OBSERVATIONS OF PULSAR RADIO-EMISSION WITH MICROSECOND RESOLUTION

by

Robert Kingsley Allen, B.Sc.(Hons.)

A thesis submitted in fulfilment of the
requirements of the Degree of
Doctor of Philosophy
in the

UNIVERSITY OF TASMANIA

HOBART

December, 1979
(accepted June 1980)

This thesis contains no material which has been accepted for the award of any other degree of diploma in any university, and to the best of my knowledge and belief, the thesis contains no copy or paraphrase of material previously published or written by another person, except when due reference is made in the text of the thesis.

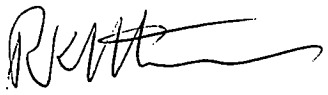

.....

R.K. Allen.

ACKNOWLEDGEMENTS

Many people have given me help during the course of the project described in this thesis. Firstly thanks must go to my supervisor, Dr. P.A. Hamilton, and also to Dr. P.M. McCulloch for assistance with the observations and for continuing interest and advice. Dr. M.I. Large arranged for use of the Molonglo radiotelescope, assisted with the observations and gave some useful advice.

I should like to give special thanks to the Director, Colin Wyld, and the Systems Manager, Tony Hancock, of the University Computing Centre and to my workmates without whose cooperation and assistance this thesis could not have been completed.

Most of this thesis was typed by Sue McQuaid and Robin Wilson and I received invaluable help from Yvonne Shaw and Jenny Sinclair. I greatly appreciate the work of these ladies.

I am deeply grateful for the patience and support of my wife through the years in which I have been occupied with this project.

Finally I acknowledge the financial assistance I received through a Commonwealth Post-graduate Research Award.

CONTENTS

SUMMARY

This thesis describes a new technique for observing pulsar microstructure, which is a class of random fluctuations in pulsar radio-emission on timescales less than one millisecond. Using this technique observations were made with a time resolution of 0.5 μs. No new phenomena were revealed by the improvement in time resolution over that used by other workers. However the first measurement of the microstructure timescale of pulsar PSR 1749-28 was obtained.

The design of the equipment and the computer processing involved in the method are described in the thesis. In this method pulsar signals in a bandwidth of 2 MHz are mixed to baseband and recorded on videotape. Subsequently selected parts of the tape are digitized using the stop action facility of the recorder and an interleaved sampling scheme. The effects of interstellar dispersion are removed by computer processing yielding the maximal time resolution 1/B where B is the bandwidth, that is, 2 MHz. This method avoids the normal restriction, imposed by the speed of the digitizer, that applies when digitizing is done in real time, namely that the bandwidth must be restricted to half the sample rate.

Recordings were made of signals received by the Molonglo radiotelescope at 408 MHz from 15 pulsars. A total of 22 pulses from four different pulsars were strong enough and otherwise suitable for analysis. Each pulse involves 86000 samples, that is, a time window of 20 ms sampled at an effective rate of 4.5 MHz.

The pulses were analyzed at time resolutions down to 0.5 μs and the dynamic spectrum of each in an observing window 15 ms × 2 MHz was obtained. A search was made for frequency-time structures such as those produced by moving, radiating sources and by interference effects in the pulsar magnetosphere. Such structures were not observed.

Auto-correlation and cross-correlation analyses were carried out on

the data. The results for PSR 0950+08 agree with those obtained by other workers. A characteristic microstructure timescale of 250 $\pm$ 50 μs was obtained from the auto-correlation of the power spectrum of PSR 1749-28.

All of the observations are consistent with the amplitude-modulated noise (AMN) model of Rickett (1975) together with interstellar scintillation.

One particular pulse from PSR 0950+08 was analyzed in great detail because some large unresolved intensity fluctuations were apparent. Single intense impulses are not consistent with the AMN model and would be of some astrophysical interest if found. A series of tests were applied to the data; it is concluded that such impulses were not observed.

The relationship of the AMN model to pulsar emission theories and the direction of future research are discussed.

CHAPTER 1

INTRODUCTION AND LITERATURE SURVEY

1.1 INTRODUCTION

Pulsar microstructure is a class of random variations in the intensity of the radio emission from pulsating radio sources, pulsars, on timescales less than 1 ms and typically a few microseconds. These timescales are less than the durations of subpulses, typically 5 ms, which in turn are less than the duration of the average pulse, typically 15 ms. It is likely that observations of microstructure, the shortest duration features in the emission, will lead to an understanding of the fundamental pulsar emission processes. Indeed observations by Hankins (1971, 1972) suggested that the radio emission from pulsar PSR 0950+08 may consist of impulses of less than 8 μs duration, this being the time resolution of Hankins' observations. Substantial improvements upon this time resolution are needed and it is also important for other characteristics of the microstructure such as its polarization and spectrum to be investigated.

This thesis describes a technique for achieving time resolutions below 1 μs, observations using this technique and investigations of the spectrum of the emission over a 2 MHz bandwidth.

Observations of pulsar radio emission with microsecond resolution requires a good signal-to-noise ratio. In addition there is the problem of dispersion of the signals during propagation through the interstellar medium. The result is that for observations below 1GHz using "conventional" receivers the best time resolution for even the nearer pulsars exceeds some tens of microseconds.

The technique described in this thesis yields a time resolution, $\Delta t = 1/B$, where the receiver bandwidth, B, is several MHz. The pulsar signals are recorded on videotape and later digitized using the stop-action replay facility of the recorder and an inter-leaved sampling scheme.

The data is then processed by a computer to remove the effects of dispersion using a method similar to that of Hankins (1971). The technique has been implemented using a domestic-quality videotape recorder giving a bandwidth of 2 MHz and hence a time resolution of 0.5 μs.

Recordings were made of signals received by the Molonglo radio-telescope at 408 MHz from 15 pulsars. A total of 22 pulses from four different pulsars were strong enough and otherwise suitable for analysis. The intensity profiles of these pulses were obtained with time resolutions down to 0.5 μs. As well the data was processed to give the dynamic spectra of the pulses in observing windows up to 15 ms x 2 MHz. In all cases the observations were found to be consistent with the amplitude-modulated noise (AMN) model of Rickett (1975). Effects of interstellar scattering were also found.

Since pulsar microstructure has been found to be an inherently random phenomenon, a statistical method is required for averaging and summarizing the data as well as for formulating theoretical predictions. Such a method is the calculation of the auto-correlation and cross-correlation functions of the data. The data for two pulsars, PSR 0950+08 and PSR 1749-28, was analysed in this way and in both cases the results are consistent with the AMN model. Further the results for PSR 0950+08 are consistent with those of Hankins. PSR 1749-28 had not previously been observed with microsecond resolution. The first measurement of the characteristic microstructure timescale of this pulsar has been obtained. It is 250±50 μs and came from calculating the auto-correlation of the power spectrum of the emission, a method anticipated by Rickett (1975) but whose use has not so far been reported.

The organization of this thesis is as follows. Section 1.2 (following) gives an overview of interstellar dispersion and the reduction of its effects by so called "dispersion removal" methods. Section 1.3 contains

a review of published observations of pulsar microstructure with the AMN model being described in sub-section 1.3.1. Section 1.4 contains a review of theories concerning microstructure (sub-section 1.4.3) preceded by a summary of their background, i.e. models for the pulsar magnetosphere, and radio emission mechanisms.

Chapters 2, 3, 4 and 5 all deal with the equipment and techniques used in the data acquisition. Chapter 2 describes the Molonglo telescope, the videotape recorder and the "recording chassis". Chapter 3 describes the use of this equipment during the observing session and preliminary analysis of the tapes. Chapter 4 deals with the principles and design of the equipment and computer programs used to digitize the data during stop-action replay of the tapes. Chapter 5 describes calibration and measurement of the performance of the system.

A detailed description of the effects of interstellar dispersion and of the numerical methods used for "dispersion removal" appear in Chapter 6. The final section of this chapter, Section 6.4, deals with factors which may prevent the time resolution 1/B being achieved. These include the effect of interstellar scattering of the signals.

Chapter 7 describes the high time resolution and dynamic spectrum observations of pulses from PSR 0950+08, 1749-28, 1133+16 and 1642-03. One of the pulses from PSR 0950+08 was investigated in some detail to determine if some large intensity spikes were single impulses or just statistical fluctuations consistent with the AMN model. These investigations are described in Chapter 8.

Chapter 9 is concerned with the auto- and cross-correlation analyses of intensity and power spectra which have already been mentioned. In conclusion Chapter 10 contains some suggestions for further research.

1.2 DISPERSION AND ITS REMOVAL

The dispersive effect of the interstellar medium on pulsar radio emission is depicted in Figure 1.2A. An idealized dispersed pulse is shown sweeping downward in frequency according to the "dispersion law"

$$t = t_o + k\ DM/f^2$$

where t is the pulse arrival time at observing frequency f, k is a constant and DM is the dispersion measure of the pulsar. DM is the integral along the line of sight of the interstellar electron density, n_e, that is

$$DM = \int_o^z n_e(z)\ dz$$

where z is the distance from the pulsar. DM is usually expressed in the hybrid units cm^{-3}parsec (cm^{-3}pc). The pulsar with least DM, hence probably the nearest to Earth, is PSR 0950+08 for which DM = 2.97 cm^{-3}pc; the median DM of the known pulsars is over 100 cm^{-3}pc.

To observe intensity structure on a timescale, Δt, requires a receiver bandwidth $B > 1/\Delta t$. However as B is increased so is the smearing effect due to dispersion. This is measured by the bandwidth sweep time, t_s, which is the difference in pulse arrival time between the upper and lower edges of the passband. It is given by

$$t_s = 2k\ B\ DM/f^3 = [8.3\ ms]\ B_6\ DM/f_8^3$$

where B_6 is the bandwidth in MHz and f_8 is the observing frequency in hundreds of MHz. Hence even for low DM pulsars the smearing effect is important at frequencies below 1 GHz and severe below 100 MHz. It can be minimized by adjusting the bandwidth so that $t_s = 1/B$. The time resolution, t_{min}, that can be obtained using the optimum bandwidth is

$$t_{min} = (2k\ DM/f^3)^{\frac{1}{2}} = [91\ \mu s]\ DM^{\frac{1}{2}}\ f_8^{-3/2}$$

Since this generally means narrow bandwidths there is a loss in signal-to-noise ratio unless multiple receiving channels are used.

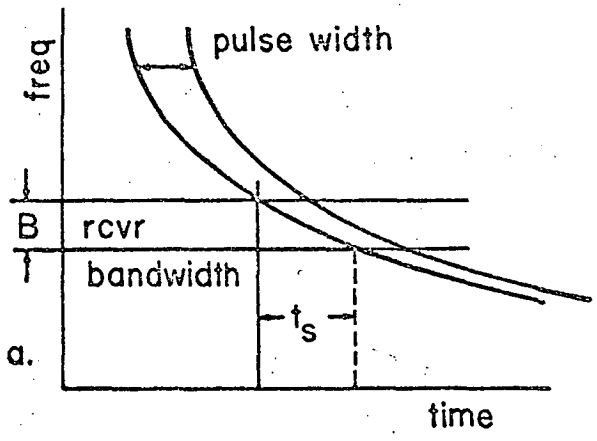

a Schematic of a typical dispersed pulse sweeping through a receiver bandwidth in a time t_s.

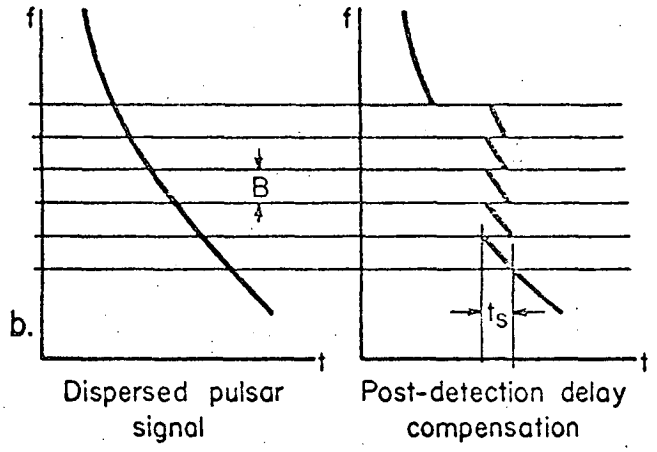

b An idealized dispersed pulse received by a bank of narrow band filters whose outputs are detected and delayed to compensate for dispersion.

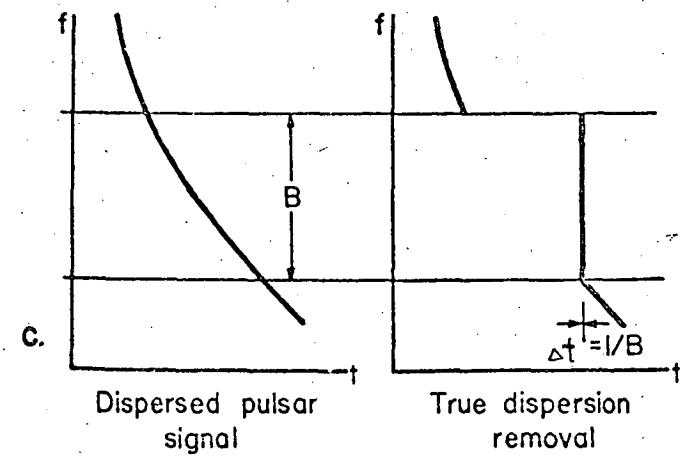

c For an idealized dispersed pulse processed by the true dispersion removal filter the minimum time resolution is the reciprocal of the total receiver bandwidth.

FIGURE 1.2A Illustration of dispersion removal by the post-detection method and by the "true" pre-detection method (from Hankins 1974).

The post-detection method of dispersion removal involves use of a large total bandwidth divided into many narrow-band channels. The detector outputs from these channels are added together after each has been delayed to compensate for the f^{-2} dispersion law. This is illustrated in Figure 1.2A(b). The delay scheme, which must be adjustable for different combinations of DM and centre frequency, is easily implemented with computer software (e.g. Taylor 1974) and has been implemented with hardware (e.g. Orsten 1970, Boriakoff 1976). However any particular set of filters that defines the channels will be far from optimum for many combinations of DM and observing frequency. Furthermore the optimum time resolution, t_{min}, exceeds 50 μs for all pulsars and $f < 1$ GHz.

With the pre-detection method of dispersion removal the time resolution is 1/B for arbitrary bandwidths. That is, the method which was developed by Hankins (1971) can provide time resolutions $\sim$ 1 μs. It is depicted in Figure 1.2A(c). Conceptually the signal is passed through a phase-shifting filter before it reaches the detector and this inserts delays appropriate to the frequency. This filter has a transfer function that is the inverse of the dispersive part of the transfer function of the propagation through the interstellar medium. In practice the filter is implemented numerically using a computer program which involves the discrete Fourier transform. The signal voltage is sampled and recorded for subsequent computer processing.

The method used by Hankins for recording the data is as follows. The signal in a 125 KHz bandwidth is split into two identical parts and each is converted to baseband, that is, to have zero centre frequency. A common local oscillator is used and its output is supplied to the two identical mixers in quadrature. Then during a 50 ms window synchronized with the pulsar the two baseband signals are sampled at 7 μs intervals,

digitized and recorded on digital magnetic tape. Variations on this system have been developed to provide a second 125 KHz channel (Cordes and Hankins 1977) and recently to provide a single channel of 1.25 MHz (Hankins and Boriakoff 1978). This latter system involves digitizing two 6-bit samples every 0.8 μs within a 26 ms window and storing them on magnetic tape using high speed tape transports.

The maximum bandwidth that can be used with Hankins' scheme is limited by the rate at which the data can be digitized and stored by the computer equipment. This means that, neglecting possible effects of interstellar scattering (see Section 6.4.2), the best time resolution that can be obtained is also limited by this hardware. I have developed a technique, which is described in Chapters 2 to 4, for avoiding the limitation. The data is not digitized during the observations but is recorded in analogue form on videotape using a domestic-quality recorder. Digitizing of the data for strong pulses takes place later using an inter-leaved sampling scheme and the stop-action facility of the recorder. The actual sampling rate is quite low, that is, samples are obtained tens of microseconds apart, but the effective rate is high, namely 4.5 MHz (0.22 μs apart). In this way pulsar signals in a bandwidth of 2 MHz at 408 MHz have been processed to remove dispersion yielding a best time resolution of 0.5 μs. The technique could readily be extended to larger bandwidths using a more expensive videotape system or a video disk as the recording medium.

1.3 OBSERVATIONS OF MICROSTRUCTURE BY OTHERS

1.3.1 Early Observations and the AMN Model

In the first account of the pre-detection method of dispersion removal Hankins (1971) reported observations of pulsar PSR 0950+08. These were made using the 305 m diameter antenna of the Arecibo Observatory (refer Bull. A.A.S., 8, 10). The receiver had a centre frequency of 111.5 MHz and a bandwidth of 125 KHz giving a time resolution of 8 μs after dispersion removal. Some 1500 pulses were processed and among them three classes were distinguished:

(i) those pulses having amorphous, *i.e.* continuous and noise-like, emission over most of the integrated profile;

(ii) those with emission confined to isolated bursts (later) called micropulses;

(iii) a few (five) showing quasi-periodic structure with periods in the range 300 to 500 μs.

Very strong (up to 40 000 Jansky, 1 Jansky = 10^{-26} W m^{-2} Hz^{-2}) intensity spikes were observed. These spikes were unresolved, that is, their widths were 8 μs. Similar spikes were also observed by Hankins (1972) in pulses from PSR 1133+16 at three frequencies, 111.5, 196 and 318 MHz. The autocorrelation function (ACF) was computed for each strong pulse and averaged over groups of about 200 pulses. The average ACF at each frequency showed three features:

(i) a spike at zero lag formed by a rapid decrease in the ACF in a time $\sim 1/B$ where B is the receiver bandwidth;

(ii) a broad feature extending to about the width of the integrated pulse profile;

(iii) a "microstructure feature" extending to times less than 1 ms and sitting atop the broad feature.

Straight lines could be fitted to the microstructure feature and the broad feature (for lags < 3.5 ms) Hankins defined the width of the microstructure feature to be the lag, Δt_μ, at which the two straight lines intersect, that is, the break in slope. For PSR 0950+08 at 111.5 MHz the microstructure timescale, Δt_μ, was found to be 175 μs while for PSR 1133+16 Δt_μ was about 575 μs independent of observing frequency.

An explanation of these ACF features was provided by the amplitude-modulated noise (AMN) model of Rickett (1975). The emitted pulsar signal is assumed to be the product of a (real) amplitude modulation function, a(t), and random noise, n(t). The noise is assumed to be Gaussian and white over a bandwidth exceeding $1/\Delta t_a$, where Δt_a is the time in which a(t) changes significantly. That is, the noise fluctuations are assumed to be rapid compared to the variation in the modulation. When the emission is observed with a receiver of bandwidth B and dispersion is removed there is no correlation between samples of the voltage entering the detector that are a time 1/B apart. Further, when the microstructure is well resolved by the receiver, that is, $\Delta t_a \gg 1/B$, the normalized ACF of intensity (measured by the output of the square-law detector) consists of a zero lag spike and a broader feature with the shape of the ACF of $a^2(t)$. If there is no post-detection smoothing then the width of the spike is 1/B and it has height 0.5. That is, at zero lag the spike and the broad component are equal. A possible deviation from the model is that the noise process may not be Gaussian, and if its modulation index is not unity then this can be detected by measuring the height of the spike.

The broader part of the ACF exhibits different features for the various modulating mechanisms present. For example, if $a^2(t)$ is the product of a stochastic function representing the microstructure (with timescale < 1 ms), a stochastic function representing the subpulses

(with timescale $\sim$ 5 ms) and a deterministic function representing the mean pulse profile then the model ACF is as shown in Figure 1.3.1A. Further details are given in Chapter 9.

Rickett's presentation of the AMN model assumes white Gaussian noise as the basic emission process. Cordes (1976 b) considered polarized shot noise as the basic process. In this case the modulation could be in the strength of each shot or in the temporal density. However, as expected from the Central Limit Theorem, for high temporal densities the predictions of the shot noise model are the same as those of the AMN model.

Now returning to the observations of Hankins (1972), it was recognised that the large intensity spikes observed in single pulses from PSR 0950+08 and 1133+16 could be statistical fluctuations. Hankins attempted to place confidence limits on the individual intensity profiles with the AMN model acting as a null hypothesis. The fact that the spikes were unresolved at best time resolution is consistent with the model. This is expected when there is no post-detection smoothing as are large random fluctuations.

In general when the receiver input is white Gaussian noise the intensity follows a chi-squared distribution with n degrees of freedom where n is determined by the relative pre- and post-detection passbands, roughly, by the product of bandwidth and post-detection time constant (Hankins and Rickett 1975). The mean square fluctuation of intensity, I, is

$$\langle (I - \langle I \rangle)^2 \rangle = (2/n) \langle I \rangle^2$$

where $\langle\ \rangle$ denotes an expectation value. When there is no post-detection smoothing n = 2 and the distribution is exponential. In particular the rms fluctuation equals the mean intensity. In the AMN model the mean and hence the distribution is not stationary. Neglecting system noise the mean is

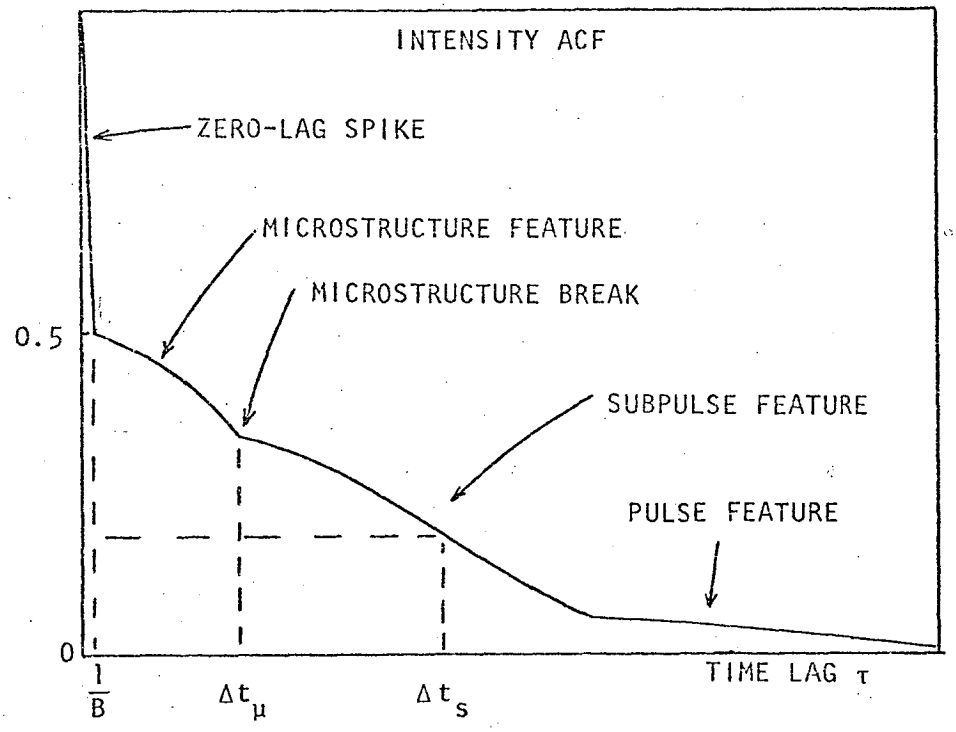

FIGURE 1.3.1A The model intensity ACF predicted by the ACF model for a pulsar whose subpulses have an average width of Δt_s. B is the receiver bandwidth. (After Rickett 1975.)

proportional to the hypothetical modulating function, $a^2(t)$, and is itself stochastic. The confidence limits against which an intensity spike is compared are therefore subject to significant estimation errors.

Hankins (1972) calculated a running mean of the PSR 0950+08 data over a window of 175 μs which is the timescale derived from the ACF results. Since many of the intensity fluctuations were many (eg. ten) times the running mean it was clear that the modulation often varied more rapidly; some micropulses clearly had widths much less than 175 μs. In fact Hankins noted that some micropulses shown in Figure 1.3.1B were "unresolved ($\leq$ 8 μs duration)". However this stronger conclusion is not supported by the published figure. A timescale $\sim$ 50 μs apparently gives a satisfactory fit to the data; loosely speaking, 8 μs is not a statistical necessity. This matter is discussed further in Chapter 8 where I describe the application of the running mean method to a single pulse observed with 0.5 μs resolution.

Another early series of observations using pre-detection dispersion removal should be mentioned. These were of PSR 1919+21 at three frequencies, 74, 111 and 318 MHz. The mean ACF for this pulsar at the two lower frequencies showed a microstructure timescale of about 1.3 ms (Cordes and Hankins 1973, Cordes 1975). However the microstructure was absent at 318 MHz, as were the deep notches in the average pulse profile that Hankins (1973) discovered at the lower frequencies. These notches are 8.3 ms apart and 1.4 ms wide at half depth. (The post-detection method for observing this pulsar at 111 MHz yields a best time resolution of about 0.5 ms). Drifting subpulses are "blanked out" as they drift through the longitude of the notches. The relationship between the drifting subpulse, the microstructure and the amorphous components of emission from PSR 1919+21 were investigated further by Cordes (1975). Observations at 196, 318 and 430 MHz, all using a 32 x 20 KHz post-detection device

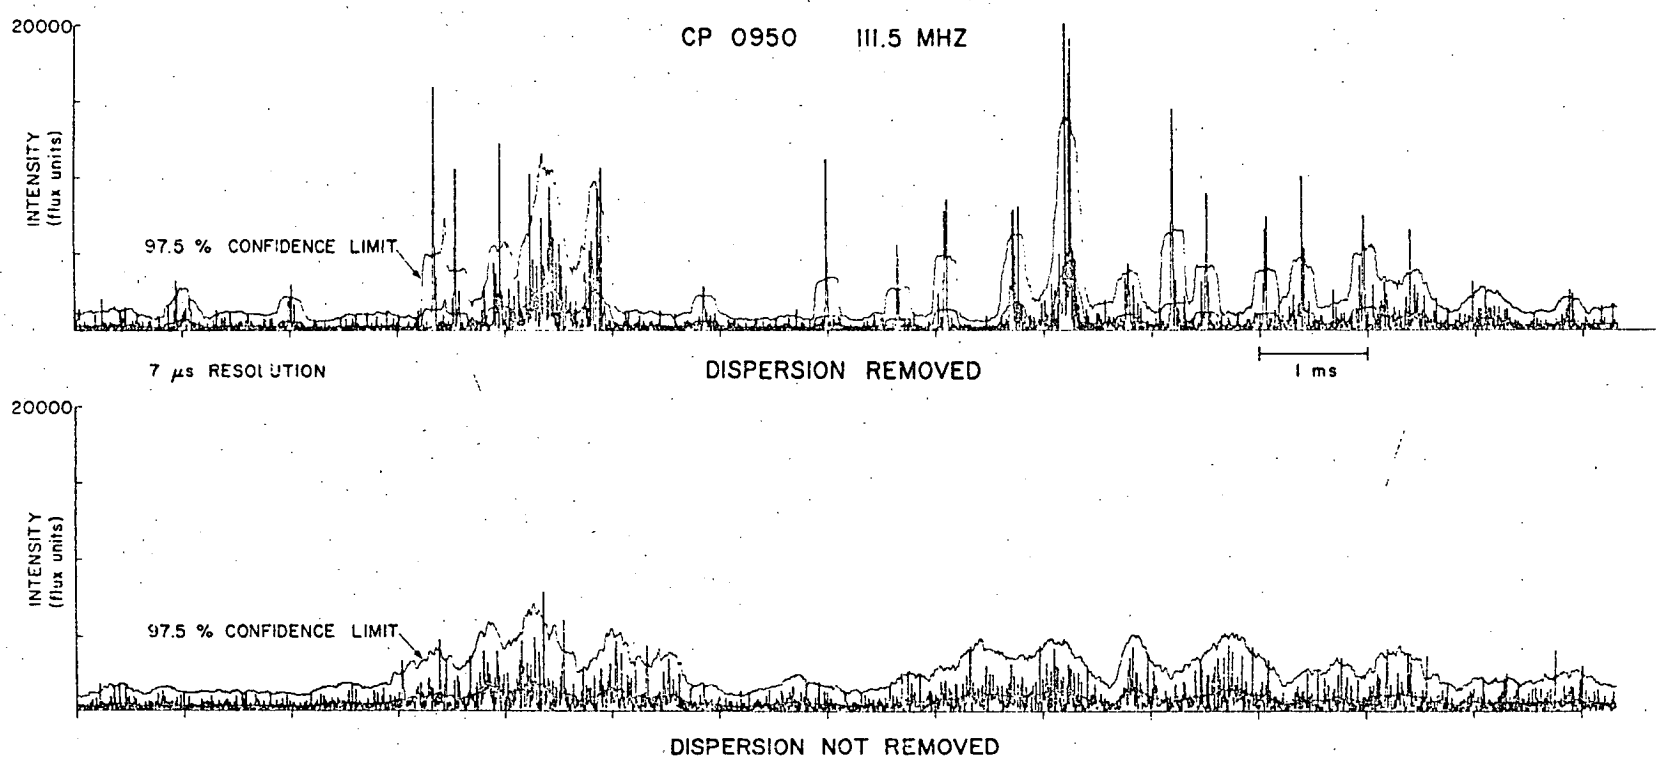

FIGURE 1.3.1B One pulse from PSR 0950+08. Bandwidth 125 kHz. (Hankins 1972)

(Boriakoff 1976), were added to the series of Hankins. The fractional notch depth was found to fall from 0.5 at 74 MHz to 0.2 at 196 MHz. The depth at 74 MHz was used to place an upper bound of 0.9 ms on the time smearing due to interstellar scattering.

1.3.2 Later Observations

This section contains summaries of all known observations of pulsar microstructure published since 1974 with the exception of polarization measurements - see next section. Unless otherwise stated all of the observations were made at Arecibo and dispersion was removed using the pre-detection method.

PSR 0950+08 was observed at two frequencies, 111.5 and 318 MHz, simultaneously by Rickett, Hankins and Cordes (1975). The time resolution was 16 μs at each frequency. The ACF of intensity at 318 MHz was found to contain a microstructure feature similar to that previously found at 111.5 MHz (Hankins 1972) and of about the same timescale, namely 175 μs. Furthermore the cross-correlation (CCF) of the intensity in the two channels averaged over 254 strong pulses also showed the feature but without the zero lag spike. The correlation between the microstructure component of the emission at the two observing frequencies was estimated to be (0.5 ± 0.2). This result came from comparing the strength of the microstructure feature in the CCF to its strength in the two ACF's. The time lag of the peak of the feature in the CCF could be measured to an accuracy $\sim$ 20 μs if model-fitting were used. It was proposed therefore that future observations could measure the dispersion of this pulsar to an accuracy of about 1 in 10^5, that is $\pm 2 \times 10^{14}$ electrons cm^{-2}. Rickett (1975 b) reported that there was no difference between the CCF calculated separately for the two components of the integrated profile of this pulsar. In particular the lag was the same to an accuracy $\sim$ 20 μs

(presumably). This is an important result since the spacing of the two components of PSR 0950+08 is reduced by about 2 ms between 111 and 318 MHz. (The profile is simple above 400 MHz.)

Rickett and co-workers also studied the dynamic spectrum of PSR 0950+08 over a bandwidth of 125 KHz at 111 MHz. This was obtained numerically with frequency-time resolutions of (140/m KHz) x (7m μs) for m = 8, 16, 32 etc. (e.g. 18 KHz x 56 μs). In each case there were 2 degrees of freedom per resolution cell so that under the AMN hypothesis the intensity distribution for each resolution cell was exponential. Each of the micropulses shown in Figure 1.3.1B was found to have a spectrum that was deeply modulated, randomly and independently from its neighbours. It was concluded that the observations were in general agreement with the AMN model.

Periodicity in the microstructure of PSR 2016+28 was discovered by Boriakoff (1976). He observed this pulsar at 430 MHz using post-detection dispersion removal. The time resolution was 60 μs and the total bandwidth was 380 KHz. The intensity ACF averaged over strong pulses had a broad peak at the lag P_μ = 0.9 ms, and single pulse records showed trains of micropulses with quasi-periods in a broad range centred on 0.9 ms. The quasi-period would vary from pulse to pulse; the Q of this periodicity was low ($\sim$ 4). PSR 2016+28 has drifting subpulses ($P_2 \sim$ 10 ms, P_3 = 3P - 15P, P = 0.558 s) and Backer (1973) had reported microstructure apparently being preserved along a drift path. Boriakoff suggested that the persistence of the period, P_μ, from one pulse to the next might explain Backer's observations.

Cordes (1976 a) also observed PSR 2016+28 at 430 MHz. This was with a bandwidth of 250 KHz and a time resolution of 12 μs (there was numerical post-detection smoothing). The microstructure quasi-period, P_μ, was found to average 0.9 ms but vary from 0.6 ms to 1.1 ms both between

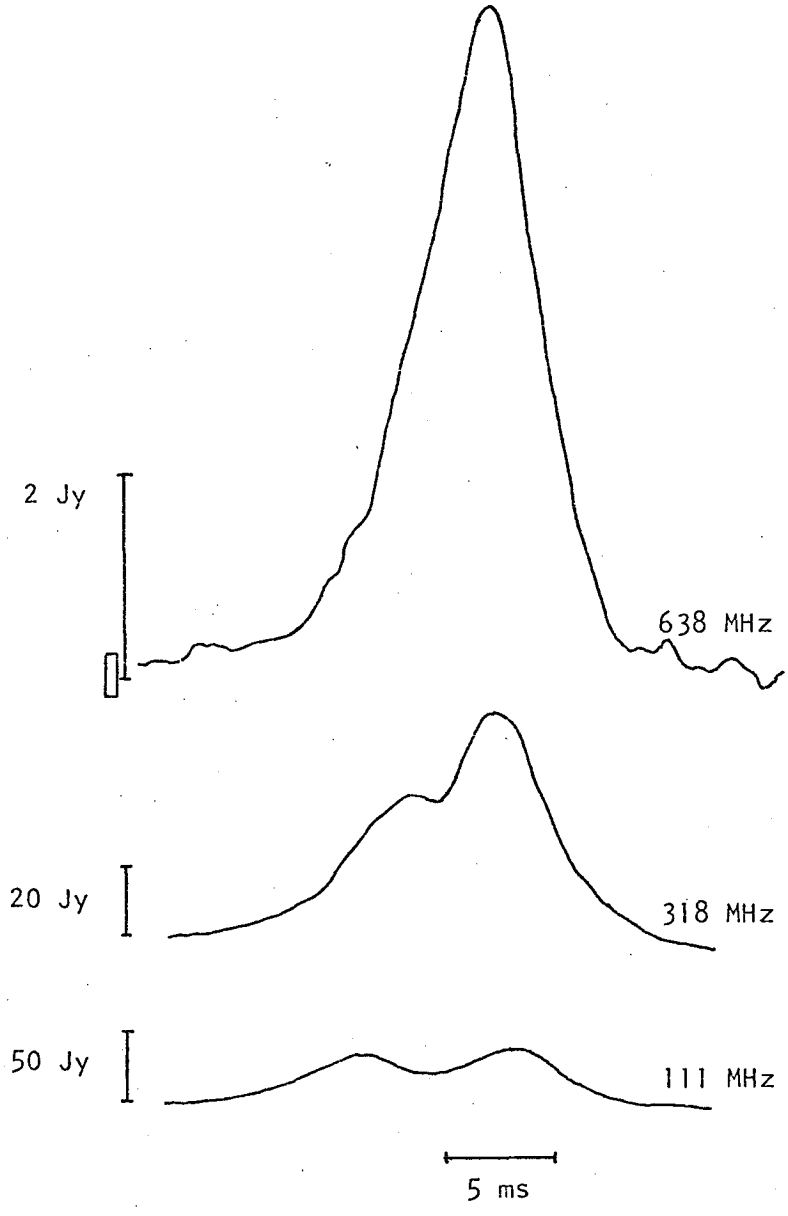

FIGURE 1.3.2A

Integrated pulse profiles of PSR 0950+08 at three frequencies: 638 MHz (P.A. Hamilton, unpublished); 111, 318 MHz (Rickett *et al* 1975).

successive pulses and between subpulses of the same pulse. There was some evidence for the quasi-period having less variation along a few drift paths. However the stronger hypothesis of microstructure being correlated down a drift path was not supported by these observations. This was tested by calculating CCFs between subpulses from one period to the next - the average CCF did not have a microstructure feature. However the average ACF did have a microstructure feature with timescale Δt_μ = 280 μs. Cordes used the height of the zero lag spike to test the AMN model. He deduced 1.02 ± 0.05 as the modulation index of the noise which agrees with the value of unity predicted by the AMN model. The modulation index, m_μ, which represents the strength of the microstructure relative to the steady part of the subpulse emission (see Section 9.2.1), was measured to be 0.5. Subsequent measurements (Cordes 1978b) over the frequency range 318 to 1410 MHz have shown that m_μ falls off faster than f^{-1} above f = 430 MHz ; no microstructure is evident at 1410 MHz.

Following observations of PSR 1133+16 at 2380 MHz Hankins *et al.* (1977) suggested that the microstructure timescale, Δt_μ, for this pulsar varies as $f^{-0.25}$, that is, with the same frequency dependence as the separation of the components in the integrated profile. However in a review Cordes (1978 b) concluded that there was no obvious frequency dependence of Δt_μ for this pulsar nor for PSR 0950+08. The conclusion for PSR 1133+16 follows from observations by Cordes and Hankins (1977) who found that at 430 MHz there were often two breaks in the ACF for this pulsar. These occurred near 370 and 700 μs for Component I and 325 and 850 μs for Component II. There was also a longitude variation in the strength of the microstructure. Modulation indices, m_μ, were greater for Component I and were greater when the longer timescale structure was present.

Hankins and Boriakoff (1978) observed PSR 0950+08 at 430 MHz with the best time resolution obtained so far by the Arecibo group, namely 0.8 μs corresponding to a bandwidth of 1.25 MHz. Autocorrelation analysis was performed

on 250 pulses and the height of the zero lag spike was found to be 0.5 in agreement with the AMN model. The timescale Δt_μ was found to be about 130 μs and independent of pulsar longitude over the main pulse. The latter result came from calculating ACF's separately over ten sectors each of 2.3^o, that is, 1.6 ms. A second breakpoint at 40 μs was found in a few ACF's and was attributed to a few strong micropulses. The interpulse of this pulsar was also observed with 0.8 μs resolution (Hankins and Boriakoff 1977) and for this Δt_μ was found to be about 75 μs, that is, half that of the main pulse. Over much of the main pulse the modulation index was about 0.37 but it was $\sim$ 0.2 near the leading edge. Periodicity was also observed in some pulses. Quasi-periods were mostly about 400 and 600 μs but ranged from 150 μs to 1.1 ms. In two pulses there was a smooth decrease in period by a factor of two across 10^o of pulsar longitude. Most pulses were reported to be of the intermittent type, that is, consisting of isolated micropulses. Some of the micropulses were said to be unresolved at the full 0.8 μs resolution. However this is open to argument and will be discussed in Section 8.6.

One other group has used the pre-detection method of dispersion removal. Kardashev *et al.* (1978) observed seven pulsars with a time resolution of 10 μs at 102 MHz using very similar techniques to those of Hankins. They used the BSA radio-telescope of the P.N. Lebedev Physics Institute which has an effective area of 20 000 m^2, comparable with that of the east-west arm of the Molonglo radio-telescope but less than half that of the Arecibo dish. The microstructure timescales obtained for PSR 0950+08 and for the two components of PSR 1133+16 and the quasi-period for the latter agree with those of Hankins and coworkers (Cordes 1978 b). However the quasi-period observed in some pulses from PSR 0950+08 was 2.1 ms, twice that observed by the Arecibo group.

A second quasi-period was observed for PSR 1133+16 namely 160 μs which agrees with the observations of Ferguson *et al.* (1976) and Ferguson and Seiradakis (1978) at frequencies ∿ 2 GHz. Kardashev and coworkers found a microstructure timescale of 90 μs for PSR 1919+21 while Cordes (1975) found 1.3 ms. They also observed a 400 μs periodicity for this pulsar. Other pulsars observed were PSR 0329+54, 0809+74, 0834+06 and 1508+55. The results are summarized in Table 1 of Section 1.3.4.

Ferguson *et al.* (1976) avoided the need to remove dispersion by observing a low DM pulsar, PSR 1133+16, at a sufficiently high frequency, 1420 MHz, using the Effelsberg 100 m radiotelescope. The bandwidth was 1 MHz for which the dispersion sweep time and hence the best time resolution is 14 μs. This means that the bandwidth-time resolution product was about 13, that is, 26 degrees of freedom per sample of intensity. Autocorrelation analyses of strong pulses were reported to confirm the results of Hankins (1972). A quasi-period of 150 μs was found both as an ACF peak and in single pulses. The intensity traces of three consecutive pulses were published; one of these showed the periodicity. The peak alignments marked in this figure are apparently not statistically significant.

Bartel (1978) also observed PSR 1133+16 using the Effelsberg radiotelescope. Data for six very strong pulses was obtained in a three hour observing session. Two receiver channels 20 MHz apart at 2700 MHz were used and these had bandwidths of 0.5 or 2 MHz for which the dispersion sweep time at 2700 MHz for this pulsar is 1 or 4 μs. No dispersion removal was performed; the sweep times were less than the post-detection time constants of 2.6 or 6 μs. Microstructure was observed in the six pulses. This occurred in both channels and was correlated to the extent that for five of the pulses the difference between the intensity traces (with those for one channel shifted by 41 μs to compensate for the dispersion delay over 20 MHz) yielded noise-like residuals. In the sixth pulse the

microstructure seemed to be much less correlated. Auto- and cross-correlations revealed microstructure timescales of about 100 and 600 μs. Some of the correlation functions also had features with half-widths in the range 6 to 30 μs. Since the correlation functions were of single pulses, not averages over many, and not normalized in the same way, these results of Bartel cannot be compared directly with the 111/318 MHz results of Rickett *et al.* (1975). However it does appear that Bartel observed greater than 50% correlation of the microstructure for a frequency separation of 20 MHz at 2700 MHz. An improvement upon this result will be described in the next paragraph.

The microstructure of PSR 1133+16 was also studied, and in some detail, by Ferguson and Seiradakis (1978). All of the observations were made without dispersion removal at frequencies of 1415, 1700 and 2700 MHz using the Effelsberg dish. The observations were of both total power and polarization (discussed in the next section) and with various time resolutions in the range 13 to 55 μs. Bandwidth-time resolution products ranged from 13 to 1600, the latter using a 30 MHz bandwidth. Average ACFs showed microstructure breaks near 400 μs at both 1700 and 2700 MHz but this time lag varied from 250 to 550 μs for averages over different groups of up to 64 pulses. Cross-correlation analyses similar to those of Rickett *et al.* (1975) were made of simultaneous observations using two channels at various spacings near 2700 MHz. It was found that the microstructure of Component II of the pulse profile was 80% correlated over 100 MHz. For channel separations of 42 MHz or less there was almost 100% correlation. That is, under the AMN model the modulating function was essentially the same in both channels.

Ferguson and Seiradakis found quasi-periodic behaviour in the microstructure of individual pulses and established a characteristic quasi-period as follows. The ACFs of single pulses were inspected for

periodic patterns in the peaks or enhancements. If at least three "harmonics" were present then the "period" was measured. From a total of 46 such pulses showing periodicity (apparently out of about 100 strong pulses) the distribution of "periods" was found to show a harmonic pattern. Hence the "preferred" quasi-periods for PSR 1133+16 are apparently multiples of about 185 μs.

1.3.3 Polarization Observations

The observations described in this Section have no parallel with those described in later chapters of this thesis. However the polarization of microstructure appears to be important to its understanding.

Cordes and Hankins (1977) made polarization observations of four pulsars (PSR 0950+08, 1133+16, 1919+21 and 2016+28) with pre-detection dispersion removal. The centre frequency was 430 MHz and the bandwidth of each channel, that is, of the left circular polarization and right circular polarization channels, was 125 KHz giving a time resolution of 8 μs. Samples 8 μs apart of the Stokes parameters (e.g. Kraus 1966 p. 116) were calculated from the dispersion-removed samples of pre-detector voltages. This data was examined in three forms, namely, average Stokes parameter pulse profiles, average ACFs of the Stokes parameters and single pulse profiles. For each pulsar there were between 130 and 200 pulses, sufficient to divide the data into longitude ranges (e.g. for PSR 1133+16: Component I, saddle, Component II). It was found that generally the polarization and strength of microstructure is dependent on longitude while the time-scale, Δt_{μ}, remains constant.

In general micropulses were found to be more strongly polarized than the subpulses in which they occur. In turn, subpulses are more polarized than average profiles, which follows from the diversity of subpulse polarization. Just as large changes in polarization often occur on the edges of subpulses, similar but faster changes tend to occur on the

edges of micropulses. These changes include transitions between orthogonal modes, that is, reversals in the sense of circular polarization and jumps in the position angle apparently by 90^{o}. However while the position angle tends to change within a subpulse at the same rate as in the average pulse profile, there is a suggestion that, at least for PSR 2016+28, it is constant within a micropulse (Cordes 1978 b, Backus and Cordes 1978).

The polarization of PSR 2020+28 was studied by Cordes, Rankin and Backer (1978) in an attempt to understand the phenomenon of orthogonal modes. The data included observations at 430 MHz with a time resolution of 8 μs. Average ACFs of Stokes parameters were obtained for three ranges of pulsar longitude (PSR 2020+28 has a double profile) from 400 pulses. A microstructure timescale, Δt_{μ}, of 110 μs was found (Cordes 1978 b). The correlation between position angle and sense of circular polarization and hence the existence of orthogonal modes as fundamental processes was firmly established. At any instant the emission from PSR 2020+28 is in only one mode, while the emission from the pulsars observed by Cordes and Hankins (1977) is apparently a superposition of modes. The transitions between modes tend to occur on subpulse boundaries but not only there. No conclusion was drawn concerning micropulse boundaries.

The study of the microstructure of PSR 1133+16 by Ferguson and Seiradakis (1978) included polarization observations at 1720 and 2650 MHz. Dispersion was not removed and the time resolution was about 40 or 50 μs, with time-bandwidth products 210-1560. The four Stokes parameters were obtained electronically and digital samples of these were recorded on magnetic tape. The observations were analysed as single pulse and average profiles; a few individual pulses were discussed in detail. A fit of a model was made to all the Stokes parameter data of one particular pulse, which consisted of a quasi-periodic train of micropulses.

The model was the generalized relativistic single vector model of Ferguson (1976 a), that is, a "light cylinder" geometric model. A separate "single vector" was allowed for each micropulse. The inferred position angle at the source of emission was orthogonal for successive micropulses, and the source was inferred to be moving at more than 0.99 the speed of light.

1.3.4 Summary and Discussion

Microstructure has been observed in the emission from twelve pulsars (Table 1). Only one pulsar, PSR 1237+25, is known not to show microstructure (Cordes 1978 b). The phenomenon is sporadic in that it does not appear in every pulse; many pulsars are amorphous. Furthermore the widths of individual micropulses are variable, that is, except for occasional periodicity the "modulation" is stochastic. The ACF of intensity when averaged over sufficiently many pulses shows a microstructure feature whose width is apparently a stable property of the particular pulsar. However the width, Δt_μ, is subject to measurement errrors especially because of its definition in terms of the intersection of trend lines and because an ACF can be dominated by a few strong pulses (their intensity squared enters the summation). The situation for PSR 1133+16 is especially confused in that multiple timescales may be present.

Microstructure is broad-band in that correlations exceeding 50% have been observed for PSR 0950+08 between 111 and 318 MHz and for PSR 1133+16 over 100 MHz at 2700 MHz. However the spectrum of microstructure is not the same as the average pulse spectrum. For two pulsars, PSR 1919+21 and PSR 2016+28 (Cordes 1978 b) the modulation index, m_μ, decreases as f^{-1} or faster where f is the observing frequency. For both pulsars the microstructure is associated with drifting subpulses and at high frequencies the emission is entirely composed of a non-drifting, amorphous, noise-like

PSR	PERIOD P (s)	MICROSTRUCTURE BREAK Δt_μ (μs)		MICROSTRUCTURE PERIOD P_μ (μs)	
		Cordes 1978b	Kardashev *et al* 1978		Cordes 1978b
0950+08	0.253	180 ± 30 (below 100)	200	2100	500 ± 100 (150-1100)
2020+28	0.343	110 ± 30			
1944+17	0.441	300 ± 50			950 ± 250
0823+26	0.531	550 ± 150			
2016+28	0.558	280 ± 50			900 ± 200 (300-1400)
1749-28	0.563		[250 ± 50]†	[105 ± 20]†	
0329+54	0.715		200	340	
1508+55	0.740		80	430	
1133+16	1.188	580 ± 150 (below 100)	600,560*	160, 1300	1100 ± 300 [185]**
0834+06	1.274	1050 ± 200	none	160-700	
0809+74	1.292		100	200-700	
1919+21	1.337	1300 ± 200	90	400	
0525+21	3.75	3000 ± 700			

TABLE 1 : Summary of published microstructure observations in order of pulsar period. Observed ranges are in parentheses.

† described in Chapter 9 of this thesis.

* Component 1 has 600 μs, Component 2 has 560 μs.

**Ferguson and Seiradakis (1978).

component. However microstructure seems to make up much of the emission of PSR 1133+16 at 2700 MHz.

The position as regards variation of Δt_μ with frequency is not clear. For PSR 2016+28 it decreases with increasing frequency according to the law $f^{-\alpha}$ where $\alpha \simeq 0.3$ (Cordes 1978 b). As mentioned previously the situation for PSR 1133+16 is confused; Cordes concluded that $\alpha = 0$ and $\alpha = 0.33$ were both compatible with the data. For PSR 0950+08, $\alpha \simeq 0$.

What is the relationship of micropulses to subpulses? The polarization observations suggest that they are similar. In both cases transitions between orthogonal modes tend to occur on the boundaries. However the observations of PSR 1133+16 and 2016+28 suggest that there are separate components to the emission: microstructure and an amorphous component. Many subpulses are amorphous. Ferguson and Seiradakis (1978) argued that a subpulse is not constructed of micropulses although the occurrence of the two is highly correlated. Both may be produced by angular beams. Taylor *et al.* (1975) found that subpulse widths are positively correlated with pulsar period, P, so that they could be interpreted as angular beams with an average width $\sim 5^\circ$ (i.e. duration $\sim$ P/70). Cordes (1978 b) described a similar relationship for microstructure, namely, Δt_μ = P/1000. The scatter of the data considered by Cordes about this regression line was very large. When the observations of Kardashev *et al.* (1978) and Ferguson and Seiradakis (1978) are included the correlation is greatly reduced. This can be seen by inspection of Table 1 where a summary of the published observations is arranged in order of P.

The AMN model appears to be an adequate description of the observations. The relationship to the model of observations of very narrow and intense micropulses is as follows. The model provides a framework to test the statistical significance of the observed narrow features. The model does not preclude narrow features in the modulation but such features may be

difficult phenomena (e.g. when compared to the average timescale) for a theory of the modulation to explain. However it may be concluded that a particular feature is incompatible with the AMN model, that is, it cannot be interpreted as Gaussian noise. One alternative is that the emission may consist of low density shot noise with each shot, or broad-band impulse, being very strong and having a duration or separation just less than the best time resolution. Therefore it is important to increase the time resolution in the hope of resolving these shots. Another possibility is that there could be some underlying complicated, non-random structure to the "noise". This could be the case even if the AMN model is an adequate description of the observations made using a single receiving channel, but the dynamic spectrum may show non-random structure. Both aims, improvement in time resolution and observation of the dynamic spectrum, are served by the technique described in later chapters of this thesis.

To conclude this Section I will review some implications about radiation energies which have been drawn from the microstructure observations. Hankins (1972)discussed an intensity spike of 35000 Jy, 8 μs wide, assuming it was not a statistical fluctuation. The argument was repeated by Hankins and Boriakoff (1978) for a 1000 Jy, 0.8 μs spike. Equating 8 μs (or 0.8 μs) to the time for light to travel across the emission region leads to a spatial scale of 2.5 km (250 m) and hence an equivalent brightness temperature of 10^{30} K (3×10^{31} K). Alternatively the spikes can be interpreted as statistical fluctuations within the AMN model. Taking 40 μs (4 μs) as a rough value for the micropulse duration leads to a temperature over 10^{28} K (10^{29} K). In either case the emission process must be coherent.

A number of workers have used the short duration microstructure observations in a continuing debate as to the feasibility of light cylinder

models. In these models the radiation is produced by particles at the outer edge of the co-rotating pulsar magnetosphere where the tangential speed is near that of light. The argument was originally presented by Manchester *et al.* (1973) and concerns the relative values of the energy density, U_B, of the magnetic field and the energy density, U_P, of the plasma. Because integrated pulse profiles are stable U_B must exceed U_P. The energy density, U_R, of the radiation is used as a lower bound for U_P so that the requirement is $U_B > U_R$ in the emission region. The duration of a micropulse is used to estimate the dimensions of the emission region. If βc is the co-rotation speed with $\gamma = (1-\beta^2)^{-\frac{1}{2}}$ the corresponding Lorentz factor, then

$$U_R = \pi^{-1}\ \gamma^4\ (1-\beta)^4\ Sd^2\ \Delta t\ \Delta\nu\ (c\Delta t)^{-3}$$

where S is the observed flux density, Δt its timescale, $\Delta\nu$ its bandwidth, c the speed of light and d the distance to the pulsar. Also

$$U_B = B^2/8\pi \simeq 2\pi^3\ I\ \dot{P}/P^5c^3\beta^6$$

where B is the magnetic field strength in the emission region, I the moment of inertia of the neutron star, P its rotation period (i.e. the pulsar period) and $\dot{P}$ the period derivative with time. (Here B is estimated from the observed pulsar slow-down).

Using the observations of PSR 0950+08 by Hankins (1972), $S \sim 35000$ Jy and $\Delta t \sim 10\mu s$ and assuming a bandwidth $\Delta\nu = 10$ MHz, Manchester and co-workers calculated that $U_B > U_R$ if $\gamma > 5$. Zheleznyakov and Shaposhnikov (1974) have attacked the basis of the argument, while Ferguson (1976 b) has argued that there is an implicit assumption of high polarization and has rejected the particular values of S, Δt and $\Delta\nu$ used. Certainly $S < 20000$ Jy and $\Delta t > 30$ μs are consistent with the observations when interpreted by the AMN model and these lead to values of $\gamma \sim 1.5$ being allowed. However Rickett *et al.* (1975) suggested that $\Delta\nu$ should be 200 MHz because of the 111/318 MHz correlation. From his observations at 2700 MHz Bartel (1978) chose $S = 500$ Jy, $\Delta t = 6$ μs and $\Delta\nu = 20$ MHz leading to $\gamma > 20$ (or $\gamma < 1.01$).

Ferguson and Seiradakis (1978) took $S = 200$ Jy, $\Delta t = 200$ μs and $\Delta\nu = 100$ MHz leading to $\gamma > 8$ (or $\gamma < 1.04$), that is, $\beta > 0.99$ (or < 0.27) which is similar to the β they obtained from a model fitted to their polarization data. This means that the outer part of the pulsar magnetosphere is supposed to co-rotate at a speed within 1% of that of light.

1.4 THEORY OF PULSAR RADIO EMISSION

1.4.1 The Magnetosphere

Much of the material in this Section is derived from the review of Manchester and Taylor (1977, p.169). I have concentrated on the Ruderman-Sutherland model because it has generally been accepted as the most comprehensive pulsar theory to date and forms the basis of some theories of microstructure.

A pulsar is believed to be a neutron star with a radius $\sim 10^6$ cm and a magnetic field with a strength at the surface of $\sim 10^{12}$ gauss. It rotates with a period, P, equal to the observed pulse period. Observed values of P range from 33 ms to about 4 s with a median of 0.65 s. The observed pulse is produced by a "lighthouse beam" of emission sweeping through the line of sight. The beam and other properties of the emission are determined by the magnetic field which co-rotates with the star. There must be an asymmetry to define the azimuth of the beam. Therefore the magnetic axis cannot be aligned with the rotation axis, however for theoretical tractibility an aligned, that is, axi-symmetric model is often assumed.

Although the surface gravitational field is $\sim 10^8$ times the terrestrial field, matter near the neutron star is dominated by magnetic and electrostatic forces. Goldreich and Julian (1969) showed that (for the axi-symmetric case) the spinning pulsar acts as a homopolar generator. Electric fields parallel to the magnetic field with magnitudes $\sim 10^{11}$ volt cm^{-1} would remove charges from the stellar surface and so the magnetosphere must acquire a space charge. Furthermore there must be a larger density of neutral plasma (Okamoto 1974). The particles are constrained to move along the magnetic field lines with essentially zero pitch angles since the lifetime against synchrotron radiation is very short due to the high strength of the field. Further they are constrained to co-rotate with

the star, but clearly this co-rotation cannot extend beyond the "light cylinder" where the tangential speed would be that of light.

The radius, R_{LC}, of the light cylinder is

$$R_{LC} = c/\Omega \simeq (5 \times 10^9 \text{ cm s}^{-1})\ P$$

where $\Omega = 2\pi/P$ and c is the speed of light.

In the class of models for pulsar emission known as "light cylinder models" radiating particles are assumed to co-rotate at relativistic speeds. Whether or not the radiation is isotropic in the co-rotating frame it is beamed into the forward direction, tangential to the cylinder (see Figure 1.4.1A). The angular width in longitude is about $1/\gamma_\phi^3$ where γ_ϕ is the co-rotation Lorentz factor. Hence it is important to know whether co-rotation can extend very close to the light cylinder, where for the slower pulsars the magnetic field strength is only a few Gauss. Henricksen and Rayburn (1974) argued that the limit occurs where the co-rotation velocity equals the Alvén velocity rather than c which restricts γ_ϕ to a (model dependent) limit $\lesssim 2$. However according to Hinata and Jackson (1974) co-rotation ceases at only about 5% of R_{LC}.

The magnetic field configuration is assumed to be dipolar as a first approximation but two mechanisms cause, possibly severe, deviations from this even in the axi-symmetric case. These are particle inertia and the presence of currents flowing along the "open" field lines. The latter are those magnetic field lines originating near the poles that penetrate the light cylinder. Due to the presence of the currents the field acquires toroidal components near the light cylinder. Assuming co-rotation extends that far, cusps form at the light cylinder in both the aligned and orthogonal cases (Henricksen and Norton 1975, Michel 1973). The cusps could provide centres where particles collect and hence allow a light cylinder mechanism to produce pulsed rather than continuous emission.

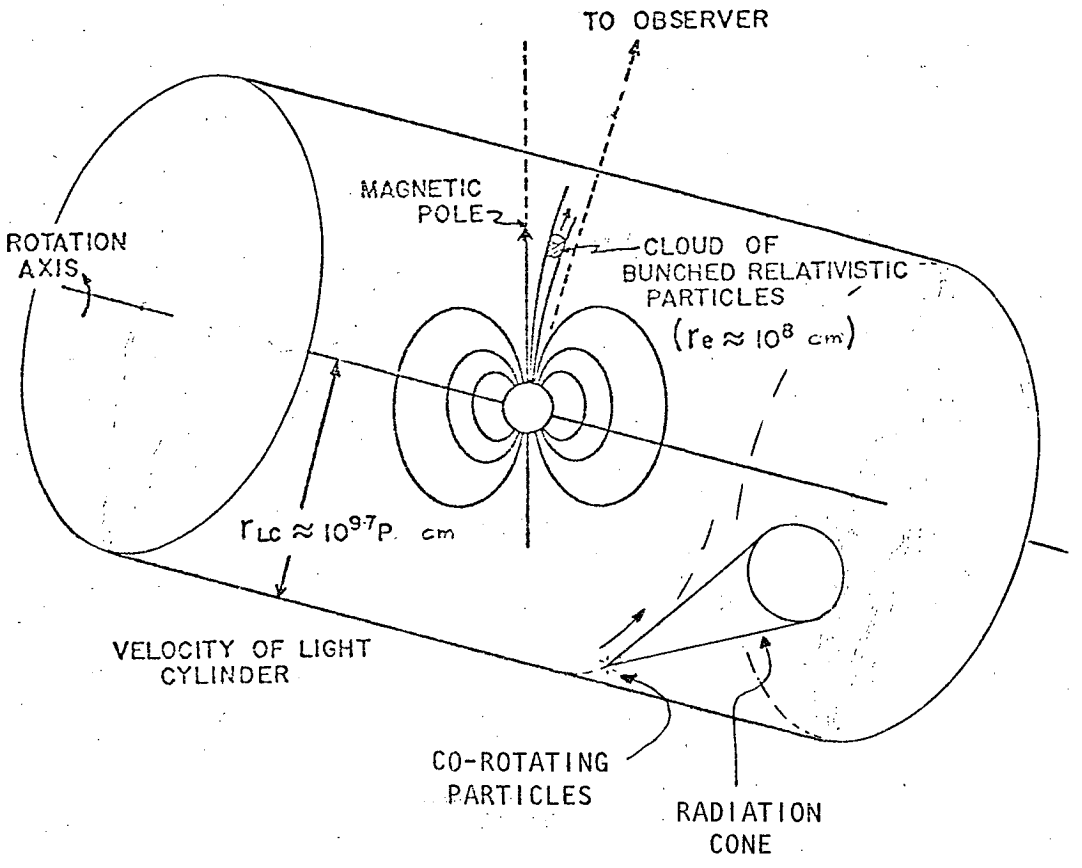

FIGURE 1.4.1A Illustration of polar cap and light cylinder models of pulsar emission. Relativistic particles streaming outward from the polar cap of the neutron star emit radiation in a beam roughly parallel to the magnetic axis.

Alternatively, particles near the light cylinder, co-rotating at relativistic speeds, emit radiation beamed tangential to the light cylinder. (Figure derived from Cordes 1978 .)

Goldreich and Julian (1969) and others assumed that charged particles were free to leave the stellar surface to form currents in the magnetosphere. However the model of Ruderman and Sutherland (1975, hereafter RS) is built on the result (Ruderman 1972, Chen *et al.* 1974) that positive ions (mostly of ^{56}Fe) are tightly bound to the surface and not free to form an outward-flowing positive current. (It is assumed that electrons cannot flow inward along the open field lines from beyond the light cylinder.) Therefore in pulsars in which the magnetic moment, $\underline{M}$, and the rotational angular velocity, $\underline{\Omega}$, are antiparallel (or in the non-aligned case $\underline{M}.\underline{\Omega} < 0$) a particle free gap develops above each polar cap. The cap is defined by the feed of the open field lines and is, in the dipolar axi-symmetric case, a circle centred on the magnetic pole of radius $\sim 10^4$ cm. The height, h, of the gap is $\sim 5 \times 10^3$ cm for P $\sim$ 1 s. The conductivity along the open field lines is very high so that most of the homopolar potential difference $\sim 10^{12}$ volt is assumed to appear across the gap. Thus within the gap the electric field, E, has a component parallel to the magnetic field, B, although $\underline{E}.\underline{B} = 0$ elsewhere in the magnetosphere (except in "outer gaps", Cheng *et al.* 1976).

When the potential difference across the gap above the polar cap grows to about 10^{12} volt an electron-positron avalanche forms a spark. The mechanism of this avalanche was described by Sturrock (1971) and depends on the magnetic field being very strong and curved. It is as follows. A gamma ray crossing a field line produces an electron-positron pair. These two particles are accelerated by the electric field to relativistic speeds, the electron downward and the positron upward along the particular magnetic field line. On average, each emits a curvature radiation photon, in fact, a gamma ray, before it leaves the gap. These gamma rays produce further pairs as they cross neighbouring field lines. (The magnetic field is assumed to have multipole components so that the curvature of the field

lines is sufficient, i.e. $\sim 10^{-6}$ cm^{-1} which is much greater than for a dipole field.) The result is that in the region above the gap there is an upward primary current of positrons with Lorentz factor $\gamma \sim 10^6$. Although no further acceleration occurs above the gap ($\underline{E}.\underline{B} = 0$), pair production continues and so there is a much denser particle flux of electrons and positrons with $\gamma \sim 10^3$. Thus the neutral plasma that fills the open flux tubes consists of electrons and positrons and it is flowing outward at relativistic speeds. It is this plasma that is involved in a number of theories of pulsar radio emission.

The phenomenon of drifting subpulses is explained by the RS model in terms of the "precession" of sparks around the polar cap. Because $\underline{E}.\underline{B} \neq 0$ in the gap the open flux tubes are not constrained to co-rotate exactly with the rest of the magnetosphere; there is a relative rotation and the sparks are presumed to share in this rotation. The rate of this is comparable to observed rates of drift in drifting subpulses. The direction of drift is also predicted (Ruderman 1976). However the observed variation in the rate of drift from one group of a few pulses to the next is not explained (Oster and Sieber 1976). In fact the picture of the gap and the sparks given by RS is probably over-simplified. Hillebrandt and Müller (1976) and Flowers *et al.* (1977) have calculated values for the binding energy of the ions to the surface of the star that are much less than those of RS. Therefore ions are likely to take part in the discharge which is then space charge limited and the particle free gap does not form. In addition Cheng and Ruderman (1977 b) have described an alternative to the Sturrock mechanism in which gamma rays form pairs in the Coulomb fields of the moving ions (or, less effectively, of the electrons and positrons). This mechanism has the advantage that large magnetic curvatures are not required in the polar cap region and avoids

the "embarrassment" of the Sturrock mechanism that the spark migrates rapidly to the magnetic pole (Hinata 1979 calculated a speed of 2000 cm s^{-1}). There are also other effects to consider such as surface heating and the "splashing" of photons from the surface.

1.4.2 Emission Mechanisms

There are two main classes of theories of pulsar radio emission, namely polar cap models and light cylinder models. I will deal first and briefly with the latter.

Proposals for the basic light cylinder emission mechanism, that is, the process occurring in the co-rotating frame, include synchrotron emission (Eastlund 1968), cyclotron emission (Smith 1970) and a maser mechanism.† According to Ferguson (1977) the biggest problem with light cylinder models is that they are underdeveloped. However the co-rotation problem is clearly important. The radiation density argument was described in Section 1.3.4 and this requires co-rotation to extend to within 1% of the light cylinder. Reviews of this and other arguments for and against light cylinder models have been made by Ginzburg and Zheleznyakov (1975), Ferguson (1976 b), Manchester and Taylor (1977, p. 214) and Smith (1977,p199).

In polar cap models the radio emission is directed in narrow beams tangential to the open field lines from regions near the polar caps. Higher frequencies are then predicted to be emitted nearer to the neutron star where the plasma density is higher. For example, in the RS model the emitted frequency varies as $r^{-3/2}$ where r is the radial distance from the star. From the absence of aberration and retardation effects in the observed arrival times of the centroids of integrated profiles at different frequencies, Cordes (1978 a) has argued that all the radio emission comes from radii which are less than a few percent of the light cylinder radius, that is, $r \lesssim 10^9$ cm.

† (Ginzburg *et al* 1969)

The emission must be coherent to give the large observed brightness temperatures (up to 10^{30} K). There are two main groups of coherent mechanisms. First radio waves could be generated by the maser effect given a suitable non-equilibrium energy distribution in the plasma. The radiation would be highly beamed (Ginzburg and Zheleznyakov 1975). This mechanism is more likely to occur in the polar or "outer" gaps where $\underline{E}.\underline{B} \neq 0$ than in the rest of the magnetosphere (Buschauer and Benford 1976). The second group of mechanisms are the "antenna mechanisms" in which the electromagnetic waves are radiated from particle bunches or sheets of charge.
The bunching is often attributed to plasma waves which grow by the two-stream instability. In the theories of Hardee and Rose (1976) and Kawamura and Suzuki (1977) transverse waves are generated at frequencies below the plasma (and gyro) frequency. The two streams in these cases may be the primary positrons ($\gamma \sim 10^6$) and the secondary electron-positron plasma ($\gamma \sim 10^3$) of the RS model. Cheng and Ruderman (1977 a) have shown that the curvature of the open field lines produces charge separation of the secondary plasma. In the centre-of-momentum frame the electrons and positrons separate at velocities $\pm c$. This relative streaming causes rapid growth (i.e. much more rapid than the original RS plasma mode) of unstable electrostatic plasma waves. Then the moving "antenna" of bunched charges generates coherent curvature radiation.

Curvature (i.e. synchrotron) radiation from single particles is broad-band with a high frequency cut-off at the critical frequency, f_c, given by (Jackson 1975, p. 676)

$$2\pi f_c = \omega_c = 3\gamma^3 c/\rho_c \qquad (1)$$

where ρ_c is the radius of curvature of the particle trajectory and the particles have energy $\gamma\, mc^2$, with m the particle rest-mass. However when the curvature radiation is coherent it is narrow-band centred on the plasma frequency, ω_p. The general theory of coherent curvature radiation has

been presented by Buschauer and Benford (1976). They found that two cases must be considered depending on the phase velocity of the bunching plasma wave. If the wave is non-convective, that is, stationary with respect to the background plasma, then the coherent spectrum factors into a product of the single particle spectrum and a resonance function centred on ω_p whose width is inversely proportional to the length of the "antenna". However when the wave is convective, that is, with phase velocity near c, there is no simple functional form. In particular there is no strong cut-off at f_c. In both cases the radiation is beamed forward along the field line into a cone of width $\sim 1/\gamma$.

Assuming that a cut-off would occur at f_c, RS used the geometry of the magnetic field, assumed dipolar, and the law $\omega_p \propto r^{-3/2}$ to delimit the regions in which emission could occur. This leads to the overall pulsar beam being a hollow cone centred on the magnetic pole. This gives a qualitative explanation for the various single and double pulse profiles of pulsars. Oster and Sieber (1976,1978) have found that the frequency dependence of the cone width inferred from observations of two pulsars does not agree with the RS model in that the width is constant above some frequency which under the model would be interpreted as the cut-off frequency. However the spectrum is not cut off there although there is a break in the spectrum. I suggest that this means that if coherent curvature radiation is indeed the emission mechanism then the plasma waves are convective. [It is interesting that Benford and Buschauer (1977) and Cheng and Ruderman (1977 a) seem to disagree as to whether the plasma wave of the RS model (1975)is convective or not.]

The overall picture of radio emission in the RS model is one of incoherent addition of coherent curvature radiation from many charge bunches. The involvement of many bunches explains the noise-like nature of the pulsar signals and, in particular, the AMN description of the data.

The bunches are produced by a smaller number of sparks in the polar gap. The inclination of the magnetic axis and the geometry of the polar field lines, which may be very complicated near the surface of the star lead to an angular beam which produces the integrated pulse profile.
The position angle of the radiation, when it is linearly or elliptically polarized, is determined by the orientation of the field lines in the emission region to the line of sight. A subpulse is associated with a cluster of radiating bunches which has a significant spread in azimuth around the magnetic axis. [This is suggested by the fact that the position angle within a subpulse tends to vary in a similar way to the way it varies in the integrated profile (Cordes 1978 b).] Each subpulse is caused by a cluster of sparks (or perhaps a single spark) and drifting subpulses are caused by "precession" of the sparks around the polar cap. Finally each micropulse may be associated with a single spark or with temporal modulation of the spark or the emission mechanism.

Oster and Sieber (1978) have made a detailed comparison of this polar cap sparking model and their analyses of observations of four pulsars (PSRs 0329+54, 1237+25, 1919+21 and 2016+28). Their method of analysis (as applied to this model) is built on the assumption that pulses are composed of subpulses and that these are associated with single drifting sparks. They found that

(i) the drift rate varies, but

(ii) at any one time all the sparks around the polar cap drift at about the same rate,

(iii) there are concentric hollow cones (which would explain the notches in the integrated profile of PSR 1919+21),

(iv) the model is not yet detailed enough to explain the observations, and (v) more single pulse observations at different frequencies (preferably acquired simultaneously) are needed.

I close this section by mentioning a problem area in theories of pulsar radio emission. This is the origin of the transition between orthogonal modes of polarization. A related problem is the occurrence of high percentages of circular polarization. In the centre of the emission cone of curvature radiation the polarization is linear, parallel to the plane of the curved field line. On either side of the plane the radiation is elliptically polarized in opposite senses but the incoherent addition of the radiation from many bunches in a cluster is likely to be linearly polarized or unpolarized (Kirk and ter Haar 1978). Near the emission region (at radius $\sim 10^9$ cm for the RS model) the gyrofrequency is much greater than the observing frequency, so that the propagation is quasi-transverse and the natural modes are linear, however near the light cylinder and in the interstellar medium the natural modes are circular. Thus there is some scope for changes in the polarization during propagation out of the pulsar magnetosphere. Recently Cheng and Ruderman (1979) have described a propagation effect called "adiabatic walking" which explains the occurrence of high percentage polarization and of orthogonal modes. Other theories of the polarization have been reviewed by Manchester and Taylor (1977, p. 224) and Cordes (1978 c).

1.4.3 The Origin of Microstructure

In this Section I shall discuss some theories for the mechanism producing pulsar microstructure; all assume a polar cap model and propagation effects are ignored. It is assumed that the AMN model is an accurate description of the observations so that there are two aspects to consider. These are the Gaussian noise process and its modulation. In the preceding section it was mentioned that incoherent addition of the coherent radiation from a number of charge bunches (i.e. "antennas") leads to a noise-like signal. This follows from the Central Limit Theorem.

Providing there is a sufficient number of radiating units then the sum will be Gaussian noise no matter the details of the fundamental emission. The density of radiating units must be high in some sense. For example, there could be a high angular density with many units radiating into the line of sight, or there could be a high temporal density as in shot noise (Cordes 1976 b). A similar choice between angular and temporal variation pertains to the modulation process of the AMN model. Temporal modulation could have some advantage in explaining the correlation of microstructure between different frequencies if all frequencies were emitted in the same region, that is, not in the RS framework. However the only developed theories involve angular modulation.

Temporal modulation was suggested by Ruderman and Sutherland (1975) to explain microstructure. They calculated the spark duration, Δt, to be about 30 h/c where h is the gap height, that is, $\Delta t \sim 10$ μs. However Hinata (1979) in a theoretical investigation of spark structure could not calculate a duration for lack of an effective quenching mechanism. He did calculate that each spark took the form of a thin ribbon with a sub-millimetre cross-section and suggested that the dimensions might be related to microstructure.

Cordes (1978 b) suggested shot noise with either temporal or angular modulation. He considered curvature radiation from relativistic particles travelling along the magnetic field lines from the polar cap towards the observer. With curvature/synchrotron radiation the observer sees a series of narrow pulses and Cordes (incorrectly) took these to have duration $\delta t \sim 1/f_c$ where f_c is the critical frequency given by equation (1). For $\gamma = 10^3$ and $\rho_c = 10^{10}$ cm this leads to $f_c \sim 3$ GHz and $\delta t \sim 1$ ns. Cordes suggested that the modulation timescale might be the light travel time for the radial distance over which the radiation occurs. Alternatively the modulation could be due to the sweeping of the curvature emission

cone through the line of sight. The angular width, $\Delta\phi$, of this cone perpendicular to the plane of the field line is given by

$$\Delta\phi \sim \gamma^{-1} (f/f_c)^{-1/3} \quad (2)$$

where f is the observing frequency. Then assuming the line of sight does not pass very near the magnetic pole (with some rearrangement of Cordes' equation)

$$\Delta t_\mu \simeq \Delta\phi (P/2\pi) \simeq (P/2000) (\rho_c f)^{-1/3} \quad (3)$$

With $\rho_c = 10^{10}$ cm this gives a very good fit to Cordes' observations of the microstructure of PSR 2016+28 between 111 MHz and 606 MHz. However this treatment is invalid - since the radiation must be coherent the single particle theory of curvature radiation does not necessarily apply.

Benford (1977) and Kirk and colleagues (Elsässer and Kirk 1976, Kirk and ter Haar 1978) have published models of microstructure involving coherent curvature radiation. In both cases the modulation is produced by an angular beam and the "noise" comes from incoherent addition of the radiation from separate charge bunches within a cluster. Part of the difference between the models is related to the transverse extent of a cluster.

The general situation in the Benford model is shown in Figure 1.4.3A. Each bunch radiates at the local plasma frequency, ω_p, and therefore, using the RS model of the magnetosphere, the radiation received at a given frequency comes from the emission surface whose edges are defined by the hollow cone. Since the beamwidth, $\Delta\phi$, of the emission from a single bunch, or "filament", is small the "observed" part of this surface at an instant is small. Its boundary is defined by the field lines diverging by $\Delta\phi$ from the line of sight. Benford implicitly assumes that $\Delta\phi$ is less than the angle θ (see Figure 1.4.3A(ii)) which is approximately the angle subtended at the rotation axis by the intersection of the flux tube

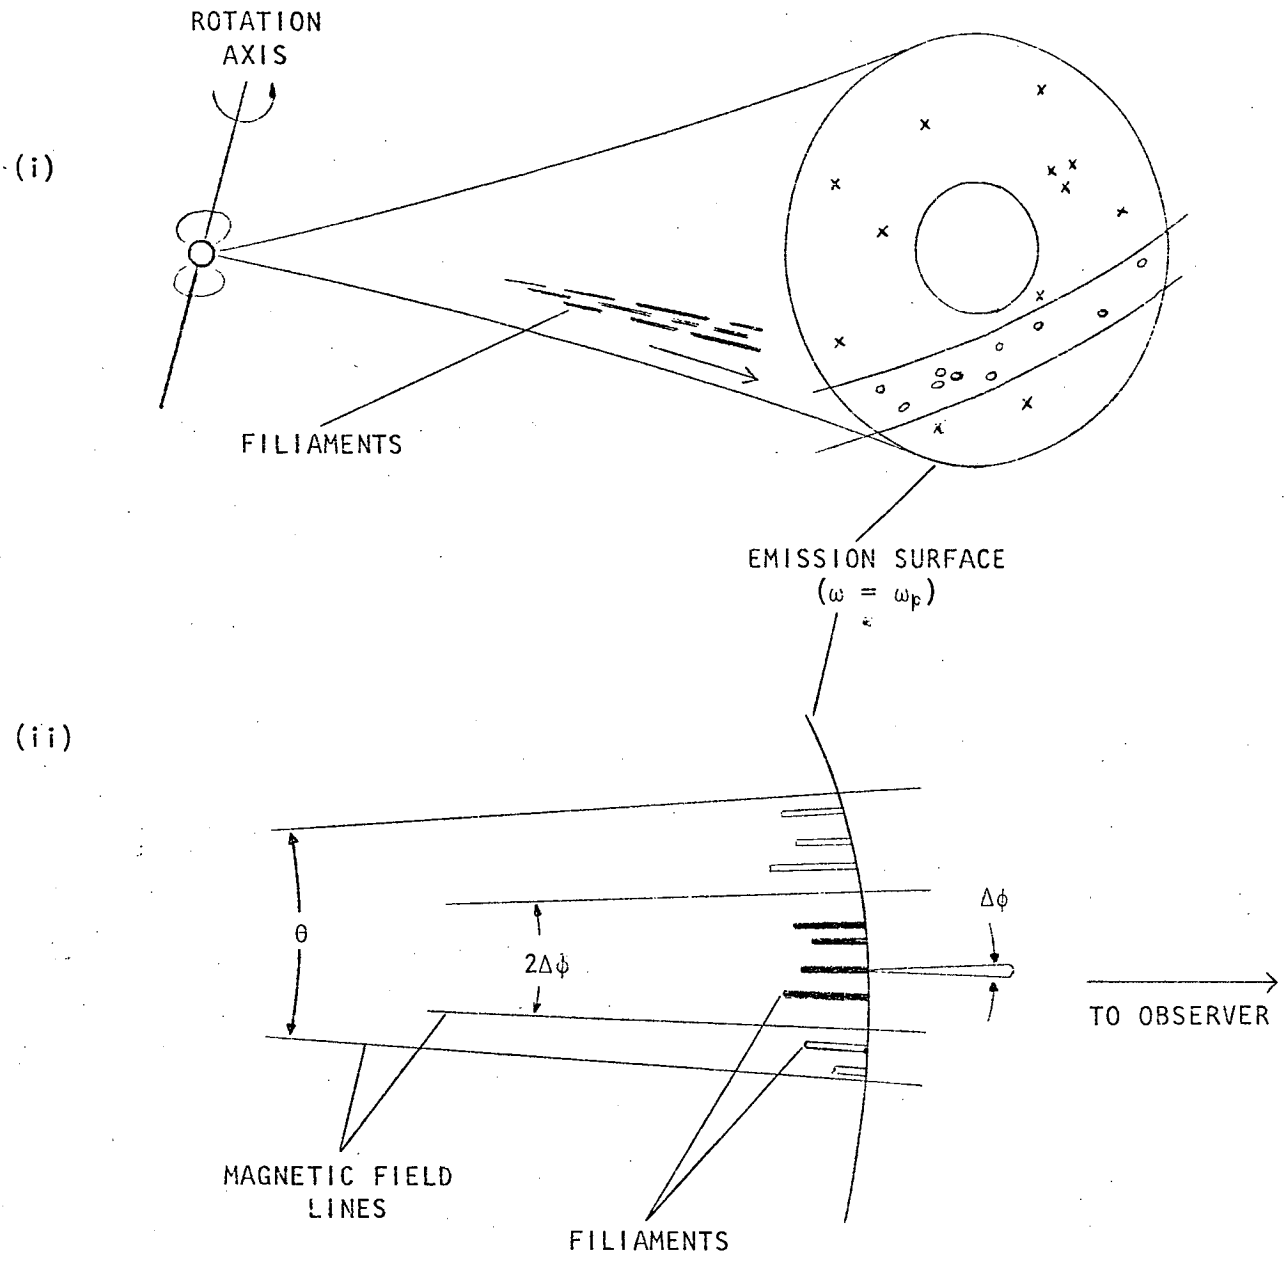

FIGURE 1.4.3A

Schematic of Benford's (1977) model of microstructure. (i) Each circle and cross represents a cluster of filiaments. Emission from the 'circle' clusters is visible as micropulses as the hollow-cone rotates past the observer. (ii) View from the rotation axis showing one 'circle' cluster. All the filiaments are radiating but at one time only those shaded are visible. Further details are in the text.

containing the cluster with the emission surface. For coherent curvature radiation $\Delta\phi < 1/\gamma$ and with the RS model, $\gamma \sim 10^3$. The angular extent of the cluster is rotated through the line of sight so that the duration, Δt, of a micropulse is

$$\Delta t = \theta P/2\pi$$

Hence for PSR 0950+08 at 111 MHz Benford estimated that $\Delta t_\mu = 175$ μs would correspond to a lateral width for the typical cluster of $\sim 10^7$ cm at 2×10^9 cm from the star. (Note that here $\theta > 4\Delta\phi$.)

