good match, or they are corrupted in some way and will have very large mean square errors. A value corresponding to an error of five pixels in each vertical searched is satisfactory.

There is no need to sample every vertical in calculating the error. In theory only every second needs to be sampled if the waveform is oversampled by a factor of two, etc. In practice sampling every five irrespective of the oversampling ratio gives a satisfactory compromise of speed of execution and seems not to be too infrequent at any level.

ISA continues searching for the required number of acceptable waveforms (given by the iteration count) and so will loop indefinitely if the events stop or the signal is lost, etc.

X.4.28. ITC

This routine executes an Interpolation To Centre, converting a double valued function such as an edge waveform to a single valued function, forcing a valid sample value at all 512 horizontal locations.

It can be executed if called from the keyboard but is automatically called in place of the digitiser's atc routine if an edge waveform is already present in the memory and an ATC is requested at the main command level. Thus it replaces the 7912 atc routine if executed after an EDGe command. It can also operate on an edge data set obtained using a routine replacing the 7912 routine. Further, it is called as part of the Actual Beam Path routine.

It executes an algorithm similar to and slightly superior to the one found in the digitiser firmware. A description of the algorithm follows.

The routine executes five steps. First it searches the two input functions for missing data. Any horizontal address missing either an upper or lower limit or with limits reversed or out of range is flagged as absent. Secondly, the simple sum of valid pairs is formed. (These two steps are usually executed conveniently in one move.) If all points are invalid, the routine exits clearing the data store, but without halting, which is automatic if an error is reported, and undesirable in unattended operating situations.

Next, the routine looks to see if the end samples are valid. If they are not, the endmost valid sample is copied to the end, in both directions. This differs from the 7912 routine, which can leave invalid samples (as different from zeroes) in the data. This may affect further routines. The copy count is included when determining the maximum number of iterations, though no interpolation is undertaken.

Next all gaps are filled with interpolated sample points. The slopes at the two ends of the intervals are determined, and these are then used to fill the gap by **cubic curve fitting.** This contrasts with the 7912 routine which uses linear interpolation only. This affects the waveform more adversely. In practice the cubically interpolated waveform appears more natural, whereas the linear one is visibly discontinuous when more than 3 or so samples have been interpolated. A cubic interpolation of over ten steps often goes unnoticed, and the spectrum reflects the improvement. This is especially true when interpolating over a step which has been lost. (This is the most common location of a momentary beam loss, since the writing rate is changing most rapidly there.) The cubic approach has been found to be misleading, however, when the waveform is noisy or patchy near the endpoints of the gap, as the two end points are used to determine a slope at the point of beam loss, and this is corrupted by the final point being out of place.

Note that if the gap has only one endpoint available at either end (that is there are two gaps separated by a small valid area, the cubic fit must be aborted in favour of a linear one. If the gap is only one point wide, linear interpolation is used directly. If the single endpoint was the last or first valid point, the slope is automatically set to 0.

The method of determining the endslopes may be improved. It has been observed that increasing the order of the fit beyond cubic opens the possibility for erratic fits (as is a well known phenomena in curve fitting to an arbitrary order on an arbitrary number of points). However, keeping the main fit at cubic, the endslopes can be determined by taking a weighted average of the slopes suggested by three of more points at the gap end. Such efforts have been tried, and have been found not to be rewarding.

The largest gap size is accounted for in the determination of the largest number of points calculated in a single interpolation.

X.4.29. KLW

This command allows a waveform file to be killed, erasing it from primary mass storage. Deletion from secondary mass storage is not permitted except under the controller's operating system.

The whole and complete waveform file name, including prefix and suffix with the correct case, must be specified. There is no method even under the controller's operating system of recovering a killed file.

A fatal error is reported if the file is not found.

A delay of six seconds is provided during which the name of the file to be killed is displayed and during which it is possible to cancel the kill command by pressing the reset key, which aborts DIOS operation.

X.4.30. KSW

KSW causes the controller to abort the execution of a routine which is currently hung in a wait state. The main use is to exit from a SSW command that is hung because no valid trigger has been received.

The KSW command (for Kill SWeep) is executed out of order; that is, it is executed immediately it is entered into the input stream. Thus it is a kernel command. The mnemonic remains in the input stream, but causes no action when encountered in correct serial fashion. Only the currently hung command is aborted, so that further commands will be executed. The command merely has the effect of falsely convincing the waiting routine that the wait is over. If a further routine in the stack hangs, it will not be cancelled by the KSW command that was first entered, but would have to be aborted with a second KSW command if required.

X.4.31. LCL

This routine returns local control of the digitiser's front panel functions. No operational parameters are altered.

The routine also ensures that the front panel button for requesting return to remote operation is enabled and can thus interrupt the processor.

The complementary function is REM.

X.4.32. LHN

LHN invokes the calculation of the Left half of a Hanning window. Along with RHN, the routine calculates and applies a Hanning window function to a single valued time domain waveform. (See HAN.)

The two halves of the window function are separately accessible as user commands to allow a degree of flexibility in selecting the window used on a given function. Situations arise where more 'information' is found in a localised part of a waveshape, because the function is faster moving at that point. When that information is close to one end of the captured time window, the use of a plain Hanning window could seriously attenuate the significant portion of the data. In this case it is very beneficial to apply one of the "half Hanning windows", LHN or RHN.

These operate by using half the bell curve of the applied window to pull the waveform at the lower information end of the trace towards the sample value at the other end of the data set. Thus rather than bringing the samples closer and closer to the centreline of the screen as they approach the ends of the window, as does HAN, the half functions compress the waveform at their end (Left for LHN and Right for RHN) to the line defined by the other endsample (Right for LHN and Left for RHN). This leaves half the data completely unaltered. The resulting frequency domain effect is not the same as a plain Hanning window, but the degradation is small, and worthwhile for preserving resolution of components betrayed end localised signal activity.

X.4.33. LIT

The LIT command allows the passing of literals to the digitiser. The purpose of the command is to make it possible for the user to send any set of characters to the digitiser, thus accessing any command in digitiser firmware that may not be directly supported by DIOS.

When called the input stream is redirected until a <cr> is entered. All text typed up to that point (up to 80 characters) is sent over the HPIB to the digitiser. After transmission, the digitiser is polled to check for errors, and if there are any, the error reporting routine is called. As with all input redirecting commands, the LIT command cannot be stacked.

X.4.34. LRx and LSx

These commands cause a Left Rotate or a Left Shift of a single valued time domain function, respectively. The complementary functions are RRx and RSx.

An error is flagged if there is not an appropriate waveform in the current store.

The variable x is an integer. It indicates the length of the shift or rotate in horizontal divisions of display. (The display is taken as having the customary 10 divisions of width. Since there are 512 samples, the shifts are approximate in length.)

Rotate commands move the sample values in a circular fashion, so that the same samples are present before and after the rotation. The shift commands, on the other hand, destroy the samples 'shifted' off the end of the array. These are made up by duplicating the sample at the end away from which one is shifting. (That is the first sample is duplicated in right shifts and vice versa.) These commands are useful for positioning or truncating a waveform before windowing.

X.4.35. MEN

This command displays a short form menu of DIOS commands on the graphics output device.

The menu is stored on the mass storage device as a data-file with entries separated by <cr/[lf]> characters. The MEN routine sorts the entries into alphabetical order before displaying them. Thus, additions may be made to the file at any time, and the added information will be sorted when displayed.

MEN halts execution of the command stack. Keyboard input is redirected, and a key press is required to proceed from screen to screen, or to return keyboard input to the command stack once all screens have been viewed.

X.4.36. NAM

This command permits the entry of a new Jobname.

The jobname is a label of up to 8 characters which is associated with the state of the machine. Default control settings are stored in a file whose name uses the first four characters of the jobname. This file is used for setting the digitiser during the powerup sequence. The first four characters of the label must conform to the system requirements for mass storage labels.

The jobname is displayed at powerup, and is used to identify secondary mass storage media. It is held on the system data-file "DATEK". The system file is updated every time the jobname is specifically accessed.

It is possible, by means of changing media without accessing the jobname, then re-initialising the controller, to change the usage data stored on the media in the system file. This is not advisable. The controller keeps track of the number of "cold starts" to which the tube has been subjected, using the system data file. This allows the machine usage to be gauged. Thus the jobname should be accessed after media change on the primary mass storage device (main disk drive).

X.4.37. OFF

The OFF command shuts down the digitiser in an orderly fashion.

OFF requires that the digitiser and the Measurement and Control Processor, which has control of the external activation line of the 7912, be on line. A fatal error is flagged if this is not the case.

An orderly shutdown involves the update of the machine usage file DATEK, and the shutting down of the main digitiser power supply. A restart, which will return power to the Scan Converter Tube, will thus increment the usage counter in DATEK.

The complementary function (on) is not executable directly, but is called by the INI routine.

X.4.38. PAC

The PAC command packs a time domain waveform into the (left) half of the storage space.

PAC requires a single valued time domain function, and will flag an error if this condition is not met when it is called.

PAC replaces the first 256 samples in the array with every second sample from the original array. The remaining 256 samples are filled with the value of the 256th sample.

Because of the inherent oversampling which occurs in the Scan Converter Tube, little information is typically lost. The process modifies the timebase setting record to give the correct time location of events in the trace. The purpose of the routine is to prevent the data set appearing as oversampled, which tends to render it impossible to calculate the Cepstrum, due to small magnitude values occurring in the higher FFT output samples.

PAC typically can be used twice, with minimal loss of features. Windowing may be carried out before or after the packing. It is best carried out afterwards if the original signal starts and ends at similar levels, since less information is lost.

X.4.39. PDR

PDR prints the directory of DIOS user commands stored in RAM. (That is, the commands available both with or without primary mass storage.)

The advantage of PDR over MEN, apart from the hardcopy output, is that it actually searches the programme and extracts the labels of valid commands from the programme text. The commands are recognised by a their label format, which is two capital letters and one number or three capital letters. No other type of routine is allowed this label format. Thus the PDR directory is continually and automatically updated. It has two main uses:

1) It allows accurate checking of MEN without manual searching of listings where errors could arise;

2) It is independent of mass storage.

PDR does not see external (disk-resident) commands. These commands are however, already virtually listed in that they are also identified by their label form, which is unique in the mass storage directory.

PDR calls PON. If the printer is down, a fatal error is flagged.

X.4.40. PLT and PLx

PLT instructs the system to copy the current waveform from memory to the system plotter with default options. PLx, where x is in the range 1 to 9, instructs the system to copy the current waveform to the system plotter with special options.

RAW data cannot be plotted.

An error is reported if PLT or PLx is invoked and the system plotter is absent or does not identify itself as an acceptable device.

The bits of the number x specify options according to the following rules. Bit 0 prevents expansion of data in the frequency domain, displaying the full FFT output. (Where expansion is used, Lagrangean interpolation is used to smooth the output for hardcopy.) Bit 1 specifies that the plotter is to use the second half of the platten or advance the paper as appropriate. Bit 2 specifies that the image is to be superimposed over the last, exactly as if a SUP request was pending. (See SUP.) Bit 3 requests that a plot of phase be added to the frequency domain magnitude. The option is ignored if the current store is not in the frequency domain.

The plotted result contains the same information as a DSP request, but adds a graticule and the programme version number, and deletes any intensity data. (See DSP.)

PLT and PLx assume an A4 size piece of paper, left hand bottom justified on the plotter platen. A3 plotters without autofeed can thus support two sheets without operator intervention, using the bit 1 option.

X.4.41. POF and PON

These commands turn the automatic printing of diagnostic, progress and warning messages off and on respectively.

POF will automatically be called if there is a printer error, such as an out of paper error, by the error handling routine. PON will not execute if there is a printer error, but no halt to execution is caused.

The automatic printing is designed to allow automatic documentation of operation for fault finding or as a record of unattended activity. With PON active, long routines such as CEP will print progress reports. Any warning messages which are associated with non-fatal problems are also printed. PON and POF may be stacked to obtain a record of the operation of certain routines within a group.

X.4.42. PTM

PTM prints the current time on the system printer.

PTM, like PDR, does not require the printer to be active (PON) to operate. PTM will abort by itself, therefore, in the event of a printer error, without halting execution.

PTM is useful for logging times in unattended operation.

X.4.43. RAW

The RAW command transfers raw target data from the digitiser to the controller.

RAW data is the least processed form of target writing information available from the 7912. It consists of two arrays of indeterminate size. The first is a set of pointers into the second array. The values contain in coded form enough information to determine whether any pixel on the target was written or not. The information is presented in the form of locations on each vertical line of pixels where the state of pixels (written or unwritten) toggles. Up to 14 toggles per vertical can be stored by the 7912. For further details, the user is referred to the 7912 operators manual.

In order to economise memory usage the arrays are stored within the RUNTEK programme by packing the fields of the floating point arrays used for data storage. The 'red' kernel utility subroutine, described in Appendix E, covers this procedure. Nevertheless, it is possible that the data transferred in response to a raw data request will overflow the (fixed) storage space available. In general, however, this event is rare, because a 'clean waveform' will require at most 2 toggles per vertical, thus giving an amount of data comparable to an EDGe waveform. Immediate transfer to a default file on disk could simply be implemented in a virtual memory scheme if this ever became a problem.

X.4.44. RCL and RCx

The RCL command loads the controller's memory and then the digitiser with the control settings stored in the primary setup memory file on mass storage.

RCx, where x is a digit from 0 to 9, loads the controller and digitiser with the contents of the secondary setup stores.

Attempting to load a non-existent file will cause a fatal error, halting operation.

The complementary functions are STO and STx. See the STO and STx command descriptions for details of file naming, and availability considerations.

X.4.45. RCV

This command ReCoVers the system from a fatal error in the digitiser firmware. It basically turns the instrument off, clears the bus, turns it on and then passes control to the RCL subroutine which loads the last saved control settings.

This command has been found necessary because of the 7912's susceptibility to crashing if exposed to EMI. It can also become irreversibly tied up if there is an interruption to a data transfer while it is in progress. (The IEEE bus can be freed up, permitting an external device such as an HP Measurement And Control processor to remove and re-assert the external activate signal on the appropriate 7912 rear panel connector.)

X.4.46. RCW and RWx

The RCW command recalls a time domain waveform on the primary or secondary mass storage device to the controller and establishes it as the current waveform store. SWx, where x is a digit from 0 to 9, operates similarly on a separate group of user files. For a discussion of file naming, see the STW and STx command descriptions.

Transfer of files between secondary and primary mass storage and between media must be done with TAP and DSC, or in the controller's operating system. Media swapping is not advised - see the explanation under the command NAM.

Four characters only must be entered to define a waveform when using RCW. These are taken from the normal keyboard data stream, and thus the name is effectively embedded in the command stack. The programme flow is not halted unless the stack is empty.

Because only the suffix part of a filename is read in, the routine does not normally allow specification of the file type. The routine will always load the most processed one of a group of files with the same identification suffix. That is, if there is a d type (single valued) waveform, this is loaded. Any edge or raw data files are not loaded in preference to the atc file. ATC (or ABP, etc) files are selected first, then edge, and finally raw waveform files. To access a precursor file, such as the raw data from which an atc waveform was derived, the processed file must be killed or renamed. This system has been found to be very convenient in practice. When a file is being recalled from secondary mass storage for display (a more rare requirement) the file type specification is allowed and demanded.

When loaded from secondary mass storage, a file is automatically displayed by control passing to the DSP routine.

X.4.47. REM

This routine removes control from the digitiser's front panel, without altering any operational parameters.

The routine sets local lockout mode, which prevents return to local mode without the consent of the processor. Thus, no local adjustments will be possible should the controller go off line.

The routine also ensures that the front panel button for requesting return to remote operation is enabled and can thus interrupt the processor.

The complementary function is LCL.

X.4.48. RHN

RHN invokes the calculation of the Right half of a Hanning window. Along with LHN, the routine calculates and applies a Hanning window function to a single valued time domain waveform. (See HAN.)

The two halves of the window function are separately accessible as user commands to allow a degree of flexibility in selecting the window used on a given function. Situations arise where more 'information' is found in a localised part of a waveshape, because the function is faster moving at that point. When that information is close to one end of the captured time window, the use of a plain Hanning window could seriously attenuate the significant portion of the data. In this case it is very beneficial to apply one of the "half Hanning windows", LHN or RHN.

These operate by using half the bell curve of the applied window to pull the waveform at the lower information end of the trace towards the sample value at the other end of the data set. Thus rather than bringing the samples closer and closer to the centreline of the screen as they approach the ends of the window, as does HAN, the half functions compress the waveform at their end (Left for LHN and Right for RHN) to the line defined by the other endsample (Right for LHN and Left for RHN). This leaves half the data completely unaltered. The resulting frequency domain effect is not the same as a plain Hanning window, but the degradation is small, and worthwhile for preserving resolution of components betrayed end localised signal activity.

X.4.49. RRx and RSx

These commands cause a Right Rotate or a Right Shift of a single valued time domain function, respectively. The complementary functions are LRx and LSx.

An error is flagged if there is not an appropriate waveform in the current store.

The variable x is an integer. It indicates the length of the shift or rotate in horizontal divisions of display. (The display is taken as having the customary 10 divisions of width. Since there are 512 samples, the shifts are approximate in length.)

Rotate commands move the sample values in a circular fashion, so that the same samples are present before and after the rotation. The shift commands, on the other hand, destroy the samples 'shifted' off the end of the array. These are made up by duplicating the sample at the end away from which one is shifting. (That is the first sample is duplicated in right shifts and vice versa.) These commands are useful for positioning or truncating a waveform before windowing.

X.4.50. RSE

RSE processes cepstrum data to Remove set Echoes arising from a shorted Stub on the signal line. The data indicating the primary echo delay and relative magnitude is stored on the file "tekcep" on primary mass storage. (See SEP.)

RSE requires that the primary mass storage device be on line, and that the current waveform be in the cepstral domain. An error is flagged if these conditions are not met.

A stub matches only at a single frequency. A wideband signal will be corrupted by such a configuration, the corruption taking the form of a series of echoes. The first echo is delayed by the time it takes to traverse the stub twice, and will have a magnitude defined by the amount of signal which is deflected into the stub, multiplied by that fraction which is transmitted forward upon the inverted return from the short. The result in the cepstrum is a series of delta functions, inverting in sign and decreasing in magnitude at each step, and spaced by the stub delay. RSE is a very specific routine for dealing with this situation. It is also, however, a model for other such specific routines, as would be required to deal with a specific mismatch situation.

X.4.51. SAx

SAx, where x is an integer in the range from 0 to 9, basically invokes the 7912 built in signal averaging (SA) routine. SAx is included only for reasons of speed, since it is completely obsoleted by the ISA (Intelligent Signal Averaging) routine, which is external to the 7912.

SAx instructs the digitiser to take 2^x separate digitisations of the input signal using normal triggering. The 7912 firmware computes an ATC waveform for each set of samples, and adds the sequential sets directly, column for column. It then transfers the single valued sum function. The SAx routine divides the results by 2^x but retains floating point format, rather than effecting a bit shift type power of two division as is intended. (One location in the 7912 manuals indicates that the 7912 is supposed to do the bit shift division, but the unit used here did not. There has presumably been a change in a ROM revision which has not been noted in the manual.)

The usefulness of the SA routine is severely limited at the higher sweep speeds, owing to trigger jitter. For this reason the output of the SAx routine cannot be trusted nearly as highly as the ISA routine. Users are advised to use ISA unless there is a need only for a rough result to be quickly determined.

In all other respects, SAx acts like SSW and ATC combined.

X.4.52. SEP

The SEP command allows the setting of Echo Parameters for use by RSE and any other (external or resident) cepstral related commands.

SEP merely stores the delay and magnitude of the mismatch echo on the data file "tekcep", used by RSE. Up to 16 data pairs can be stored on this datafile for reference by processing routines. While RSE uses only one pair, others may be organised to use pairs other than the first, or to use many pairs.

X.4.53. SET

This command displays all the digitiser mainframe and horizontal plugin control settings on the system graphic display and allows alteration of them, remotely from the instrument. It is useful mainly for adjusting display intensity or similar fine tuning in situations where the instrument is located remotely from the controller.

Operation halts when the routine waits for user input, and the keyboard input stream is diverted around the normal key handling routine.

Settings are updated immediately after each control adjustment is entered, so that errors in setting are reported at once. A null entry terminates the routine and returns control to the command stack.

X.4.54. SPM

The SPM command uses a previously stored correction array to be used to modify the magnitude and phase of the frequency domain data held in the controller. The correction array is held on mass storage. (The correction data file is called "TEKfrq" in the HPL implementation listed in this thesis. It is calculated as described in chapter 5 by the utility programme "CALCOM".)

It is probable in any situation that the frequencies of the elements of the correction array will not coincide with the samples resulting from the Fourier transformation. Thus sixth order Lagrangean interpolation is used separately in the orthogonal axes of the complex elements of the array to develop values between those available.

The correction array is extended with unit elements as required. Thus the interpolation algorithm smoothly shifts to applying no correction beyond the maximum frequency available from the correction array.

Once the spectrum correction has been executed, control is passed to the SPT routine.

SPM will report an error if invoked when there is no spectral data in the controller memory.

X.4.55. SPT

The SPT routine smoothly truncates the spectrum of a set of samples. The shaping is Gaussian, with an order of one over the standard deviation, and introduces a 3 dB attenuation at either 3GHz or half the Nyquist frequency, whichever is the smaller.

The routine has two main uses. It can be used to remove out of band noise which occurs in the input waveform. Such signals result from the methods of processing the signal inherent to the 7912 series instruments, as well as numerical processes such as interpolation and fourier transformation. These various effects have been discussed in their respective locations in this text.

It is also used to continue the spectrum modification which is invoked by SPM. Here, it overcomes the problem that the correction array is limited to the bandwidth of the instrument, and thus cannot dictate correct action beyond that point. (See SPM.)

SPT will report an error if invoked when there is no spectral data in the controller memory.

X.4.56. SSW

This command causes the digitiser to arm and await a trigger, and thus to acquire a single waveform. Several procedures are involved in this process.

If there is no identification text present (see IDT) identifying the signal, a warning is issued.

If the horizontal plugin is not in single sweep mode, this mode is set. If the digitiser mainframe is not in digitise mode, this mode is set. The mainframe is instructed to set the operation complete flag (OPC) when the sweep is acquired.

The routine then loops until the operation complete flag is detected as set. The loop checks for an abort instruction (see KSW).

When the sweep is completed and the OPC flag set, the requirement to set the flag is cancelled, the current time appended to the identification text is updated, and control is returned.

X.4.57. STO and STx

The STO command transfers the current digitiser control settings to the system controller and then to the primary setup memory on mass storage. There is a primary setup memory which is maintained on all mass storage units, as well as a set of 10 secondary setup stores which are available only with the main mass storage unit.

The STx command, where x is a digit from 0 to 9, operates similarly to the STO command, but stores into one of the 10 secondary locations available only on primary mass storage.

Both commands will overwrite the previous contents of a store if it exists already. The primary storage location on secondary mass storage exists inherently (see TAP). A "new file" message is displayed when the store is created. Both commands may used with different media, thus expanding the available number of stores by changing media.

Transfer of files between secondary and primary mass storage and between media must be done with TAP and DSC, or in the controller's operating system. (See TAP command.)

On primary mass storage, primary setup files are named as "cs" followed by the first four letters of the jobname (see JOB). Secondary files are named "csTEKx". Thus primary files are associated with a jobname, and several can co-exist on one media.

On secondary mass storage, there is provision only for one store, on the assumption that media will be exchanged with job.

On primary mass storage, a fatal error will be reported if the media becomes full. This problem is overcome in the case of control setting stores on secondary mass storage. This decision stems from the intention that secondary mass storage will be used for field work, where unattended operation is probable. The only shortcoming is that a trace of settings under automated control cannot be kept in a looping command sequence without primary mass storage.

X.4.58. STW and SWx

The STW command stores the current time domain waveform on the primary or secondary mass storage device. The file is given a six letter name. The name has a two letter (lower case) prefix formed by combining the letter w with a type identifying letter, and a four letter suffix which is the first four letters of the ID Text (see IDT) associated with the waveform. The type letters are d for single valued waveforms, e for edge data and r for raw data. SWx, where x is a digit from 0 to 9, stores the current time domain waveform in a temporary user store. The temporary stores are named with the same prefixes as the primary stores, but have suffixes of the form TEKx.

The primary mass storage device will thus hold (space permitting) one single valued, one edge and one raw data file for each different ID Text used, as well as ten temporary user files of each type. The secondary mass storage device is not permitted to store raw data, nor are temporary user files allowed. This latter decision is basically a space, time and tape wear based one.

Quefrency data is considered to be a time waveform for storage purposes, because of the effort involved in its determination.

SWx will overwrite the previous contents of a store if it exists already. A "new file" message is displayed when the store is created. STW will NOT overwrite existent data. The store is aborted and execution continues with a warning message if an overwrite attempt occurs. This prevents accidental data destruction. A consequence of this is that a looping task cannot create a file for each loop that is executed, without either killing the old file or updating the identification text. (The latter would typically be done in a task routine. See TAS.)

If necessary, both commands may used with different media, thus expanding the available number of stores by changing media. Transfer of files between secondary and primary mass storage and between media must be done with TAP and DSC, or in the controller's operating system. Media swapping is not advised - see the explanation under the command NAM.

A fatal error will be reported if the media becomes full.

X.4.59. SUP

The SUP command indicates that the next display or plot command is to produce an image superimposed upon the image already present on that device.

The command operates by setting a flag which is checked by the next display or plot command. The flag has one effect on both the graphic display device and the system plotter, as well as whatever functions are specifically required by the device in question. In all cases, the writing on the image is displaced so that it does not collide with the previous image's writing.

On the graphics device, erasure of previous images is inhibited. The writing is displaced downwards, but the trace location is unaltered. Trace coincidence is marked by highlighting.

On the plotter, pen colour is changed. Colour identifies the writing with its appropriate trace.

X.4.60. TAP

TAP initialises secondary mass storage media, and then stores the requisite files on the tape in order to permit operation of the digitiser using it alone. The complementary function is DSC.

All data on the secondary media is lost. TAP requires that both primary and secondary mass storage devices be on line. A fatal error is flagged if this is not the case. The files stored are 1) a blank autostart file which may be used to initialise the digitiser automatically, or by the user for any purpose in the field, 2) the machine usage record file DATEK, 3) the frequency compensation array, 4) RUNTEK, the DIOS kernel and resident commands, 5) the menu data file, 6) the control setting file and 7 onwards) the data (waveform) stores.

The purpose of differentiating secondary and primary mass storage devices stems from the availability of both compact inboard (tape) storage and superior outboard (disk) storage mechanisms in the HP9825. It is intended that the digitiser be usable for a specific task (once prepared) with only the minimum of equipment. The transfer of files to the secondary and portable storage system would be undertaken just before switching to the portable mode of operation and going out into the field. Upon returning the files (acting as a control token) would be transferred back to primary storage.

The routines relating to the two systems could easily be altered to reflect small (microfloppy) and large (host system) mass storage facilities, if say, an Integral computer was substituted for the 9825 in later versions. If desired, the restrictions on use of certain (external) commands can be relaxed with improved secondary storage capability. Nevertheless, it is recommended that the distinction be kept (requiring the use of TAP to produce new secondary media) in order to ensure that the usage file DATEK remains associated with the digitiser, by the effective passing of a 'token' (the data files sought) between storage levels.

X.4.61. TAS

TAS is the label given to any user written routine that tests for satisfaction in a looping task, and decides whether to loop or continue. (EG, whether the set task in a remote and unattended application has been completed satisfactorily. This might be constituted by the mass storage unit running short of space, or by a waveform of some magnitude being captured on the last digitisation, etc.)

There is a TAS routine provided as a model, and it can be modified to produce the desired effect, or it can be replaced by the user's own routine. The action to be taken is specified by a flag (which is HPL flag 5 in V3.0) identified in the DOC file. (See DOC.) The flag must be set to loop. When cleared, the routine does not loop to the first command in the stack. The flag is not altered by DIOS at any time.

It is for use in building TAS routines that the DOC facility is mainly provided.

X.4.62. TIM, TMx and TPx

The TIM command allows the setting of the system clock. The TPx and TMx commands, where x is a digit from 1 to 9, advance and retard the system clock x integral hours.

Execution of the stack is halted for entry by TIM, but not for the other commands. Time is always entered in 24 hour format, with a decimal point as the delimiter, so that it is a single legal real number. (This is the numeric return format of the internal 'time' subroutine.)

If the logging printer is active when any of these commands is executed, the old time is printed along with the updated time. (See PON, POF and PTM.)

The TP and TM commands are useful in systems (such as the 9825) where the clock is external to the processor, and has limited counting range. They allow swift adjustment of the time if an integral number of hours have elapsed and rollover was not detected by the processor.

X.4.63. TRD

TRD allows the marking of specific waveform X and Y coordinates on plotted output. It operates in a fashion analogous to CUR, but on plotted output rather than volatile graphic display. Its main use is in the hardcopy determination of risetimes ("Tr").

When called, the numeric keypad is used to issue commands to TRD. A message is placed in the system display to remind the user what the keys do.

TRD moves the red pen over the paper along the track of the waveform in the current store. When so commanded, it places a small red marker on the plot, and writes the time coordinate next to the mark. Numbers are printed in engineering notation, that is, with the exponent forced to a multiple of 3, and deleted if zero.

When called, the input stream is redirected. Execution is terminated by entry of any character that is not recognised. When such a key is pressed, it is entered into the command stack, the plotter is returned to its idle state, the keyboard input stream is restored to the normal handler, and control is returned to the calling routine or kernel. This means that typing any valid DIOS command automatically terminates TRD and executes that command. Thus TRD cannot be stacked. (It is an interactive command.)

X.4.64. TWx

TWx, where x is an integer from 1 to 9, causes a Tukey window to be imposed upon the data. The integer defines the size (in tenths of full width) of the central and hence unaltered, flat section of the window.

TWx like all window functions will flag a fatal error if there is not a single valued time domain (ATC type) waveform in the central waveform store.

A Tukey window consists of two half raised cosine bell shaped sections to be applied to the waveform. However, unlike the Hanning function from which it is derived, the two bell halves do not necessarily meet in the middle. A central plateau is left untouched, while the end sections are attenuated. This arrangement is useful because it allows a tradeoff between attenuation of interesting parts of the waveform in the centre of the array, and the degree of corruption of the frequency spectrum.

If the two routines RHN and LHN have been set up to accept the passing of extent parameters, the routine can be comprised simply of two calls to these routines with suitable parameters to indicate the data range to be processed.

A. OSCILLOSCOPE PERFORMANCE CHARACTERISATION

The specifications typically quoted for oscilloscopes today are inadequate to correctly and completely characterise these instruments. The standards setting process appears to have been out paced. This may be a result of proliferation in the number of basically different designs of instruments capable of giving a voltage/time graph, and the long time acceptance of the basic CRO format. This appendix does not attempt to set this right, but to explain some of the fundamental concepts required to appreciate the limits of these devices. These may seem intuitively obvious once described, but are often not fully understood.

A committee has been set up by the Institute of Electrical and Electronic Engineers (IEEE) to draft a standard for the specification of sampled systems (digitising instruments). The crucial concepts are, however, the same in both the continuous and sampled situations. This standard is expected to be prepared in draft form before the end of 1985.

The important properties are simply those which would be important on a graph made with pen and paper, added to the figure of response speed, or bandwidth. These are namely:

(1) The range of values available in both axes, which would be reflected in the vertical sensitivity and horizontal timebase speed;

(2) The resolution and accuracy in both axes, which depend as

shall be noted upon a variety of factors including noise, linearity, spot size, and number of samples;

(3) The frequencies of signals which can be handled, (possibly a level dependant property), defined by amplifier bandwidth and slew limits, writing rate, etc.

These things, though fundamental, may be defined by very different mechanisms in different instruments. A description of each follows.

A.1. Vertical Sensitivity

This is perhaps the most straightforward specification. It is usually quoted in units per division, however, which requires a knowledge of how many divisions are available. A de facto standard is eight. (Surprising in a decimal world.) A more useful way of describing this is to quote the signal required for full deflection. This is being done to a larger extent with the newer instruments which use conventional Analog to Digital Converters (ADCs). It is a straightforward flow on from the converter designer's jargon.

In many instruments the vertical sensitivity can be extended down to a smaller figure (higher actual sensitivity) by sacrificing another specification. In a plain analog CRO, this is usually small signal bandwidth. In an 'equivalent time' instrument (Sampling CRO or SCRO - see Chapter 1) it is usually the time required to develop a complete sweep, because signal averaging is performed as described in the discussion of sampler operation in

Chapter 2, Section 1. This built-in tradeoff possibility should not be overlooked as it can lead to expectations beyond what is available from the instrument.

A.2. Horizontal Sensitivity

This specification is equivalent to timebase speed, when the X axis is in units of time. Again there is a need to account for the number of divisions on the screen (usually ten). In discrete (sampled) systems the sensitivity may be implied by the total number of samples taken in one acquisition cycle, together with the rate of sampling.

Note, however, there is less need to be concerned with an instrument's sweep speed specification. This is because an adequate lower limit (maximum sweep speed) is set once other factors are known. A properly designed instrument will be able to achieve this limit.

This limitation on the requirements of the timebase arises from the limited bandwidth and resolutions. Consider a waveform containing all frequency components in the pass band, displayed so that it is (only just) completely resolved. This waveform on a continuous instrument is faithfully displayed, and the samples in a discrete instrument permit complete reconstruction of the original band limited signal. A faster timebase speed would be equivalent to oversampling, which offers no new information. In the continuous case, the image is expanded, and in the discrete case the oversampling might produce the reconstruction for the eye, but the extra samples are not necessary for that function. Indeed, information will have been lost, because there is a section of the image now effectively off screen, which would be on screen at the lower speed. Assessment of this requirement is relatively simple, once the resolutions and the factors defining them are understood.

A.3. Resolution and Accuracy

Although often different in each axis, the questions of resolution and accuracy in both may be addressed simultaneously. The resolution is a measure of how small a feature can be resolved or separated, or equivalently, of how different two values must be in order to tell them apart. The accuracy is a measure of how reliably the reported value reflects the true, infinitely precise value. The former of these values is a relative property, while the latter is absolute. In certain applications either may be the important factor, or in general the worse figure may be quoted.

Both can be reported in the units concerned with the particular axis, or as fractions of the full scale value. In this thesis, both will be measured as fractions of the Full Scale Deflection factor, (FSD), since this will vary frequently without affecting the resolution fraction. They will in general be reported as a binary fraction, expressed as 'bits', possibly fractional. This is referred to as the concept of 'effective bits'.[46,42] The actual resolution or accuracy will be equal to the full scale value divided by two to the power of the number of bits quoted. This is a clearly digital jargon, but of course will perfectly well describe an analog situation where the resolution, though not so precisely pegged, is none-the-less real. The term Dynamic Range, familiar already, is synonymous with the concept of effective bits.

A.3.1. Resolution

Resolution, in either axis, has been defined above as **the difference which must exist between two values in order that they should be reported as different.** Initially this might seem to be set purely by the smallest unit which can be read out, expressed as a fraction of the full scale reading. This is the meaning which is most usually taken. However, a little careful thought will indicate that noise should be taken into account also. If the random uncertainty is sufficient that a single value may be reported as two separate ones, or more importantly two values differing by an amount which the noise free system would resolve, may be reported as identical, then the resolution is **not** primarily set by the smallest unit read out, but by the maximum random uncertainty (or noise or error value) acting in conjunction with it. In most oscillographic systems, noise does not contribute significantly to the resolution. Sampling Oscilloscopes operating without signal averaging form the major exception to this rule.

In a digitally sampled instrument, the resolution is set by the ADC, and is usually equal to the number of bits of resolution of the ADC itself. In an analog instrument, the resolution is not as well defined, and must be estimated in the light of the information presentation mechanism. For example, a typical CRO screen offers 10 cm horizontally and 8 cm vertically. Resolution will be effected by the trace width, or more precisely the spot size. An experienced operator could resolve as little as 0.2 or as much as 0.05 of a division, if the tube is of high quality, the focus good and the signal sufficiently free of noise. The resolution normally expected is 0.1 division. This corresponds to 6.3 effective bits.

A.3.2. Accuracy

Accuracy,[43-47] or more precisely the errors affecting it, can be divided into three categories: Random, Systematic, and Signal-related. Random errors are noise, which implies that such errors cannot be removed or negated once added to a signal. These arise in most stages of a system, and are usually linked to the physics of the devices and processes involved. The theory of noise is well developed, and it is not discussed further here.[48]

Systematic errors are repeatable and thus predictable. In general a source of systematic errors distorts the output (in one or other axis) according to a transfer function which is well defined, though not necessarily either well known or elementary. The transfer function may be viewed as that equation which would transform a set of ideal values (ignoring other errors) into the set actually obtained. In a practical situation the error might be approximated by a function which is linear, exponential, polynomial, or a combination of the above.

In simpler analog instruments, such as the stereotyped low frequency CRO, these systematic distortions are lumped with the random errors. An accuracy figure is then quoted, usually as a fraction of the full scale value. In more recent instruments which digitise their results, there exists the possibility of applying systematic correction, where the systematic error function is known or can be determined. This has successfully been applied in some devices to date, particularly digital voltmeters and low frequency transient recorders.[40,41]

The third kind of error has only recently been brought to light.[43,44,46] It appears to have been investigated only in digitising systems. The basic observation is that accuracy degrades from the static performance situation as the signal is made to present greater rates of change. It is usually reflected in the instrument specification by a reduction in the number of effective bits for higher frequencies at full scale input. Whether the error is random or systematic, and if systematic what its transfer function might be, has yet to be reported specifically for any instrument. The error is regarded in one instance as arising from uncertainty in the aperture time of the sampler, and it is implied that this error is random.[46] The author feels that this error may in time be characterised to the degree that some software correction could be applied to a group of samples considered together.

A.4. Bandwidth

The bandwidth of an oscillographic instrument is defined as that range of frequencies which are displayed, with the specified accuracy, to within another specified limit, usually 3 dB. Many assumptions in, and implications of this definition, are overlooked even by frequent users of such instruments.

One implication of limited bandwidth affects accuracy. The normal 3 dB point selected to define bandwidth means that frequency components at that point will have their amplitudes reduced by about 30%. If the response falloff were due to a single real pole only in a continuous system, there would be a 2% reduction even at 1/5 of the nominal cutoff frequency, and an error of more than 1 part in 256 at 1/11 of the nominal cutoff. This implies that the 2% accuracy figure quoted for a typical 10 MHz CRO is doubled at 2 MHz, and that the 1 LSB error of an instrument with 8 bit resolution and having its bandwidth of 100 kHz defined by analog amplifiers rather than the sampling process, is doubled at about 9 kHz!

Frequently, when bandwidth is quoted, an instrument is assumed to have a smooth extended response. That is to say the frequency response curve is assumed to approach the Bode asymptotes not only in the pass band but also for some way beyond. In addition, it is assumed to have a single pole response. This is rarely true in practice, except for machines well below the state of the art in frequency response. Most oscilloscopes have an overshoot specification, which is an admission of non-Gaussian response.[66] One particularly fast analog instrument specifying a 1 GHz bandwidth is quoted as being "...usable to 1.4 GHz...", which is much less than one would expect from the initial claim[5], if simple response falloff is assumed. The author has observed a similarly drastic falloff in a reputable company's instrument rated at only 2 MHz, though with a very impressive sensitivity specification in exchange.

In a single pole situation, bandwidth (which equals the maximum frequency) is approximately related to risetime by the equation:

$$\text{Risetime (in pS)} = 350/(\text{Maximum Frequency in GHz}) \qquad \text{(A1)}$$

This equation proves to be reasonably accurate even when the response is not due to a single pole.[29]

When discussing bandwidth, an engineer tends to assume that the **small signal** bandwidth is intended. Indeed, since specification is usually made at or near full signal capacity, any false assumption would be in the pessimistic direction. Such might occur if the deflection amplifiers were to have slew limits that set the full power bandwidth below the small signal bandwidth. This is avoided in oscillographic equipment, as far as has been ascertained,[48] for the following very good reason: Seriously misleading results could be obtained if the slew limitation was present in, say, a deflection amplifier. This is because the Slew Induced Distortion (SID) does not prevent the processing of signals beyond the quoted bandwidth, but does grossly distort the waveform. The result of attempting to display a large signal beyond the power bandwidth would be a waveform totally unlike the original signal, with the operator being potentially unaware of the extreme distortion.

(It should be noted here that the term distortion can have two meanings: A waveform may be distorted in one fashion by having various spectral components increased or decreased in amplitude with respect to the others, and in another fashion where spectral components not previously present in the signal are introduced. These two effects both distort the original waveshape, potentially deceiving the observer, but the second is considerably more objectionable. A simple analogy may be taken from audio engineering, where tone controls are introduced to produce the former type deliberately, but where great pains are taken to prevent the latter from occurring at all.)

The matter of slew limitation is present in oscillographic instruments in a rather safer form. Cathode Ray Tubes (CRTs) are specified as having a maximum writing rate. As the beam moves across the target, some impression must be made to record the path. The electron gun can deliver only a finite beam current. If the beam traverses beyond some critical speed, it may be that insufficient charge or energy is conveyed to the target to leave a detectable record of the trace. The maximum speed is the writing rate. This is similar in principle to the slew rate of an amplifier, but of course depends upon the signals in both axes, and the degree of retracing present in the spot path. It is a common limiting mechanism in very fast CRT based instruments. This limitation is 'safer' because the mode of failure is not distortion of, but loss of, the image of the signal.

B. THE EFFECT OF CASCADED SAMPLING GATES ON SIGNALS

B.1. Introduction

Rigorous determination of the minimum disturbance which might be caused to a transmission line by a sampling gate of the preferred two diode type is extremely difficult. This appendix recaps on the earlier description of a sampling gate circuit, develops an equivalent circuit which a sampling gate presents to the signal transmission line, considers a practically soluble, pessimistic simplification of this circuit, and presents an analysis of the effect of a cascade of such simplified discontinuities on a wideband signal.