Benford made no estimate of the number of filaments concentrated in the flux tube of diameter $\sim 10^7$ cm nor of the lateral extent of each, but noted that, assuming a dipole field, the foot of the tube would have a diameter $\sim$ 1m. The length of each filament was estimated to be $< 3\times10^6$ cm. This was from the fine structure in the dynamic spectrum published by Rickett *et al.* (1975). (The pulse was shown in Figure 1.3.1B). However in accordance with the AMN model this fine structure consists of random fluctuations whose width in frequency is set by the resolution of the (software) spectrum analyzer. Therefore the length of the filaments is an open question. Benford did note that the clusters could be much shorter than 10^9 cm which is the radial separation of the 111 MHz emission surface from the 318 MHz emission surface. This is because the filaments move at relativistic speeds and because the beamwidth $\Delta\phi$ is so small this must be almost directly towards the observer. Therefore the signals at the two frequencies appear to be emitted simultaneously (± a few microseconds).

In the emission theory of Kirk and coworkers the cluster angular extent, θ, is implicitly assumed to be less than the beamwidth, $\Delta\phi$, and in fact it is not clear whether they envision any transverse extent to the cluster. Hence the microstructure more closely reflects the basic emission beamwidth than in Benford's model. However there are complications

and the simple relationship of equation (3) only applies when the line of sight scans across the plane of the field line at right angles. The light cylinder radius, R_{LC}, is taken as a typical value for ρ_c (accurate for the axi-symmetric dipole model near the star and about halfway from the pole to the last open field line) and hence, from (1) and (3),

$$\Delta t_\mu = \Delta t_\perp = (P/2\pi)(fP/3)^{-1/3} \qquad (4)$$

In the converse case the line of sight remains in the plane of the field line as the pulsar rotates (a y-scan in Figure 1.4.3B(ii)). Kirk and ter Haar give

$$\Delta t_\mu = \Delta t_{||} = \rho_c/\gamma c = (P/2\pi)\gamma^{-1} \qquad (5)$$

where $\rho_c = R_{LC}$ as before. Values of γ were derived from observed pulsar spectra using a theory for coherent curvature radiation (Elsässer and Kirk 1976) in which the spectrum factors into the single particle form and a resonance function centred on ω_p (the wave is convective; cf Buschauer and Benford 1976). The derived values of γ for 9 pulsars lie in the range 350 to 800. The choice of whether to use (4) or (5) depended on the observed polarization. It was argued that the fractional linear polarization should be high for the parallel scan and low for the perpendicular scan (Figure 1.4.3B) Comparisons between the predictions and observations are shown in Table 2. In general they agree to within a factor of 2 and the fit for PSR 2016+28 at four frequencies is very good. (This was found by Cordes; equations (3) and (4) are very similar.) However the classification of this pulsar and of PSR 0950+08 as showing zero linear polarization is open to question. Furthermore when applied to the hollow-cone model this theory predicts that pulse profiles should have zero linear polarization in their centres. This is because the orientation of the planes of the field lines relative to the locus of the line of sight varies along that locus [Figure 1.4.3B(iii)]. One counter example is

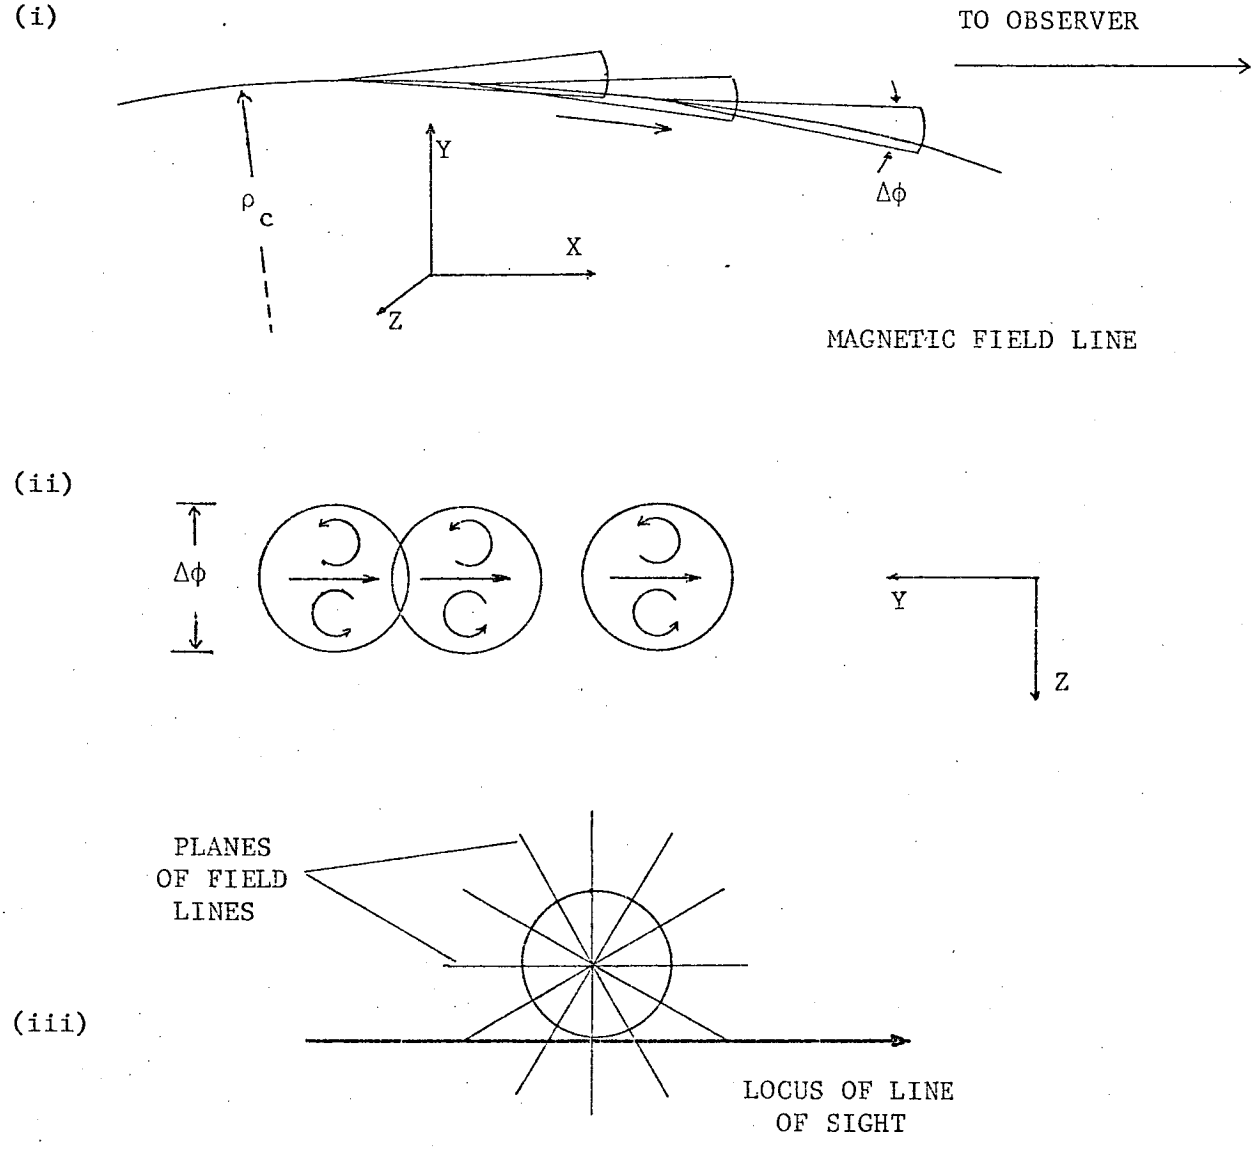

FIGURE 1.4.3B Geometry of the microstructure model of Kirk and co-workers.

(i) Particle bunches travelling along a field line with radius of curvature ρ_c emit curvature radiation into a cone of width $\Delta\phi$.

(ii) The radiation from each bunch is linearly polarized in the centre of the emission cone, i.e. in the plane of the field line, and elliptically polarized on either side.

(iii) View down to the magnetic pole showing that the line of sight makes a z-cut at the centre of the pulse.

PSR 2020+28 which at 430 MHz shows 65% linear polarization in the saddle region of its double profile (Cordes *et al.* 1978). (Of course these arguments ignore possible propagation effects as do those of Kirk and his colleagues.) Similarly the microstructure timescale should vary across the profile and this has not been observed (Cordes and Hankins 1977, Hankins and Boriakoff 1978). Kirk and ter Haar claim that their theory makes a prediction (unspecified) about the ACF of the power spectrum of microstructure (see Chapter 9). This is apparently related to the fact that the emission is narrow-band, a feature that their theory shares with those of RS and Buschauer and Benford (1976). However in this case the spectral width is $\omega_p/2\pi\gamma$ rather than $\sim c/s$ where s is the length of a filament (Buschauer and Benford 1976). Using Kirk and ter Haar's figures the spectral width for PSR 0950+08 at 408 MHz is about 1 MHz. Could this be observed or more generally, in the RS framework, is it likely that the narrow-band nature of coherent curvature radiation would be observed?

No mechanisms have been proposed for particle bunches suddenly to start radiating at a particular radius and there are apparently many bunches in a cluster. Therefore it seems that a cluster should radiate energy at a smoothly varying rate as it travels outward along a flux tube. Therefore the narrow-band nature of at least the microstructure emission should not be observed directly although some changes of the statistics with frequency are likely. On the other hand Udal'tsov and Zlobin (1974, 1976) have reported observations of sporadic narrow-band emission from four pulsars.

Finally, very few theories have been advanced to explain periodicity in microstructure. In discussing the 0.9 ms quasi-period of PSR 2016+28, Boriakoff (1976) explored the possibility of non-radial vibrations of the neutron star coupled to the emission cones by angular displacements of the

magnetic field. Benford (1977) suggested a spatial effect, namely periodic spacing of the sparks around the polar cap. He also considered a temporal effect, namely some relaxation process for the gap. Clearly we require more detailed knowledge of the gap processes.

PSR	observed		predicted			
	Δt_μ	α	$\Delta t_\perp$	$\alpha_\perp$	$\Delta t_\parallel$	$\alpha_\parallel$
0031-07					300	0
0329+54	200		380	.33		
0950+08	175	$\sim$0	190	.33	(80	0)
1133+16	575	$\neq$0	(575	.33)	235	0
1919+21	1300		575	.33	(400	0)
1929+10					65	0
2016+28	280	$\sim$0.3	325	.33		

TABLE 2: Comparison of observations (from Table 1) of microstructure width Δt_μ (μs) with the predictions of Kirk and ter Haar (1978). α is the spectral index, i.e. $\Delta t_\mu (f) = kf^{-\alpha}$. All of the values of Δt_μ in the table are for 111 MHz. Values in parentheses are not predictions of Kirk and ter Haar but are simply calculated from the alternative formulae.

CHAPTER 2

THE RECORDING SYSTEM

2.1 INTRODUCTION

The following is an outline of the recording method. A 2 MHz band of radio-frequency data is heterodyned to the frequency range 0 - 2 MHz and recorded with a domestic videotape recorder (VTR). Subsequently the videotape is replayed and segments of the data are digitized for analysis with dispersion removal by computer. To assist in the digitization a reference tone and synchronization (sync) pulses are added to the data during recording. Note that all three signals, the noise-like data, the reference tone and the sync pulses, share the same recording channel. There is some multiplexing in frequency and time to segregate the three signals as described in the following two paragraphs.

Since sampling takes place during digitization at the zero-crossings of the reference tone, i.e. at twice the reference frequency, the reference frequency is the Nyquist frequency of the system. This is about 2.25 MHz so that the maximum useful bandwidth of the system is less than 2 MHz depending on the sharpness of the filters in the system. The zero-crossings of the reference tone will be referred to as "clock ticks", the clock in question being a notional observatory clock. In order that the clock ticks can be recovered during replay as accurately as possible they are recorded as cleanly as possible. That is a notch filter tuned to the reference frequency is inserted in the data signal path to augment the attenuation provided by a 2 MHz low-pass filter in that part of the spectrum. A sketch of the recorded spectrum is shown in Figure 2.1B(a).

The helical scan VTR divides the data naturally into 20 ms segments which I will refer to as "stripes" according to the diagonal locus of each recording head across the tape. A sync pulse is recorded at each end of each stripe. The first indicates the zeroth clock tick of the stripe. The second is used for checking that on any one replay of the stripe (some hundreds are needed for complete digitization) the clock ticks are correctly

recovered. So that these two pulses are as clean as possible on replay, the data signal is switched off i.e. "squelched" for their duration. This is depicted in Figure 2.1B(b). The reference tone continues throughout the stripe.

This chapter contains discussion of each part of the recording/observing equipment. It commences with a brief description of the Molonglo radiotelescope the east-west arm of which was used for observations. Section 2.2 describes the videotape recorder. The bulk of the Chapter is concerned with the "recording chassis" which conditions the signal from the Molonglo 2nd i.f. outputs for recording together with the reference tone and sync pulses. A simplified schematic of the system is shown in Figure 2.1A with more detail in Figure 2.4B.

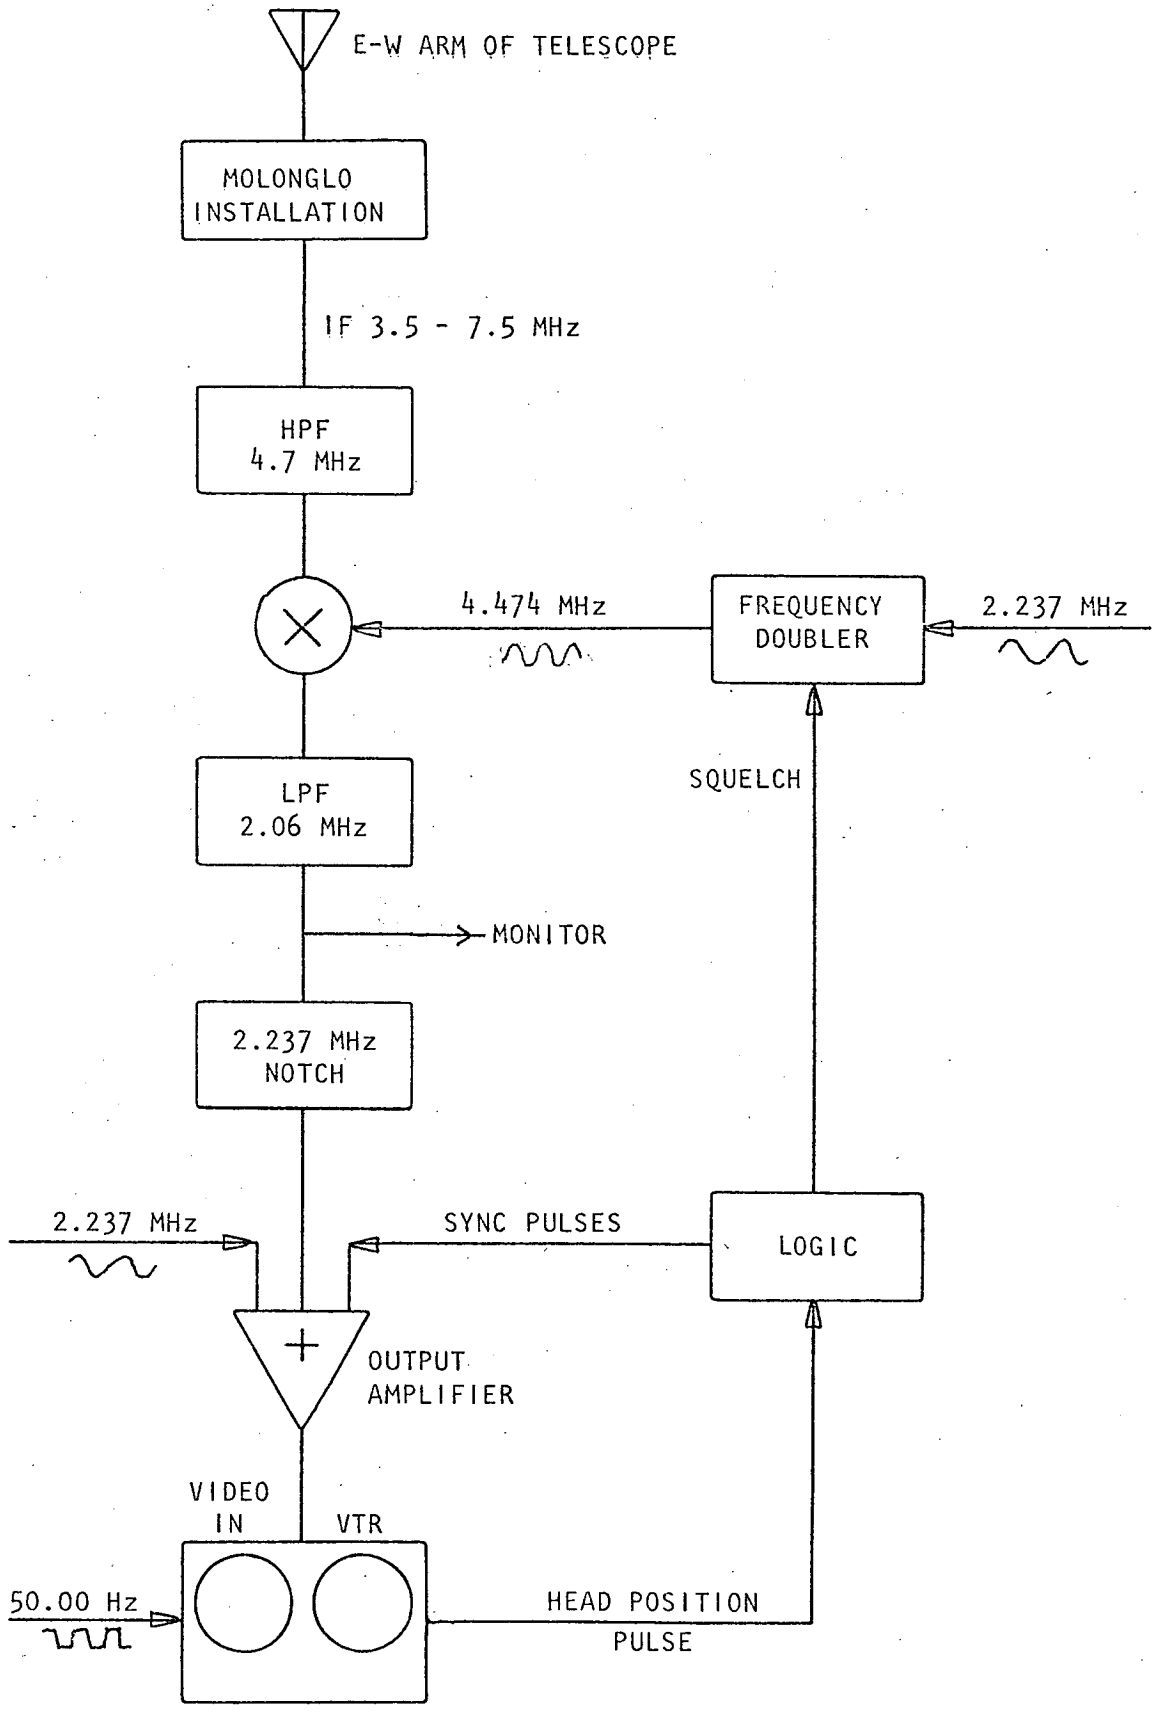

FIGURE 2.1A Simplified schematic of the recording system.

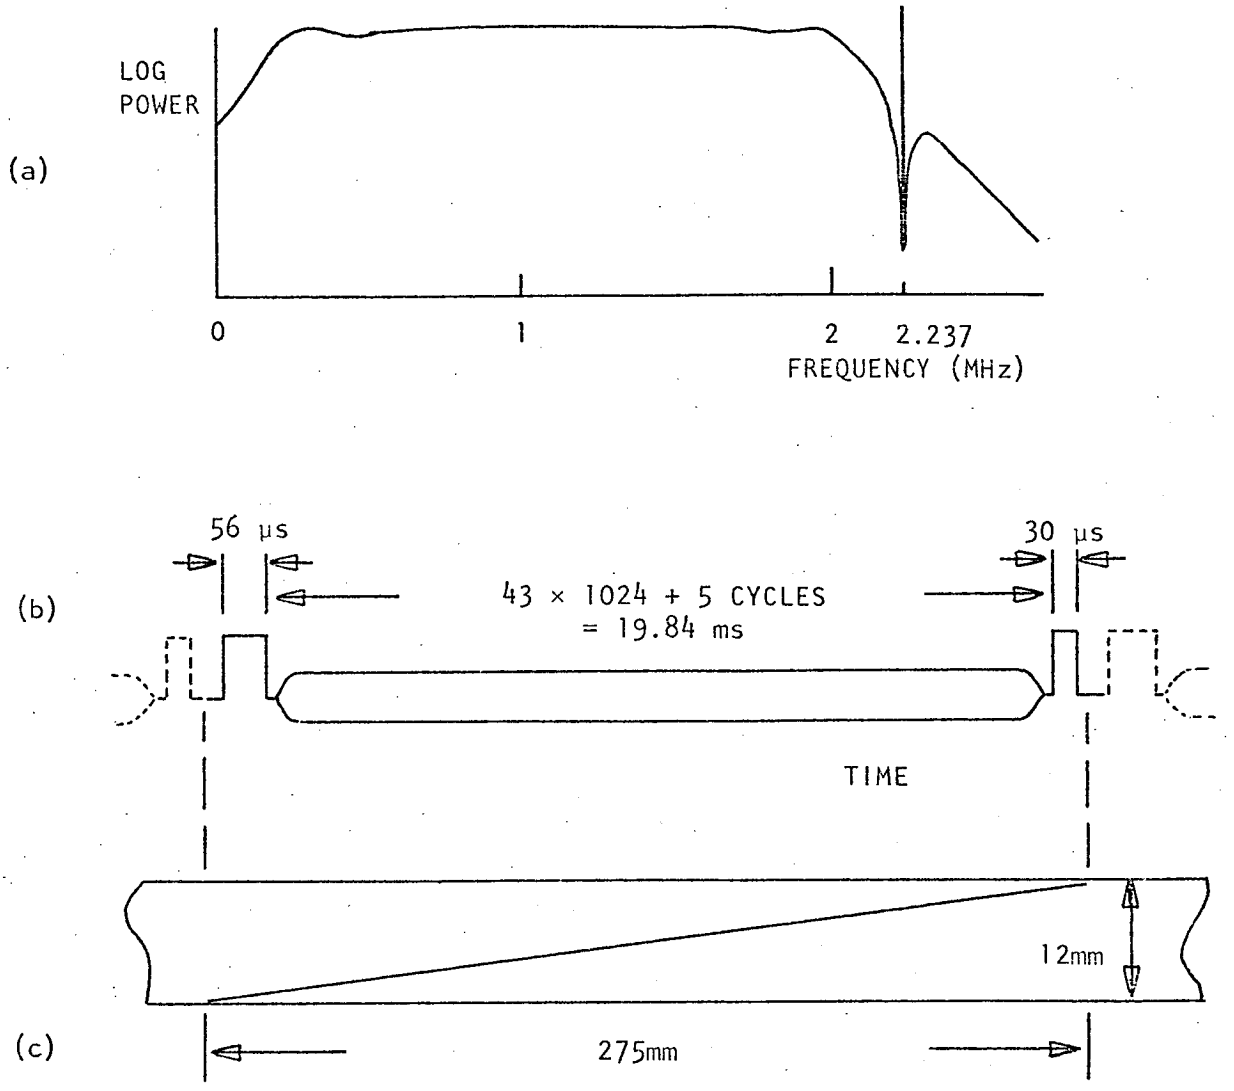

FIGURE 2.1B Idealized recorded spectrum and waveform.

(a) Recorded spectrum, showing the reference tone as a line at 2.237 MHz.

(b) Time domain view, showing the sync pulses preceding and following the data envelope.
The dotted lines denote the preceding and following stripes.
The reference tone is not depicted - it is always present.

(c) Corresponding locus of a recording head across the videotape, i.e. one stripe.

2.2 MOLONGLO RADIOTELESCOPE

The cross radiotelescope at Molonglo, near Canberra, operated by the University of Sydney, has been described by Mills *et al* (1963) and Large and Frater (1969). There are two arms each 1560m long and 12m wide with a total collecting area of 19000m^2. Each arm has the form of a cyclindrical parabolic reflector with dipoles at the focus polarised east-west. The east-west arm can be tilted mechanically for steering in declination. When used for total power observations at 408 MHz the reflector is fully illuminated, and the instrument has a fan beam 4°.2 north-south by 1°.5 east-west (3 db) (Sutton 1968). Three simultaneous fan beams are generated of which the "early" and "late" are separated by 2.8 min arc.

The receiving system uses single-sideband double conversion from a centre observing frequency of 408 MHz to a second intermediate frequency of 5.5 MHz. For total power observations using the early and late beams of the east-west arm the bandwidth is 4 MHz (-1 db). For correlation observations using the whole cross a bandwidth of 2.5 MHz is used. Therefore the system response across the central 2.5 MHz of the 4 MHz is uniform in power to ±0.5 db and linear in phase to ±1° approximately.

The system noise temperature during 1975 was over 500°K (total power, east-west arm). It was equivalent to a source flux of (119 ± 3) Jansky (M.I. Large, private communication). The power level at the 2nd i.f. output was about -22 dbm (i.e. 22 db below 1 mW).

2.3 THE VIDEOTAPE RECORDER

The videotape recorder (VTR) employed is a Sony model CV-2100 ACE. This is a domestic quality helical-scan recorder using 12.5mm tape and suitable for 50 field-per-second black-and-white operation. Playing time of 40 minutes is obtained with a 720m reel of tape, the tape speed being 29.14cm per second. Two recording heads diametrically opposed rotate

diagonally past the tape at 1500 rpm. Under normal operation with the tape moving each head traverses a stripe about 28 cm long inclined 2.5° to the edge of the tape. Each stripe corresponds to a single 20 ms television field. Successive stripes are displaced 5.8 mm along the tape.

Along the edges of the tape are recorded tracks of (quasi) audio information. As designed one would contain the television sound and the other head synchronization pulses. During normal replay the video head rotor is synchronized to these pulses so that the heads retrace the stripes that were recorded. During the preceding recording the rotor is locked to the composite video. Neither of these modes of synchronization was applicable to the radio-astronomy application, so requiring one of a number of small modifications to the VTR used in the project.

The particular VTR used was modified to synchronize its head rotor onto an external 50 Hz square wave. A switch was installed to enable synchronization during running replay onto the pulses in the "audio" track, that is, as the designers intended. This is necessary, for example, for finding the pulsar pulses to be digitized. A negative-going head rotor position pulse (50 sec^{-1}) is available at the connector panel of this model VTR. No modification was required to make use of this by the recording chassis. However during digitization it is necessary to select one of the heads and a TTL monostable was installed to generate a head position pulse (25 sec^{-1}) for synchronizing the replay electronics. Modification of the head resonance response of the chosen head was also required as described in Section 4.3.2.

The "television sound" track was used to record stripe identification in the form of a nine-bit binary number. Another bit noted whether a pulse had just been detected by the "pulse logger" circuitry at Molonglo. The ten bits were encoded as bursts of a 10 kHz tone together with mark pulses. This identification scheme will not be described further as it is not very

useful. It is unavailable when the tape is advanced by hand with stop-action replay. Advancement by hand has been the only successful technique for finding the stripes to digitize. It should be noted that only about 1000 stripes are involved in each transit of a pulsar and there were only three days of observing. Therefore the manual scheme is feasible though tedious.

The video input to the VTR is unbalanced 75 Ω with a rated level 1 - 3 Vpp, the output is unbalanced 75 Ω, 1.4 Vpp. There is a potentiometer to control the input level. The automatic gain control (AGC) has not been used for any recording.

The video is recorded in this type of VTR using high deviation frequency modulation (FM) of a carrier near 3 MHz. The manufacturers' line resolution figure implies a video cut-off frequency of over 4 MHz and this model has been used to record radio astronomy data up to 3 MHz (Ellis 1973). The manufacturer also quotes a signal-to-noise ratio of 40 db. However when bench tested with sinusoidal signals between 0 and 2 MHz at a normal level (-24 dbm input) harmonics and beats were observed only 28 db below the fundamental. The dynamic range using the noise-like inputs involved in this project is less than 30 db. This figure is the difference between near-saturation and the replay noise obtained with stop-action replay (see Section 4.4.2).

2.4 PASS-BAND DEFINITION

The 4.5 - 6.6 MHz band of the Molonglo east-west arm i.f. output is selected for recording by a combination of a high-pass filter (HPF) before and a low-pass filter (LPF) after the final mixer. The corner frequency (i.e. -0.28 db cut-off) of the HPF is 4.74 MHz which is 0.27 MHz above 4.474 MHz, the frequency of the final local oscillator. The LPF has a corner frequency of 2.06 MHz. Both filters have a designed pass-band ripple of 0.28 db and are of the Cauer-Cauchy type tabulated by Zverev (1967).

The design of these filters involved a compromise between the sharpness of the amplitude cut-off and the phase response of the system. The latter is best expressed in terms of its derivative with frequency, namely the group delay. The sharper the cut-off of each filter the greater the excursion in group delay at nearby frequencies. The system was designed to keep such group delay peaks below 1μs as shown in Figure 2.4A(v).

Minimizing the group delay variation across the pass-band has two advantages : the likely errors in estimation (or measurement) of the phase response are reduced, and compensation for phase response becomes unnecessary for most dispersion removal analyses. At the time the recording system was designed I did not know whether measurement of the phase response of the whole system would be feasible. However a technique was developed - see Chapter 5. Compensation for the phase response is needed only when the resolving time of each receiving channel simulated by software is sufficiently small relative to the group delay variation. The minimum resolving time for a channel of bandwidth B is 1/B. It is attained for pulsar signals with dispersion removal and no post-detector smoothing. Variation of group delay both within and between channels must be considered. Thus, according to the design, phase compensation is not needed for channels with $B < 0.5$ MHz nor for channels with $B < 1$ MHz provided their pass-bands lie within the central 1 MHz range of the data (total range 0 - 2.237 MHz).

Figure 2.4A shows the spectrum at various points through the record-replay-digitize system. No attempt was made in the design to completely eliminate imaging and aliasing. (In this context "completely" is -40 db.) Signals at frequencies immediately below that of the local oscillator, namely 4.474 MHz, are attenuated by 16 db only. However signals further into the lower sideband, namely those which image to 0.25 MHz or higher (recorded or "video" frequency) are attenuated by at least 40 db. Similarly signals that are aliased from above 2.474 MHZ to below 2 MHz are attenuated by 40 db or more. Data at frequencies below about 0.25 MHz and above about 2 MHz are rejected during computer analysis.

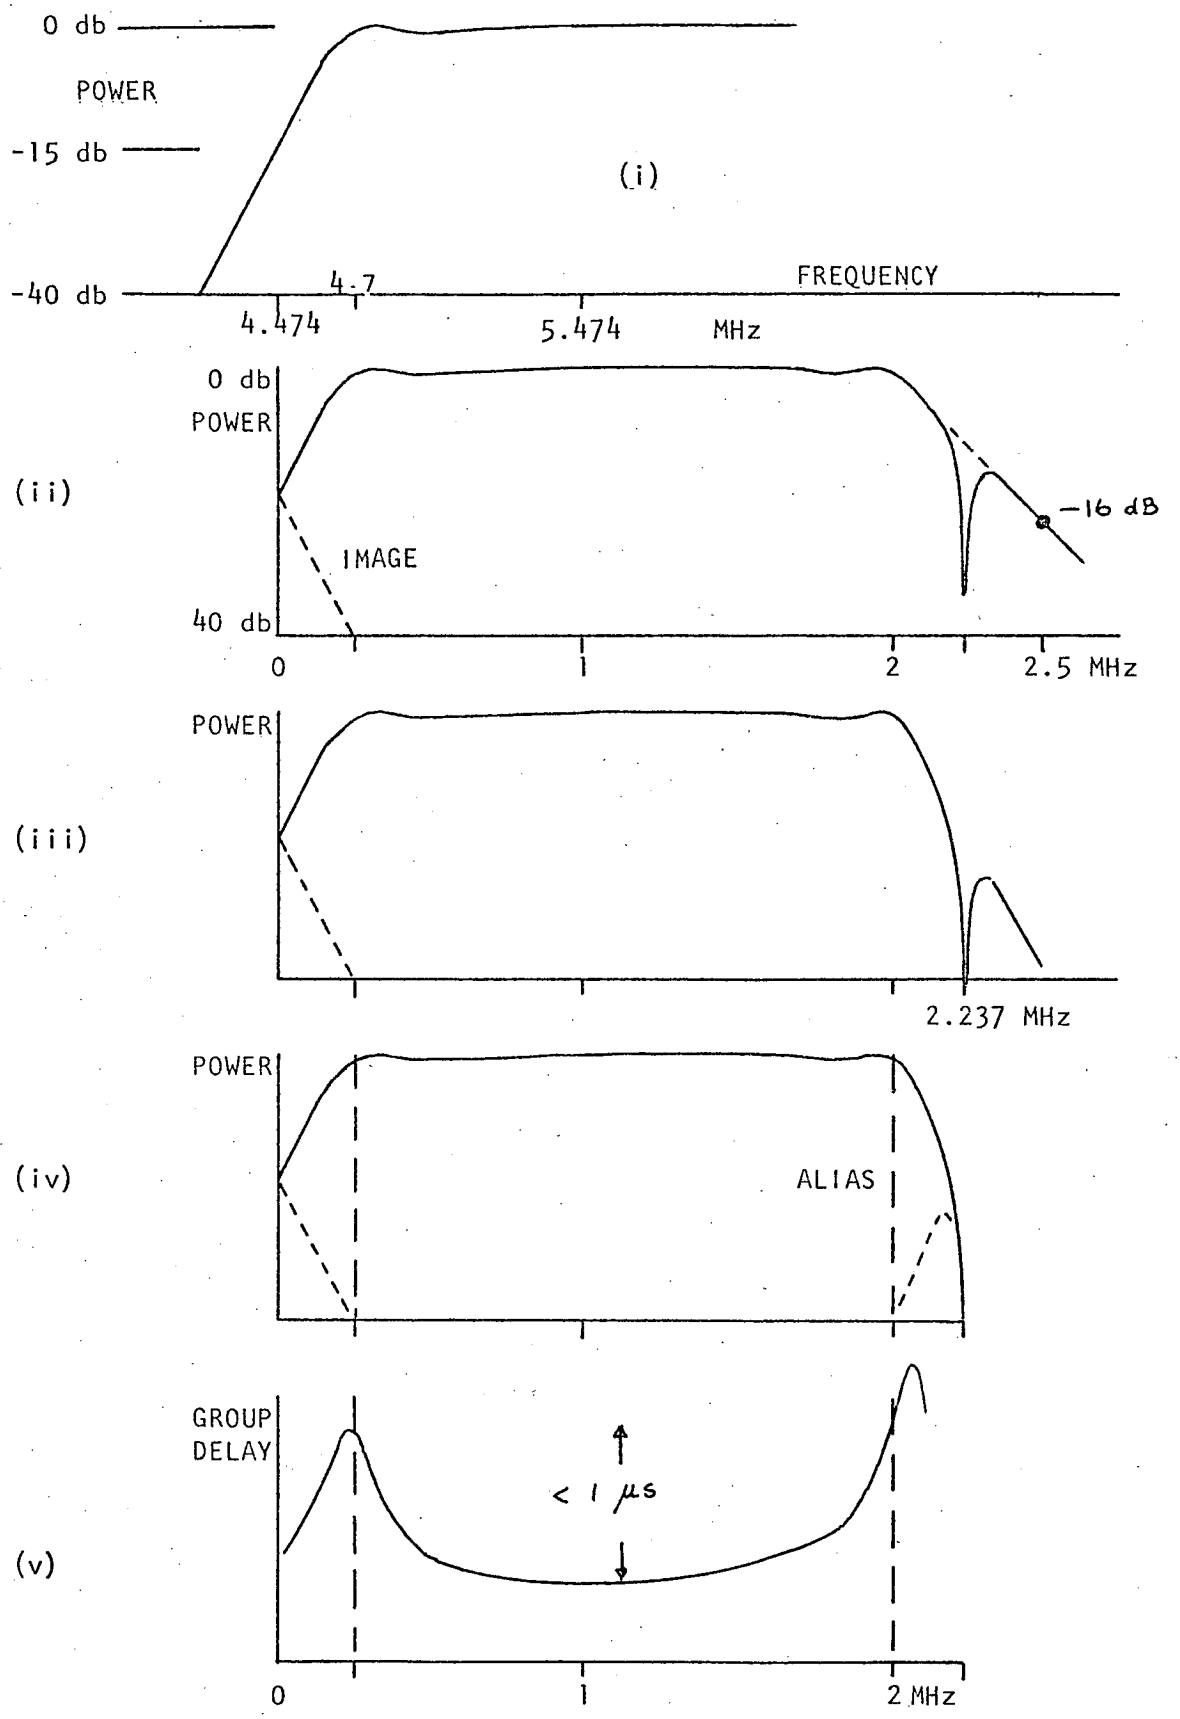

FIGURE 2.4A The Frequency response of the system:

(i) power after the 4.7 MHz high-pass filter; (ii) at the input to the VTR (dashed line below 0.25 MHz indicates imaging);
(iii) on replay, at the input to the sample-and-hold;
(iv) as sampled (dashed line above 2 MHz indicates aliased signals);
(v) sketch of group delay response.

Data outside the broken vertical lines is rejected numerically for 40 db signal to image/alias noise ratio.

The 2.06 MHz LPF limits the bandwidth of the "video" recorded on tape and hence reduces the generation of spurious signals within the VTR. There is another LPF in the replay equipment with a turnover frequency of 2.0 MHz. These two filters together control the amount of aliasing. Therefore the 2.06 MHz LPF does not need a sharp cutoff (a very sharp cutoff is precluded by the group delay constraint). It has an attenuation of 16 db at 2.44 MHz, i.e. 0.2 MHz above the reference frequency. Note that this LPF is not required to provide significant attenuation at the reference frequency itself, i.e. at 2.237 MHz. There is a notch filter for this in the following stage of the recording chassis.

The 4.7 MHz HPF was designed to be sharper than the 2.06 MHz LPF. It is the only filter defining the low frequency edge of the pass-band. It has an attenuation of 40 db at 4.2 MHz, i.e. 0.25 MHz below the local oscillator.

Both filters are isolated from the mixer by fixed attenuators (see Figure 2.4B). To make up for these losses and to provide some additional gain the HPF is preceded by a broadband amplifier (50Ω, gain 25.0 db flat 4 - 7 MHz). Overall system gain is controlled by a switched attenuator preceding the amplifier, that is, at the input of the recording chassis.

2.5 MIXER AND FREQUENCY DOUBLER

The mixer is double-balanced, a Hewlett-Packard type 10514A. It is driven at +3 dbm by a frequency doubler which generates 4.474 MHz from the 2.237 MHz output of the reference tone crystal oscillator. There is a junction field-effect transistor (FET) across the output of the doubler to provide a squelch. Switching the FET "on" gives a 30 db reduction in conversion gain. The mixer is driven somewhat below saturation, viz. conversion gain down 1 db, to maximize the squelch effect. The squelch is asserted during the sync pulses at the start and finish of each 20 ms stripe (see Figure 2.1B(b). It can also be operated manually to place timing marks on the videotape by switching off the data signal for the whole stripe.

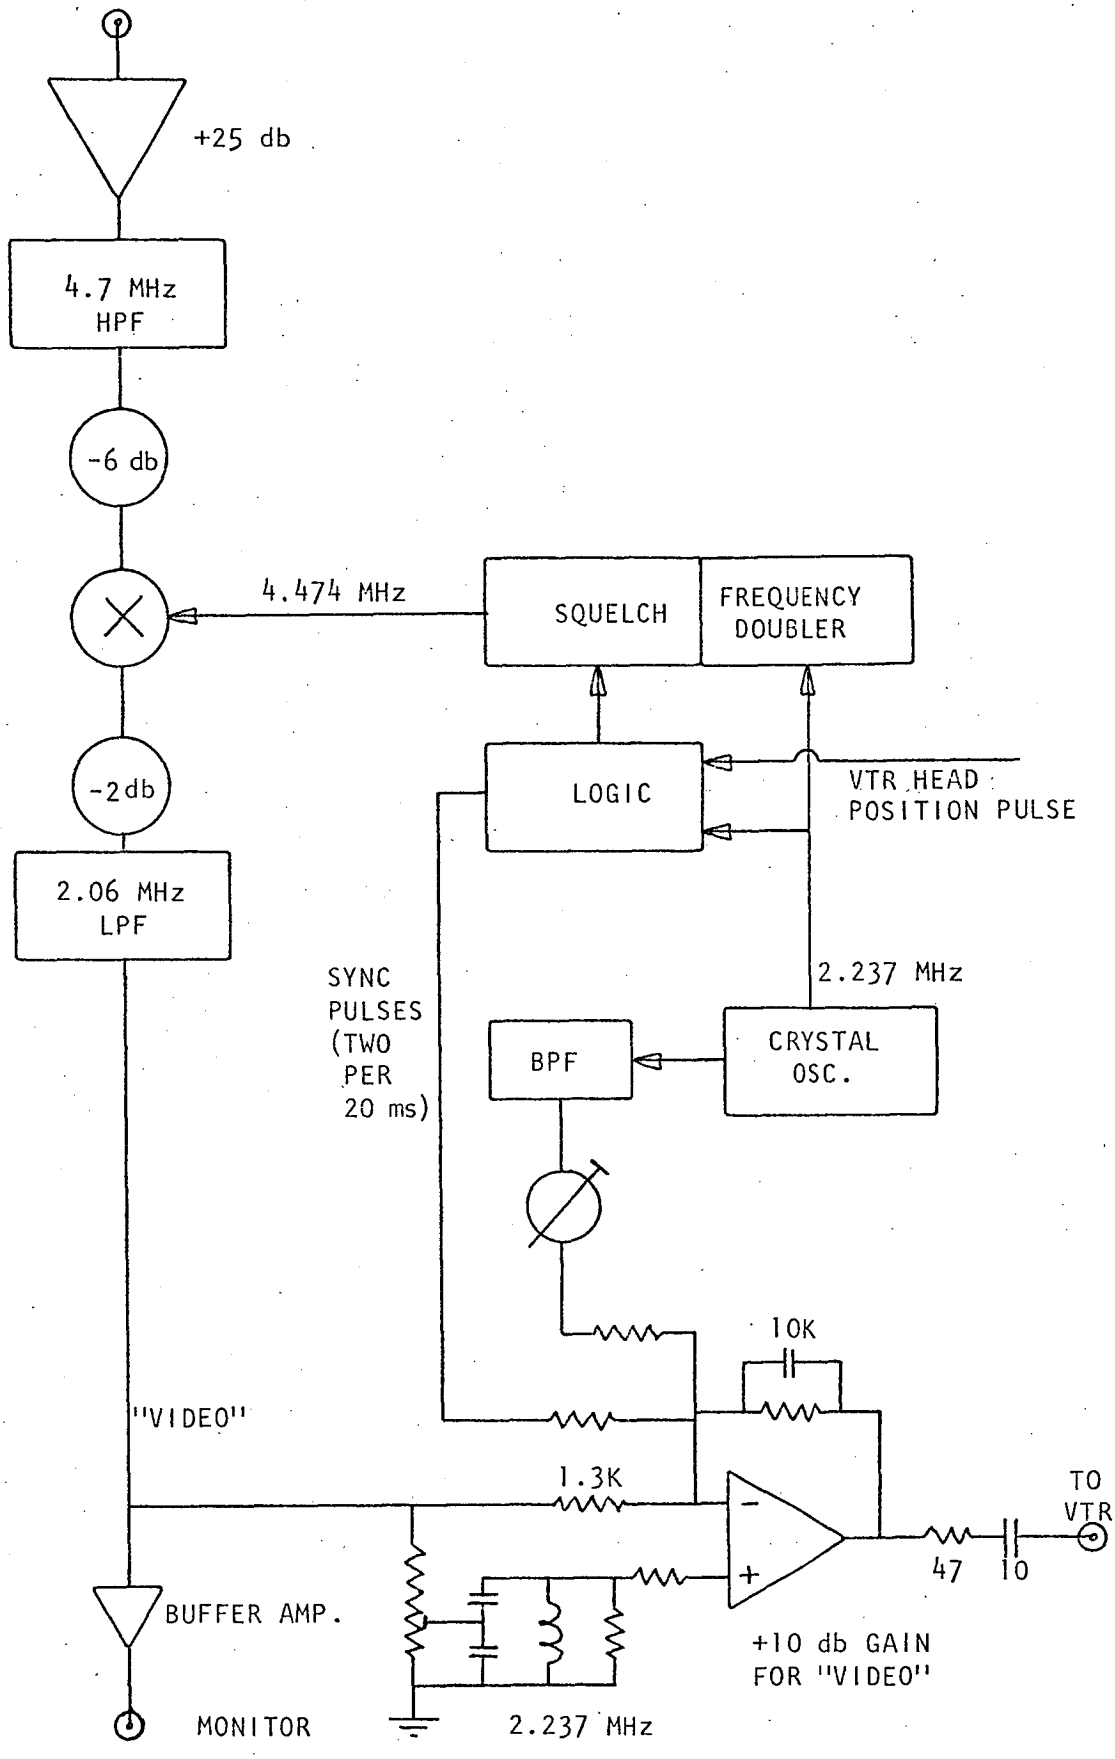

FIGURE 2.4B Schematic of recording chassis.

2.6 REFERENCE TONE

The reference tone is a 2.237 MHz signal that is recorded with the data. On replay sampling takes place at twice that rate so that the reference frequency, 2.237 MHz, is the Nyquist frequency of the system. Since it is convenient to double the reference frequency to provide a local oscillator for the mixer, the reference frequency determines the exact centre frequency of the observations as well as the maximum bandwidth. That is, 4.474 MHz in the Molonglo i.f. output is heterodyned to zero frequency while 6.711 MHz (= 3 × 2.237 MHz) becomes 2.237 MHz. In fact there is a frequency inversion in the preceding Molonglo receiving system so that on videotape zero frequency corresponds to 409.025 MHz at the antenna while 2.237 MHz corresponds to 406.788 MHz. Note that the exact value of the reference frequency is not critical. Any frequency near 2.25 MHz would suffice.

The reference tone is derived from a TTL-output crystal oscillator (Hi-Q type Q06) via a band-pass filter which attenuates the harmonics to 50 db below the fundamental. Thus an almost pure sinusoid at 2.237 MHz reaches the VTR (where various distortions and beats occur). The level of this signal is adjusted with a switched attenuator between the band-pass filter and the output amplifier. This adjustment is critical to about 6 db. The normal setting is 6 db (measured with a 30 kHz bandwidth) above the video spectrum continuum near 2 MHz.

2.7 OUTPUT AMPLIFIER

A schematic of the final stage in the recording chassis, the output amplifier, is shown in Figure 2.4B. In this stage the reference tone and sync pulses (discussed below) are summed to the "video" from the 2.06 MHz LPF. The levels of the reference tone and sync pulses are adjustable. The stage has a fixed gain for the video of +10 db (at 0.2 MHz, drooping to 8 db at 2 MHz). The active element is a high speed operational amplifier

(National Semiconductors type LH0032). Its output is capacitively coupled to the videotape recorder's video input.

To reconstruct the reference tone during replay with minimum jitter in its zero-crossings, i.e. in the "clock ticks", it is necessary to keep the spectrum in its vicinity as "clean" as possible. To this end there is a 2.237 MHz notch filter in the video signal path. Had the preceding 2.06 MHz LPF been designed for a sharper cutoff such a notch would have been unnecessary. However as previously discussed (Section 2.4) a greatly sharper cutoff would have meant a large group delay peak. Furthermore while locked the phase-locked-loop used to reconstruct the reference tone during replay has a very narrow bandwidth in its response to noise. Therefore only a narrow notch in the spectrum is needed.

The notch filter is implemented by a subtractive arrangement. At 2.237 MHz a signal applied to the video input of the output amplifier appears at both the inverting and non-inverting terminals, i.e. "pins", of the op. amp. equal in amplitude and phase. The signal arrives at the inverting pin directly and at the non-inverting pin via a parallel-tuned circuit, suitably damped, with a voltage "step-up" to compensate for losses. There is no effect on the signals entering the other two inputs of the output amplifier, namely the reference tone and sync pulses. The arrangement works very well. The maximum rejection is over 40 db with a 0.3 MHz bandwidth at -3 db. There is negligible effect on the video frequency response below 2 MHz, neither in amplitude nor in phase. Some field adjustments are necessary but easy.

The spectrum of the total signal as recorded is shown in Figure 2.1B(a). The reference tone appears as a spectral line centred in the notch.

2.8 SYNCHRONIZATION PULSES

A time domain view of the total recorded signal is depicted in Figure 2.1B (b). There is a sync pulse at each end of each stripe. The first pulse defines the zeroth clock tick. That is, the first rising zero-crossing of the reference tone following the trailing edge of the first sync pulse is uniquely marked. This marking is preserved during stop-action replay. The second sync pulse is for error checking. During stop-action replay the number of cycles of the reference tone, as reconstructed by a phase-locked-loop, between the falling edges of the two pulses has to be repeatable.

The sync pulses are generated by TTL gates and capacitively coupled into the output amplifier. There is a screw-driver adjustment of their amplitude; it should be 0.5V at the VTR input. The first sync pulse is generated after an adjustable delay of about 1 ms triggered by a VTR rotor position pulse (one each 20 ms signifying that one of the two recording heads is about to start traversing a stripe). Thus the first sync pulse occurs about 0.1 ms after the start of the stripe. It is 56 μs wide, the width being determined by a monostable. It rises and falls on a rising and falling zero-crossing of the reference tone. In fact this detail of the design is unnecessary. It is only necessary that the falling edge bears a fixed phase relationship to the reference tone. On replay the falling edge presented to the discriminator extends over a few cycles of the reference tone. Changing the discriminator level selects different zero-crossings to define the zeroth clock tick.

The second sync pulse occurs 0.1 ms before the end of the each stripe. It commences exactly (43 × 1024 + 5) cycles of the reference tone, that is 19.84 ms, after the fall of the first sync pulse. Its width is determined by a monostable, and is 30 μs, that is, about half the width of the first sync pulse so that the two sync pulses can be distinguished during replay.

The video is squelched during the sync pulses so that in the course of digitization the sync pulses are presented to a pulse discriminator as cleanly as possible. The squelch circuit is in the frequency doubler described in Section 2.5. The video is switched "on" about 1 μs after the trailing edge of the first sync pulse and "off" about 1 μs before the start of the second sync pulse. Both transitions are smooth, taking place in a time ~1 μs. This is depicted in Figure 2.1B(b).

CHAPTER 3

THE RECORDING SESSION

3.1 OBSERVATIONS

Observations were made at Molonglo over three days, 1975 June 18-20. Pulses from 15 different pulsars were recorded. There were 36 separate recordings each of about one minute duration. The useful part of each was the central 10-20 seconds in which the pulsar passed through the meridian. By manually switching the input to the recording chassis from the "early" beam to the "late" beam of the east-west arm at the nominal time of meridian transit, the effective observing time was doubled. The observing time per pulsar per day measured to half power of the antenna beams was thus (12 sec δ) seconds where δ is the pulsar declination.

The equipment configuration is shown in Figure 3.1A. Two VTRs were used, one "slaved" to the other so that the two head rotors rotated in phase with each other and with a crystal-derived 50 Hz signal. Both VTRs had the same input signal. This arrangement was to provide duplicate tapes without transcription noise. Duplicates were needed because of the tape wear occurring with stop-action replay. A time mark was recorded simultaneously on both tapes by momentarily pressing a push-button on the recording chassis to squelch the data for the whole of a number of consecutive stripes not just during the sync pulses. This time mark was recorded when switching from early to late beam for each pulsar.

The early and late beam outputs were monitored separately using a chart recorder and standard Molonglo radiometers. The output from the recording chassis was also monitored on the chart using a receiver which had wide bandwidth (switchable from 60 to 500 kHz) and could be tuned over the range 0-2 MHz. This receiver was also used to check the recording system by replaying the tapes "on the spot", and later for "pulse finding" for digitization.

Calibration recordings for amplitude and observing frequency were made. In the amplitude calibration recordings the total system noise was modulated by 1 db with 10 ms period. The frequency calibrations involved injecting signals near 408 MHz into the antenna system by the simple

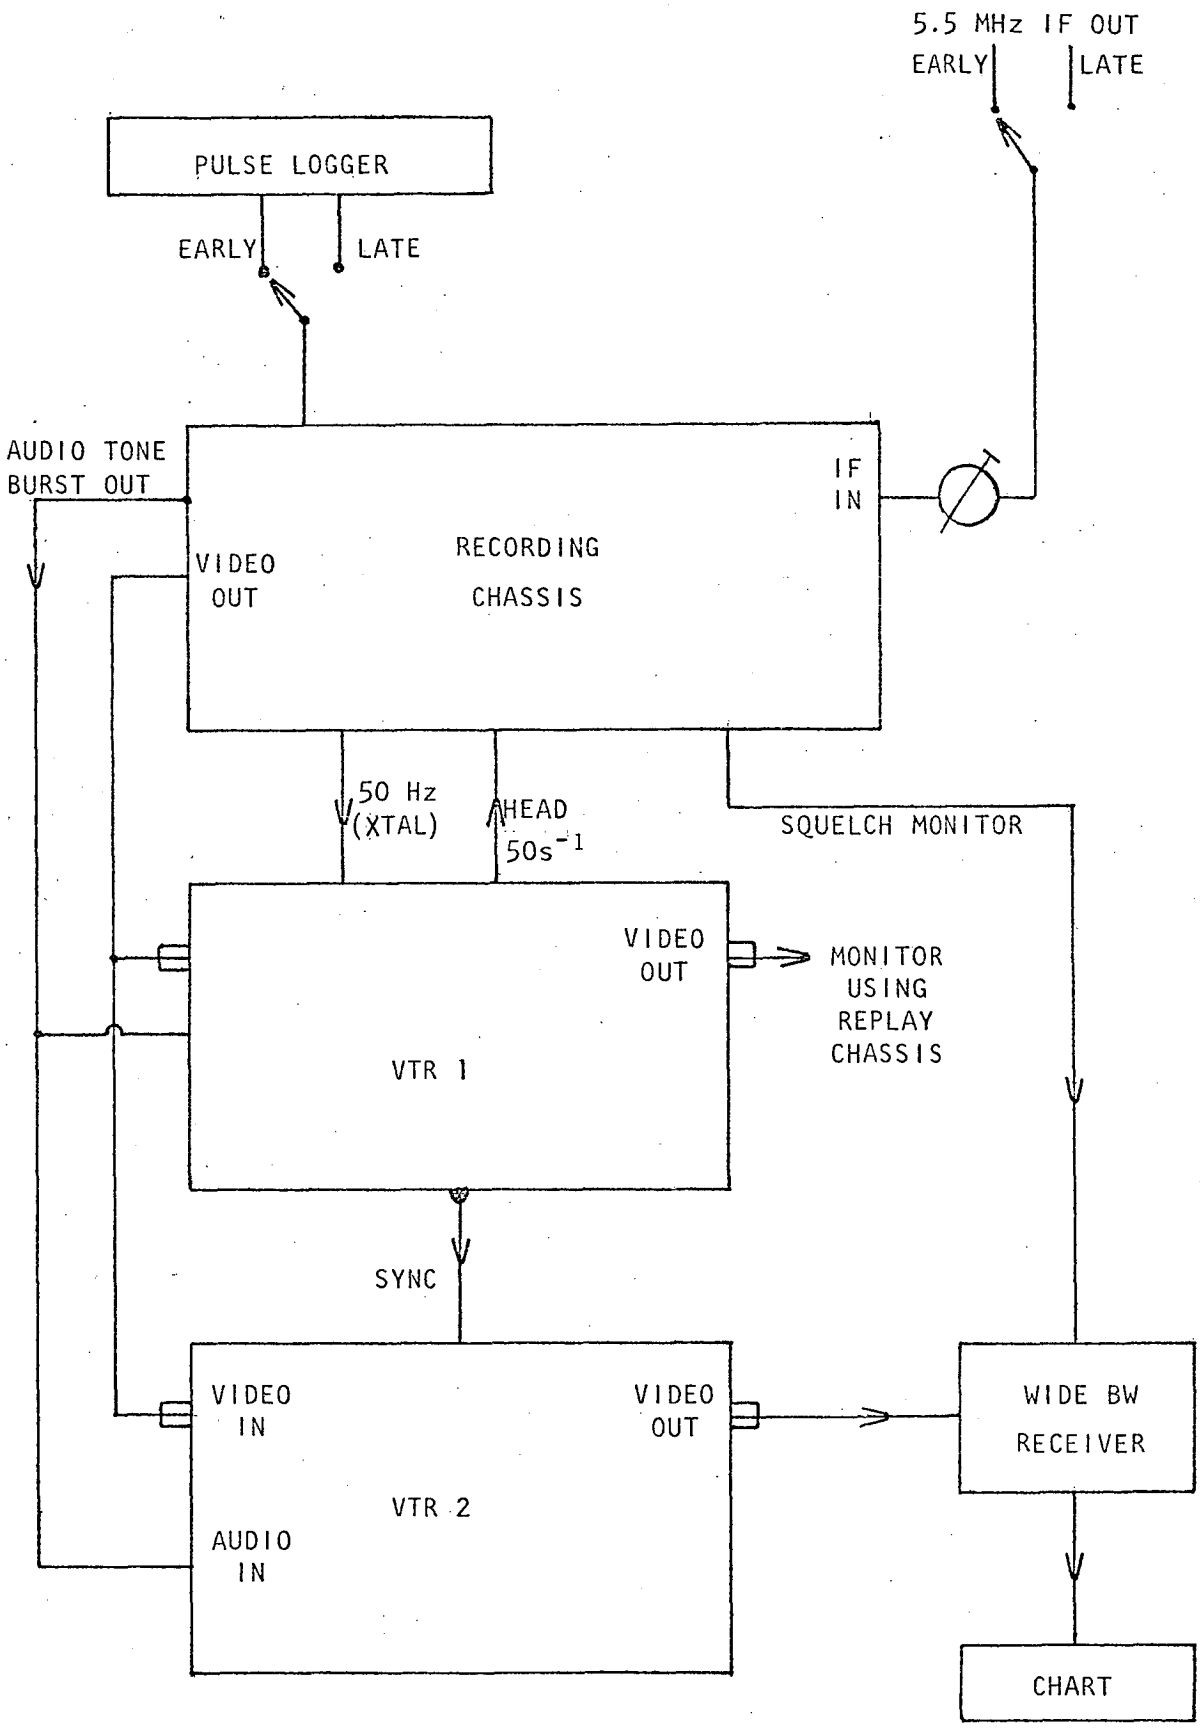

FIGURE 3.1A Equipment configuration during observations.

expedient of draping a wire from a signal generator out the window of the screened control room. Signals at 407.0, 407.5, 408.0, 408.5 and 409.0 MHz were recorded at two different levels in separate recordings. The frequencies were measured to ± 10 Hz.

3.2 PRELIMINARY REPLAY OF THE TAPES

The videotapes recorded at Molonglo were replayed into a 256-channel frequency analyzer at the University of Tasmania. This was normal running replay, i.e. not stop-action. The paper charts made at Molonglo were used as a guide as to which pulsar transits contained strong pulses. Individual strong pulses could be seen momentarily on the CRT frequency analyzer display. When a strong pulse was seen the tape transport was stopped and, with the heads still rotating, i.e. stop-action replay, the tape was wound back by hand to the appropriate stripe. Figure 3.2A shows a typical dynamic spectrum display of one pulse from PSR 0833-45. The pulse appears as a fuzzy sloping line across the 2 MHz (vertical) by 20 ms (horizontal) display. The slope is upward, i.e. rising in video frequency with increasing time, because of the frequency inversion occurring at the second mixer in the Molonglo installation. The magnitude of this slope depends on the dispersion measure of the pulsar. For a low dispersion pulsar, e.g. PSR 0950+08 with DM = 2.969 $cm^{-3}pc$, the line is nearer vertical than for PSR 0833-45 with DM = 69 $cm^{-3}pc$. Of interest in the Figure is the sudden onset of the pulse at all frequencies in the pass band. The pulse profile is approximately a truncated exponential curve with a characteristic time ~ 3 ms. (Komesaroff *et al* 1972). It is dominated by interstellar scattering.

Later (in 1977) I used a fast chart recorder with ultra-violet sensitive paper to systematically analyze selected transits. These chart recordings conveniently displayed the strength of each pulse (strictly speaking, sub-pulse) and where it was placed relative to a stripe boundary.

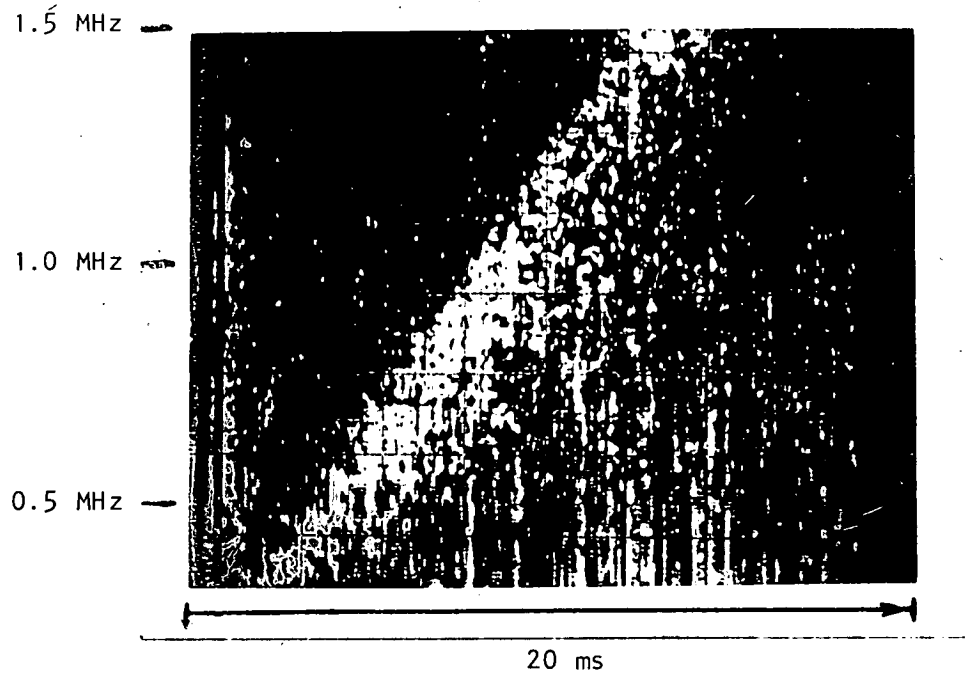

Figure 3.2A Spectrum Analyzer Display of one pulse from PSR 0833-45 Replayed from Video-tape.

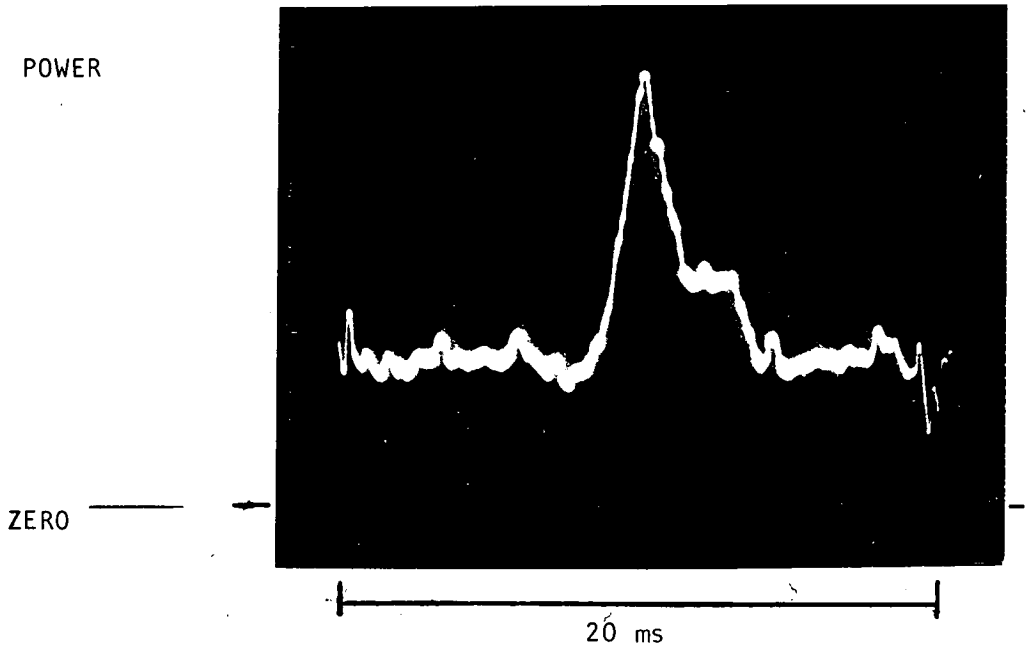

Figure 3.2B Replay of pulse "D" of PSR 0950+08 through wide bandwidth receiver (F = 200 kHz, BW = 120 kHz, RC ≒ 1 ms).

Having found a stripe containing a pulse suitable for digitizing the place along the videotape was marked by inserting a slip of paper into the reel. Later when digitization was performed using a minicomputer at a field station the marking of the stripe was confirmed using the wide bandwidth receiver. The detector output from this receiver for one of the stronger pulses recorded from PSR 0950+08 (pulse "D") is shown in Figure 3.2B. In this case the receiver was tuned to 0.2 MHz, the bandwidth was 120 kHz and the detector time constant 1 ms. The system noise is more than doubled, i.e. signal-to-noise ratio ~ 1.

There are two criteria which a (sub-) pulse had to meet for selection to be digitized. Firstly it had to be strong enough, viz. approximately double the system noise on a timescale of 1 ms. Secondly the pulse had to be centred in the stripe, i.e. not overlap the ends and appear in the adjacent stripes. This criterion is very restrictive for pulsars with high DM in which the bandwidth sweep time is more than 10 ms, i.e. for pulsars with DM > 40 $cm^{-3}pc$ (for 2 MHz bandwidth at 408 MHz). The case of the mean pulse profile being greater than 20 ms did not preclude any individual pulse from that pulsar from analysis, since an individual sub-pulse could be wholly within a 20 ms stripe. It was that sub-pulse that was analyzed.

A total of 30 stripes containing pulses were selected for digitization and 22 of these were completely analyzed with dispersion removal. The details are given in Table 1. On average one off-pulse stripe was digitized for each on-pulse stripe. Four different pulsars are involved. The Vela pulsar PSR 0833-45 is not one of these in spite of the fact that good recordings of strong pulses were made on two days. Except for some preliminary tests pulses from this pulsar were not digitized for dispersion removal. This is because of the large amount of interstellar scattering suffered by the emission from this pulsar. At 408 MHz the decorrelation bandwidth is about 50 Hz (Backer 1974) so that pre-detector dispersion removal which presumes phase coherence across the whole bandwidth, e.g. 2 MHz, has no advantage over the post-detector dispersion removal methods of other observers.

PSR	Transit	Pulses Analyzed	Comments
0950+08	1975 June 18	1	only one strong pulse during that transit
	19	5	7 others digitized, 1 too strong to digitize (PLL failed)
	20	2	other pulses on stripe boundaries
1133+16	1975 June 19	3	weak; 3 consecutive pulses; 2 from component 1, the other from component 2
1642-02	1975 June 19	2	weak
1749-28	1975 June 19	9	weak

TABLE 1 : Summary of pulses analyzed

CHAPTER 4

DIGITIZATION

4.1 INTRODUCTION

This Chapter is concerned with the hardware and software used to digitize each stripe.

Section 4.2 describes the inter-laced sampling scheme while the implications of this scheme on noise within the system are discussed in Section 4.4.2. The major problem area of over-modulation dropouts is introduced in the discussion relating to the videotape recorder (VTR) in Section 4.3. The long Section 4.5 deals with all the circuitry of the "replay chassis", stage by stage. Particular attention is given to the operation of the phase-locked-loop which is used to reconstruct the clock ticks for sampling. Section 4.6 deals with the computers and the programs used to control the digitization. Measurement of the effective transfer function of the record-replay-sample system is covered in the next chapter.

Figure 4.1A shows the main parts of replay equipment and the main flows of data and logical control. All components shown except the VTR and the computer (including ADC) are assembled into the "replay chassis". Most of these components are described in Section 4.5. Notable omissions from the diagram are a wide bandwidth receiver for finding the pulses on the tape and oscilloscopes for various monitoring tasks.

4.2 INTERLACED SAMPLING SCHEME

Adequate sampling of a bandwidth B ~ 2 MHz requires samples at intervals $t_s \leq 1/2B \sim 250$ ns. The direct approach, viz. sampling and digitizing in real time at intervals of t_s, has two main difficulties. These are: performing each analogue-to-digital (A/D) conversion in a time $< t_s$; providing fast enough mass storage for the many numbers that are generated within the time window e.g. 80,000 in 20 ms. The answer to both problems is an interlaced sampling scheme which can be used if the data in the time window can be recorded in analogue form.

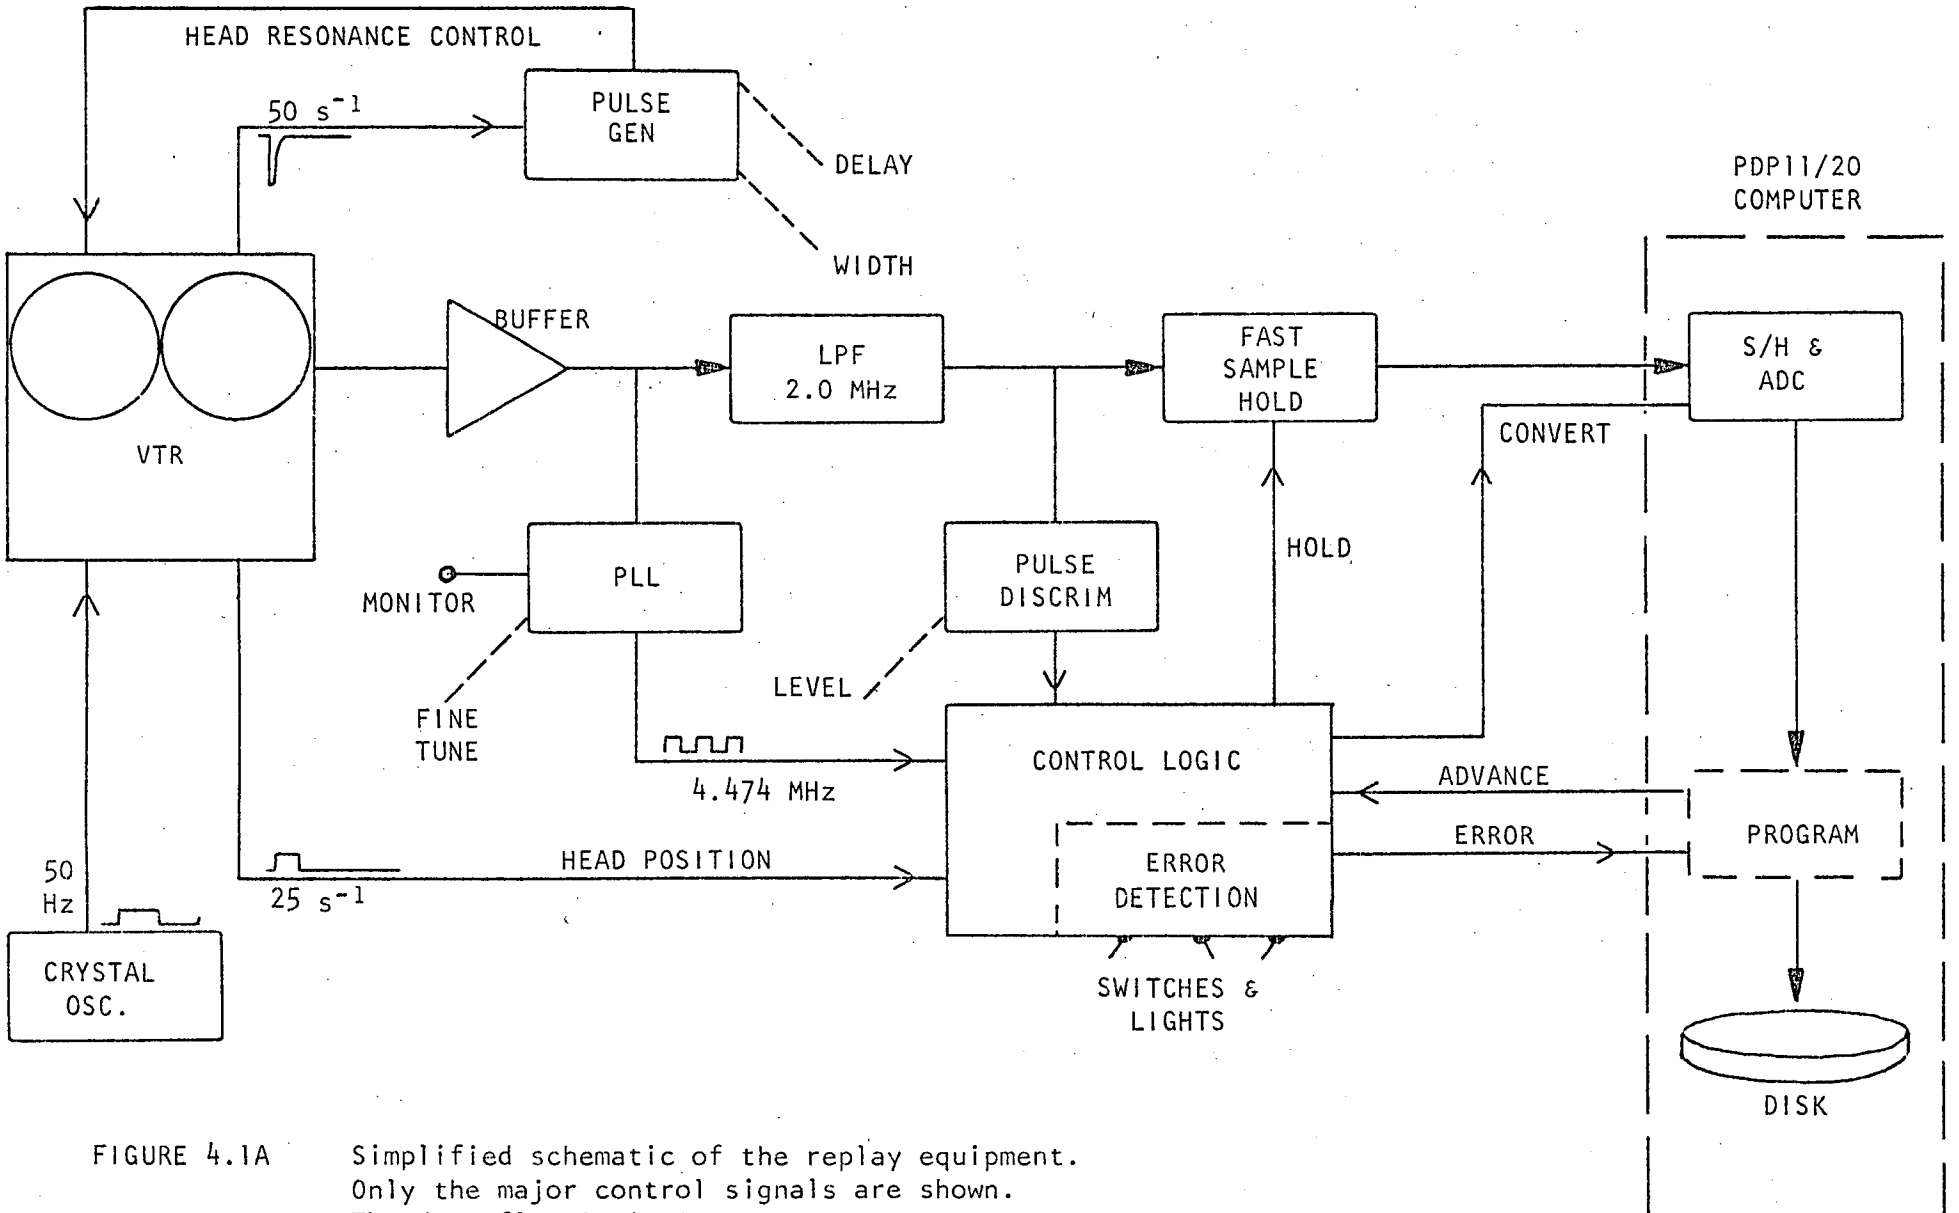

FIGURE 4.1A Simplified schematic of the replay equipment.
Only the major control signals are shown.
The data flow is indicated by black arrow-heads.

During one replay of the recorded data segment, e.g. of one videotape stripe or one revolution of a video disk, as many samples are taken as can be handled by a conventional A/D convertor. These samples are equally spaced throughout the data segment at an interval which is an exact multiple, M, of the basic sampling interval t_s. On the next replay of the segment the whole set of sampling instants, the sampling "comb", can be advanced in time by exactly t_s. In this way the digitization process is completed in M replays of the data segment. Details specific to this project are covered in the next few paragraphs.

The instants within a videotape stripe at which samples are required are termed "clock ticks". These are the zero-crossings of the 2.237 MHz reference tone, therefore t_s = 224 ns. The zeroth tick is defined to be the first upward zero-crossing of the reference tone after the synchronization pulse at the start of the stripe (see Figure 4.2A). The first set of samples are taken at clock ticks

$$256\ i \quad , \ i = 1..344$$

the interval between them being

$$256\ t_s = 256/(2 \times 2.237 \times 10^6) \simeq 57\ \mu s$$

The A/D convertor used has a conversion time of 20 μs giving enough time between samples for housekeeping by the PDP11/20 computer. (The corresponding time for the PDP8/e system was 35 μs.) In subsequent replays of the stripe the sampling comb is advanced earlier by one clock tick at a time i.e. by t_s = 224 ns. After j advancements the set of sampling ticks is

$$256\ i - j \ , \ i = 1..344, \ j = 0..255$$

To completely digitize the 19.84 ms data part of a 20 ms stripe requires 256 replays of that stripe, producing 88064 numbers. In principle this process could be completed in 5.12 seconds but in practice it takes at least one minute for the following reasons:

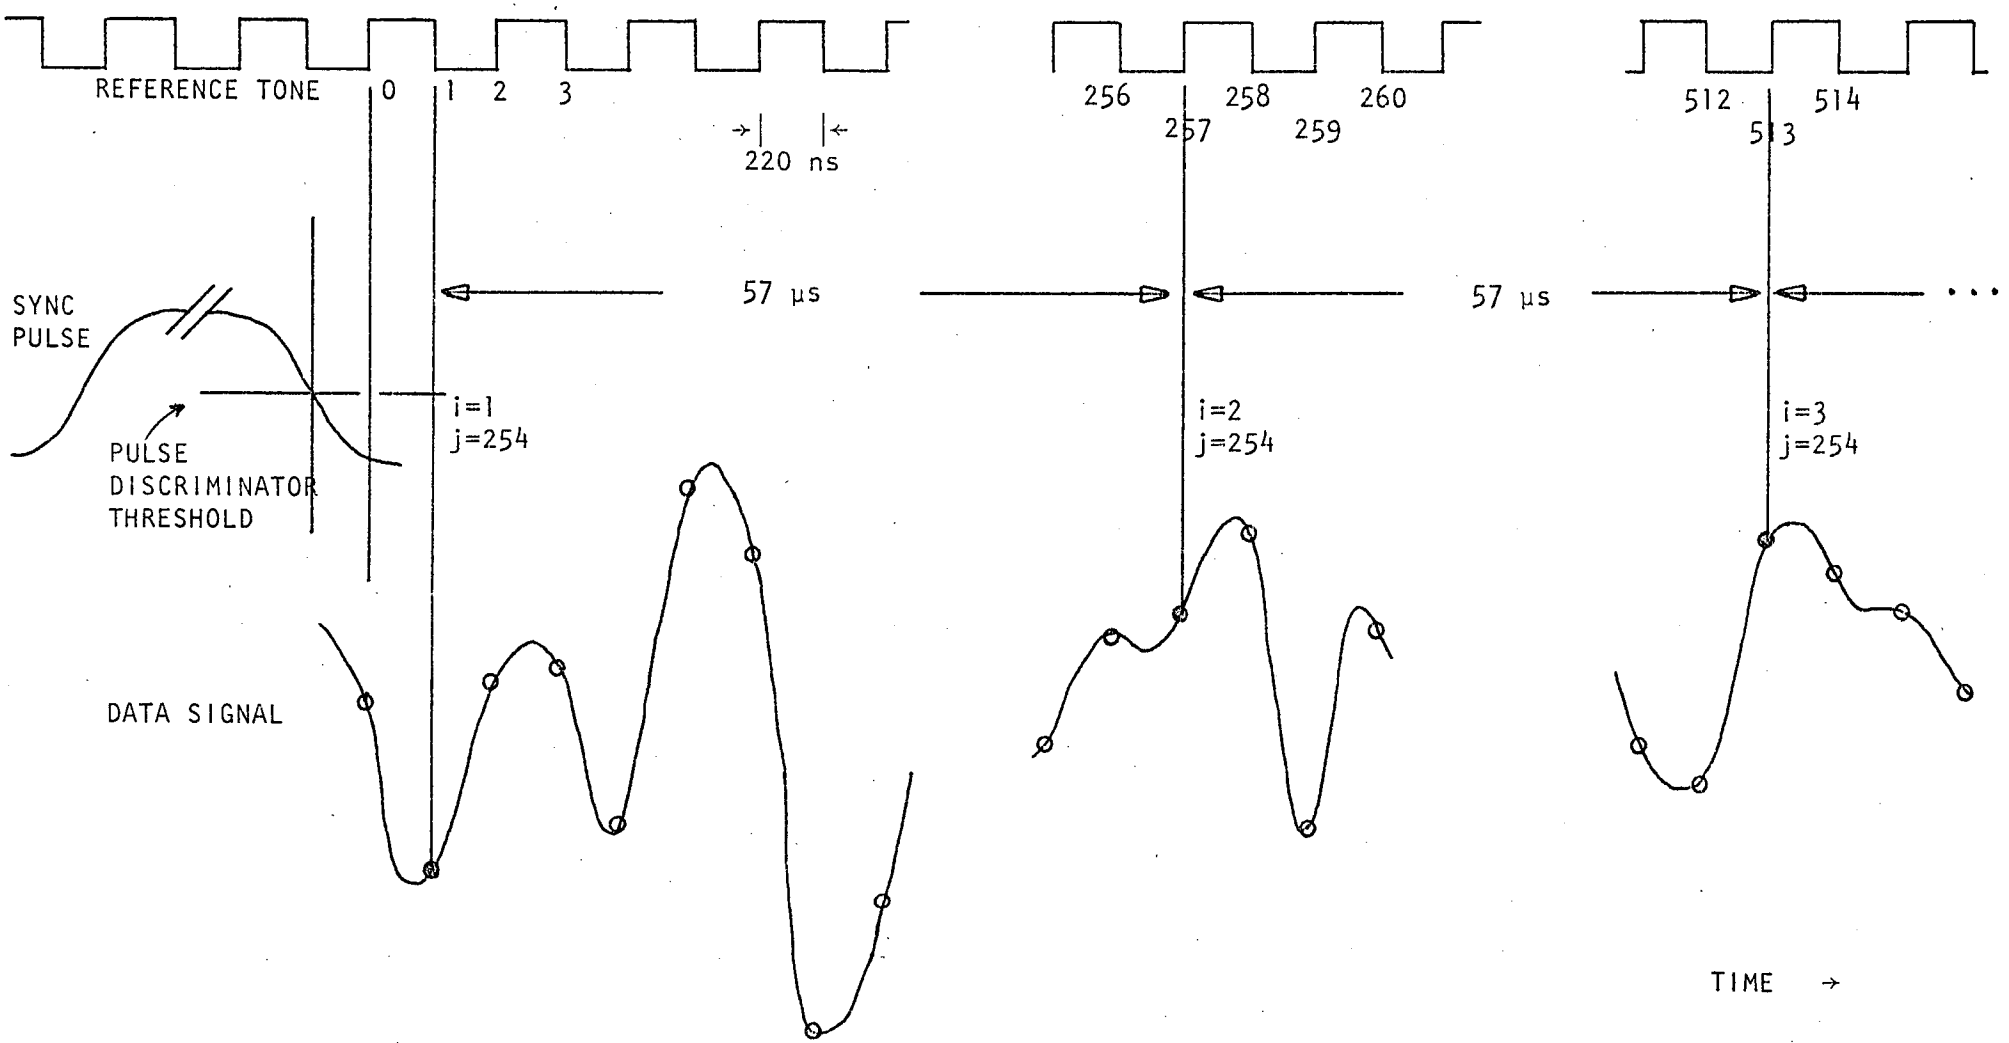

FIGURE 4.2A The interlaced sampling scheme. The first 3 "teeth" in the sampling comb are shown and the comb has advanced 254 times, i.e. this is pass number 254. Samples taken on this pass are denoted by black circles, others by white circles.