B.2. Equivalent Gate Loading Circuit

Figure 2.3 presented the circuit of a sampling gate. As indicated in Chapter 2, the pulse shaper which developed the triangular or trapezoidal pulse presented to the capacitors consists primarily of a shorted length of transmission line. When not sampling, the diodes appear as non-linear capacitors with some series resistance. Assuming that the bias supplies and pulse amplifier input appear as resistances, as is the case in current designs,[6,8] the circuit perceived from the line can be drawn as in Figure B.1. In that circuit C_{jd} and R_s represent the capacitance and series resistance of the sampling gate diodes, R_p is the equivalent resistance of the bias resistance network, C_s is the value of the sample storage capacitors which form two arms

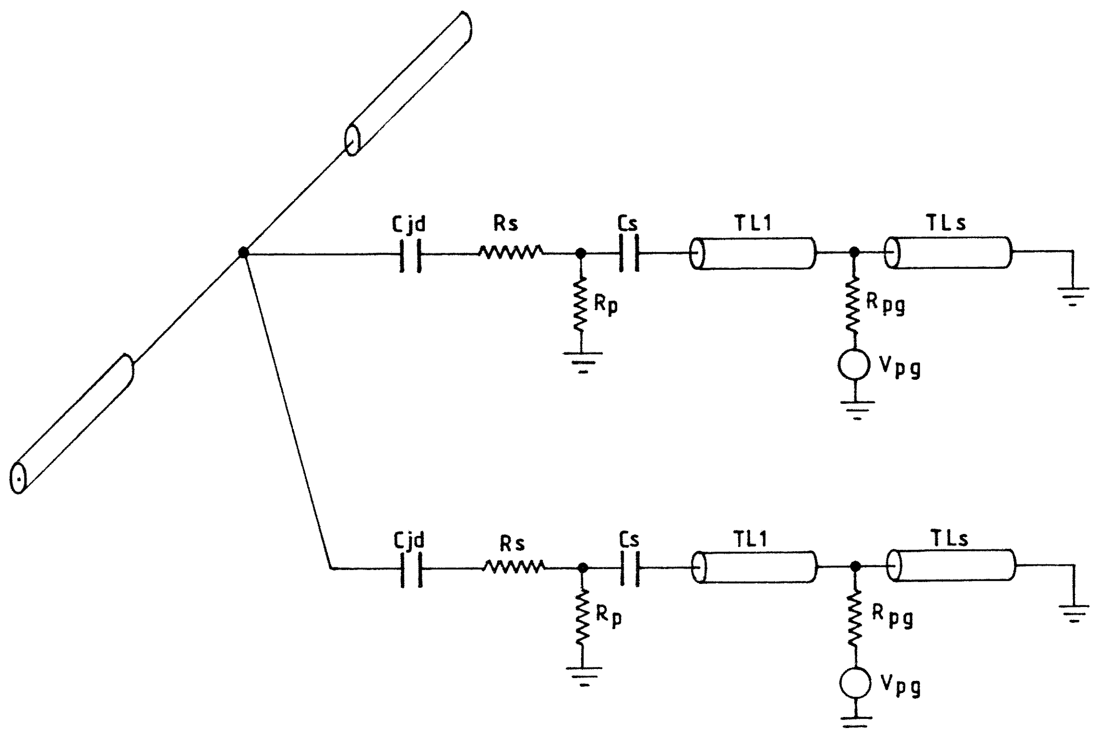

FIGURE. B1

of the bridge. V_{pg} and R_{pg} form the step generator equivalent circuit. TL_1 is the line conveying the pulse to the bridge, and TL_s is the shorted length of line used to produce the required pulse width. When not sampling, V_{pg} will be zero.

In a practical circuit, it is safe to assume that R_p and R_{pg} will be negligibly large. Also, C_s will be much larger than C_{jd}. Any error introduced by assuming that C_s swamps C_{jd} will be much smaller than the change in C_{jd} experienced as a result of changing reverse bias levels with signal fluctuation. The circuit thus reduces to that of Figure B.2.

Consider the situation where a step wave front is propagating down the transmission line formed by the series connection of many gates. When the signal arrives at a junction between the line and the circuit representing a passive gate, it encounters a frequency dependant discontinuity. To frequencies at which the junction capacitors C_{jd} appear to be an open circuit, there will be no discontinuity; to ones at which they appear as a low impedance, there will be a discontinuity, with a second transmission line leading to ground from the junction. In reality all components of the signal will fall between these extremes. The energy in the signal will be divided, some being reflected back from whence it came, some being passed, and some being diverted through the capacitor. That diverted through the capacitor but not absorbed by the resistance, will be reflected (or nearly so) at the short circuit, and will return to the junction. Here the same three way process will occur, producing

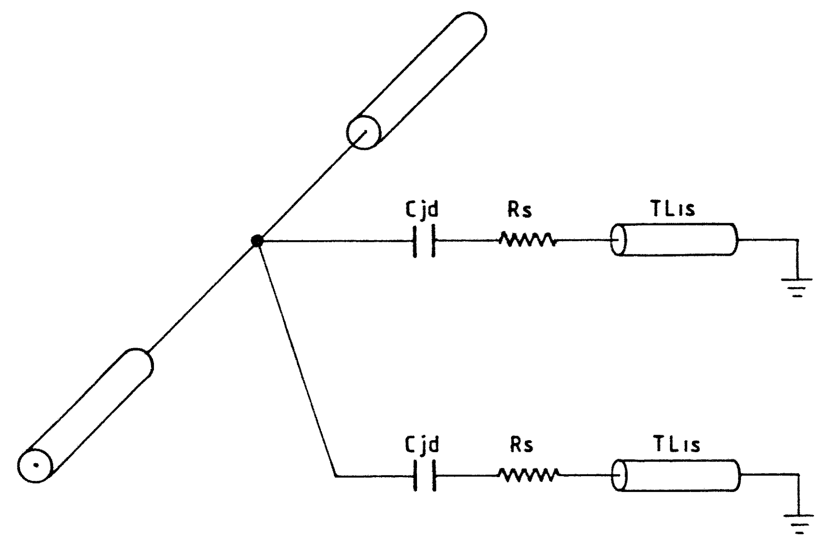

FIGURE. B2

artifactual "echoes" travelling both forward and backward down the line, and retaining some energy within the side arm for the process to repeat. This process will of course produce a (theoretically infinite) set of partially reflected waves of decreasing magnitude, added to the original signal on the main line, causing distortion. Precisely how significant the secondary reflected returns will be compared to the original division is not obvious.

As indicated before, some improvement of the situation is to be had from removal of distributed capacitance of the main transmission line in the vicinity of the junction. Were the capacitance of the diodes to be directly connected to ground, a maximum benefit would be gained from the technique. However, the delaying effect of TL_{1s} and the energy capturing effect may need to be taken into account. The second problem is that any subtracted capacitance is actually distributed, and it turns out that it is not possible to regard it as being a lumped subtraction. Finally, the capacitance of the diode junctions, C_{jd}, is not constant, but will of course vary with reverse bias of each diode, and therefore signal level. This is the reason why the two side arms of the circuit in Figure B.2 have not been combined at this time.

B.3. Diode Junction Capacitance

Figure B.3 shows the typical variation of junction capacitance with applied voltage for HP diode type HSCH-5311.[38,51,52] The graph includes a package capacitance allowance. The curve is described by the equation:

$$C = C_p + C_0(1+V_j/V_p)^{-m} \qquad (A1)$$

where

C is the total capacitance,

C_p is the constant package/mounting capacitance,

C_0 is the zero bias junction capacitance,

V_j is the applied potential,

V_p and m are constants defined by diode construction.

The constant m varies between 0.5 and 0.3 for abrupt and graded junctions respectively, and cannot be precisely predicted for any given junction. A simple small-capacitance measurement instrument was constructed with the aim of determining m with more accuracy, and determining if it varied within a batch.[53] The very small change in the **junction capacitance variation** with m coupled with the small capacitances involved prevented accurate determination with the instruments available, although the curve shape itself is verifyable. More complex instruments[54] offer little improvement in resolution or accuracy for values in the ranges concerned here, so the determination of m was abandoned, and the optimistic or pessimistic end of the theoretical range is taken where appropriate.

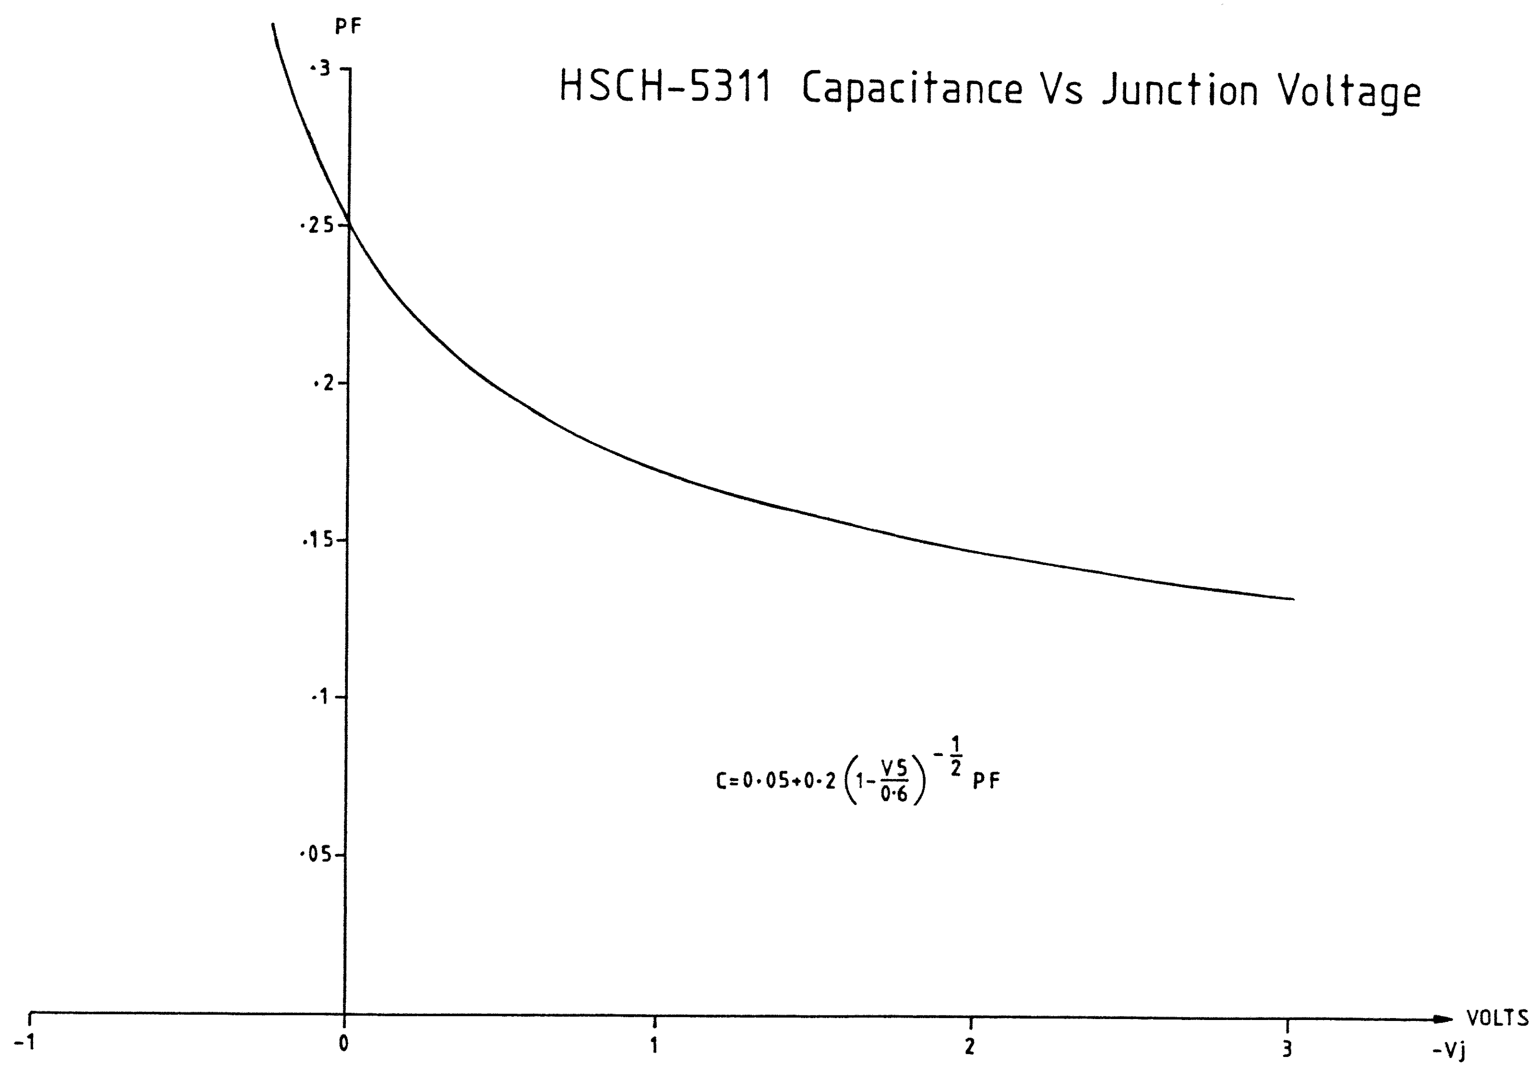
HSCH-5311 Capacitance Vs Junction Voltage
PF
·3
·25
·2
·15
·1
·05
-1
0
1
2
3
VOLTS
-Vj
$C=0\cdot05+0\cdot2\left(1-\frac{VS}{0\cdot6}\right)^{-\frac{1}{2}}$ PF

FIGURE. B3

The reverse breakdown voltage of the diode HSCH-5311 is 4 volts. This compares favourably with the values of 2 volts estimated or specified for the diodes in other sampling systems.[6,8] This implies that a signal peak of somewhat below 2 volts of either polarity will be handled safely by a gate using these diodes. It also implies that a reverse bias variation of approximately zero to 3 volts or so is to be expected. (This could be reduced by limiting the signal input to much less than the reverse bias level, but this would directly compromise sampler sensitivity, as the noise floor would be unaffected.) A simple calculation using Kirchoff's laws, the sampler circuit and equation A1 above, indicates that the total capacitance will have a minimum of approximately 0.3 pF, and will vary by up to 0.05 pF with varying line level, at least.

B.4. The Complete Model

The circuit of Figure B.2 lacks any representation of the removal of distributed capacitance. Because a transmission line is defined by the distributed elements which form it, and because it is described (in the lossless case) by its characteristic impedance, length and group velocity, it is possible to model the removal of distributed capacitance from a small line section by substituting a suitable length of line of different impedance for the line which has had capacitance removed.[55] The calculations are not difficult. A remarkably complete model of the gate discontinuity is thus afforded by the circuit of Figure B.4.

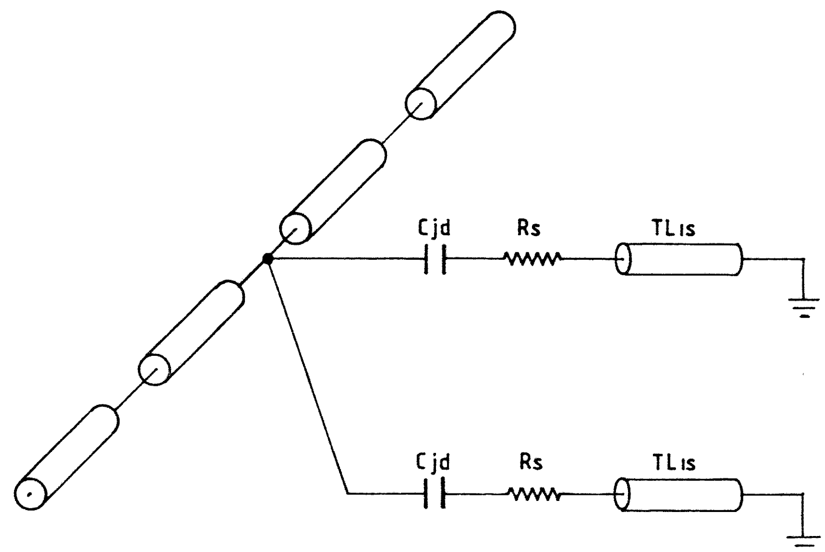

FIGURE. B4

B.5. The Simplified Model

Attempts to analytically solve for a transfer function to describe the output signal which would result from an arbitrary input signal passing an arbitrary length cascade of the approximation cells of the form of that in Figure B.4 have not been successful. This is not surprising in view of the complexity. It was thought that an advanced application of Phase Integral Methods might afford some progress in the absence of the side transmission lines if required, but such approaches have been abandoned in favour of numerical ones.

Initially, the circuit of Figure B.4 can be modelled in SPICE.[52,56] In order to reduce computation time, which becomes critical when many samplers are cascaded in the model, some simplifications have been considered. By comparison of the results of modelling with measurements made on a single simulated gate assembly, and by comparison between the simulation results of different models, the simplest **acceptable** circuit for modelling will be selected. The model will be expanded to obtain a prediction of the effect of cascading.

The pair of side arms differ in only a minor way, namely that their junction capacitances may not be equal. Rather than modelling the capacitances as variable, non-linear elements (diodes), they can be replaced by constants. Following this, the side arms may be combined into a single arm. It is simple to verify that the results are not significantly different, but that the computational effort is greatly reduced. As will be evident later, this is a necessary as well as a convenient simplification.

The next point to be taken into consideration is that the loading capacitance is not directly connected to ground, but is grounded via a transmission line and a series resistor. Thus it is not likely to be the case that the optimum value of distributed capacitance for removal is the same as the junction loading capacitance. An iterative procedure on the single gate model reveals that the minimum discontinuity actually results when considerably less equivalent capacitance is removed - around one quarter. This suggests that subtractive removal will be of little assistance in alleviating the discontinuity.

The line discontinuity (without the side arm echo) predicted with the model has been compared to that measured on a jig carrying a pair of beam lead diodes of the appropriate type. (HP type 5082-2716 diodes were used. These are the type superseded by the 5311, and substantially identical to it.) Measurement of so small a discontinuity proves to be difficult, as it is not only barely visible on the most sensitive scale of a Tektronix 7S12 Time Domain Reflectometer, (TDR), but of considerably smaller magnitude than the discontinuity found at the SMA connectors of the jig. However, the 12.4 GHz bandwidth permits the features to be resolved, and the actual discontinuity is found to be comparable to the model value as shown in Figure B.5. As the

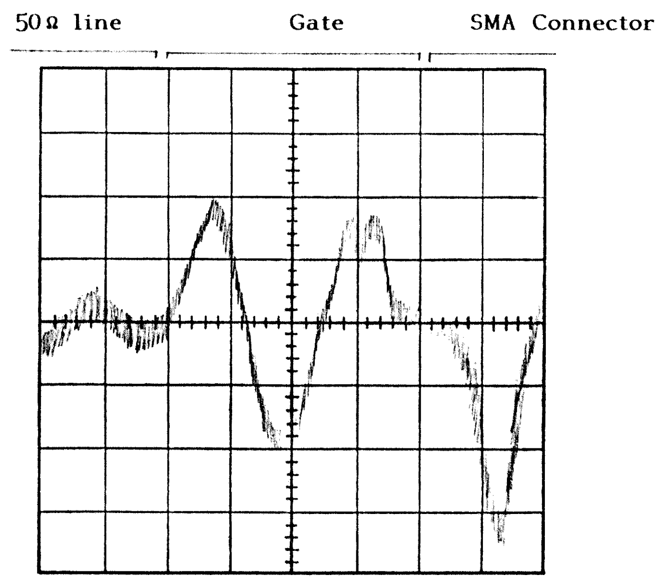

Measured reflection of Sampling Gate Model
(Scale: Vertical 0.005ρ/div, Horizontal 50ps/div)

Comparison with SPICE prediction

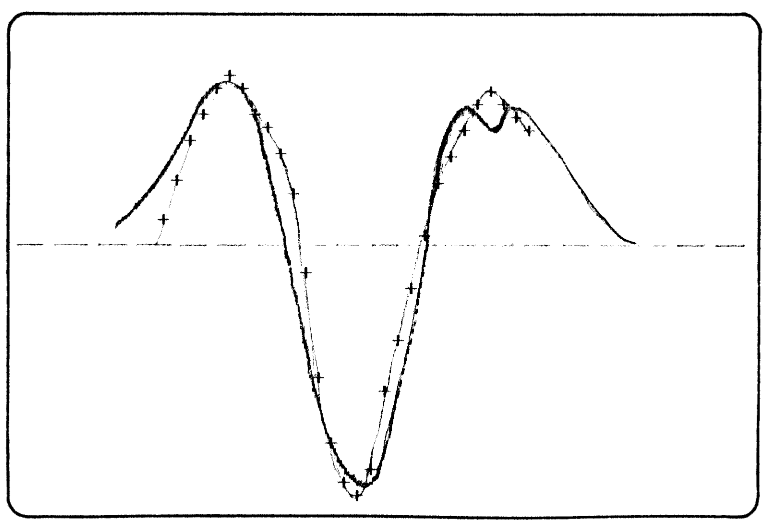

FIGURE. B.5

instrument resolves the small patches of higher impedance line both sides of the central connection, the image from the TDR is actually less flat than that from the model with a nominal 2 GHz signal bandwidth. Note that the small peaks of the TDR display resulting from the removed capacitance are approximately equal in area to the trough due to the loading arms. Also note that the echoes are not recorded on either trace. This problem obviously cannot be dealt with by altering the circuit components, and must be dealt with otherwise, if at all.

Comparison has been made between the results of the above model, in which there is some basis for confidence, and other ones. Discarding the transmission line from the side arm to simplify it, apart from altering the optimum value of distributed capacitance required, produces very different results. These suggest a much less distorted result when capacitance deletion is re-optimised. Substitution of a graded set of three shorter transmission lines, in an attempt to model the gradual transition which would in reality occur where capacitance subtraction is implemented, offered little improvement. A great simplification, consisting of replacement of the model section with a simple capacitance equal to the uncertainty (0.05 pF) in true capacitance, placed in shunt across the line, proved too crude. Although it allowed a larger number of sections, the discontinuity was reduced and the result deceptively clean (though not ideal). On the whole it appears that the model cannot be simplified without compromising the reliability, nor do greater complexities offer improved accuracy.

B.6. Limitations of the Results

The modelling has ignored certain points which would produce a result less acceptable than that which has been obtained here. These are not incorporated because the results obtained even in this reduced case have proven sufficient to discourage the author from proceeding with the design. No connectors have been introduced to the cascade, although they have been measured to produce larger disruptions than a gate. Normal line losses have been ignored. The diode junction capacitors have been taken as constant, which is not the case.

A signal containing all frequencies of concern will clearly suffer some distortion from each mechanism. Thus a step is chosen to test the model. An ideal step function has all frequencies present, with their amplitudes falling off as the inverse of frequency. The step used was less than ideal, and was passed through a single pole filter with a turnover at 2 GHz. The resulting waveform exhibited rounded corners and a risetime commensurate with the design bandwidth of 2 GHz. In short it represents a typical, but testing, worst case example.

When cascaded, it is found that SPICE 2G.7 implemented on a VAX 11/780 with 2 Megabytes of core memory will not accept 100 samplers without exceeding the programme's capacity. It will accept 50, however, taking typically in excess of **15 CPU hours** to complete an analysis of duration equal to the line's total length. This gives some idea of the difficulty involved in analysing the line performance.

Before presenting the results, it is necessary to indicate that the **particular** results in this appendix are **representative only** of what would actually be the case. Considerable variation in the details can be obtained with minor changes. Examples might be altering the length or characteristic impedance of the pulse delivery transmission line, or the value of diode series resistance. Such differences would arise between samplers beyond the designers control, as well as under it. While values which have been derived from actual experience have been used, these might change or be changed in the course of realisation. A very considerable amount of investigation has been carried out as to what alternatives could be employed to advantage. The conclusion has been that the general implications are not subject to easy change.

B.7. Cascade Model Results

A typical input file for SPICE is reproduced in figure B.6. The input step and output step after 50 gates with 250 ps temporal spacing, have the following characteristics:

Attribute	Step at Input	Step at Output
10-90% risetime:	200ps	425ps
Overshoot:	<< 2% (0)	approx 30%
Settling time (2%)	375ps	> 2ns
Guaranteed Monotonic?	Yes	No

The rising edge of the step has clearly suffered considerable degradation. (The "overshoot" which may be visible on the input pulse is a consequence of reflected energy from the cascade, and is not present on the original input.) If the pulse were processed by a passive network without delay, the flat of the step would be expected to settle. In fact this is not the case here. The level is corrupted or "unsettled" for some nanoseconds after the step, by energy delayed in the cascade and side arms. This effect corrupts the waveform with small re-injected signals. These signals, added to the actual waveform, make determination of the true level to more than approximately 4 bits of accuracy impossible. The relative magnitude of the summed reflections is sufficient to effectively hide detail in the original signal.

```
SPICE INPUT FILE for cascade modelling
* Subcircuit to model a single sampler
.subckt sam 1 2
TL1 1 0 3 0 Z0=50 TD=200PS
TL2 3 0 4 0 Z0=60 TD=25PS
TL3 4 0 2 0 Z0=60 TD=25PS
CDIODES 4 5 .31PF
RDIODES 5 6 15
TLSIGGEN 6 0 0 0 Z0=50 TD=150PS
.ENDS
*
* Pulse generator arrangement
VSRC 102 0 AC 2 PULSE(0 2 0 100PS)
RSRC 102 101 .1
CSRC 101 0 800PF
RS 101 1 50
*
* The array of (50) samplers defined by subcircuit SAM
X1 1 2 SAM
X2 2 3 SAM
X3 3 4 SAM
X4 4 5 SAM
    +
---------- This sequence is abbreviated here for convenience
    +
X49 49 50 SAM
X50 50 51 SAM
*
* Control cards
*
.OPTIONS ITL=8000 LIMPTS=1001 NODE ACCT
.TRAN 25P 15N
.PLOT TRAN V(51) V(1) V(26)
.END
```

Figure B.6

B.8. A Best Case Result

One modification to the pulse generator circuit permits the exchange of pulse drive requirement for reduced side arm reflection. The basic observation is that the relatively high level of re-injected signal depends upon the side arm transmission line terminating in a short circuit. While a shorted section is required for the development of a pulse from a step, this section need not be connected directly to the delivery line. A resistive network could in principle be introduced, providing a better termination as viewed from the sample delivery line, but wasting drive pulse amplitude.

In existing fast samplers[8,24] the gating pulse is used at full amplitude. However, where the pulse need not be of so short a duration, a step of sufficient speed and greater amplitude can be developed.[20,21] This can be divided to feed several samplers. Division and subsequent delivery by transmission lines (which was discussed in Chapter 2) does not relieve the immediate problem, but indicates that use of certain pulse sources offers amplitude above that needed for the gating process. Reference to the device specifications and technology available[28,38,30] indicates that there is again a tradeoff between step magnitude and cost. In order to avoid further assessing cost details until the technique is proven potentially useful, an optimistic estimate is made of the excess amplitude which could be made available without advancing to a significantly more costly arrangement. A factor of 5 is taken. (There is a sensitivity tradeoff involved as well, which is taken to its worst case.)

The cascade model can be adjusted to simulate the situation by removing the shorted termination at the end of the side arm transmission line, TL_{1s}, and replacing it with a slightly mismatched resistive one. The mismatch introduced is obtained after some calculation, and of course depends upon what attenuation is used. (50 millirho, or a reflection coefficient magnitude of 0.05, is to be expected for the case above.) Again, the input step and output step after 50 gates with 250 ps temporal spacing, have the following characteristics:

Attribute	Step at Input	Step at Output
10-90% risetime:	200ps	400ps
Overshoot:	<< 2% (0)	approx .25%
Settling time (2%)	375ps	approx 600ps
Monotonic?	Yes	Yes

As expected, the rising edge is not improved significantly. The overshoot problems arising from delayed energy are all but eliminated, and the degree of corruption is reduced. In this case the bandwidth is not saved, but the resolution within the pass band is not compromised.

B.9. Conclusion

This appendix has indicated that signal distortion in a cascade of gates is difficult to assess exactly. It has shown, however, that even in the absence of certain distorting effects, making optimistic assumptions where there was uncertainty in the design, and assuming only half the number of samplers which are desired in a cascade, that the resultant signal quality will in all probability not exceed a particular level. Whether this degree of distortion can or need be accepted is discussed in Chapter 2.

Comment is made in Chapter 5 regarding the feasibility of removal of mismatch induced reflection signals. It is concluded there that the technique of Cepstral Analysis is not capable of achieving this in a general situation with few reflections. It cannot, therefore, be considered as feasible in this Appendix.

C. A LOW COST TECHNIQUE FOR BONDING BEAM LEAD DEVICES

C.1. Introduction

In the course of investigating the feasibility of designing a sampling gate suitable for an MSGTD instrument, the fabrication and performance measurement of microwave circuit assemblies with beam lead devices was undertaken. The surface mounting of these devices was required for the investigations. The equipment required for such fabrication is a costly capital outlay, not justified in the early stages of such an investigation.

The author, faced with a need to use beam lead devices, but lacking access to suitable equipment, investigated a number of alternative techniques. One such method, devised in the Air Navigation Laboratories by the author in conjunction with Mr Jim Bruderer, has proven so reliable, is so remarkably free of the need for capital outlay and is of sufficiently unusual nature that it is reported here. The technique, although unsuitable for either fast or mass production, is employed still for research purposes well above the frequency of a few GHz for which it was originally intended.

C.2. Bond Requirements

A beam lead device has a useable lead length typically measured in the low hundreds of microns in length and less in breadth. A beam lead diode may be only 700 microns in its largest dimension.

Figure C1 (referred to later) depicts a typical device placed over the groove in an etched substrate, which might be finline or a thick film MIC surface.

The bonds on the leads have four important requirements: They must physically hold the device in place; they must provide "good" electrical contact; and they must not risk harming the device when being made. Finally, of course, the bond must be durable, not corroding through or aging to an extent where the integrity of the contact is lost with time.

Most practical bonds that produce adequate electrical contact also provide sufficient physical strength to restrain the device. Thus the first requirement is readily satisfied once the second is satisfied.

C.3. Currently Available Techniques

The third of the four requirements is responsible primarily for eliminating conventional bonding proceedures such as soldering. Apart from the temperature which the silicon would have to withstand, great hazard exists for the device due to the surface tension if volumes of solder comparable to the device itself are present. In practice it is impossible to prevent the solder running over the device or "swallowing" it whole when employing conventional soldering methods. The difficulty involved in placing small solder pellets at appropriate positions and then applying sudden controlled heat to effectively solder the minute leads eliminates typical brazing techniques altogether.

Thermo-compression bonding is the preferred technique, but requires considerable equipment outlay. Other commonly employed methods require even larger outlays.

The author has investigated the use of wirebonding machines of the kind used for bonding flying leads to IC pads. These prove to have insufficient energy delivery capacity to effect a satisfactory weld of the beam lead to the substrate, probably due to the mass of the substrate. Attempts to attach a lead to the device and then to the substrate, mimicking the lead from an IC pad to the pin of its package, also failed. The substrate has large thermal mass and requires significant surface preparation, while the device is not easily restrained and supported for the bond without effectively increasing its heatsinking, which impedes bonding.

Micro-spotwelding has been investigated. This presents some hazard to the diode not only from the voltages available at the point of bonding, but also from sheer physical damage done to the device when the probes place pressure on the lead. The pressure required to prevent arcing is quite large. A considerable mechanical jig would be required to apply the pressure, restrain the device, and yet not harm it. If arcing occurs, the lead is eaten away rather than being partially melted. A clear analogy is the welding of thin sheet metals by means of an arc welding machine. The currents required to bond successfully and the size of the tool needed to apply them are analogous to using a large rod and high current settings on a thin sheet. Electronic switching of welding impulses was investigated, with the aim of eliminating energy delivery without good contact, thus eliminating arcs, but this proved impractically complex in itself.

C.4. The New Technique

The novel method reported here consists basically of gluing the device in place with a conductive adhesive. Epoxy glues are available which surprisingly produce very satisfactory physical and electrical bonds. The advantages are clear: there is neither heat nor electrical energy applied throughout the whole operation; the bond is effected slowly (as the epoxy cures) so that there is an opportunity to alter positioning, unlike flash methods; there is no requirement for pressure or for the lead even to mate perfectly with the bonding (supporting) surface; the equipment required consists of no more than the microscope and probes required for inspecting and moving the devices anyway. The disadvantages are that the application of the glue requires a little patience on the part of the operator, and the glue requires some time to cure. Neither presents a significant problem in a research situation.

The conductive glue consists of an epoxy liquid with an apparent suspension of metallic particles. The metallic particles are small in comparison to the dimensions of the lead width. Some property of the glue suspension ensures continuity across the glue masses. It is assumed that it somehow guarantees contact between suspended particles.

(Although it was believed for some years that the technique was original, it is now known that it has been used at one other research establishment, Telecom Research Laboratories in Clayton, Victoria.)

C.5. Physical Layout

Figure C.1 shows a beam lead diode bonded to a substrate (in this case etched printed circuit board). The board is produced so as to have a channel wide enough to take the diode body, but narrow enough that its leads can span the channel. The glue and device are applied to the board with suitable fine tools, as described in the manufacturer's instructions for handling.

Figure C.2 shows a similar situation, but here the device is initially attached to the board using a small amount of insulating silicon putty. This approach, devised later by Dr J G Lucas, offers the advantage that the diode cannot become accidentally shorted by stray droplets of the conducting adhesive. In addition it holds the diode in place while the conducting epoxy is applied to the delicate leads.

DIAGRAM of DEVICE-SUBSTRATE BONDS

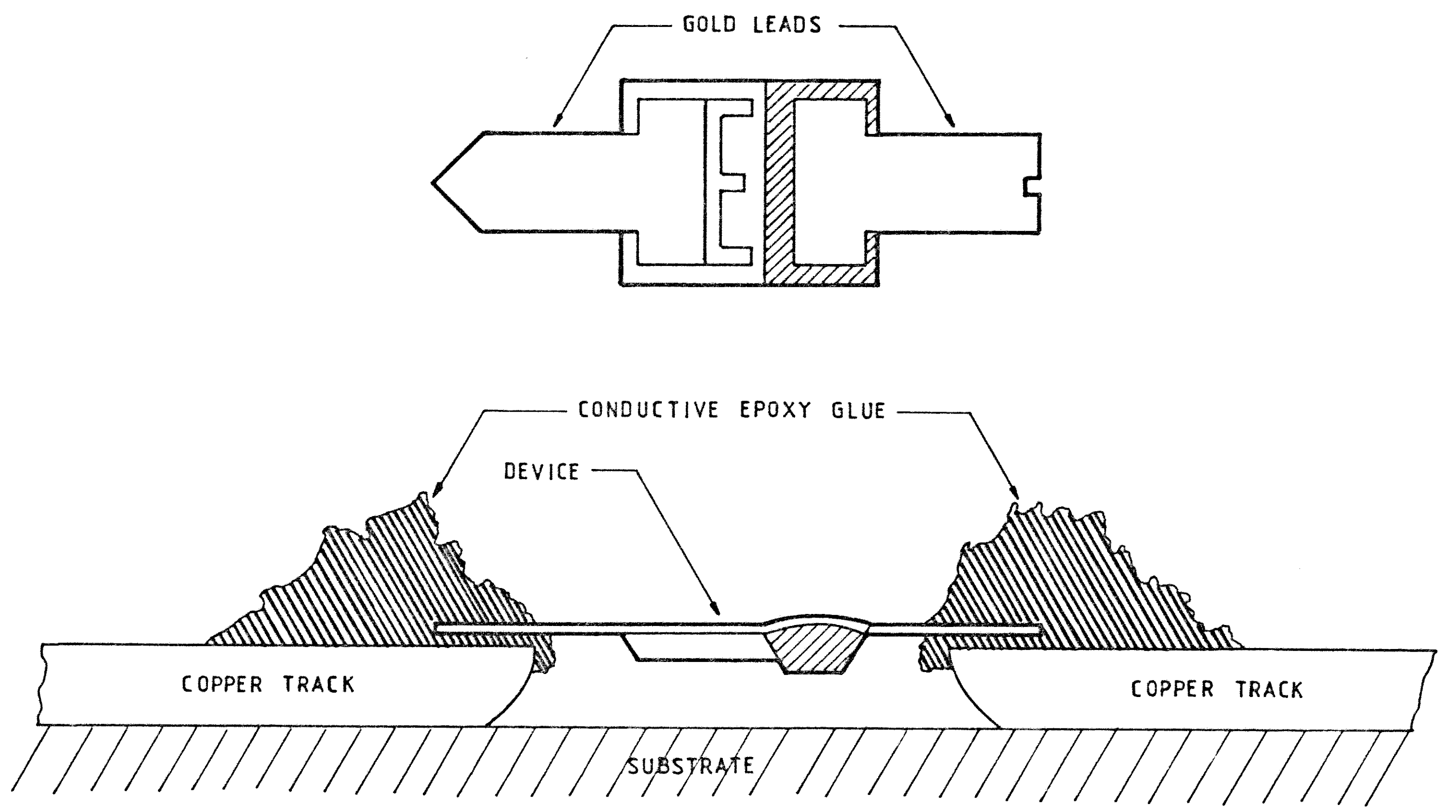

FIGURE. C.1

DIAGRAM of DEVICE-SUBSTRATE BONDS

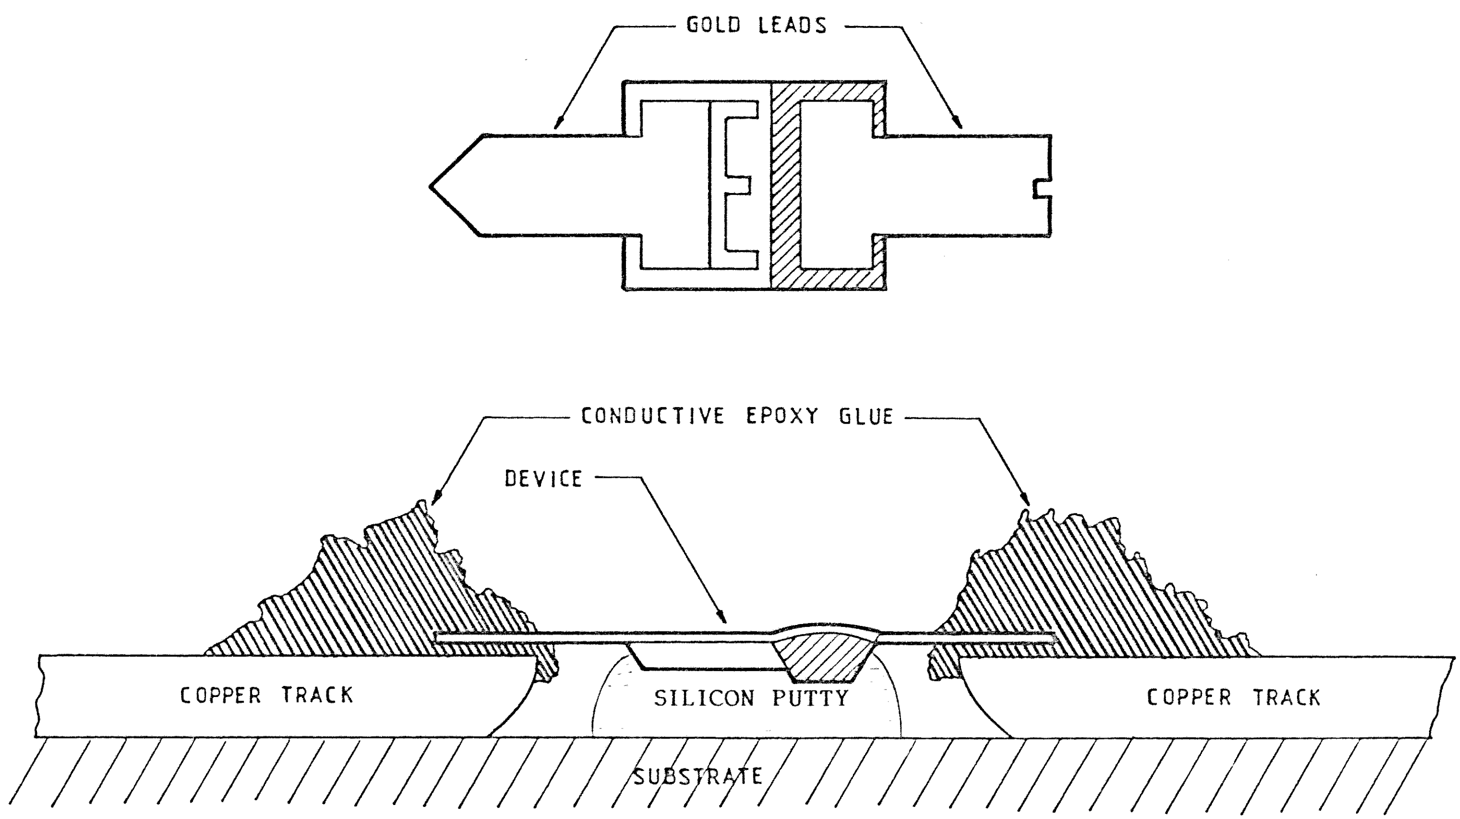

FIGURE. C.2

C.6. Performance

The resistance of the bond pair associated with a diode is satisfactorily small. The diodes (2837 type, refer to Chapter 2) have a natural bulk and contact ohmic resistance of a few tens of ohms. The measured ohmic series resistance of diodes bonded in the manner described here is not detectably larger than that measured using nickel probes on the diodes in their shipping container. It is concluded that the contact mechanism introduces less than one ohm of additional ohmic parasitic impedance.

The reactive parasitics associated with the bond are more difficult to measure, but will be limited by the physical size of the bond area and thickness. The diode package and leads offer sub-nanohenry reactance, while being much larger in physical size than the bond. Thus the bond will not be capable of worsening the package parasitics by a significant margin.

Diodes bonded by the method described have been tested two years after original fixing. The bond characteristics were found not to have degraded. Finally a (destructive) test of bond strength was carried out. The leads sheared off the diode body before the bond surface in one case, and in all cases the diodes were observed to be secure. Particles of glue resisted removal even from the smooth substrate.

1. SPECIFICATIONS OF ACTIVE COMPONENTS

1.1. Introduction

This appendix contains the following data and specifications:

1) Data Sheets for HP BJTs used in the 7912ADM;

2) Data Sheets for the WJ modules used in the 7912ADM.

The specification sheets are reproduced from manufacturer's data catalogues current at the time of design, with permission.

HEWLETT *hp* PACKARD COMPONENTS	**LINEAR POWER TRANSISTOR** — **2N6701 (HXTR-5101)**

BIPOLAR TRANSISTORS 2

Features

HIGH P_{1dB} LINEAR POWER
23 dBm Typical at 2 GHz
22 dBm Typical at 4 GHz

HIGH P_{1dB} GAIN
13 dB Typical at 2 GHz
7.5 dB Typical at 4 GHz

LOW DISTORTION

HIGH POWER-ADDED EFFICIENCY

MATCHING CONDITIONS INDEPENDENT OF OUTPUT POWER

INFINITE SWR TOLERANCE ABOVE 2 GHz

RUGGED HERMETIC PACKAGE

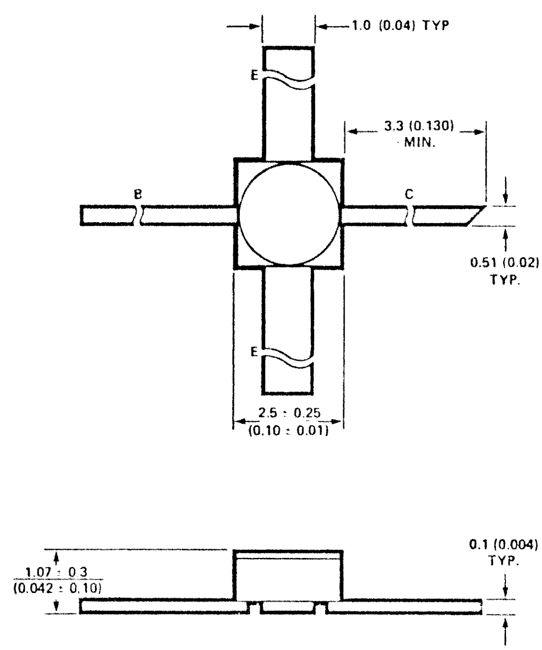

DIMENSIONS IN MILLIMETERS (INCHES).