The time to digitize a stripe is doubled by the fact that only one of the two VTR heads is used. The 20 ms interval between replays is convenient for error checking and some data validation. Another factor increasing the digitization time is the limited computer store available for use as a buffer. Four or five times during the digitization process sampling and advancing must be suspended while the contents of the buffer are written to magnetic tape or disk. Furthermore the data validation method adopted with the PDP11/20 requires between 2 and 30 (typically 5) replays at each sampling position before advancement.

4.3 USING THE VTR WITH STOP-ACTION REPLAY

4.3.1 Mechanical Details

I use a stop-action replay technique little different from that recommended by the VTR manufacturers. Tape movement is stopped but the heads are left rotating. The desired stripe is found by rotating the tape spools by hand while monitoring the output with a wide bandwidth receiver (typically set to 250 kHz bandwidth at 1 MHz). The tape is held in place against the moving heads only by the weight and friction of the spools. I tried to develop a more positive method of tape tensioning but obtained a negligible improvement in performance. Replaying the same stripe for long periods, say for 30 minutes total time which is typical, wears the tape significantly. Dropouts become more prevalent and replay noise increases. There is wear not only from the video heads but also from the capstan which rotates continuously, with the pinch-roller retracted, against a nearby part of the tape.

For stop-action replay it is necessary to adjust the slant angle of the tape feed so that the heads follow along each stripe. This is because the effective slant angle was different during recording, i.e. with the tape moving. The adjustment consists of raising a guide peg (on the left in Figure 4.3.1A) by 0.25mm. If this is not done then the heads run

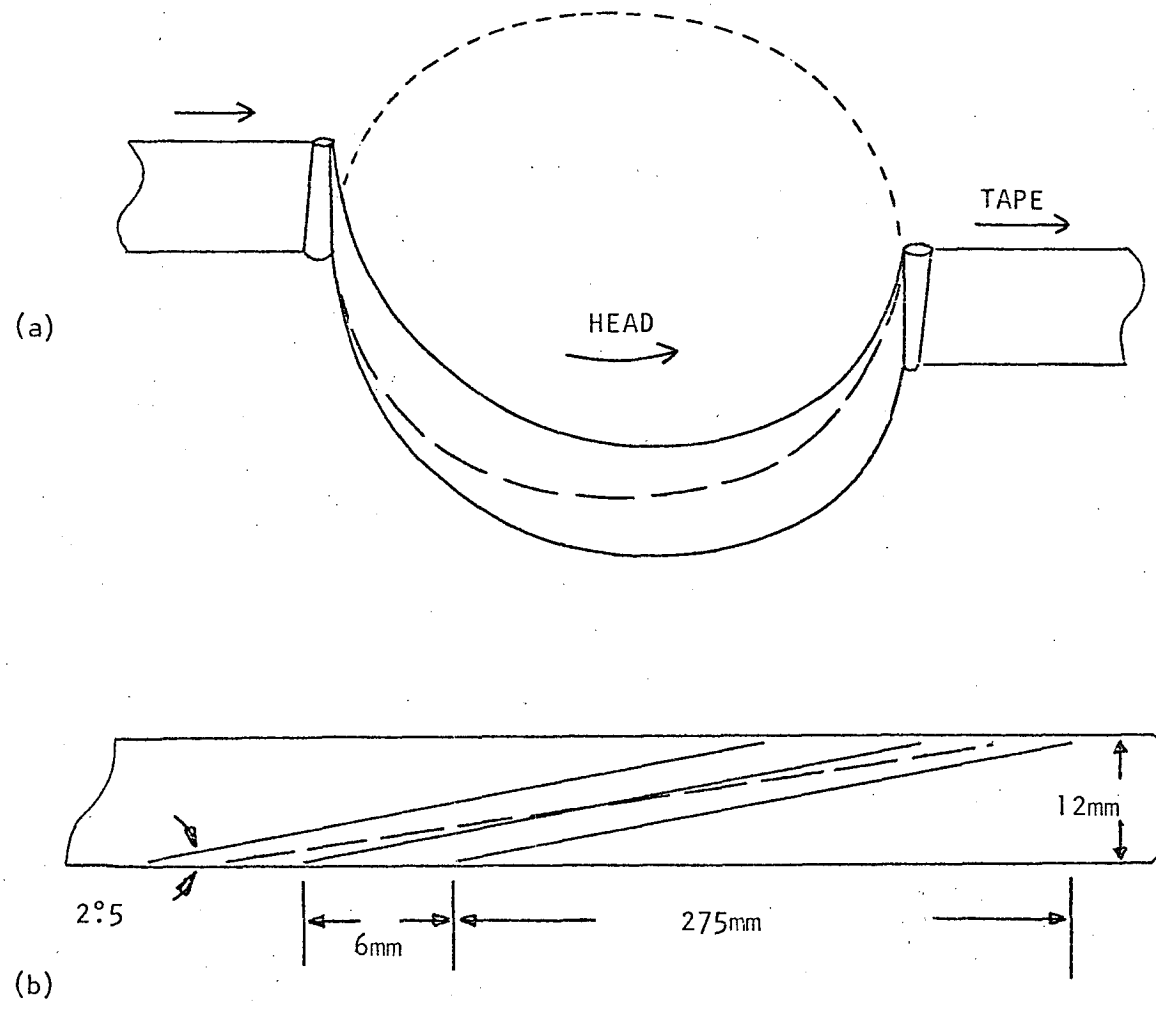

FIGURE 4.3.1A VTR Geometry

(a) Sketch of tape path with a dashed line representing the locus of a head rotating in a horizontal plane.

(b) Sketch of tape showing three stripes (continuous diagonal lines denote their centres) and the locus of a head (dashed line) during stop-action replay before adjustment to the tape path.

off the stripe, at least partially, at some point along its length (see Figure 4.3.1A(b)). The level of the FM carrier drops as a result, causing an increase in replay noise and the occurrence of a group of dropouts, that is, bursts of large amplitude noise due to complete loss of carrier.

4.3.2 Head Resonance and Over-Modulation Dropouts

Only one of the two VTR heads is used for sampling. This is because of the interlaced sampling scheme. Even small differences in the responses of the heads are significant, e.g. a small DC difference might express itself as a component at the Nyquist frequency. A TTL pulse is brought out of the VTR to signal to the digitizing system that the chosen head is about to start along the stripe. Furthermore the "resonance" network for the chosen head is optimised dynamically to better cope with "over-modulation dropouts".

Over-modulation dropouts are the result of recording a flat noise spectrum at too high a level. Each such dropout appears as a large negative spike in the replayed video instead of the fast-rising positive peak that was recorded. The mechanism is that two cycles of the FM carrier are recorded so close together that one is missed by the zero-crossing discriminator within the FM demodulator. Such dropouts are unfortunately prevalent during the stronger pulsar pulses I recorded. They are also discussed in Sections 4.4.1 and 4.6.3.

The head resonance modification to counter over-modulation dropouts is as follows. During the data part of each stripe about 10 pF extra capacitance is switched across the network at the input of the head pre-amplifier. This raises by about 10 db the resonance peak in the head frequency response. This peak is near 3.8 MHz, the free-running frequency of the FM carrier. During the sync pulses at each end of the stripe the head response reverts to its original flatter shape. The extra capacitance is switched into circuit by a 2N3819 junction FET whose gate is controlled by a pulse generator mounted in the replay chassis. A negative-going

pulse to switch off the FET is adjusted in width (~ 1 ms) and phase so that it overlaps both ends of each replay of the stripe (Figure 4.3.2A). It is triggered from the 20 ms head position pulse coming from the VTR.

The demodulated waveform of a typical over-modulation dropout is shown in Figure 4.3.2B. This particular dropout is intermittent as is suggested by the double trace in curve 'a'. The corresponding waveform produced by the modified head circuit is shown as curve 'b'. Clearly the frequency response of the VTR record-replay system is modified. The change in power spectrum, namely a 6-9 db attenuation between 1 and 2 MHz leaving a broad peak around 0.4 MHz, can be seen in Figure 4.3.2C. The phase response is discussed in Chapter 5.

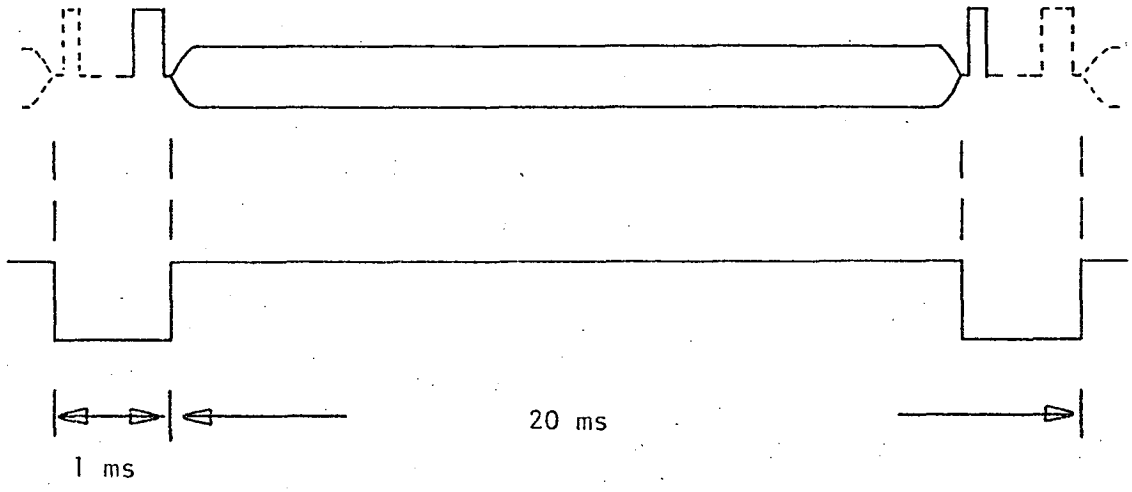

FIGURE 4.3.2A Head resonance control waveform in relation to the replayed video. Dotted lines denote the same stripe replayed by the other head.

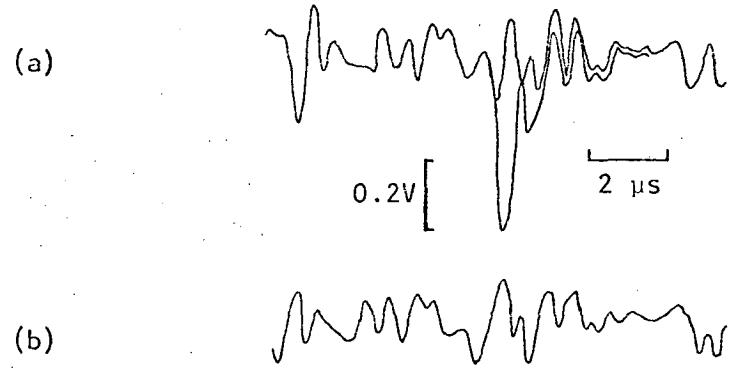

FIGURE 4.3.2B Sketches of oscillograms showing an over-modulation dropout:

(a) un-modified head resonance;

(b) head resonance modified to counter dropouts.

Note: monitoring point is at input to the sample-and-hold, i.e. after the 2 MHz LPF.

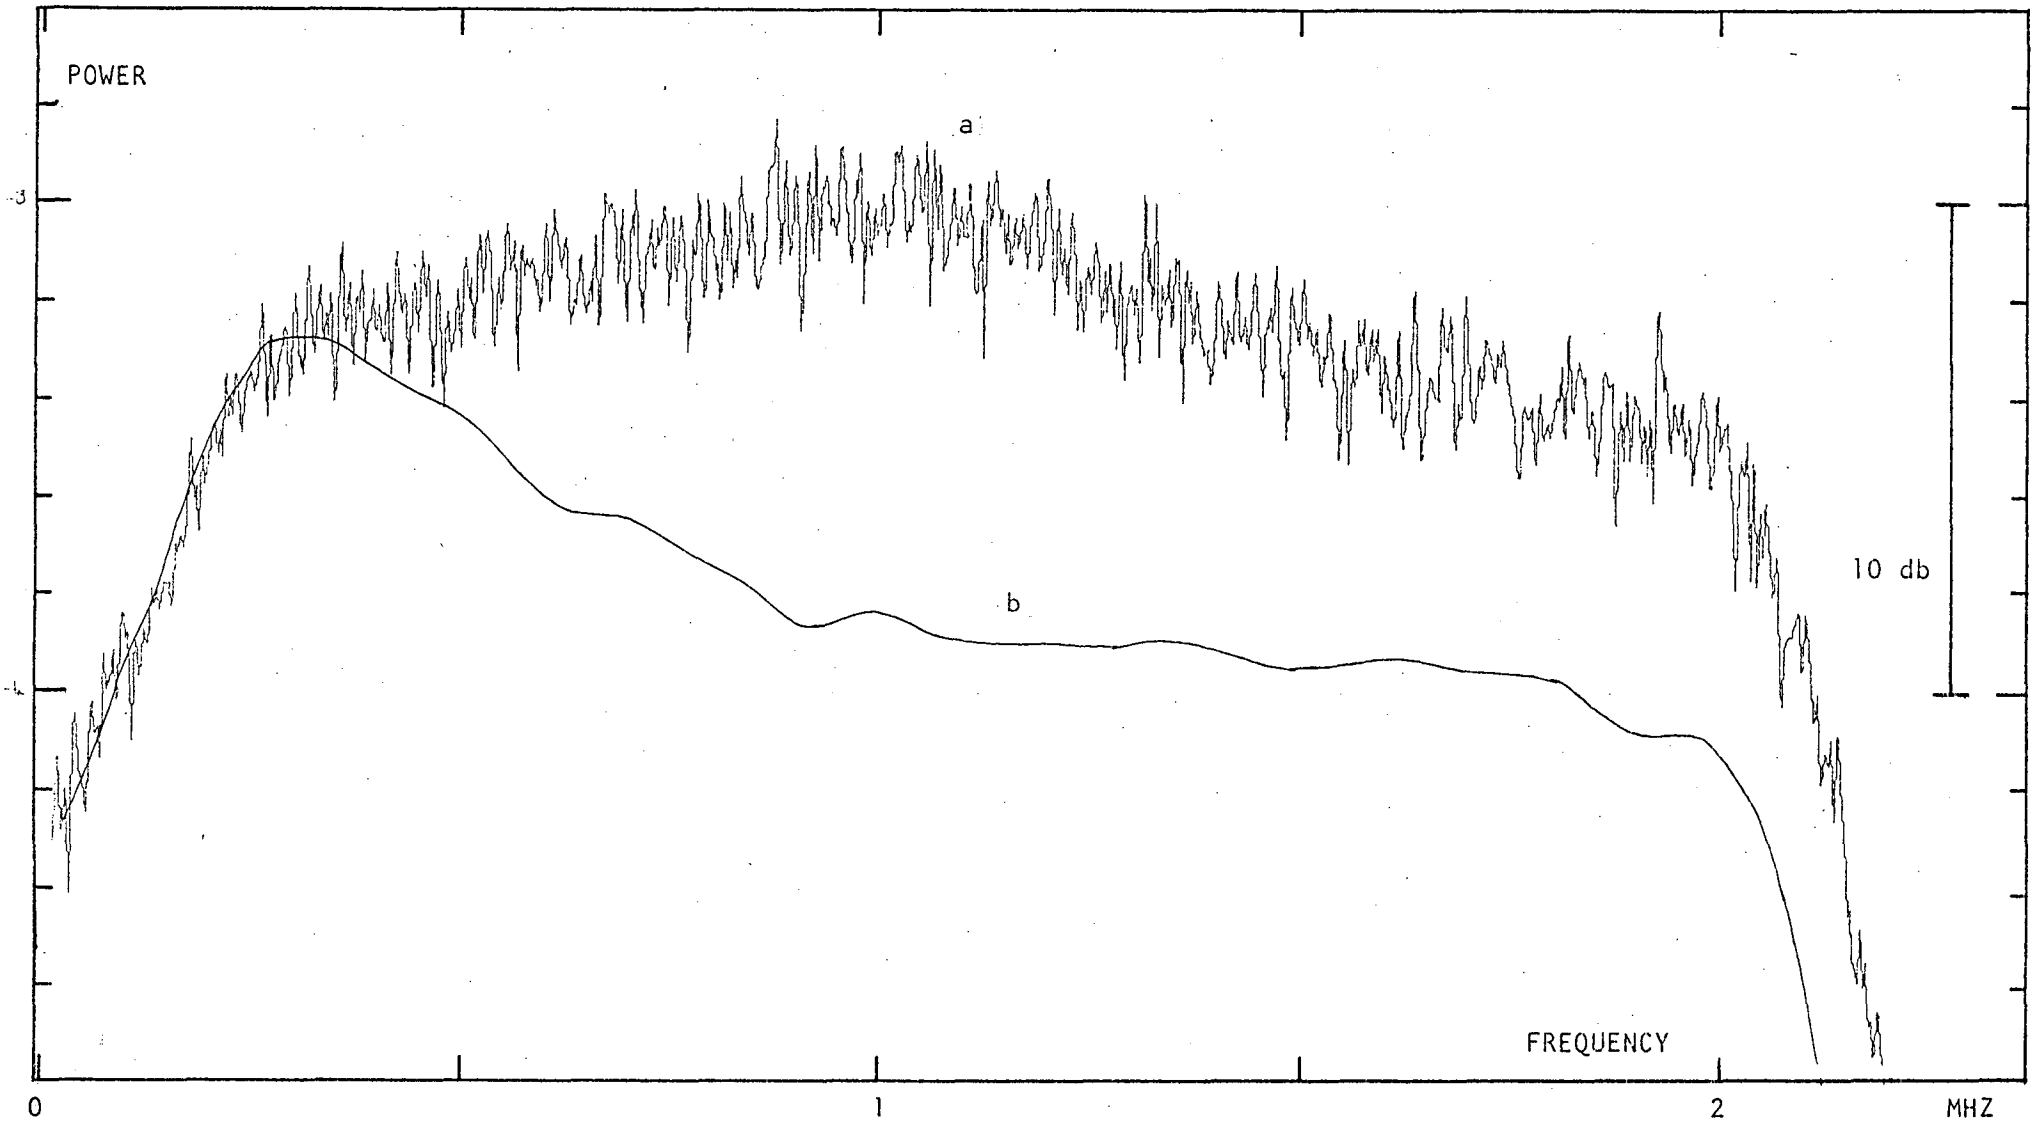

FIGURE 4.3.2C Off-pulse power spectrum with the VTR head resonance

(a) un-modified (approx.), and

(b) modified to reduce over-modulation dropouts, i.e. as used for all pulsar observations.

Curve a is plotted with 500 Hz resolution and curve b (a different stripe) with 70 kHz resolution.

4.4 NOISE

4.4.1 Sources of Noise

The sources of the noise that is present with the pulsar signal during replay include (with characteristics in parentheses):

(i) radiotelescope system noise (Gaussian);

(ii) spurious signals generated at record time in

  (a) the record chassis, e.g. images below the 4.474 MHz local oscillator (quasi-Gaussian),

  (b) the VTR, e.g. intermodulation products (a multiple line spectrum);

(iii) reproducible artefacts of the replay, namely

  (a) reproducible tape noise, e.g. due to magnetic grain (quasi-Gaussian),

  (b) over-modulation dropouts (negative impulses);

(iv) non-reproducible replay noise, namely

  (a) mechanical vibration modulating the FM carrier (low frequency random),

  (b) head pre-amplifier noise (quasi-Gaussian),

  (c) dropouts (bipolar impulsive);

(v) interactions between (iii)(b) and (iv)(a) and (b) leading to intermittent over-modulation dropouts.

Noise is introduced from two more sources during digitization:

(vi) sample jitter i.e. random phase modulation of the sampling comb which is derived from the reconstructed reference tone (quasi-Gaussian);

(vii) quantization noise (random, discussed in Section 4.6.1).

Over-modulation dropouts were described in Section 4.3.2. Ordinary dropouts occur when the level of the FM carrier drops below that of the quasi-thermal noise present at the discriminator, the input stage of the

FM demodulator. These dropouts appear in the video output as a burst of approximately Gaussian noise with peak amplitudes up to 0.7V at which level "soft" clipping occurs. On the other hand over-modulation dropouts are always negative-going spikes of about 0.5V amplitude with a 1-2 μs exponential recovery. There is always an episode of ordinary dropouts during 0.2 ms at the ends of each stripe due to the changeover of the heads and to loss of carrier.

The interactions (v) are due to the non-linearity of the FM discriminator. Usually low level noise, i.e. categories (iii)(a), (iv)(a) and (iv)(b), present at the discriminator is converted to phase jitter in the discriminated FM carrier and hence to quasi-Gaussian noise in the video output. However if a particular cycle of the carrier is near the discriminator threshold the low-level noise can control whether it is lost or not. Latent over-modulation dropouts can appear and others vanish, that is, they are intermittent.

Within the data band 0.2 - 2 MHz the dominant noise sources before sampling are the radiotelescope system noise and the over-modulation dropouts. Within the noise passband of the phase-locked loop (PLL) there is no contribution from the system noise nor from the pulsar due to the notch filter in the output amplifier of the record chassis. However all of the noise sources (ii) to (v) affect the PLL and hence introduce jitter in the sampling comb. This jitter produces noise which is proportional to the amplitude of the jitter and to the slope of the data waveform at each point. The analysis of this is greatly complicated by interactions. For example an intermittent dropout affects both the waveform being sampled and the sampling comb. However the jitter-produced noise can as usual be separated into a reproducible and a non-reproducible component according to the characteristics of the jitter, and the interlaced sampling scheme spreads the latter component uniformly across the spectrum (see next Section).

4.4.2 Effect of the Interlaced Sampling Scheme

Since samples at successive clock ticks are taken at least 40 ms apart their non-reproducible noise components are uncorrelated. This assumes of course that this noise has a low frequency cut-off above 50 Hz, however all details of the shape of its spectrum are lost. The result is a spectrum that is white from zero to the Nyquist frequency.

This effect is well illustrated by Figure 4.4.2A, which shows a typical power spectrum obtained with a 1 MHz LPF before the sample-and-hold. Although the filter attenuation is 35 db above 1.35 MHz the background above 1.35 MHz is flat and 24.6 db below the 1 MHz level.

When a 2 MHz LPF is used the background level is probably increased by 3 db. (There is some uncertainty in this because the spectrum of the non-reproducible noise before sampling is unknown.) Therefore defining the non-reproducible noise plus quantization noise as "noise" and all other signals as "signal", the signal-to-noise ratio between 1 and 2 MHz is 21.5 ± 0.5 db. Below 1.0 MHz the S/N ratio is higher with a maximum of 27.5 db at 0.3 MHz. Averaged over the band 0.2 - 2 MHz the S/N is 24.5 db. This corresponds to an average error of 6% in each voltage sample. Note that these figures apply to stripes that are digitized using the PDP11/20 computer. When the PDP8/e was used quantization noise was more important so that the S/N for some pulses was reduced by 4 db. Refer to Section 4.6.1(ii).

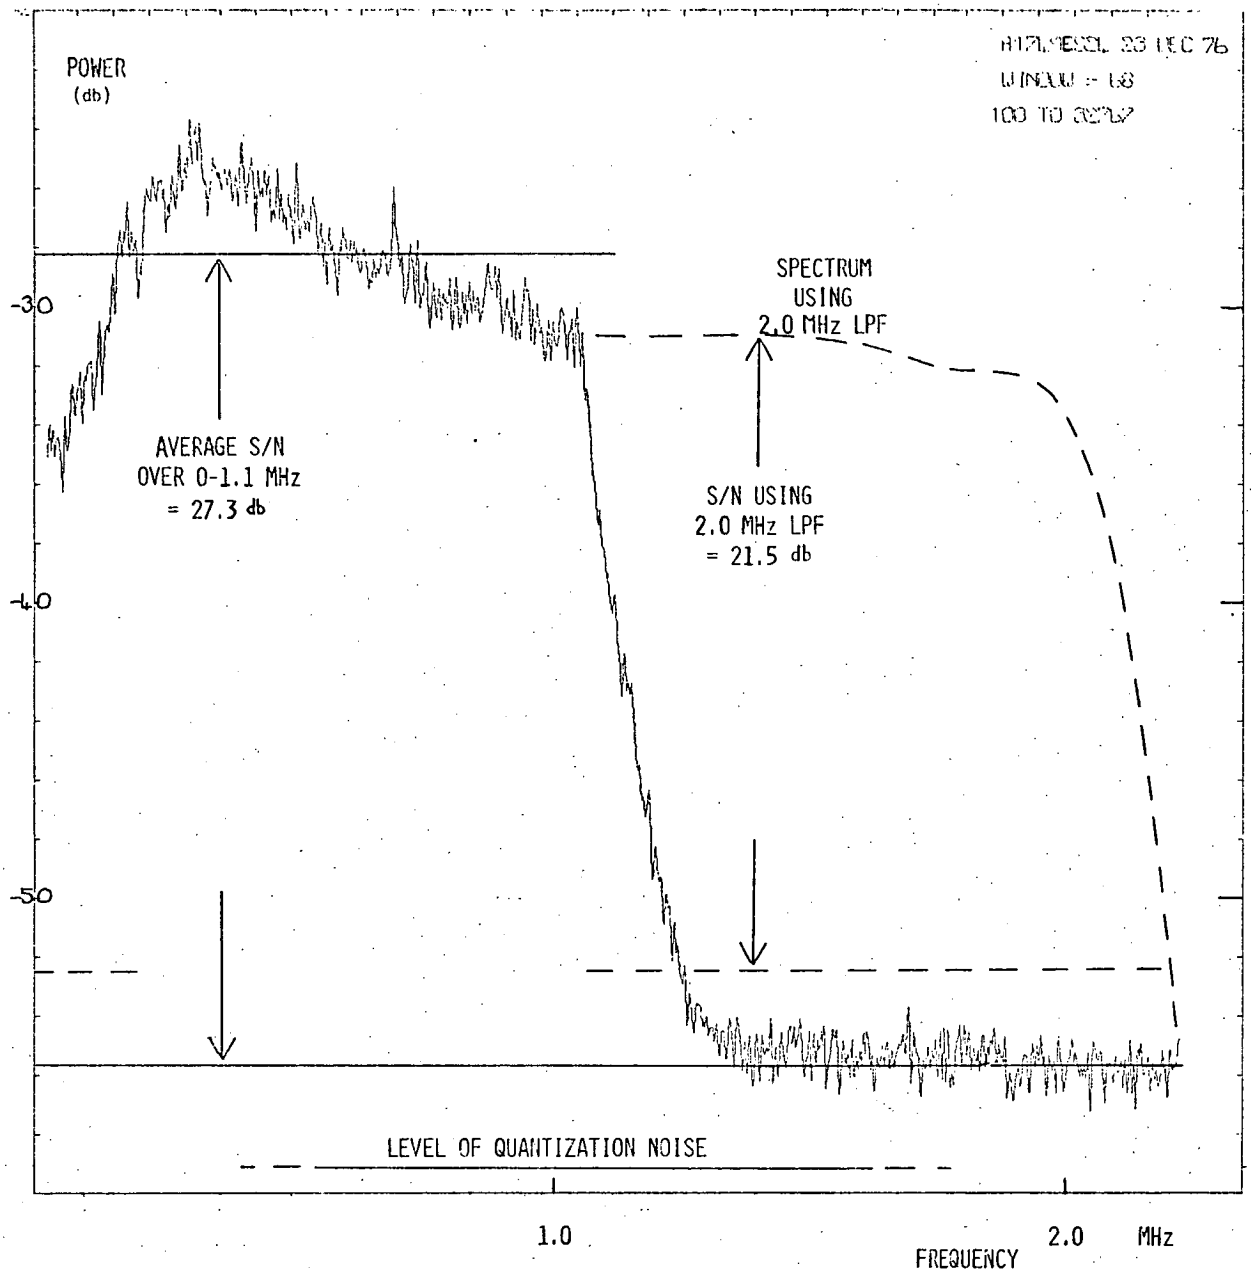

FIGURE 4.4.2A Power spectrum of a weak pulse from PSR 1749-28 sampled via a 1 MHz LPF to illustrate the flattening effect of interlaced sampling on the replay noise spectrum (above 1.35 MHz).

4.5 REPLAY CHASSIS

4.5.1 Buffer and Signal Splitter

The replay circuitry which first treats the video coming from the VTR is shown in Figure 4.5.1A. The video which is constrained by the VTR to be 1.4 Vpp into 75 Ω requires no amplification, however it is buffered by a unity-gain two transistor (complementary pair) amplifier before the signal splitter (three 50 Ω resistors). Logically a three-way split is required to separate the data, the reference tone and the sync pulses. However in the normal case of the built-in LPF (discussed in Section 4.5.6) being used this filter is common to the signal paths of the data and sync pulses. Therefore one of the three splitter outputs is normally available for connecting a monitoring device, usually the broad-band receiver, without affecting the rest of the system.

In normal use the signal from the built-in LPF feeds the sample-and-hold and the pulse discriminator section, the latter having high input impedence. There is an LM318 operational amplifier providing an inverting gain of 2 before the pulse discriminator proper (described in Section 4.5.3). There is some roll-off in the spectrum of this amplifier but the rise times of the pulses reaching the pulse discriminator are determined by the 2.0 MHz cut-off of the built-in LPF. The presence of this filter in the signal path from the VTR to the pulse discriminator significantly reduces the noise reaching the discriminator.

At the input to the buffer is a squelch system. This is used during the tuning of the phase-locked-loop, not during the actual sampling operation. It involves a junction field-effect transistor in series with the AC-coupled input to the buffer. There is a small but acceptable signal loss due to this FET which is normally in its conducting state.

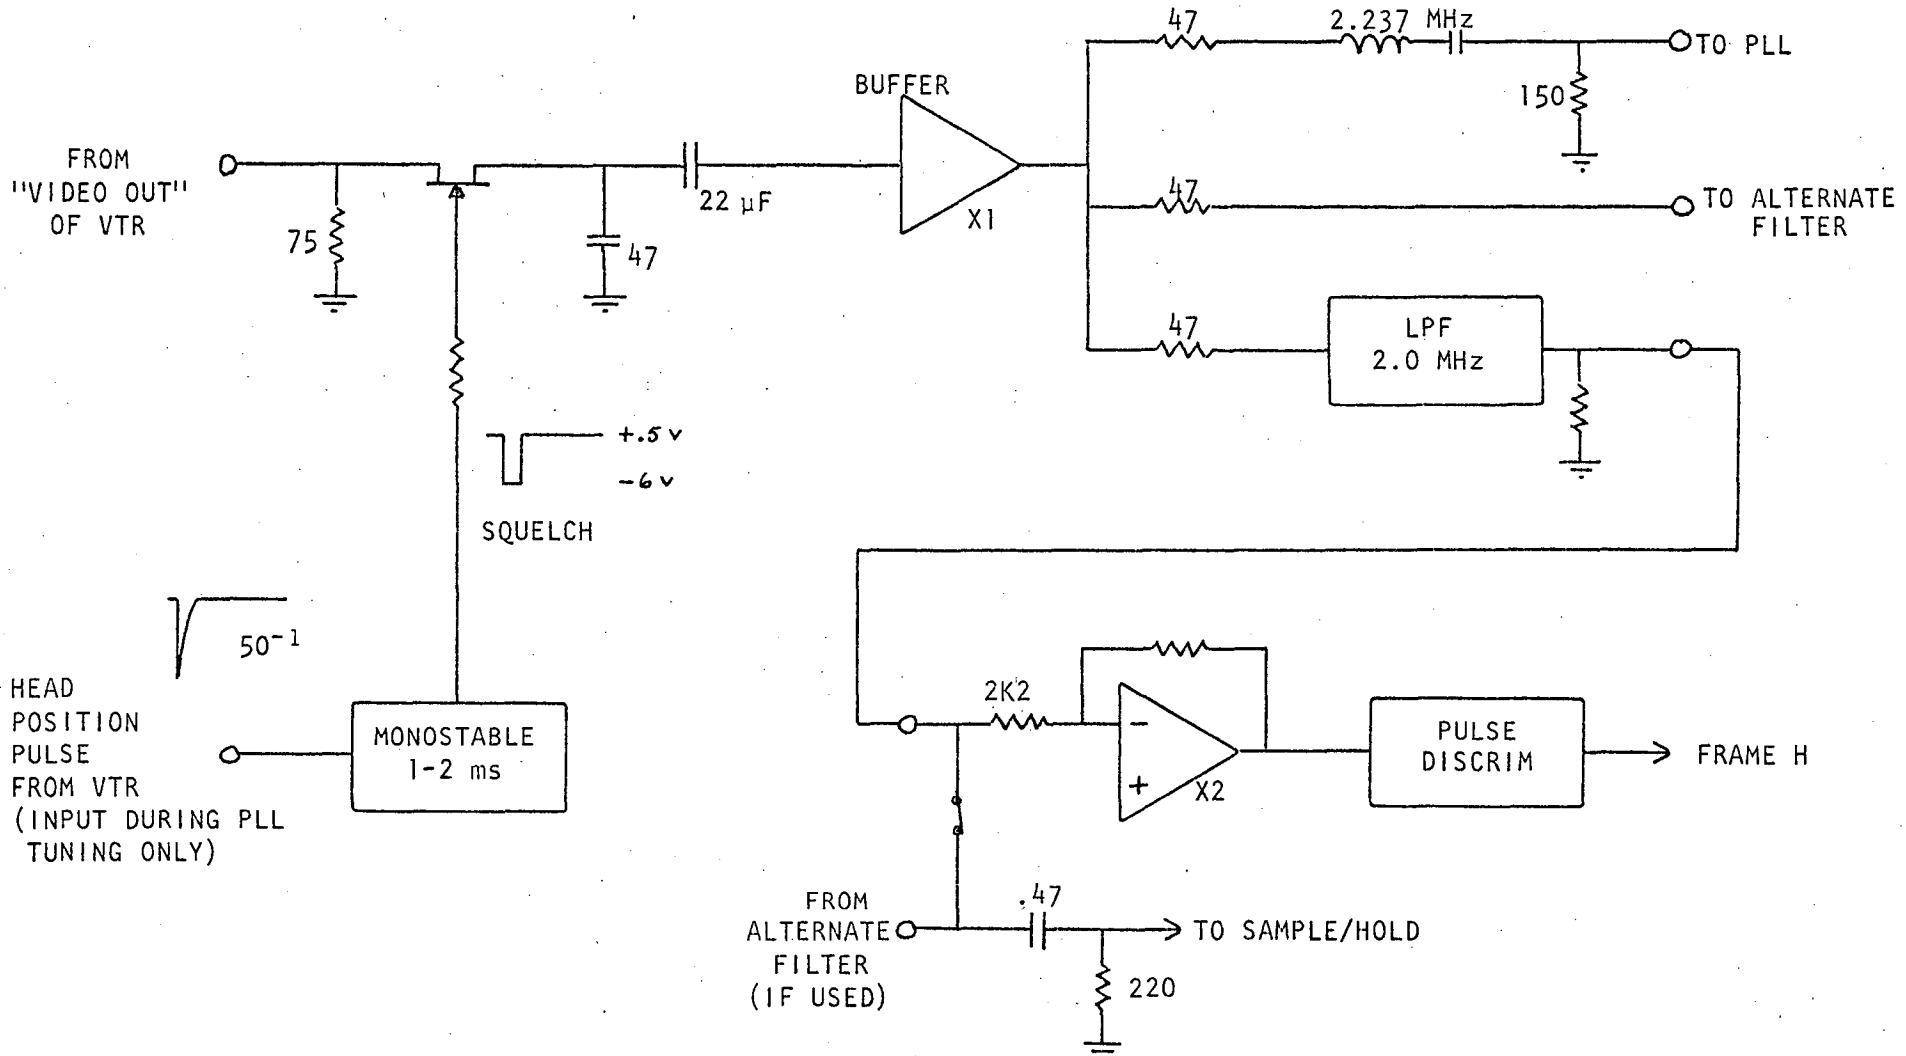

FIGURE 4.5.1A Simplified schematic of buffer and signal splitter.

4.5.2 The Phase-Locked-Loop

The purpose of the phase-locked-loop (PLL) in the replay system is to reconstruct the reference tone and hence define the clock ticks, the instants at which samples are taken. The PLL is used as a narrow band-pass filter to separate the reference tone from the data. In this role it has a number of special advantages:

(i) It is a tracking filter, coping with small variations in the VTR head rotation rate.

(ii) The output is a constant amplitude square-wave suitable for driving the following logic.

(iii) Phase is averaged over a time of the order of $1/B_L$ where B_L is the loop bandwidth. (I am using the notation of the Signetics Analog Data Manual 1977.) When the input is absent the PLL "remembers" the phase through a time at least $1/B_L$. The memory is longest when the free running frequency of the loop equals the reference.

(iv) It provides an indication at the phase comparator output of when lock has been lost.

The actual PLL used is the Signetics integrated circuit NE560B. A simplified schematic is shown in Figure 4.5.2A. The reference tone is capacitively-coupled into one side of the input to the phase comparator via a broad band-pass filter. The purpose of this BPF is to prevent the possibility of overloading the PLL with large amplitude signals outside the pass-band of the PLL.

The balanced output of the voltage controlled oscillator (VCO) within the PLL is delivered to a μA710 differential comparator via a pair of emitter followers. These buffers improve the mark-space ratio of the VCO output and hence of the 2.237 MHz square-wave produced by the μA710. The rising edges of this square-wave trigger one monostable with 110 ns duration, while the falling edges trigger another. The logical OR of the outputs of the two monostables forms an approximately square-wave at twice

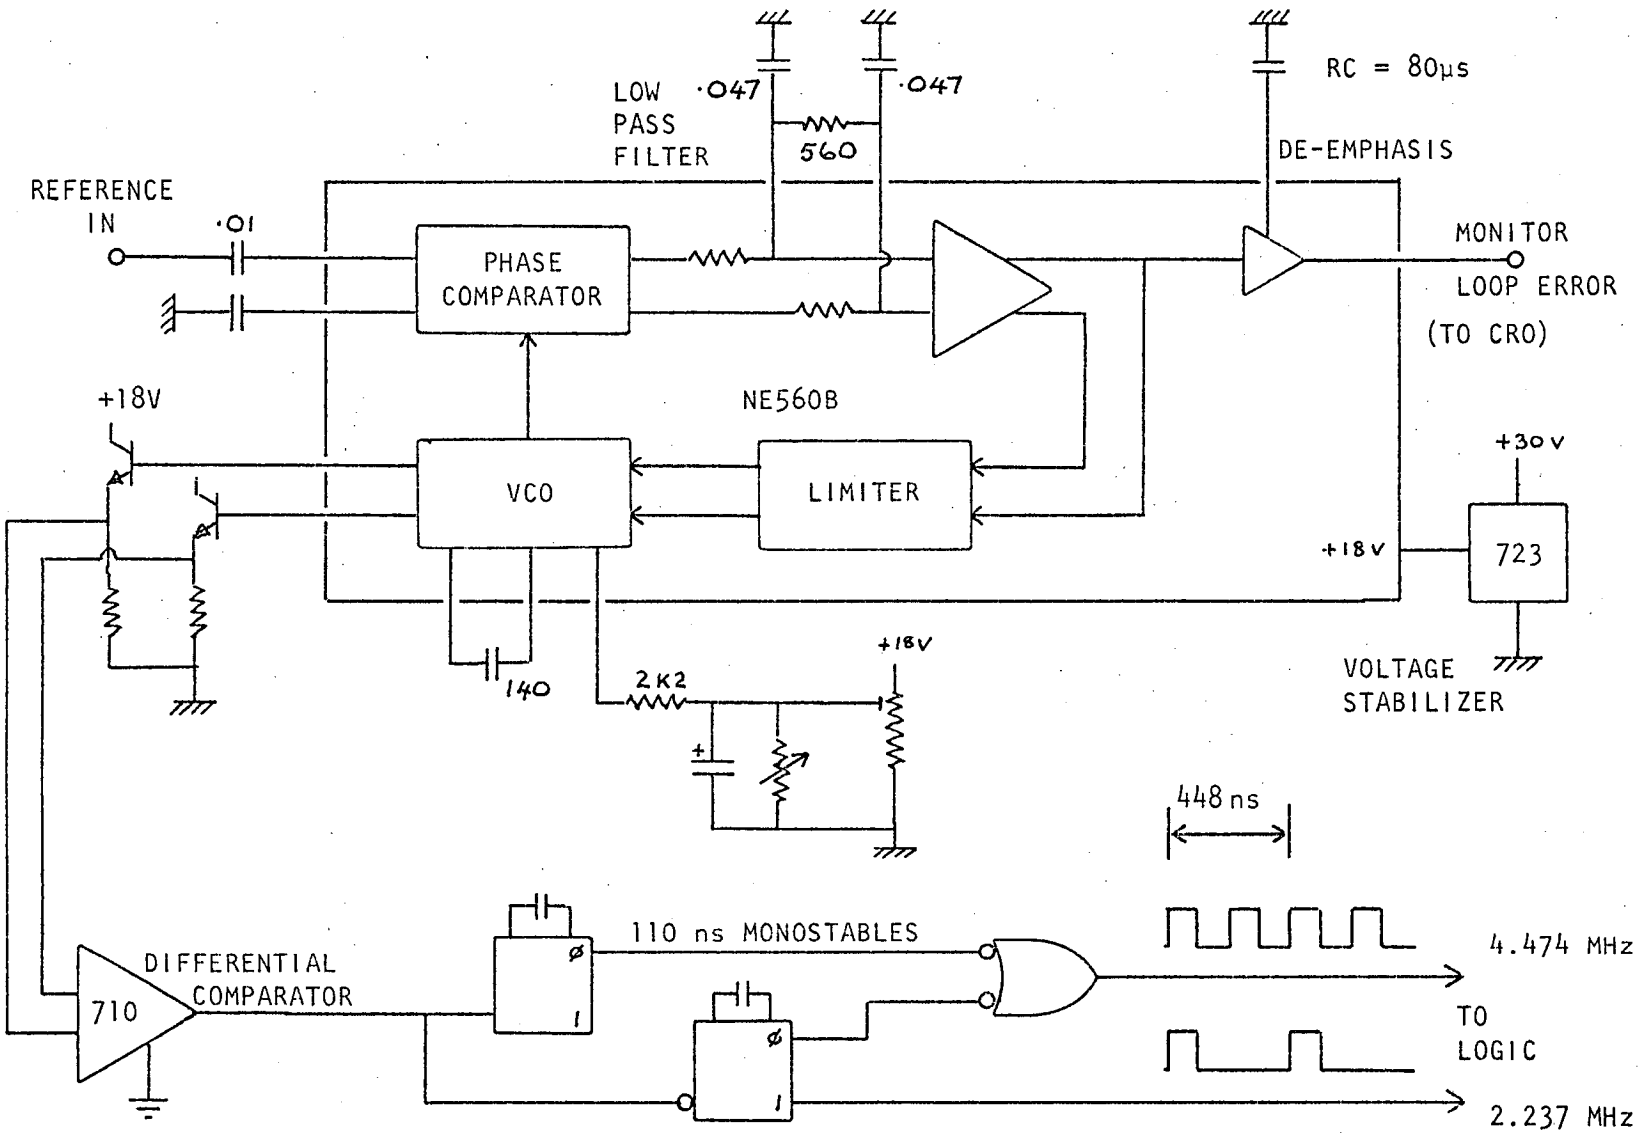

FIGURE 4.5.2A Phase-locked-loop

the reference frequency. The falling edges of this square-wave are approximations to the clock ticks. The errors in the approximation are of two types: alternate rising edges are not precisely equispaced, and there is phase jitter.

The first type of error, the spacing of alternate edges, is due to the method of doubling the reference frequency. Care was needed so that the mark-space ratio of the VCO output was near unity. The problem was not serious but it would have been more convenient to have used a divide-by-two stage within the PLL (but the NE562 was unavailable).

The presence of phase noise i.e. "jitter" is a much more complicated matter. Essentially noise in a bandwidth B_L about the reference frequency appears as phase noise in the VCO output. However there is also a contribution from loop instability. The following factors were found to be important:

(a) the signal-to-noise ratio of the reference;

(b) the loop filter;

(c) the free-running frequency;

(d) the level of the input to the PLL.

(a) The reference S/N ratio

There are two reasons for limiting the recorded amplitude of the reference tone:

(i) the possibility of intermodulation and other overload effects within the VTR; and

(ii) the fact that the reference tone itself is sampled along with the data because of jitter in the sampling comb.

If the sample jitter is reproducible the latter effect is not a problem since the result is a narrow-band noise signal in the vicinity of the Nyquist frequency f_N, in fact just below f_N. However if the jitter has a component which is not the same for each replay of the stripe then part of the energy of the (hypothetically) large reference tone is spread uniformly

over the spectrum because of the interlaced sampling scheme. One counter to this effect is a low-pass filter in front of the sample-and-hold. This filter is required because of other noise emanating from the VTR, however the tightest constraint upon it is to attenuate the reference tone.

Given that the reference tone could not be recorded at a very high level its S/N ratio is improved greatly by attenuating the data in its vicinity. This is done with the notch filter in the output stage of the recording chassis. There remain other sources of noise as discussed in Section 4.4.1.

(b) The loop low-pass filter

As shown in Figure 4.5.2A the loop low-pass filter consists of a pair of .047 μF capacitors and a 560 Ω resistor, giving a time constant τ of about 80 μs. The 560 Ω resistor dominates the internal 6 kΩ resistors in determining τ but its prime purpose is to reduce the total loop gain, in fact by a factor of 20. This particular configuration was found by trial and error, there being very little in the literature of practical guidance for using a PLL in the presence of noise. However the following few paragraphs give some explanation of why the particular low-pass filter configuration gives acceptable results.

When locked the PLL with a single pole low-pass filter as here acts as a second order feedback system. Such a system has a frequency response that is low-pass with a peak just below the cut-off. The frequency of this peak is known as the "natural frequency" ω_n given by:

$$\omega_n = \sqrt{K_d K_o / \tau}$$

where

K_d = gain of phase comparator (volt rad $^{-1}$),

K_o = gain of VCO (rad sec^{-1} volt^{-1}),

τ = time constant of loop low-pass filter (sec).

It is assumed that the gains of other parts of the loop are lumped in with K_o. Note that K_d is proportional to the input reference amplitude for small amplitudes. The damping factor of the loop which determines the height of the peak at ω_n is

$$\zeta = 1/\sqrt{4\ K_d\ K_o\ \tau}$$

The noise bandwidth of the loop which determines the amount of jitter is of the order of $\omega_n\ /\ \pi$. Thus a low loop gain K_dK_o or a high τ should reduce the noise bandwidth. The former, a low loop gain, is preferable as it increases ζ if a major contribution to jitter is being made by noise-driven instabilities.

Substituting into the equations the manufacturer's typical values of K_o and K_d together with the actual loop low-pass filter parameter gives

noise BW ~ 3 kHz

ζ ~ 0.5

assuming 0.1 mV rms reference input.

(c) The free-running frequency

Keeping the free-running frequency of the VCO as close as possible to that of the reference is the most important factor in satisfactory operation of the PLL in this application. It allows the loop gain and hence the noise bandwidth to be reduced without regard to capture and lock range. With the loop filter described above capture at the start of each stripe takes place satisfactorily i.e. before the trailing edge of the first sync pulse.

To keep the frequency of the replayed reference near constant the VTR rotor is synchronized to a crystal-derived 50 Hz signal, just as it is during recording. The replayed frequency varies a little within each stripe according to the characteristics of the servo loop controlling the rotor. It also varies between stripes at different parts of a videotape

according to the actual slant-track geometry used during recording and replay. Note that for the average slant angle the actual frequency of the reference during stop-action replay is ~ 2.5% higher than 2.237 MHz, and a 20 ms stripe takes 19.6 ms to replay. However I ignore this frequency shift everywhere else in this thesis.

Tuning the VCO

The VCO is coarsely tuned by a fixed capacitor (140 pf) and finely tuned by current injection into the VCO from a potentiometer with a front panel knob. The PLL error signal is brought out to the front panel for monitoring with an oscilloscope (AC coupled). Figure 4.5.2B shows sketches of oscillograms obtained during the following tune-up procedure:

1. warm up all equipment for 30 minutes,
2. run the VTR rotor up to speed (20-30 sec.),
3. tune the VCO into the locked range as indicated by having error spikes either side of that tuning position,
4. using the 50 Hz head position pulse from the VTR, squelch the input during the start of the stripe and tune the VCO so that there is no step in the error signal. At this stage the free-running frequency (as indicated by the DC level of the error signal during the squelch time) equals the reference frequency,
5. disable the squelch for normal operation.

(d) The level of the input to the PLL

It was found that least jitter occurred when the level of the reference tone into the PLL was small. This level reduction was made by fixed pads at the PLL input (not of course by recording the reference tone at a lower level). The effect of reduced input level is to reduce K_d and hence the total loop gain and the noise bandwidth and to raise the damping factor as described in (b) above.

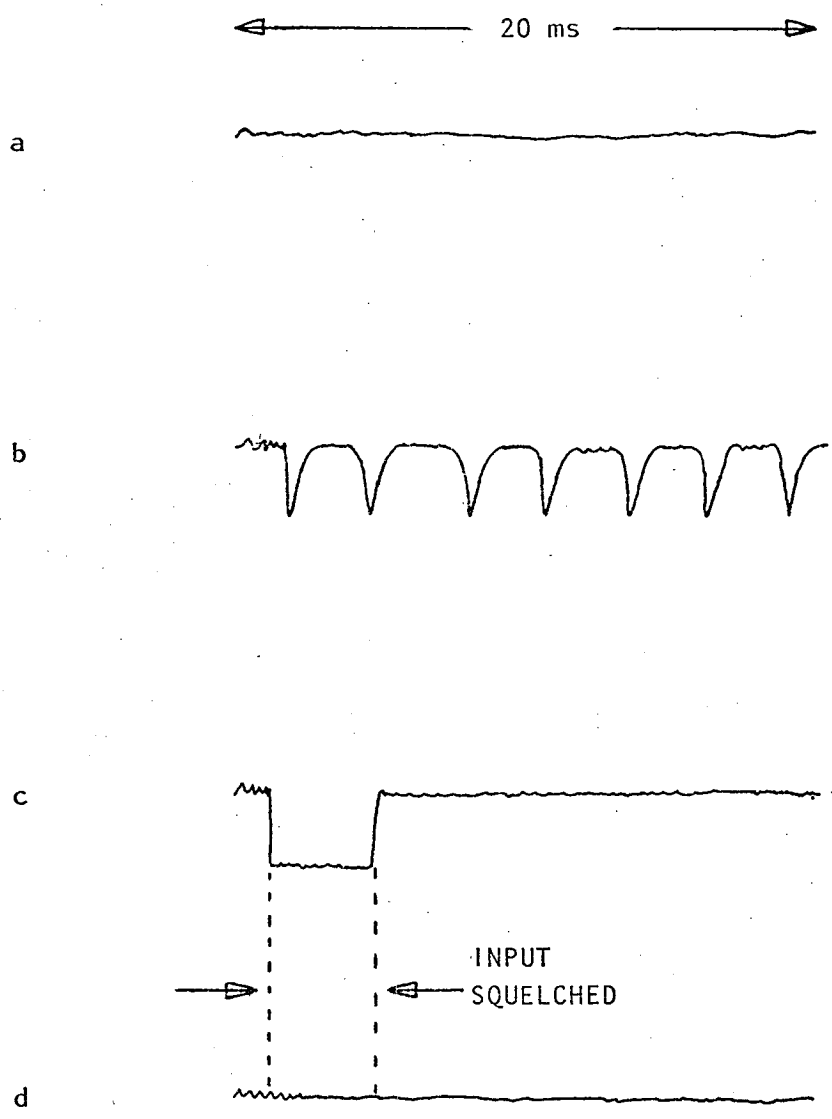

FIGURE 4.5.2B

Sketches of oscillograms of the PLL error signal as the free-running frequency, f_o, is adjusted towards that of the reference, f_r.

(a) f_o far from f_r, loop unlocked

(b) f_o nearer f_r, loop partially locked

(c) f_o near f_r, loop locked, input squelched near start of stripe

(d) $f_o = f_r$, still squelched as in (c).

This method of adjusting the loop characteristics depends on the free-running frequency being adjusted correctly. If this were to drift sufficiently away from the reference frequency, due to equipment warm-up for example, then lock would be lost. However before that stage there would be a significant shift in the phase of the VCO relative to the reference (this is 90° normally). The size of this shift for a given frequency difference ("loop strain") is inversely proportional to K_d and hence to reference level. A phase shift away from 90°, say by 45°, would have severe effects on the interlaced sampling particularly near clock ticks that are multiples of 256 from the zeroth. Therefore having the reference input at a low level places tight limitations on the amount of drift that can be tolerated in the free-running frequency during the time taken to completely sample one stripe, namely 2 to 5 minutes. Hence the standard practice of warming up the equipment for half an hour before sampling.

4.5.3 Pulse Discriminator

The purpose of the pulse discriminator is to identify as accurately as possible the occurrence of the first sync pulse. Since its trailing edge marks the start of the data part of each stripe, it is important that no mis-identifications occur. The main output of this circuit is a logic signal, FRAME H. The zeroth clock tick is, less a constant time displacement, the first edge of the 4.474 MHz square-wave produced by the PLL which follows the rise of FRAME H. A secondary function of the pulse discrimator is to produce logic transitions at each end of the second sync pulse. These transitions are used by the error detection system.

The pulse discriminator circuit is shown in Figure 4.5.3A. The incoming video is compared to an adjustable voltage by a μA710C voltage discriminator. The output of this (called SYNC L) passes to a pair of retriggerable monostables in series which produce a signal GATE H centred on the expected fall of the first sync pulse i.e. the rise of SYNC L. The

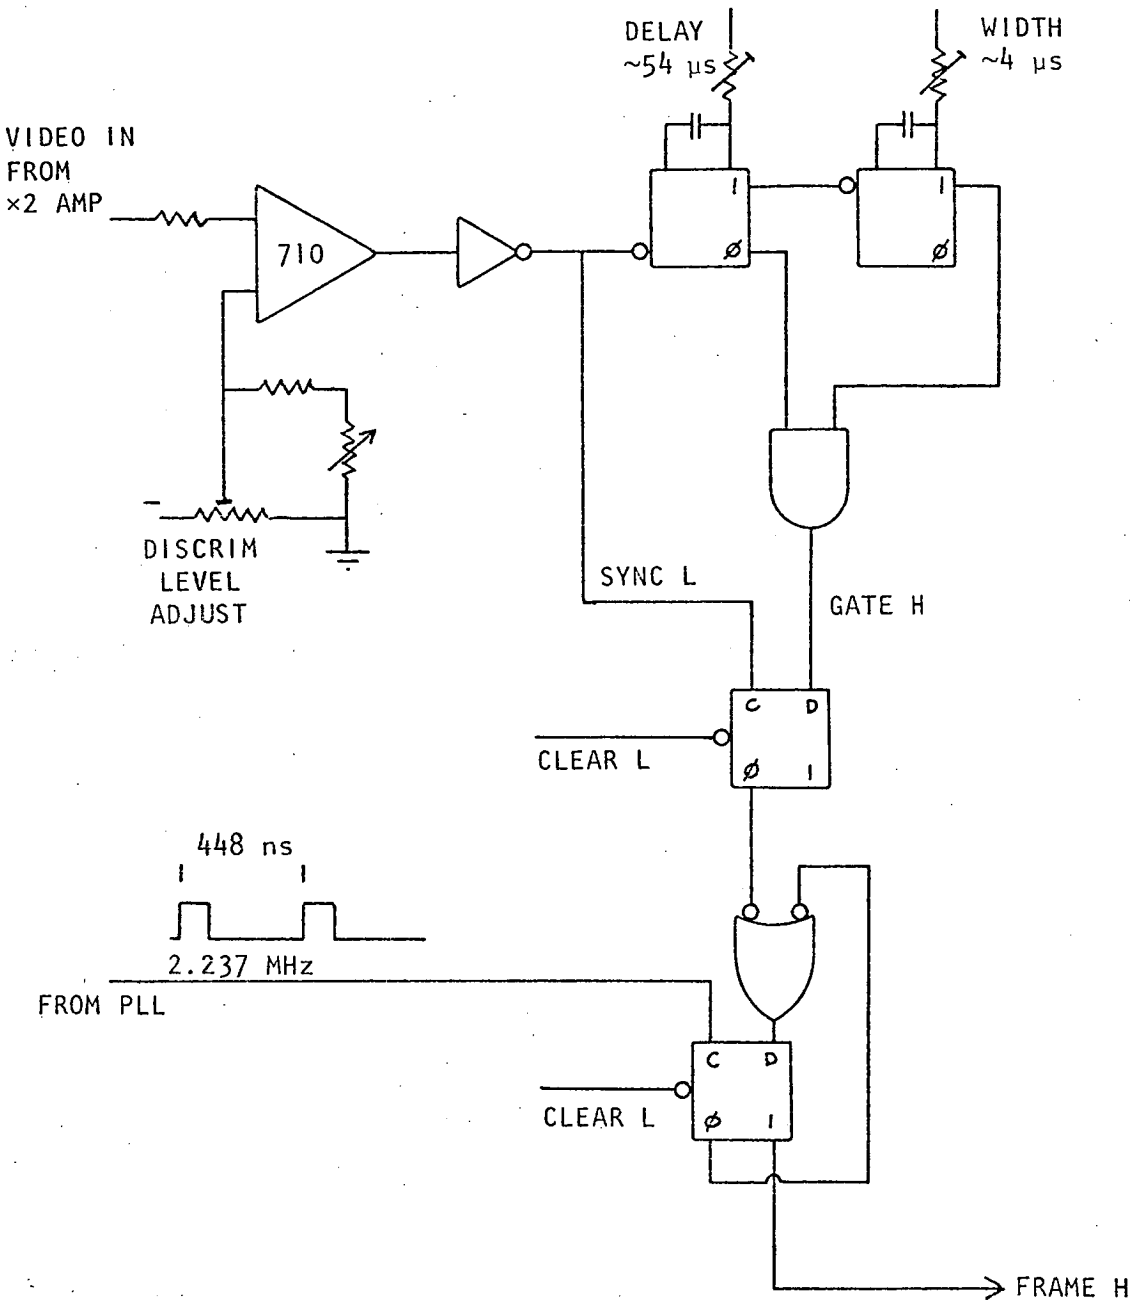

FIGURE 4.5.3A Schematic of pulse discriminator.

occurrence of this rise together with GATE H clocks a D-type flip-flop and a second D-type is clocked on the next rising edge of the reference tone as reconstructed by the PLL. The timing diagram for these signals is Figure 4.5.3B.

The timing diagram shows how the pulse discriminator is sensitive to the first sync pulse. The video is required to rise and stay up for 56 ± 2 μs (assuming the delay and width monostable times set as shown). Because the first monostable is retriggerable the noise burst which invariably precedes the first sync pulse at the start of each stripe does not affect the operation. Furthermore the second sync pulse is not confused with the first because of its different duration.

The ability to adjust the reference voltage for the discrimination is useful for the stop-action replay of some stripes. Some stripes have noise superimposed on the top of the first sync pulse (this also occurs when the tape tension or position is incorrect) requiring the reference voltage to be reduced. Also the reference voltage adjustment can choose which of two possible cycles of the reference tone is the zeroth. This is because the fall time of the sync pulses is ~ 500 ns while the reference period is 448 ns. The reference voltage is adjusted so that noise on the trailing edge has least tendency to cause a jump in the identified zeroth clock tick. This jump is by 2, i.e. 2 × 224 ns.

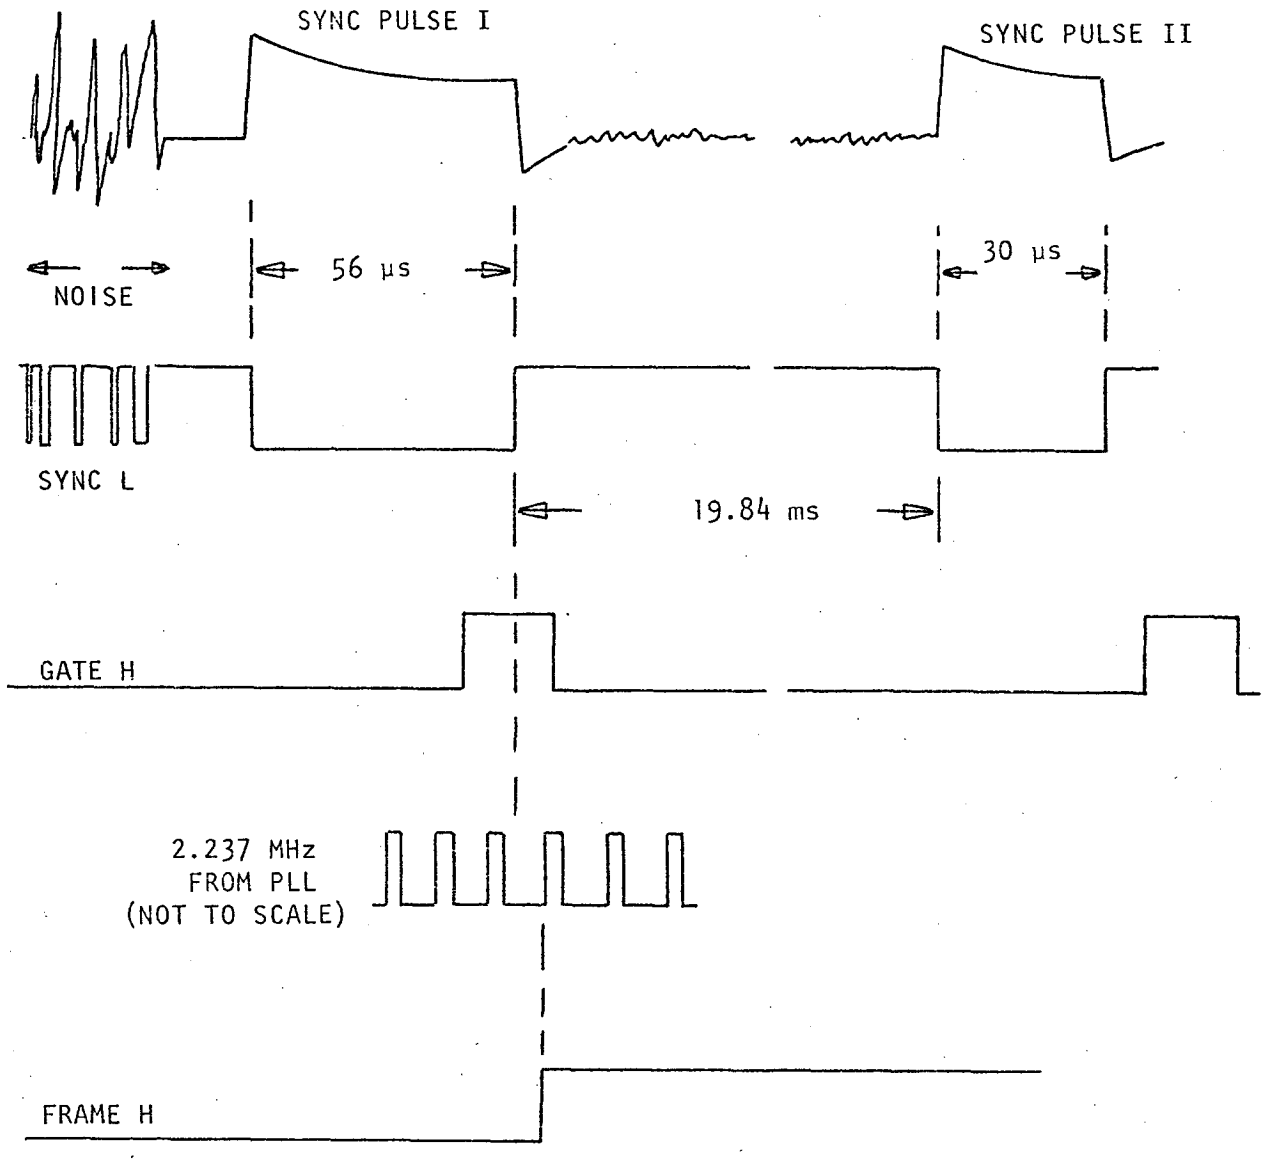

FIGURE 4.5.3B Timing diagram for pulse discriminator.

4.5.4 Control Logic

The function of the control logic circuitry is to produce pulses which cause sampling and digitizing of the data, all under supervision by a computer. A simplified schematic appears as Figure 4.5.4A while the associated timing diagram is Figure 4.5.4B.

The pulse HOLD H causes the fast sample-and-hold to "hold" the data for its 9 μs duration. The leading edge of this pulse is the instant of sampling, supposedly on an ideal clock tick. Actually the sample is an average over a window extending perhaps up to 50 ns after the leading edge. After a settling time of 3 μs the leading edge of the CONVRT H pulse triggers the slow sample-and-hold in the computer into taking a copy of the data sample and digitizing it. The controlling computer program is informed of the completion of digitizing say 25 μs later by the analog-digital (A/D) converter.

Successive HOLD H pulses start (ideally) exactly $256 = 2^8$ clock ticks apart, i.e. 256 × 224 ns = 57 μs apart. Each such sampling instant is a constant time interval after a falling edge of the 4.474 MHz square wave produced by the PLL. This time interval is propagation delay mainly in the 8-bit ripple counter and amounts to a few tens of nanoseconds. Neglecting this delay the first falling edge after the rise of the signal FRAME H produced by the pulse discriminator represents the zeroth clock tick. Clock ticks 0, 1 and 2 are marked in the timing diagram. The rise of the first HOLD H pulse follows (256 - j) ticks later where j is count pre-set in the 8-bit ripple counter before the start of FRAME H. The time t marked in Figure 4.5.4B is (256 - j) × 224 ns. The count j is kept in the 8-bit counter shown lower in the diagram. Count j is incremented by the controlling computer program to advance the sampling comb when the current set of samples are satisfactory.

The number of samples taken in one replay of a stripe, i.e. in one

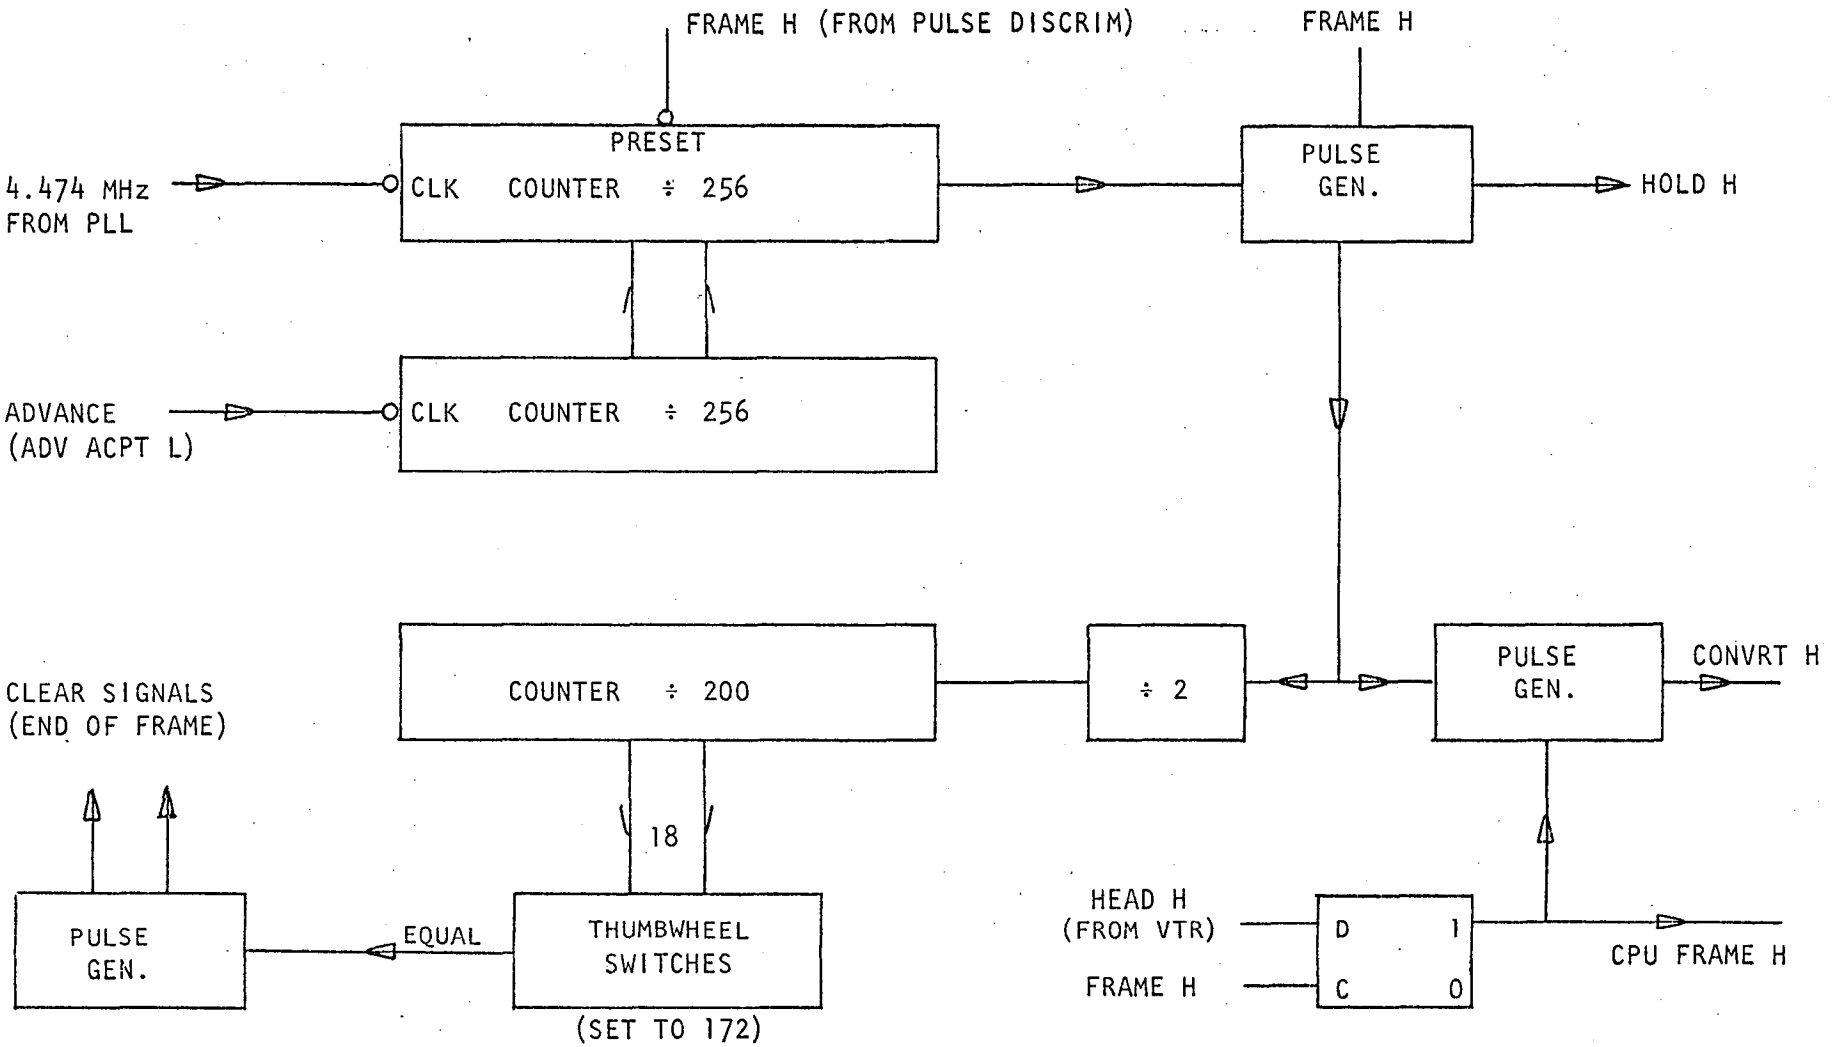

FIGURE 4.5.4A Simplified schematic of the control logic.

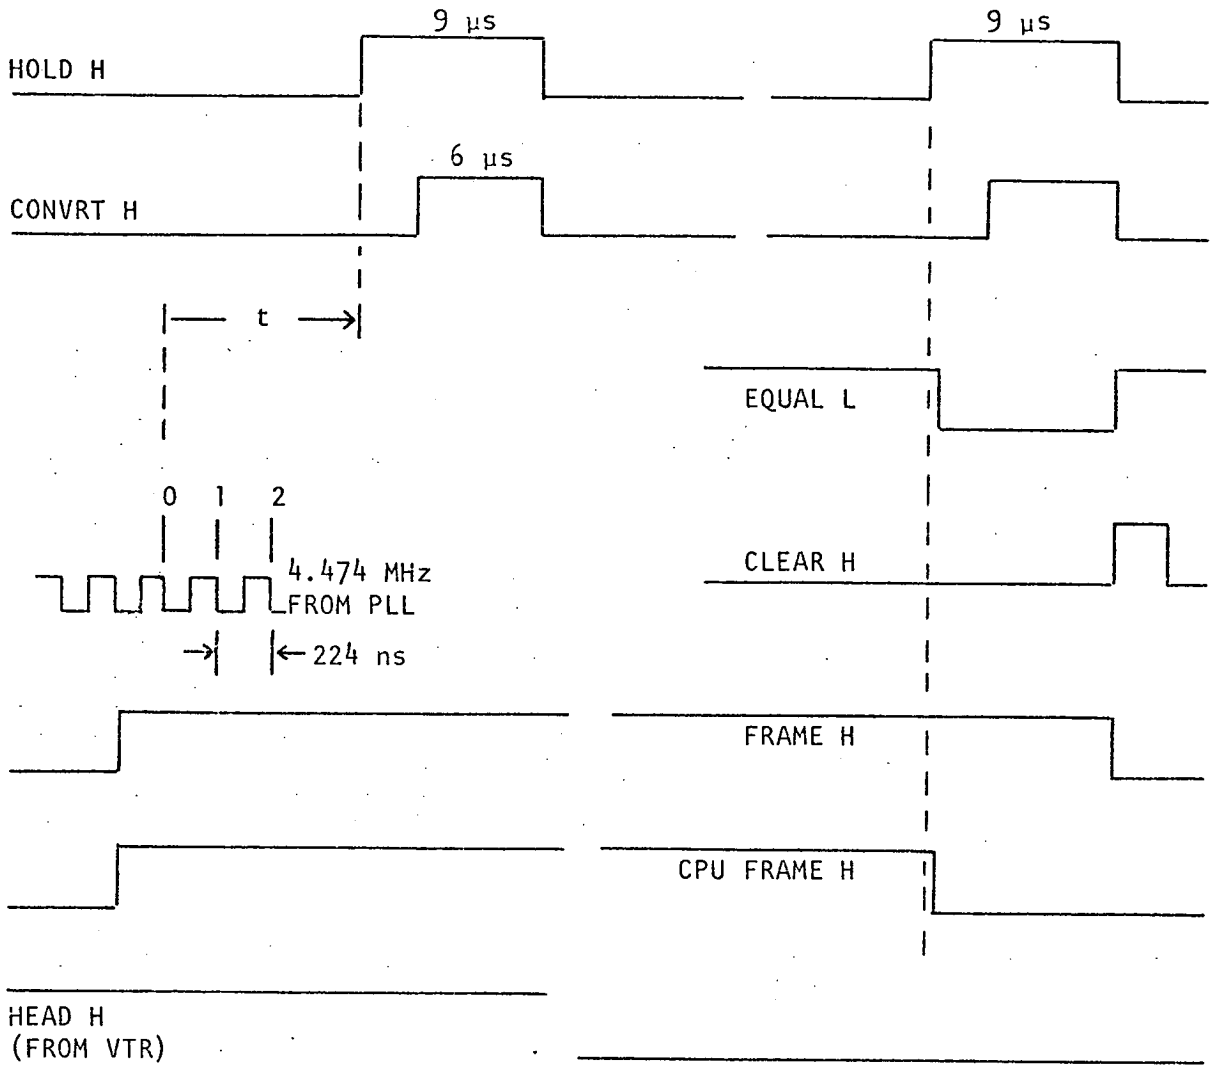

FIGURE 4.5.4B Simplified timing diagram for control logic (only the first and last hold and convert pulses are shown).

"frame", is determined by a front-panel thumbwheel setting. Usually this is set to 172 so that there are 344 samples in the comb i.e. obtained in each frame. This gives a complete coverage of the 19.84 ms of data in a stripe. The signal FRAME H which is ubiquitous throughout the control logic is cleared at the end of the 344th HOLD pulse. A derived signal CPU FRAME H is sent to the computer to indicate that sampling is underway. It gives a warning up to 80 μs before the first A/D conversion is complete. This signal differs mainly from FRAME H in that it is only asserted for one of the two VTR heads. The CONVRT H pulses are suppressed for the other head.

Many of the details left out of Figure 4.5.4A are concerned with initiallization of the circuit. There are also some interlocks on the "advance" pulse sent to the lower 8-bit counter. These are concerned with latching the pulse from the computer and ensuring that each "advance" request gets obeyed once only.

4.5.5 Error Detection Logic

The function of the error detection logic circuitry is to check that the PLL has remained in lock throughout each replay of a stripe. It does this by counting the cycles of the reconstructed reference tone between the trailing edges of the two synchronization pulses and comparing this count (modulo 128) with a reference count. If there is a mismatch then:

(i) the controlling computer program is told via a signal WAIT H;

(ii) the operator is told by the illumination of a light;

(iii) any "advance" request pending is kept pending;

(iv) the error detection logic is frozen until the operator or the computer clears the error state;

(v) the least significant three bits of the offending count is displayed by lights.

A greatly simplified schematic and the associated timing diagram appear as Figures 4.5.5A and 4.5.5B.

An 8-bit count of the rising edges of the 2.237 MHz signal from the PLL is compared with OXOOOSSS (binary) where X denotes "don't care" (actually 1) and SSS, the three least significant bits, are set up on toggle switches by the operator. The ability to vary these bits is mainly to cope with variations in the duration of the second sync pulse due to drift in a monostable at record time i.e. on different days. However it is also required by small variations in the trailing edges due to noise, different settings of the pulse discriminator reference voltage, and the use of different VTRs for recording. To make setting the switches practical the 3 least significant bits of the count are latched at the instant of comparison and displayed on LEDs. This display is updated every 40 ms, only one of the two VTR heads being used.