HPAC-100 Package Outline

Description/Applications

The 2N6701 (HXTR-5101) is an NPN bipolar transistor designed for high output power and gain up to 5 GHz. To achieve excellent uniformity and reliability, the manufacturing process utilizes ion implantation, self-alignment techniques and Ti/Pt/Au metallization. The chip has a dielectric scratch protection over its active area and Ta_2N ballast resistors for ruggedness.

The superior gain, power, and distortion performance of the 2N6701 commend it for applications in radar, ECM, space, and commercial and military telecommunications. The 2N6701 features both guaranteed power output and associated gain at 1 dB gain compression.

The 2N6701 is supplied in the HPAC-100, a metal/ceramic hermetic package, and is capable of meeting the environmental requirements of MIL-S-19500 and the test requirements of MIL-STD-750/883.

Electrical Specifications at T_{CASE} =25°C

Symbol	Parameters and Test Conditions		Test MIL-STD-750	Units	Min.	Typ.	Max.
BV_{CBO}	Collector-Base Breakdown Voltage at I_C = 3mA		3001.1*	V	40		
BV_{CEO}	Collector-Emitter Breakdown Voltage at I_C = 15mA		3011.1*	V	24		
BV_{EBO}	Emitter-Base Breakdown Voltage at I_B = 30μA		3026.1*	V	3.3		
I_{EBO}	Emitter-Base Leakage Current at V_{EB}=2V		3061.1	μA			2
I_{CES}	Collector-Emitter Leakage Current at V_{CE}=32V		3041.1	nA			200
I_{CBO}	Collector-Base Leakage Current at V_{CB}=20V		3036.1	nA			100
h_{FE}	Forward Current Transfer Ratio at V_{CE}=18V, I_C = 30mA		3076.1*		15	40	75
P_{1dB}	Power Output at 1dB Gain Compression	f=2GHz		dBm		23	
		4GHz			21	22	
G_{1dB}	Associated 1dB Compressed Gain	2GHz		dB		13	
		4GHz			6.5	7.5	
P_{SAT}	Saturated Power Output (8dB Gain)	2GHz		dBm		25.5	
	(3dB Gain)	4GHz				25	
η	Power-Added Efficiency at 1dB Compression	2GHz		%		35	
		4GHz				24	
IMD	Third Order Intermodulation Distortion (Reference to either tone), at P_O(PEP)= 22dBm	4GHz		dB		-30	
	Tuned for Maximum Output Power at 1dB Compression V_{CE}=18V, I_C=30mA						

*300μs wide pulse measurement at ≤2% duty cycle.

Recommended Maximum Continuous Operating Conditions [1]

Symbol	Parameter	Value
V_{CBO}	Collector to Base Voltage [2]	40V
V_{CEO}	Collector to Emitter Voltage [2]	24V
V_{EBO}	Emitter to Base Voltage [2]	3.3V
I_C	DC Collector Current [2]	50 mA
P_T	Total Device Dissipation [3]	700 mW
T_J	Junction Temperature	200° C
T_{STG}	Storage Temperature	-65° C to +200° C

Notes:
1. Operation of this device in excess of any one of these conditions is likely to result in a reduction in device mean time between failure (MTBF) to below the design goal of 1 x 10^7 hours at T_J = 175°C (assumed Activation Energy = 1.5 eV). Corresponds to maximum rating for 2N6701.
2. T_{CASE} = 25°C.
3. See Figure 7 for derating conditions.

Absolute Maximum Ratings *

Symbol	Parameter	Limit
V_{CBO}	Collector to Base Voltage	45V
V_{CEO}	Collector to Emitter Voltage	27V
V_{EBO}	Emitter to Base Voltage	4V
I_C	DC Collector Current	100 mA
P_T	Total Device Dissipation	1.1 W
T_J	Junction Temperature	300° C
$T_{STG(MAX)}$	Maximum Storage Temperature	250° C
—	Lead Temperature (Soldering 10 seconds each lead)	+250° C

*Operation in excess of any one of these conditions may result in permanent damage to this device.

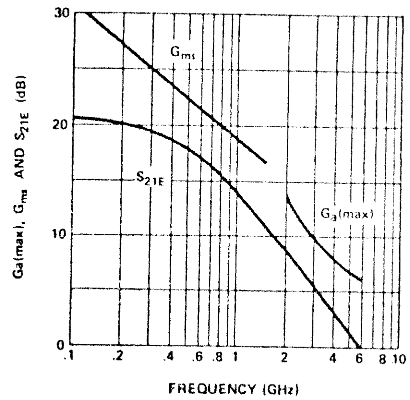

Figure 1. Typical $G_a(max)$, Maximum Stable Gain (G_{ms}), and S_{21E} vs. Frequency at V_{CE} = 18V, I_C = 30mA.

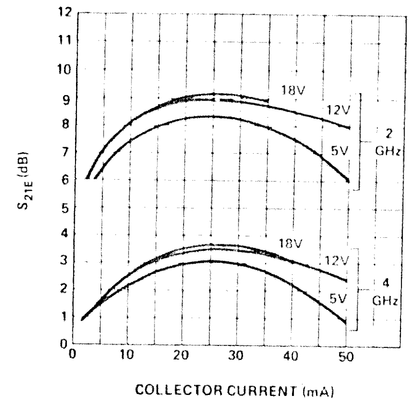

Figure 2. Typical S_{21E} vs. Current at 2 and 4GHz.

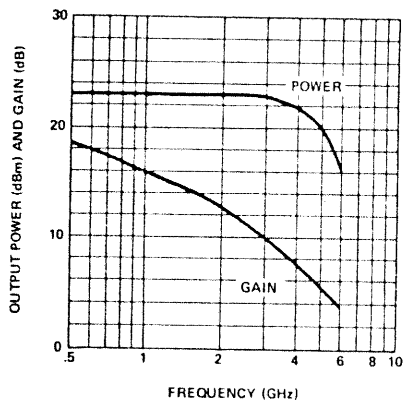

Figure 3. Typical P_{1dB} Linear Power and Associated 1dB Compressed Gain vs. Frequency at V_{CE} = 18V, I_C = 30mA.

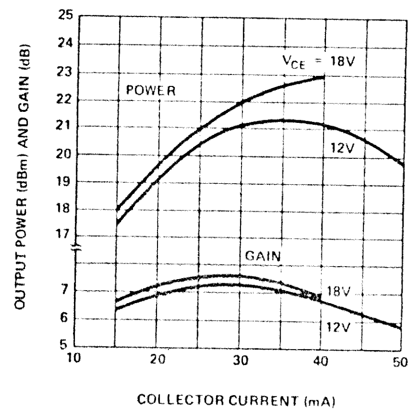

Figure 4. Typical P_{1dB} Linear Power and Associated 1dB Compressed Gain vs. Current at V_{CE} = 12 and 18V at 4GHz.

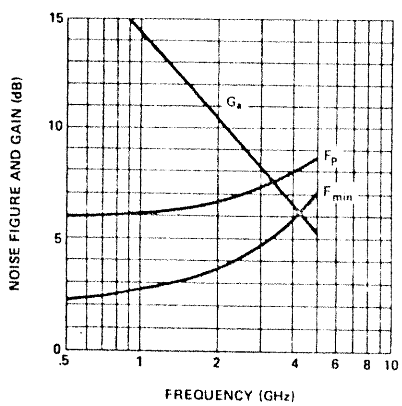

Figure 5. Typical Noise Figure (F_{min}) and Associated Gain (G_a) when tuned for Minimum Noise vs. Frequency at V_{CE} = 18V, I_C = 10mA. Typical Noise Figure (F_p) when tuned for Max P_{1dB} at V_{CE} = 18V, I_C = 30mA.

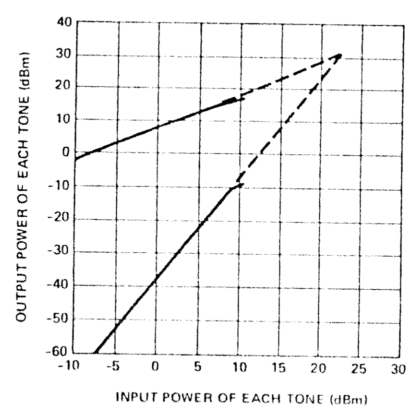

Figure 6. Typical Two Tone 3rd Order Intermodulation Distortion at 4GHz for a frequency separation of 5MHz at V_{CE} = 18V, I_C = 30mA.

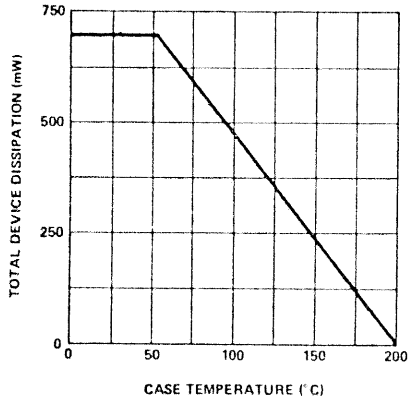

Figure 7. Maximum Power Dissipation Curve for θ_{jc} = 210° C/W, T_{jMAX} = 200° C.

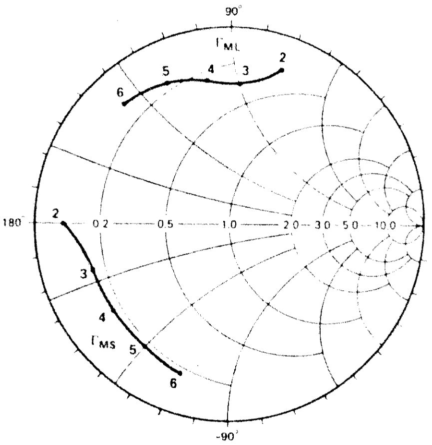

Figure 8. Typical Γ_{MS}, Γ_{ML}, (calculated from the average S-parameters) in the 2 to 6GHz frequency range, at V_{CE} = 18V, I_C = 30mA.

Typical S-Parameters V_{CE} = 18V, I_C = 30mA

	S_{11}		S_{21}			S_{12}			S_{22}	
Freq. (MHz)	Mag.	Ang.	(dB)	Mag.	Ang.	(dB)	Mag.	Ang.	Mag.	Ang.
100	0.80	-19	20.6	10.7	165	-37	0.01	77	0.98	-8
200	0.78	-37	20.1	10.2	154	-31	0.03	67	0.94	-15
300	0.75	-53	19.5	9.44	143	-28	0.04	60	0.88	-21
400	0.72	-68	18.7	8.63	133	-27	0.05	53	0.83	-26
500	0.68	-81	17.9	7.87	124	-26	0.05	47	0.78	-30
600	0.66	-92	17.0	7.15	117	-25	0.06	42	0.73	-33
700	0.64	-102	16.2	6.52	110	-24	0.06	39	0.69	-36
800	0.62	-111	15.5	5.96	104	-24	0.07	36	0.66	-38
900	0.61	-119	14.8	5.49	99	-23	0.07	33	0.64	-41
1000	0.60	-126	14.1	5.08	94	-23	0.07	31	0.61	-43
1500	0.56	-151	11.2	3.64	75	-23	0.08	25	0.55	-51
2000	0.55	-169	8.9	2.80	59	-22	0.08	22	0.52	-61
2500	0.56	179	7.2	2.29	45	-21	0.09	21	0.53	-72
3000	0.55	168	5.7	1.93	33	-21	0.09	21	0.52	-79
3500	0.56	158	4.5	1.69	21	-20	0.10	20	0.55	-89
4000	0.54	148	3.5	1.50	10	-19	0.11	19	0.58	-96
4500	0.54	137	2.5	1.33	0	-19	0.11	18	0.58	-106
5000	0.52	128	1.6	1.21	-11	-18	0.13	16	0.62	-113
5500	0.54	115	1.0	1.12	-23	-17	0.14	14	0.60	-122
6000	0.54	108	0.0	1.01	-32	-17	0.15	11	0.64	-132

Typical S-Parameters V_{CE} = 15V, I_C = 15mA

	S_{11}		S_{21}			S_{12}			S_{22}	
Freq. (MHz)	Mag.	Ang.	(dB)	Mag.	Ang.	(dB)	Mag.	Ang.	Mag.	Ang.
100	0.80	-18	19.4	9.35	166	-37	0.01	78	0.98	-7
200	0.78	-35	19.1	9.07	155	-31	0.02	69	0.95	-14
300	0.76	-50	18.5	8.44	145	-28	0.03	61	0.91	-20
400	0.73	-64	17.8	7.79	135	-26	0.04	55	0.86	-25
500	0.69	-77	17.1	7.16	127	-25	0.05	49	0.81	-29
600	0.67	-88	16.3	6.56	119	-24	0.06	44	0.76	-32
700	0.64	-97	15.5	6.02	113	-23	0.06	40	0.72	-35
800	0.62	-107	14.8	5.54	107	-23	0.06	37	0.69	-38
900	0.60	-115	14.2	5.13	101	-23	0.07	34	0.66	-40
1000	0.60	-122	13.5	4.76	96	-23	0.07	32	0.63	-43
1500	0.57	-148	10.8	3.47	76	-22	0.08	24	0.57	-53
2000	0.55	-166	8.6	2.69	60	-21	0.08	21	0.54	-63
2500	0.56	-178	6.9	2.21	46	-21	0.09	19	0.55	-75
3000	0.56	171	5.1	1.80	36	-20	0.09	21	0.50	-85
3500	0.56	160	4.3	1.65	21	-20	0.10	18	0.56	-91
4000	0.53	151	3.3	1.47	10	-19	0.11	18	0.59	-99
4500	0.53	141	2.3	1.30	0	-19	0.11	17	0.59	-108
5000	0.50	130	1.5	1.18	-10	-18	0.12	15	0.62	-116
5500	0.52	118	0.8	1.10	-22	-17	0.14	13	0.61	-124
6000	0.53	110	0.0	0.99	-31	-16	0.15	11	0.64	-135

WJ-DA31

10 TO 2000 MHz RF DIP AMPLIFIER

- LOW NOISE: 3.6 dB (TYP.)
- HIGH GAIN: 23 dB (TYP.)
- LOW VSWR: < 1.6:1 (TYP.)
- ULTRA BROADBAND

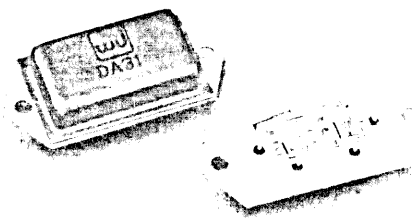

Guaranteed Specifications*

Characteristic	Typical	0°-50°C	-54°C-+85°C
Frequency (Min.)	5-2050 MHz	10-2000 MHz	10-2000 MHz
Small Signal Gain (Min.)	23.0 dB	21.0 dB	20.0 dB
Gain Flatness (Max.)	±0.6 dB	±1.0 dB	±1.2 dB
Noise Figure (Max.)	3.5 dB	4.2 dB	4.7 dB
Power Output at 1 dB Compression (Min.)	+1.0 dBm	-1.0 dBm	-1.5 dBm
VSWR (Max.) Input/Output	<1.6:1	2.0:1	2.0:1
DC Current (Max.) at 15 Volts	23 mA	27 mA	31 mA

*Measured in a 50-ohm system at 15 Vdc +1% Nominal

Outline Drawing

DA31

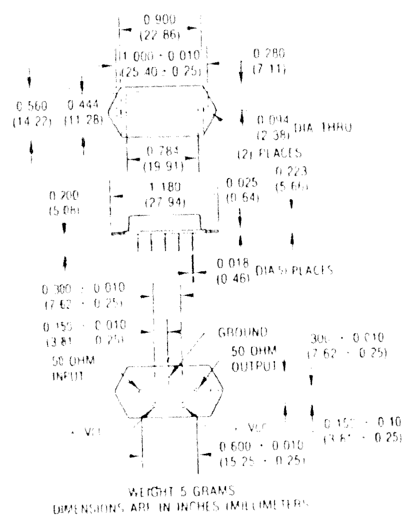

Typical Intermodulation Performance at 25°C

Second Order Harmonic Intercept Point .... +28 dBm (Typ.)
Second Order Two Tone Intercept Point .... +22 dBm (Typ.)
Third Order Two Tone Intercept Point .... +12 dBm (Typ.)

Absolute Maximum Ratings

Ambient Operating Temperature .... -54°C to +100°C
Storage Temperature .... -62°C to +125°C
Maximum Case Temperature .... 125°C
Maximum DC Voltage .... +17 Volts
Maximum Continuous RF Input Power .... +6 dBm
Maximum Short Term RF Input Power (1 Minute Max.) .... +50 Milliwatts
Maximum Peak Power .... 0.5 Watt (3 μsec Max.)

"S" Series Burn-in Temperature .... 85°C

Weight 2.27 grams (0.08 oz.) maximum

Typical Performance at 25°C

Gain

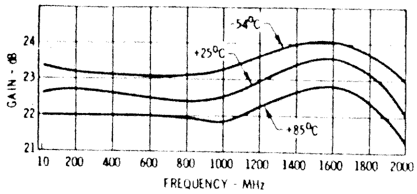

Noise Figure

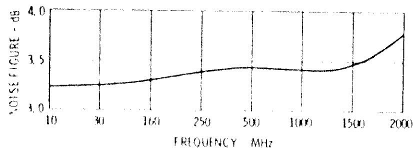

Power Output*

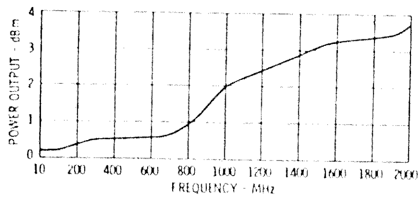

*1 dB Gain Compression

VSWR

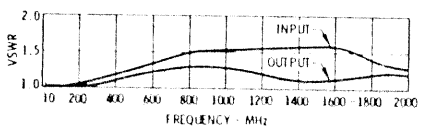

Typical Automatic Test Data

V_{CC} = 15 V

FREQ MHZ	VSWR IN	VSWR OUT	GAIN DB
100.	1.0	1.1	22.6
200.	1.0	1.1	22.7
300.	1.1	1.1	22.7
400.	1.2	1.1	22.6
500.	1.3	1.2	22.4
600.	1.2	1.2	22.5
700.	1.4	1.2	22.3
800.	1.5	1.2	22.4
900.	1.5	1.4	22.4
1000.	1.5	1.2	22.5
1100.	1.5	1.2	22.6
1200.	1.5	1.2	22.9
1300.	1.5	1.2	23.1
1400.	1.5	1.1	23.4
1500.	1.5	1.1	23.4
1600.	1.6	1.1	23.6
1700.	1.6	1.2	23.5
1800.	1.4	1.2	23.2
1900.	1.3	1.2	22.7
2000.	1.3	1.2	22.0

Linear S-Parameters

FREQ MHZ	S11 MAG	S11 ANG	S21 MAG	S21 ANG	S12 MAG	S12 ANG	S22 MAG	S22 ANG
100.	.08	110.7	13.78	-20.8	.02	-5.[illegible]	.04	140.7
200.	.02	-51.2	13.66	-41.8	.02	-14.2	.04	92.9
300.	.04	-102.8	13.60	-62.7	.02	-20.9	.05	47.9
400.	.08	-124.8	13.41	-84.3	.02	-28.9	.06	1.4
500.	.12	-162.8	13.22	-106.6	.02	-35.4	.08	-29.9
600.	.14	166.1	13.28	-127.7	.02	-44.4	.10	-80.4
700.	.16	122.7	13.09	-148.4	.02	-52.1	.12	-117.0
800.	.19	88.8	13.16	-169.2	.02	-59.2	.14	-151.9
900.	.20	52.5	13.22	170.9	.02	-67.7	.15	171.8
1000.	.20	13.9	13.28	151.7	.02	-77.2	.15	127.5
1100.	.20	-28.2	13.50	130.8	.02	-82.0	.12	101.1
1200.	.20	-71.7	13.91	108.4	.02	-91.2	.11	63.9
1300.	.19	-122.5	14.35	84.6	.02	-98.9	.07	19.1
1400.	.20	-175.6	14.80	60.6	.02	-108.2	.04	-46.4
1500.	.21	131.5	14.86	35.1	.02	-116.7	.04	145.5
1600.	.22	80.0	15.17	9.2	.02	-129.1	.06	155.2
1700.	.22	31.2	14.96	-19.2	.02	-140.6	.08	118.2
1800.	.18	-11.1	14.47	-48.1	.02	-151.4	.07	88.1
1900.	.12	-41.4	13.67	-78.1	.02	-162.8	.07	73.5
2000.	.12	-57.8	12.52	-109.5	.02	-177.1	.09	54.4
2100.	.22	-47.4	11.22	-145.1	.02	165.8	.11	5.8
2200.	.30	-95.5	9.[illegible]	177.2	.02	146.4	.10	-63.5

Deviation from Linear Phase, Gain and Group Delay

FREQ MHZ	DEV LIN PH DEG	REL PH DEG	GAIN DEV DB	REL GAIN DB	GROUP DELAY N-SEC
100.	-9.18	.00	-.01	22.78	.60
200.	-7.01	-21.61	-.08	22.71	.59
300.	-5.75	-42.52	-.12	22.67	.59
400.	-4.38	-64.13	-.24	22.55	.61
500.	-3.08	-86.41	-.36	22.43	.60
600.	-1.79	-107.50	-.32	22.47	.58
700.	.47	-128.22	-.45	22.34	.58
800.	2.58	-149.09	-.40	22.39	.57
900.	5.74	-168.91	-.36	22.43	.54
1000.	9.51	-188.12	-.33	22.46	.56
1100.	11.59	-209.02	-.18	22.60	.60
1200.	12.09	-231.49	.08	22.87	.64
1300.	11.32	-255.24	.35	23.14	.66
1400.	10.27	-279.27	.62	23.40	.69
1500.	7.79	-304.72	.65	23.44	.73
1600.	3.86	-331.63	.83	23.62	.76
1700.	-.62	-359.09	.71	23.50	.79
1800.	-6.53	-387.38	.42	23.21	.82
1900.	-13.49	-417.91	-.07	22.72	.85
2000.	-21.97	-449.27	-.83	21.96	.87

EQUIV LENGTH IN AIR CM = 19.1482
TIME DELAY IN N-SEC = .638275

TECHNICAL DATA — WATKINS-JOHNSON

10 TO 2000 MHz CASCADABLE AMPLIFIER WJ-A35

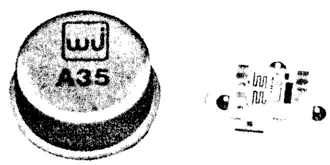

- MEDIUM OUTPUT LEVEL: +9 dBm (TYP)
- WIDE POWER SUPPLY RANGE: +8 TO +20 VOLTS
- SMALL SIZE: TO-8

GUARANTEED SPECIFICATIONS*

Characteristic	Typical	0°-50°C	−54°C-+85°C
Frequency (Min.)	1-2050 MHz	10-2000 MHz	10-2000 MHz
Small Signal Gain (Min.)	10.2 dB	9.0 dB	8.5 dB
Gain Flatness (Max.) 500-2000 MHz	< ±0.3 dB	±0.6 dB	±0.8 dB
Gain Flatness (Max.) 10-2000 MHz	< ±0.5 dB	±0.8 dB	±1.0 dB
Noise Figure (Max.)	5.0 dB	6.5 dB	7.0 dB
Power Output at 1 dB Compression (Min.)	+9 dBm	+7 dBm	+ 6.5 dBm
VSWR (Max.) Input	<1.5:1	2.0:1	2.2:1
VSWR (Max.) Output		2.2:1	2.2:1
DC Current at 15 Volts (Max.)	24 mA	27 mA	29 mA

Second Order Harmonic Intercept Point: +38 dBm (Typ.)
Second Order Two Tone Intercept Point: +33 dBm (Typ.)
Third Order Two Tone Intercept Point: +21 dBm (Typ.)

*Measured in a 50 ohm system, at +15 Vdc Nominal

ABSOLUTE MAXIMUM RATINGS

Ambient Operating Temperature . . . . . . −54°C to +100°C
Storage Temperature . . . . . . −62°C to +125°C
Maximum Case Temperature . . . . . 125°C
Maximum DC Voltage . . . . +21 Volts
Maximum Continuous R.F. Input Power . . . . +13 dBm
Maximum Short Term R.F. Input Power (1 minute Maximum) . . . . +50 Milliwatts
Maximum Peak Power . 0.5 Watt (3 μsec maximum)
"S" Series Burn-In Temperature . . 125°C

*Supersedes WJ-A35 Technical Data Sheet Dated Nov. 1977

NOTE: All testing per applicable internal Watkins-Johnson test procedures which are available upon request and are also subject to change without notice.

WATKINS-JOHNSON COMPANY
3333 Hillview Avenue, Palo Alto, California 94304
(415) 493-4141, TWX: 910-373-1253, TELEX: 348-415, CABLE: WJPLA

192 A

JULY 1980*

Specifications subject to change without notice.

WJ-A35

TYPICAL PERFORMANCE AT 25°C

GAIN

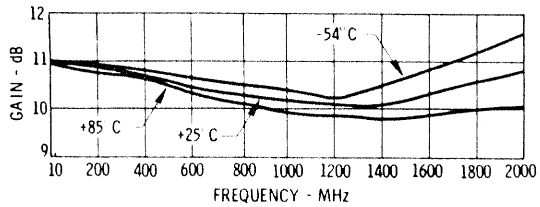

NOISE FIGURE

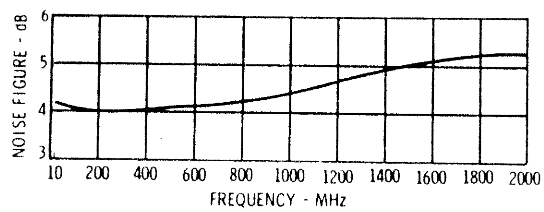

POWER OUTPUT*

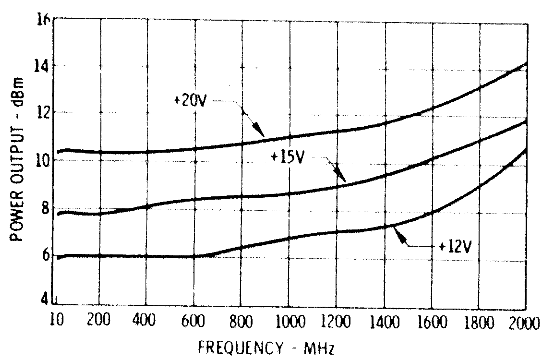

*at 1 dB Gain Compression

VSWR

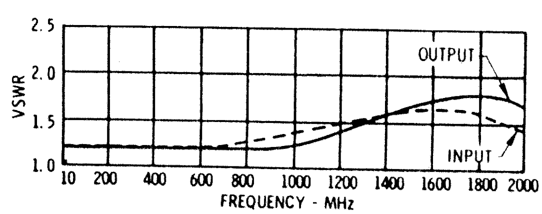

TYPICAL AUTOMATIC TEST DATA V_{cc} = 15V

FREQ MHZ	VSWR IN	VSWR OUT	GAIN DB
100.	1.2	1.5	10.6
200.	1.1	1.4	10.5
300.	1.1	1.3	10.5
400.	1.1	1.3	10.5
500.	1.1	1.2	10.5
600.	1.1	1.2	10.3
700.	1.1	1.2	10.3
800.	1.1	1.2	10.3
900.	1.2	1.2	10.3
1000.	1.2	1.2	10.2
1100.	1.3	1.2	10.2
1200.	1.3	1.2	10.2
1300.	1.4	1.2	10.3
1400.	1.4	1.3	10.3
1500.	1.4	1.3	10.5
1600.	1.4	1.3	10.7
1700.	1.4	1.4	10.9
1800.	1.2	1.4	11.1
1900.	1.2	1.3	11.2
2000.	1.5	1.2	11.1
2100.	2.4	1.2	10.2

LINEAR S-PARAMETERS

FREQ	S11		S21		S12		S22	
MHZ	MAG	ANG	MAG	ANG	MAG	ANG	MAG	ANG
100.	.09	160.6	3.38	171.0	.15	-1.8	.19	168.2
200.	.05	-166.8	3.35	156.7	.14	-7.0	.16	164.0
300.	.06	179.3	3.37	144.3	.14	-12.3	.15	149.8
400.	.04	166.2	3.36	131.7	.14	-18.3	.12	134.8
500.	.03	-167.2	3.35	119.8	.14	-24.3	.10	114.6
600.	.04	-164.7	3.29	106.9	.14	-31.1	.08	100.7
700.	.03	-124.2	3.26	93.5	.14	-36.4	.08	68.2
800.	.07	-126.9	3.29	79.9	.14	-43.4	.07	40.9
900.	.07	-136.1	3.26	66.4	.14	-50.2	.07	17.7
1000.	.10	-127.4	3.25	52.4	.14	-55.7	.09	-3.9
1100.	.13	-139.3	3.25	39.2	.13	-62.9	.10	-28.1
1200.	.14	-151.6	3.25	25.7	.13	-69.0	.09	-44.5
1300.	.18	-158.8	3.26	13.4	.13	-76.7	.10	-63.5
1400.	.17	-169.6	3.28	-.8	.13	-84.6	.12	-81.8
1500.	.18	176.6	3.33	-16.2	.13	-91.5	.13	-107.2
1600.	.16	160.5	3.43	-32.1	.13	-99.3	.14	-120.6
1700.	.15	143.9	3.51	-48.3	.14	-108.3	.17	-140.3
1800.	.10	101.1	3.59	-66.4	.14	-118.2	.16	-163.5
1900.	.10	31.0	3.65	-85.8	.14	-129.5	.14	174.2
2000.	.21	-22.3	3.61	-107.2	.14	-144.1	.09	145.5
2100.	.42	-67.3	3.23	-141.9	.13	-168.6	.08	20.9
2200.	.57	-93.8	2.71	-167.0	.12	173.6	.17	-33.5

OUTLINE DRAWING

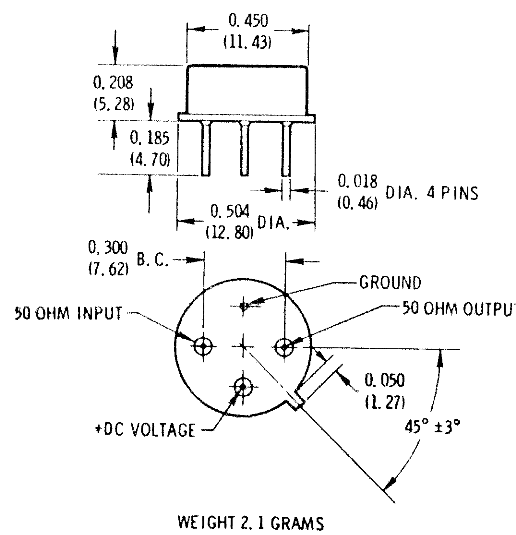

WEIGHT 2.1 GRAMS
DIMENSIONS ARE IN INCHES (MILLIMETERS)

JULY 1980

TECHNICAL DATA

WATKINS-JOHNSON

10 TO 2000 MHz CASCADABLE AMPLIFIER WJ-A38

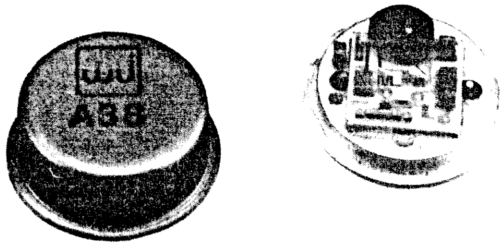

- HIGH OUTPUT POWER: +19 dBm (TYP)
- HIGH THIRD ORDER I.P.: +30 dBm
- WIDE BANDWIDTH: 10 - 2000 MHz
- SMALL SIZE: TO-8

GUARANTEED SPECIFICATIONS*

Characteristic	Typical	0°C - 50°C	-54°C - +85°C
Frequency (Min.)	5-2050 MHz	10-2000 MHz	10-2000 MHz
Small Signal Gain (Min.)	7.5 dB	6.5 dB	6.0 dB
Gain Flatness (Max.)	±0.3 dB	±0.7 dB	±1.0 dB
Noise Figure (Max.)	7.0 dB	8.7 dB	9.2 dB
Power Output at 1 dB Compression (Min.)	+19.0 dBm	+18.0 dBm	+17.5 dBm
VSWR (Max.) Input/Output	1.7:1	2.0:1	2.2:1
DC Current Max at 15V	65 mA	69 mA	72 mA

*Measured in a 50 ohm system at 15 Vdc Nominal.

TYPICAL INTERMODULATION PERFORMANCE AT 25°C

Second Order Harmonic Intercept Point:	+52 dBm (Typ.)
Second Order Two Tone Intercept Point:	+45 dBm (Typ.)
Third Order Two Tone Intercept Point:	+30 dBm (Typ.)

ABSOLUTE MAXIMUM RATINGS

Ambient Operating Temperature . . . . -54°C to +100°C
Storage Temperature . . . . -62°C to +125°C
Maximum Case Temperature . . . . 105°C
Maximum DC Voltage . . . . +17 Volts
Maximum Continuous RF Input Power . . . . +50 Milliwatts

Maximum Short Term CW Input . . . . +100 Milliwatts (1 Minute Maximum)
Maximum Peak Power . . . . 0.5 Watt (3 μsec maximum)
"S" Series Burn-In Temperature . . . . 100°C

WATKINS-JOHNSON COMPANY
3333 Hillview Avenue, Palo Alto, California 94304
(415) 493-4141, TWX: 910-373-1253, TELEX: 348-415, CABLE: WJPLA

192A

NOVEMBER 1979

Specifications subject to change without notice.

WJ-A38

TYPICAL PERFORMANCE AT 25°C

GAIN

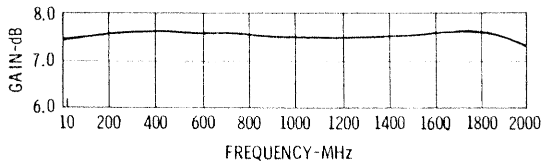

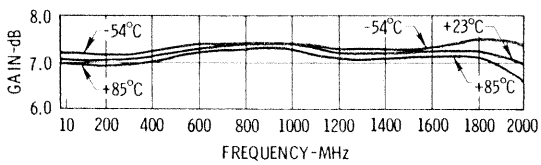

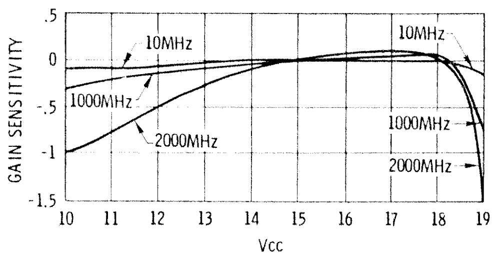

POWER OUTPUT

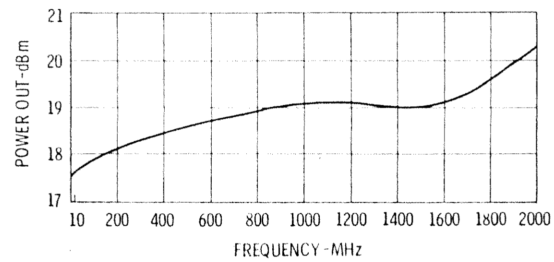

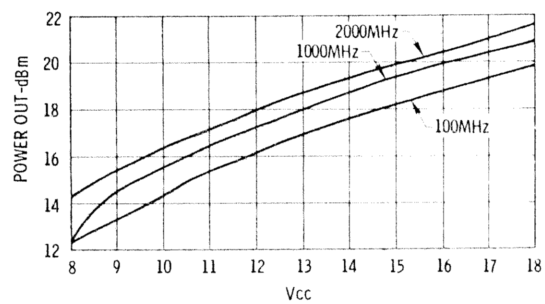

NOISE FIGURE

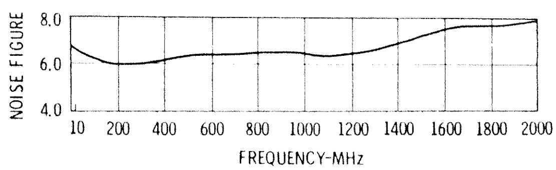

INTERCEPT POINT

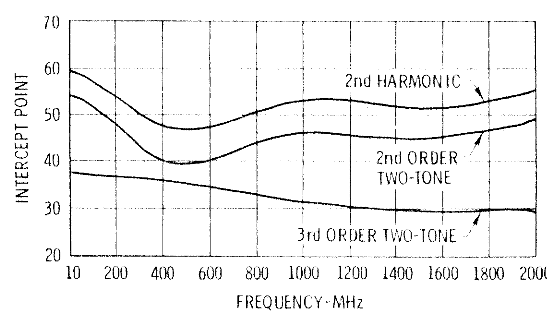

VSWR

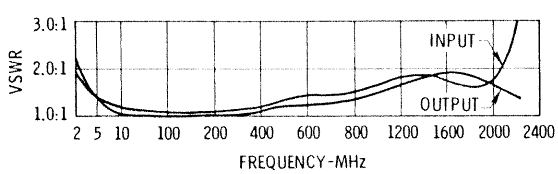

TYPICAL AUTOMATIC TEST DATA

V_{cc} = 12V

FREQ MHZ	VSWR IN	VSWR OUT	GAIN DB
100.	1.1	1.0	7.4
200.	1.1	1.0	7.4
300.	1.2	1.1	7.4
400.	1.2	1.1	7.4
500.	1.3	1.2	7.3
600.	1.4	1.2	7.3
700.	1.4	1.3	7.3
800.	1.5	1.4	7.3
900.	1.5	1.5	7.3
1000.	1.6	1.5	7.2
1100.	1.7	1.6	7.1
1200.	1.7	1.7	7.1
1300.	1.7	1.8	7.0
1400.	1.7	1.8	7.0
1500.	1.7	1.9	7.0
1600.	1.7	1.9	7.0
1700.	1.6	1.9	6.9
1800.	1.6	1.8	6.9
1900.	1.6	1.7	6.9
2000.	1.8	1.6	6.5
2100.	2.1	1.5	6.0

V_{cc} = 15V

FREQ MHZ	VSWR IN	VSWR OUT	GAIN DB
100.	1.1	1.0	7.5
200.	1.1	1.0	7.4
300.	1.2	1.1	7.4
400.	1.2	1.1	7.4
500.	1.3	1.2	7.4
600.	1.3	1.2	7.4
700.	1.4	1.3	7.4
800.	1.5	1.4	7.4
900.	1.5	1.4	7.4
1000.	1.6	1.5	7.3
1100.	1.6	1.6	7.3
1200.	1.7	1.7	7.3
1300.	1.7	1.8	7.2
1400.	1.7	1.8	7.2
1500.	1.7	1.9	7.2
1600.	1.7	1.9	7.2
1700.	1.6	1.9	7.2
1800.	1.6	1.8	7.1
1900.	1.5	1.8	7.2
2000.	1.7	1.7	7.0
2100.	1.9	1.5	6.6

LINEAR S-PARAMETERS

V_{cc} = 15V

FREQ	S11		S21		S12		S22	
MHZ	MAG	ANG	MAG	ANG	MAG	ANG	MAG	ANG
100.	.03	-94.0	2.37	163.9	.20	-8.6	.00	-170.6
200.	.06	-105.0	2.35	146.9	.20	-18.0	.01	-104.8
300.	.08	-117.8	2.36	130.0	.20	-27.4	.03	-98.0
400.	.10	-129.2	2.35	113.2	.20	-37.1	.05	-103.3
500.	.12	-139.1	2.34	94.6	.21	-47.0	.07	-115.0
600.	.14	-149.1	2.34	77.1	.21	-56.8	.10	-128.4
700.	.16	-158.4	2.34	59.9	.21	-66.7	.12	-140.2
800.	.19	-168.9	2.35	43.1	.21	-76.1	.15	-154.2
900.	.21	-179.8	2.35	26.4	.21	-85.8	.18	-167.5
1000.	.23	168.5	2.33	9.8	.21	-95.5	.21	178.4
1100.	.24	155.9	2.32	-6.9	.22	-105.8	.23	163.3
1200.	.26	143.0	2.31	-23.3	.22	-116.1	.26	148.8
1300.	.27	129.2	2.29	-40.6	.22	-126.3	.27	132.9
1400.	.27	113.4	2.28	-57.8	.23	-137.2	.29	118.1
1500.	.27	96.4	2.28	-74.8	.24	-147.6	.30	101.8
1600.	.26	76.2	2.28	-92.6	.24	-159.0	.31	84.5
1700.	.24	50.4	2.28	-110.3	.26	-170.4	.31	66.5
1800.	.22	18.5	2.27	-129.0	.27	176.8	.30	47.0
1900.	.21	-23.6	2.29	-148.8	.28	162.9	.28	23.9
2000.	.25	-69.6	2.23	-169.9	.30	147.5	.25	-3.3
2100.	.31	-109.6	2.13	169.0	.30	130.4	.20	-40.4
2200.	.41	-143.3	2.00	146.6	.30	113.8	.16	-83.9

DEVIATION FROM LINEAR PHASE, GAIN AND GROUP DELAY

V_{cc} = 15V

FREQ MHZ	DEV LIN Ø DEG	REL Ø DEG	GAIN DEV DB	ABS GAIN DB	GROUP DELAY N-SEC
100.	-.87	.00	.25	7.48	.47
200.	-.64	-17.06	.21	7.43	.47
300.	-.22	-33.92	.22	7.45	.47
400.	.26	-50.73	.21	7.44	.49
500.	-1.06	-69.34	.16	7.38	.50
600.	-1.20	-86.77	.16	7.38	.48
700.	-1.11	-103.97	.17	7.40	.47
800.	-.63	-120.77	.20	7.43	.47
900.	-.07	-137.51	.18	7.40	.46
1000.	.62	-154.11	.12	7.34	.46
1100.	1.18	-170.84	.08	7.30	.46
1200.	2.05	-187.25	.05	7.27	.47
1300.	2.04	-204.55	-.03	7.19	.48
1400.	2.21	-221.67	-.05	7.18	.48
1500.	2.43	-238.74	-.05	7.17	.48
1600.	1.97	-256.49	-.07	7.16	.49
1700.	1.54	-274.21	-.07	7.15	.51
1800.	.15	-292.88	-.11	7.11	.54
1900.	-2.39	-312.71	-.02	7.20	.57
2000.	-6.24	-333.86	-.25	6.97	.59

HEWLETT PACKARD

GENERAL PURPOSE TRANSISTOR

2N6679 (HXTR-2101)

Features

HIGH GAIN
10.5 dB Typical at 4 GHz

WIDE DYNAMIC RANGE

RUGGED HERMETIC PACKAGE

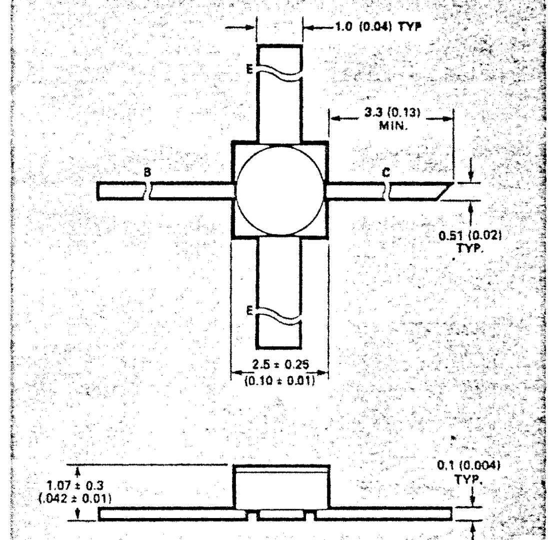

DIMENSIONS IN MILLIMETERS (INCHES).