The main complications of the circuit are concerned with determining the instant of comparison, viz. of finding the trailing edge of the second sync pulse. Obviously the best approach to this problem would be to use

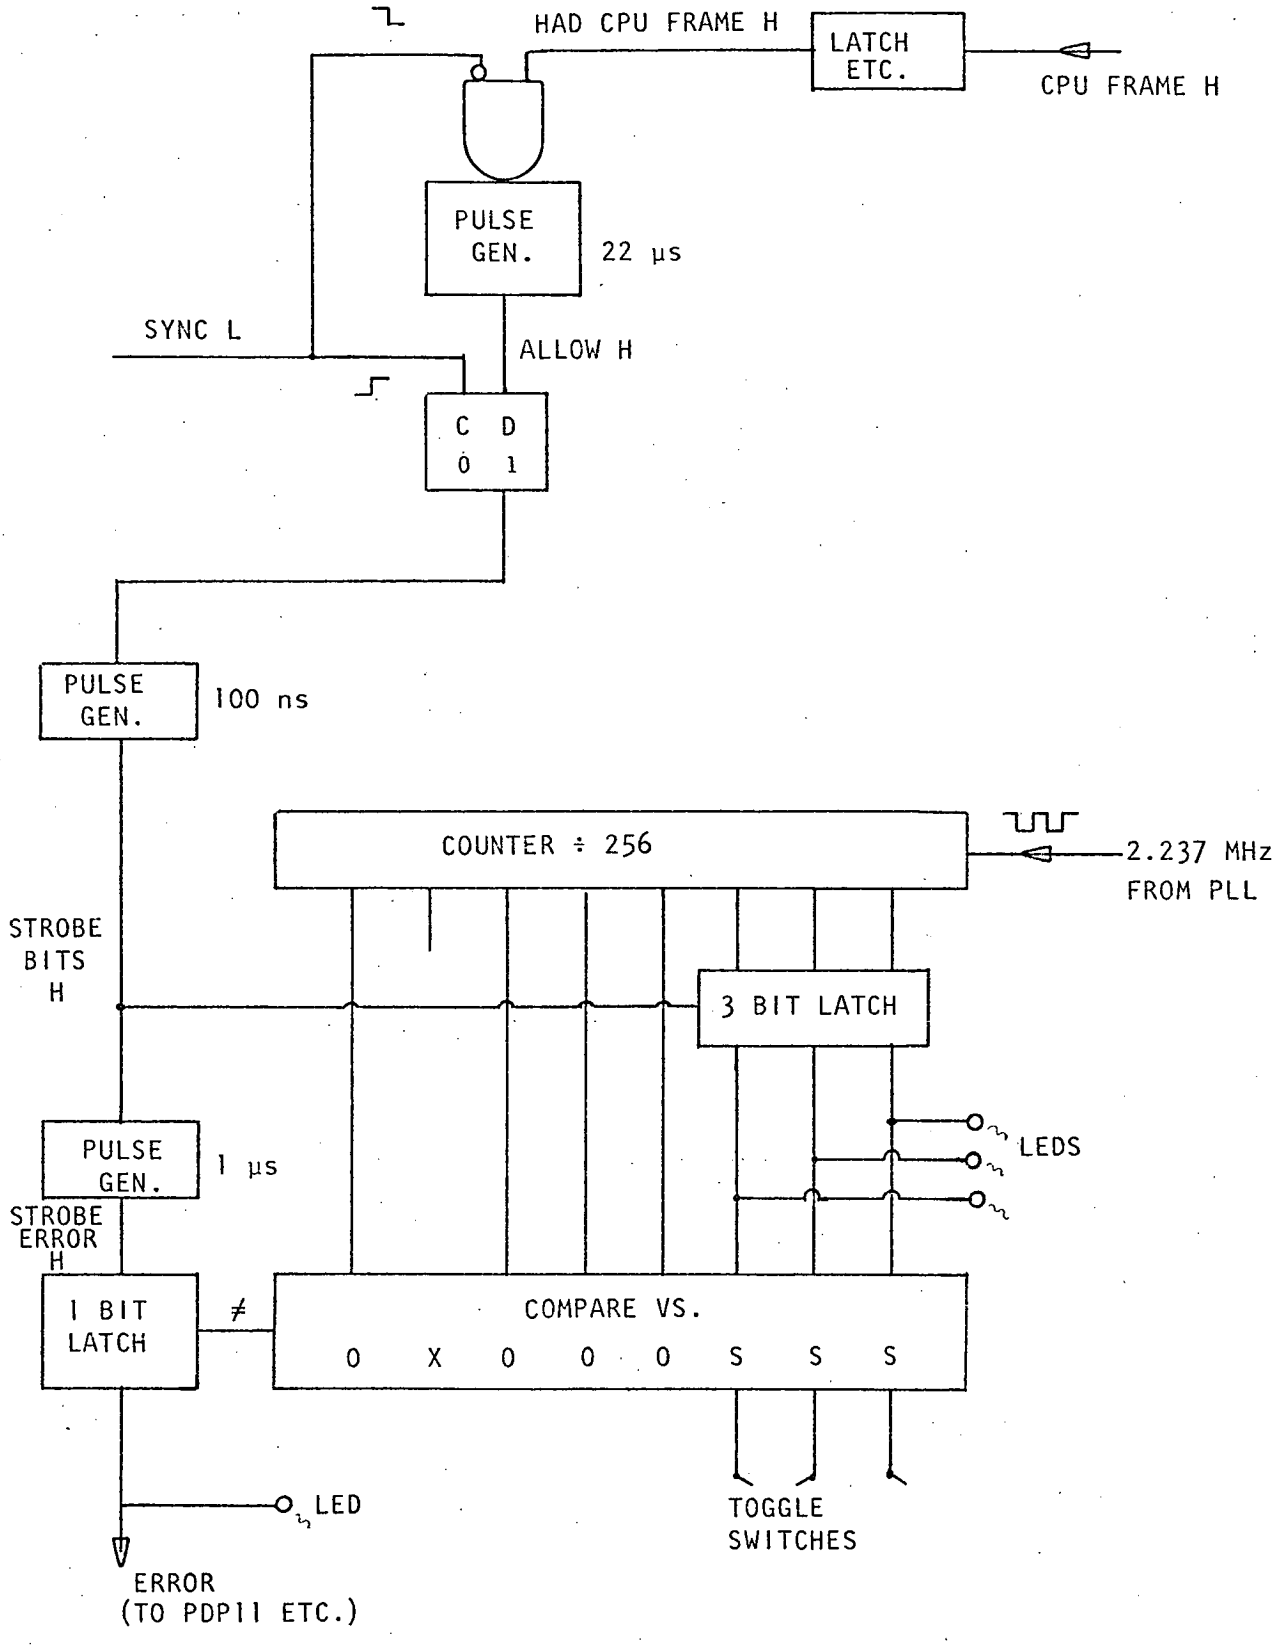

FIGURE 4.5.5A Simplified schematic of error detection logic.

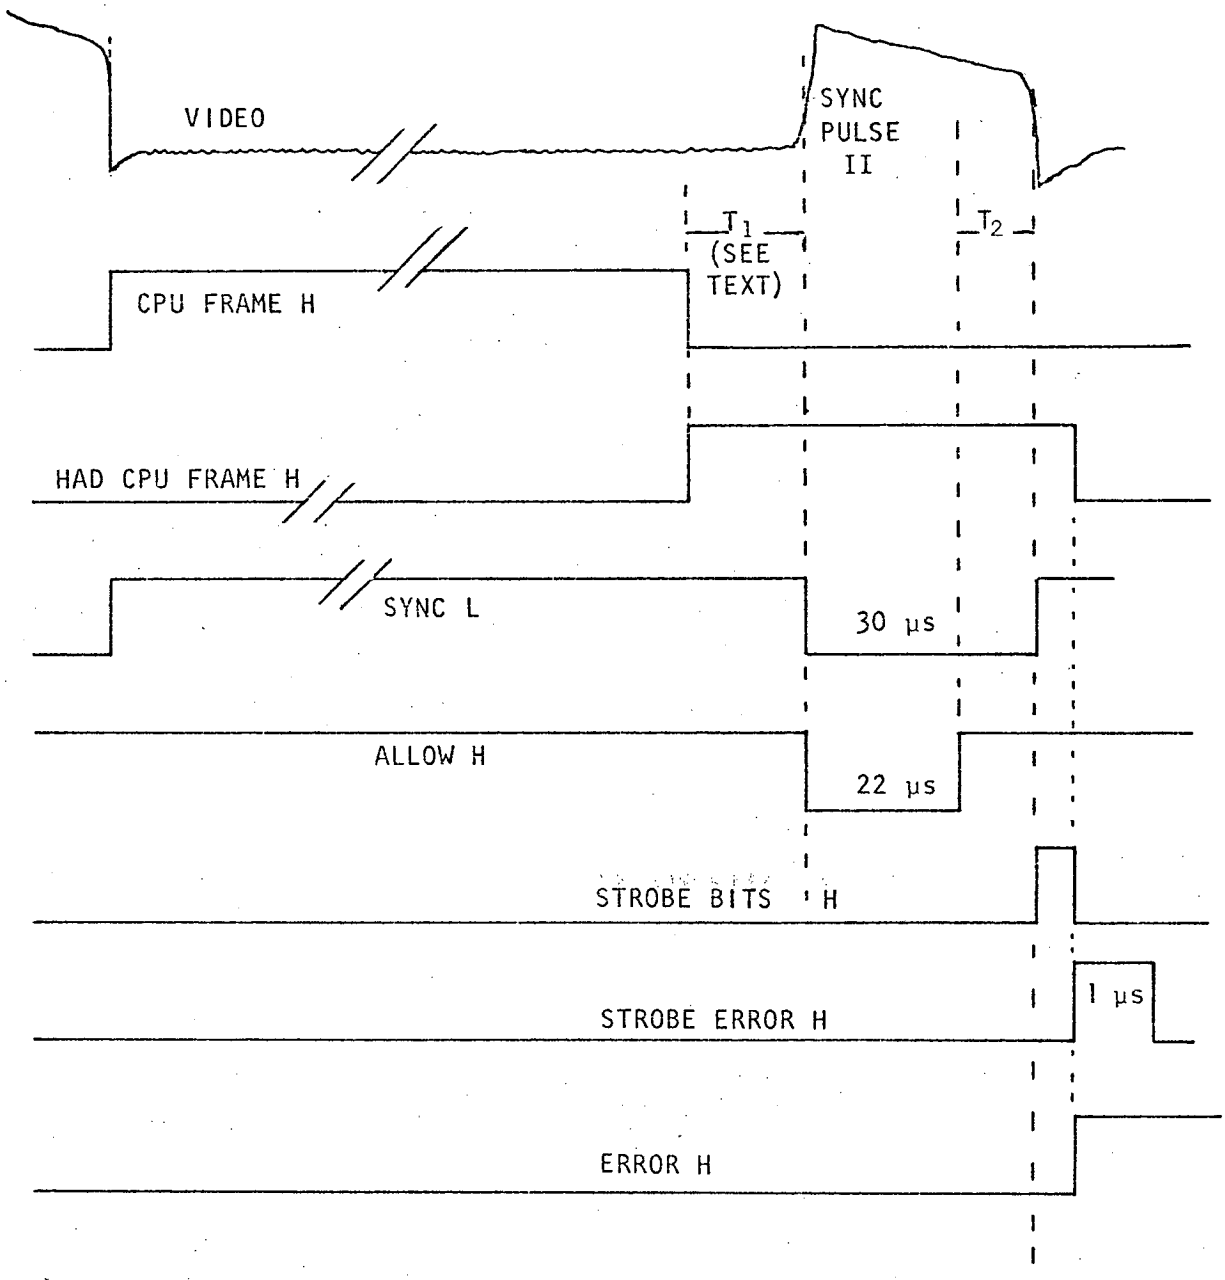

FIGURE 4.5.5B Timing Diagram for Error Detection Logic.

the same circuits used to record the sync pulses but clocked by the output of the PLL instead of the original crystal oscillator. However the approach adopted for "historical" reasons does work although it is not entirely satisfactory. It is based on the supposition that the first pulse after CPU FRAME H falls, i.e. after the control logic thumbwheels' count is satisfied, is the second sync pulse. The method is overly sensitive to rising noise spikes during time interval T1 (see Figure 4.5.5B) and to falling spikes in interval T2. Note that there is no need to be as careful about mis-identification of the second sync pulse as of the first. However in any design T2 should not be less than a few microseconds. In this case it is 8 μs assuming no noise spikes occur in T1. The duration of T1 varies from 57 μs to 0 depending on the stage the interlaced sampling process has reached.

The recovery from an error can be made either by the controlling computer or by the operator. The former does so by asserting a signal called FIX, the latter by switching a front-panel toggle to MANUAL and then, possibly after a minor tape tension adjustment, back to AUTO.

4.5.6 Low-Pass Filter

A low-pass filter (LPF) is needed before the fast sample-and-hold

(i) to reduce corruption of the data by aliasing,

and (ii) to reduce the total amount of non-reproducible noise.

Sampling causes signals at a frequency (f_N + f) where f_N = 2.237 MHz is the Nyquist frequency to appear as an alias at f_N - f. Therefore if the input to the LPF had a white spectrum then the attenuation at f_N + f would equal the relative level of the alias noise at f_N - f (assuming this is the corner frequency of the LPF). However the data itself is recorded through the 2.06 MHz LPF giving an attenuation of 17 db at 2.474 MHz (alias 2.000 MHz). Thus one of the functions of the LPF preceding the fast

sample-and-hold is to improve this attenuation.

For the second function, that of reducing non-reproducible noise, the LPF has no assistance from filters in the recording chassis. Due to the interlaced sampling scheme non-reproducible noise is spread evenly across the spectrum 0 - 2.237 MHz. A general reduction in this noise is obtained by reducing the pass-band of the LPF, e.g. to cut off at 2 MHz if the data is only to be analyzed below 2 MHz. Because the reproducible noise above 2.237 MHz is already 10 to 20 db down relative to the data at 2 MHz, the required attenuation is not great. However the major signal above 2 MHz is the reference tone and when sampled with a sampling comb containing non-reproducible jitter is itself a source of non-reproducible noise (spread uniformly). This suggests a sharper cut-off than the consideration of aliasing. However against this must be balanced the need for a moderate phase derivative, that is, for no large excursion in group delay below 2 MHz.

The 2.0 MHz LPF built into the replay chassis has Cauer-Cauchy design CC 05 25 53 (Zverev 1967). Its design properties are:

pass-band ripple (0 - 2.0 MHz) : 0.28 db

stop-band attenuation (above 2.5 MHz) : 36 db

group delay variation across pass-band : 1.0 μs.

Measured attenuation (installed) is 1 db at 1 MHz, 2 db at 2 MHz, 12 db at 2.237 MHz and 24 db at 2.474 MHz (see Figure 5.3A). The performance is poor due to a mismatch at the output (75 Ω not 50 Ω) but adequate, except possibly for the attenuation of the reference tone. An additional notch filter should have been incorporated in the replay chassis. The total response of all the filters in the equipment is discussed in Chapter 5.

Two alternative filters were built both with a 1 MHz cut-off. One is a "Gaussian to 6 db" design (Zverev 1967) with a 3 db attenuation and an almost linear phase spectrum below 1 MHz. The other (CC 05 20 47, op cit) has a sharper cut-off being 35 db down at 1.35 MHz but a large group delay variation near 1 MHz. Both filters have been used for testing and for

digitizing pulses from high dispersion pulsars in which the pulse above 1 MHz crosses a stripe boundary i.e. leaves the observing window.

4.5.7 Fast Sample-and-Hold

The fast sample-and-hold has to take samples of the data waveform with an aperture, defined as the time over which the held voltage represents an average of the data waveform, that is much less than δt = 224 ns. It does not have to hold the sample voltage for more than a few microseconds since its output is sampled by the sample-and-hold belonging to the A/D convertor. This second sample-and-hold maintains its sample voltage constant, i.e. within 1 mV, for the duration of A/D conversion, i.e. for 20 μs with the PDP11/20.

The circuit and a sketch of the output waveform is shown in Figure 4.5.7A. Its performance depends strongly on that of the National Semiconductor AM1000 "analog switch". During the "hold" state, viz. the time interval in which logic signal HOLD H is asserted, the sample is the charge on the .001 μF capacitor. This charge is isolated by the "off" resistance of the AM1000 and by the high input resistance of the LH0042 FET-input operational amplifier. The output of the sample-and-hold, i.e. the output of the LH0042, shows a switching transient with a settling time of about 1 μs at the start of the "hold" state. Therefore the control logic waits 3 μs before triggering the computer's sample-and-hold. There is also a constant offset of 220 mV between the "following" and the "hold" state. This offset is cancelled out in the LH0042 stage so that the digitized samples have an average of zero volts.

The frequency response of the LH0042 rolls off above 1 MHz so that the output waveform in the "following" mode is a smoothed version of the waveform being sampled. This low-pass response does not effect the sampling itself, the settling time being satisfactory. However the frequency response of the LH0002 buffer is important : it is flat to well above 2 MHz.

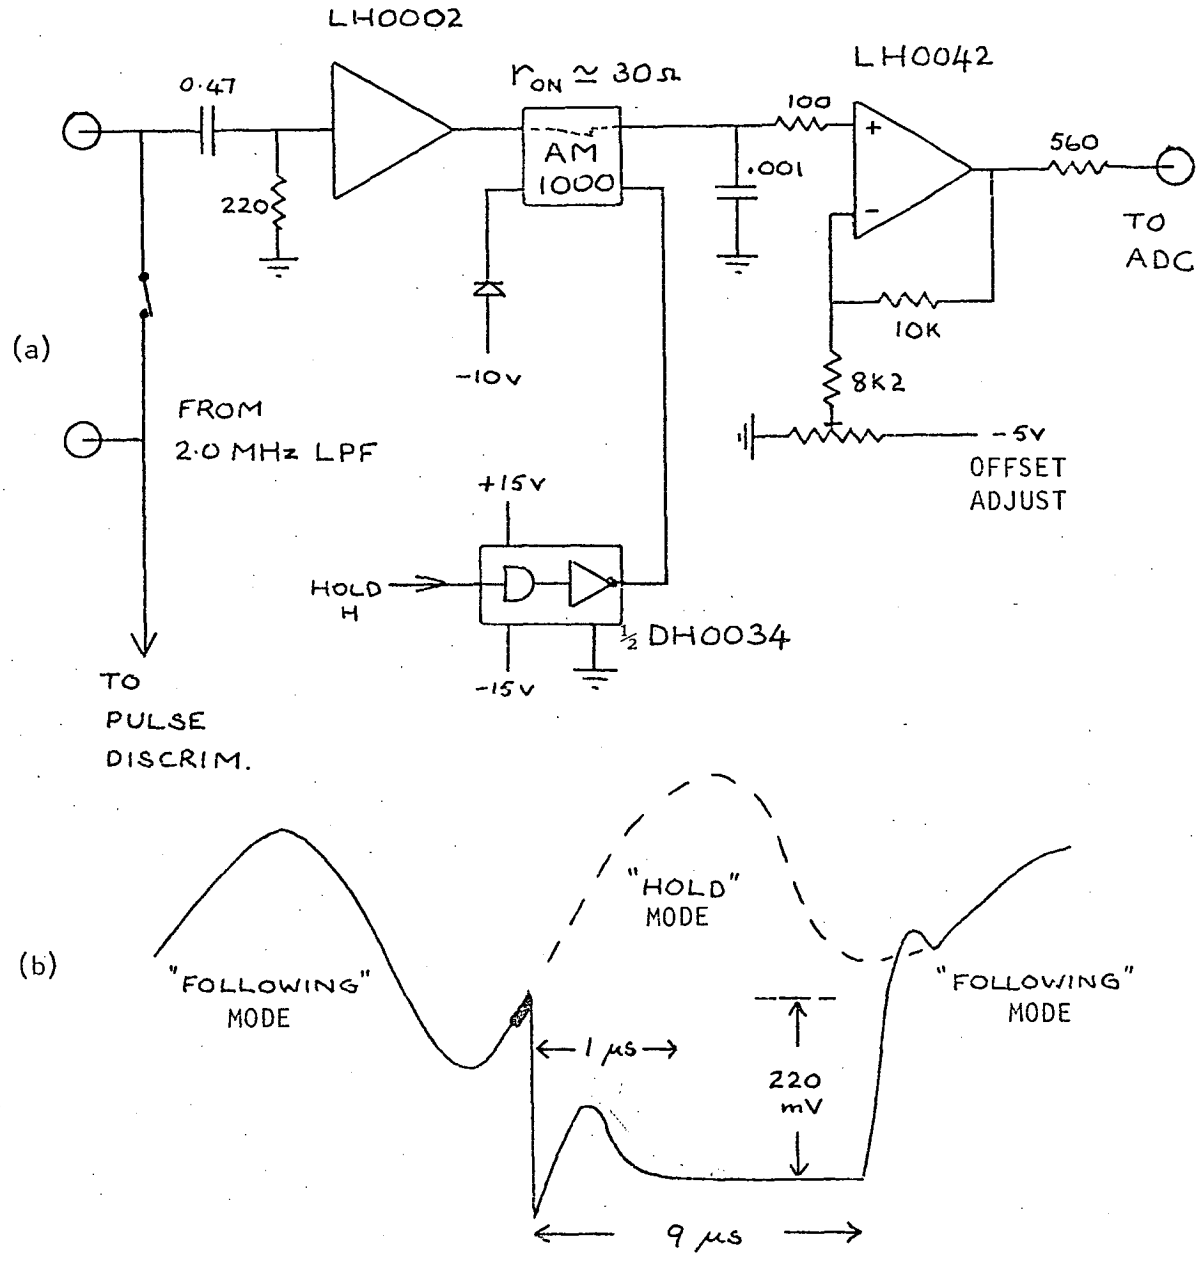

FIGURE 4.5.7A Fast sample-and-hold
(a) schematic
(b) sketch of output waveform (sample aperture shaded).

The aperture time has not been measured accurately. However an approximate lower bound is 35 ns, being the RC time-constant of the output impedence of the LH0002 buffer plus the "on" resistance of the AM1000 together with the .001 μF capacitor. Bench observations showed apparent fast action i.e. < 100 ns. Furthermore computer-drawn plots of the samples compared well with oscilloscope traces of the data waveform at the same part of the stripe.

The effect of a non-zero aperture time is to produce a roll-off in the system frequency response. The aperture time is in fact the equivalent width of the impulse response of a low-pass filter. This impulse response is assumed to be well-behaved, i.e. not oscillatory which would correspond to slope-sampling. Assuming for simplicity that it is rectangular with a width $\alpha\delta t$ where δt = 224 ns is the sampling interval, then the power frequency response of the sample-and-hold is $\text{sinc}^2(\alpha f/f_N)$ where f_N = 2.237 MHz. For α = 0.3 corresponding to 70 ns this gives a roll-off by 1.0 db at 2 MHz.

4.6 A/D CONVERSION AND COMPUTER CONTROL

In this section I describe the programs controlling the digitizing process and the mini-computers in which they were run. The signal distortion, i.e. noise introduction, due to quantization is discussed in the first sub-section. This distortion is always present when analogue to digital conversion occurs, however, in this case it is made important by the small number of bits devoted to storing each sample value.

The discussion is complicated by the fact that two different mini-computers have been used. More space is devoted to describing the system in current use. This involves a Digital Equipment Corporation PDP11/20 computer installed at the University of Tasmania High Frequency Array, Llanherne. The PDP11 is well known. This particular installation has 28K 16-bit words of core memory, a phase-encoded 9-track magnetic tape drive, an RK05 disk cartridge, numerous display peripherals and various interfaces for analogue and digital signals. Prior to December 1976 a PDP8/e was used. This was situated at the 2000 Foot Array, Llanherne. It had 12K 12-bit words of core, two TD8E DECtape transports, and analogue and digital interfaces.

The editing to remove over-modulation dropouts from the files of digitized samples is described in sub-section 4.6.3.

4.6.1 Quantization Noise

(i) Using the PDP11/20

The analog-digital convertor attached to the PDP11 produces a result of 10 bits with a separate sign bit. It is preceded by a pre-amplifier with selectable gain, which has usually been set to give a range ± 2.50 volt. The digitizing program VM2 rejects the least significant bit and stores each sample into an 8-bit byte. Therefore sample values have a quantum of

$$2.50 \div 1023 \times 2 = 4.9 \text{ mV}$$

and a possible range

$\pm$ 127 quanta i.e. $\pm$ 0.62 V.

A typical stripe replayed through the usual 2 MHz LPF presents a 60 mV rms signal to the ADC, i.e. it has a Gaussian amplitude distribution of rms width, σ = 12 quanta. Since $\sigma >> 1$ quantization noise can be assumed to be uniformly distributed over the range -½ to +½ quantum, and hence to have a power proportional to $\frac{1}{3}(0.5)^2$. This leads to a signal-to-noise ratio of

$(12)^2 / \frac{1}{3}(0.5)^2 \simeq 1700$, i.e. 32 db

averaged over the whole 2.237 MHz bandwidth.

Figure 4.4.2A showed the power spectrum of a stripe digitized using a 1 MHz LPF. The corresponding signal : quantization noise ratio is 27.8 db which gives a level of quantization noise 3 db below the observed background, i.e. the spectrum above 1.3 MHz. This confirms the identification of that background as replay noise. Furthermore for the more usual case in which the built-in 2 MHz LPF filter is used the relative amount of quantization noise is reduced even further. This is because the total signal power is approximately doubled by doubling the filter bandwidth and the replay noise is doubled while the quantization noise remains constant.

(This argument ignores the fact that the spectrum as replayed is not flat - the power below 1 MHz is more than twice that above because of the head resonance modification. A minor compensating effect is that the jitter noise is proportional to waveform slope and hence is more sensitive to the high end of the spectrum.)

In conclusion quantization noise in files created using the PDP11/20 computer is not important. Replay noise dominates so that the signal-to-noise ratio of the entire record-replay-sample system is about 24 db.

(ii) Using the PDP8/e

Quantization noise occurring when the PDP8/e computer controlled the digitization was relatively 12 db greater than with the system developed later using the PDP11/20. This was mainly because 6-bit storage was used instead of 8-bit as later adopted with the PDP11/20. Sample values were constrained to the range -32 to +31, the quantum being 20 mV. A typical "off-pulse" stripe showed after digitizing a Gaussian amplitude distribution of rms width, σ = 2.8 quanta. This leads to a signal to (quantization) noise ratio of

$$(2.8)^2 \; / \; \frac{1}{3} \; (0.5)^2 \simeq 100 \quad \text{i.e. } 20 \text{ db}$$

averaged over the whole 2.237 MHz bandwidth. Thus in contrast to the situation with the PDP11/20, quantization noise dominates replay noise so that the signal-to-noise ratio of the entire record-replay-sample system is about 20 db.

4.6.2 The Control Programs

(i) The PDP11/20 Digitizing Program

The PDP11 program VM2 controls the digitizing process. It interacts with the replay chassis via the following TTL signals (firstly those generated by the replay chassis):

(i) CONVT L, a falling edge which triggers the A/D convertor, which sets a "done flag" that is "seen" by VM2 some 20 μs later;

(ii) CPU FRAME H, asserted by the replay chassis to indicate that sampling is underway;

(iii) WAIT H, asserted when the operator has intervened (MANUAL) or when a PLL error has been detected in the preceding replay;

(and those generated by VM2):

(iv) RESET H, to initiallize the replay chassis;

(v) FIX H, to clear a WAIT H error state;

(vi) ADVANCE H, to move the sampling comb by one clock tick.

I define a "pass" as a replay of the stripe as indicated by CPU FRAME H and the corresponding 344 A/D conversions. There is one pass lasting 20 ms every 40 ms. During one pass 344 sample values are stored ready for validation during the following 20 ms. If an error is detected, i.e. WAIT H is asserted when checked soon after the fall of CPU FRAME H, then the pass is ignored, i.e. the 344 values are rejected, and an attempt is made to clear the error state using FIX H. The validation process is as follows.

Each sample has to be confirmed individually by measurements within a small tolerance (usually 3 quanta, i.e. 15 mV) on successive passes. Typically 5 to 10 error-free passes are needed to confirm all 344 samples for one position of the sampling comb. (In fact the two samples at each end of the comb are not checked. These are retained in the file but not used because they are affected by the VTR head resonance control.) Most of the samples are confirmed in the second pass and most of the remainder in the next. Samples at dropouts need more passes. If more than 30 passes are required the program aborts. However the normal result is that after fewer than 10 passes no unconfirmed samples remain, the value adopted for each being the average of the two successive values that are within the tolerance. Then the sampling comb can be advanced.

The sampling comb has to be advanced 255 times to completely digitize a stripe. To check the advancement process itself the comb is advanced one extra bringing it back to its initial position (modulo 256). Part of the final set of samples is compared with the first : the values should be the same to within the tolerance.

Four times during digitization of one stripe the 20000 byte buffer storage becomes full and is written to the RK05 disk cartridge while actual digitization is suspended. About one minute in total or half the typical run time of program VM2 is spent on this transfer.

For each stripe the program VM2 creates a disk file of 88064 sample values each of 8 bits. These samples are initially in order of sampling but are rearranged into time order by a program called VMORDR. This takes about 2 minutes. VMORDR also accumulates and stores in the file a histogram of the occurrence of each of the 255 values that are possible. This is the amplitude distribution of the data. It is invariably well fitted by a slightly skewed Gaussian curve except for

(i) depressions of 2 - 5% in alternate bins due to slight irregularities in the A/D convertor;

(ii) a long tail at the low end when over-modulation dropouts are present.

The low end tail usually involves one to about 50 points out of 86000; if more over-modulation dropouts are present then the stripe cannot be digitized at all due to PLL errors. This tail is roughly flat from -5σ to -10σ. The standard deviation, σ, of the Gaussian is usually 10 - 13 quanta. This represents the rms signal averaged across the whole 2.237 MHz bandwidth.

(ii) The PDP8/e Digitizing Program

The PDP8/e digitizing program was similar to the PDP11/20 program described above, but no validation was done. Only one error-free pass was used for each position of the sampling comb, hence most of the 1 minute run time of this program was spent in writing to DECtape. A duplicate digitization was done immediately. Thus two or more files were created for each stripe, each file consisting of $257 \times 344 = 88408$ 6-bit numbers and filling one quarter of a DECtape reel. Each file was rearranged into time order using the Elliott 503 at the University of Tasmania. (This computer was shut down in 1977.) The duplicate files were then compared and the file needing least editing of dropouts was selected for analysis.

One stripe, pulse "A" of PSR 0950+08, has been digitized both with the PDP8 system and with the PDP11 system. The raw files are very similar apart from the difference in quantum. After dispersion removal the details of the pulsar emission are virtually identical.

4.6.3 Editing Over-Modulation Dropouts

(i) Editing Files created by the PDP11

Prior to transcribing the file of samples for each digitized stripe from PDP11 disk to magnetic tape, the file is checked and, if needed, edited using the interactive program VMEDIT. This program scans through the file looking for points below a specified threshold. If any are found the 128 points in the vicinity are displayed on a CRO screen. The suspect section of the file can then be identified and edited. The threshold itself is set by the operator: the histogram calculated by program VMORDR is displayed and when suitably scaled up (there are about 4000 counts in the central bin) any outliers at the low end of the distribution are visible. As mentioned in the last section there are at most about 50. They are overmodulation dropouts.

The total number of overmodulation dropouts in a given stripe is reduced by the head resonance control. Those remaining which are intermittent are probably, but not certainly, eliminated by the confirmation algorithm of VM2. VMEDIT is another weapon against them.

Recall that in an overmodulation dropout a peak in the waveform is replaced by a large negative spike. It is not possible to reconstruct the missing peak nor can any structure superimposed on the exponential recovery be relied upon, due to the combination of the interlaced sampling scheme and the intermittent nature of some dropouts. Therefore I replace the suspect section of the file, usually about 10 successive samples, with a section randomly chosen from elsewhere in the file. The end-points of the "patch" are joined in smoothly.

For all off-pulse stripes and most on-pulse stripes analysed no editing was required. For those which were edited the effect was merely of a small reduction in S/N ratio compared to the case of no overmodulation occurring. No structure on a microsecond scale can be destroyed by a single dropout whether edited out as described or left in the file. This is because dispersion removal (not post-detection) smears out the disturbance over a time equal to the bandwidth sweep time. For a 70 kHz channel analyzing data for PSR 0950+08, the pulsar with least dispersion, this time is 25 μs. For a 2 MHz bandwidth the corresponding time is 0.73 ms.

(ii) Editing Files Created by the PDP8

Files created by the PDP8 were edited in a sequence of batch runs using the Elliott 503 computer. In general much more editing was needed than would have been the case if the PDP11 had been used. The files contained intermittent dropouts of both the random and the overmodulation type which are removed by the PDP11 validation method. Samples in error were detected as being points out-lying from the Gaussian amplitude distribution or by comparison of duplicate files. For intermittent dropouts the "correct" value for each sample in error was provided by one or other of the files. Reproducible overmodulation dropouts were usually "patched" to a constant waveform, i.e. constant voltage through the suspect section. As discussed in (i) above the exact nature of the "patch" matters little since it is smeared out by dispersion removal.

CHAPTER 5

CALIBRATION AND SYSTEM TESTS

5.1 INTRODUCTION

This Chapter deals with the measurement of three aspects of the total system response : the phase response of the equipment; the amplitude response; the exact observing frequency. Only in the case of the latter was the Molonglo radiotelescope itself included in the measurements. Assumptions and approximations were required for the others.

Measurements of the phase response of the record-replay-digitize system are described in Section 5.2. It is assumed that the Molonglo response is phase linear to a tolerance ±5° across the passband of 2.24 MHz. The published tolerance is ±1° (Mills *et al* 1963).

Measurements of the amplitude response of the record-replay system are described in Section 5.3. The results are presented for reference in regard to the estimation of aliasing and replay noise. They have not been used for calibration. Instead the standard method of analysis of pulses has involved numerical adjustment of the spectrum to flatten the off-pulse spectrum. The assumptions behind this are

(i) that the Molonglo front-end noise is dominant,

and (ii) that this noise is proportional to the telescope gain function. Of course the second assumption would have been unnecessary had the "sky" noise been dominant.

It should be noted that the VTR part of the system is assumed to be a linear system so that its performance can be expressed by a (complex) gain function. It is not linear for signal levels comparable to those of the recorded data. This is demonstrated by its square wave response. At an amplitude exceeding 150 mV (measured on replay) there are significant differences (10%) between the rising and falling edges, both in slope and overshoot. The rms amplitude of the data is typically 60 mV. Therefore treating the VTR as a linear system is a significant departure from reality.

Furthermore the data should have been recorded at a lower level.

Unless otherwise stated, the LPF used during replay was the one built into the replay chassis, that is, 2.0 MHz cut-off (0.3 db).

5.2 PHASE RESPONSE

5.2.1 Theoretical Response of the Filters

The phase response of a filter system can be calculated by integrating the group delay time, tg, defined by

$$tg = \frac{d\phi}{d\omega} \tag{1}$$

where ϕ (radian) is the phase at angular frequency $\omega = 2\pi f$. For a Cauer-Cauchy filter tg is given by formulae of the form

$$tg = \sum_i \frac{\sigma_i}{\sigma_i^2 + (\omega \pm \omega_i)^2} \tag{2}$$

where the constants σ_i and ω_i have been tabulated for normalized low-pass filters by Zverev (1967). For the 7-pole 4.7 MHz HPF in the record chassis five terms of equation 2 are required ($i = 0, \pm 1, \pm 3$). The transformation from normalized low-pass to high-pass values has been described well by Hansell (1969). For each of the two 5-pole filters, the 2.06 MHz LPF in the record chassis and the 2.0 MHz LPF in the replay chassis, three terms of equation (2) are needed.

The three separate group delay responses were added together giving an estimate of the group delay response for the whole system. This is shown in Figure 5.2.1A. In performing the addition the response of the HPF, which precedes the mixer in the record chassis, was translated in frequency by the local oscillator frequency 4.474 MHz. That is, the group delay calculated for this filter for the frequency (f + 4.474) MHz contributes to the total group delay at f MHz. As shown in Figure 5.2.1A the total group delay varies by 1.8 μs across the bandpass (0 - 2.2 MHz) with maxima at both ends

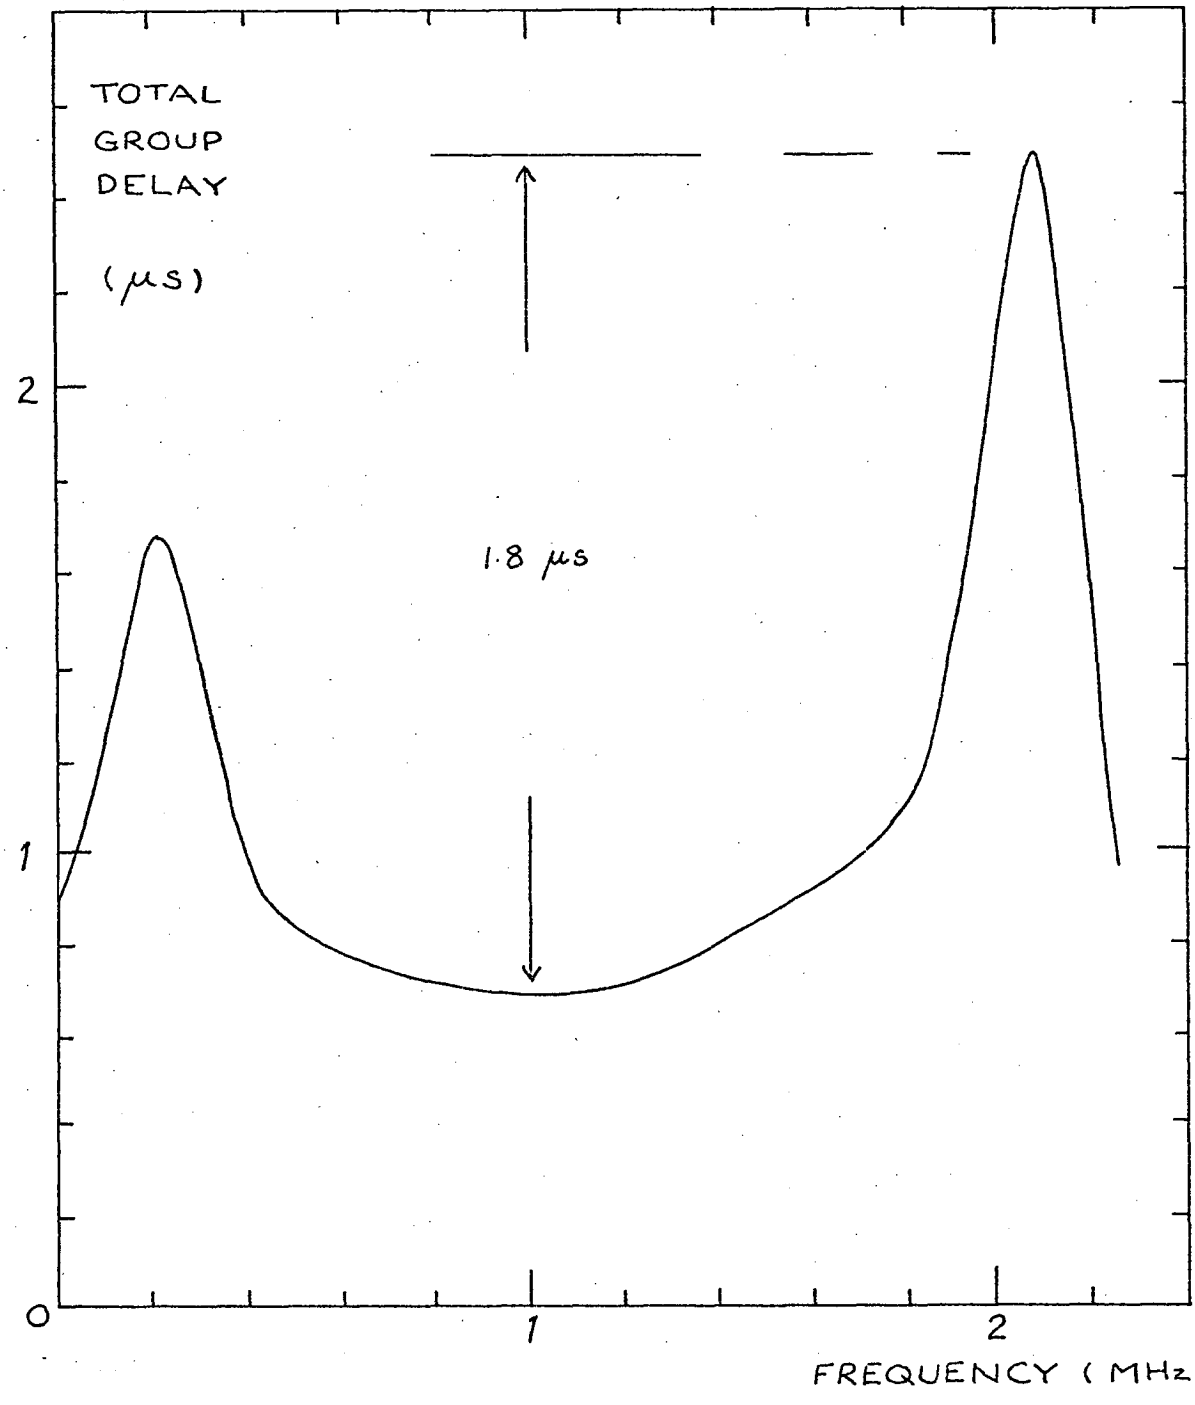

FIGURE 5.2.1A Total theoretical group delay curve obtained by adding the separate curves for the 4.7 MHz HPF, 2.06 MHz LPF and 2.0 MHz LPF.

i.e. in parts of the spectrum effected by imaging and aliasing. Only in the central 1 MHz is the group delay variation less than the sampling interval 224 ns. However the variation is much less than the resolution time for a 1 MHz channel, namely 1 μs, so that compensation for the phase response is not needed. The compensation is needed for a channel bandwidth greater than 1 MHz, e.g. a (rectangular) bandwidth of 1.8 MHz centred on 1 MHz, not only because of the reduction in resolution time but also because the regions of greater group delay variation at each end of the spectrum must be included in the pass-band.

5.2.2 Bench Measurement of Phase

The phase and amplitude responses of sections of the record-replay equipment were measured with a Hewlett-Packard model 8405A vector voltmeter. The signal path was broken only at the mixer in the record chassis. The VTR was operated in the "record monitor" mode, i.e. the FM carrier is modulated and immediately demodulated so that the actual tape and the heads are not involved. In this mode the VTR has an amplitude response that is flat ±0.5 db and a phase response that is linear, i.e. uniform delay, to ±3° from 0 to 2.2 MHz.

The responses of the sections preceding and succeeding the mixer were added using the frequency translation described in Section 5.2.1. The combined phase response is represented by the dashed curve in Figure 5.2.4A (with an arbitrary linear, i.e. delay, component removed). The curve below 1 MHz is less accurate because the vector voltmeter could not be used at actual frequencies below this. The measurements of the equipment following the mixer at frequencies below 1 MHz were made using an oscilloscope. A systematic error ±10° is likely near zero frequency.

5.2.3 Phase Response using Square-Waves

The following indirect method of measuring the phase response of the system is clearly superior to the direct bench method as it includes the effect of mis-matches at the mixer and the effect of the record-replay cycle itself.

Square-waves with fundamental frequencies in the range 10 to 100 kHz were applied to the input of the record chassis and recorded on videotape. Five stripes were digitized and Fourier transformed and the complex spectrum used to evaluate the phase. The input signal level was set below the onset of overload as indicated by an increase in even harmonic distortion. Actually it was those harmonics of the square-waves in the band from 4.474 to 6.701 MHz that were recorded; they were translated downward in frequency to 0 - 2.237 MHz. However there was some leakage at each end of the band, i.e. imaging at the low end and aliasing at the high end.

The harmonics were observed as peaks in the power spectrum, i.e. spectral lines, each with a sinc-squared profile. The centre frequency of each peak could be found by interpolation to an accuracy 0.1 δf where δf is the tabulation interval (δf = 546 Hz for an 8192 sample Fourier transform). However interpolation could not be used to find the phase at the centre of each peak due to the rapid change in phase by 180° at the centre. Therefore the stratagem of taking the nearest tabulated phase value was adopted. Each such value had an ambiguity of 180°. Furthermore each phase spectrum had an arbitrary linear (delay) component producing a phase difference between successive harmonics of up to ±180°, i.e. ±360° between successive odd harmonics.

A computer program was developed to aid in the analysis of the spectra. With the aid of a graphics terminal to display plots of phase versus frequency, solutions for the linear component could be found quickly through iteration. Furthermore once a smoothed phase response had been adopted it

was refined iteratively by compensating each new spectrum accordingly and plotting the residual phase response. The phase-frequency plot for the first spectrum analyzed in this way (without compensation) is shown in Figure 5.2.3A. The overall shape of the curve is somewhat arbitrary due to the freedom of choice in subtracting a linear component. The noise in this plot is mainly due to taking the tabulated values nearest the centre of each spectral peak. However some of the peaks were weak enough to be affected by replay noise.

The assumption behind this method of phase response measurement is that all the harmonics involved were in phase at the input of the equipment. This assumption is open to doubt for the very high harmonics, i.e. harmonic number $n \sim 500$. One of the signal generators used produced both odd and even harmonics with about equal amplitude for $n > 400$, i.e. above 4.474 MHz with a fundamental of 11.3 kHz. The normal amplitude law 1/n was not obeyed, there being an additional cosine factor in the envelope for the odd harmonics with a complementary sine factor for the even harmonics. This is the usual behavior when the mark-space ratio of a square-wave is not precisely unity. As expected the phase of each even harmonic was midway between the phases of the neighbouring odd harmonics. The presence of the even harmonics was convenient in that twice as many data points were obtained from each spectrum and the choice of which linear component to subtract was easier. The fact of neighbouring odd and even harmonics being in phase lends support to the assumption that all the harmonics were in phase at the input to the equipment. Further, and convincing, support is provided by the fact that using a different signal generator (with no significant even harmonics) at a number of different fundamental frequencies led to the same phase response.

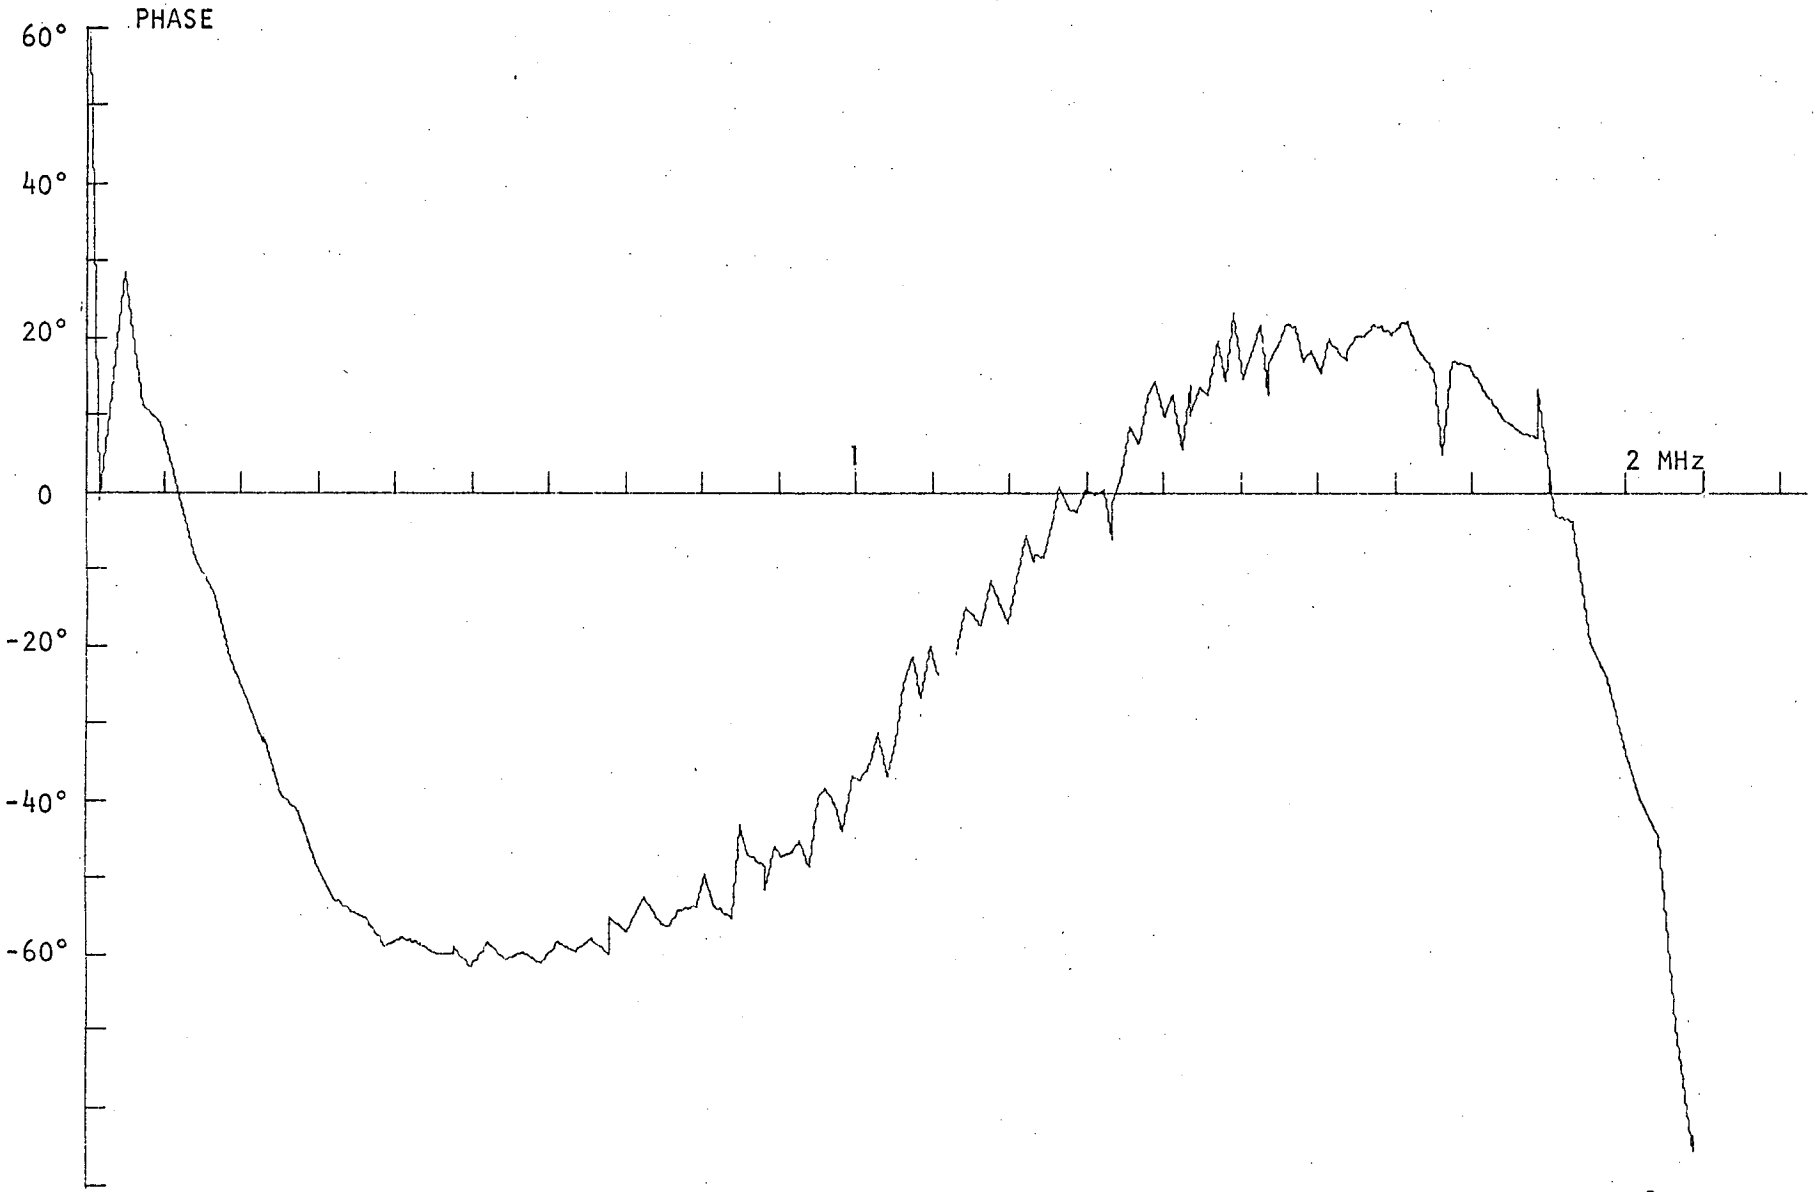

FIGURE 5.2.3A System phase response as measured by square-wave harmonics using one 8192 point Fourier transform. The fundamental frequency is 11.26 kHz. Harmonic numbers 398 to 583 are present (both odd and even).

5.2.4 Comparison of the Phase Measurements

Figure 5.2.4A shows the phase response

(i) as calculated from the filter theory group delay,

(ii) as measured directly,

and (iii) as obtained from the square-wave spectra.

Figure 5.2.4B shows the residuals (with linear components removed):

(a) of (ii) - (i) representing the effect of impedence mis-match at the mixer and elsewhere, and of various amplifiers in the system;

(b) of (iii) - (ii) representing the effect of the record-replay cycle including the head resonance.

Both sets of residuals are satisfactorily small, viz. ±10°.

Figure 5.2.4C shows the curve adopted as the best estimate of the instrumental phase response. It is based entirely on the square-wave measurements. (Note that it differs from curve (iii) in Figure 5.2.4A mainly by a linear component.) The errors in this curve are less than ±10° between 100 and 300 kHz, and ±2° above 300 kHz. This phase response was subtracted from the data for all dispersion removal analyses (performed after August 1977) involving single channel bandwidths exceeding 0.5 MHz.

The effect of the adjustment of the VTR head resonance response on phase can be seen in Figure 5.2.4D. The adopted phase response has been subtracted for this plot so that the remaining departure from linearity represents the "wide FM bandwidth" setting relative to the usual "minimum dropouts" setting. Note that no pulses were analyzed after digitizing with the "wide" setting.

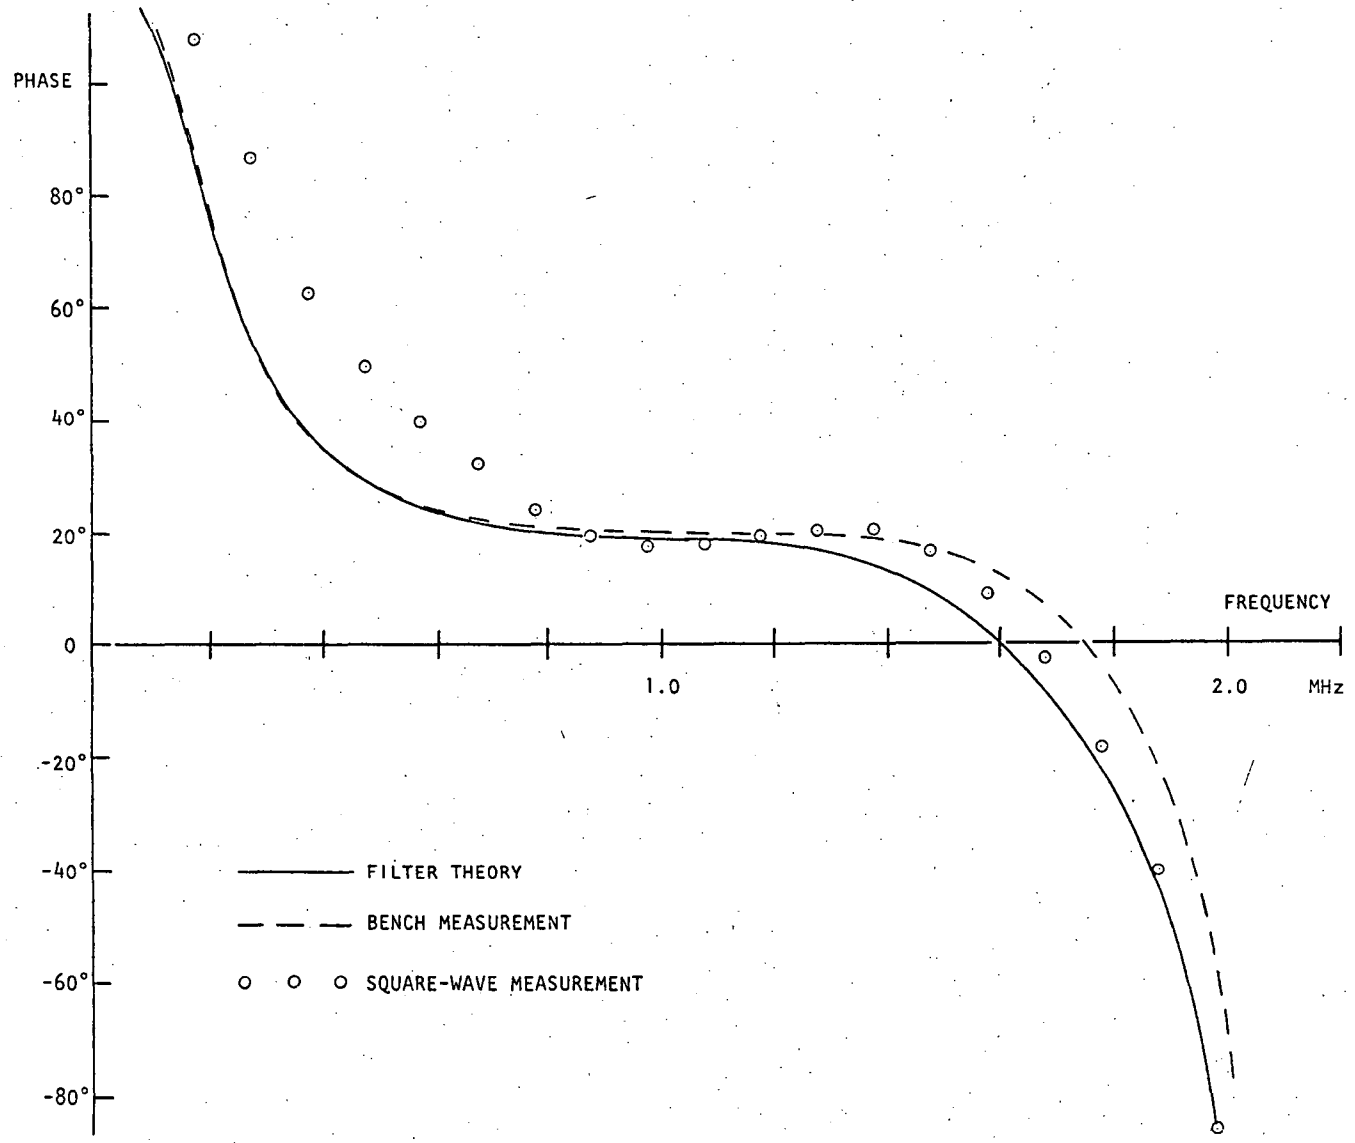

FIGURE 5.2.4A Comparison of phase responses.

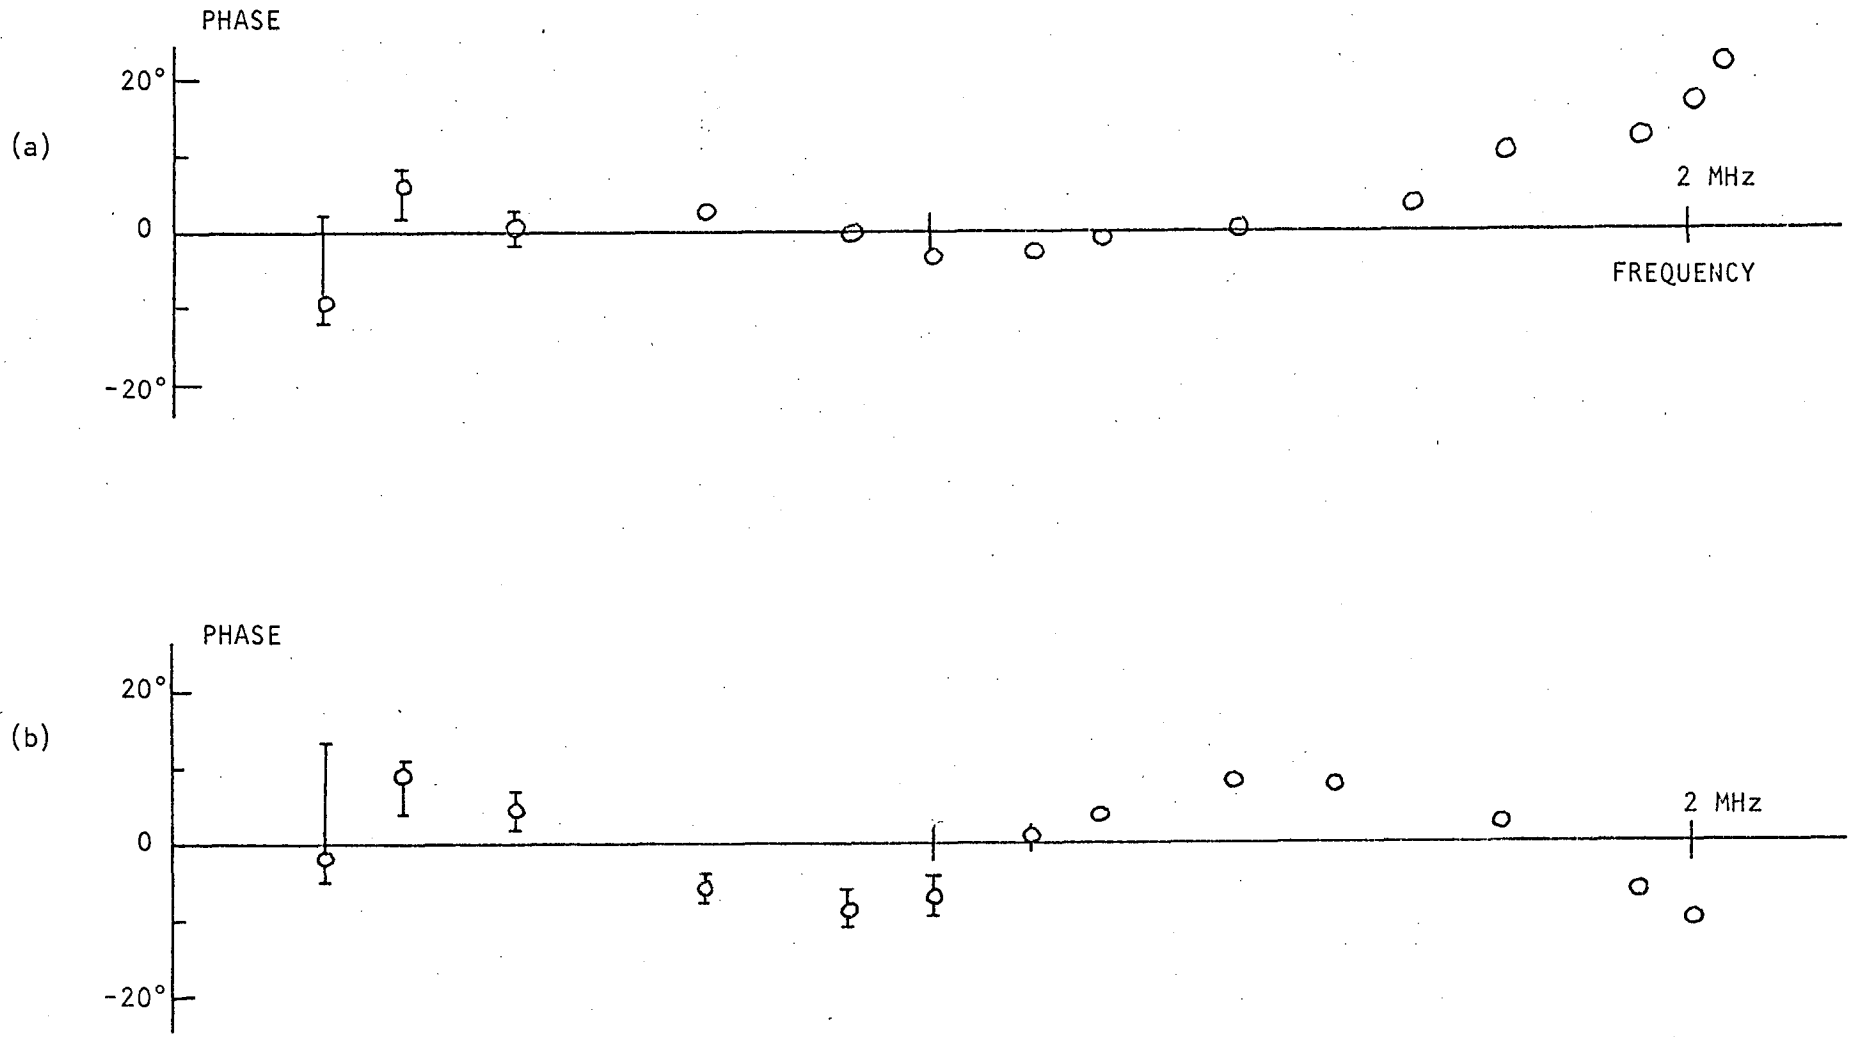

FIGURE 5.2.4B Phase residuals: (a) bench measurements less filter theory (effect of mis-matches etc.); (b) square-wave harmonic measurements less bench measurements (mainly effect of VTR).
Arbitrary linear components removed.

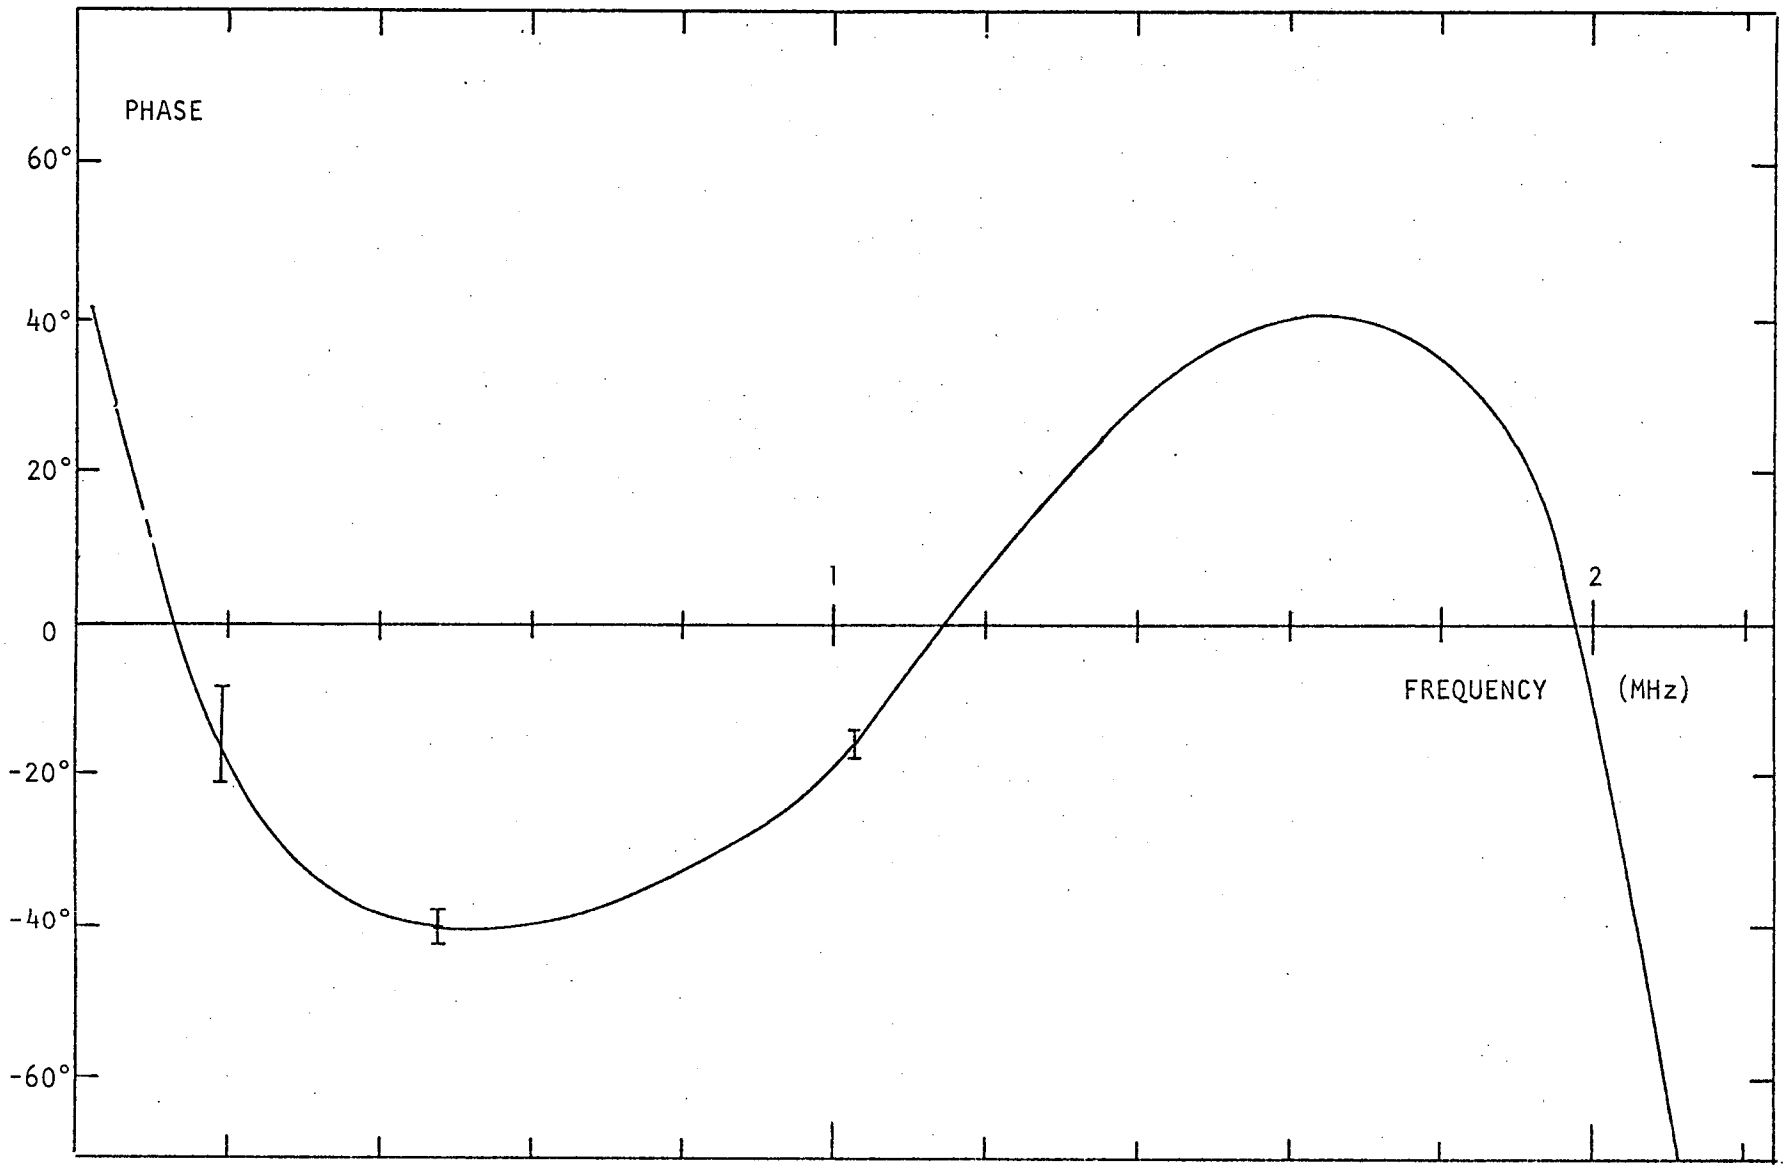

FIGURE 5.2.4C Curve adopted as best estimate of instrumental phase response. (An arbitrary linear component has been removed.) Typical error bars are marked.

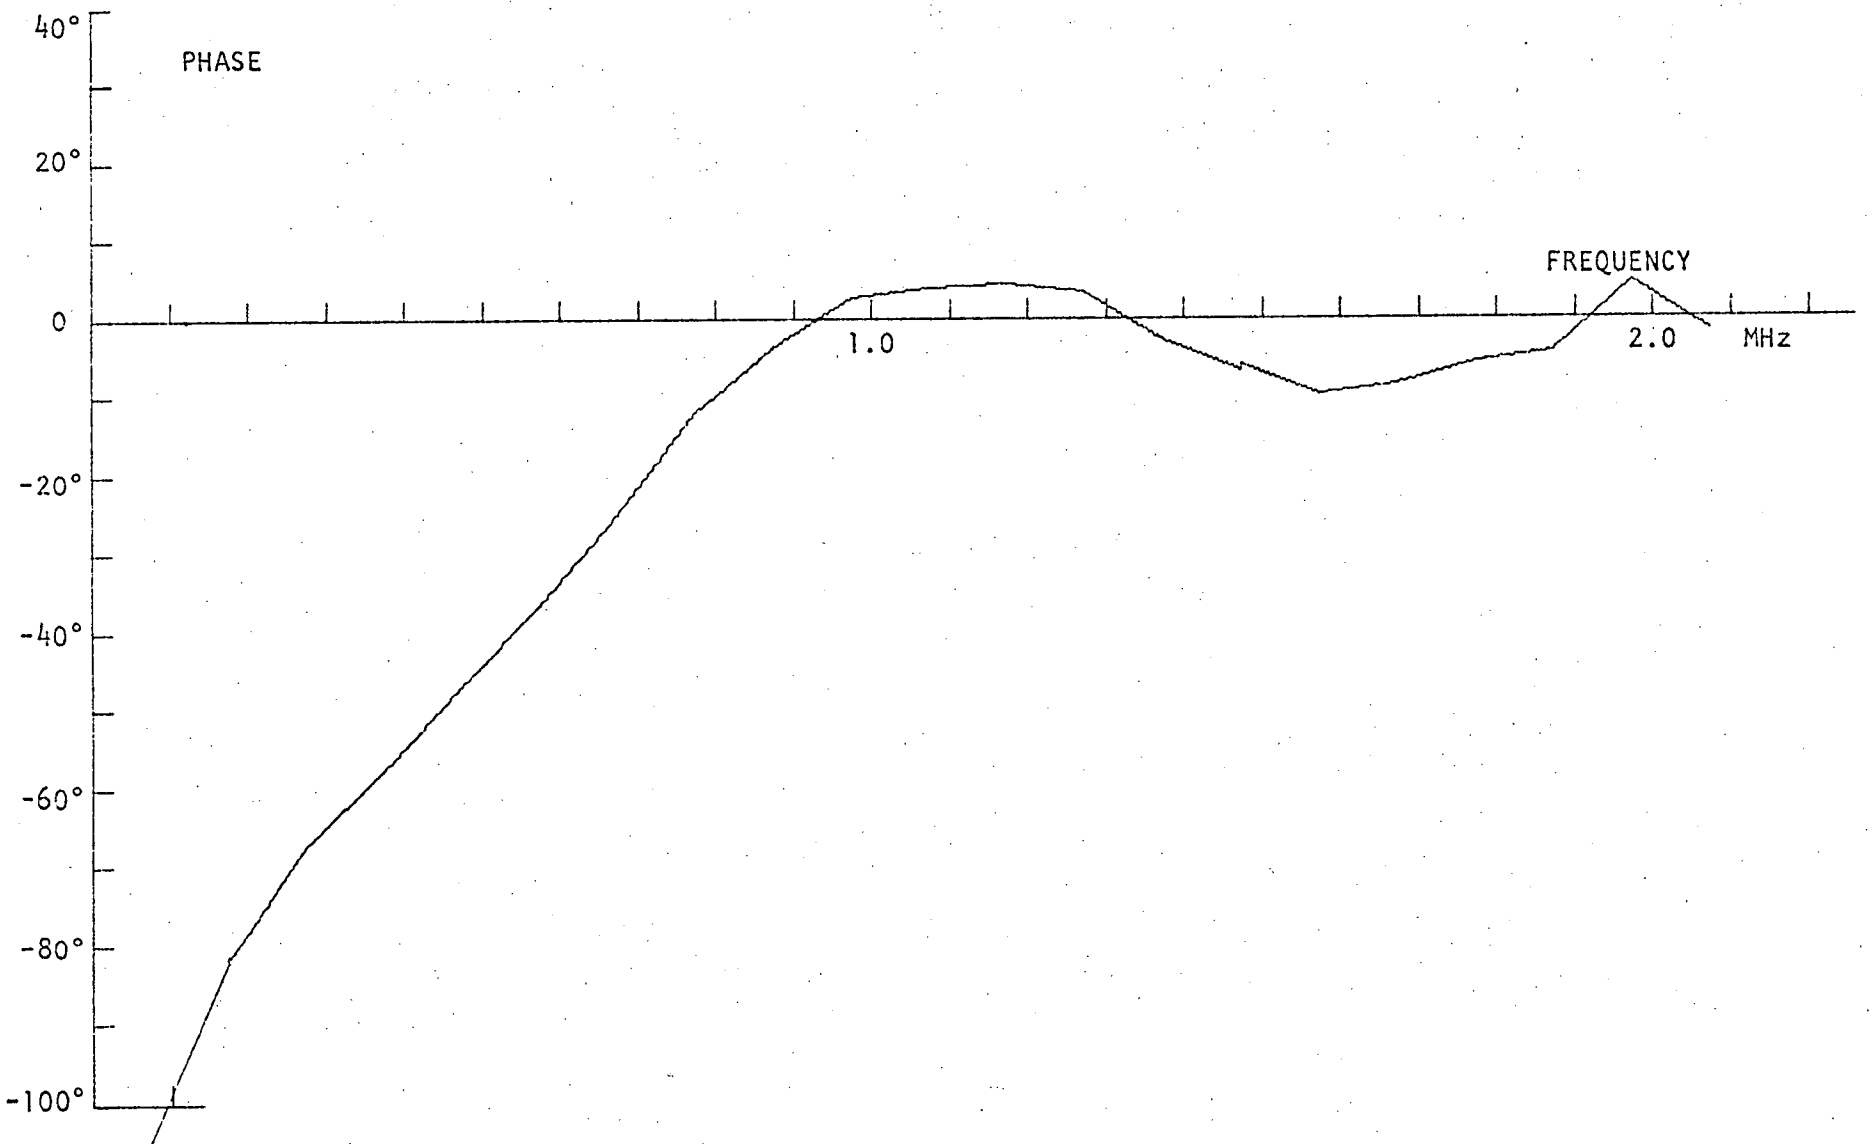

FIGURE 5.2.4D The change in the phase response of the system due to the head resonance adjustment (±20°). See text.

5.3 AMPLITUDE RESPONSE

The square-wave method used for phase was not reliable for amplitude. However bench tests were easy. The amplitude response of the record-replay equipment excluding the VTR is shown in Figure 5.3A. The image and alias curves shown are simply the reflections of the relevant filter cut-off slopes about 0 and 2.237 MHz respectively. The notch in the response at 2.237 MHz which is imposed on the data in the record chassis is shown in the figure. However the reference tone itself is not subject to this notch only to attenuation by the 2.0 MHz LPF, whose cut-off slope is also shown. Replay noise is also subject to this LPF only.

Figure 5.3B shows an observed off-pulse spectrum smoothed to about 100 kHz resolution. The main difference between this curve and the corresponding curve in Figure 5.3A is the VTR record-replay cycle. (The Molonglo spectrum was assumed flat to ±0.5 db.) The head resonance adjustment has produced an attenuation of 6 db, constant above 0.8 MHz, with little change below 0.3 MHz. The increase near zero frequency is in accordance with the 3 db predicted due to imaging. At the other end of the spectrum the cut-off slope levels off near 2.2 MHz, apparently due to sampling of the reference tone because of reproducible jitter in the sampling comb.

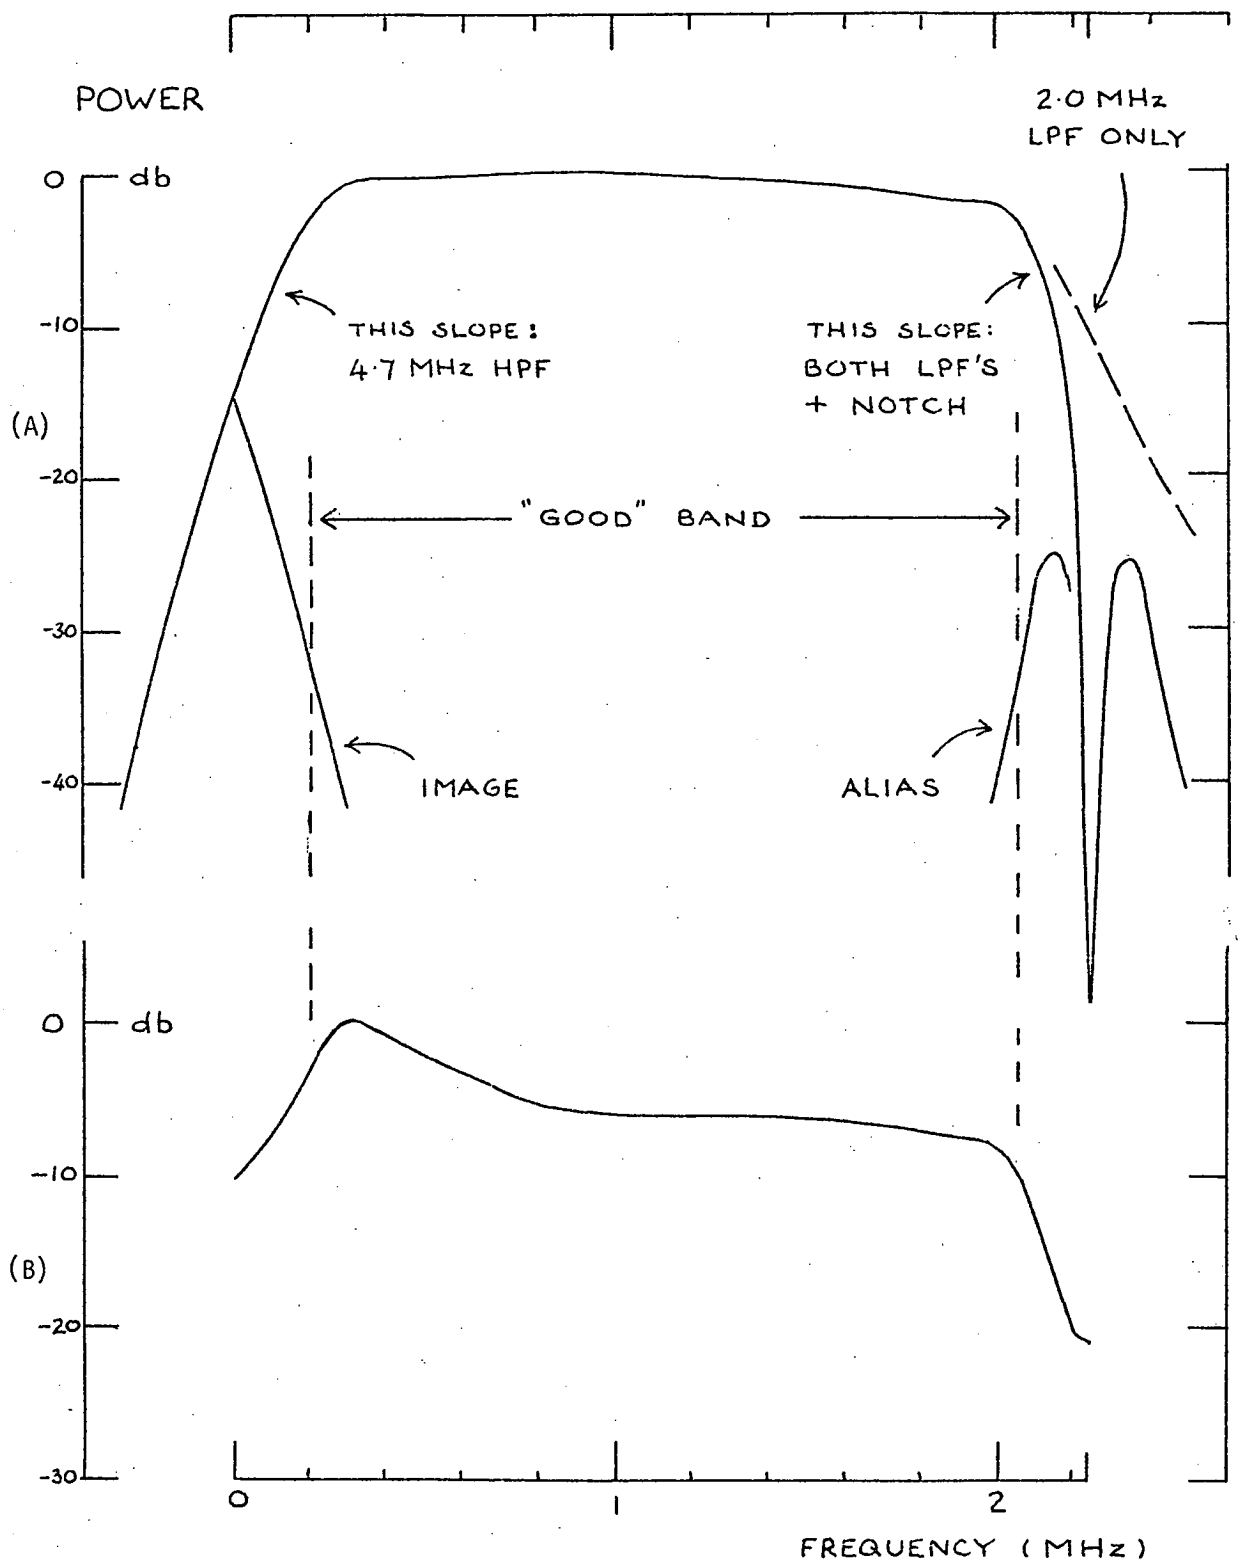

FIGURE 5.3A & B

A. System amplitude response excluding VTR. The "good" band is that frequency range in which image and alias signals are below -30 db.