HPAC-100 Package Outline

Description

The 2N6679 (HXTR-2101) is an NPN bipolar transistor designed for high gain and output power at 4 GHz. The device utilizes ion implantation techniques and Ti/Pt/Au metallization in its manufacture. The chip is provided with a dielectric scratch protection over its active area.

The 2N6679 is supplied in the HPAC-100, a rugged metal/ceramic hermetic package, and is capable of meeting the environmental requirements of MIL-S-19500 and the test requirements of MIL-STD-750/883.

Electrical Specifications at $T_{CASE}=25°C$

Symbol	Parameters and Test Conditions	MIL-STD-750 Test Method	Units	Min.	Typ.	Max.
BV_{CES}	Collector-Emitter Breakdown Voltage $I_C=100\mu A$	3011.1*	V	30		
I_{CEO}	Collector-Emitter Leakage Current at $V_{CE}=15V$	3041.1	nA			500
I_{CBO}	Collector Cutoff Current at $V_{CB}=15V$	3036.1	nA			100
h_{FE}	Forward Current Transfer Ratio $V_{CE}=15V, I_C=15mA$	3076.1*	—	50	120	220
G_T	Tuned Gain		dB	9.0	10.5	
P_{1dB}	Power Output at 1 dB Compression Bias Conditions for Above: $V_{CE}=15V, I_C=25mA$, Frequency = 4 GHz		dBm		18.5	

*300 μs wide pulse measurement ≤2% duty cycle.

Recommended Maximum Continuous Operating Conditions[1]

Symbol	Parameter	Value
V_{CBO}	Collector to Base Voltage[2]	25V
V_{CEO}	Collector to Emitter Voltage[2]	16V
V_{EBO}	Emitter to Base Voltage[2]	1.0V
I_C	DC Collector Current[2]	35mA
P_T	Total Device Dissipation[3]	450 mW
T_J	Junction Temperature	200° C
T_{STG}	Storage Temperature	-65° C to +200° C

Notes:

1. Operation of this device in excess of any one of these conditions is likely to reduce device median time to failure (MTTF) of 3.5 x 10^6 hours at T_J = 125°C (based on Activation Energy = 1.1 eV).
2. T_{CASE} = 25°C.
3. Derate at 4.8 mW/° C, $T_C \geq$ 106° C.

Absolute Maximum Ratings*

Symbol	Parameter	Limit
V_{CBO}	Collector to Base Voltage	30V
V_{CEO}	Collector to Emitter Voltage	20V
V_{EBO}	Emitter to Base Voltage	1.5V
I_C	DC Collector Current	70 mA
P_T	Total Device Dissipation	900 mW
T_J	Junction Temperature	300°C
$T_{STG(MAX)}$	Maximum Storage Temperature	250°C
—	Lead Temperature (Soldering 10 seconds each lead)	+250°C

*Operation in excess of any one of these conditions may result in permanent damage to this device.

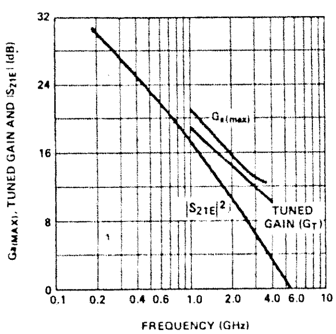

Figure 1. Typical $G_{a(MAX)}$ and Tuned Gain vs. Frequency at V_{CE}=15V, I_C=25 mA

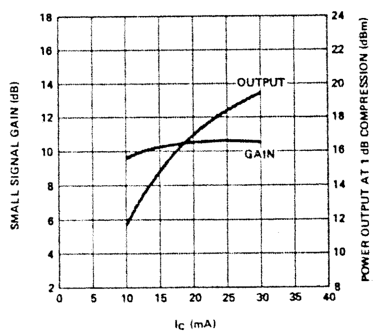

Figure 2. Typical Power Output at 1 dB Compression and Small Signal Gain vs. Collector Current at 4 GHz for V_{CE} = 15V.

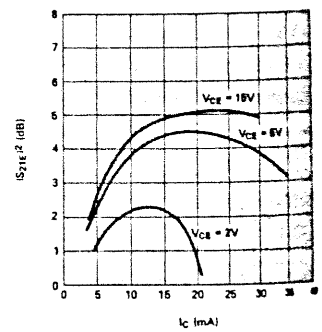

Figure 3. Typical $|S_{21E}|^2$ vs. Bias at 4 GHz.

Typical S-Parameters V_{CE} = 15V, I_C = 25mA

	S_{11}		S_{21}			S_{12}			S_{22}	
Freq. (MHz)	Mag.	Ang.	(dB)	Mag.	Ang.	(dB)	Mag.	Ang.	Mag.	Ang.
100	0.59	-66	30.8	34.6	146	-40.0	0.01	69	0.86	-18
500	0.58	-150	22.1	12.7	96	-33.2	0.02	44	0.51	-27
1000	0.59	-175	16.7	6.86	78	-30.5	0.03	51	0.44	-32
1500	0.59	173	13.3	4.61	64	-28.0	0.04	55	0.45	-39
2000	0.60	162	11.0	3.53	53	-25.7	0.05	55	0.44	-49
2500	0.61	156	8.9	2.79	43	-24.2	0.06	55	0.47	-60
3000	0.62	146	7.3	2.32	33	-22.6	0.07	56	0.48	-67
3500	0.63	139	5.9	1.96	22	-21.2	0.09	53	0.52	-79
4000	0.62	131	4.8	1.73	11	-19.7	0.10	50	0.55	-84
4500	0.61	123	3.5	1.50	1	-18.8	0.12	48	0.59	-93
5000	0.60	116	2.6	1.35	-9	-17.0	0.14	44	0.65	-102
5500	0.62	109	1.8	1.23	-19	-15.9	0.16	36	0.66	-113
6000	0.62	103	0.9	1.11	-28	-15.6	0.17	32	0.68	-123
6500	0.62	93	0.0	1.02	-37	-13.7	0.20	28	0.67	-131

E. ESTIMATION OF ERROR IN FAST SWEPT MEASUREMENT

E.1. Introduction

A frequency response graph of any system ideally catalogs that system's response to an infinite sinewave occurring at each frequency for which the graph provides a value. It is indicated in the main text that any measurement made in finite time must be an imperfect representation of this response, and furthermore, that the degree of imperfection is dependent upon the ideal system response, and cannot thus be calculated a priori. This limitation is rather less stringent than its pure description suggests. This appendix provides some discussion of the problem and suggests ways that the severity of the errors may be estimated.

The situation is set out in Figure E.1. The problem is to determine H(w), the frequency (w) domain response of the system. In this case x(t) is a signal which is a sinewave whose frequency moves from one limit of the range in which it is desired to know H(w), to the other limit. Such a signal is often called a "chirp" function in the context of radar, or a "glide tone" in the context of audio engineering. It is assumed in making a simple swept response measurement that H(w) is the quotient of ENV(y(t)), the envelope of y(t), and ENV(x(t)), the envelope of x(t). (The input envelope is commonly held constant, simplifying determination of the result.)

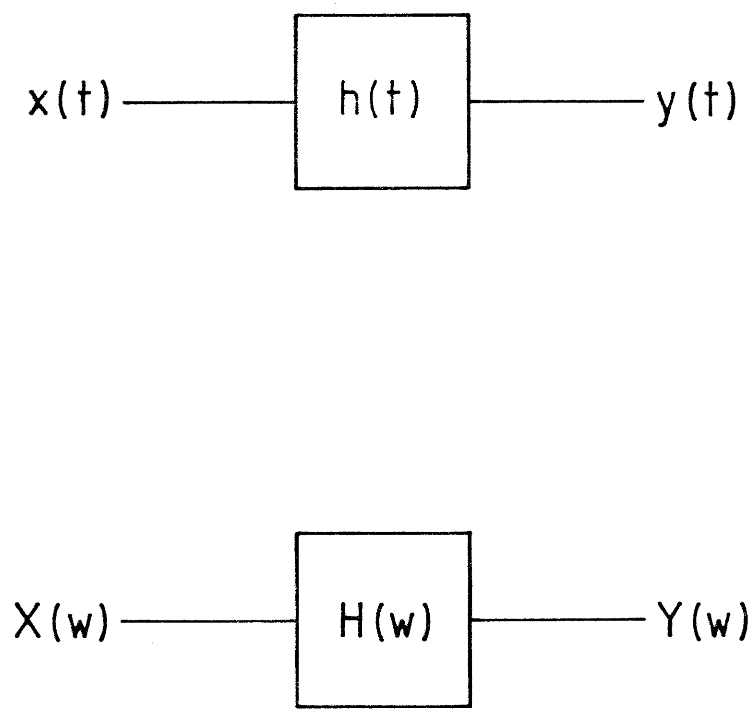

FIGURE. E.1

It is known that this assumption is not precisely true,[82] but it is intuitively obvious that it is accurate to any arbitrary degree, provided that the rate of sweep of the chirp is sufficiently low. It is highly desirable to have some estimate of where this limit falls, particularly with respect to the Single Trace method of determining H(w), discussed in the main text.

E.2. The Spectrum of a Sweep Signal

The Fourier spectrum of the swept input signal has been worked out, regarding the signal as a carrier with sawtooth frequency modulation.[83] (With this definition the carrier frequency, w_c, is the average of the start and finish frequencies of the signal, and the deviation or angular dispersion, D, the difference. This is a particularly odd instance of frequency modulation, as the deviation often approaches twice the carrier in frequency. This large deviation demands careful treatment.) The Fourier spectrum cannot be a continuum of frequencies, but must consist entirely of components which are multiples of the repetition frequency of the chirp. The complex Fourier coefficients C_N are:[83]

$$C_N = (1/T).(Pi.T/D)^{0.5}.e^{-i(T/2D)(NW)(NW)}$$

$$[C(w_1)+C(w_2)+iS(w_1)+iS(w_2)] \qquad (E1)$$

where

- T is the period of the chirp repetition
- W $= 2.Pi/T$
- w_1 is the start frequency of the sweep
- w_2 is the final frequency of the sweep
- D is the angular dispersion, w_2-w_1
- $C()$ is the Fresnel cosine integral
- $S()$ is the Fresnel sine integral

and i $=(-1)^{0.5}$.

Figure E.2 from [83] depicts the sideband coefficients C_N at one edge of the section of the spectrum bounded by w_1 and w_2. Call this section the sweep range interval. Broadly summarising, the spectrum of a wide deviation signal (a sweep with large sweep range interval) has discrete components which are large and of nearly constant amplitude in the sweep range interval, and which fall off away from that window. Furthermore, the ripple in the amplitudes of the components near the edge of the sweep range interval is small, and can be confined to an arbitrarily small interval near the edge frequencies, analagous to Gibbs' phenomenon.[83]

To relate equation (E1) to our understanding of a sweep signal, we observe that it indicates that the signal can be produced by summation of a large number of sinewave signals of almost identical amplitude. The resultant will have **constant amplitude**

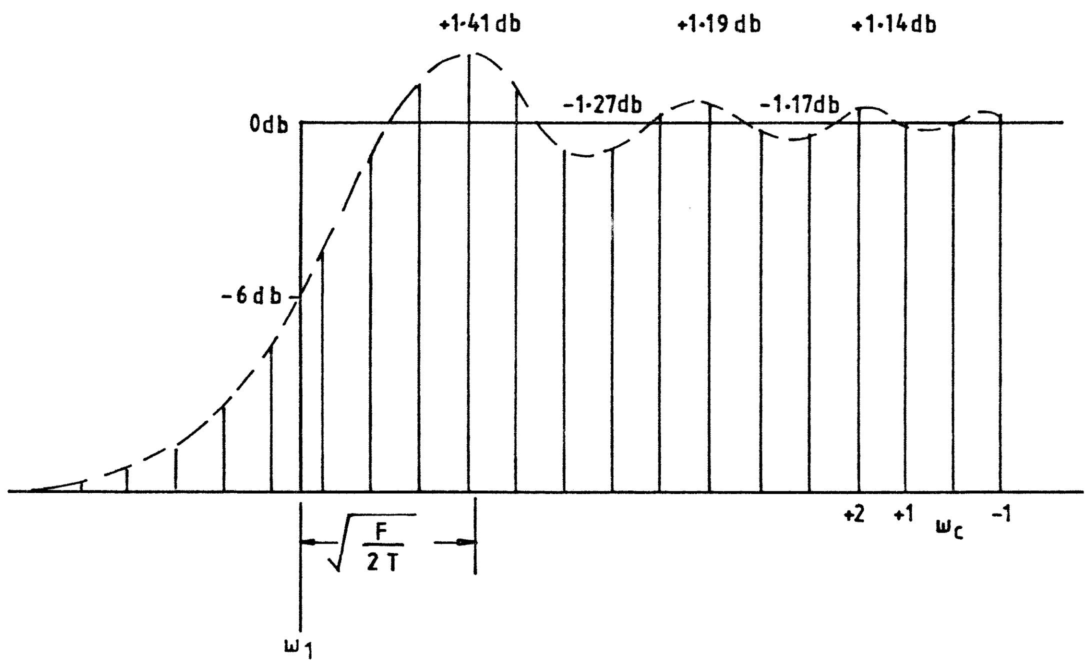

LEFT HALF OF THE SPECTRUM OF A PERIODIC CHIRP FUNCTION

FIGURE. E.2

equal to the amplitude of each of the generators. The reason for this is that the signals will cancel out, almost entirely, in any part of the spectrum except that corresponding to what we expect as the instantaneous frequency of the swept signal. About that frequency, there will be a "bandwidth" of uncancelled components. This "effective bandwidth" is of the order of:[83]

$$EBw = (Pi.D/T)^{0.5} \qquad (E2)$$

A further useful observation is that the effective bandwidth is of the same order as the ripple period in the amplitudes of the components at the edges of the sweep range interval, so that an analyser will not perceive the ripples, as is observed.

E.3. Discussion

Armed with this knowledge of the spectrum of the chirp signal, it is possible to decide what sweep speed will allow a good representation of the ideal response curve. In equation (E1) the sweep speed is inversely proportional to the interval T and thus the spacing of the components of X(w), the chirp spectrum, is inversely proportional to T.

The signal we observe as the system output, y(t), has spectrum:

$$Y(w) = X(w).H(w) \qquad (E3)$$

However, X(w) is discrete. To observe a feature of H(w), which is theoretically continuous, two conditions must be met. One: There must be sufficient components of the exciting signal present in the frequency band of the feature in order that the curve be sampled sufficiently often. This is equal to saying that the sweep must be sufficiently slow that the spacing of components of X(w) is small enough to sample H(w) at its Nyquist rate. Two: Since the envelope of the chirp signal used represents the amplitude of the uncancelled components in the "effective bandwidth", this too must be sufficiently narrow to permit observation of the features of H(w). These conditions are exemplified in Figure E.3. Thus to resolve features of bandwidth B in H(w) we must satisfy the conditions:

$$B \gg 2.Pi/T \qquad (E4)$$

and

$$B \gg (Pi.D/T)^{0.5} \qquad (E5)$$

It must of course be remembered that these constraints operate in addition to others, and concern only the problem of speed of sweeping. For a 2GHz bandwidth swept in 10ms, the constraint of (E5) dominates, and suggests that the resolution cannot exceed approximately 0.7 MHz.

RESOLUTION LIMITS ON H(w) FOR THE TWO CASES DISCUSSED

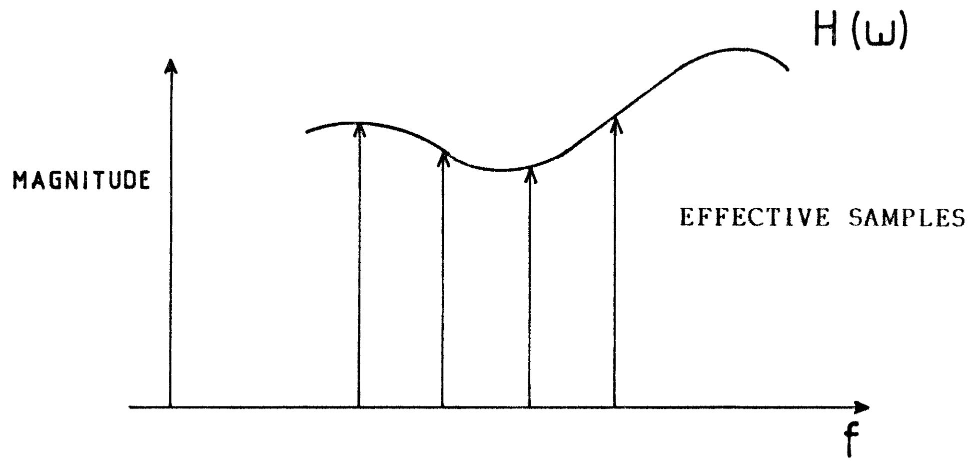

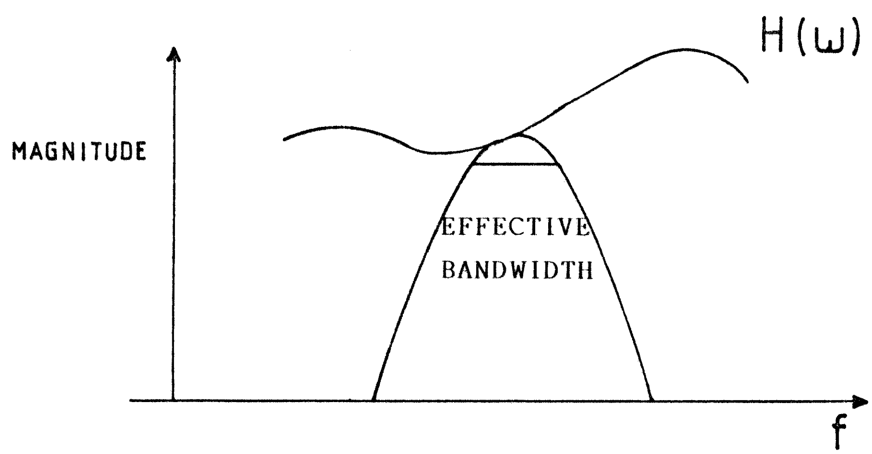

FIGURE. E.3

F. DESCRIPTION OF SOFTWARE FILES and UTILITIES

F.1. Introduction

Any given disk used in conjunction with the 7912 will carry a number of files. These fall into several categories:

1) The DIOS kernel, usually named "RUNTEK";

2) The DATEK data file, which carries data unique to a 7912;

3) Expected general data files, named "TEK..." or "tek...";

4) Waveform data files, starting "wd....", "we...." or "wr....";

5) Control setting data files, starting "cs....";

6) Disk resident DIOS command files, with 3 capital letter names;

7) System and utility files.

Each group is discussed below. The purpose of this appendix is to describe the function of all the files which will be encountered so that the user will not disturb correct operation by illegal renaming, deleting, etc. In addition, the uitiities will greatly assist any further work undertaken.

F.2. Required Files

1. RUNTEK is arbitrarily named. It is the DIOS kernal and resident commands in a single file. It is the programme which is run in order to get into DIOS.

2. The DATEK file should not be copied except for backup purposes, and it is best that this involve changing its name. This is because it is where DIOS records information such as the usage of the digitiser.

3. There are five "tek..." files. Each is associated with a certain DIOS command. Reference should be made to the appropriate command in order to determine the contents of each. They are: "tekcep", associated with the SEP command; "tekfrq", associated with the SPM command; "tekDCW", associated with the DCW command; "tekdoc", associated with the DOC command; and "tekmen", associated with the MEN command. The respective commands expect the files to be present and a system error will occur if any is absent at the time the command is exectued.

F.3. Operating Data Files

4. Waveform data files store acquired arrays of points and their accociated information. These files are created by the STW and STx commands. They may be copied, replicated and deleted without affecting DIOS operation, but can only be read with the appropriate two character beginning of their name in place. For this reason it is advisable to preserve the type.

5. Control setting files record the internal status of the 7912. They can only be accessed by DIOS with the "cs" prefix in the name. Beyond this there are no restrictions on their duplication and alteration. They are stored as ASCII files in Tektronix format.

F.4. Command Overlays

6. Disk resident DIOS commands may be replicated, but will cause a system error if their names are changed and the new name invoked. Only three character names can be invoked, so any name with more may be used without fear as a copy or backup. Their absence merely deletes the command from the repertoire available in DIOS. The format of a disk resident command under HPL is simply any executable set of statements, the first of which is a label with the same name as the command. That is, the file name must equal the label of the first line. All handling is looked after by the kernel, which searches mass storage for a file corresponding to any otherwise unknown mnemonic. One further file, "ender", must be present for the handler to operate. This file contains the null programme (length 0 bytes) which is loaded in order to overwrite an overlayed command file, recovering the memory it occupied.

F.5. Utilities

7. There are some 18 utility programmes provided. These are described briefly below.

F.5.1. "GTK"

This is a special function key assignment file. It is an acronym for Graphics Translator Keys. These keys permit a programme to be listed onto the system graphics display device. The list can be advanced or retreated in 20, 40, 80 and 160 line steps, or can go from a current line number. This permits display of a particular subroutine, by first 'gto-ing' its call label. They also allow the disk catalog to be displayed or printed with files in alphanumeric order within type groups, i.e. programme, data, key and memory types, by invoking "catdsp" and "catprt", and f7 invokes the "resave" programme, for safe resaving of programme files.

F.5.2. "discup"

This programme copies all files on a disk to one or more tapes, storing each one in a form which can be individually accessed. This is in contrast to the HPL "dump" command, which stores programmes and data files in a form which can only be decoded by the inverse operation "load". Further, "dump" unnecessarily stores all disk HPL bootstraps on the tape. This is very wasteful, and "discup" will mostly consume one fewer tapes.

F.5.3. "redisc"

This is the inverse of "discup", corresponding to the HPL command "load". It is used where all files on the tape(s) are required back on disk.

F.5.4. "catdsp"

This programme displays the disk catalog in a logical fashion on the display device. Files are first grouped into types - data, programme, key and memory - and then they are displayed in alphanumerical order. The programme assumes that B$ has been dimensioned. (This is in order to make it compatible with other GTK-invoked files, and in order to allow it to operate while all the RUNTEK variables are intact. This latter is very useful in debugging and stepping operations with DIOS resident.)

F.5.5. "cat prt"

This programme is identical to "catdsp" except that the output is sent to the line printer.

F.5.6. "TESTwe" and "TESTwd"

These programmes allow the production of a data file that DIOS can read (as an edge or atc record). The waveform is produced by a small subroutine at the end of the files which contains the equation which it is desired represent the waveshape to be produced. This subroutine can quickly be tailored for any experimental application. The result is rounded to the appropriate accuracy, and the ancilliary record-keeping fields are appropriately filled to indicate that the waveform is an example.

F.5.7. "S-plt" and "S-prt"

These files plot on the system plotter, or print on the line printer, the contents of the corrector array stored in the file "tekfrq".

F.5.8. "PLIST"

This programme (developed by Alan Young and modified by the author) plots a section of programme on the plotter. The advantage over plain listing is that the result is letter/graphic quality and inclides special symbols encountered in the HPL language.

F.5.9. "MENadd", "MENsub" and "MENabc"

This suite of three programmes permit the menu record file "tekmen" to be updated. "MENadd" adds an entry to the file. "MENsub" removes an entry from the file. "MENabc" sorts the file into abcderian (correct alphabetical) order.

F.5.10. "EDITOR"

This is a programme editing programme. Its main advantage is that amongst its editing facilities, it provides global search and replace, which is not otherwise available on the HP9825. The HP9836 series provides global search only. It is conjectured that the global replace is not provided because of the HPL property of store-time syntax and validity checking. "EDITOR" merely halts if a requested replace would cause a store-time error.

When a logical exit is executed, "EDITOR" renames the original file as "backup", and stores the new (edited) file as the original. If "backup" existed, it is killed.

"EDITOR" is simple, and thus has been made substantially self documenting, and requires no system entities apart from a single disk and the built in 9825 peripherals.

F.5.11. "resave"

This programme replaces the HPL command 'resave'. The HPL command requires the user to provide the name of the file on disk that is to be overwritten by the file in memory. This a typographical error or a wrong user decision can totally obliterate a disk file.

The "resave" programme, when invoked by GTK, chains itself to the resident HPL programme. From this position it checks that the programme in memory has its name given as a label on the first line. (This is advocated HP practice.) It checks the disk to see that there is a file of that name. If there is not, an error is given. If there is, it is renamed as "reback", any old "reback" file being killed, and the programme (resident before "resave" was invoked) is stored in its place. "resave" then erases itself, by reloading the recently updated programme.

This programme has been proven over a number of years, and has saved much work both of the author's origination, and of his collegues.

F.5.12. "DCWgen"

This programme is used to generate the data file "tekDCW" which contains the disk resident window accessed by the DCW command.

When run it calculates the Dolph-Chebyshev (Harris' spelling) window, but can easily be modified to produce and window the user desires. The window calculation is done in a subroutine within the programme. The subroutine merely fills an array. Calculation of the Dolph-Chebyshev window is a rather intensive process. The remainder of the programme merely formats the data as the DIOS routine expects to see it.

F.5.13. "WRTdoc"

This file contains a small executable routine and a large number of comment lines. Because of the interpretive nature of HPL, and the memory restrictions present in the HP9825, few comments are included in "RUNTEK". Some are included in "WRTdoc", specifically those which give the function of variables within "RUNTEK".

When "WRTdoc" is run, the comments are formatted and stored in the data file "tekdoc". This file is used by the DOC command in DIOS, to provide basic indication of the contents of various arrays, simple variables, strings and flags.

REFERENCES

1. N S Nahman, "The measurement of baseband-pulse risetimes of less than 10^{-9}s", Proc. IEEE, vol. 55, pp855-864, June 1967.

2. N S Nahman, "Picosecond-domain waveform measurements", Proc. IEEE, vol. 66, pp441-454, April 1978.

3. N S Nahman, "Picosecond-domain waveform measurements: Status and future directions", IEEE Trans. Instr. & Meas. vol IM-32, No. 1, March 1983.

4. G Andrieux and C Loty, "A high speed oscilloscope for real-time use", Phillips Tech. Rev., vol. 30, Nos 8/9/10, pp256-261, 1969.

5. Tektronix Co, User's Manual for Model 7100 1GHz phosphor target oscilloscope with 7A29 plugin. Tektronix Co of Beaverton, Oregon, USA.

6a. Hewlett Packard Co., User's Manual for Model 1430C 18 GHz Sampler with 1815B control plugin, Hewlett Packard, California, USA.

6b. HP journal, vol. 24, No. 8, April 1978.

7. W M Grove, "Sampling for oscilloscopes and other RF systems: DC through X-band", IEEE Trans. on Microwave Theory and Tech., vol. MTT-14, No. 12, December 1966.

8a. G Frye, "A new approach to fast gate design", Tektronix Service Scope, No. 52, October 1968.

8b. S-6 Sampler plugin Service Manual, Tektronix Co., Beaverton, Oregon, USA.

9. R A Kiehl, et al., "Selectively doped heterostructure frequency dividers", IEEE Electron Device Letters, vol. EDL-4, No. 10, October 1983.

10. N A Robin and R Ramirez, "Capture fast waveforms accurately with a 2-Channel programmable digitiser", Elect. Des. 3, February 1, 1980.

11a. T P Dagostero and M R Turner, "100 MHz oscilloscope displays innovations in digital storage", Electronics, May 7, 1980, pp161-167.

11b. Le Croy Research Systems Corporation, CAMAC System User's Manual, with TR-8828 ADC plugin, Le Croy International.

12. James R Andrews, "Precision Picosecond-pulse measurements using a high quality superconducting delay line", IEEE Trans on Instr. and Meas., vol. IM-23, No. 4, December 1974.

13. A J Cummins and A R Wilson, "Cryogenic nanosecond pulse recirculator", Proc IEEE, vol. 52, p1749, December 1974.

14. Michael Lampton, "The microchannel image intensifier", Sci. Am., vol. 245, No. 5, November 1981, pp46-55.

15. G Clement and C Loty, "Channel plate electron multipliers in cathode ray tubes", Acta Electron., vol. 16, No. 1, pp102-111, January 1973.

16. R T Gallagher, "Scan conversion CRT shows 7 GHz signals", Electronics, May 31, 1983, pp86-87.

17. Tektronix Co., Service Manual for the 7912AD programmable digitiser, Tektronix Co., Beaverton, Oregon, 97077, April 1978.

18. R W Comerford, "Measurement Science is catching up", Electronics, July 31, 1980, pp86-87.

19. Ray Smith, Lockheed Missile and Space Corporation, Microelectronics Centre, Sunnyvale, California, private communication.

20. R Schwarte, "New results of an experimental sampling system for recording single events", Electron. Lett., vol. 8, No. 4, pp95-96, February 24, 1972.

21. L Stanchi, "Spatial sampling for fast single events", IEEE Trans. Nucl. Sci., p107, 1969.

22. P M Cashin and D A H Johnson, "A novel method for analysing single occurring pulses with nanosecond resolution", The Radio and Electronic Engnr., vol. 38, No. 1, July 1969.

23. M I Skolnik, 'Introduction to RADAR SYSTEMS', McGraw-Hill, 1962.

24. Hewlett Packard Co., 'Model 1815/17 Sampling System service manual', Section IV, VIII, HP, Palo Alto, California.

25. W Michael Henebry, "Avalanche transistor circuits", Rev. Sci. Instr., vol. 32, No. 11, November 1961.

26. Paolo Spirito, "Static and dynamic behaviour of transistors in the avalanche region", IEEE Journal of Solid State Physics, No. 6, April 1971, pp83-86.

27. R H Vandre, "An ultrafast avalanche transistor pulser circuit", Electronic Engineering Applied Ideas, 1977, pp101, Morgan-Grampian.

28. Hewlett Packard Application Note 918, 'Pulse and Waveform Generation with Step Recovery Diodes', HP Palo Alto, California, 1968.

29. T Dreher, 'Fast pulse techniques', E-H Research Laboratories, Inc., Oakland, California, February 1968.

30. J L Moll, S Krakauer and R Shen, "P-N charge storage devices", Proc. IRE, vol. 50, January 1962, pp43-53.

31. H A Wheeler, "Transmission line properties of parallel strips separated by a dielectric board", IEEE Trans on Microwave Theory and Tech., vol. MTT-13, pp172-185, March 1965.

32. C P Wen, "Coplanar Waveguide: a surface strip transmission line suitable for non-reciprocal gyromagnetic device applications", IEEE Trans on Microwave Theory and Tech., December 1969, pp1086-1090.

33. T Kitazawa, Y Hayashi, and M Suzuki, "A coplanar waveguide with thick metal coating", IEEE Trans on Microwave Theory and Tech., September 1976, pp604-608.

34. D A Rowe and B Y Lao, "Numerical analysis of shielded coplanar waveguides", IEEE Trans on Microwave Theory and Tech., vol. MTT-31, No. 11, November 1983.

35. M K Krage and G I Haddad, "Frequency dependant characteristics of microstrip transmission lines", IEEE Trans. on Microwave Theory and Tech., vol. MTT-20, No. 10, October 1972.

36. R E Stegens, "Coplanar waveguide FET amplifiers for satellite communications", Comsat Tech. Rev., vol. 9, No. 1, Spring 1979, pp256.

37. P Silvester and P Benedek, "Equivalent capacitances of microstrip open circuits", IEEE Trans on Microwave Theory and Tech., vol. MTT-20, No. 8, August 1972.

38. Hewlett Packard Co., 'HP Diode and Transistor Designer's Catalog', HP San Hose, 1983.

39. Alpha Industries, 'Microwave Semiconductors', product catalog and application notes, Woburn, MA, 1983.

40. R J Lewandowski, "Automated signal generator level characterisation and verification", IEEE Trans on Instr. and Meas., vol. IM-33, No. 3, September 1984, pp147-154.

41. R Goyal and B T Brodie, "Recent advances in precision AC Measurements", IEEE Trans on Instr. and Meas., vol. IM-33, No. 3, September 1984, pp164-167.

42. T E Linnenbrink, "Effective bits: Is that all there is?", IEEE Trans on Instr. and Meas., vol. IM-33, No. 3, September 1984, pp184-187.

43. M Neil and A Muto, "Tests unearth A/D converter's real-world performance", Electronics, February 24, 1982, pp127-132.

44. R A Malewski, T R McComb and M M C Collins, "Measuring properties of fast digitisers employed for recording HV pulses", IEEE Trans on Instr. and Meas., vol. IM-32, No. 1, September 1984.

45. T M Souders, "A dynamic test method for high resolution A/D converters", IEEE Trans. on Instr. and Meas., vol. IM-31, No. 1, March 1982.

46. K Gardener and M Story, "A test technique for high speed sampling systems", Electronic Engineering, March 1982, pp44-51.

47. A van der Ziel, 'Noise', Englewood Cliffs, Prentice Hall, 1956.

48. R Orwiller, 'Oscilloscope Vertical Amplifiers', Tektronix press, Beaverton Oregon, 1969.

49. "ECM/EW Designs to 40GHz and above accomplished in coax", Microwave Systems News, vol. 14, No. 6, June 1984, p48.

50. Hewlett Packard Co., Semiconductor Division, California, USA, Private communication.

51. Paul R Grey and Robert G Meyer, 'Analysis and Design of Analog Integrated Circuits', 2nd Ed., John Wiley and Sons, New York, 1977.

52. Ian E Getreu, (Tektronix, Inc.), 'Modelling the Bipolar Transistor', Elsevier Publishing Co., New York, 1978.

53. S Kuroda, "A simple stray-free capacitance meter by using an operational amplifier", IEEE Trans. Instr. and Meas., vol. 32, No. 4, December 1983, pp512-513.

54. G Popkirov and N Tabov, "Method for measurement of capacitance, series and shunt resistances of semiconductor junctions and their voltage dependance", Rev. Sci. Instr., vol. 53, No. 6, June 1982, pp864-866.

55. Robert A Chipman, 'Theory and Problems of Transmission Lines', Ch. 5, McGraw-Hill, New York, 1968.

56. A Vladimirescu, K Zhang, A R Newton and D O Pederson, 'SPICE 2G User's guide', University of California Berkeley, 1981.

57. Tektronix Inc, '7912AD Transient Digitiser Operator's Manual', Tektronix, Beaverton, Oregon, 1979.

58. F Guterl, "Instrumentation", IEEE Spectrum, vol. 21, no. 11, November 1984, pp56-57.

59. L Besser, "Microwave circuit design", Electronic Engineering, October 1980, pp103-112.

60. Watkins-Johnson Co., "Application information for thin film cascadable amplifiers", WJ Microwave devices Catalog, 1980, pp220-234.

61. K B Niclas, "Multi-octave performance of single ended microwave solid state amplifiers", IEEE Trans. on Microwave Theory and Tech., vol. MTT-32, no. 8, August 1984, pp896-908.

62. Y Ayasli, "Decade bandwidth amplification at microwave frequencies", Microwave Journal, April 1984, pp71-79.

63. W C Peterson, A A Gupta and D R Decker, "A monolithic GaAs DC to 2 GHz feedback amplifier", IEEE Trans on Microwave Theory and Tech., vol. MTT-31, no. 1, January 1983, pp27-29.

64. F Perez and J Obregon, "Low frequency limitations of ultra-broadband matched microwave amplifiers", Electronics letters, 7th January 1982, vol. 18, no. 1, pp31-33.

65. Calculator Programme Library for the HP-9845 System, Hewlett Packard Company, August 1980.

66. Bob Orwiler, 'Oscilloscope Vertical Amplifiers', Tektronix Inc., Beaverton Oregon, 1969.

67. Tektronix Inc., 'Tek Products Catalog', Tektronix, Beaverton, Oregon.

68. Ray Smith, Lockheed Missile and Space Corporation, Microelectronics Centre, Sunnyvale, California, private communication.

69. William Farnbach, "A scrutable sampling oscilloscope", Hewlett Packard Journal, pp2-8, 1971.

70. Larry Waller, "GaAs waits in the wings", Electronics Week, April 1, 1985, p24.

71. N K Osbrink, R B Levitsky and S B Moghe, "Surface-mounted MICs and MMICs play vital role in military communications", Microwave Systems News, January 1985, p72.

72. S B Moghe, Horng-Jye Sun, T Andrade, C C Huang and R Goyal, "A monolithic direct-coupled GaAs IC amplifier with 12 GHz bandwidth", IEEE Trans on Microwave Theory and Tech., vol. MTT-32, no. 12, December 1984, pp1698-1704.

73. A V Oppenheim and R W Schafer, 'Digital Signal Processing', Prentice Hall, 1975, Chapter 9.

74. W R Bennett, "Spectra of quantised signals", Bell Syst. Tech. Journal, Vol. 27, 1948, pp446-472.

75. K B Niclas and R R Periera, "Feedback applied to balanced FB amps", Microwave Systems News, pp66-69, November 1980.

76. K B Niclas, "On design and performance of lossy match GaAs MESFET amplifiers", IEEE Trans on Microwave Theory and Tech., vol. MTT-30, November 1982, pp1900-1907.

77. K B Niclas, W T Wilser, R T Kritzer, and R R Pereira, "On theory and performance of solid-state microwave distributed amplifiers", IEEE Trans on Microwave Theory and Tech., vol. MTT-31, June 1983, pp447-456.

78. E W Strid, K R Gleason and J Addis, "A dc-12 GHz GaAs FET distributed amplifier", Res Abstracts 1981 Gallium Arsenide Integrated Circuit Symposium, October 1983, p47.

79. User's Manual, Super Compact, 1983.

80. Hewlett Packard Dynamic Signal Analyser HP3561A Service Manual, Hewlett-Packard Company, 1983.

81. A V Oppenheim and R W Schafer, 'Digital Signal Processing', Prentice Hall, 1975. Chapter 10.

82. Athanasios Papoulis, 'The Fourier Integral and its applications', McGraw-Hill, 1962, Chapter 8.

83. Richard C Heyser, "Acoustical measurements by time delay spectrometry", Journal of the Audio Engineering Society, vol. 15, no. 4, October 1967, pp370-382.

84. Jerome Kristian and Morley Blouke, "Charge-coupled devices in astronomy", Sci. Am., vol. 247, no. 4, October 1982.

85. John K Scully, "Frequency response through impulse excitation", Electro-Technology, July 1964, pp53-56.

86. Athanasios Papoulis, 'The Fourier Integral and its applications', McGraw-Hill, 1962, Chapter 10.

87. A M Nicolson, "Forming the fast Fourier transform of a step response in time domain metrology", Electron. Letters, vol. 9, pp317-318, 1973.

88. H A Samulon, "Spectrum analysis of transient response curves", Proc IRE, vol. 39, pp175-186, 1951.

89. Jurg Waldmeyer, "Fast Fourier transform for step-like functions: The synthesis of three apparently different methods", IEEE Trans on Instr. & Meas., vol. IM-29, no. 1, March 1980.

90. Yves Balcou, "A method to increase the accuracy of fast Fourier transform calculations for rapidly varying functions", IEEE Trans on Instr. & Meas., vol. IM-30, no. 1, March 1981.

91. Wen-Sheng Liu and Sheng-He Sun, "The transient response of a sampling oscillographic system excited by a step like pulse having overshoot", IEEE Trans on Instr. & Meas., vol. IM-30, no. 1, March 1981.

92. G D Bergland, "A guided tour of the fast Fourier transform", in 'Digital Signal Processing', ed. L R Rabiner and C M Rader, IEEE Press, New York, 1972, pp228-239.

93. Adly A Girgis and Frederic M Ham, "A quantitative study of pitfalls in the FFT", IEEE Trans. on Aerospace and Electronic Systems, vol. AES-16, no. 4, July 1980.

94. Bidyut Parruck and Sedki M Riad, "Study and performance evaluation of two iterative frequency-domain deconvolution techniques", IEEE Trans. on Instr. and Meas., vol. IM-33, no. 4, December 1984.

95. Frederick J Harris, "On the use of windows for harmonic analysis with the discrete Fourier transform", Proceedings of the IEEE, vol. 66, no. 1, January 1978, pp51-83.

96. Bruce P Bogert, M J R Healy, and John W Tukey, "The quefrency alanysis of time series for echoes: Cepstrum, pseudo-autocovarience, cross-cepstrum and saphe cracking", Proc. Symp on Time Series Analysis, ed. M Rosenblatt, Wiley, 1963, Chapter 15, pp209-243.

97. J M Tribolet, "A new phase unwrapping algorithm", IEEE Trans. on Acoustics, Speech and Signal Processing, vol. ASSP-25, pp170-177, April 1977.

98. F Bonzanigo, "An improvement of Tribolet's phase unwrapping algorithm", IEEE Trans. on Acoustics, Speech and Signal Processing, vol. ASSP-26, pp104-105, February 1978.

99. J B Scott, "Improving confidence in the phase unwrapping algorithm", IEEE Trans. on Acoustics, Speech and Signal Processing, vol. ASSP-32, December 1984.

100. Roger G Cox, "Window functions for spectrum analysis", Hewlett-Packard Journal, September 1978, pp10-11.

101. H Renders, J Schoukens and G Vilain, "High accuracy spectrum analysis of sampled discrete frequency signals by analytical leakage compensation", IEEE Trans. on Instr. and Meas., vol. IM-33, no. 4, December 1984.

102. J Bunton, "Audio testing with an on line computer", IREE Convention Digest, Melbourne, 1977, pp228-230.

103. R W Schafer and R L Rabiner, "A digital signal processing approach to interpolation", Proc. IEEE, vol. 61, June 1973, pp692-702.

104. B Parruck and S M Riad, "An optimisation criterion for iterative deconvolution", IEEE Trans. on Instrumentation and Measurement, vol IM-32, no. 1, March 1983.

105. "A better way to measure transistor speed", Electronics, January 27, 1986.