B. System amplitude response including VTR obtained from the observed power spectrum assuming Molonglo system noise is "flat".

5.4 OBSERVING FREQUENCY

Sinusoidal signals at 0.5 MHz intervals from 407 to 409 MHz were recorded through the Molonglo system. The videotape was replayed and the resulting spectral lines observed with a spectrum analyzer. The observed frequencies were divided by a factor ~1, in fact 1.027, to compensate for the time compression due to the change in geometry of the VTR stripe using stop-action replay. This factor was obtained from the observed frequency of the reference tone, namely 2.298 MHz in this case. Hence was established the mapping

$$(409.0254 - f)\ \text{MHz} \rightarrow f\ \text{MHz},\quad 0 \leq f \leq 2.237$$

The crystal derived reference frequency was measured as

$$2\,236\,958 \pm 1\ \text{Hz}.$$

The frequency mapping was established to a higher accuracy, namely ±60 Hz, by digitizing and Fourier transforming a stripe and locating the line in the power spectrum. This is a much higher accuracy than needed since an error of 1 kHz is equivalent to a fractional error in dispersion measure of 5×10^{-6}. The dispersion measures of only a few pulsars are published to more than 4 figures.

CHAPTER 6

DISPERSION REMOVAL METHOD

6.1 INTERSTELLAR DISPERSION AND ITS EFFECT ON TIME RESOLUTION

This Section presents the equations which describe the distortion of pulsar signals due to interstellar dispersion. There is little original material in this Section; it is presented as a basis for Section 6.2 following. Much of the material was presented in more detail by Hankins and Rickett (1975) and I have generally used the same symbols.

Let the emitted pulsar signal, that is, the voltage that would be induced in an antenna near the pulsar, be s(t). This can be expressed in narrow-band form:

$$s(t) = \mathrm{Re}[v(t)\ \exp(i2\pi f_o t)] \rightleftharpoons S(f) \qquad (1)$$

$$S(f) = \tfrac{1}{2}\ [V(f - f_o) + V^*(-f-f_o)]$$

where the asterisk denotes the complex conjugate and $\rightleftharpoons$ denotes a Fourier transform relationship, that is, the complex envelope of the signal is

$$v(t) = \int_{-\infty}^{+\infty} V(f)\ \exp(i2\pi ft)df$$

(In this Chapter I shall use upper-case letters to denote functions of frequency that are Fourier transforms of the functions of time denoted by the corresponding lower-case letters.) Frequency f_o is the "carrier" frequency and $V(f) = 0$ for $|f| > B$ where B is the receiver bandwidth.

In propagation through the interstellar medium the signal s(t) is convolved with the impulse response of the medium, h(t). The corresponding transfer function is derived from the dispersion law for a tenuous plasma and can be expanded in a Taylor series about some reference frequency, f_1, thus

$$H(f) = H(f_1 + \Delta f)$$

$$= \exp\{-i(\phi + 2\pi t_1\ \Delta f + \pi/\alpha\ [-\Delta f^2 + \Delta f^3/f_1 - \Delta f^4/f_1^{\,2} + \ldots])\} \qquad (2)$$

where $|\Delta f| \ll f_1$,

ϕ is effectively randomly distributed between 0 and 2π,

t_1 is the group delay at frequency f_1, i.e., the pulse travel time,

and α is the bandwidth sweep rate. Now

$$t_1 = z/c + D/f_1^2$$

where z is the distance from the pulsar, c is the velocity of light and D is the dispersion constant. The delay term ($i2\pi t_1 \Delta f$) can be eliminated by re-defining the time origin and the phase constant term $i\phi$ can be ignored. The dispersion constant, D, is proportional to the dispersion measure, DM. Since the constant of proportionality involves physical constants which are not known accurately, the working definition of DM in terms of the observable, D, is (Manchester and Taylor 1972)

$$DM\ (cm^{-3}pc) = 2.410000 \times 10^{-16} D\ (sec\ Hz^2). \qquad (3)$$

I define the bandwidth sweep rate, α, by

$$\alpha = \alpha(f_1) = -\ (df_1/dt_1) = f_1^3/2D \qquad (4)$$

(I shall ensure that α is always positive; in the next Section when f_1 is taken to be negative, minus signs will be introduced into the equations.) The term in Δf^2 is known as the linear sweep term while the higher order terms are known as the curvature terms. A criterion for the omission of the latter, that is, the assumption that the signal sweeps linearly through the receiver pass-band, is that they should not represent a phase shift of more than 1 radian from band centre to band edge. The maximum receiver bandwidth for which this is true is

$$B_m = 2(\alpha f_1/\pi)^{\frac{1}{3}} = 2(f_1^4/2\pi D)^{\frac{1}{3}} \qquad (5)$$

For $f_1 \simeq 408$ MHz,

$$B_m \simeq (2.0\ MHz)/[DM(cm^{-3}pc)]^{\frac{1}{3}}$$

Therefore the "linear sweep approximation" is not accurate enough for this project and I have retained two curvature terms.

The bandwidth sweep time, t_s, is defined, neglecting the curvature terms, by $t_s = B/\alpha$. Values of t_s for a 2 MHz channel at 408 MHz and the pulsars of this project are given in Table 1.

Consider a receiver whose transfer function in narrow-band form is

$$E(f) = R(f - f_o) + R^*(-f - f_o) \tag{6}$$

where the complex frequency response R(f) is non-zero only near zero frequency and has width B. Then the complex amplitude, y(t), of the received pulsar signal is

$$y(t) = v(t) * r_a(t)$$

where the asterisk denotes convolution and $r_a(t)$ is the impulse response of the system consisting of the interstellar medium and the receiver. Let H(f) be split into positive and negative frequency halves,

$$H(f) = H_+(f) + H_-(f),\ H_+(f) = 0 \text{ for } f<0$$

then

$$r_a(t) = h_+(t) * r(t) = \int_{-\infty}^{+\infty} R(f)\ \exp(i\pi f^2/\alpha)\ \exp(i2\pi ft)df \tag{7}$$

Equation (7) uses the linear sweep approximation. If the pass-band is rectangular then equation (7) becomes

$$r_a(t) = \int_{-B/2}^{+B/2} \exp(\ i\pi f^2/\alpha + i2\pi ft)df \tag{8}$$

which is the same form as that describing the Fresnel diffraction pattern of a slit. If $\alpha << B^2$ then the pattern is αB wide to quarter power (see also Section 6.4.1). The condition $\alpha << B^2$ states that the dispersion smearing is substantial. For a smooth pass-band under this condition it can be shown that

$$|r_a(t)| \simeq |\alpha^{\frac{1}{2}}\ R(\alpha t)| \tag{9}$$

The signal reaching the square law detector is

$$x(t) = \mathrm{Re}[y(t)\ \exp(i2\pi f_c t)] \tag{10}$$

where f_c is either f_o or the intermediate frequency of the receiver and frequencies near $2f_c$ are filtered out. Alternatively a base-band quadrature detector can be used. In either case the detector output can be written (neglecting gain factors)

$$I(t) = |y(t)|^2 = |v(t) * r_a(t)|^2$$

If the emitted pulse is incoherent in the sense that successive samples of v(t) are independent, as in amplitude-modulated noise, then using (9)

$$I(t) = |v(t)|^2 * |r_a(t)|^2 \simeq \alpha|v(t)|^2 * |R(\alpha t)|^2 \tag{11}$$

Thus the receiver bandshape is, in effect, converted to a function of time by the frequency sweep and convolved with the pulse. In particular an emitted impulse produces a detected pulse of width $t_s = B/\alpha$. This was depicted in Figure 1.2A.

It was assumed in the last paragraph that $B >> \alpha^{\frac{1}{2}}$. If however $B < \alpha^{\frac{1}{2}}$, i.e. $t_s < 1/B$, then the detected pulse width for an emitted impulse is 1/B. Therefore if post-detector dispersion removal is to be used the optimum choice for the bandwidth, b, of each channel is $b = \alpha^{\frac{1}{2}}$ and the best time resolution is $1/b = \alpha^{-\frac{1}{2}}$. At 408 MHz, $b \simeq (90\ \mathrm{kHz})/\mathrm{DM}^{\frac{1}{2}}$. For PSR 0950+08 at 408 MHz the optimal bandwidth is 50 kHz giving a resolution of 20 μs. For all other pulsars the bandwidth must be narrower and hence the resolution poorer (Table 1).

PSR	Dispersion Measure DM ($cm^{-3}pc$)	Sweep-rate α (MHz/ms)	2 MHz Sweep-time t_s(ms)	Optimum post-det. BW (kHz)
0950+08	2.969 ± .001	2.756	0.726	53
1133+16	4.8479 ± .0006	1.688	1.18	41
1642-03	35.71 ± 0.01	0.229	8.73	15
1749-28	50.88 ± 0.14	0.161	12.43	13
0833-45	69.08 ± 0.01	0.118	16.88	11

TABLE 1: The dispersion measures of five pulsars and derived quantities for a centre frequency of 408 MHz. The values of DM are from Taylor and Manchester (1975).

6.2 PRE-DETECTION DISPERSION REMOVAL FOR THE VTR SYSTEM

This Section presents the theory of the specific dispersion removal scheme used for the videotape system. Continuous Fourier representation is used; the properties of the discrete Fourier transform are introduced and specific data processing formulae are given in the next Section.

From equations (1) and (6) the spectrum of the received pulsar signal is

$$X(f) = \tfrac{1}{2}V(f - f_o)H_+(f)R(f - f_o) + \tfrac{1}{2}V^*(-f-f_o)H_-(f)R^*(-f-f_o)$$

At this point specific definitions are required for f_o and R corresponding to the Molonglo lower sideband down-conversion. Define f_o = 409.025 MHz which is the frequency that is converted to zero video frequency. With some idealization of the real situation, namely assuming there is no imaging or aliasing, f_o is the upper end of the pass-band. Therefore R(f) = 0 for $f < -B$ or $f > 0$, where B = 2.237 MHz.

The received signal, x(t), is mixed to baseband:

$$x(t) \cos 2\pi f_o t \rightleftharpoons [X(f + f_o) + X(f - f_o)]/2$$

$$\simeq \tfrac{1}{4}V(f)R(f)H_+(f + f_o) + \tfrac{1}{4}V^*(-f)R^*(-f)H_-(f - f_o) \qquad (12)$$

High frequency terms have been omitted from (12); these are filtered out. The subscripts on H_+ and H_- can be dropped at this stage. Note that H is hermite, that is, $H(-f) = H^*(-f)$. Thus the video spectrum can be written

$$U(f) = V^*(-f)R^*(-f)H(-f_o+f) + V(f)R(f)H^*(-f_o-f)$$

$$= V_1(f)\ R_1(f)\ H(-f_o+f) + V_1^*(-f)R_1^*(-f)H^*(-f_o-f) \qquad (13)$$

where $V_1(f) = V^*(-f)$ and $R_1(f) = R^*(-f)$. $V_1(f)$ is effectively the reflection in frequency of the emitted spectrum $V(f + f_o)$. A similar relationship

holds for the receiver characteristics expressed by $R_1(f)$. From the definition of R(f) the second term vanishes for $f > 0$ and the first vanishes for $f < 0$. This is simply the condition that there is no imaging in the down-conversion. It is also different from the double-sideband system of Hankins.

Using the Taylor series expansion (2) about $f_1 = -f_o$, neglecting the delay term gives

$$H(-f_o + f) = \exp[i\phi + i(-\pi/\alpha)(f^2 + f^3/f_o + f^4/f_o^2 + \ldots)] \qquad (14)$$

where $\alpha = f_o^3/2D$. During computer processing the positive half of the spectrum is multiplied by the inverse transfer function while the negative half is replaced by zero:

$$U'(f) = U(f)\ \exp[+i(\pi/\alpha)(f^2 + f^3/f_o + f^4/f_o^2)]\ ,\ f > 0 \qquad (15)$$

$$= 0\ ,\ f < 0$$

Calculation of the inverse Fourier transform of U'(f) gives an estimate not of $v_1(t) * r_1(t)$ but of the analytic signal (Bracewell 1965 page 268) corresponding to that function. The modulus of the analytic signal is the envelope of v(t). This method of obtaining it is a simulation of a base-band quadrature detector. The squared modulus is proportional to the intensity that would be measured by a radiometer with a pass-band defined by (6) in the vicinity of the pulsar. That is

$$I(t) = |u'(t)|^2 = |v(t) * r(t)|^2 \qquad (16)$$

and, if the radiation is incoherent,

$$I(t) = |v(t)|^2 * |r(t)|^2 \qquad (17)$$

That is, the emitted pulse is convolved only with the receiver resolution

function of width ~ 1/B, assuming of course that there is no "post-detection" smoothing.

The exact shape of the resolution function is determined by R(f). The responses of all parts of the system both hardware and software are incorporated in R(f). In particular R(f) can be controlled numerically during the calculation of U'(f) from U(f), the "raw" Fourier transform of the video. This control is of three types:

(i) to compensate for instrumental defects, that is, to flatten the pass-band and remove the instrumental dispersion whose measurement was discussed in Chapter 5;

(ii) to reject those regions of the spectrum U(f) affected by imaging and aliasing, i.e. below about 0.2 MHz and above about 2 MHz;

and (iii) to define narrow-band channels.

These three adjustments to the spectrum are depicted in Figure 6.2A.

Narrow-band channels are needed for observing the dynamic spectrum of the pulsar signals. That is, various configurations of multi-channel spectrum analyzers are simulated. This is done by many inverse Fourier transforms to obtain $|u'(t)|^2$ each with a different R(f). As will be discussed in the next Section the computation is efficient for narrow-band channels.

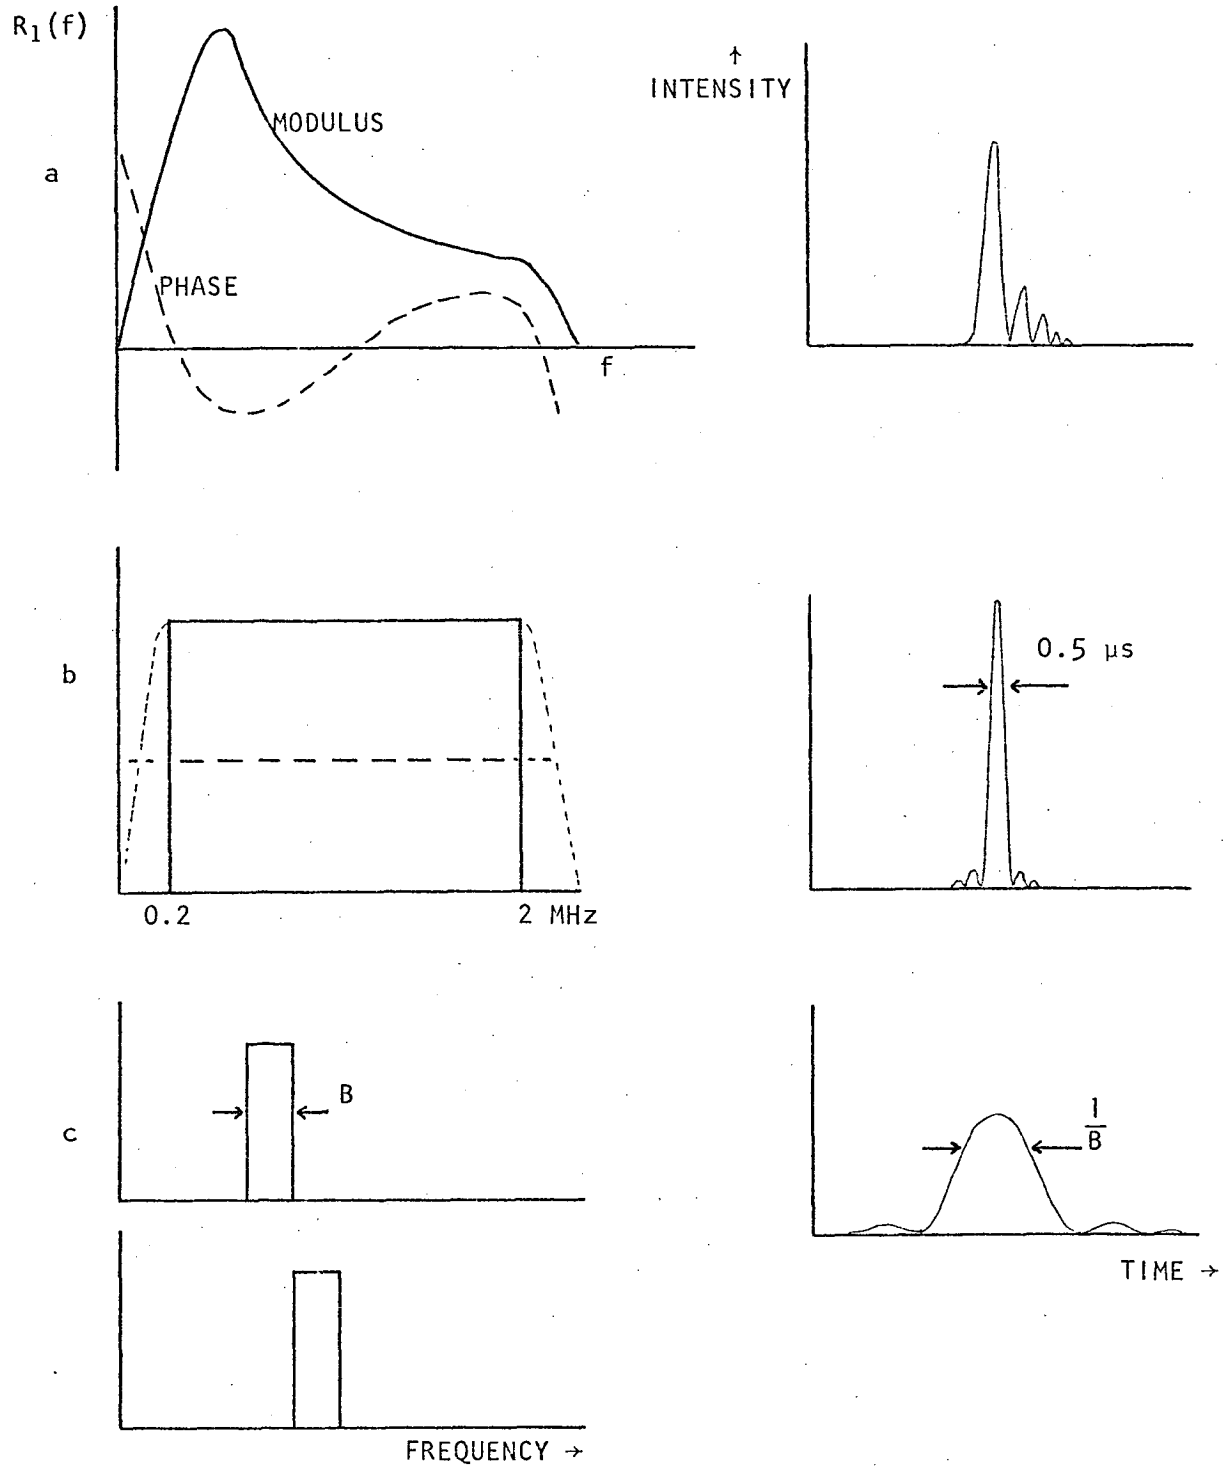

FIGURE 6.2A Numerical control of passband and the resulting intensity impulse responses (after dispersion removal) for:

(a) passband as defined by the hardware;

(b) after flattening the passband, correcting for instrumental dispersion, and rejecting the image and alias parts of the spectrum;

(c) the definition of two (for example) channels of a multi-channel spectrum analyzer. They have the same intensity impulse response.

6.3 NUMERICAL FORMULAE AND PROCEDURES

The formulae presented in Section 6.2 apply to samples of the video, u(t), taken at the recorded clock ticks, that is, at intervals δt = 224 ns. The continuous Fourier transform is approximated by the discrete Fourier transform (DFT) defined for a function of time g(t) by

$$g_j \rightleftharpoons G_k = \frac{1}{N} \sum_{j=0}^{N-1} g_j \exp(-i2\pi jk/N) \qquad (18)$$

$$g_j = \sum_{k=0}^{N-1} G_k \exp(i2\pi jk/N)$$

where $g_j = g(j\delta t)$. The spectrum is tabulated at intervals of $\delta f = 1/N\delta t$ and terms G_k, k = 0..N/2, represent frequencies 0 to f_N where $f_N = 1/2\delta t$ is the Nyquist frequency, 2.237 MHz in the present case. $T = N\delta t$ is the time window. Since the system is a sampled system it is periodic in the frequency domain with period $1/\delta t$. Therefore the terms G_k, k = N/2 .. N-1 represent the negative frequency spectrum from $-f_N$ to $-\delta f$. None of the numerical processes involved in this project require the calculation of the negative half of the spectrum.

Each DFT calculation is performed using the Fast Fourier Transform (FFT) algorithm (Cooley and Tukey 1965). With this algorithm there is a great saving in program complexity and some in computer time if N is a power of two. The dispersion removal is performed using N = 16K (K = 1024), 32K or 64K corresponding to observing windows of 3.6, 7.3 or 14.6 ms respectively. The sequence of samples (u_j, j = 0 .. N-1) is selected from within the 86K samples for one 20 ms stripe. The DFT of the N real points is performed using an extension of the original Cooley-Tukey FFT (procedure 4 of Cooley, Lewis and Welch 1970) which halves the computer time. The result is the set of (N/2) + 1 complex spectrum points U_k, k = 0 .. N/2.

6.3.1 The Dispersion Removal

Corresponding to equation (15), in which the spectrum is multiplied by the dispersion removal transfer function, we have

$$X_k = C_k U_k \exp[i(\pi/\alpha)(k^2\delta f^2 + k^3\delta f^3/f_o + k^4\delta f^4/f_o^2)] \qquad (19)$$

$$k = 0 \,..\, N/2$$

where x_j and X_k are the sampled analogues of u'(t) and U'(f) and C_k is a complex passband correction factor. C_k provides the control over R(f) that was discussed at the end of the last Section, namely, passband flattening and phase correction, alias/image rejection and the definition of narrow-band channels. C_k can be written

$$C_k = F_k \exp(-i\theta_k) \text{ for } n_1 \le k \le n_2, \qquad (20)$$

$$= 0 \text{ elsewhere.}$$

The flattening coefficients, F_k, are determined by off-pulse spectra averaged together and smoothed to 70 kHz resolution. The off-pulse data is generally taken from stripes two removed from the pulse, i.e. 40 ms away, and outside the mean pulse profile. The passband is supposed to be flat in the central 1.8 MHz with a roll-off to zero at each end. The phase coefficients, θ_k, are calculated by interpolation from the measured instrumental phase (Figure 5.2.4C). n_1 and n_2 depend on the desired total bandwidth which depends on the alias/image level that is acceptable. For N = 64K and the rejection of 210 kHz at each end of the spectrum leaving 1.8 MHz, n_1 = 3K = 3072 and n_2 = 29K.

If the data is to be treated as a single channel of bandwidth ~2 MHz then, for N = 64K, 32K complex zeroes representing the negative frequency part of the spectrum are placed after the N/2 = 32K terms X_k giving a total

workspace of 64K complex numbers. Then the performance of the inverse FFT followed by calculation of the squared-modulus of each term x_j gives the intensity time series I_j, j = 0 .. (N-1), that is, j = 0 .. (64K-1).

The intensity impulse response is

$$\mathrm{sinc}^2(t/\tau), \quad \tau = 1/B \tag{21}$$

where B ~ 2 MHz is the channel bandwidth. For B = 1.8 MHz, τ = 0.55 μs. The samples of intensity, I_j, are δt = 224 ns apart.

6.3.2 The Narrow-band Algorithm

For a multi-channel analysis of the data each narrow-band channel could be simulated using the method described above. That is, the range over which C_k is zero could be increased and many large inverse FFTs could be performed. However the following method is computationally more efficient by a factor ~(f_N/b) where b is the bandwidth of each narrow channel. Suppose such a channel is defined formally by setting $C_k = 0$ except in the range $k_o \leq k < k_o + n$, then the band-shape is rectangular, the bandwidth is $b = n\delta f$ and the centre frequency (as video) is $(k_o + n/2)\delta f = (2k_o + n)f_N/N$. Re-define the frequency origin as $k_o\delta f$, which can be done since information about the centre frequency will be lost in the baseband quadrature detector. Take the n terms X_k, follow them with n zeroes and perform a 2n complex point inverse FFT. (It is assumed for simplicity that n is a power of two and thus a submultiple of N.) The resulting intensity sequence I''_j consists of 2n points with a tabulation interval $\Delta t = (N/2n)\delta t$ and the time resolution is $\tau = 1/b = 2\Delta t$.

For example if N = 64K and n = 1024 then b = 70 kHz, Δt = 7.2 μs and τ = 14.3 μs.

Each intensity value I''_j is the same apart from a constant factor as

every (N/2n)th, i.e. every (f_N/b)th value I_j that would be obtained by the first method (N point inverse FFT). This result is an application of the "stretch theorem" (theorem 8 of Cooley, Lewis and Welch 1969). If necessary the missing points can be obtained to a good approximation by cubic interpolation, but this is not warranted for routine use because of the quantity of data. If interpolation is not used then the "picket-fence effect" (Bergland 1969) operates. That is, for the worst placement of an impulse relative to the sampling comb (with spacing Δt) the highest sample is 80% of the true peak intensity.

6.3.3 The Observing Window

The multiplication of the spectrum for dispersion removal which was expressed by equation (19) is equivalent to convolution in the time domain, and by the properties of the DFT this convolution is cyclic with period $T = N\delta t$. This means that the sequence I_j , $j = 0 .. N-1$ has a wrap-around by the duration of the impulse response of the dispersion, that is, by the bandwidth sweep time $t_s = f_N/\alpha \simeq B/\alpha$. One way of considering the effect is depicted in Figure 6.3.3A. This shows the dynamic spectrum of an impulse emitted by the pulsar before and after dispersion removal. Before dispersion removal the window onto the dynamic spectrum is T × B, e.g. 14.6 ms × 2 MHz, where B is the total bandwidth of the supposed multi-channel receiver. The dispersion removal produces a skew on this window of t_s, which amounts to 0.7 ms for PSR 0950+08 but 12.4 ms for PSR 1749-28.

If the data is analyzed with a single channel of bandwidth B then clearly a section of its simulated output of duration t_s must be rejected. With the choice of 409.025 MHz for f_o (the frequency about which the Taylor series expansion of H(f) is performed) the dispersion removal "pivots" about zero video frequency. Hence the rejected section lies at the end of the

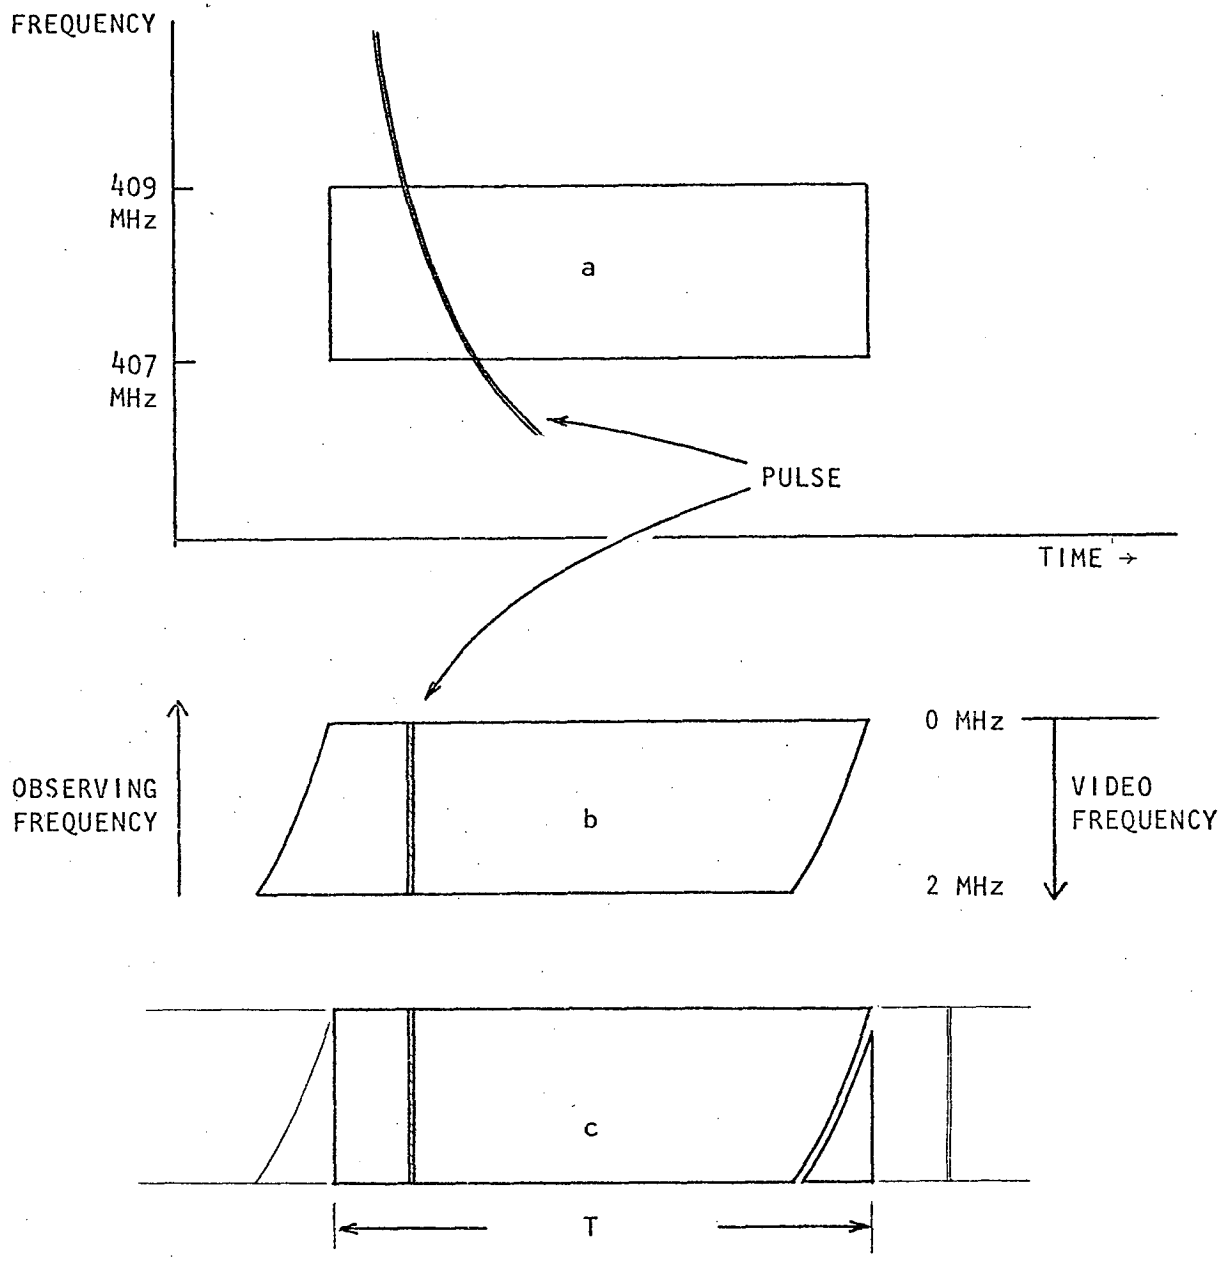

FIGURE 6.3.3A The dynamic spectrum observing window

(a) before dispersion removal (dispersion curvature exaggerated),

(b) after dispersion removal,

(c) as (b) but showing the cyclic scale of time due to the DFT ($T \leq 20$ ms).

sequence I_j or I''_j, that is, $[(T - t_s) \,..\, T]$.

When a multi-channel analysis of the data is performed the amount of data to reject is much less. For each channel the duration of the rejected section must exceed $t_s = b/\alpha$. The position of this section changes from channel to channel, following the boundary of the skewed observing window which is "folded up" as shown in the figure. (It can of course by "unfolded" during display of the data.)

6.3.4 Comparison with the Method of Hankins

Table 2 compares the dispersion removal methods described above with those of Hankins and co-workers (Hankins 1971, 1972, Rickett, Hankins and Cordes 1975, Hankins and Rickett 1975). Dual channel, including polarization, measurements have been omitted.

Described in this thesis	At Arecibo
*bandwidth $B \leq 2.25$ MHz	*bandwidth B = 125 kHz (1.25 MHz in 1976)
*sampling off-line	*sampling in real-time
*20 ms window with no control over its phase relative to the pulsar	*50 ms window (later 26 ms) centred on the mean pulse profile
*single-sideband down-conversion to video	*two channels in quadrature mixed to zero centre frequency
*sample rate 2B	*sample rate B for each of 2 channels
*dispersion removal by multiplication of the spectrum	*dispersion removal by cyclic overlapped convolution (via FFT)
*$I(t) = \|u'(t)\|^2$ available only	*$u'(t)$ and $I(t)$ available for secondary analysis
*quadrature channel defined numerically by zeroing the negative spectrum	*quadrature channel produced by hardware
*rectangular passband(s)	*passband defined by hardware only(?) †
*intensity impulse response has sidelobe level 0.05	*sidelobe level ≤ 0.01(?) †
*multiple channels defined in the frequency domain	*dynamic spectrum by calculating the power spectrum of short blocks of $u'(t)$

TABLE 2 Comparison of the observing system described here with that used at Arecibo (based on published accounts).

† It is not clear what passband flattening/shaping is done in practice.

6.3.5 Steps in the Data Processing

The following is a summary of the sequence of steps in the basic computer processing of the data for each pulse. Most of the steps have been described in the preceding Section. Brief descriptions of the programs involved are given in Appendix A.

1) Load the work-file with 62K samples representing a time window of 14 ms and subtract the approximate mean. Fill the rest of the work-file with 2048 zeroes (this is to make obvious in the plots any errors in treating the cyclic nature of the time windows rather than anything necessary for the analysis).

2) Treating the 64K numbers as 32K complex pairs compute the FFT. Then unscramble the result to give 32K +1 complex spectral values.

3) Calculate and plot the power spectrum smoothed to a resolution of 4.7 kHz. Also write to a file 32 numbers representing the power spectrum averaged in 70 kHz blocks (in the case of off-pulse data these are used for passband flattening).

4) Correct for instrumental phase by multiplying the complex spectrum by exp[-iP(f)] where P is a cubic interpolation between measured phase values.

5) For each simulated channel, say of the bandwidth of 70 kHz, perform the following:

   (i) Load a block of 2048 complex spectral values into an array.

   (ii) Flatten the passband by multiplying by factors obtained from off-pulse data.

   (iii) Remove dispersion by multiplying by exp[-iK(f)] where K is a fourth degree polynomial.

   (iv) Zero out the first element of the array (representing the DC

term) and fill the top half of the array with 2048 complex zeroes (the negative frequencies).

(v) Perform a 4K inverse FFT.

(vi) Calculate the squared modulus of each element giving 4096 intensity values.

(vii) Calculate the mean, second moment and maximum (for scaling and general information).

(viii) Plot, possibly smoothed with an 11 point window, the 4096 points of the trace. Accumulate the average trace, i.e. add across channels.

(ix) Calculate the running-mean over a rectangular window (with 2-point tapering at each end) and plot a multiple of this on top of the trace. Accumulate a histogram of relative intensity, i.e. relative to the local mean.

(x) Write the trace to a file and various housekeeping and summary information to a separate smaller file.

6) Write the average trace to a file, i.e. one trace.

7) Plot the average trace, possibly smoothed, which completes the basic multi-channel calculation for a pulse.

For dispersion removal of a single channel of bandwidth ~ 2 MHz step 5 involves the loading of about 28K complex spectral values (the exact number depends on how much of the image and alias regions are to be rejected) and a 64K inverse FFT. The result is a time sequence of 64K points. It is never necessary to plot all of it as a section of duration equal to the bandwidth sweep time of the pulsar is invalid. Note that the plotter used has a resolution of .01 inch so that unsmoothed plots can be very long.

For the case of off-pulse data, steps 4 to 7 are only needed if an auto-correlation analysis calibration is required. The plotting in steps 5 (viii)

and (ix) can be omitted. Most off-pulse runs are only to obtain the off-pulse spectrum. For this purpose windows of 16K sample points are sufficient. Spectra tabulated at 32 points (i.e. at 70 kHz intervals) as a result of step 3 are averaged between windows within a stripe and between stripes. These stripes are near the pulses being analyzed. Cubic spline coefficients are calculated for the average off-pulse spectrum and used for flattening the central 1.8 MHz of the passband (step 5(ii)).

Secondary analyses include auto- and cross-correlations of the traces of a pulse, auto-correlation of the power spectrum (strictly speaking the periodogram) which is obtained from the complex spectrum (step 2) by band-pass flattening and taking the squared modulus. These correlation analyses are described in Chapter 9. The dynamic spectrum is usually obtained with a resolution of 140 kHz × 7 μs. To give a display more acceptable to the eye there must be some overlap between neighbouring channels. Therefore a standard multi-channel configuration is to divide a total bandwidth of 2.0 MHz into 56 channels of bandwidth 140 kHz and a 75% overlap, that is, centre frequencies displaced by 35 kHz. The display method is a "z-plot" on a lineprinter, that is, character positions on the page represent the point in frequency-time space with the intensity indicated by the blackness provided by overprinting. This multi-channel configuration and display method has not been satisfactory for PSR 1749-28, in particular, since the bandwidth sweep time requires a 14.6 ms window and hence charts 13m long.

6.4 IMPEDIMENTS TO THE DISPERSION REMOVAL

6.4.1 Errors in Assumed DM and Phase

There are three aspects of the observing and analysis system which can impair the time resolution to be obtained by pre-detection dispersion removal. These are instrumental dispersion, errors in the assumed dispersion measure (or equivalently errors in the centre frequency of the observations), and numerical errors in the software.

The dispersion removal programs were tested by numerically generating samples of a dispersed impulse. This signal was a sinusoid whose frequency, f_j, at time $t_j = j\delta t$ was obtained from a series expansion of

$$f_j = -f_o[(1 + f_o^2 t_j/D)^{-\frac{1}{2}} - 1]$$

where f_o = 409.025 MHz. This formula includes the curvature in the dispersion, and f_j rises with increasing time in the same way as signals recorded at Molonglo. The amplitude of the sinusoid was tapered below about 0.1 MHz and above about 2.1 MHz. I performed single and multi-channel analyses of this test data and much of the material in this Section is based upon them.

Figure 6.4.1A shows the result of a simulation of an impulse from PSR 0950+08 passing through four adjacent rectangular channels of 280 kHz bandwidth. For this bandwidth, and in fact for all bandwidths where the error in DM is being considered, the linear sweep approximation is adequate and so the intensity impulse response is accurately described by a Fresnel diffraction pattern. Simulations such as shown in the Figure were compared with published single-slit diffraction patterns (e.g. Stone 1960). Note that 5% of the energy falls outside the "geometric" width of the pattern, which is the bandwidth sweep time $t_s = B/\alpha$. When the dispersion removal is perfect the intensity impulse response is a sinc^2 curve (equation (21)) and

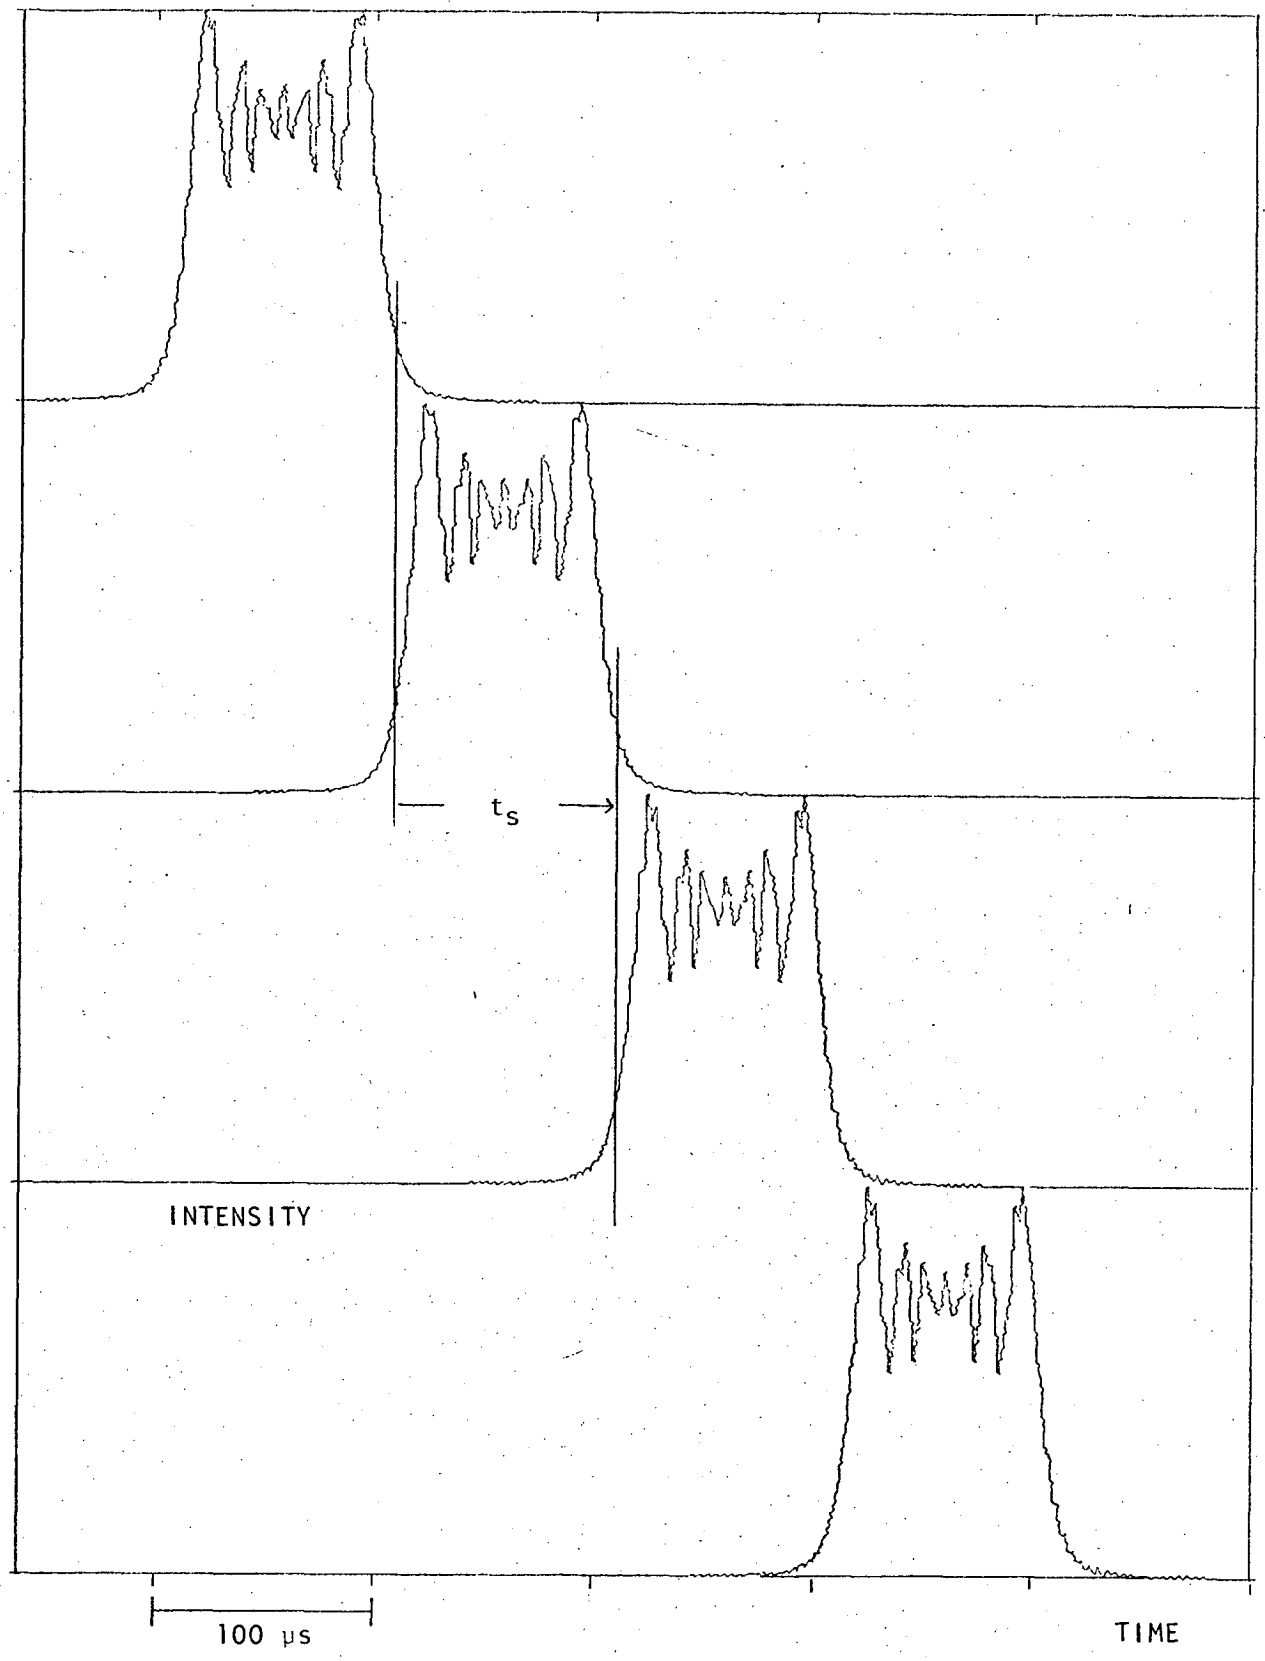

FIGURE 6.4.1A

Simulation of a dispersed impulse (DM = 2.97 $cm^{-3}pc$) passing through four contiguous channels with rectangular bandshapes and bandwidths B=280 kHz. Observing frequency f_c = 408 MHz. The bandwidth sweep time is t_s = 102 μs. (The results scale as B^2DM/f_c^3 = constant, i.e. Bt_s = constant.)

93% of the energy is in the central lobe, i.e. within a time interval 2/B. The response for a 1.1 MHz channel is shown in Figure 6.4.1B(i).

When there is an error, ΔDM, in the assumed value for DM then there is a residual frequency sweep time $\Delta t_s = 2B\Delta D/f_c^3$ where f_c is the centre observing frequency and ΔD is proportional to ΔDM according to equation (3). As ΔDM is increased from zero the intensity impulse response changes from a narrow $sinc^2$ curve to a multi-peaked Fresnel pattern. This is illustrated in Figure 6.4.1B(ii) for a 2.1 MHz channel. For the smaller values of ΔDM the most obvious effect is a decrease in peak height. This is expressed by a "DM sensitivity curve" (Figure 6.4.1C). The tolerance on DM for a 20% reduction in peak height, which is comparable with the picket-fence effect (Section 6.3.2), is ± 0.0037 $cm^{-3}pc$ and there is a 50% reduction for $\Delta DM = \pm 0.0065$ $cm^{-3}pc$. These figures relate to B = 2.1 MHz and f_c = 408 MHz. The results scale according to $B^2\Delta DM/f_c^3$ = constant.

The 20% reduction in peak height corresponds to a quadratic phase error of $\pi/2$ from band centre to edge. The tolerance can be written $|\Delta t_s| < 2/B$ and may be compared with the best time resolution for post-detection dispersion removal, quoted in Section 6.1, namely $t_{min} = t_s = 1/b$. Since the 20% reduction in height is associated with an increase in width, i.e. resolution time, of less than 5%, the channel bandwidths given in Table 1 can in practice be doubled to approximately halve the resolving time of the post-detection method.

The condition $|\Delta t_s| < 2/B$ is the tolerance suggested by Hankins and Rickett (1975) (however they did not express this clearly). Their "ambiguity diagram", Figure 5, shows in detail the degradation of the intensity impulse response as $|\Delta DM|$ increases. However the dispersion coefficient scale on this diagram is misleading and the caption incorrect. The diagram is in fact specific to observations of PSR 0950+08 at 111.5 MHz.

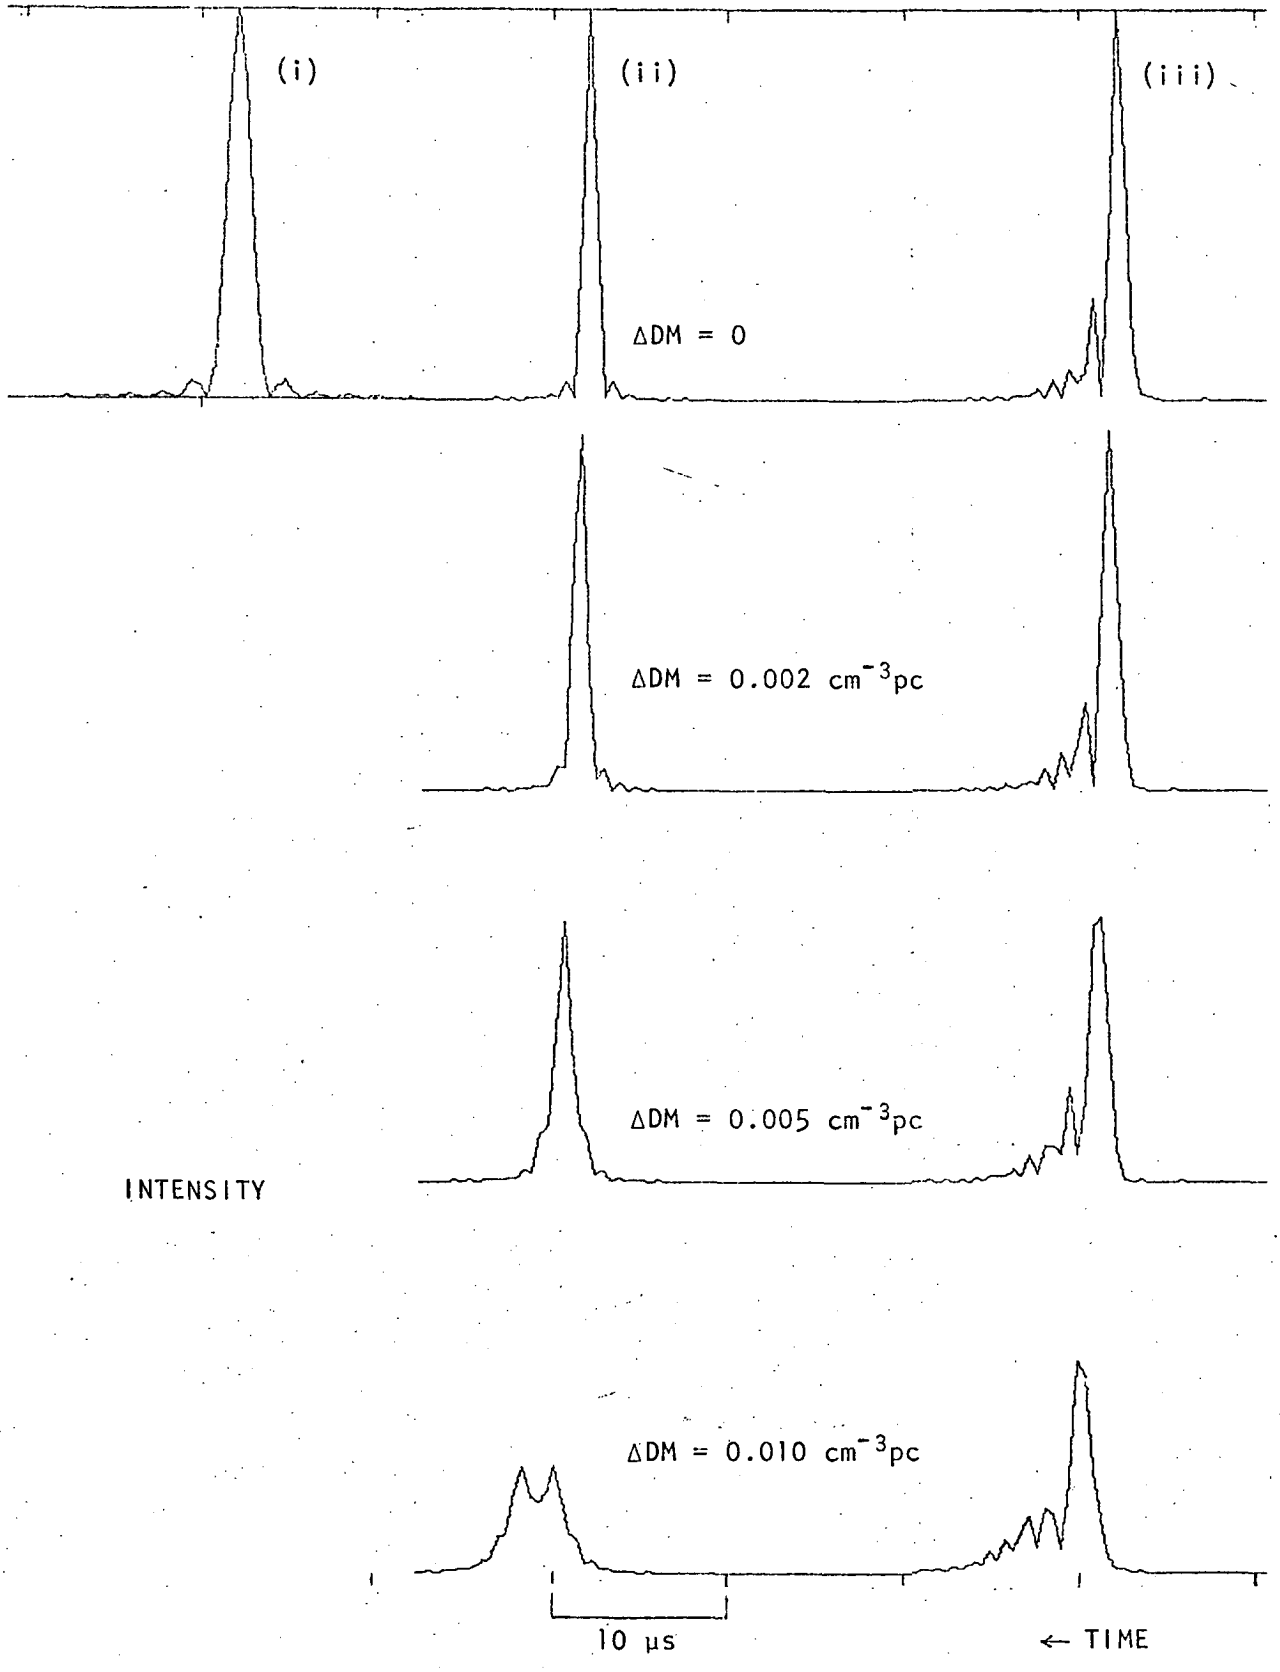

FIGURE 6.4.1B

Theoretical intensity impulse responses for a rectangular passband and various errors, ΔDM, in the assumed dispersion measure of the pulsar signal. Observing frequency f_c = 408 MHz.

(i) bandwidth 1.1 MHz; (ii) bandwidth 2.1 MHz;
(iii) as (ii) but with a different vertical scale and with the instrumental phase response imposed on the spectrum.

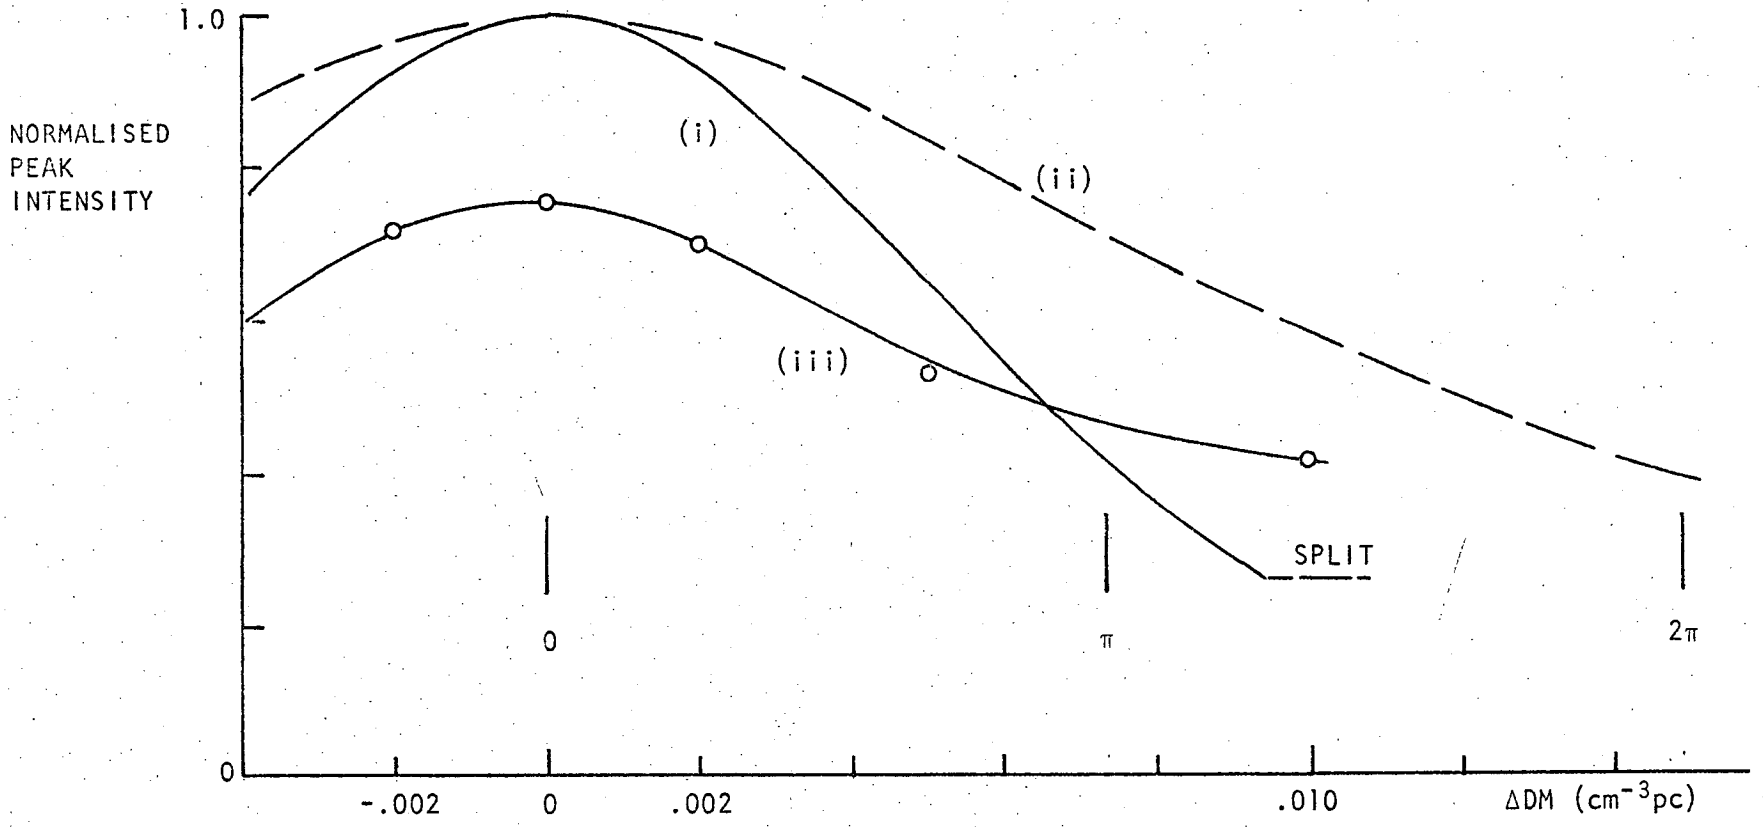

FIGURE 6.4.1C DM sensitivity curves, i.e. variation in the peak of the intensity impulse response with error, ΔDM, in the assumed dispersion measure. The curves are for bandwidth B = 2.1 MHz, f_c = 408 MHz (they scale as $B^2\Delta DM/f_c^3$ = constant).

(i) rectangular bandshape;
(ii) cosine taper;
(iii) rectangular with the measured instrumental phase (same normalisation as (i)).
Phase errors due to ΔDM from band centre to edge are shown.

There is an absolute restriction on ΔD not a relative one, $\Delta D/D$, as implied by the scale. The figure caption is also interesting in that it quotes a best peak-to-sidelobe ratio of 67 which is much better than the ratio 21 applicable to a sinc^2 function. This difference is apparently due to truncation of Hankins' test data and inverse filter in the time domain which produces spectra that are not strictly confined to a rectangular band. In contrast my test data and dispersion removal transfer function, which is defined by equation (19), are cut off sharply in the frequency domain by the definition of a rectangular channel. A sinc^2 function is the result.

The DM sensitivity can be reduced by tapering the passband. For example, Klauder *et al* (1960) suggested the Taylor window, $1 + 0.84 \cos 2\pi(f - f_c)/B$. This approximately doubles the allowed ΔDM (see curve (ii) of Figure 6.4.1C), and reduces the sidelobe level, but increases the resolving time by 30%. However I have consistently used rectangular channels for all analyses of the observations. This is partly for convenience but also because when averaging over channels no data is wasted and the channels are statistically independent (on a microsecond scale) if the channels are rectangular and contiguous. In retrospect, this decision was wrong because the autocorrelation analyses described in Chapter 9 would have been easier if there were no sidelobes.

Figure 6.4.1C(iii) shows the effect upon the intensity impulse response of not correcting for instrumental phase. One effect is the appearance of a series of fringes or sidelobes on the trailing side of the main lobe. (Effectively these curves are plotted with time reversed - the instrumental effect *is* causal.) Since part of the energy is transferred to these sidelobes, the peak height is reduced by about 25% for $\Delta DM = 0$ (curve (iii) in Figure 6.4.1C). The 25% reduction is for a 2.1 MHz channel; for a 1.1 MHz channel centred in the observation band, i.e. $f_c = 408$ MHz, there is no

significant effect in not correcting for instrumental phase.

In conclusion, pre-detection dispersion removal with a 2 MHz rectangular passband at 408 MHz requires careful control over instrumental dispersion and accurate knowledge of the dispersion measure, namely $|\Delta DM| \sim 0.005$ cm^{-3}pc. This DM accuracy is attainable by normal pulse-timing procedures and has been attained for the two low dispersion pulsars PSR 0950+08 and PSR 1133+16. The values for the other three pulsars in Table 1 have not been published with this accuracy. However for these pulsars interstellar scattering (next Section) precludes meaningful dispersion removal with maximum time resolution over a 2 MHz bandwidth and the accuracy of the DM values is more than adequate for the bandwidths to which dispersion removal is applicable.

In principle pre-detection dispersion removal with a 2 MHz bandwidth could be used to measure DM. However the method is not practical as it requires the existence of isolated impulses from the pulsar. Such emissions have yet to be observed from any pulsar and the amplitude-modulated noise model assumes that they are unobservable - see Chapter 8.

6.4.2 Interstellar Scattering and Faraday Rotation

The effect of interstellar scattering of pulsar signals has been reviewed by Manchester and Taylor (1977). The topic is important for pre-detection dispersion removal because the time resolution is degraded to a time characteristic of the scattering by, for example, the loss of phase coherence across the passband.

As a result of the scattering the spectrum becomes modulated by an irregular pattern of interference fringes with frequency scales generally less than 1 MHz. The fringes vary on timescales of minutes and therefore can be assumed to be constant for each pulsar transit at Molonglo. The mean fringe width is measured by the decorrelation frequency, Δf, defined as

follows. If S(f) is the complex modulation function representing the interstellar scintillation then fluctuations in its power spectrum $|S(f)|^2$ are characterized by the autocorrelation function $\rho = \rho_{|S|^2}(\nu)$. Δf is the frequency difference, ν, at which $\rho = 0.5$. For pulsars for which the scattering is prominent $\Delta f \propto f_0^4$, where f_0 is the observing frequency. Although defined in terms of the power spectrum Δf is also the characteristic width in frequency of fluctuations in phase, provided that S is a well-behaved random variable (Hankins and Rickett 1975). This assumption is almost certainly violated for low dispersion pulsars for which the scattering process involves a small number of rays.

Whether the modulation is of phase or amplitude or both the effect on pre-detection dispersion removal must be the same : energy is transferred from the main part of the impulse response to some kind of sidelobe structure that follows in time (i.e. it is a causal system). This structure is constant over minutes of time but is random, being determined by the particular modulating function, S(f), that is present. Except in the case of little phase modulation and one fringe in the passband which behaves like a channel of narrower bandwidth, the structure in the impulse response has a characteristic timescale ~ 1/B, where B is the bandwidth of the receiver channel. In the case of strong, multiple scattering the mean impulse response, averaged over many changes in the fringes or seen as an envelope when there are many fringes in the passband, is the same as the pulse broadening function. This is the function which is convolved with the mean pulse profile and is, for most pulses for which it has been observed, a truncated exponential, that is, $\exp(-t/\tau)$ where $\tau \simeq 1/2\pi\Delta f$ (Sutton 1971).

To predict the likely effect of the scattering on dispersion removal the bandwidth, B, may be compared with Δf. Values of Δf at 408 MHz for five pulsars are shown in Table 2 and for two of these pulsars the situation is

clear. Firstly with $\Delta f \sim 50$ Hz pulsar 0833-45 is completely incompatible with pre-detection dispersion removal. Secondly, interstellar scattering has no effect on dispersion removal for PSR 0950+08 if $B \sim 2$ MHz. PSR 1133+16 is more problematic with the best estimate of Δf being 2 MHz but with a large variation in Δf values between workers. Such variation would be expected if a small number of rays were involved. Therefore for low dispersion pulsars like PSR 1133+16 the power spectrum observed at a particular epoch must be inspected to see if there are in fact any scintillation fringes in the passband.

For the higher dispersion pulsars PSR 1642-03 and PSR 1749-28 many scintillation fringes are expected across 2 MHz. A suitable strategy would be to divide the total bandwidth into a number of channels each containing at most one fringe. This suggests a channel bandwidth of $b \sim 2\Delta f$ giving a time resolution $\sim 1/2\Delta f$, however this time is longer than the scattering time τ. Therefore a revision of b upwards to $\sim 6\Delta f$ is suggested. This should be compared with the decorrelation bandwidth, B_h, which is defined as that bandwidth at which the modulation index of the scintillation is half that for zero bandwidth. Sutton (1971) used the relationship $B_h \simeq 10\Delta f$. For PSR 1749-28, $6\Delta f \sim 90$ kHz and I have used only 70 kHz channels for analyzing this pulsar. The same bandwidth was used for two very weak pulses from PSR 1642-03 for which $6\Delta f \sim 400$ kHz.

Since the Molonglo antenna is linearly polarized the observations could be affected by interstellar Faraday rotation. However, from the rotation measures published by Manchester and Taylor (1977), the expected fringe spacing for all five pulsars is greater than 10 MHz near 408 MHz so that Faraday rotation can be ignored.

PSR	DM($cm^{-3}pc$)	Δf	comments
0950+08	2.97	24 MHz	
1133+16	4.85	2 MHz	range 0.7 - 5
1642-03	35.71	~ 70 kHz	Manchester + Taylor (1977)
1749-28	50.88	16 kHz	
0833-45	69.08	50 Hz	Backer (1974)

TABLE 3 Values of DM and decorrelation frequency Δf at 408 MHz. Except as noted the values of Δf are those adopted by Sutton (1971) scaled from 318 MHz according to $\Delta f \propto f^4$. Probable errors in Δf are factors of 2.

CHAPTER 7

GENERAL, HIGH TIME RESOLUTION AND DYNAMIC SPECTRUM OBSERVATIONS

This Chapter describes the high time resolution and dynamic spectrum observations of 22 pulses from four pulsars. Section 7.1 deals with one pulse from PSR 0950+08, designated pulse "A", to illustrate the processing and the forms of presentation of the observations as well as their typical content. This was the only strong pulse from this pulsar recorded on June 18th, 1975. The observations of the other seven (H, K, P, T and W recorded on June 19th; and C and D on June 20th) are presented in Section 7.2.

The observations of nine pulses from PSR 1749-28 are discussed in Section 7.3. These were all recorded in a single transit and some are consecutive. Strong interstellar scintillation fringes were found in their spectra.

Three pulses from PSR 1133+16 and PSR 1642-03 are discussed in Section 7.4. Since these were very weak the observations have not been analyzed in full.

Section 7.5 contains discussion of the observations in general while Section 7.6 deals with the performance of the observing system as revealed by the observations.

7.1 ONE PSR 0950+08 PULSE IN DETAIL

Figure 7.1A shows the power spectrum of pulse A from PSR 0950+08. The spectrum was calculated from a time window of 3.6 ms, that is, from 16000 samples. Like all power spectra shown in this Chapter it is plotted against video frequency - zero corresponds to an observing frequency of 409 MHz and 2 MHz corresponds to 407 MHz. This spectrum is about 3 db above the off-pulse spectrum when averaged across the central 2 MHz. It is smooth in the sense that there are no structures on a scale <0.5 MHz which cannot be attributed to the response of the equipment or statistical fluctuations. For example, each point in the curve having 20 kHz resolution has about 150 degrees of freedom. (Gaussian noise is assumed which leads to chi-squared statistics.) Therefore each point has an rms error of 0.15 db. The smoother curve in Figure 7.1A has a 70 kHz resolution and is shown on a linear scale in Figure 7.1B. This Figure also shows the spectrum after passband flattening.

The flattening coefficients were derived from spectra of off-pulse data and in this case two off-pulse stripes were used. For each of these, five power spectra were calculated with time windows of 3.6 ms and all ten spectra were smoothed to 70 kHz and averaged. This gave 5000 degrees of freedom per off-pulse spectral point which would suggest an rms error of 0.09 db. However there were differences in trend between the ten spectra so that the mean off-pulse spectrum and hence the power spectrum of pulse A after flattening may have systematic errors of ±0.15 db or 3.5%.

The flattened spectrum shown in Figure 7.1B has a strong rising trend with increasing video frequency. This trend is typical of all the pulses and will be discussed in Section 7.6. The flattened spectrum of pulse A exhibits a signal-to-system noise ratio of 1.0 averaged across the central 2 MHz. This ratio is for a time window of 3.6 ms. As shown in Figure 7.1C the pulse is in fact about 5 ms wide.

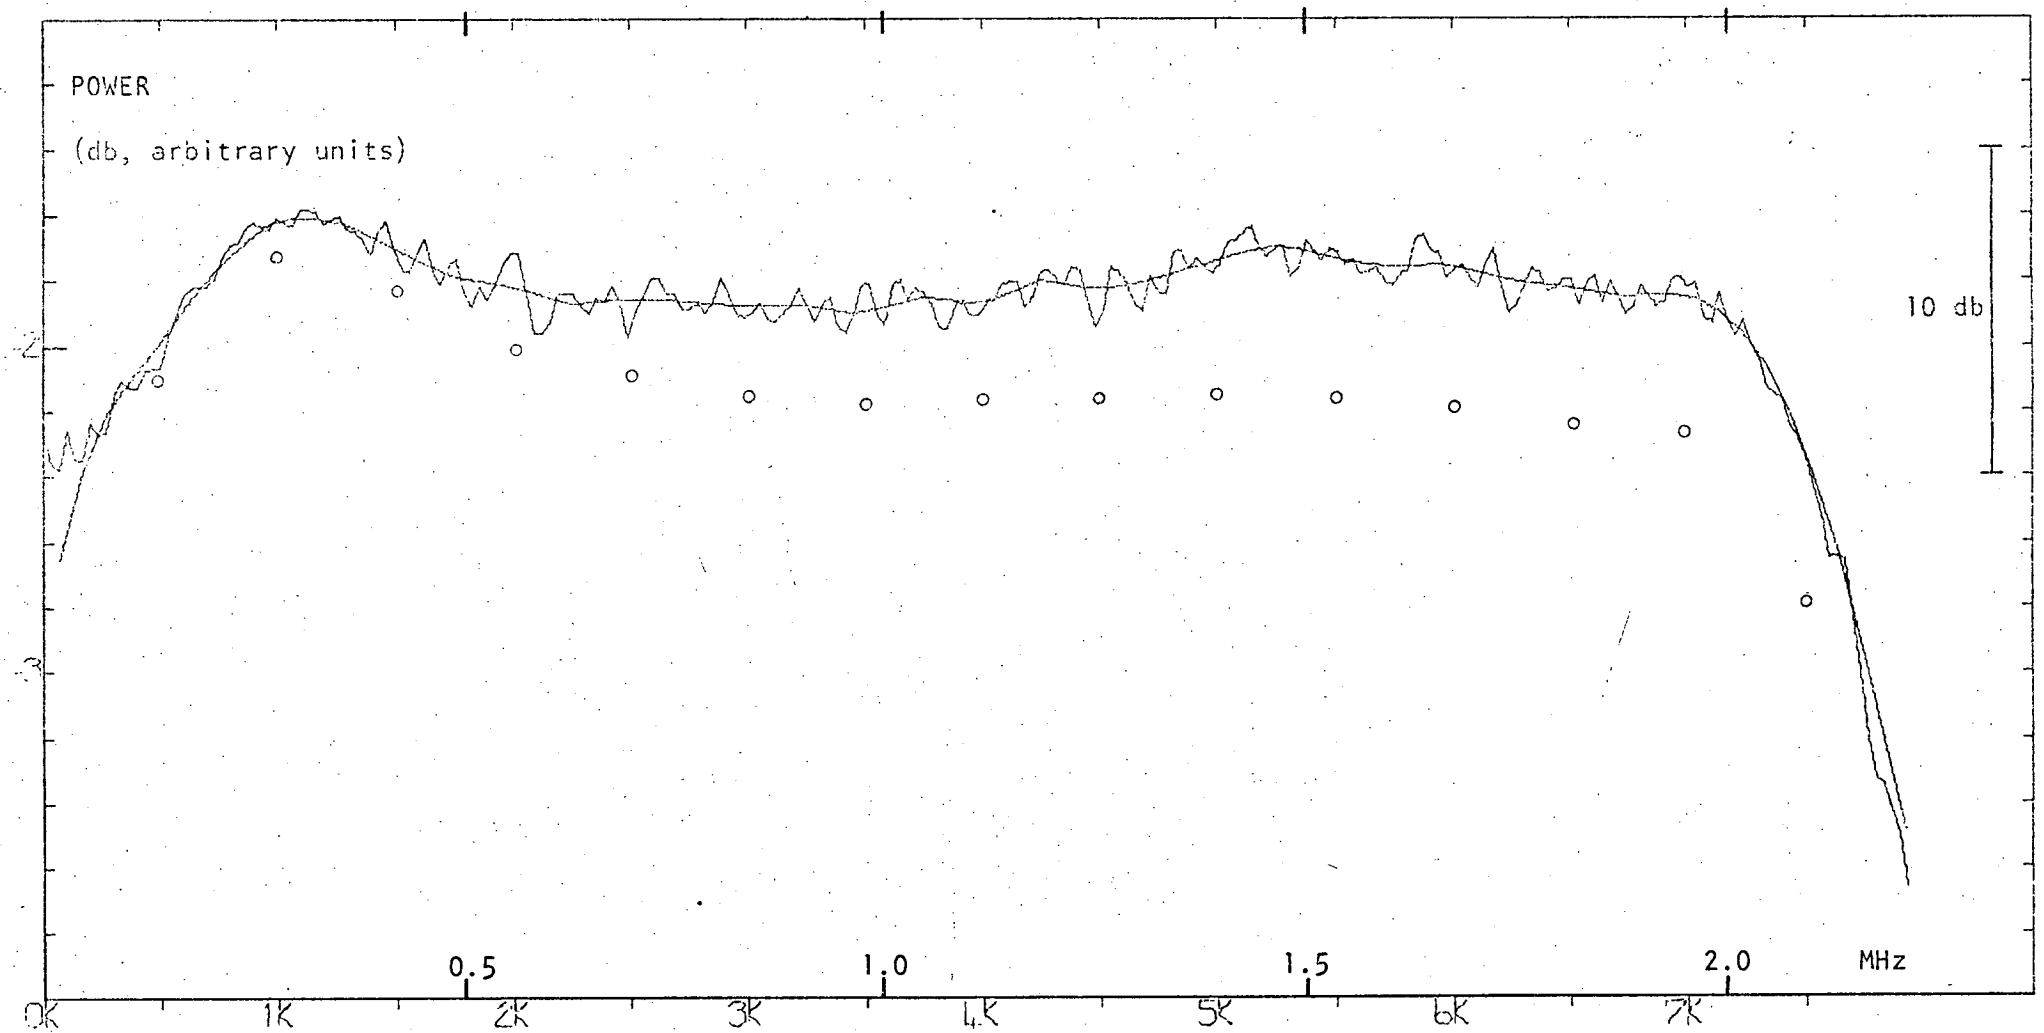

FIGURE 7.1A Power spectrum of pulse A of PSR 0950+08 at resolutions of 20 kHz and 70 kHz (curves). The spectrum was calculated from a time window of 3.6 ms. The mean off-pulse spectrum is indicated by circles.

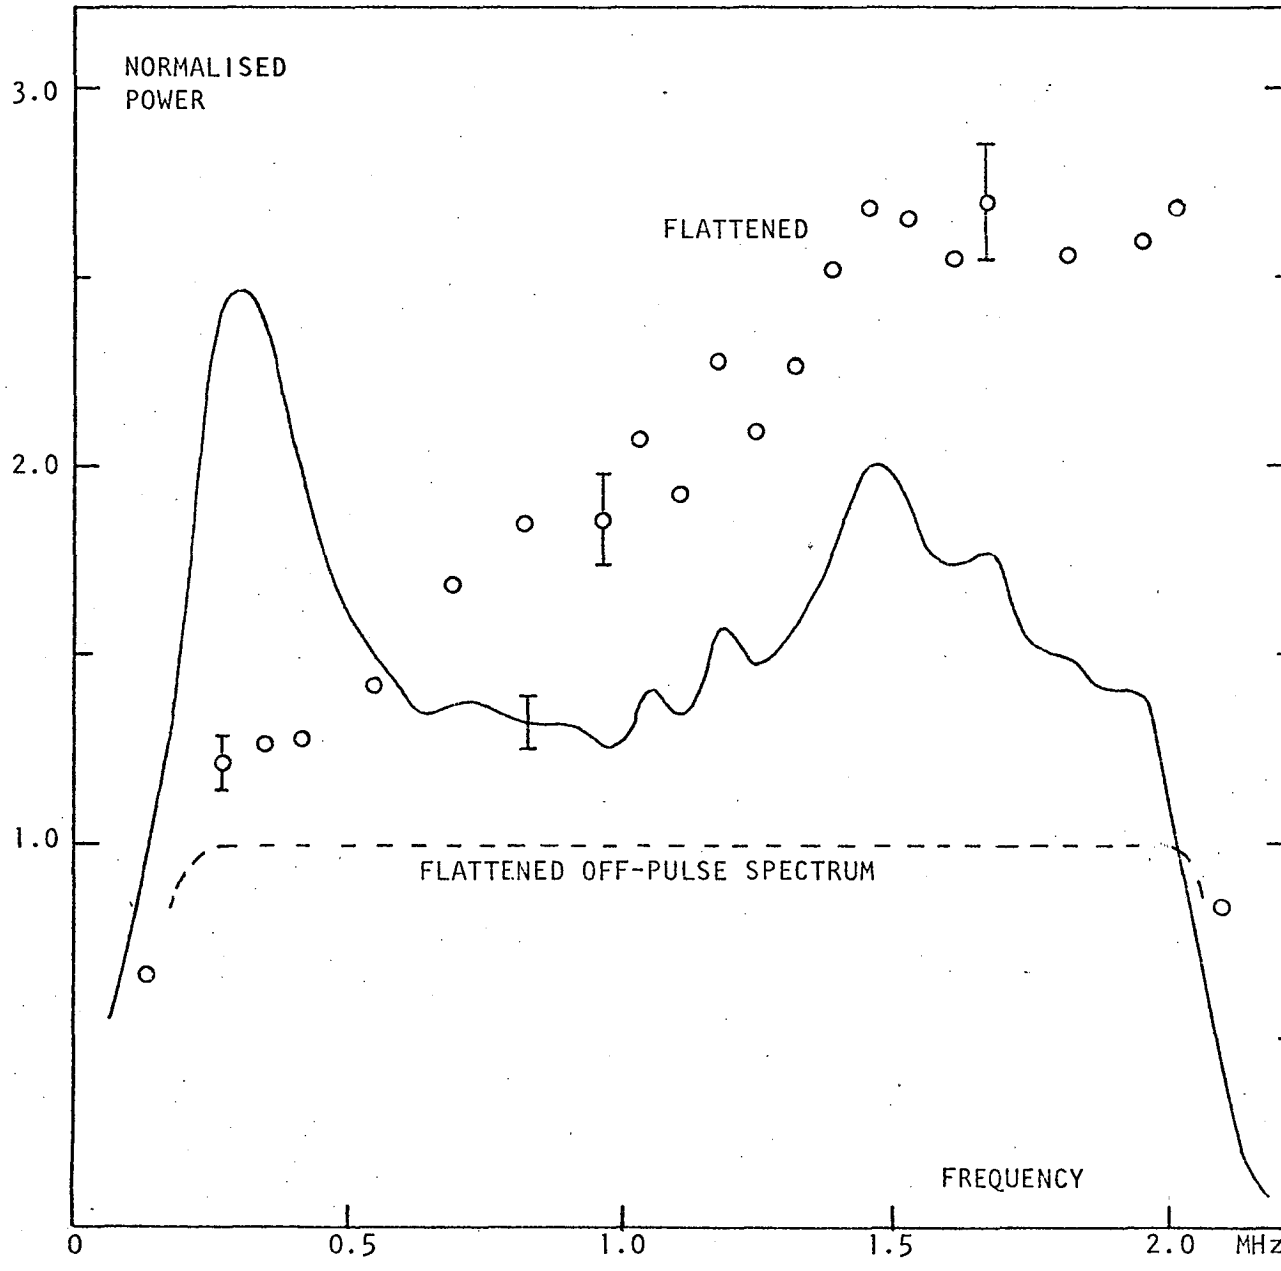

FIGURE 7.1B Power spectrum of pulse A of PSR 0950+08 plotted on a linear scale before (curve) and after (circles) passband flattening. Time window = 3 ms. Typical 1σ error bars are marked (512 degrees of freedom).

Figure 7.1C shows the average over channels of two multi-channel "observations" of the pulse, that is, two intensity profiles of pulse A. In one case dispersion has been removed as described in Chapter 6 while in the other it has not (i.e. assumed DM = 0). In both cases thirteen 140 kHz channels were used. They were rectangular and contiguous, that is, each had a rectangular bandshape and they were non-overlapping so that the total bandwidth was 1.8 MHz. The "detector output" of each was smoothed with a bell-shaped window to a resolution of 16 μs at half intensity. This is a standard configuration of the simulated receiving system, except that for PSR 0950+08 it is more usual, for reasons of computing cost, to use a time window of 7.3 ms rather than the 14.6 ms window shown in the Figure. Both intensity profiles in the Figure are accompanied by smoother curves which are multiples of the running-mean over a 220 μs window. As will be discussed, similar curves are used as confidence limits on the intensity fluctuations. The level of off-pulse intensity can be gauged from the last third of each trace (but not the last millisecond which contains artefacts). This level provides the only measurement of absolute intensity in the observing system; it is assumed to be 125 ± 5 Jansky.

The intensity profile for which the dispersion has been removed has a peak whose flux density is 480 ± 20 Jansky and with a duration of about 70 μs. When dispersion is not removed the energy of this peak is expected to be smeared over the sweep time for a bandwidth of 1.8 MHz, namely 660 μs and this is consistent with the other intensity profile in the figure. Since the peak has, in fact, an intrinsic width of 70 μs post-detection dispersion removal would be expected to produce similar results to those of the pre-detection method. This has been tried by reconfiguring the simulated multi-channel receiver. The best resolution for the post-detection method is ~20 μs using channel bandwidths in the range 50 to 100 kHz (Table 1 of Chapter 6 and Section 6.4.1). Using 140 kHz channels I obtained the expected result, namely a peak duration of about 120 μs.

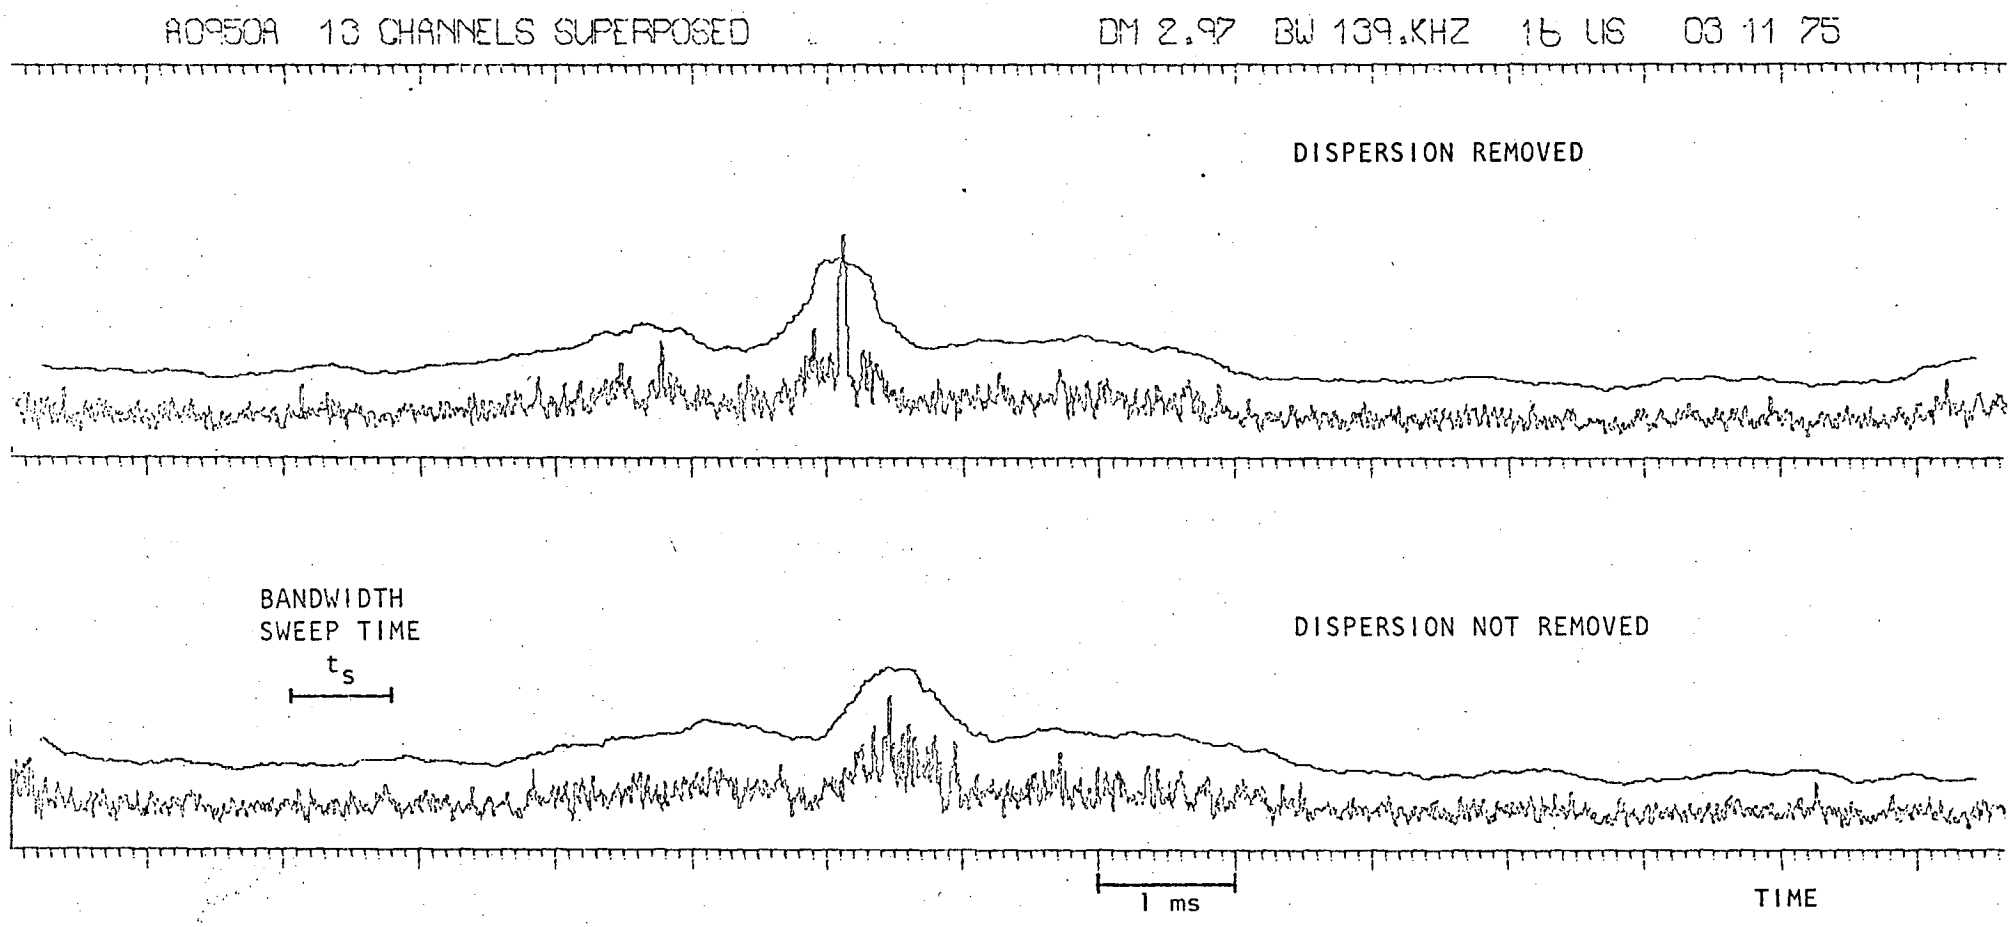

FIGURE 7.1C Intensity profiles of pulse A of PSR 0950+08 with and without dispersion removal. Each is the average over 13 channels of bandwidth 140 kHz with smoothing to a time resolution of 16 μs. The smoother curves are multiples of the running-mean over 220 μs.

The emergence of the central peak in pulse A with dispersion removal is one of three observations which indicate that the dispersion removal is approximately correct. The others are the observations of similar peaks in the intensity profiles (smoothed to a time resolution ~ 10 μs) of pulse T and to a lesser extent of pulse W, both of PSR 0950+08. However the peak of pulse A is more isolated and has sharper edges so that the observations are easier to interpret.

Figure 7.1D shows the central peak of pulse A observed with a single 2.1 MHz channel. Dispersion has been removed (unless otherwise stated this is so for all the plots). There are two curves (traces) in this Figure, one showing the best resolution of 0.5 μs and the other showing the same data smoothed to 1.8 μs resolution (half intensity). The pulsar emission is clearly consistent in appearance with amplitude-modulated noise. The modulating waveform, or function, appears to be a rectangular pulse of duration (70 ± 5) μs.

As a test of whether the data is well described by the AMN model I have routinely calculated the running-mean of every unsmoothed intensity trace and compared the intensity fluctuations with this. This is the method used by Hankins (1972) and is discussed further in Chapter 8. Its basis is briefly as follows. The AMN model predicts that the distribution of relative intensity, that is, intensity divided by the modulating function, would be exponential. The probability of a particular point having a relative intensity value, I', exceeding x is $\exp(-x)$. There are two problems in using the running-mean as an estimate of the modulating function: first, the running-mean must be over a sufficiently wide window that it has enough degrees of freedom for the statistical fluctuations to be small and for the contribution of the point being tested to the mean to be small; second, the window must not be wider (in duration) than the actual fluctuations in the modulating function. The requirements conflict

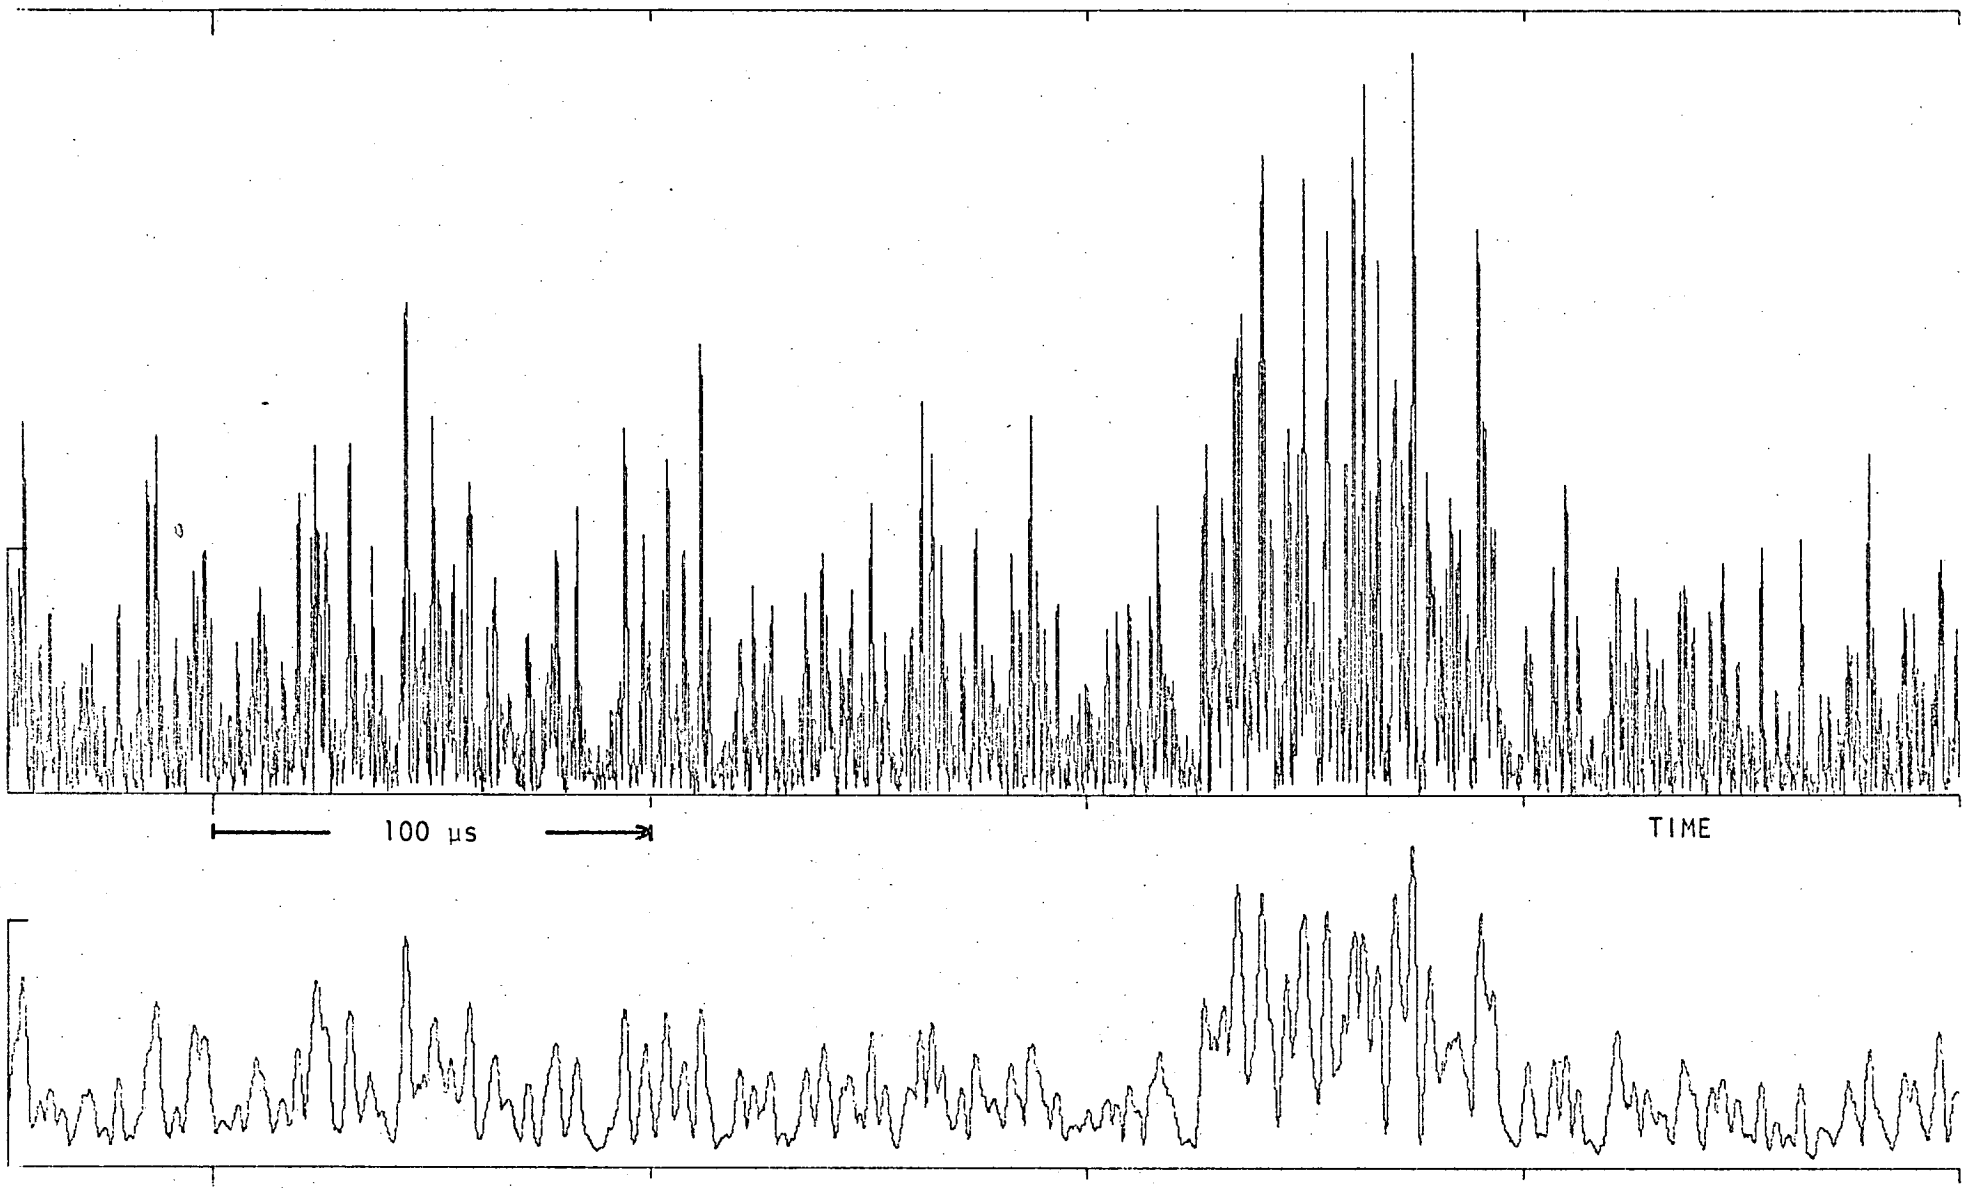

FIGURE 7.1D The strongest part of the intensity profile of PSR 0950+08 pulse A observed with a 2.1 MHz bandwidth. The upper trace has resolution 0.5 μs; the lower (on same scale) is smoothed to 1.8 μs. Samples are 0.22 μs apart.

and, at least for PSR 0950+08, can only be satisfied with total bandwidths >1 MHz. In the case of pulse A the window should not be more than ~100 μs wide because the strong central peak is about 70 μs wide. When, for example, a running-mean is calculated over 50 μs the 2.1 MHz data has a relative intensity which fits the exponential distribution well. In particular there are no large impulses and the highest I' value is 8. Thus this test shows the data to be consistent with the AMN model.

Figure 7.1E shows twelve of the thirteen intensity traces that were averaged to form the "dispersion removed" profile in Figure 7.1C. Before dealing with the peak alignments in this Figure I will continue the discussion of the last paragraph, that is, discuss the data in the framework of the AMN model.

Each channel has a bandwidth, B, of 140 kHz so that the resolution was 1/B or 7.2 μs but has been smoothed to 16 μs. The sample interval is 1/2B or 3.6 μs so that every second unsmoothed point was independent. Therefore each smoothed point has ~8 degrees of freedom and so has a 1σ uncertainty of about 50%. The running-mean over 100 μs has been calculated for each trace and a multiple (~4) of it is shown in the Figure. Since the modulating function is supposed to be independent of frequency the running-mean should be strongly correlated from one channel to another. This is generally so, but there are some anomalies such as event 'X' marked in the Figure (7.1E). However these anomalies can be explained as statistical fluctuations, because each running-mean has about 30 degrees of freedom and hence has a statistical uncertainty of about 25% (1σ). For this reason alone it is a poor estimate of the modulating function irrespective of any argument about the correct timescale. Nevertheless the running-mean of each trace was used to check the distribution of relative intensity (of unsmoothed data) and the number of points with $I' > 7$ was found to be consistent with off-pulse data (but not with the exponential distribution

FIGURE 7.1E Pulse A of PSR 0950+08 as observed with 12 contiguous channels of 140 kHz bandwidth and rectangular bandshapes. The 12 intensity traces have been smoothed to 16 μs resolution (about 4 degrees of freedom per sample point) and samples are 3.6 μs apart. The observing frequency increases up the page with the channel whose trace is at the bottom having centre frequency 407.1 MHz and at the top 408.6 MHz. The smoother curves are 3.7 times the running-mean over 220 μs. The vertical scale of the traces are not the same; the scale of each can be judged from the (apparently) off-pulse background in the first 0.5 ms of each trace. The feature marked "X" is mentioned in the text.

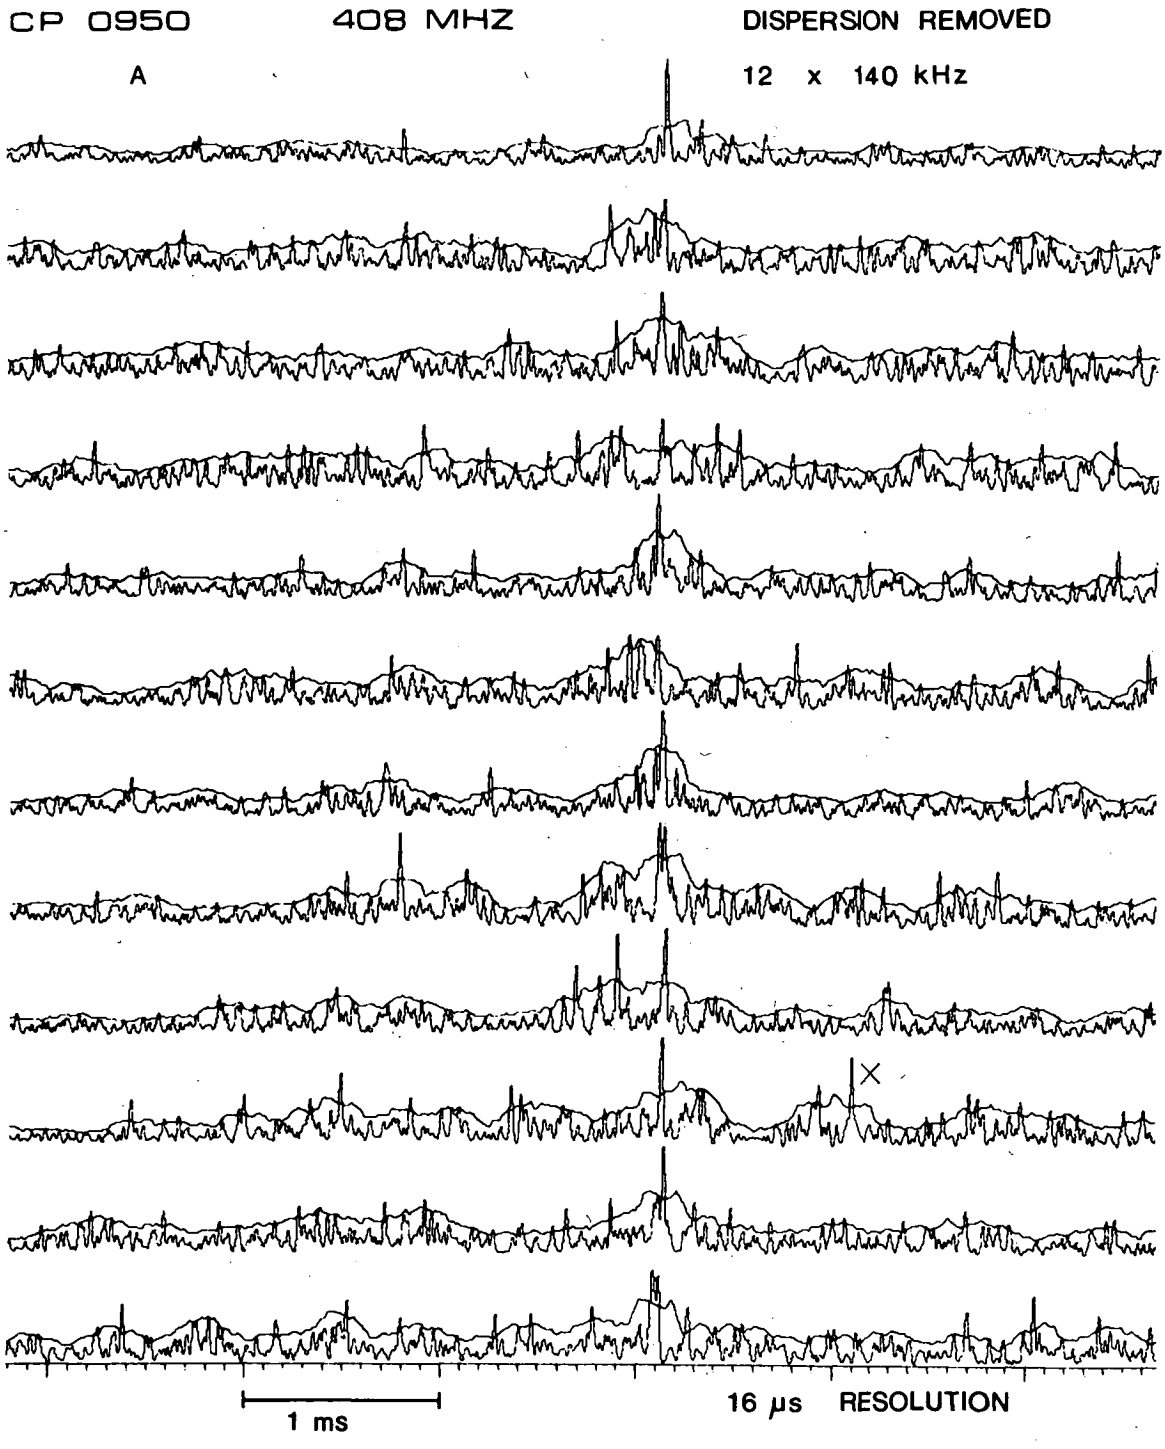
CP 0950
408 MHZ
DISPERSION REMOVED
A
12 x 140 kHz
1 ms
16 μs RESOLUTION

FIGURE 7.1E

because when there are few points in the running-mean the point being tested makes a significant contribution to the mean and depresses the tail of the distribution - see Chapter 8).

The most interesting thing about Figure 7.1E is the alignment of the main peak from one channel to the next. It is not linear and so is not consistent with an ideal impulse or with a dispersed impulse which might mean the presence of an error in the dispersion removal. The alignment is apparently in the form of an S-curve with a peak-to-peak variation of about 50 μs. The nature of this structure is revealed by the display of the dynamic spectrum.

Figure 7.1F shows the dynamic spectrum of the central part of pulse A in a window 600 μs × 1.96 MHz with a resolution of 7.2 μs × 140 kHz. The simulated multi-channel receiver had 56 channels of 140 kHz bandwidth with centre frequencies displaced by 35 kHz. (Thus the set of channels involved in Figure 7.1E is a subset of the 56, however for the dynamic spectrum the intensity traces are unsmoothed.) It can be seen that an S-curve is not the only suitable description of the feature. Others include three bands of emission rising at 45° (i.e. ~10 MHz/ms), and the AMN model interpretation. This is that there is a vertical structure (i.e. constant time) about 70 μs in duration and this modulates a noise process. The apparent S-curve is due to chance. This hypothesis is tenable - although no off-pulse data is shown in Figure 7.1F inspection of off-pulse data shows that there is no qualitative difference. The dynamic spectra always show the same "random speckle" that appears in most of the Figure. Furthermore the human eye is very good at seeing structure in random data, e.g. the "canals" of Mars. Other dynamic spectra are shown in the next Section.

Assuming that the structure in the dynamic spectrum of pulse A is due to the modulation of noise then this structure confirms that the dispersion removal is approximately correct. It appears that the edges of the

FIGURE 7.1F The dynamic spectrum of the central part of pulse A of PSR 0950+08 at 408 MHz in an observing window 630 μs × 2.0 MHz. Time increases horizontally and observing frequency vertically (from 406.9 to 408.8 MHz). The intensity at each point is indicated by the blackness of the point - the same data is shown with two different contrast scales. The intensity is derived from channels with 140 kHz bandwidth and is unsmoothed, so that the time resolution is 7 μs. The channels have a 75% overlap in their passbands. Therefore every second point in time and every fourth in frequency is independent.

FIGURE 7.1G As 7.1F except that the channel bandwidth is 70 kHz so that every fourth point in time and every second in frequency is independent.

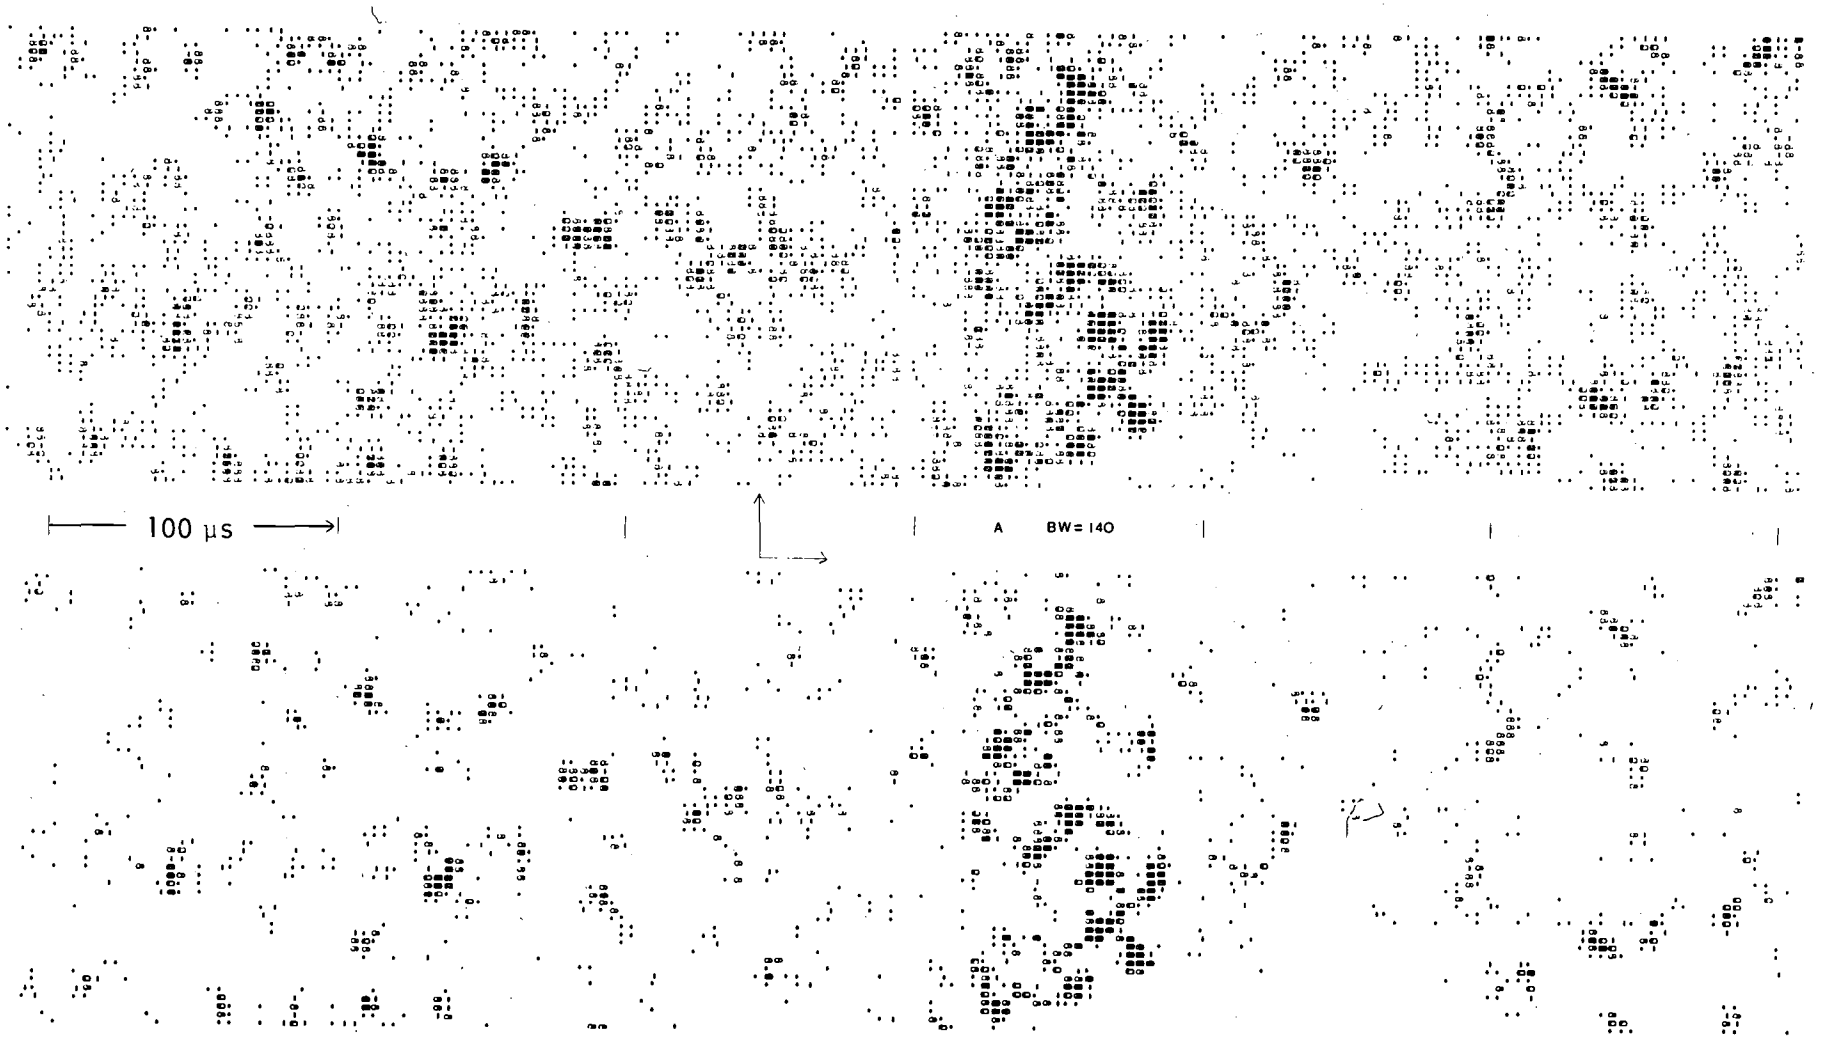

FIGURE 7.1F

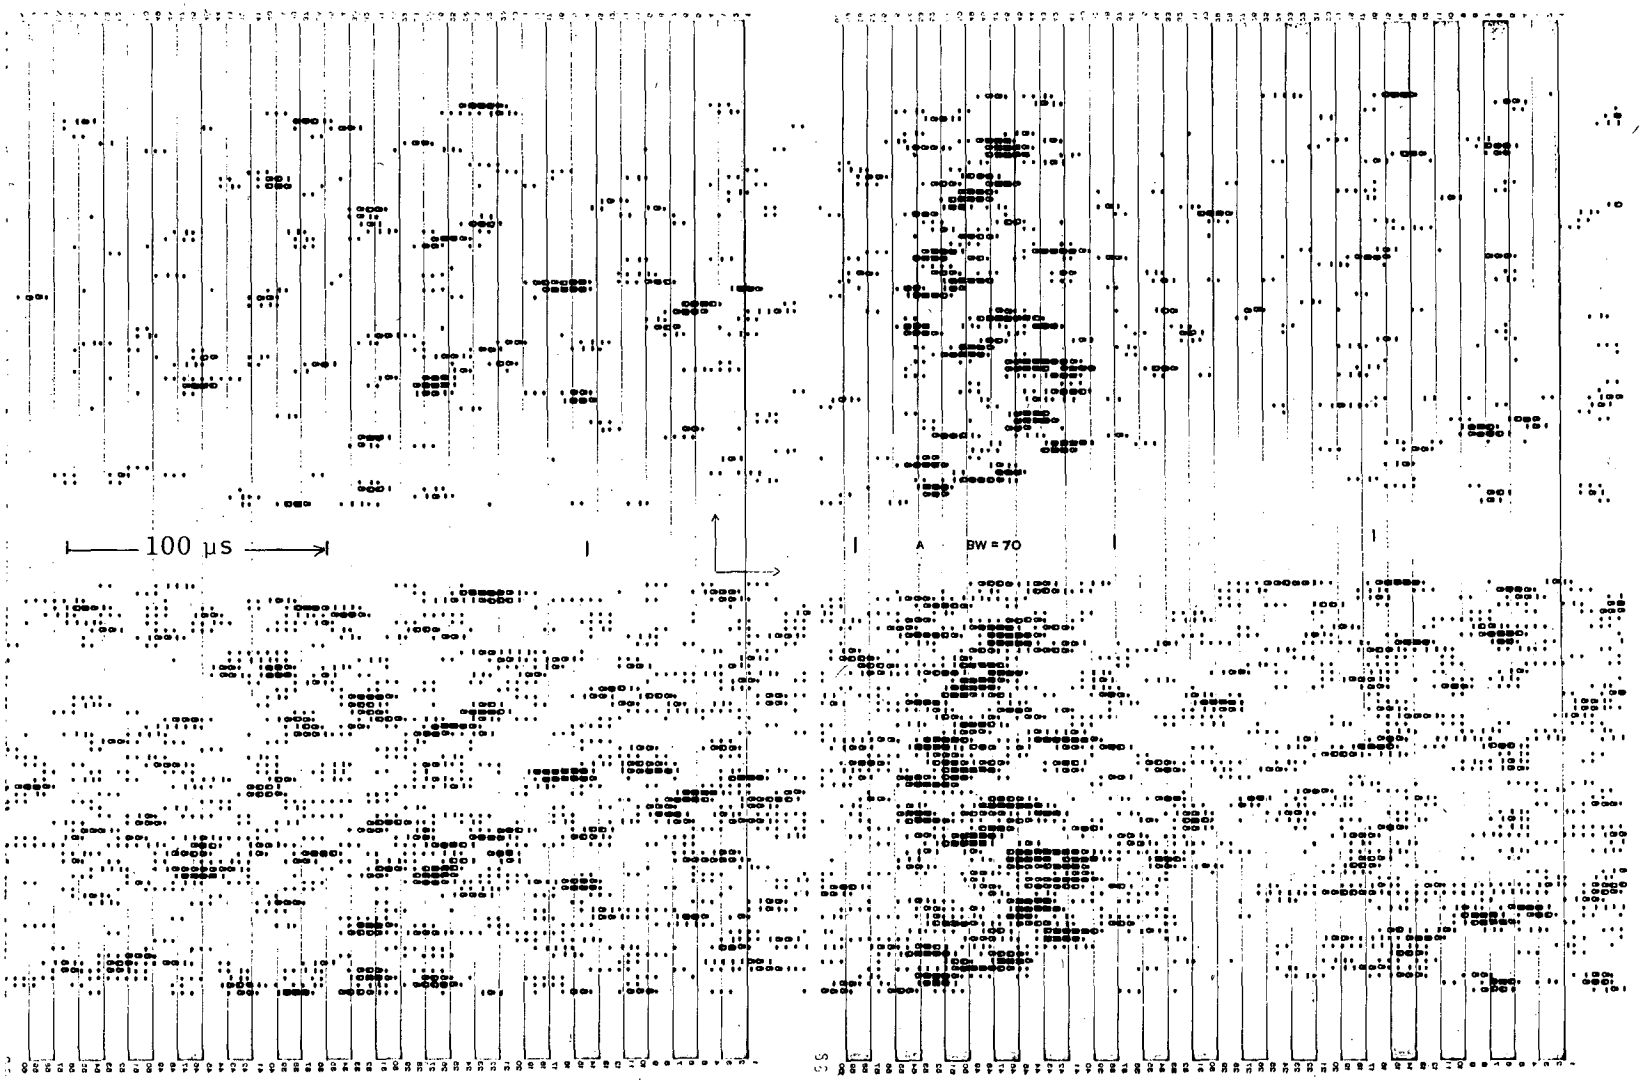

FIGURE 7.1G

enhanced area in the frequency-time plane are vertical (i.e. constant time), which is consistent with but not implied by the observed 2.1 MHz profile. That is, an upper limit can be placed on the residual bandwidth sweep of about 25 μs in 2 MHz which corresponds to $|\Delta DM| \lesssim 0.1$ cm^{-3}pc.

7.2 OBSERVATIONS OF PSR 0950+08

Figure 7.2A shows the intensity profile of seven of the eight pulses from PSR 0950+08 that were analyzed (pulse P is not shown). Each profile is for a total bandwidth of 1.8 MHz (in fact it is the average intensity of thirteen 140 kHz channels) and has a time resolution of 32 μs. Dispersion has been removed and unless otherwise stated this is true of all the plots. Since none of these pulses have strong emission outside a time interval of 5 ms and the integrated profile of PSR 0950+08 at 400 MHz has a width of ~ 20 ms at 10% peak, the pulses are identified as sub-pulses. However the pulsar longitude of these sub-pulses is not known as no absolute timing information was recorded.

In most cases weak emission can be seen preceding the sub-pulse and at least in the case of pulse W there is some following. The peak flux density ranges from (220 ± 30) Jy for pulse K to (1150 ± 200)Jy for pulse W while the pulsar signal to off-pulse noise ratio averaged over 3 ms ranges from 0.7 (pulse P) to 2.5(W).

Three timescales can be discerned in these profiles. First, there is the overall duration of the sub-pulse which ranges from 2 to 5 ms. Second, there is structure on a scale ~ 300 μs, for example, the last two components of pulse H and the three components of T. Third, there are spikes about 100 μs wide, for example, the peak of A, the spikes in the centre of W and the two peaks of T. The latter are seen clearly to be separate when the profiles are not smoothed, only averaged across channels (Figure 7.2B). It is not certain that there is any real division between the latter two timescales for this group of pulses. Certainly when a large number of

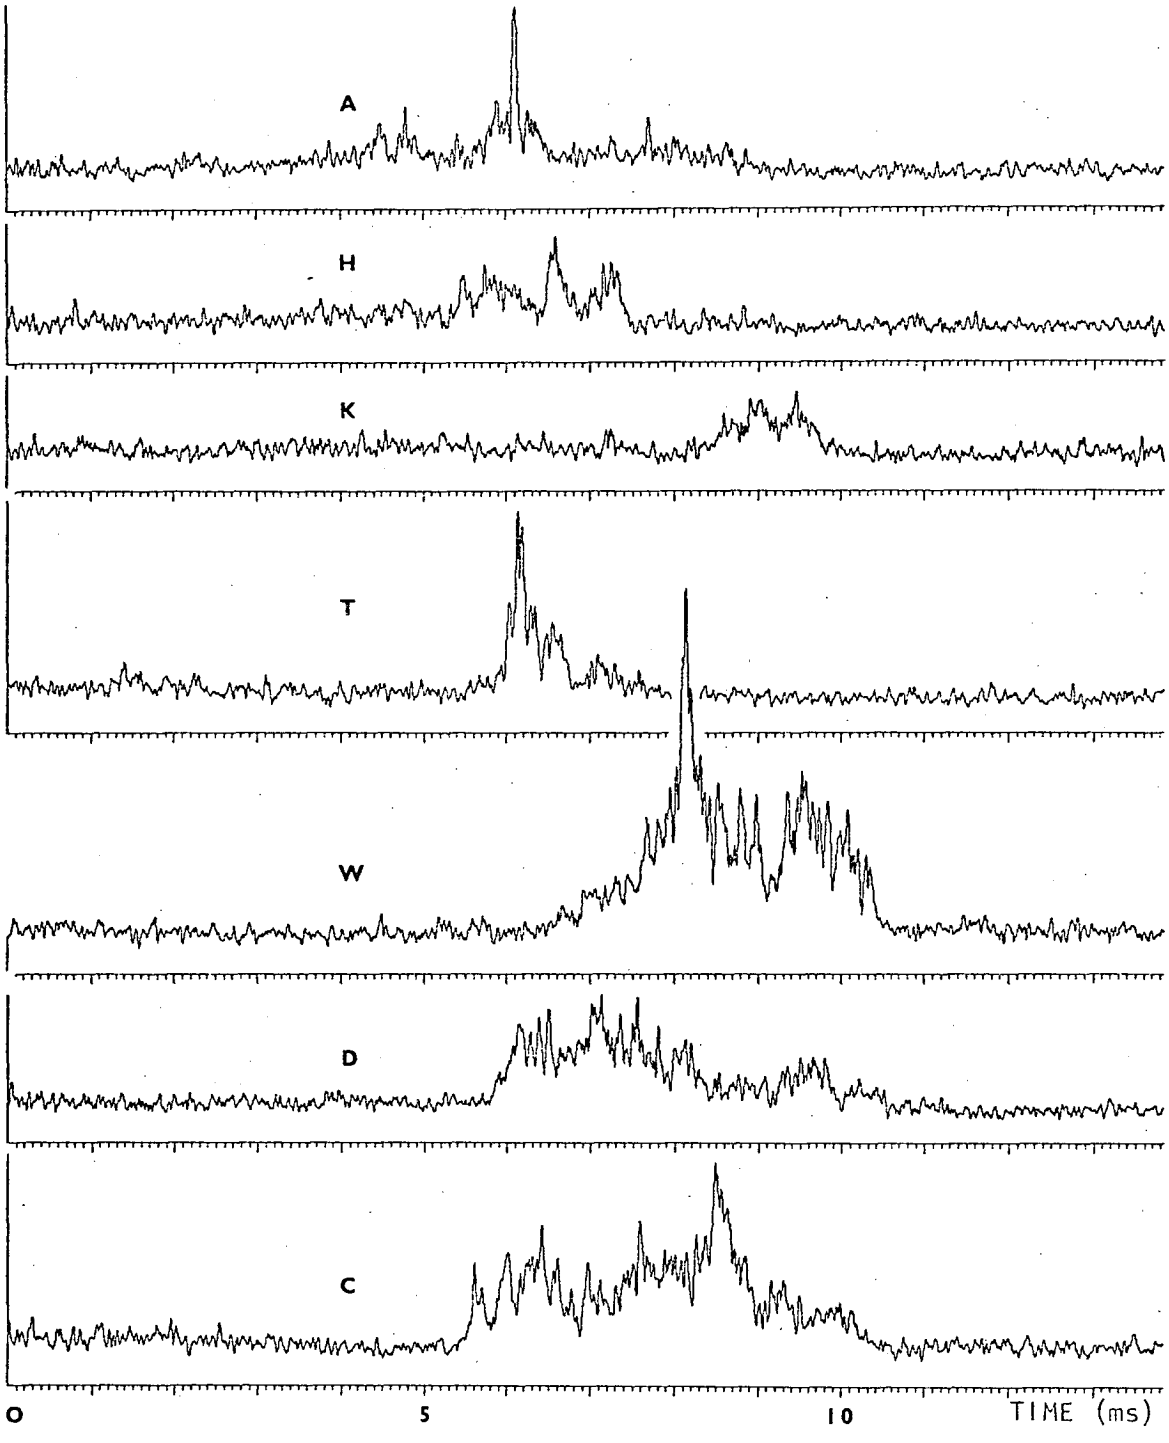

FIGURE 7.2A Intensity profiles of 7 pulses from PSR 0950+08 at 408 MHz. Each is the average over thirteen contiguous 140 kHz channels with smoothing to 32 μs (116 degrees of freedom, 1σ uncertainty 13%). The profiles are aligned arbitrarily. They are scaled so that the apparent off-pulse level is the same ±10%. In each case the zero of intensity is indicated by the time axis which covers 14 ms. Dispersion has been removed.

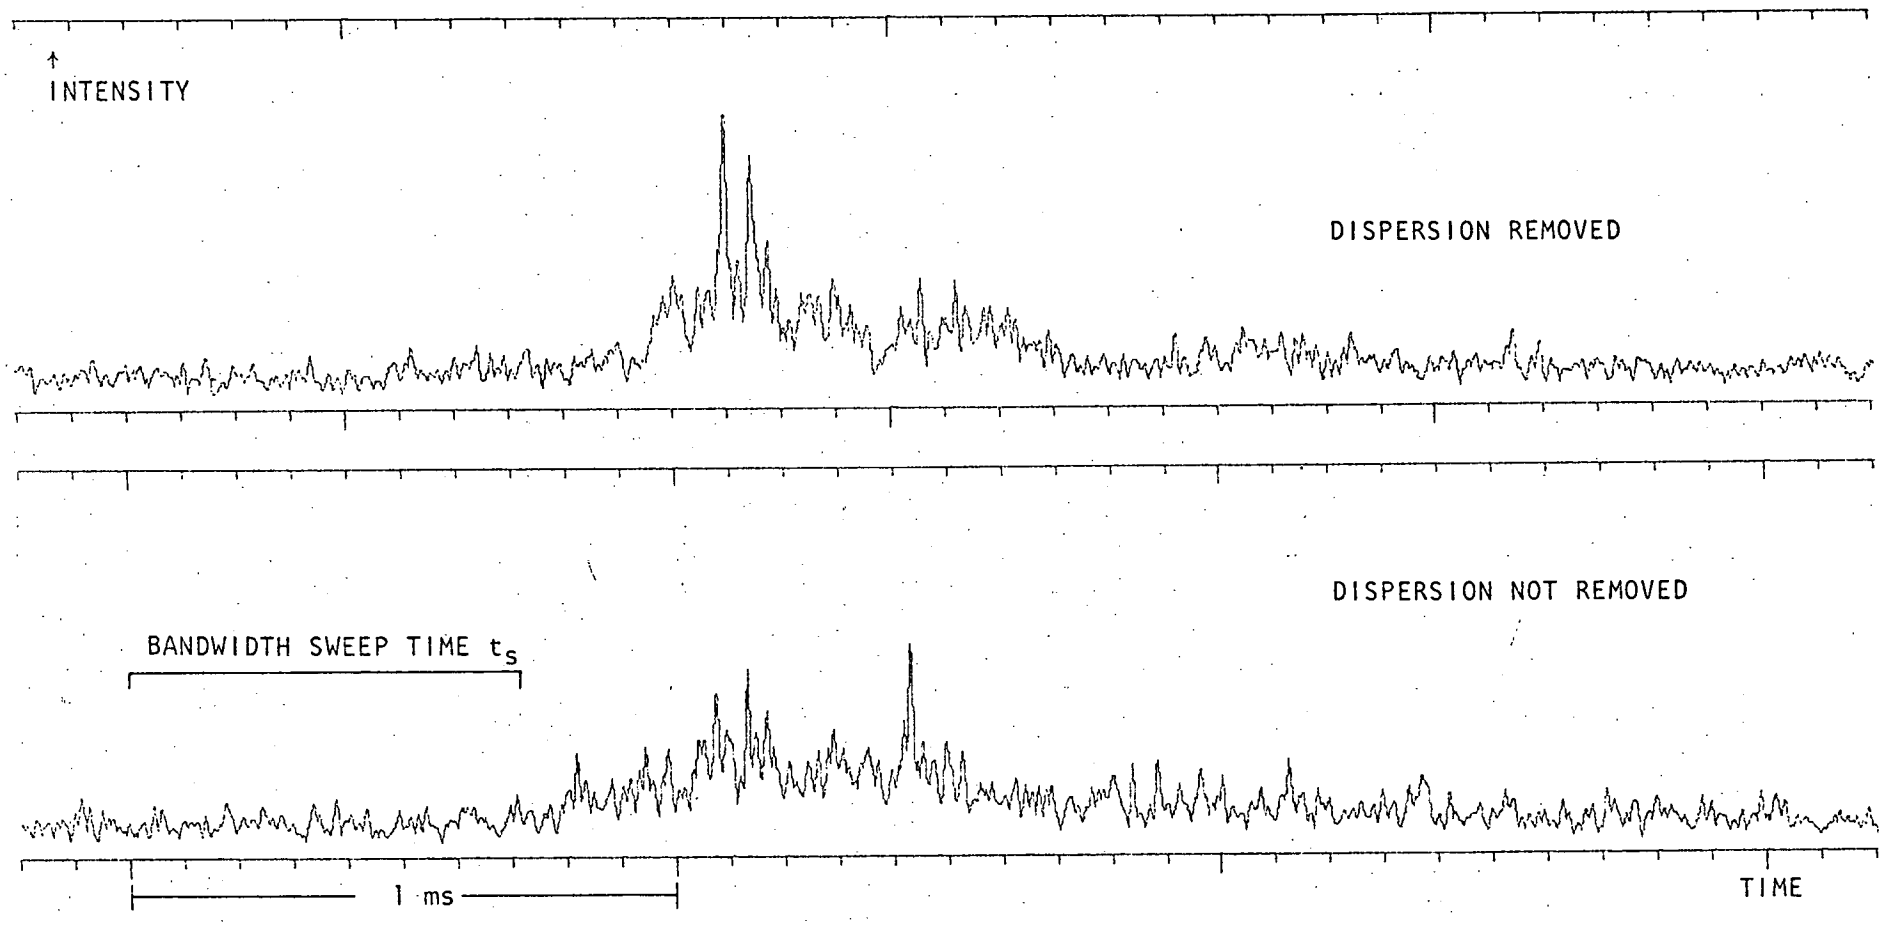

FIGURE 7.2B The intensity profile of pulse T of PSR 0950+08 with and without dispersion removal. Each trace is the average of 56 overlapping channels of 140 kHz bandwidth with a total bandwidth of 2 MHz centred at 408 MHz. The time resolution is 7 μs, i.e. unsmoothed, and each point has 28 degrees of freedom. The first 0.6 ms of the upper-trace is off-pulse background.

pulses are considered there is a continuous distribution of timescales between ≤ 10 μs and 175 μs (microstructure) and between 175 μs and ~10 ms (sub-pulses). This was shown by the auto-correlation functions obtained by Hankins (1972). However, it can be said of this group of eight pulses that in no case do isolated micropulses (< 300 μs) occur. That is, there is always a continuum component in the sub-pulse.

The stronger pulses, that is, A, T, W, C and D, have all been analyzed at full time resolution with bandwidths ~ 2 MHz. In all cases the observed distribution of relative intensity has been satisfactorily fitted by the predicted exponential curve using a running-mean window width of 50 or 53 μs. However in the case of pulse T a better fit to an exponential would be obtained with a width < 50 μs (Chapter 8).

Typical unsmoothed intensity traces are shown in Figures 7.2C and 7.2D. These show the strongest part of pulse C where the signal to off-pulse noise ratio exceeds 3.0. Close inspection shows that the details of the intensity spikes change considerably as the bandwidth changes from 1.1 MHz to 1.8 MHz, although in this case the strongest spike retains its position. Furthermore the appearance is clearly like noise.

Figure 7.2E shows pulse C as observed with thirteen 140 kHz channels with a time resolution of 32 μs. A number of periodic trains of four or more spikes may be seen but the periods are different from channel to channel. It must be concluded that the apparent periodicity is due to chance. Indeed an auto-correlation analysis of these thirteen channels does not show any significant periodicity (Chapter 9). The vertical scale of the thirteen traces in the Figure has been set to ±10% by the total intensity. The signal-to-noise ratio can be judged by the apparently off-pulse noise in the two milliseconds preceding the pulse. It can be seen that the pulsar signal increases with decreasing observing frequency, that is, with increasing video frequency. In fact the S/N ratio averaged over

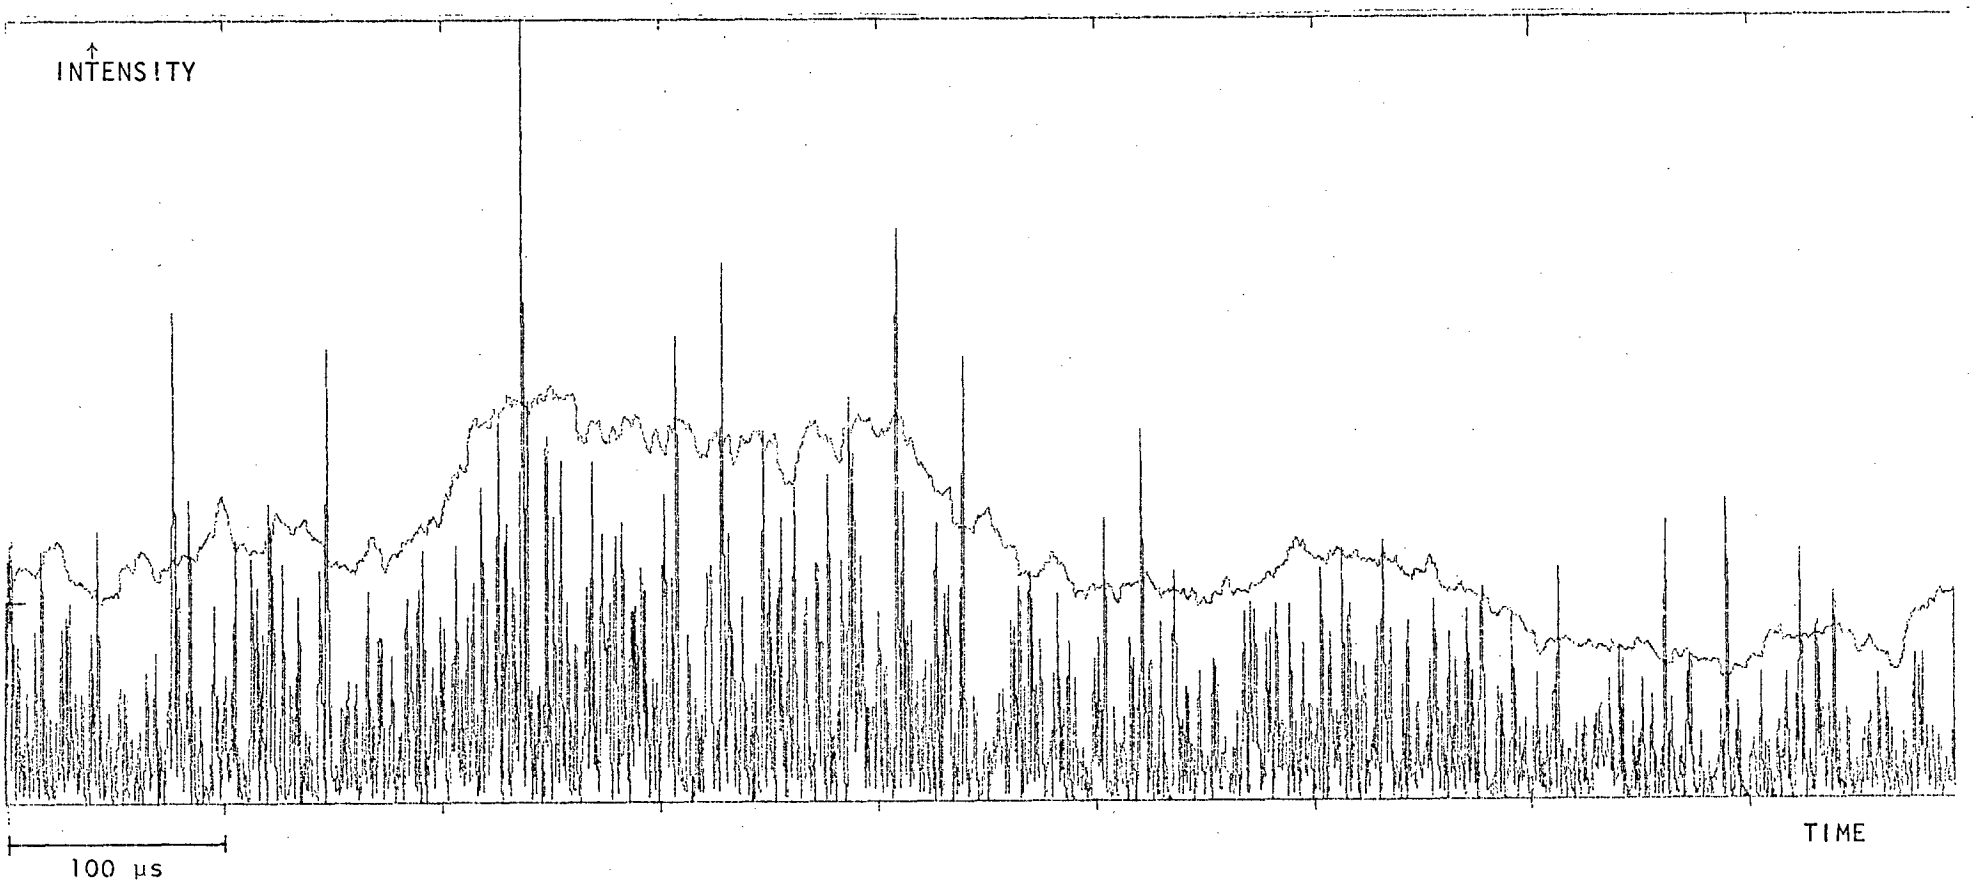

FIGURE 7.2C The strongest part of PSR 0950+08 pulse C as observed with a single channel of bandwidth 1.1 MHz and time resolution 0.9 μs with samples at intervals of 0.45 μs. The smoother curve is 3.7 times the running-mean over 50 μs and forms a 97.5% confidence limit. Figure 7.2D also covers the first half of the data shown here.

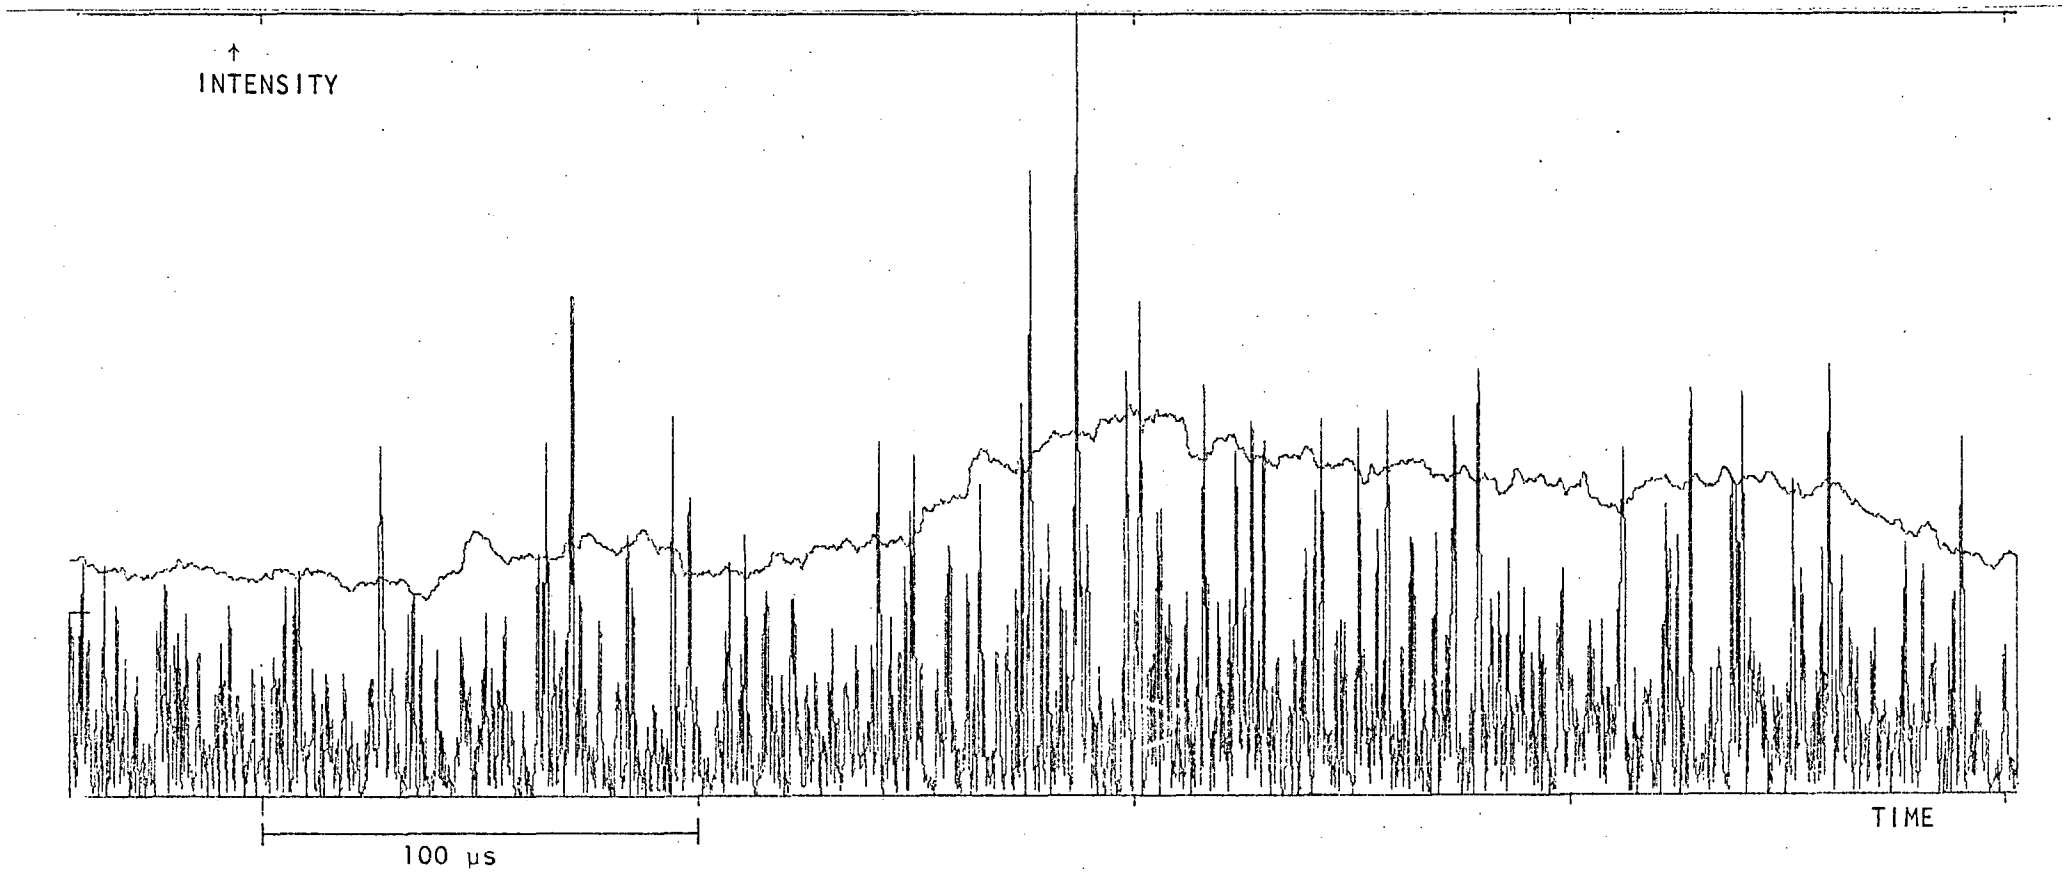

FIGURE 7.2D The strongest part of PSR 0950+08 pulse C as observed with a single channel of bandwidth 1.8 MHz and time resolution 0.45 μs with samples at intervals of 0.22 μs. The smoother curve is 3.7 times the running-mean over 50 μs and forms a 97.5% confidence limit. c.f. Figure 7.2C.

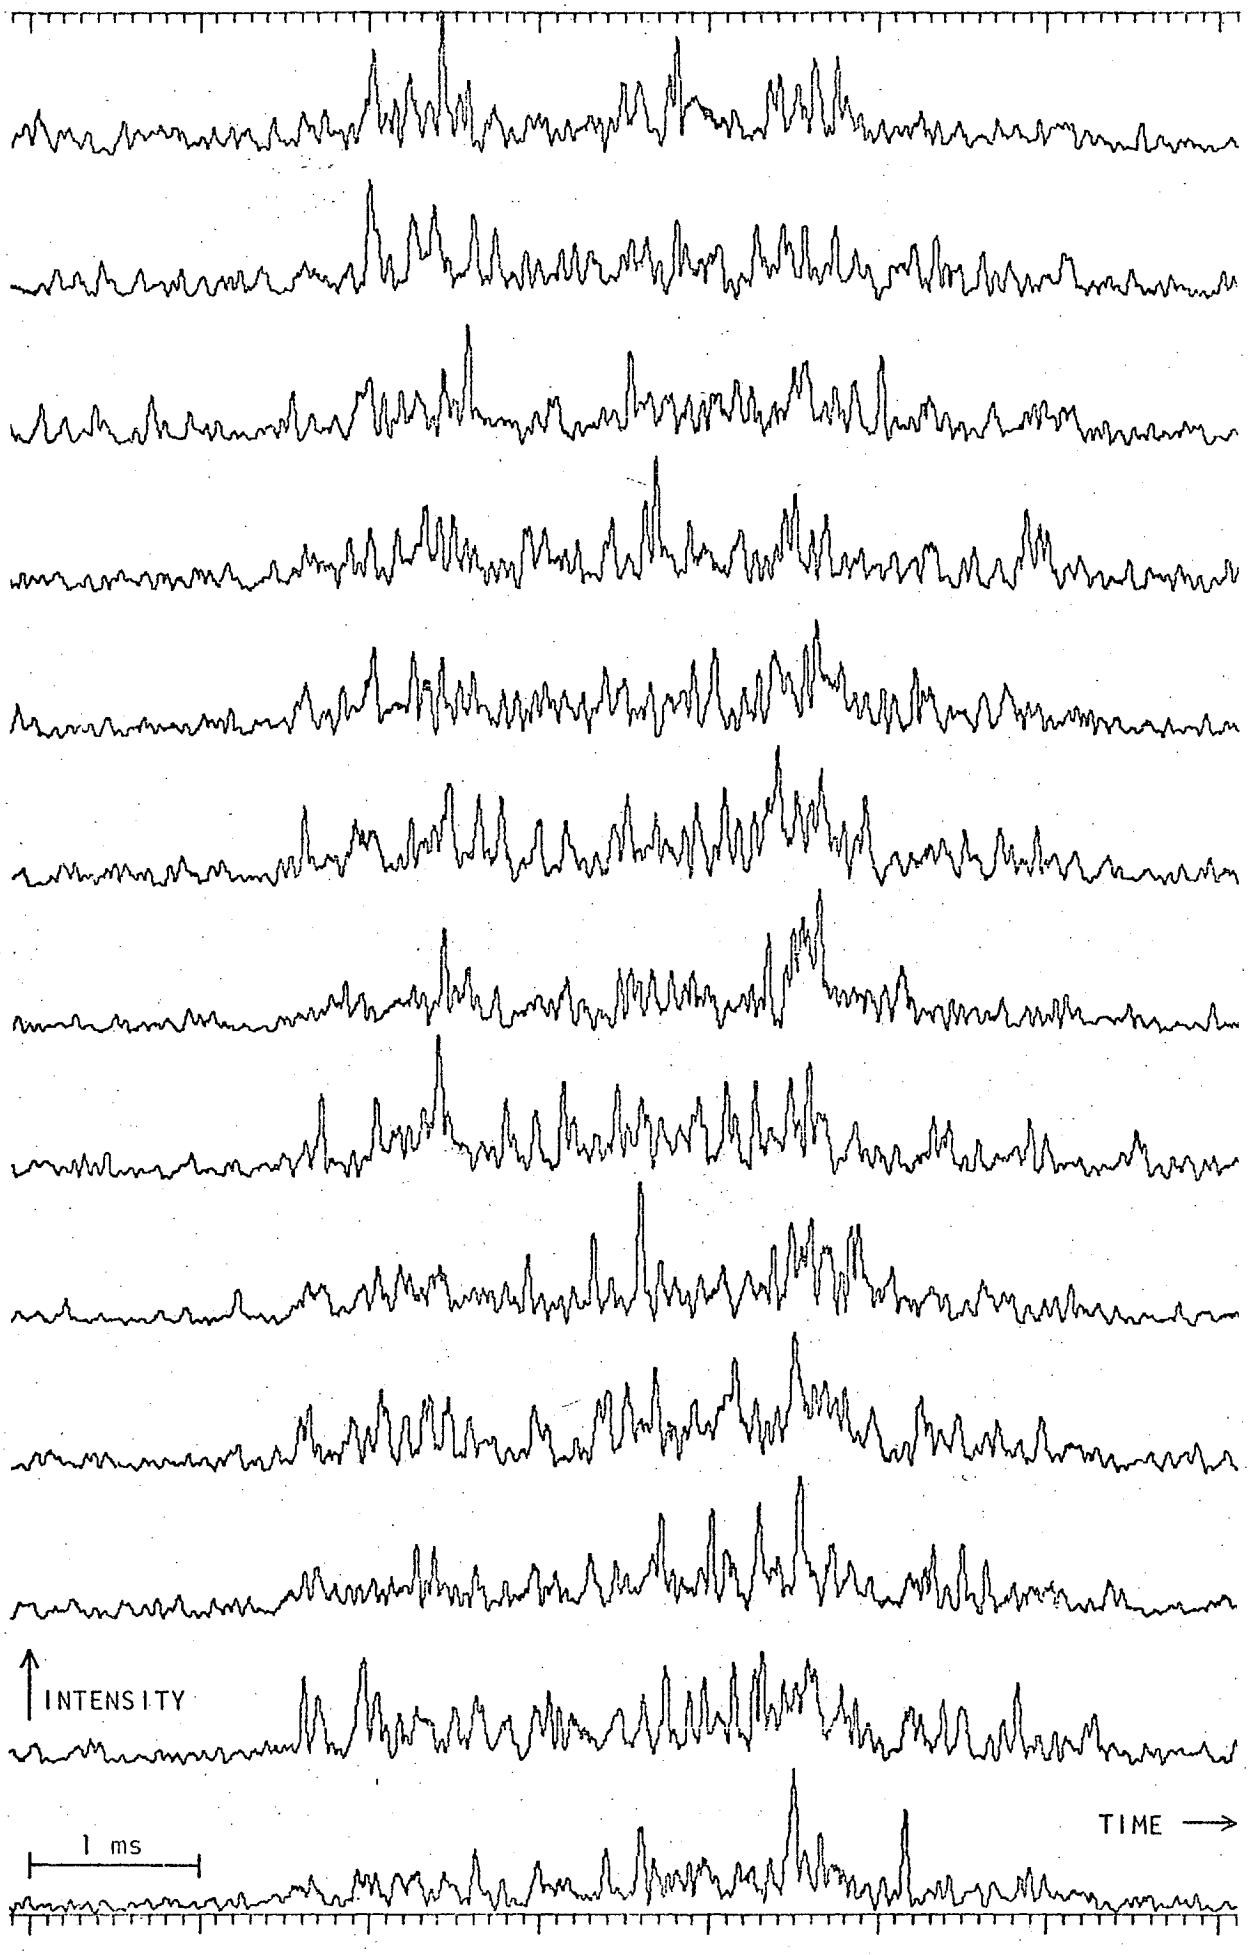

FIGURE 7.2E Pulse C of PSR 0950+08 as observed with thirteen 140 kHz channels smoothed to 32 μs. The total observing bandwidth is 1.8 MHz centred at 408 MHz. Observing frequency increases up the page.

1.4 ms in the strongest part of the pulse increases approximately linearly from 1.2 to 4.5.

A multi-channel observation of pulse W, the strongest pulse observed, is shown in Figure 7.2F. There is an interesting alignment of very strong spikes in two channels (marked in the Figure). Their strength can be seen by the depression of the vertical scale of the off-pulse background, and by comparison with the smoother curves which are 3.7 times the running-mean over 120 μs. (Since these traces are smoothed the expected distribution of relative intensity is not exponential but the smoothing is light, from 7.2 μs to 10 μs, and so the exponential distribution is a rough approximation. On the other hand 120 μs is too narrow for the points being tested not to make significant contributions to the local mean.) However when the distribution of relative intensity was obtained for these traces without smoothing it was found to be consistent with that of off-pulse data. That is, none of the peaks were so high that they could not be explained as noise fluctuations even using running-mean windows > 100 μs.

The dynamic spectrum of the strongest part of pulse W is shown in Figure 7.2G, and the position of the peaks discussed in the last paragraph is indicated. Like that of pulse A this dynamic spectrum is described well by the AMN model. The same is true of the dynamic spectra of all eight pulses from PSR 0950+08 including that of pulse T shown in Figures 7.2H and 8.3B.

In an attempt to test objectively the significance of various diagonal lanes seen very occasionally in the dynamic spectra I calculated slanted average intensities. That is, I effectively re-introduced dispersion (in fact this was the post-detection process) and varied the dispersion measure, i.e. the slant across the frequency-time plane. If a lane was significant then there would be a significant dip in the re-evaluated pulse profile. No significant dips or peaks were found except at the true dispersion measure, that is a vertical average across the dynamic spectrum, but the test did not

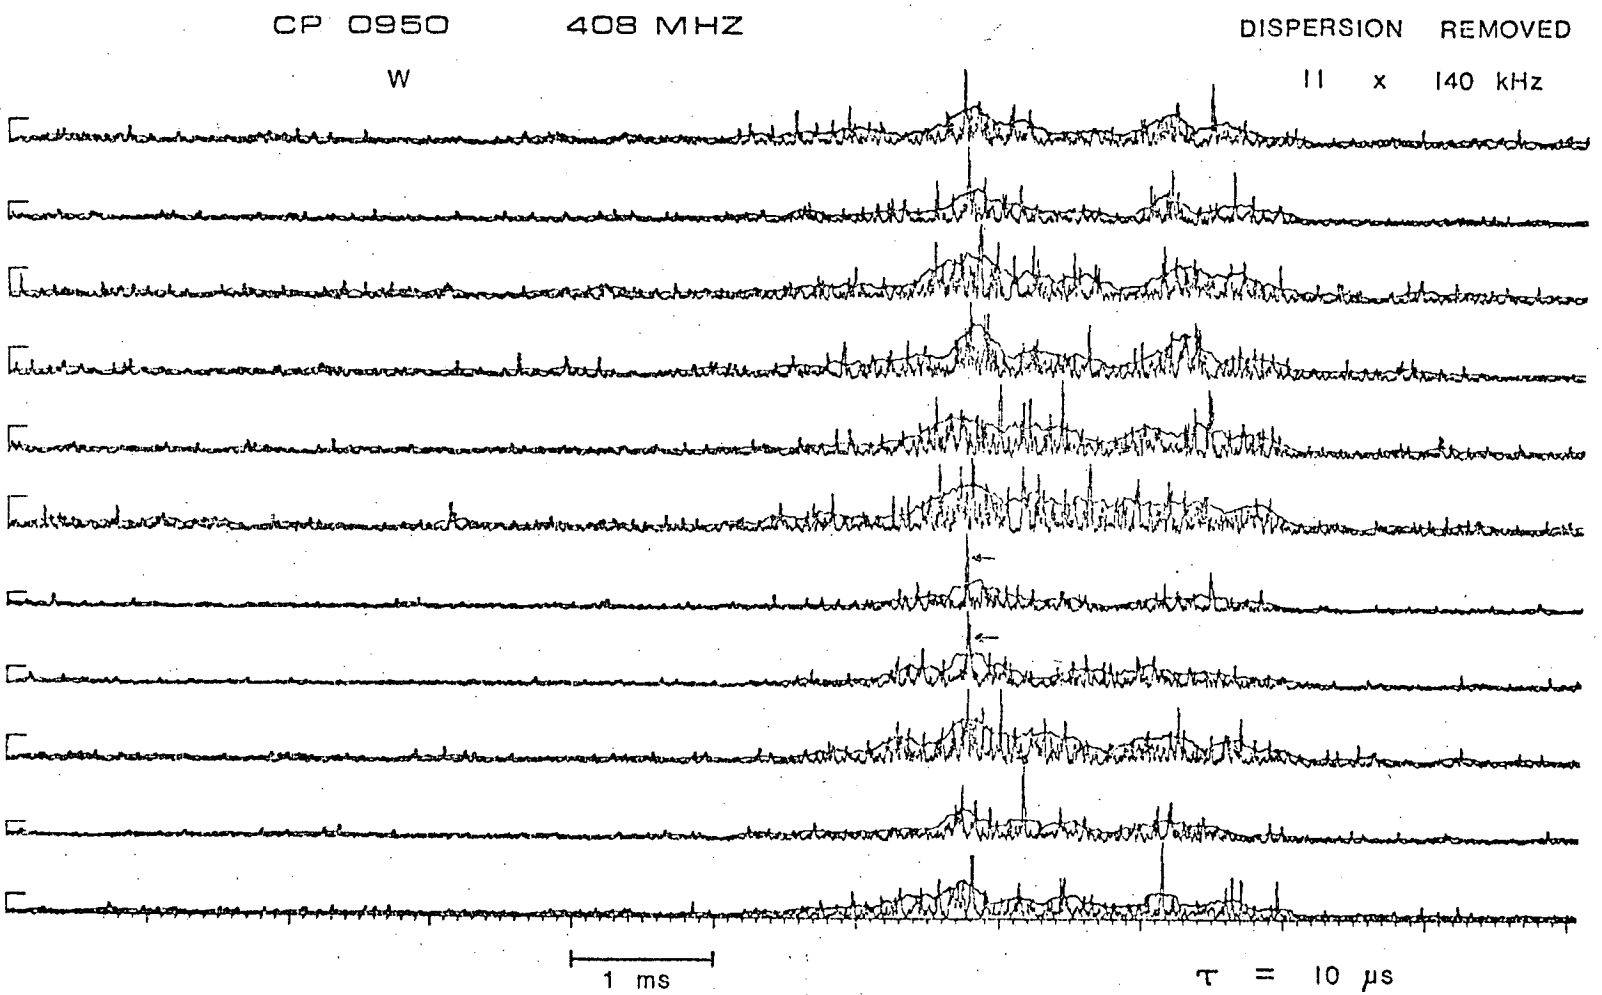

FIGURE 7.2F Pulse W of PSR 0950+08 as observed with eleven 140 kHz channels smoothed to 10 μs. Observing frequency increases up the page. The peaks marked with arrows are discussed in the text. The horizontal bars at the left indicate the total signal in each trace.

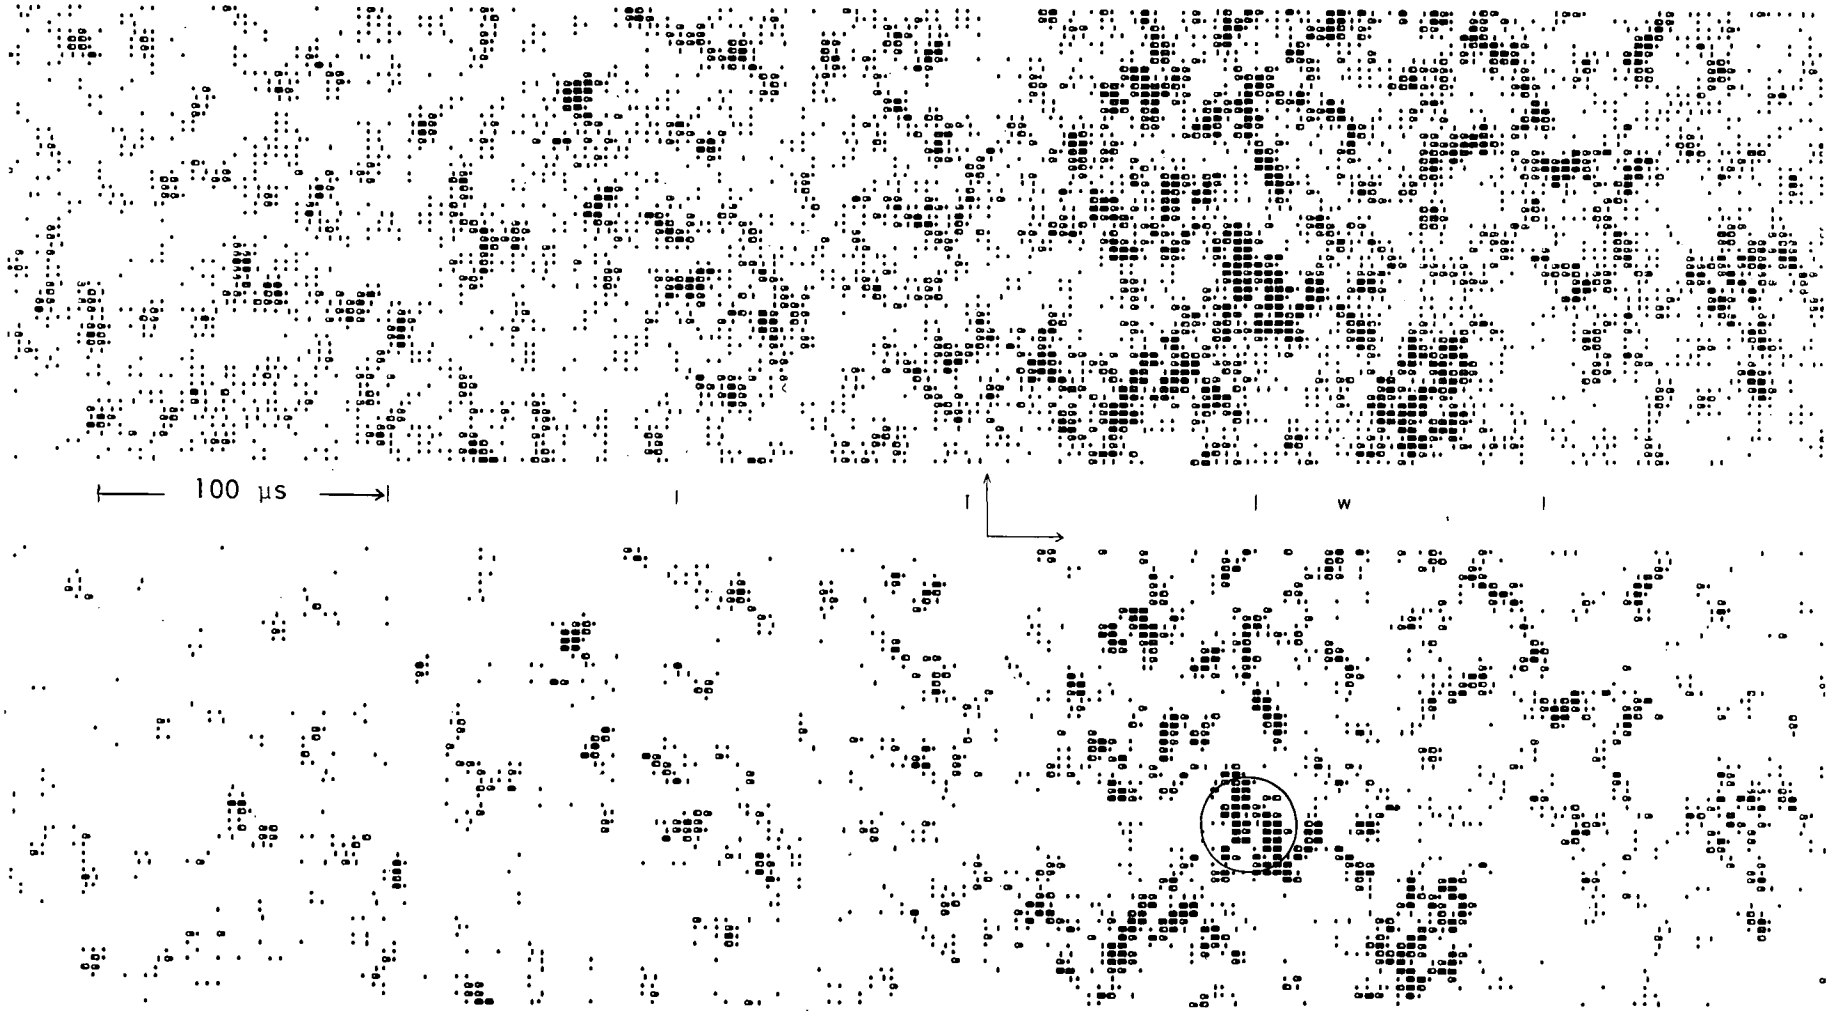

FIGURE 7.2G The dynamic spectrum of PSR 0950+08 pulse W. See caption for Figure 7.1F.

The circled feature is discussed in the text.

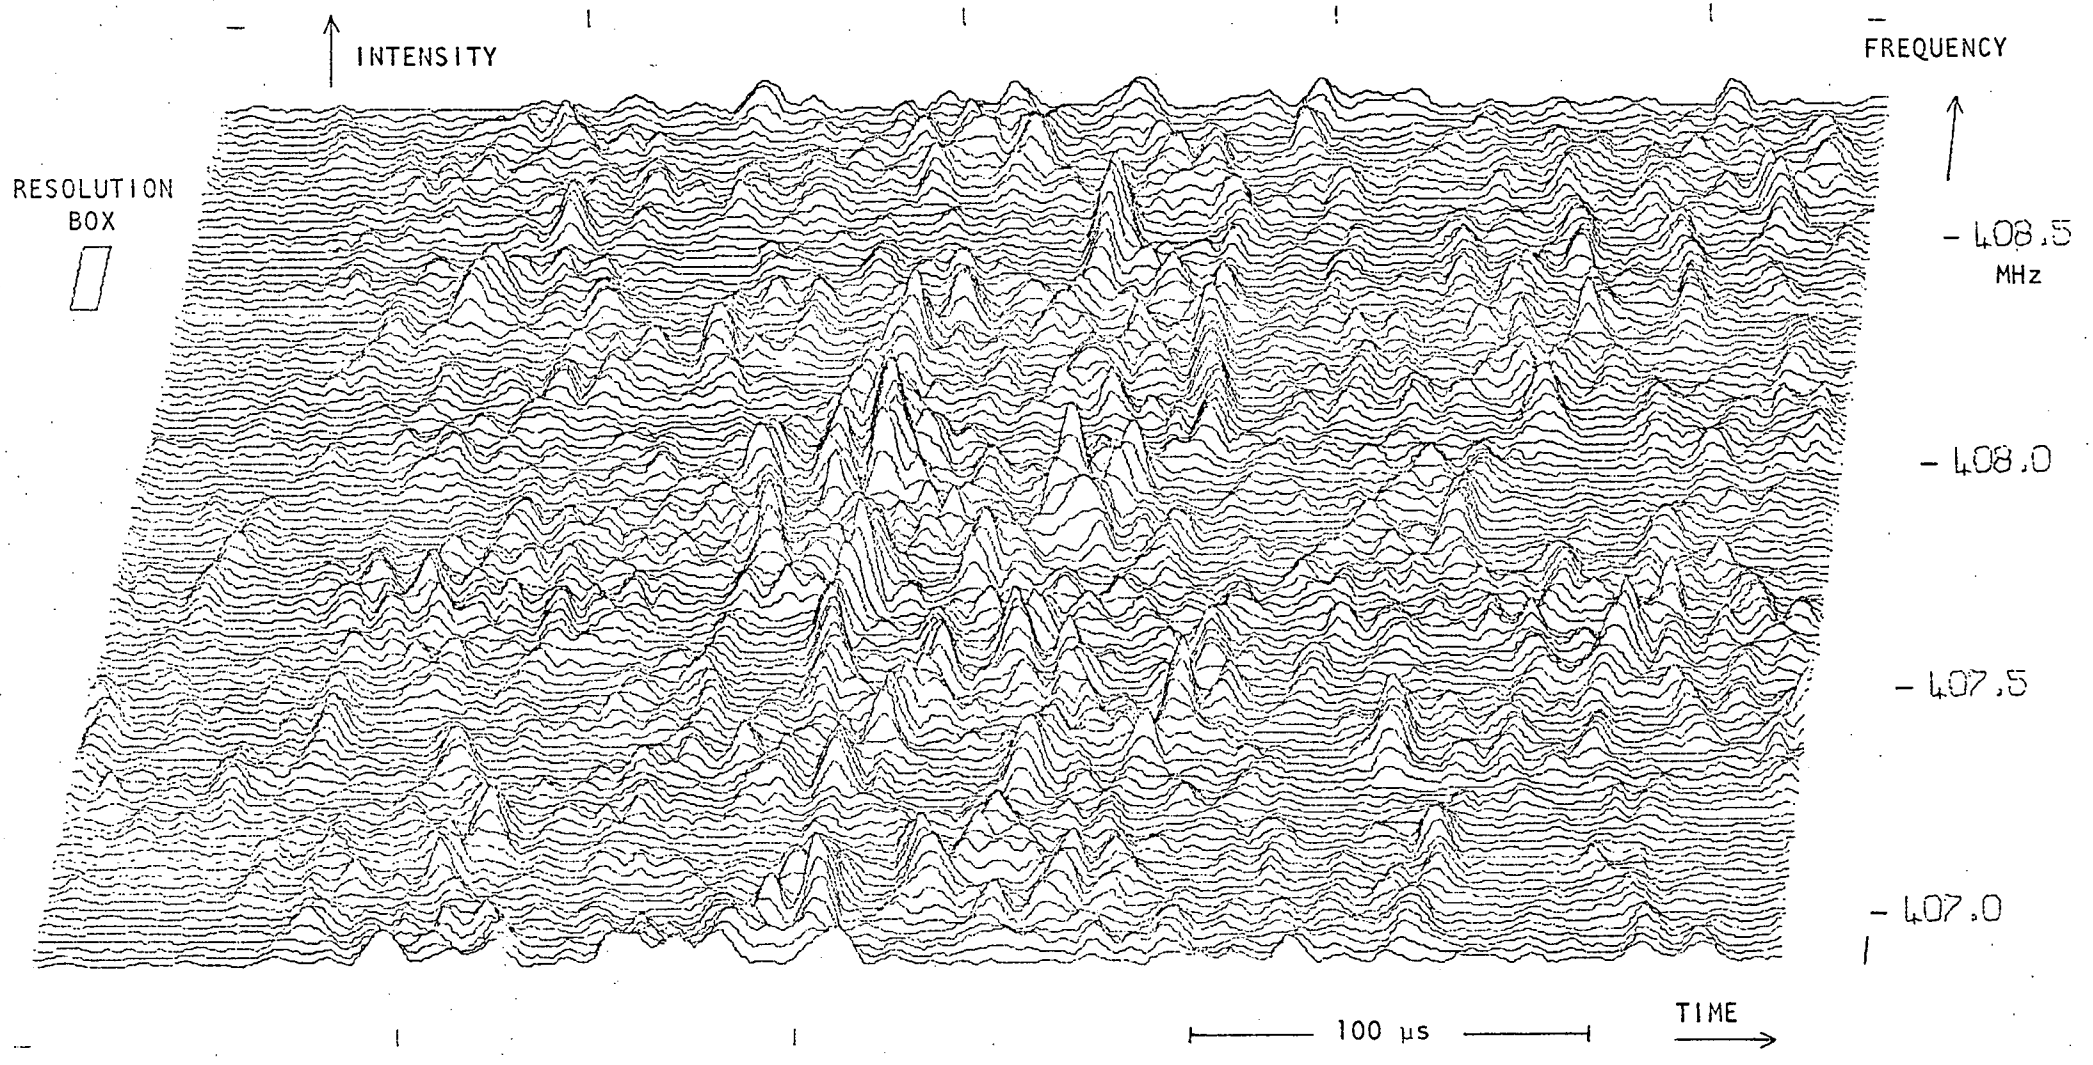

FIGURE 7.2H The dynamic spectrum of PSR 0950+08 pulse T. Each intensity point has 2 degrees of freedom.

have a high "Q" even for this. The conclusion is that the lanes were illusions or chance occurrences with no physical significance.

7.3 OBSERVATIONS OF PSR 1749-28

The intensity profiles of nine pulses from PSR 1749-28 are shown in Figure 7.3A. They were recorded in a single transit of the pulsar (on 1975 June 19); the relative pulse numbers within the transit are shown in the Figure. The widths of the (sub) pulses vary from 3 ms (pulse B) to 9 ms (A and D). Note that narrow pulses are selected preferentially for analysis because they are less likely to cross a stripe boundary. In fact pulses D and E do have emission in the following stripe at video frequencies > 1 MHz, that is, observing frequencies < 408 MHz. Therefore these two pulses were digitized using a 1 MHz low-pass filter so that the profiles shown are not truncated by the end of the time window.

The nine profiles in Figure 7.3A have been aligned according to pulsar longitude (±0.5 ms), but the absolute phase is not known. However there is emission over a time (i.e. longitude) span of about 12 ms. Since PSR 1749-28 has a simple integrated profile of equivalent width 7 ms and a width at 10% peak of about 12.5 ms (Hamilton *et al* 1977), the phase can be estimated to $\pm$ 2 ms. That is, the integrated profile is probably centred on the 14.6 ms time axis of the Figure. Therefore there is a suggestion that the narrower subpulses, that is, with durations ~ 3 ms, are more prevalent at the leading edge of the integrated profile, but the sample is too small for a definite inference.

Each profile is the intensity averaged over 56 over-lapping channels with bandwidths of 70 kHz covering a total bandwidth of 2.0 MHz (1.0 MHz for pulses D and E) with smoothing to 63 μs resolution. Thus each off-pulse point in the Figure has about 240 degrees of freedom (120 for D and E) and hence a 1σ uncertainty of about ±0.09 (or ±0.18). However each

FIGURE 7.3A Intensity profiles of nine pulses from PSR 1749-28 at 408 MHz. The pulse identification and pulse number within the transit of 1975 June 19 are marked, e.g. "A #5" indicates that pulse A was five pulsar periods (P = 0.5626 s) after the first pulse on the chart record. Pulses D and E had a total bandwidth of 1.0 MHz, the others 2.0 MHz. In all cases the profiles have been smoothed to 63 μs resolution and are averages across channels of 70 kHz bandwidth. Because of the use of the discrete Fourier transform in dispersion removal the time domain is periodic. In this case the period is 14.6 ms while the bandwidth sweep time is 12.4 ms for 2 MHz. Therefore the off-pulse background is a mixture of the background before and after the pulse and effectively the system gain is tapered outside the region between the two triangular symbols. Refer to Figure 7.3B.

The profiles have been aligned according to pulsar longitude.

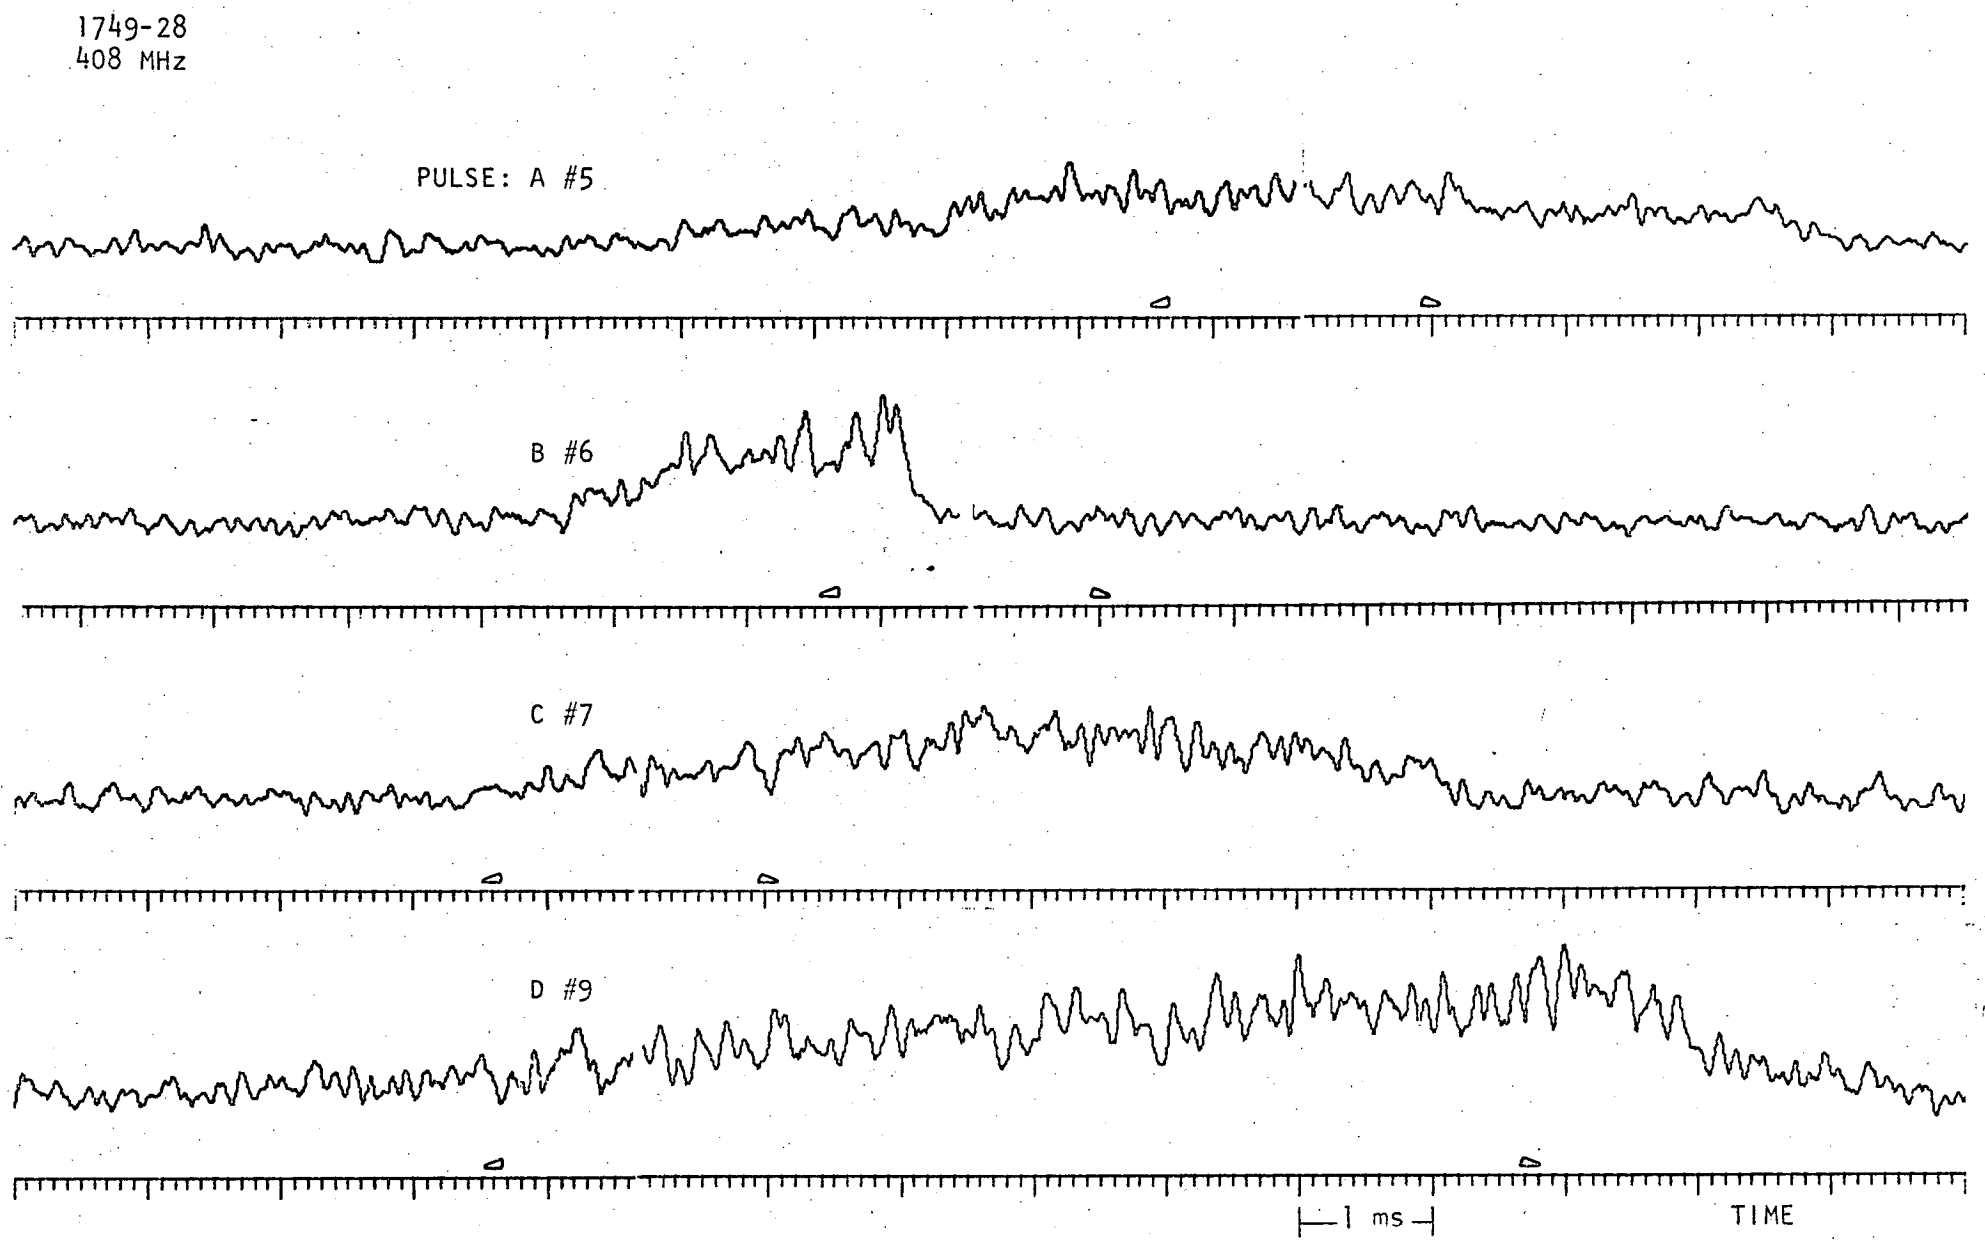

FIGURE 7.3A

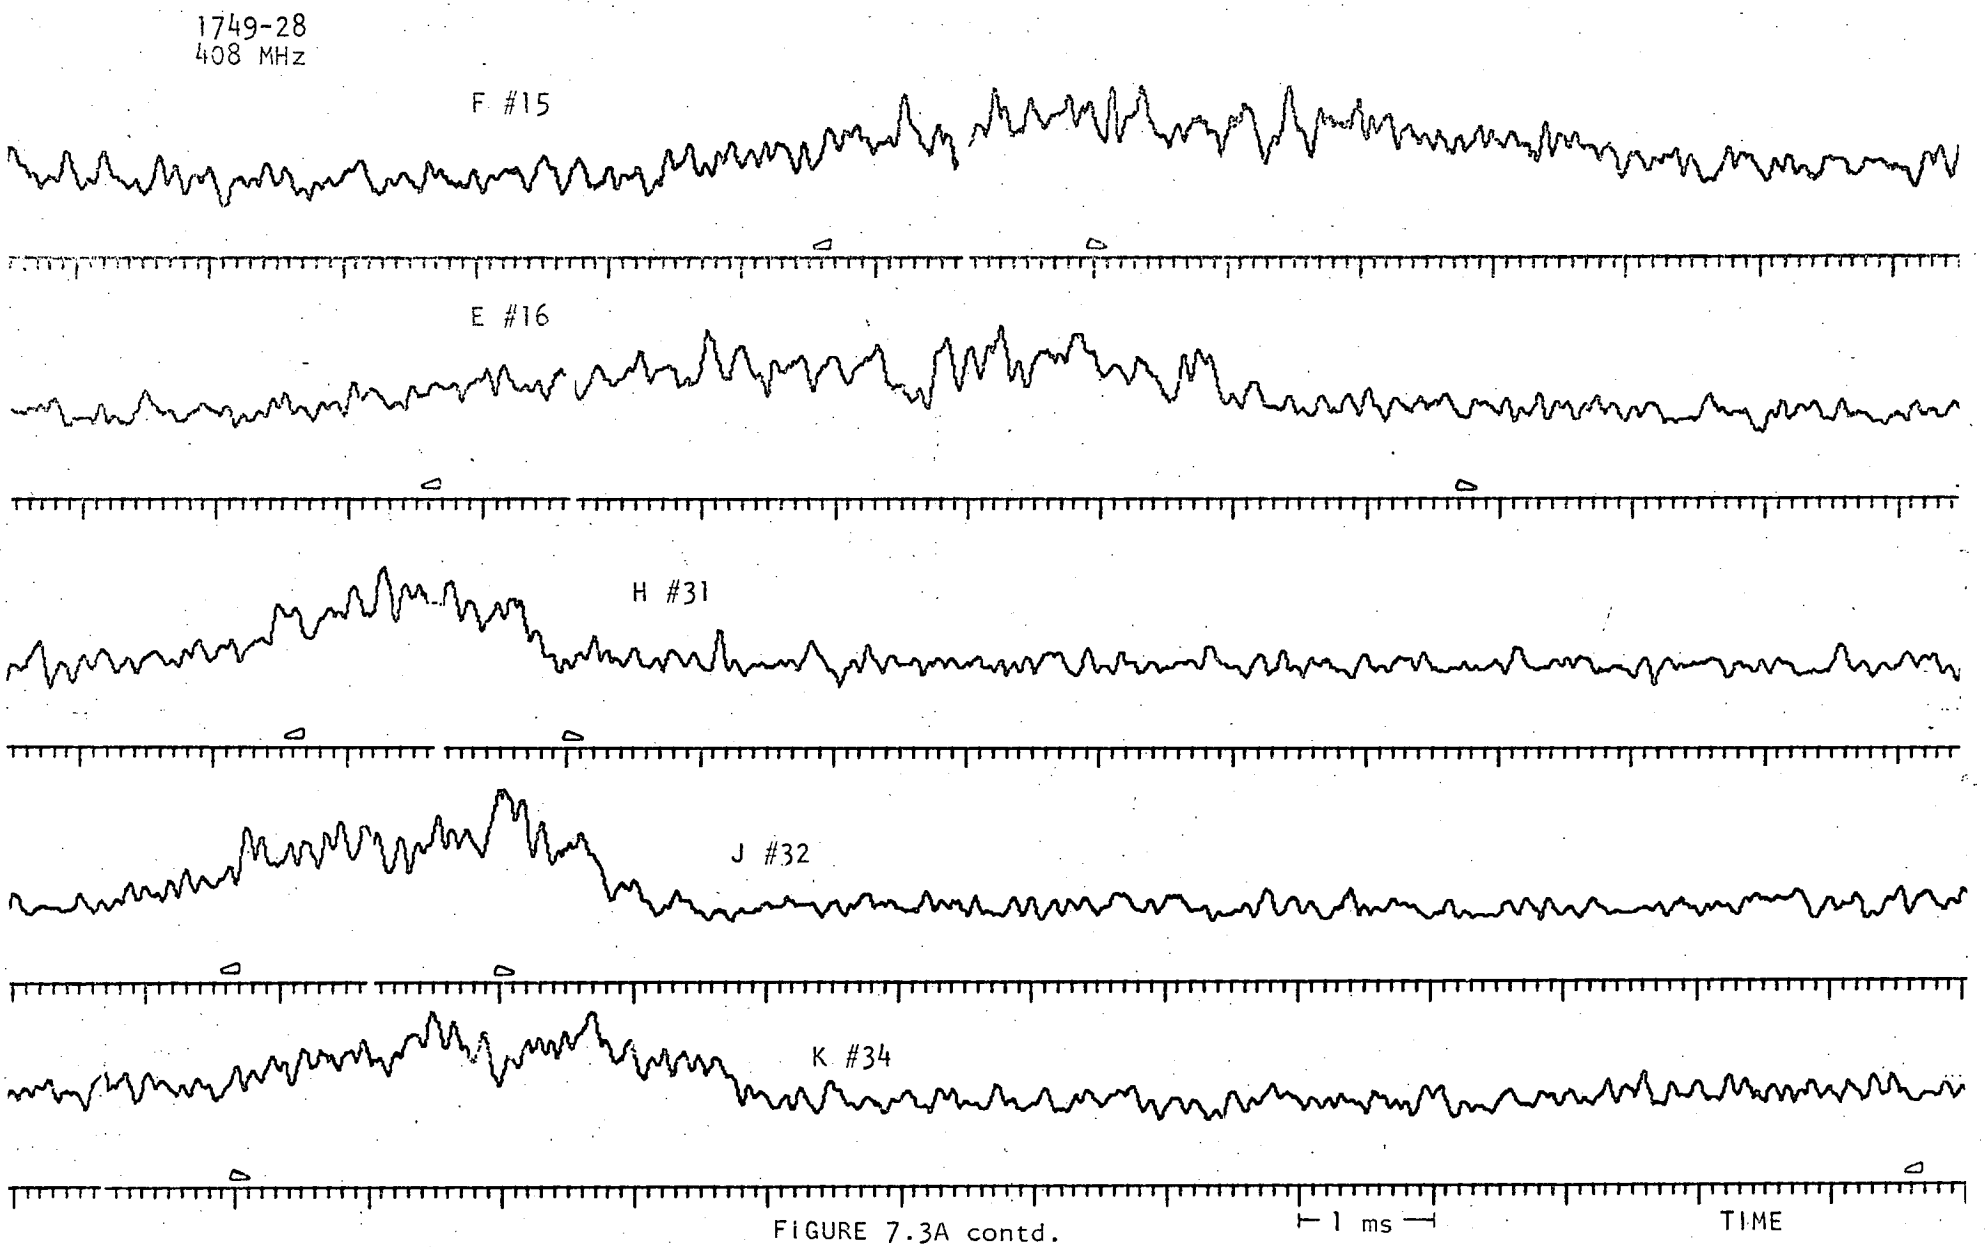

FIGURE 7.3A contd.

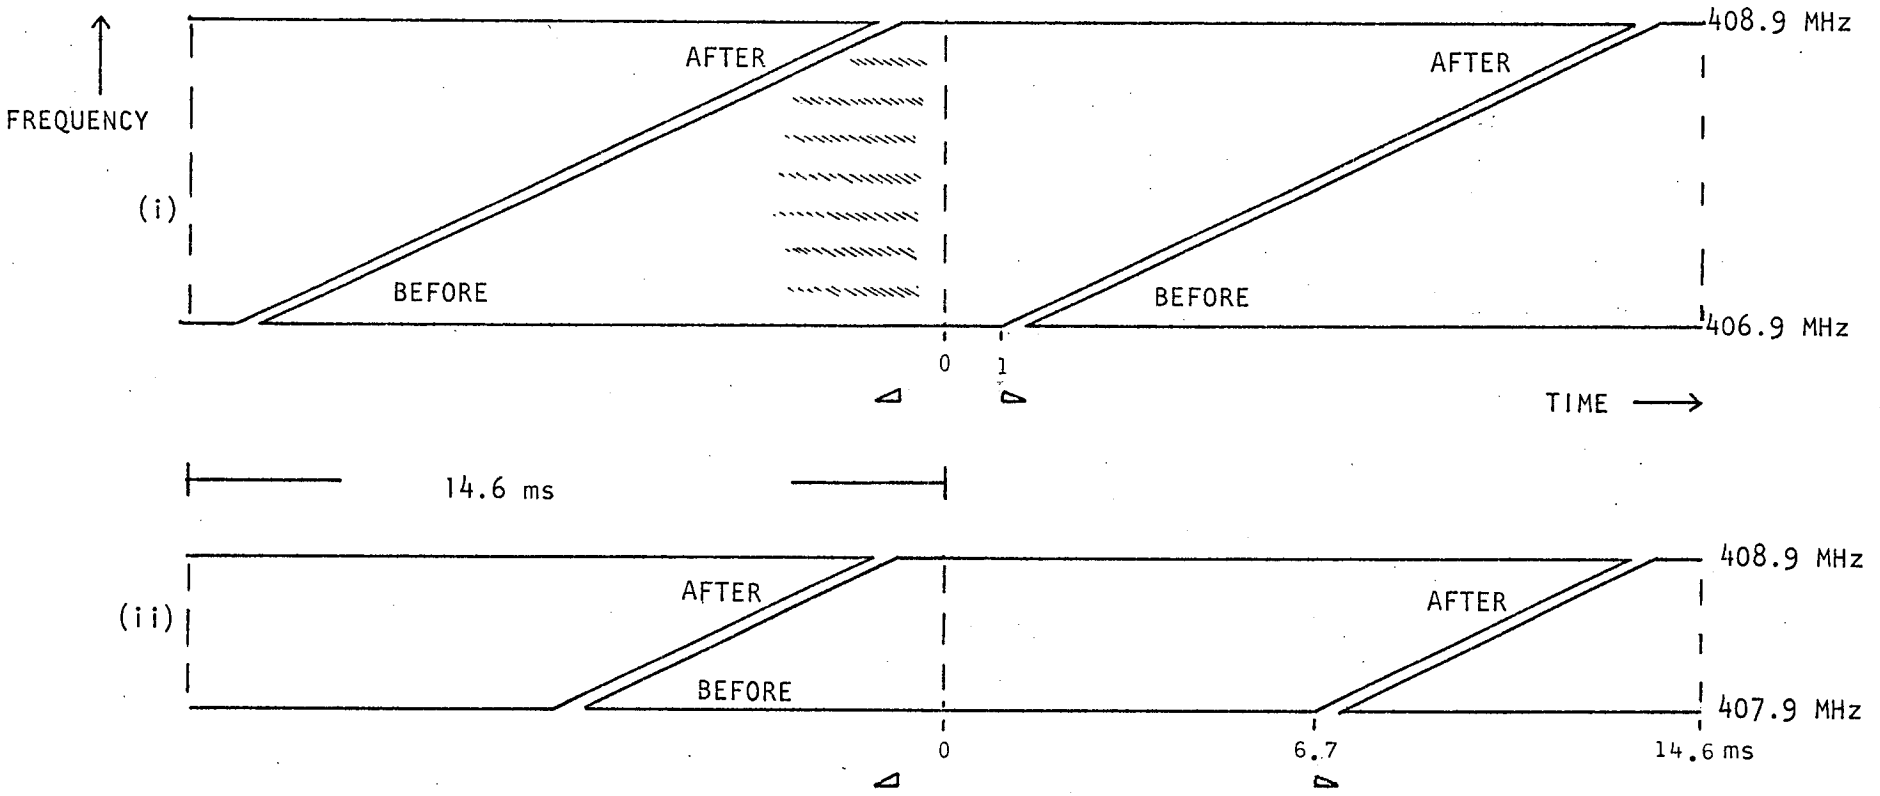

FIGURE 7.3B The observing window onto the dynamic spectrum of PSR 1749-28 pulses (i) A, B, C, F, H, J, K and (ii) D, E. The blanked-signal stripe is about 0.5 ms wide; its length is determined by the bandwidth sweep time. The triangular symbols appear in Figure 7.3A and between them the system gain is constant for summation across channels. The position of pulse B is sketched in (i) (but the fringe spacing is not to scale).

on-pulse point has fewer degrees of freedom because the effective bandwidth is approximately halved by interstellar scintillation. The peak signal-to-(off-pulse) noise ratio varies from 0.9 (pulse A) to 1.8 (pulse D), and the averages over 2 ms are about two-thirds of these ratios. Little pattern can be seen in the profiles although there seems to be a tendency to have steep trailing edges (c.f. the PSR 0950+08 pulses). There is no obvious structure on a sub-millisecond timescale.

The relative intensity distributions for the 56 unsmoothed traces for pulses A, B and C are exponential. This is with a rather wide running-mean window, namely 730 μs, which suggests that, in accordance with the appearance of the profiles, there is no strong microstructure on a timescale < 500 μs.

As mentioned previously all of the pulses have spectra which are deeply modulated by interstellar scattering. The power spectrum for pulse C is shown in Figure 7.3C and also in Figure 7.3D with pulse K after passband flattening. As would be expected from the fact that the nine pulses were observed within 16 seconds, the same fringe pattern applies to all of them. Differences in the height of the fringes can be attributed to variation in pulse intensity and statistical fluctuations. (Note that the error bars shown in Figure 7.3D are under-estimates of the uncertainty because they are based on 14 ms of assumed stationary Gaussian noise while the pulses occupy between 3 and 9 ms. The quasi-period of the pattern is about 150 kHz (over 1.8 MHz) while in the centre of the band there is a periodic group of four narrow fringes with period 165 ± 10 kHz. The scale of the pattern as a whole is consistent with, i.e. within a factor of two of the estimate of the de-correlation frequency given in Section 6.4.2, namely 16 kHz.

The dynamic spectra of the nine pulses show the scintillation fringes but no other structure is apparent. The fringes themselves appear as bands in which the noise speckle is darker and there is no correlation between

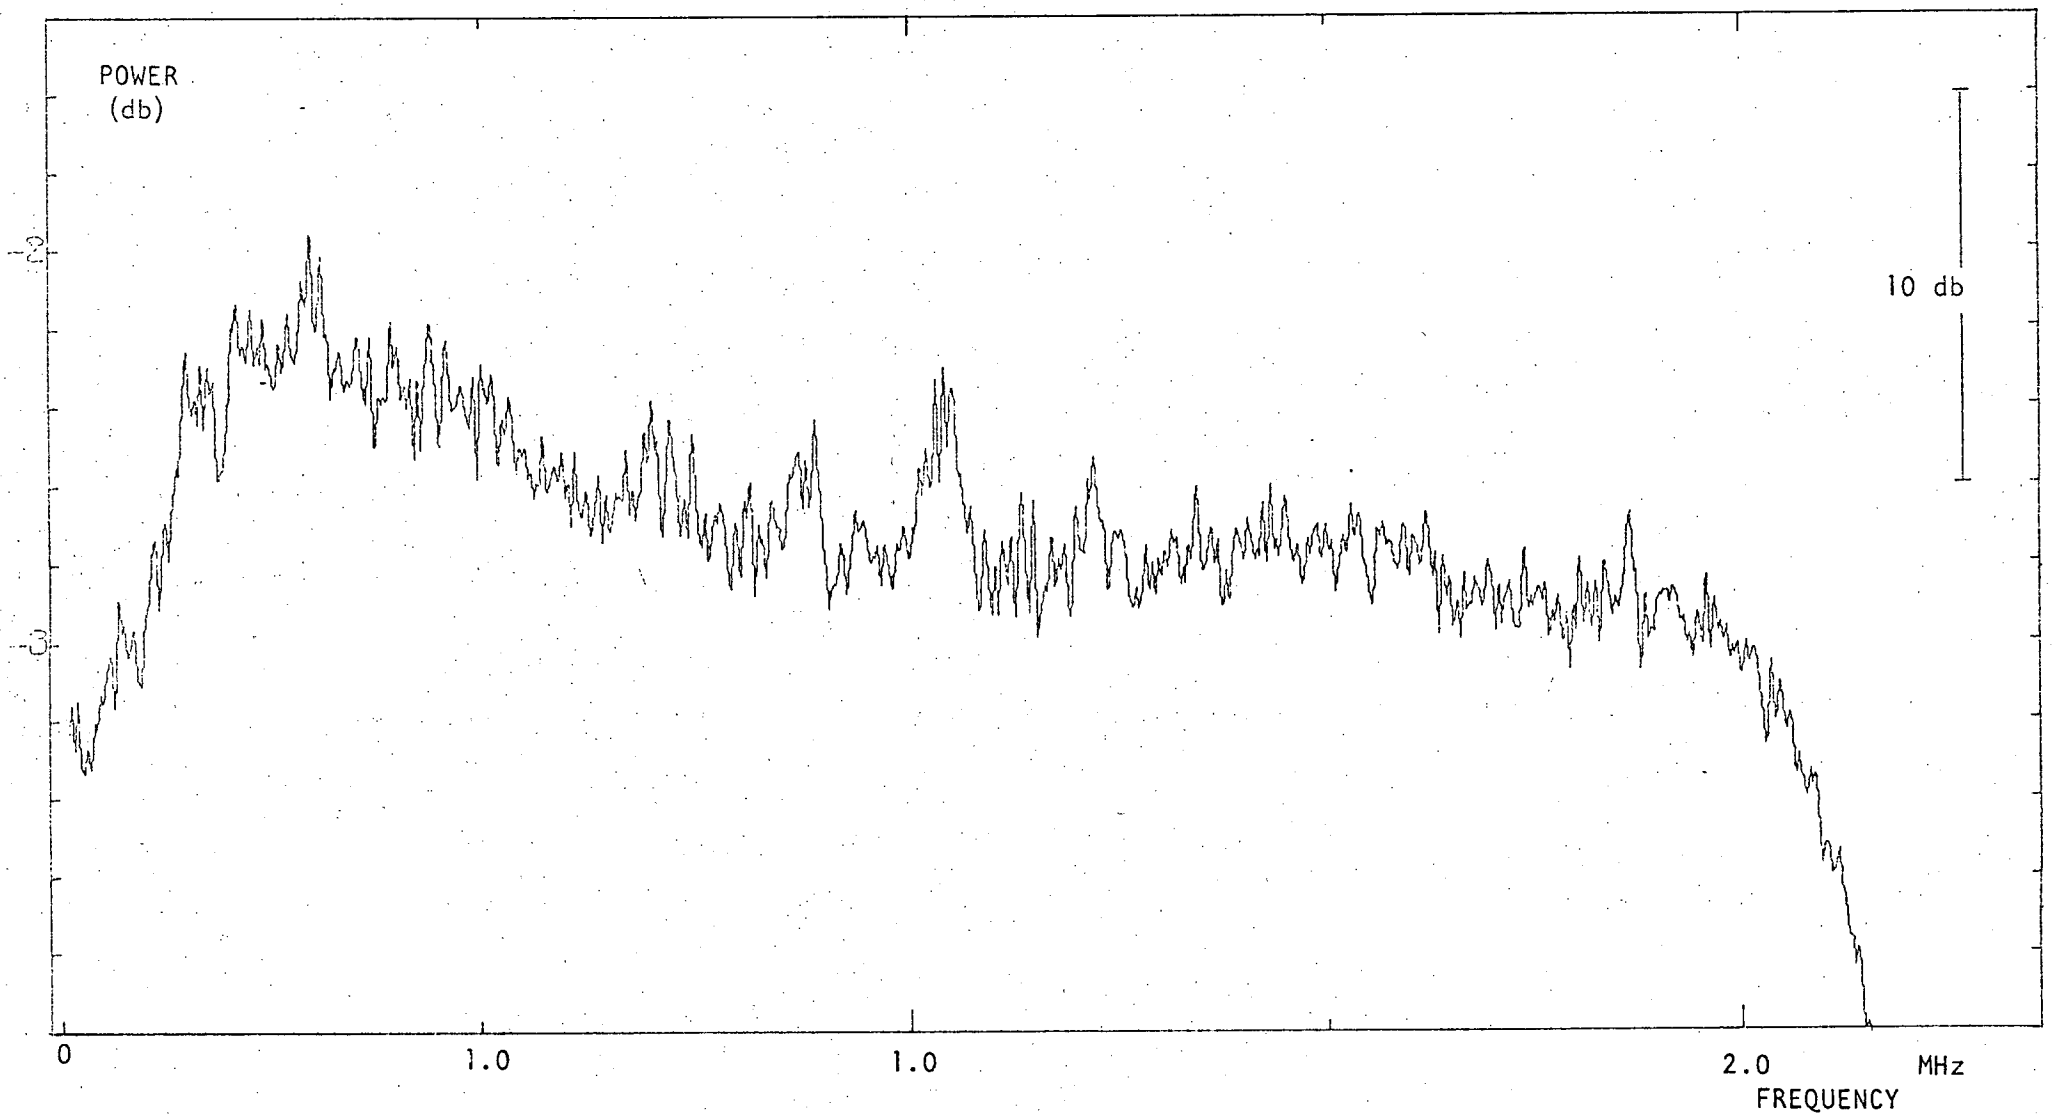

FIGURE 7.3C Power spectrum of pulse C of PSR 1749-28 with a resolution of 4.6 kHz. The spectrum was calculated from a time window of 14 ms (the pulse occupies about 6 ms after dispersion removal). Each point has a 1σ uncertainty of about 2 db.

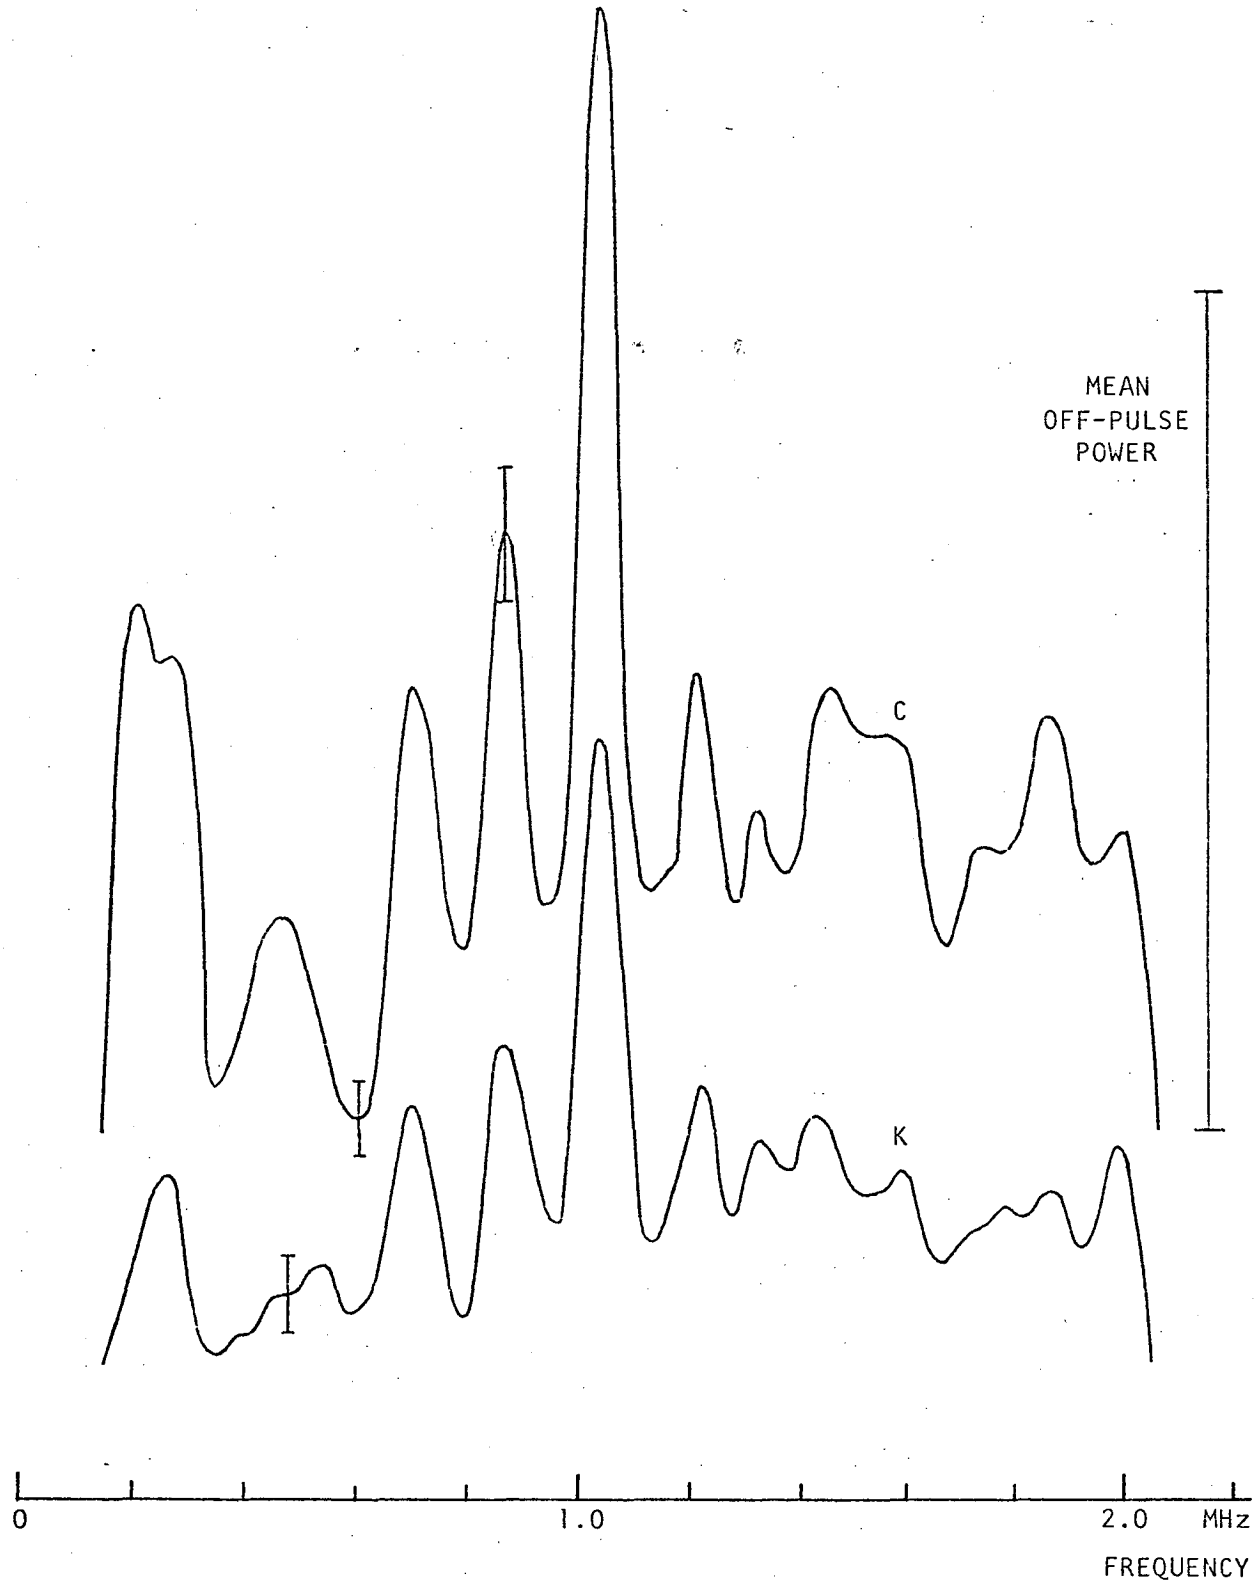

FIGURE 7.3D The power spectra of pulses C and K of PSR 1749-28 after passband flattening showing scintillation fringes. The curves are plotted on a linear scale and each has been truncated at the level of the off-pulse background. The time resolution is 70 kHz. Typical 1σ error bars are shown. Each spectrum is based on a time window of 14 ms; after dispersion removal each pulse occupies < 6 ms.

neighbouring fringes on timescales less than about 200 μs. That is the amplitude modulation is apparently smooth; and there is no obvious microstructure. (None of these dynamic spectrum plots are suitable for reproduction.)

7.4 OBSERVATIONS OF PSR 1133+16 AND PSR 1642-03

Parts of three successive pulses from PSR 1133+16 recorded on 1975 June 19 were analyzed. The first stripe, designated pulse "A", contains a subpulse within Component I of the integrated profile while the latter two, "D" and "F", contain subpulses of Component II. (At 408 MHz the centroids of the components are separated by 21 ms, so that in favourable cases the two components fall in successive stripes. However the "other halves" of these three pulses were deemed too weak for analysis.) The power spectra of the three pulses, smoothed to a resolution $\sim$ 200 KHz, are very similar. Each spectrum rises with decreasing observing frequency and has a shallow dip at 407.8 MHz (1.2 MHz video frequency). This dip is probably due to interstellar scintillation for which the decorrelation frequency was estimated to be $\sim$ 2 MHz (Section 6.4.2)

Intensity profiles reveal pulses A and D to be single amorphous subpulses about 5 ms in duration. Both are weak with signal-to(off-pulse) noise ratios of 0.7 and 0.4 respectively. Pulse F is narrower with a peak signal-to-noise ratio (averaged over 0.5 ms) of about unity. The intensity profile of this pulse is shown in Figure 7.4A. The lower curve in this Figure shows the data smoothed over 24 μs. Since it is also averaged over channels of 140 KHz bandwidth covering 2 MHz, the 1σ uncertainty is ± 0.14. Therefore some of the features in this profile are real, such as the micropulse 75 μs wide at 1.8 ms from the start of the trace. The subpulse contains micropulses ranging in duration from about 50 to 500 μs. Inspection of the dynamic spectrum of the subpulse revealed little of interest and it is well described as noise-like.

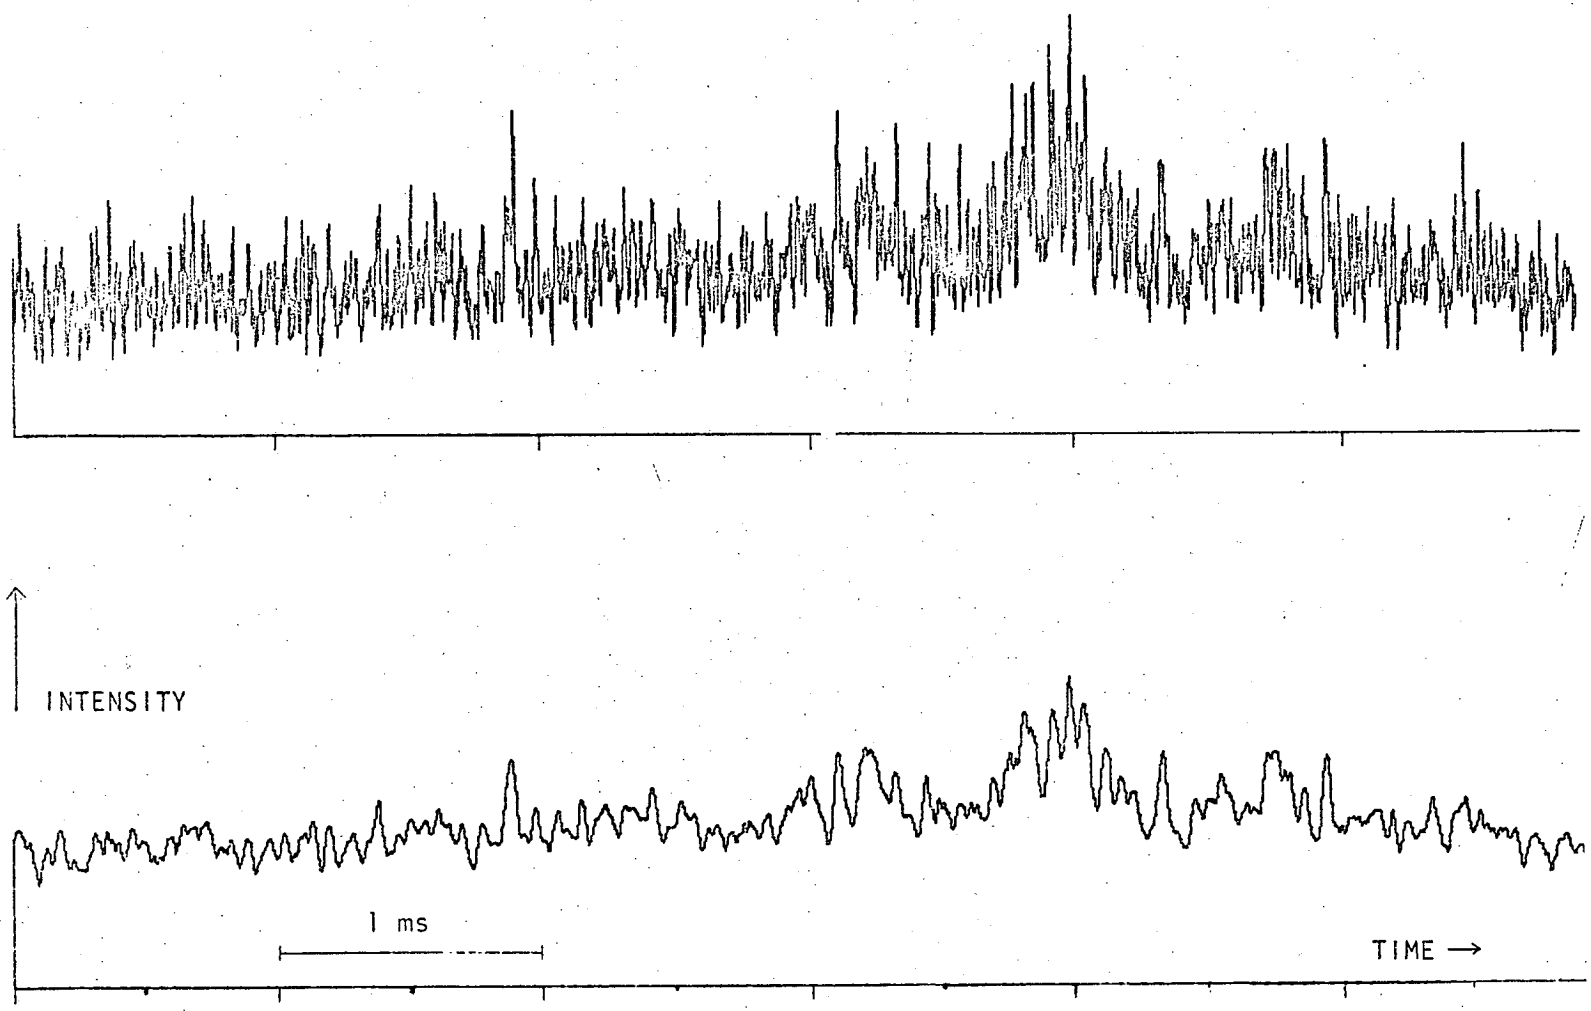

AVERAGE INTENSITY ACROSS CHANNELS OF 140 kHz BANDWIDTH. UNSMOOTHED, τ=7.3 μs. 28 DEGREES OF FREEDOM

AS ABOVE, ON THE SAME SCALE, BUT SMOOTHED TO τ=24 μs. (92 D.O.F.)

FIGURE 7.4A The strongest part of the intensity profile of PSR 1133+16 pulse "F".

The dynamic spectra of the other two pulses are almost indistinguishable from the off-pulse background.

Only two pulses from PSR 1642-03 were selected as being strong enough to be worth digitizing. These were found to have signal-to-noise ratios less than 0.5 (over 3 ms). Both pulse profiles are apparently amorphous and about 4 ms wide. Their average spectra and their dynamic spectra show weak emission fringes which are apparently due to interstellar scintillation. The signal-to-noise ratios are too low for positive identification of the fringes but their average spacing is $\sim$ 200 KHz. This is in good agreement with an estimated 70 KHz for the decorrelation frequency (Section 6.4.2).

7.5 SUMMARY

The observations described in this Chapter demonstrate that the videotape system can be used successfully. There are, of course, some problems and these will be discussed in the next Section. However the emergence of micropulses in pulses A, T and W of PSR 0950+08 indicates that the dispersion removal is effective.

Microstructure has been observed in pulses from PSR 0950+08. The signal-to-noise ratio of the observations of the other pulsars is insufficient to allow definite decisions as to whether microstructure is present or not. However, the PSR 1749-28 pulses seem to be amorphous.

No significant deviations from the AMN model have been found. In all cases the distribution of intensity relative to the running mean is consistent with the model. The appearance of the intensity profiles is noise-like especially at the highest time resolution of 0.5 μs. The dynamic spectra are also noise-like. The amplitude-modulation corresponding to each micropulse produces a general darkening (in the usual display mode) of the dynamic spectrum as a function of time.

However the "random speckle" usually makes the boundaries of the micropulse hard to discern.

The effects of interstellar scattering are visible in the dynamic spectra of PSR 1749-28 and PSR 1642-03 as "horizontal" fringes, that is, a modulation as a function of frequency. The fringes occur in the power spectra of the observations of these two pulsars as well as of PSR 1133+16. The de-correlation frequencies are approximately as predicted.

Finally I will mention some phenomena that have not been observed. First, the improvement in time resolution to 0.5 μs has not resolved any impulses or shots. If the noise process of the AMN model is shot noise then its temporal density is much higher than 2×10^6 s^{-1}. Second, the dynamic spectra have not shown any frequency-time structures that might be produced by moving, radiating sources or by interference effects in the pulsar magnetosphere. The hallmark of either would be slanting lanes or fringes (i.e. not parallel to either the frequency or the time axis). There is no strong evidence for sporadic narrow-band emission. (This would be distinguished from interstellar scintillation fringes by *not* persisting from one pulse to the next.) However the energy of the strongest micropulse of pulse W of PSR 0950+08 is concentrated in a band of 400 KHz (see Figures 7.2F and 7.2G).

7.6 INSTRUMENTAL CONSIDERATIONS

This section deals with the performance of the videotape system as revealed by the observations described in this Chapter. Two aspects are to be considered, namely the level of off-pulse noise and the system gain as a function of frequency.

In the calculation of passband flattening coefficients from off-pulse stripes (see Section 7.1 paragraph 2) it was found that the level of the background varies between different off-pulse stripes, even those recorded using the same head. Similar variations were found between 3 ms

segments within a stripe. The differences occur in the mean or slope of a smoothed power spectrum and are $\lesssim$ 0.5 db at a given frequency. After passband flattening the off-pulse intensity averaged over 2 MHz and $\sim$ 1 ms varies by $\sim$ 0.5 db. In the case of pulse C of PSR 0950+08, whose profile appears in Figure 7.2A, the apparently off-pulse background following the strong subpulse is 0.6 db below the level in the calibration stripe. (This is with the same passband flattening coefficients for each stripe.) It is interesting that six of the seven strong subpulses shown in Figure 7.2A are apparently preceded by weak continuous emission. In each case the intensity drops by about 0.6 db from one side of the subpulse to the other. This could represent a systematic decrease in system gain along the stripe. However there is no clear pattern of variation along off-pulse stripes. In general the variation may be due to videotape quality, or changes in tape-to-head contact along a stripe during stop-action replay.

The uncertainty of $\sim$ 0.5 db or 10% in off-pulse background level is important for numerical tests of the AMN model which are described in Chapter 9. In these tests the mean background intensity must be subtracted from the data. Because of the small signal-to-noise ratios the uncertainty in this intensity dominates the confidence limits of the tests.

Passband flattening has been included in the data analysis to equalize the system gain across the spectrum. However there is strong evidence that this has not been achieved. For all of the on-pulse data except for the two weak pulses from PSR 1642-03 the power spectrum shows the same underlying trend. After passband flattening the spectrum rises with increasing video frequency, that is, decreasing observing frequency. The trend can be seen in Figure 7.1B and to a lesser extent in Figure 7.3D.

Average signal-to-noise ratios vary by a factor of two across the band. That is, the "flattened" system gain increases by 3 ±1 db as the signal frequency decreases from 408.8 to 407.0 MHz.

The passband flattening has been based on the assumption that the observed off-pulse noise is proportional to the system gain at each frequency. Clearly this has not been the case. The discrepancy may arise in two places. First it may be due to the videotape recorder's FM section as modified by the head resonance adjustment (Section 4.5.2). Second the antenna pre-amplifiers at Molonglo during 1975 (replaced in 1977) may each have had a noise spectrum sufficiently different from the gain response of the antenna feed/pre-amplifier combination.

The preceding discussion has been based on the observations of 22 pulses and their calibration stripes. Of course these observations are only of pulses which were digitized successfully. A few others were so strong that the videotape recorder was over-loaded to an extent that the data could not be recovered even using the head resonance adjustment. Future observations will use a lower recording level. Other pulses were rejected because they were too weak, and this depends on the radiotelescope gain and the overall system noise. Still others were rejected because they fell on stripe boundaries. This is a defect of this application of a helical-scan videotape recorder. However it could be remedied by recording the same signal using two closely matched recorders with their heads rotating out of phase by half the stripe duration, in this case by 10 ms, and with a common reference tone. Very little extra circuitry would be required and this would be in the recording chassis only. In practice a particular pulse would usually be digitized from one stripe on one reel of videotape or the other. However with care it should be feasible to combine the data.

CHAPTER 8

DETAILED INVESTIGATIONS OF PULSE "T"

8.1 INTRODUCTION

This Chapter describes attempts to establish whether or not two large intensity spikes in the observation of pulse T of PSR 0950+08 are single near-ideal impulses emitted by the pulsar. For example, they could be large shots in a shot noise process. The null hypothesis is that they are consistent with the AMN model, that is, noise fluctuations.

The observation which initiated this work is shown in Figure 8.1A. This is a plot of the central part of pulse T obtained by analyzing the data with a 2.1 MHz channel. There was no post-detection smoothing, so that the time resolution is 0.5 μs, the samples being 224 ns apart. There are two large unresolved intensity spikes in the plot, events "a" and "b". They rise to 52 and 44 times the off-pulse mean intensity, that is, to 6500 and 5500 Jansky (Jy) respectively. There is a third spike, "c", of 3500 Jy which is important in the statistical tests.

The possibility of events a and b being single impulses is not ruled out *a priori* by the fact that they are surrounded in time intervals ~10 μs by other intensity spikes. However these other spikes contribute to the "noise" background with which a and b are compared (Section 8.2).

A number of different tests are described in the following Sections. They involve the following aspects of the data:

(i) the intensity distribution;

(ii) the dynamic spectrum;

(iii) variation of peak intensity with bandwidth;

(iv) variation of peak intensity, width and position in time with variation in the DM assumed in the dispersion removal.

The conclusion reported in Section 8.6 is that all of the data including events a, b and c are consistent with the AMN model. That is, a and b are not single impulses.

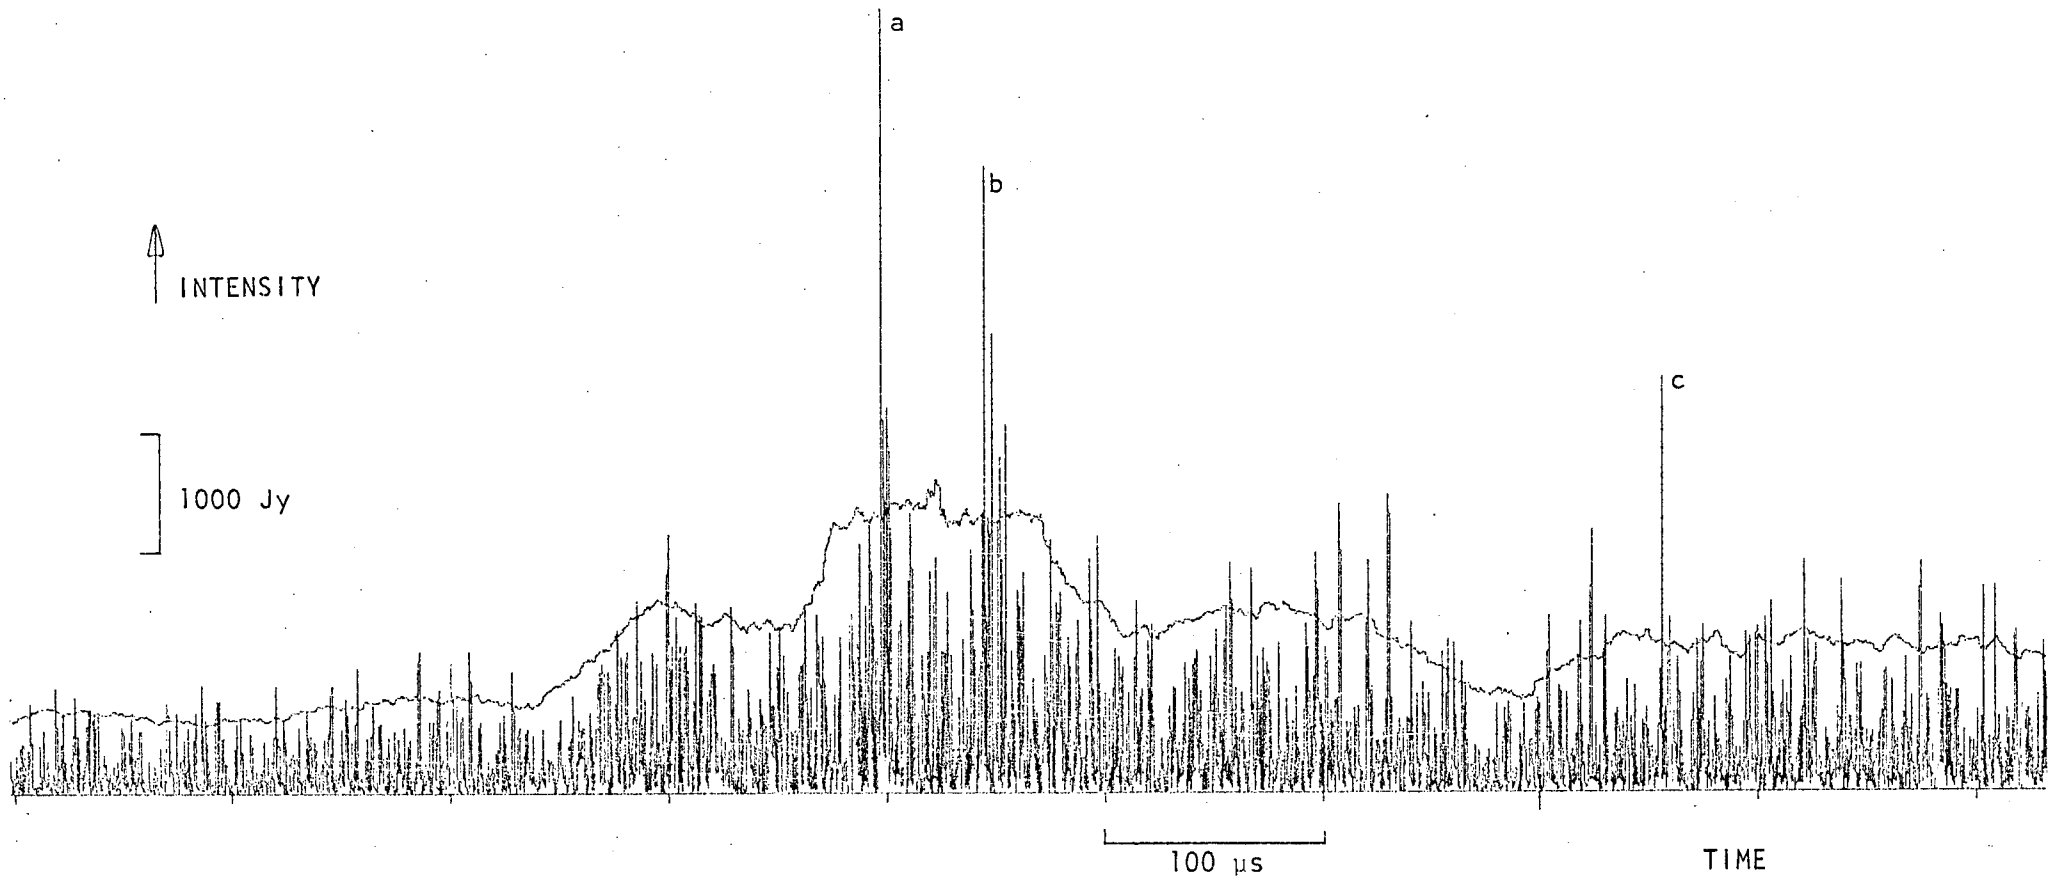

FIGURE 8.1A The strongest part of pulse T of PSR 0950+08 observed with a bandwidth of 2.1 MHz. The data is unsmoothed and so has a time resolution of 0.5 μs; samples are 224 ns apart. The smoother curve is 3.69 times the running-mean over 53 μs and represents a 97.5% confidence limit. The events a, b and c are discussed in the text.

After the work reported in this Chapter had been completed and the conclusions drawn it was found that the correction for instrumental phase had been accidentally omitted from the dispersion removal. This correction is only important for the full 2.1 MHz bandwidth and is completely negligible for bandwidths ~ 250 kHz. In the corrected 2.1 MHz intensity profile the peak intensity of event a is reduced by 20% so that the relative intensity test does not indicate a significant departure from AMN (P = 0.1, formerly 0.004). The overall conclusions of this work are not changed.

8.2 INTENSITY DISTRIBUTION

This section describes the comparison of the peak intensities of events a, b and c with the distribution predicted by the AMN model. In this model the pulsar intensity observed when there is no post-detection smoothing can be written

$$I(t) = A(t)N(t)$$

where $N(t) = |n(t)|^2$ and $n(t)$ is a complex Gaussian random variable. The amplitude-modulating function, $A(t)$, is a positive, non-stationary random variable that varies on timescales less than ~ 1 ms. $N(t)$ is a chi-squared variable with two degrees of freedom, that is, it has an exponential distribution. It can be assumed without loss of generality that $N(t)$ has unit mean and contains a contribution from the system noise with $A(t)$ having a constant component equal to the level of the system noise. The value of $I(t)$ at a fixed time is a sample from an exponential distribution of mean $A(t)$. Therefore, to determine if a particular measured value $I(t_1)$ is consistent with the distribution, the value of $A(t_1)$ must be estimated from the data. I have used the method of Hankins (1972), that of taking the running mean over some time window t_{rm} to obtain an estimate, $M(t)$, of $A(t)$. The relative intensity, I', is then defined by

$$I'(t) = I(t)/M(t).$$

This definition forces stationarity onto the data to allow conventional statistical tests. For a sampled system with a sample spacing of δt the relative intensity at time $j\delta t$ is

$$I_j' = I_j/M_j$$

$$M_j = \left(\sum_{k=-m}^{m} I_{j+k} \right)/(2m + 1)$$

where $t_{rm} = (2m + 1)\delta t$

How should t_{rm} be chosen? Clearly good accuracy requires $t_{rm} >> 1/B$ where B is the channel bandwidth since that will give an average of Bt_{rm} statistically independent values each with two degrees of freedom. Hence the standard deviation in M_j is approximately $A(j\delta t)/\sqrt{Bt_{rm}}$. However, t_{rm} must be the time over which A(t) is approximately constant. By calculating the autocorrelation function for a large number of pulses Hankins (1972) obtained 175 μs as the characteristic time, Δt_μ, for PSR 0950+08, and then used that as the value of t_{rm}. However since Δt_μ is the break point in an ACF curve it is approximately an upper limit on the timescale of variation of A(t). Furthermore it is an average over many pulses. It is not necessarily applicable to any one pulse nor to the particular structure in A(t) which underlies a particular fluctuation I_j. Therefore I chose m=120 and hence 53 μs as an initial value of t_{rm}.

The running-mean, M_j, over time t_{rm} = 53 μs was calculated and a confidence limit curve based upon it is shown in Figure 8.1A (and for each curve in Figure 8.3A). This curve is the 97.5% confidence limit (following Hankins 1972), that is, $I_j = \ell og_e(1/.025)M_j = 3.69\ M_j$. Spikes a, b and c clearly exceed this limit and are in fact 10.5, 8.4 and 10.2 times the running-mean, respectively. The probability of a given sample having a relative intensity $I' > x$ is exp(-x), for example 3.7×10^{-5} for x = 10.2. However there are a large number of data points. I took as a statistical sample the 3800 intensity points in a window of 0.85 ms containing the strongest part of pulse T and the events being investigated. The observed distribution of relative intensity is well fitted by an exponential distribution except for the two events a and c. The probability of two or more events with $I' > 10$ occurring among 2000 points (every second of the 3800 is independent) is 0.004, which implies that a and c may well be single

impulses. Note however that one event with $I' > 10$ would not be significant ($p = 0.09$) so that further tests must treat a and c together. Event b is accepted as being consistent with the AMN model.

I investigated the variation of the relative intensity of the events as t_{rm} is varied. As expected for $t_{rm} \sim 1$ ms the I' distribution is not exponential but has a greatly enhanced tail. That is, 1 ms is a very poor estimate of the timescale of variation of the modulation. Conversely the fact that the observed distribution for $t_{rm} = 53$ μs is accurately exponential shows that this time is a good estimate of that timescale. This may be confirmed by inspection of Figure 8.1A although a reduction of t_{rm} by half would seem to give a better fit near event b. However when t_{rm} is reduced to 22 μs the fit to an exponential is poor, with the tail beyond $I' = 7$ being depressed. This does not imply anything about the events since similar treatment of off-pulse data produces almost the same distribution. The depression of the tail is due to the smallness of the sample used to calculate the running mean, namely 100 points of which every second is independent (2 MHz bandwidth with $\delta t = 224$ ns). In each summation to obtain M_j the central point makes a significant contribution.

Some running-mean analyses were performed leaving out the central point and its two neighbours. This procedure changed the $t_{rm} = 22$ μs distributions approximately as expected. The revised off-pulse distribution is exponential, however, the on-pulse distribution has a significantly enhanced tail ($I' > 6$). The procedure was abandoned since it is of questionable validity.

8.3 DYNAMIC SPECTRUM

Figure 8.3A shows the same 2 MHz × 1.3 ms of pulse T data divided into two contiguous channels of 1 MHz bandwidth and four channels of 0.5 MHz bandwidth. In each case there was no smoothing so that the time resolution is 1 μs for each of the 1 MHz channels and 2 μs for the other four. Careful inspection shows that none of the events a, b and c involves good alignments of peaks in neighbouring channels. The spectrum of each event appears to be deeply modulated.

This impression is confirmed by observation of the dynamic spectrum of pulse T (Figures, 7.2H, 8.3B and 8.3C). Although some enhancement of the intensity at the times of events a and b can be discerned the dynamic spectrum is clearly best described as noise.

8.4 VARIATION OF PEAK INTENSITY WITH BANDWIDTH

One impression gained from Figure 8.3A is that as the bandwidth is increased from 0.5 MHz to 2 MHz the events a, b and c grow dramatically out of the noise. Attempts were made to quantify this impression.

I considered the set of channel bandwidths 2, 1, 0.5 and 0.14 MHz and investigated the way intensity, $\overline{I}_B$, averaged across channels of a given bandwidth, B, varied with bandwidth. $\overline{I}_B$ is defined by

$$\overline{I}_B(t) = \sum_{j=1}^{k} I^{(j)}(t)/k$$

where k is the number of channels, kB = 2.1 MHz, and $I^{(j)}(t)$ denotes the intensity at time t in the j th channel. Since only samples of intensity $I^{(j)}(t)$ exist and at a spacing $k\delta t = 1/(2B)$ the intensity at a particular time is calculated by interpolation. The particular times of interest are of course the times of the three events. From the available channel bandwidths four values of $\overline{I}_B$ are obtained for each event. They are compared

FIGURE CAPTIONS

FIGURE 8.3A A single pulse ("T") from PSR 0950+08 observed June 19 1975, following dispersion removal. The intensity obtained simultaneously in 7 channels is plotted as a function of time (1 division = 100 μs). The channels are arranged in 3 groups, each displaying the same 2 MHz of spectrum. From the top there are :

(i) a single channel of 2 MHz bandwidth, centre frequency 407.9 MHz, 0.5 μs resolution;

(ii) 2 neighbouring 1 MHz channels, centre frequencies 408.4, 407.4 MHz, 1 μs resolution;

(iii) 4 channels of ½ MHz, centres 408.7, 408.2, 407.6, 407.1 MHz, with 2 μs resolution. The smoother curves are 97½% confidence limits. Events a, b, c are discussed in the text.

FIGURE 8.3B The dynamic spectrum of pulse T of PSR 0750+08 in a window 630 μs × 2.0 MHz. Time is horizontal and observing frequency increases upward. Intensity is represented by the blackness of the dots; the data is shown twice at different contrasts. Events a and b are marked; c is another 130 μs off the diagram to the right.

FIGURE 8.3C Another display of the dynamic spectrum of pulse T in the vicinity of events a and b. The box indicates the time and frequency resolution.

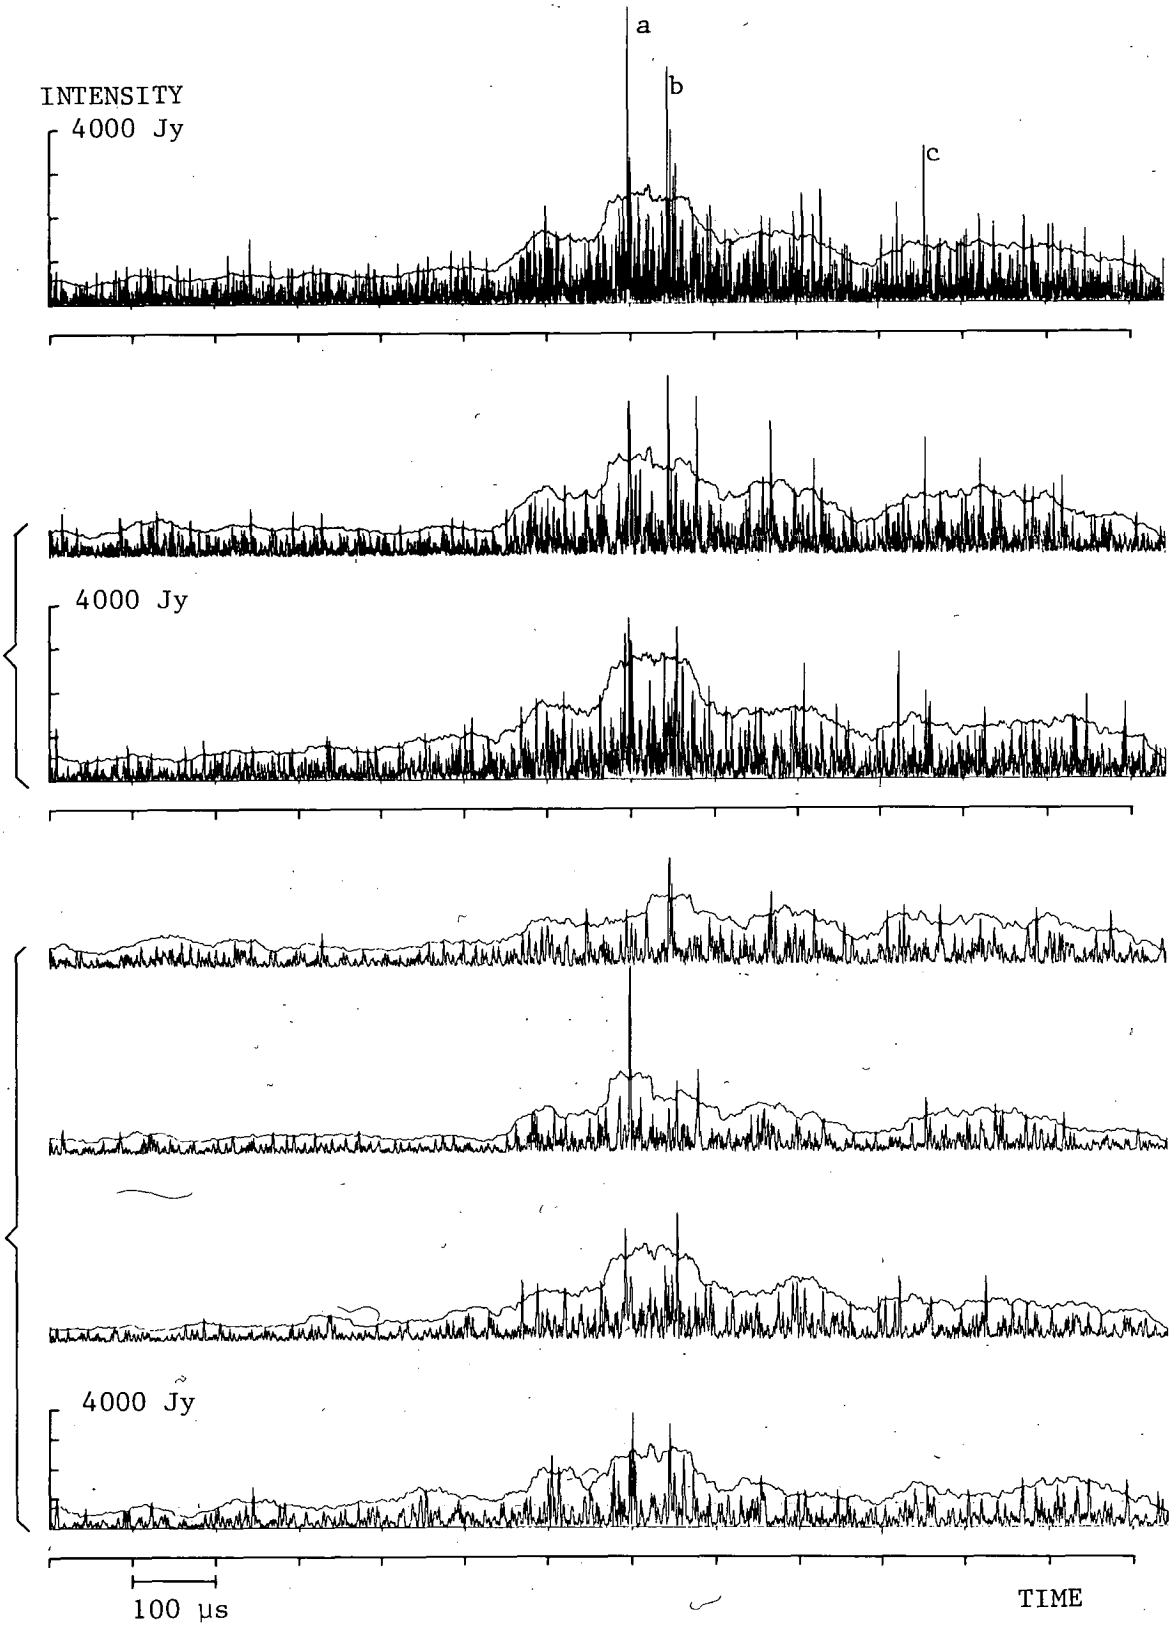

FIGURE 8.3A

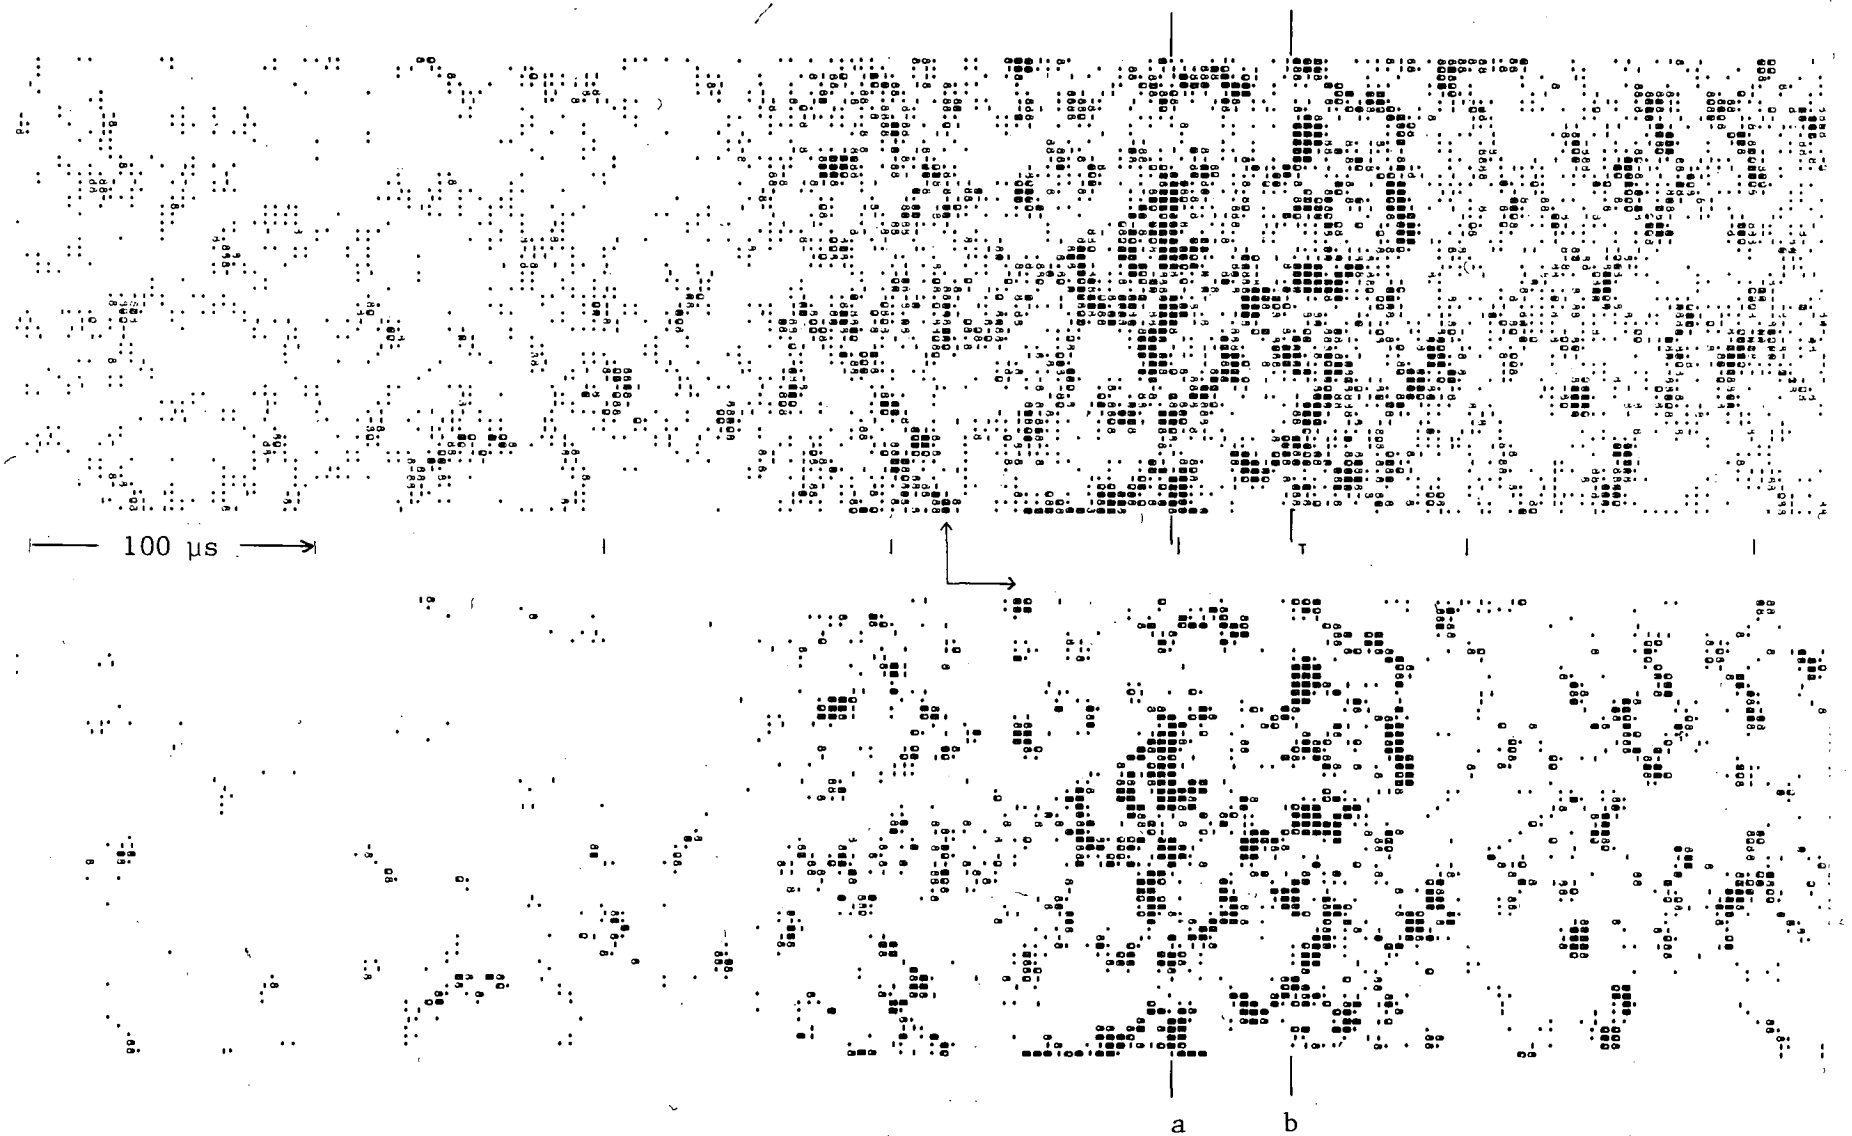

FIGURE 8.3B

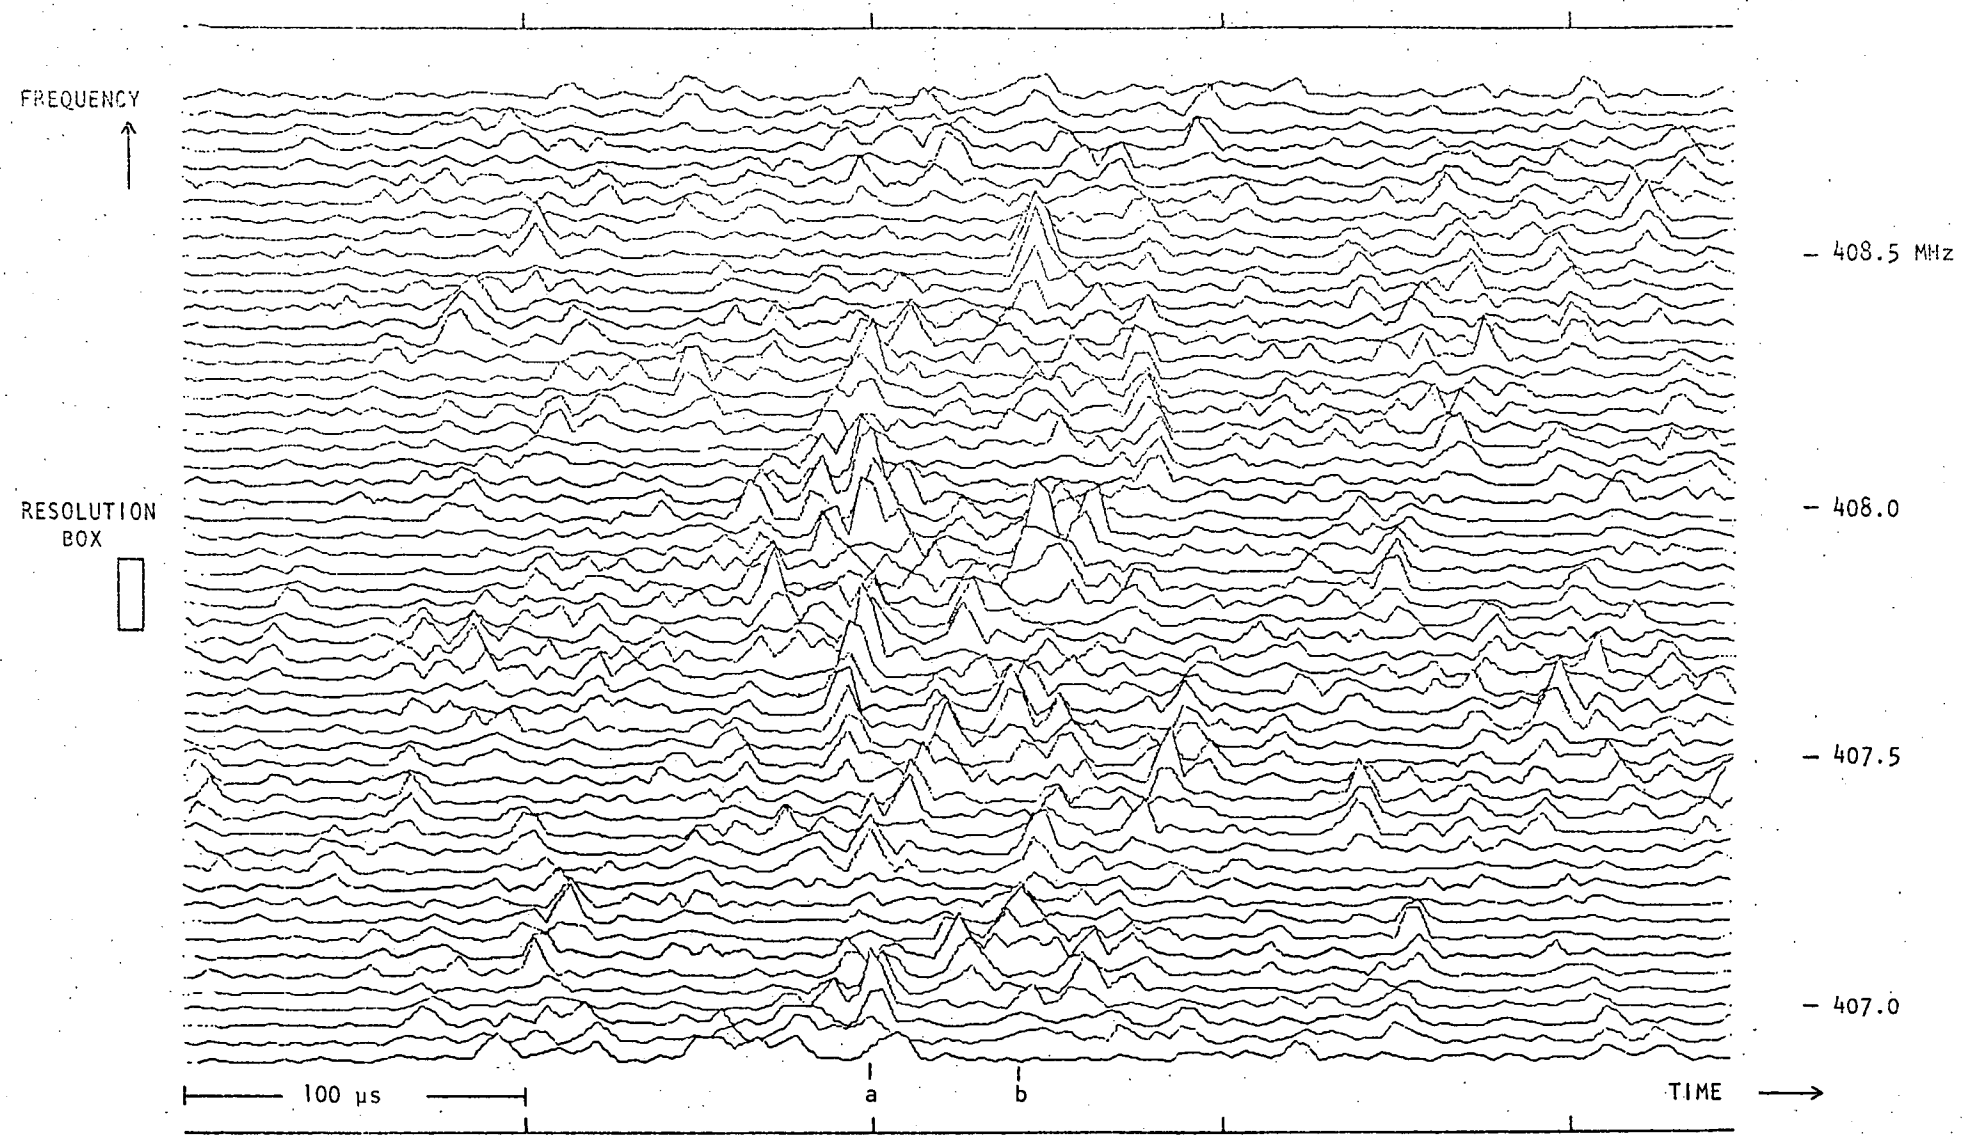
FREQUENCY
RESOLUTION BOX
408.5 MHz
408.0
407.5
407.0
100 μs
a
b
TIME

FIGURE 8.3C

with the power law model

$$\langle \bar{I}_B(t) \rangle \propto B^m,$$

that is, $$\langle \bar{I}_B(t) \rangle \propto k^{-m}, \quad k = 1, 2, 4, 14$$

where the angle brackets denote the expectation value. For Gaussian noise at a randomly chosen time one would expect m = 1. For an ideal impulse in a noise-free system at the arrival time of that impulse one would expect m = 2. For events a, b and c I obtained m = 1.60, 1.54 and 1.72 respectively. For all 4000 intensity values of the 2 MHz channel on-pulse data in a window of 0.9 ms the observed range of m is -1.0 to +1.73 and the distribution is similar to that for off-pulse data. The values of m for the three events are above the 98 percentile of the distribution; however this is not additional evidence for the events being impulses.

There are two selection effects in this power-law analysis. First, the time at which intensities are compared across channels is determined by the peaks in the 2 MHz intensity. Therefore positive values of m would be expected. This is reinforced by the second selection effect, namely that the peaks in question are in fact very large peaks, that is, $I' > 8$. There must be a positive correlation between I' and m for intensity peaks.

The correlation was investigated experimentally by considering as a statistical sample all those peaks occurring among the 4000 on-pulse intensity values, and similar off-pulse data. The two scattergrams of m versus I' obtained are shown in Figures 8.4A and 8.4B. Apart from the three events the two distributions are very similar. The curved trend line for off-pulse data, which was obtained as a fit to the medians of each univariate distribution at constant I', is an excellent fit to the on-pulse distribution as well. Furthermore the events lie close to an extrapolation of the trend line.

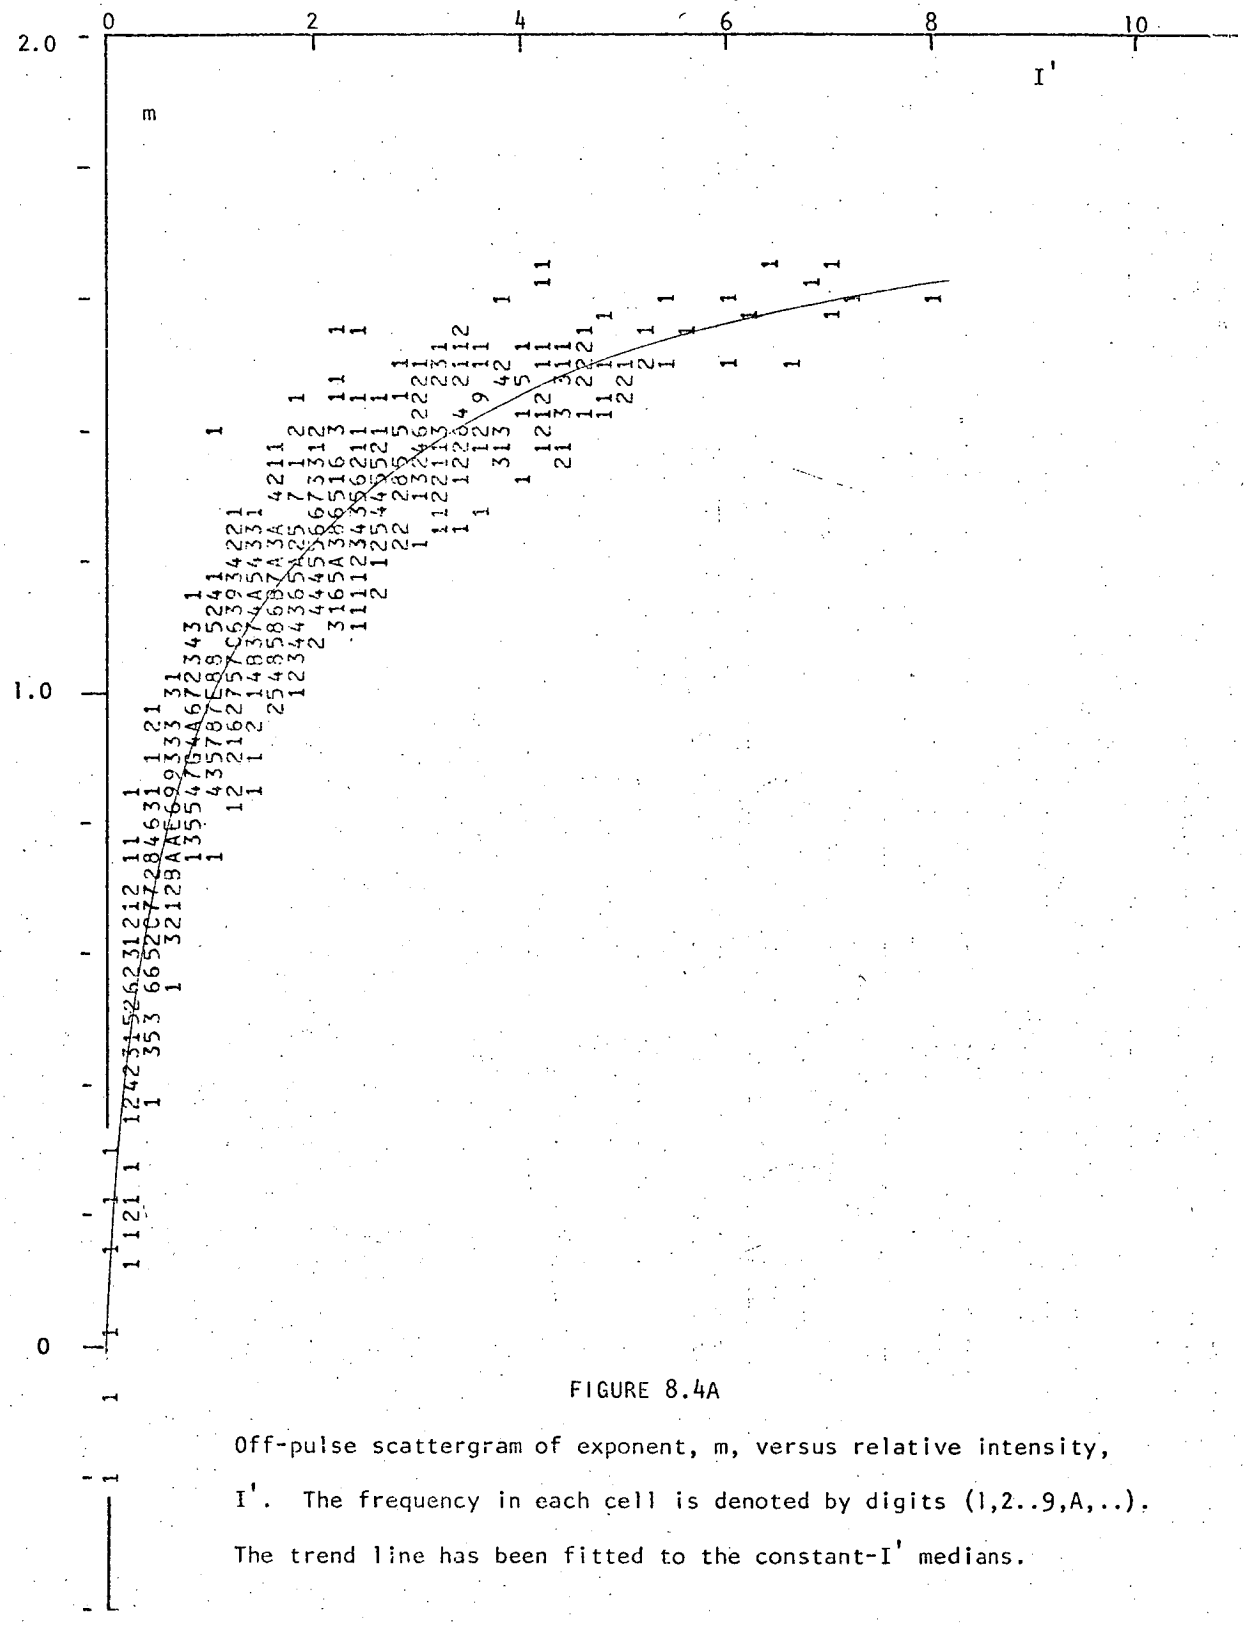

FIGURE 8.4A

Off-pulse scattergram of exponent, m, versus relative intensity, I'. The frequency in each cell is denoted by digits (1,2..9,A,..). The trend line has been fitted to the constant-I' medians.

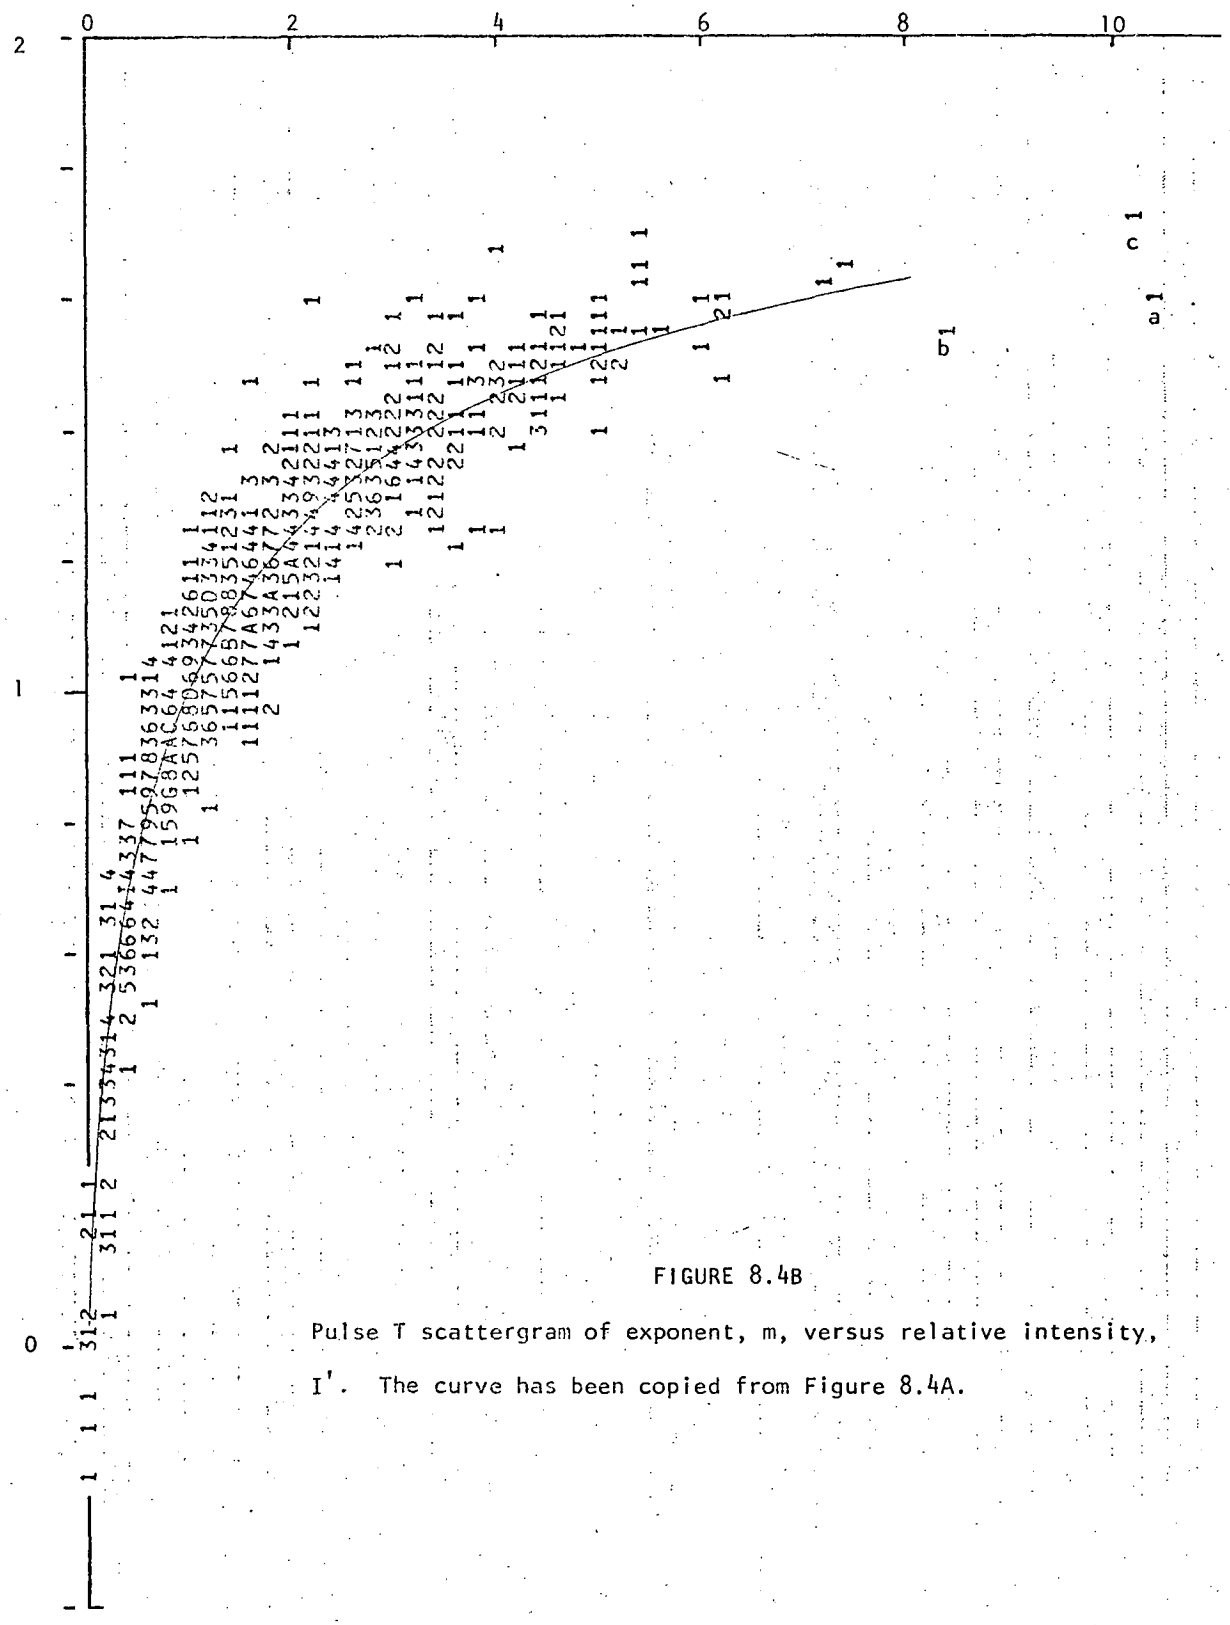

FIGURE 8.4B

Pulse T scattergram of exponent, m, versus relative intensity, I'. The curve has been copied from Figure 8.4A.

The conclusions are:

(i) this analysis of the effect of bandwidth is highly correlated with the relative intensity analysis and so provides little new evidence as to the interpretation of the events;

(ii) the fact that their m values are not abnormally high implies that the three events are consistent with AMN; and

(iii) the consistency of the 0.9 ms section of pulse T with AMN is demonstrated again.

8.5 VARIATION WITH DISPERSION MEASURE

Another approach to investigating the events a and c of pulse T is to vary the value of the dispersion measure, DM, which is assumed in the analysis. If the pulsar were to emit large impulses they would be best observed when the assumed DM matches that of the pulsar, namely $2.969 \pm .001$ $cm^{-3}pc$. Events a and c were found to have maximum peak intensity at DM = 2.970 and 2.972 $cm^{-3}pc$ respectively ($\pm.0005$). These values of DM were established by interpolation to obtain better estimates of the intensity peaks than provided by the samples of intensity at intervals of δt = 224 ns, that is, to avoid the "picket fence" effect. A small sample of off-pulse spikes selected because they were large at DM = 2.970 $cm^{-3}pc$ actually peaked at values of DM taken from an approximately normal distribution with mean 2.970 and standard deviation 0.003.

This suggests that a and c are not both impulses from the pulsar. However the test is insensitive : suppose that a and b were in fact impulses then the background noise would produce errors in each measurement of the best DM by $\sim .001$ $cm^{-3}pc$. What would be a significant result? That the best DM values agreed to $\pm.001$ and either this was an unexpected value that later proved to be the true DM for the pulsar, or there were more than two events. Therefore the evaluation of best DM does not lead to any firm conclusion about whether events a and c are consistent with noise or not.

The DM sensitivity curve, that is, the plot of peak intensity versus the error in DM, ΔDM, was discussed in Chapter 6 (see Figure 6.5.1C) for impulses in a noise-free system. For a rectangular bandwidth of 2 MHz and for small values of ΔDM (i.e. before the Fresnel pattern develops multiple peaks) the sensitivity curve is approximately Gaussian with a full-width half-maximum (FWHM) of 0.012 $cm^{-3}pc$. Event c has this FWHM to a measurement accuracy of $\pm.001$ $cm^{-3}pc$, while a has a FWHM twice as great, namely 0.025 $cm^{-3}pc$. Broadening of the sensitivity curve could be produced by modulation of an impulse with respect to frequency, that is, by having a reduced equivalent bandwidth. However a small sample of off-pulse intensity spikes were found to have FWHM values in the range $.020 \pm .009$ so that the FWHMs of a and c are consistent with noise.

A characteristic of each intensity spike so far not considered is its width in time. An ideal impulse has a width (at 80% of peak) of 1.0/B or 0.24 μs. As ΔDM increases the peak evolves into Fresnel fringes and during this process the 80% width is, for small ΔDM, a weak function of ΔDM. The width effect is second order compared to the height effect, that is, the DM sensitivity curve. Indeed no clear trend could be obtained in measurements of the width of spikes a and c as functions of ΔDM. However c has a minimum width (at 80%) of $(1.08 \pm .05)/B$ which is close to that of an ideal impulse. The corresponding width for a is 1.35/B, which is consistent with the width of its DM sensitivity curve and implies an effective bandwidth $\sim$ 1.3 MHz. A small sample of off-pulse peaks were found to have widths in the range 1.15/B to 1.4/B, that is, similar to that of a and c.

The precise time at which intensity spikes occur varies with ΔDM. This follows from considering dispersion (removal) as the progressive imposition of a time delay across the passband. If the "pivot" frequency of the dispersion removal were the centre of the rectangular passband then the sense of the drift in time of a spike would depend on which side of the passband most of its energy lies. The effect can be quantified : the

observed rate of drift in time is linearly related to the equivalent centre frequency, f_{equiv}, of its spectrum. This is conveniently expressed as a relative frequency offset, $\Delta f'$, defined by

$$\Delta f' = 2(f_{equiv} - f_c)/B$$

where f_c is the centre frequency of the channel, 407.9×10^6Hz, and f_{equiv} is obtained from the observed drift rate

$$\Delta t/\Delta DM = (f_{equiv} - f_o) \times 1.205 \times 10^{-16}/f_c^3$$

This last equation is derived from the relation $\Delta t = (f_{equiv} - f_c)/\Delta\alpha$ where $\Delta\alpha$ is the residual bandwidth sweep rate for $\Delta DM(cm^{-3}pc)$. f_o is the pivot frequency of the dispersion removal, namely, 409.025×10^6 Hz corresponding to zero video frequency. Clearly $\Delta f'$ must lie in the range -1 to +1 and has a value of zero for an ideal impulse and an expected value of zero for AMN, that is, on average the energy which makes up each spike is centred in the channel. For a sample of 39 peaks in off-pulse data the standard deviation in $\Delta f'$ is 0.22, while for events a and c values of $\Delta f'$ of 0.01 and 0.00 were obtained. This would suggest that a and c are not typical of noise spikes since their energies are remarkably well centred in the band. However the sample of off-pulse peaks is not a good reference; the sample should be of peaks with high relative intensities and such a sample would have a lower standard deviation. This selection effect can be stated loosely as that the noise fluctuations that are largest are those made up of energy from the whole passband and hence the centre of their spectra must be the centre of the passband.

8.6 CONCLUSIONS AND DISCUSSION

The only test result which implies deviations from the AMN model is that of the relative intensity. Event b is accepted as being consistent with noise fluctuations. Then event a only remains in consideration as a possible single impulse if event c is also considered. None of the other tests show a significant deviation from the model. The conclusion is therefore that a, b and c are not single impulses.

Three tests treated the whole of the strongest part of pulse T (i.e. .85 ms or more) and in each case the data was found to be accurately described by the AMN model. These were the distribution of relative intensity, I', for $t_{rm} \sim 50$ μs; the dynamic spectrum; and the joint distribution of I' and the exponent in the bandwidth effect.

A relationship of statistical inter-dependence between all of the quantitative tests (i.e. excluding the dynamic spectrum) was found. This is a selection effect in that once an intensity spike is chosen because of a large value of I' then particular values of the other parameters become very probable. For example, if a noise fluctuation is selected because it is large at some value of DM, whatever that value is, then it is likely to have a maximum for $\Delta DM \simeq 0$. This selection effect is only obvious for noise.

Another selection effect must be considered. There is a large number of on-pulse data points in this project, so that the occurrence of several large I' values is likely. Was pulse T selected for this investigation because it was the only pulse that showed large I' values? No, in fact this was the first pulse which was analyzed at the highest time resolution of 0.5 μs. Therefore a case can be made for what was done, namely the taking of only 3800 points as the statistical sample on which to calculate the probability of two events with $I' > 10$.

Of the tests that were applied to pulse T the two that seem of most value for future use are the investigation of relative intensity and,

providing there are many more than two spikes being considered, the variation of peak height with DM. The latter test is capable of measuring the DM of the pulsar provided that the spikes are in fact single impulses, in which case they will all have the same "best" DM. Examination of the dynamic spectrum is not a crucial test because the individual channels do not have the maximum time resolution and hence neighbouring spikes, resolved using the full bandwidth, can interfere with each other. That is, the voltage is convolved with the receiver response not the intensity.

For investigating relative intensity the main problem is the estimation of the modulating function, A(t). An estimate must be averaged over sufficient points, e.g. 200, so that the standard error of this estimate is small. The averaging could be over ~200 channels (at some computational cost) in which case the time resolution of the estimate is that of an individual channel. The alternatives are to use a running mean either with one channel, as described in Section 8.2, or many. In all cases some choice has to be made about the time resolution of the estimate and this raises a methodological problem : Suppose there is a strong isolated spike such as in Figure 1.3.1B (Hankins 1972). It is always possible to choose, say, a running mean window narrow enough that the I' value is insignificant. The solution to this problem can only come when a minimum timescale for A(t) is observed. This would be exhibited as a flat top to the microstructure feature in the ACF curve (Chapter 9). No such shape has been reported.

Figure 8.6A shows part of a pulse from PSR 0950+08 observed by Hankins and Boriakoff (1978) who claimed that the large intensity spike is an unresolved micropulse. The time resolution here is 0.8μs. However this spike should be interpreted as a statistical fluctuation and the micropulse duration is at least a few microseconds. I have drawn a curve, g(t), in the Figure to represent the modulating function near the spike.

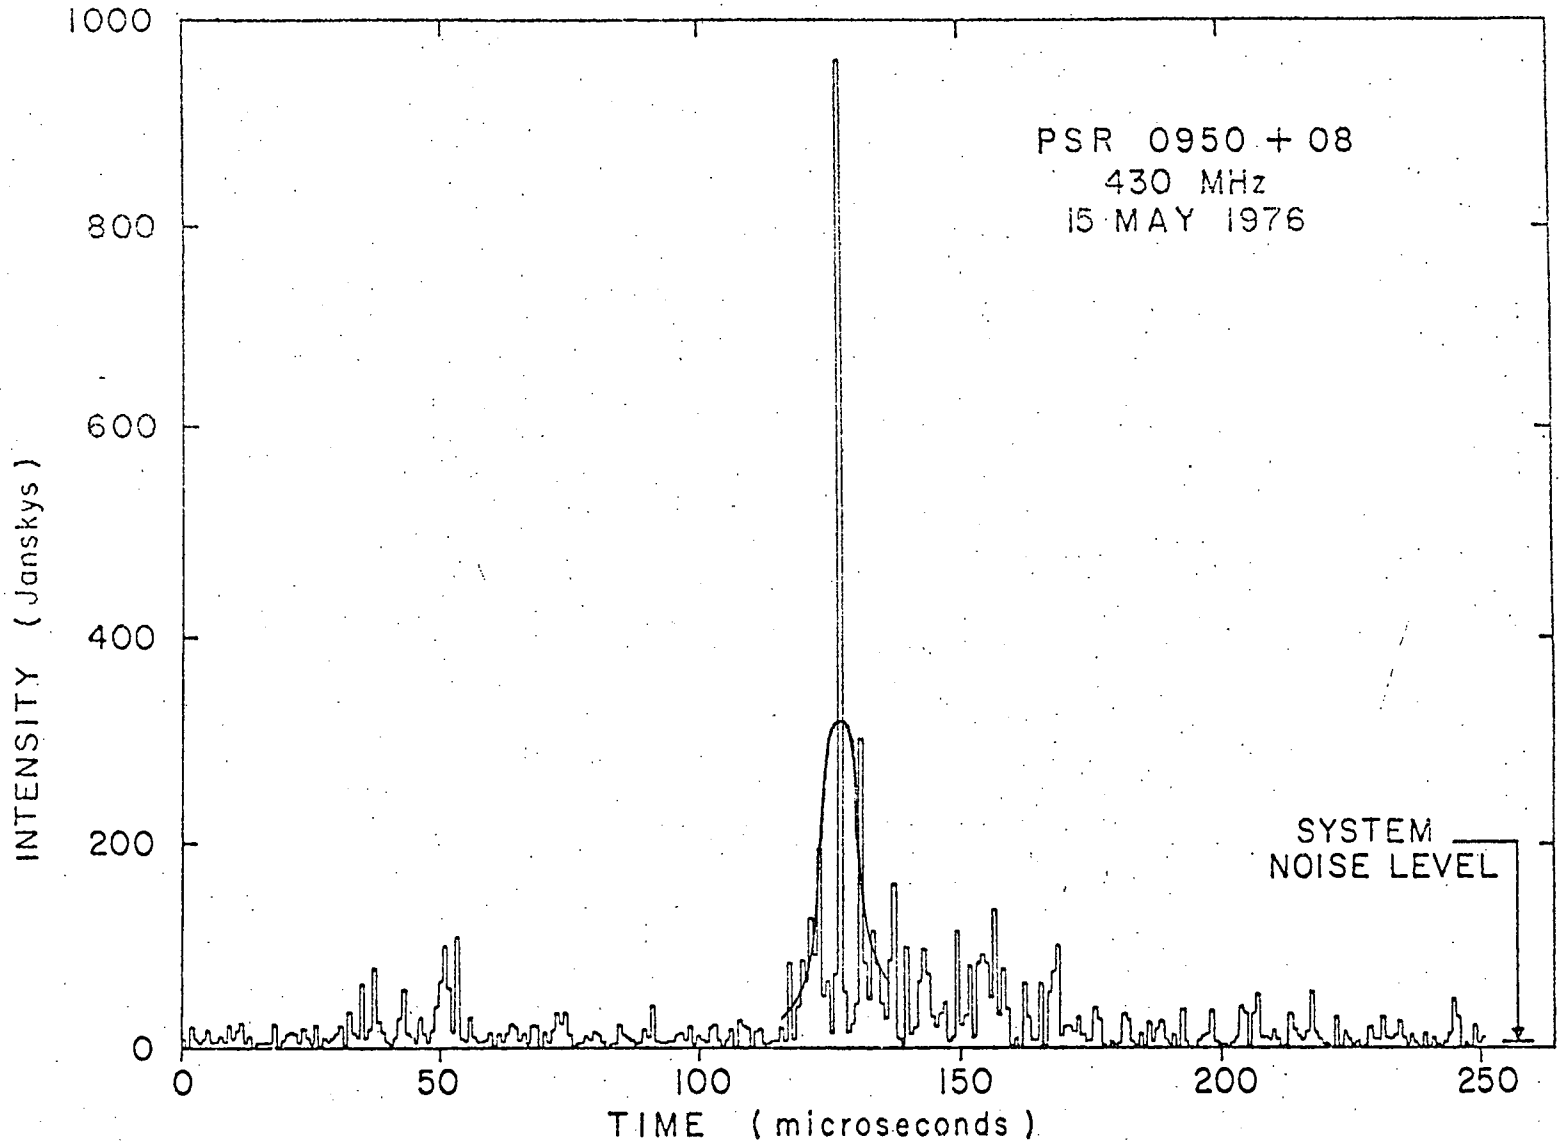

FIGURE 8.6A A pulse observed by Hankins and Boriakoff (1978) with 0.8 μs resolution. The smoother curve has been added and is discussed in the text.

In this case it has not been calculated from a running mean but it is an estimate of the ensemble mean intensity at each time. It is constrained to fit the data according to

$$0.05\ g(t) < I(t) < 3\ g(t).$$

That is, the observed intensity, I(t), must lie between the 5 and 95 percentiles of the exponential distribution with mean g(t). These are conservative constraints. One point, the 1000 Jy spike, has already been selected as unusual (i.e. a very large intensity). Therefore it is allowed to exceed the mean by a larger factor and hence the curve need not be as high or as narrow as drawn. Hence the micropulse duration may be greater than 10μs.

CHAPTER 9

AUTO- AND CROSS-CORRELATION ANALYSES

9.1 INTRODUCTION

Since pulsar microstructure involves stochastic processes the auto- and cross-correlation functions are natural ways of expressing theoretical models and of handling the data. When the AMN model was formalized (Rickett 1975) it was in terms of the auto-correlation function (ACF); the ACF and the cross-correlation (CCF) between pulses were used by Cordes (1976a) to model the radiation of PSR 2016+28.

As a data reduction tool the ACF has two main advantages. First there is a well defined way of averaging all the data. This is especially important when the signal-to-noise ratio is poor and it is usual when using bandwidths ~100 kHz to average the ACFs for many pulses together. However ACFs of strong single pulses have been published (Boriakoff 1976, Cordes 1976a, Ferguson *et al* 1976, Kardashev *et al* 1978). The second advantage of the ACF for data reduction is that no decision need be taken *a priori* about the timescale of the microstructure - the raw data can be unsmoothed and the timescales are revealed by the observed ACF (as break-points). If smoothing is needed to reduce the statistical fluctuations it can be done to the ACF. This situation contrasts strongly with the use of the running-mean, for example, where the time window has to be chosen in advance.

Because of the large amount of data for each pulse I have been able to obtain ACFs for single pulses in spite of the poor signal-to-noise ratios. This has been achieved by dividing the total bandwidth into narrow channels (70 or 140 kHz bandwidths) and averaging the ACF across channels. Of course this sets the time resolution to ~10 μs while the alternative, a 2 MHz channel and smoothing of the ACF as required, would have avoided this prejudice. However I chose the multi-channel approach for PSR 0950+08 so that the CCF between channels could be examined, and for PSR 1749-28 because the spectrum is broken up by scintillation fringes. I have also

evaluated the ACF of the power spectrum for each of these pulsars. Although a model for this ACF was given by Rickett (1975) there are no published accounts of it being evaluated. It yields equivalent information to the time domain ACF.

Note that the correlation functions obtained for single pulses represent the particular amplitude-modulating function, a(t), that is present. That is, the ACF of a single pulse or the ACF averaged over a few pulses, as here, does not necessarily characterize the microstructure of the pulsar. However in the case of PSR 1749-28 I have been able to obtain a characteristic timescale of 250 ± 50 μs. This was obtained from the power spectrum ACF and is a new result. The observations of PSR 0950+08 show structure on a similar timescale and are consistent with those of other workers.

As well as revealing timescales of the modulating function the ACF gives a specific numerical test of the AMN model. This is the height of the "zero-lag" spike which should be 0.5 for unsmoothed data. All of my observations are consistent with AMN in this respect.

The organization of this Chapter is as follows. Models of the ACF and CCF according to the AMN model are given in Section 9.2.1. These are basically as presented by Cordes (1976a) with parameters to represent deviations from the AMN model. The practical problem of background correction is dealt with in Section 9.2.2. The results for PSR 0950+08 and PSR 1749-28 are presented in Sections 9.3 and 9.4 respectively.

9.2 THEORY

9.2.1 Model Correlation Functions

In the AMN model the complex amplitude, $y(t)$, of the pulsar signal with dispersion removed is represented by the product

$$y(t) = a(t)n(t)$$

where the modulating function $a(t)$ is real and $n(t)$ is complex Gaussian noise with unit variance and a spectrum that is uniform over the receiver passband. The detected signal is the intensity

$$I(t) = a^2(t)|n(t)|^2 = S(t)b(t)|n(t)|^2 \qquad (1)$$

where $a^2(t)$ has been split into a subpulse part $S(t)$ and a microstructure part $b(t)$. It is assumed that $b(t)$ is stationary while $S(t)$ is non-stationary and has an ACF that is wider than that of $b(t)$. That is, $S(t)$ is assumed to be slowly varying, e.g. on a timescale ~5 ms, compared to $b(t)$, e.g. 0.1 ms. $b(t)$ is characterized by a modulation index, m_μ, and by the auto-correlation of its varying part. These are defined by

$$m_\mu{}^2 = [\langle b^2(t)\rangle - \langle b(t)\rangle^2]/\langle b(t)\rangle^2 \qquad (2)$$

and

$$\rho_\mu(\tau) = \langle b_1(t+\tau)b_1(t)\rangle/\langle b_1{}^2\rangle \qquad (3)$$

where $b_1(t) = b(t)-\langle b(t)\rangle$. The angle brackets denote ensemble averages.

Define the time auto-covariance (ACV) of a function $g(t)$ at a lag τ by

$$C_g(\tau) = \frac{1}{T}\int_0^T g^*(t)g(t+\tau)dt \qquad (4)$$

where the time window $[0..T]$ is such that $g(t)$ is zero outside it. Of

interest is the case where g(t) is the pulsar intensity I(t). Then the restriction above is that no pulse be truncated by the observing window. Note that all my observations involve single subpulses and no account is taken of pulsar phase so that in this context "subpulse" is the same as "pulse" and the observing window duration is much less than that of the integrated profile. Note also that, as in equation (4), I will use continuous notation for convenience in spite of the fact that in practice all the calculations must be made with sampled data.

The mean auto-correlation function, $r_g(\tau)$, of g(t) is defined by

$$r_g(\tau) = \langle C_g(\tau) \rangle / \langle C_g(0) \rangle \tag{5}$$

The intensity ACF for a single rectangular channel with no post-detector smoothing and no off-pulse noise is (from Cordes 1976a)

$$r_I(\tau) = \frac{r_S(\tau)[1 + m_\mu{}^2 \rho_\mu(\tau)][1 + \alpha \operatorname{sinc}^2 B\tau]}{[1 + m_\mu{}^2][1 + \alpha]} \tag{6}$$

where $r_S(\tau)$ is the subpulse ACF defined by (5) and α is the square of the modulation index of the noise $|n(t)|^2$. For Gaussian noise, as assumed by the AMN model, $\alpha = 1$. It is assumed that 1/B is much less than the characteristic width of ρ_μ and then the term $\alpha \operatorname{sinc}^2 B\tau$, which has a width ~ 1/B, forms what is known as the zero-lag spike. The particular shape, sinc^2, is because the passband is rectangular. The spike therefore has sidelobes with expected amplitudes of 0.023, 0.008, 0.004, etc. relative to the total height of the spike. The y-intercept at the base of the zero-lag spike is denoted by $r_I(0+)$ and defined by the extrapolation of $r_I(\tau)$ to $\tau = 0$. This extrapolation can only be done in practice if $\rho_\mu(\tau)$ has no spike component narrower than 1/B, that is, if b(t) is well-resolved using bandwidth B (Cordes 1976b). Under this condition,

$r_I(0+) = 1/(1+\alpha)$ in general and $r_I(0+) = 0.5$ for AMN in particular.

The model ACF, $r_I(\tau)$, is sketched in Figure 9.2.1A. The characteristic width, Δt_μ, of the microstructure is defined by the break in the slope where the microstructure feature meets the subpulse feature. The subpulse width, Δt_s, is defined as the half-width at half-maximum of the subpulse feature. The height of the microstructure feature is $H_1 = r_I(0+) - r_I(\Delta t_\mu)$ and this measures the strength of the microstructure relative to the steady component of the subpulses. This is expressed in terms of the modulation index, m_μ, and (Hankins and Boriakoff 1978) $m_\mu^2 = H_1/r_I(\Delta t_\mu)$. If $m_\mu = 1$ then $H_1 = 0.25$, however if the subpulses contain little microstructure, i.e. are "continuous" (Hankins 1971), then $m_\mu << 1$ and $H_1 \sim 0$.

A model for the intensity cross-correlation between channels is easily derived. Providing the channel passbands do not overlap, the noise in one is completely uncorrelated with that in the other. Therefore there is no zero-lag spike and from equation (6) we obtain the model CCF, namely

$$r_{II}(\tau) = \frac{r_s(\tau)[1 + \eta m_\mu^2 \rho_\mu(\tau)]}{[1 + m_\mu^2][1 + \alpha]} \qquad (7)$$

where $\eta \leq 1$. It is assumed that the modulating function b(t) is approximately the same for all channels observing the same pulse. However there may be some de-correlation and this is represented by the factor η. That is

$$\langle b_1'(t + \tau) \cdot b_1''(t) \rangle = \eta \rho_\mu(\tau) \langle b_1^2 \rangle$$

where the primes distinguish (adjacent) channels. The AMN model as formulated requires $\eta=1$, and in that case the CCF is the same as the ACF except for the zero-lag spike. If $\eta<1$ then the micropulse feature is reduced so that the CCF does not peak to 0.5, or $(1+\alpha)^{-1}$, at $\tau=0$. Thus

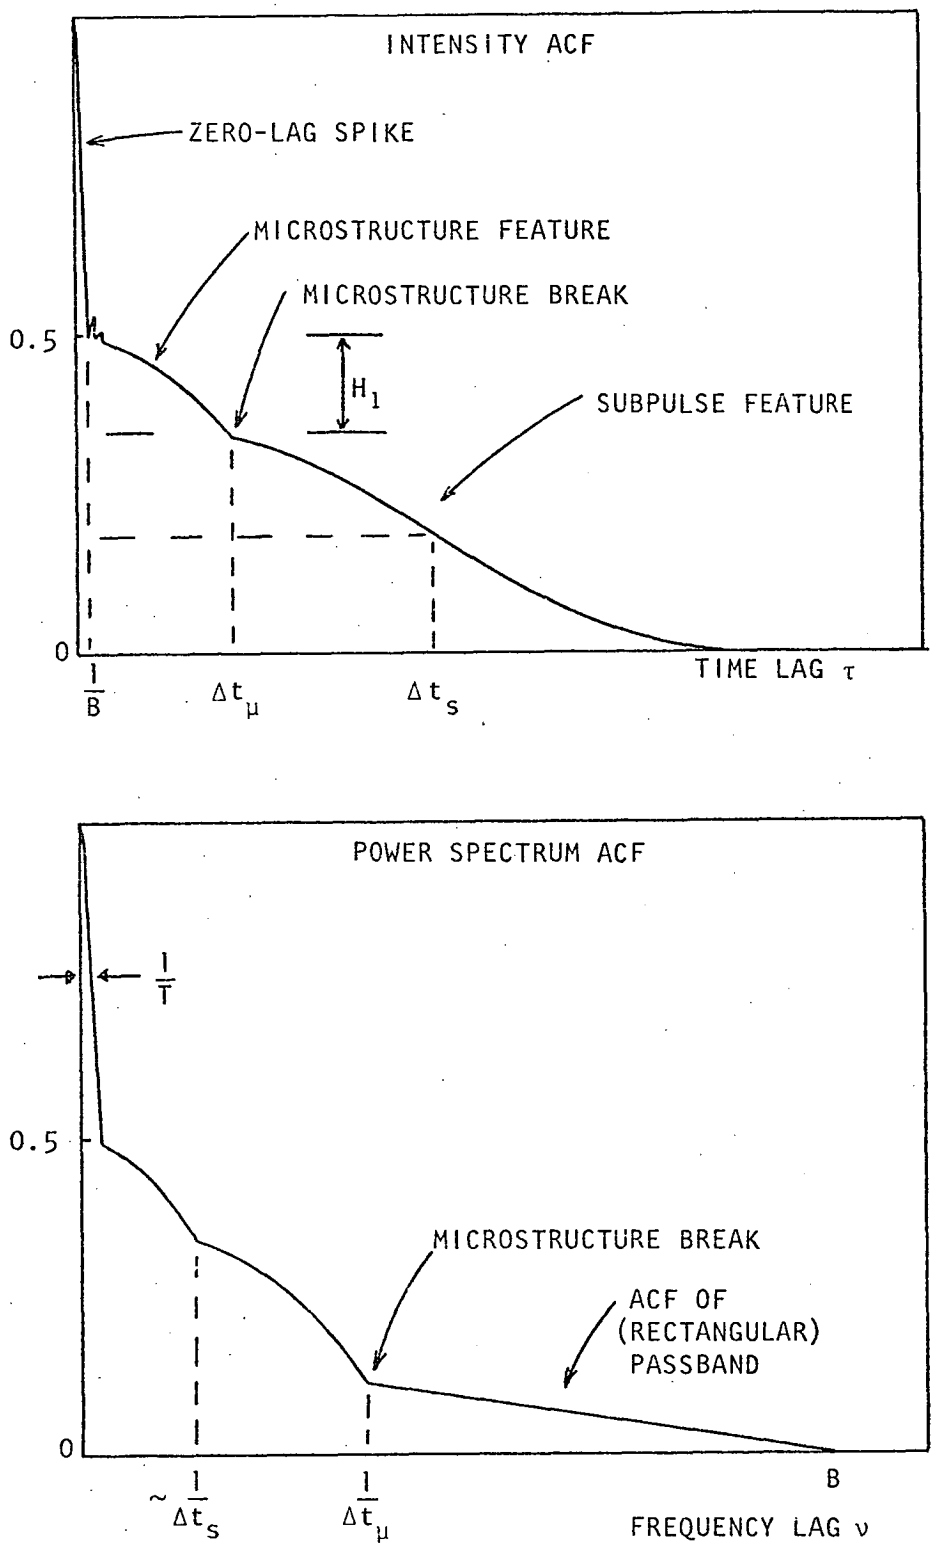

FIGURE 9.2.1A The model intensity ACF and the corresponding power spectrum ACF predicted by the AMN model. The channel bandwidth is B and the time window duration T. Sidelobes of the zero-lag spike of $r_p(\nu)$ have been omitted (see text).

the value of η is easily tested in practice. The ACF of a pulse is estimated by averaging across channels. The average CCF between pairs of these channels at some spacing (e.g. adjacent) is calculated. Then the resultant correlation functions are plotted and compared. The null hypothesis is $\eta=1$ or that the differences in the curves can be explained by statistical fluctuations and measurement errors. These correlation functions can also be averaged between pulses.

By using two channels and averaging 250 strong pulses from PSR 0950+08 Rickett, Hankins and Cordes (1975) obtained $\eta = 0.5 \pm 0.2$ for a frequency separation of 200 MHz (111 to 318 MHz) while for PSR 1133+16 Ferguson and Seiradakis (1978) obtained $\eta \sim 0.8$ for a frequency separation of 100 MHz near 2.7 GHz. For the Molonglo data frequency separations up to 1.8 MHz are available and it is likely that $\eta=1$ for all separations in this range.

As well as providing a measurement of the de-correlation, η, in the modulating function, b(t), the cross-correlation between channels gives a way of distinguishing two effects which would increase the height of the ACF zero-lag spike. These are the modulation index of the noise exceeding unity and the presence of unresolved variations in b(t). If the latter, and on its own, i.e. $\alpha=1$, then the CCF will have a zero-lag spike which rises to 0.5 Alternatively if there are a significant number of large impulses in the noise then the CCF will have a spike whose base is at 0.5

A model for the auto-correlation, $r_p(\upsilon)$, of the power spectrum, P(f), has been outlined by Rickett (1975). Define in a similar way to (4) and (5),

$$C_p(\upsilon) = \frac{1}{B'} \int_0^{B'} P(f)P(f+\upsilon)df$$

and $$r_p(\upsilon) = C_p(\upsilon)/C_p(0)$$

where $P(f) = |U'(f)|^2$ with $U'(f)$ the single-sideband video spectrum,

and $B' > B \sim 2$ MHz. It is assumed that the spectrum $U'(f)$ is well-resolved, that is, is calculated from a time window, [0..T], which does not truncate the pulse. Under these two conditions (i.e. on B' and T) a Fourier relationship holds between $r_I(\tau)$ and $r_p(\upsilon)$.

$r_p(\upsilon)$ is sketched in Figure 9.2.1A. It has a zero-lag spike of height 0.5 (assuming $\alpha=1$) and width 1/T. If the power spectrum is calculated with a DFT without any zero-fill then it is tabulated at intervals 1/T and side-lobes of the zero-lag spike do not appear because the points fall upon zeroes of a sinc^2 function. Next as υ increases comes a subpulse feature that may merge into the zero-lag spike unless $\Delta t_s \ll T$. The microstructure feature follows and the ACF curve has a microstructure break at $\upsilon \simeq 1/\Delta t_\mu$. This is where the microstructure feature meets the broadest part of the ACF which is the auto-correlation of the receiver passband. Note that the frequency of this microstructure break agrees with the intuitive idea that the spectrum ought to be modulated on a scale which is the reciprocal of the scale of modulation in the time domain.

9.2.2 Background Correction

Throughout the previous Section it was assumed that the off-pulse intensity was negligible. A procedure for correcting estimates of intensity correlation functions when this is not the case was given by Rickett (1975). It is as follows.

The mean off-pulse intensity, v_{off}, is estimated as accurately as possible. This is subtracted from the on-pulse intensity giving $I_{on}(t)$ and from the off-pulse intensity giving $I_{off}(t)$. Note that $\langle I_{off}(t)\rangle=0$. Both intensity functions are assumed to be sampled in windows of length T at a rate > 2B, where B is the channel bandwidth. The off-pulse ACV estimate is $\langle C_{I_{off}}(\tau)\rangle_q$ where the subscript q denotes averaging over channels and/or pulses. A multiple of this is subtracted from the on-pulse ACV estimate, $\langle C_{I_{on}}(\tau)\rangle_q$, giving an estimate for $\langle C_I(\tau)\rangle$ namely

$$\langle C_I(\tau)\rangle_q = \langle C_{I_{on}}(\tau)\rangle_q - \langle C_{I_{off}}(\tau)\rangle_q\,[1+2\xi] \qquad (8)$$

where ξ is the signal to off-pulse noise ration (S/N) evaluated in the observing window

$$\xi = \langle \int_0^T I_{on}(t)dt\rangle_q / v_{off}T \qquad (9)$$

Now the off-pulse ACV consists of a zero-lag spike only, i.e.

$$\langle C_{I_{off}}(\tau)\rangle = \langle I_{off}^2(t)\rangle\ \mathrm{sinc}^2\ B\tau \qquad (10)$$

so that it suffices to evaluate the second moment of the off-pulse intensity. However I have always passed off-pulse data through the same set of numerical processes as on-pulse data as a simple means of keeping track of scaling factors and for estimating the level of statistical fluctuations and artefacts in the analysis.

A method for background correction of power spectrum ACFs has been given by Rickett (private communication). The entity corresponding to v_{off} is the mean off-pulse power spectrum. When a rectangular passband is used the computations are simpler because the background is constant within the passband, because of the passband flattening, and zero elsewhere. The details are as follows: Equation (9) becomes

$$C_p(v) = \int_{f_1}^{f_2} P(f+v)P(f)df/(f_2-f_1) \qquad (11)$$

where, as usual, the integral notation is for convenience, and $f_1 \sim 0.2$ MHz and $f_2 \sim 2$ MHz. The passband is flat between f_1 and f_2, that is

$$\langle P_{(off)}(f) \rangle = p_{off} \quad \text{for } f_1 < f < f_2.$$

p_{off} is subtracted from the raw power spectra giving $P_{off}(f)$ and $P_{on}(f)$ and the evaluation of $\langle C_p(v) \rangle_q$ and hence of an estimate of $r_p(v)$ proceeds exactly as in the intensity case.

The disadvantage of this "flat background" method is that the rectangular passband truncates the power spectrum of the pulsar. Whereas the time window could be chosen to avoid cutting off the pulse, there is no such freedom in the frequency domain because the spectrum is expected to be fairly uniform. (This is always the case but there is less disruption if the passband has rounded edges.) However I have used all the data for each pulse in a single power spectrum ACF calculation so that there are only two sharp edges imposed on the spectrum (at frequencies f_1 and f_2), while the 1.8 MHz of spectrum contains ~350 individual modulation fringes. This last calculation assumes that $\Delta t_\mu \simeq 200$ μs so that the fringe spacing is $\sim 1/\Delta t_\mu = 5$ kHz. Therefore the amount of artefact introduction into the ACF is small. In fact it is smaller than in the case of the intensity ACF

where each of the many channels has a rectangular passband (which results in a sinc^2 function for the zero-lag spike).

The main practical problem I have found in background correction is related to finding accurate values for v_{off} and p_{off} which are appropriate to each pulse. The general uncertainty in background levels was discussed in Chapter 7. To estimate the resultant systematic errors in the intercept values suppose that the on-pulse system gain is actually lower than the off-pulse gain so that the background, v_{off}, is an over-estimate by an amount εv_{off}. Then to first order in ε, the intercept value will be biased, i.e.

$$r_I(0+) = (1 - 2\alpha\varepsilon\xi/\Gamma)/(1 + \alpha) \qquad (12)$$

instead of being an unbiased estimate of $1/(1 + \alpha)$. Here

$$\Gamma = \langle C_I(0)\rangle_q / v_{off}^2 \quad \text{or} \quad \langle C_p(0)\rangle_q / p_{off}^2$$

where $\langle C_I(0)\rangle_q$ or $\langle C_p(0)\rangle_q$ are calculated according to (8). For $\alpha=1$ equation (12) becomes

$$r_I(0+) = 0.5 - \varepsilon\xi/\Gamma \qquad (13)$$

Note that Γ is the second moment relative to v_{off} and so is $\sim \xi^2$. Therefore the error in the intercept is approximately inversely proportional to the signal-to-noise ratio.

9.3 CORRELATION ANALYSES FOR PSR 0950+08

9.3.1 Intensity Correlations

The intensity ACF was evaluated for eight pulses from PSR 0950+08. In each case twelve or thirteen 140 kHz channels were used. The auto-covariance (ACV) was evaluated for each and the cross-covariance (CCV) was evaluated between adjacent channels using the well-known FFT method (Bergland 1969). For each pulse the covariances were averaged and all the correlation functions were plotted.

For the group of six pulses excluding C and D, the correlation functions were evaluated to a lag of 0.72 ms (i.e. 200 points) using a time window of 3.0 ms (824 points) or 5.5 ms depending on the pulse width. Background was not removed so that for small S/N each curve has a dominant linear trend from 0.5 at zero lag to zero at ±3.0 ms or ±5.5 ms. The curves for off-pulse data are shown in Figure 9.3.1A while those for some of the pulses and the average for the six pulses are shown in Figures 9.3.1B to 9.3.1E. The results are summarized in Table 1.

It is not expected that the average correlation functions for the six pulses should be very alike since each characterizes the modulating function, a(t), for that particular pulse only. However for none of them is there a strong microstructure feature which can be clearly separated from the statistical fluctuations. In no case is the CCF curve significantly different from the corresponding ACF curve and no CCF curve contains a significant spike at zero-lag. Furthermore in no case is the intercept, $r_I(0+)$, significantly different from 0.5 which agrees with the AMN model but the test is insensitive because of the inclusion of the background.

The following are comments on individual pulses. The break in the ACF for pulse A at 65 μs apparently reflects the fact that the central peak of this pulse has a rectangular profile 70 μs wide. The corresponding

CAPTION FOR FIGURES 9.3.1A-E

Each Figure shows the autocorrelation function, ACF, averaged over 12 or 13 contiguous rectangular channels of bandwidth 140 μs, and the cross-correlation, CCF, between adjacent channels. The maximum lag is 720 μs i.e. 200 points. No background subtraction has been done so that the zero-lags intercept is expected to be 0.5 (marked with a cross for the CCF). The zero-lag spike has been cut off each ACF. In each case a straight line has been drawn over the ACF curve and then copied with reflection onto the CCF curve. Vertical displacements of the CCF from these lines can be explained by statistical fluctuations in the normalization.
Figure 9.3.1A is off-pulse data while Figure 9.3.1E shows the ACF and CCF averaged over six pulses.

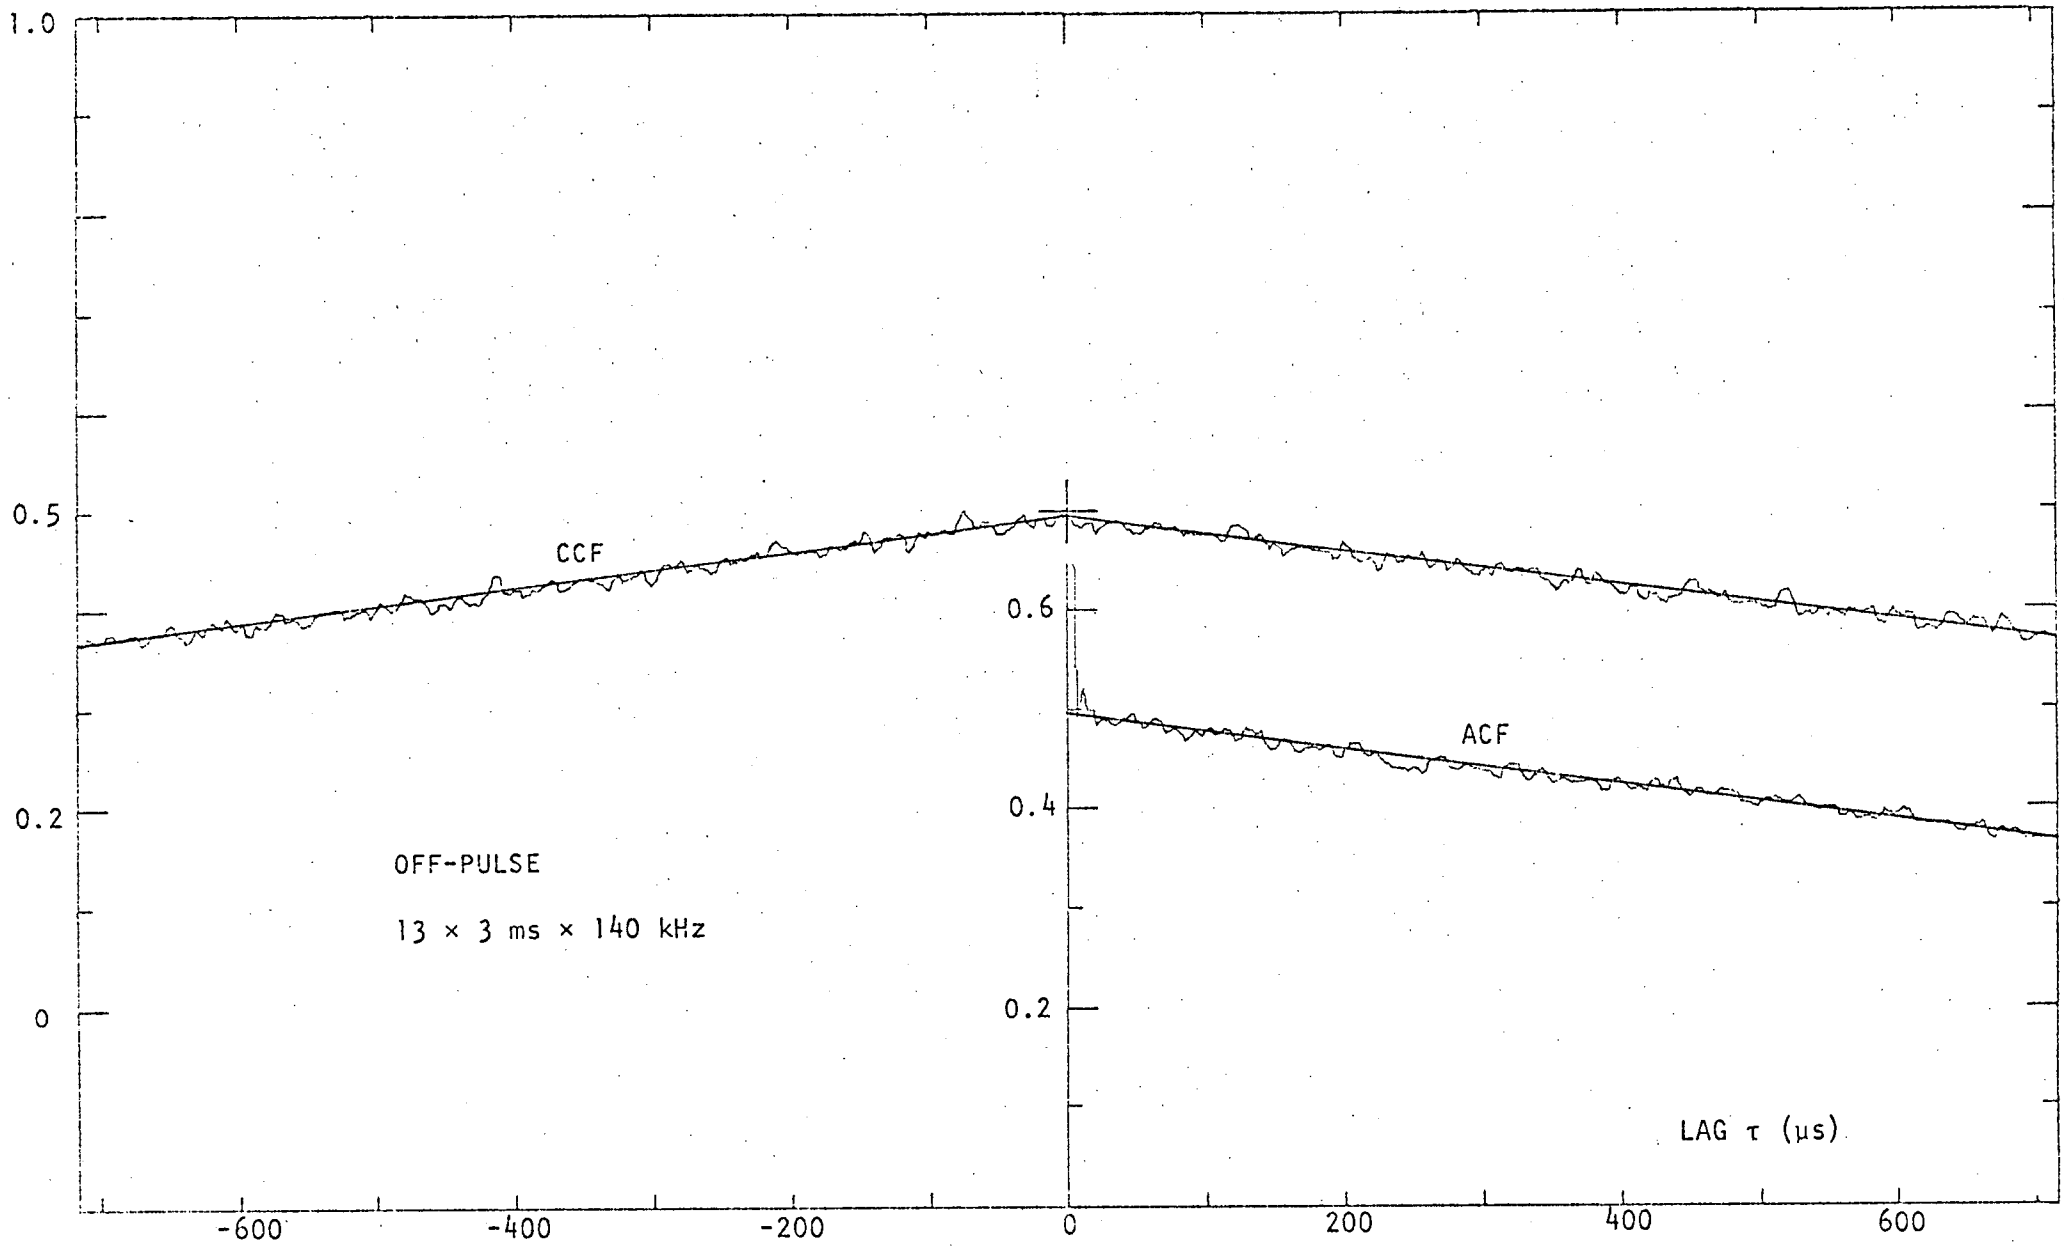
1.0
0.5
0.2
0
CCF
OFF-PULSE
13 × 3 ms × 140 kHz
-600
-400
-200
0
0.6
0.4
0.2
ACF
200
400
600
LAG τ (μs)

FIGURE 9.3.1A

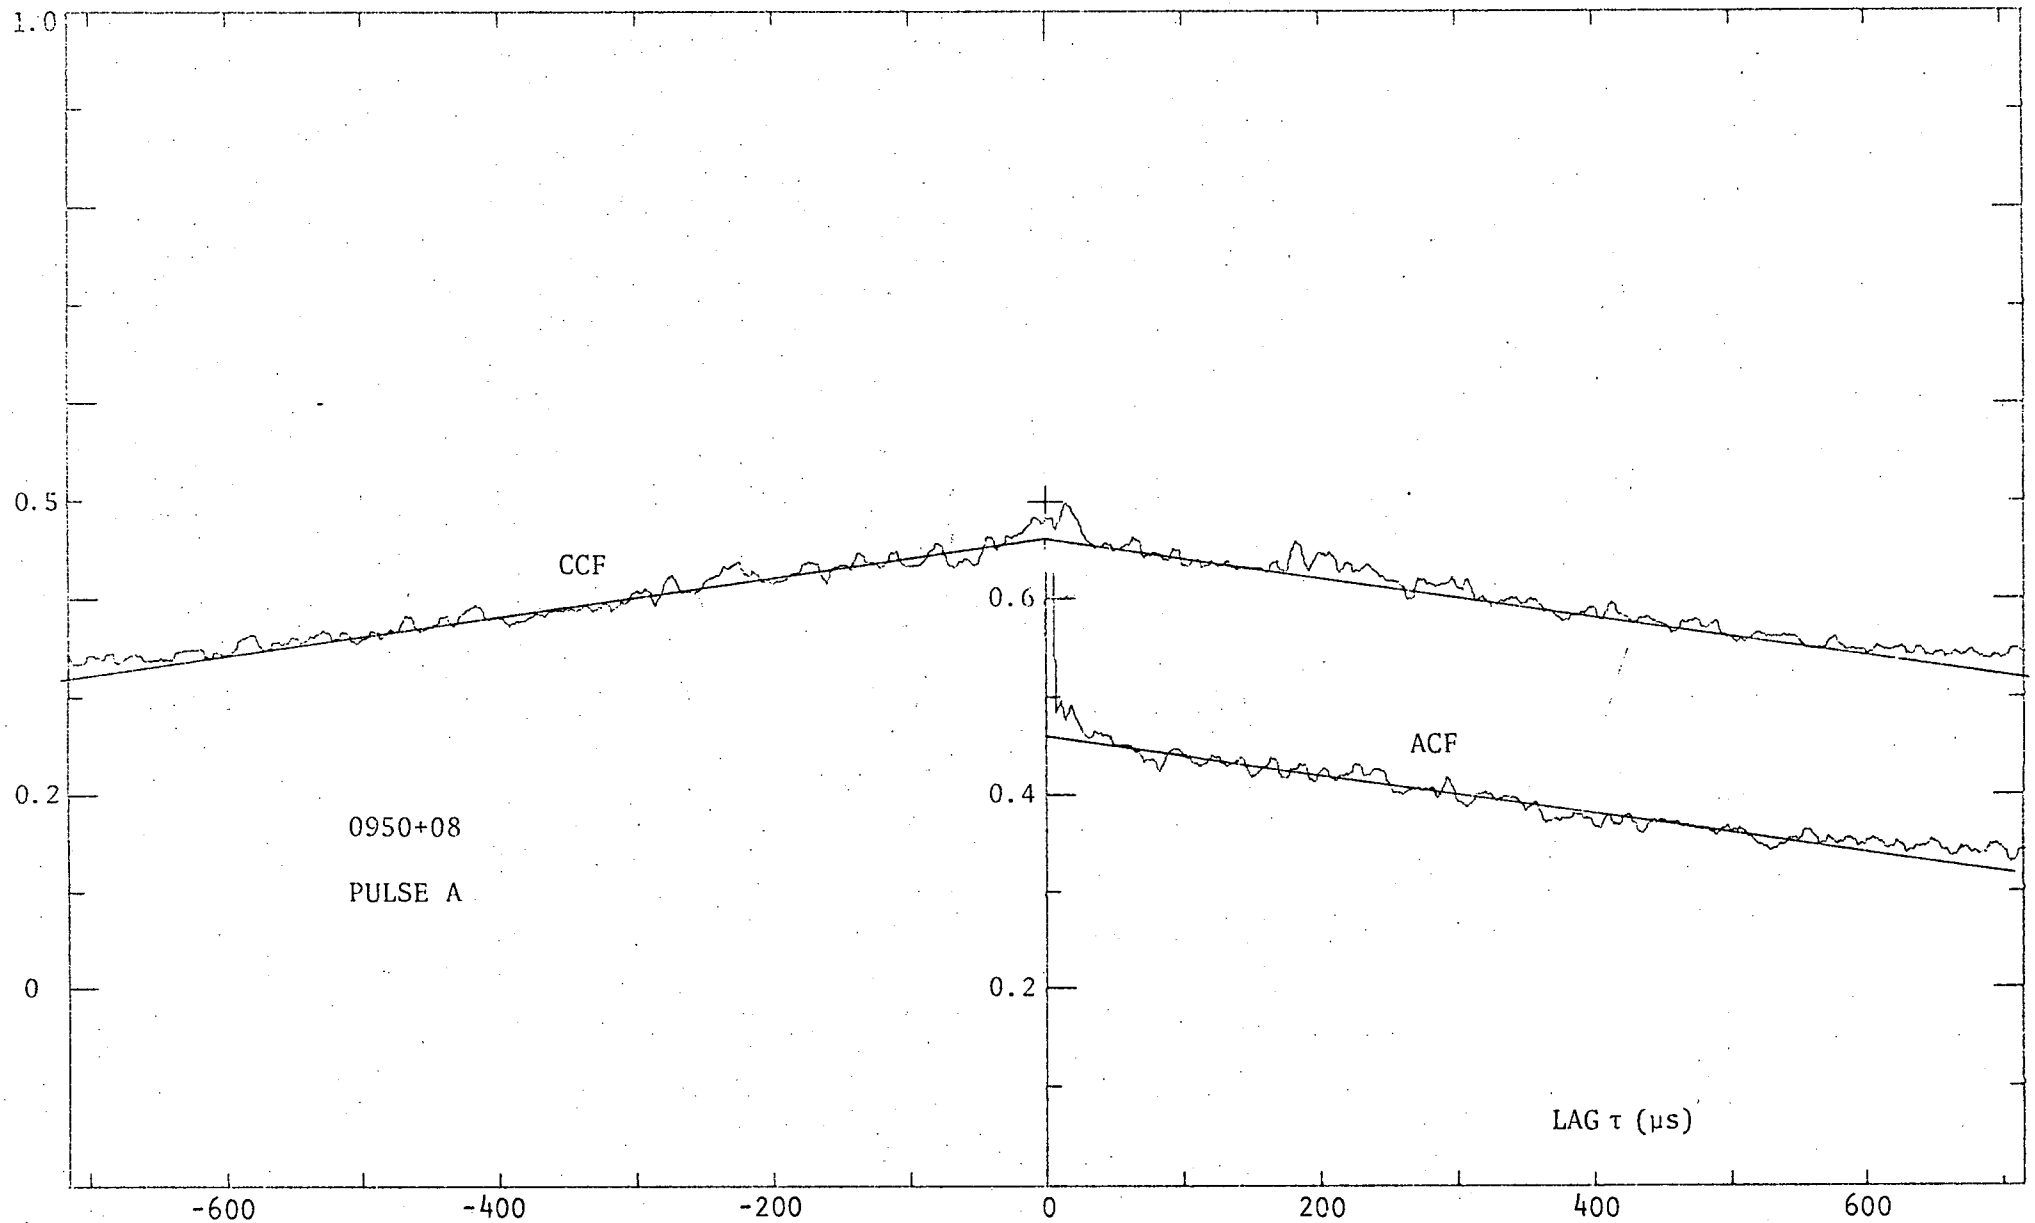
1.0
0.5
0.2
0
CCF
0950+08
PULSE A
0.6
0.4
0.2
ACF
LAG τ (μs)
-600
-400
-200
0
200
400
600

FIGURE 9.3.1B

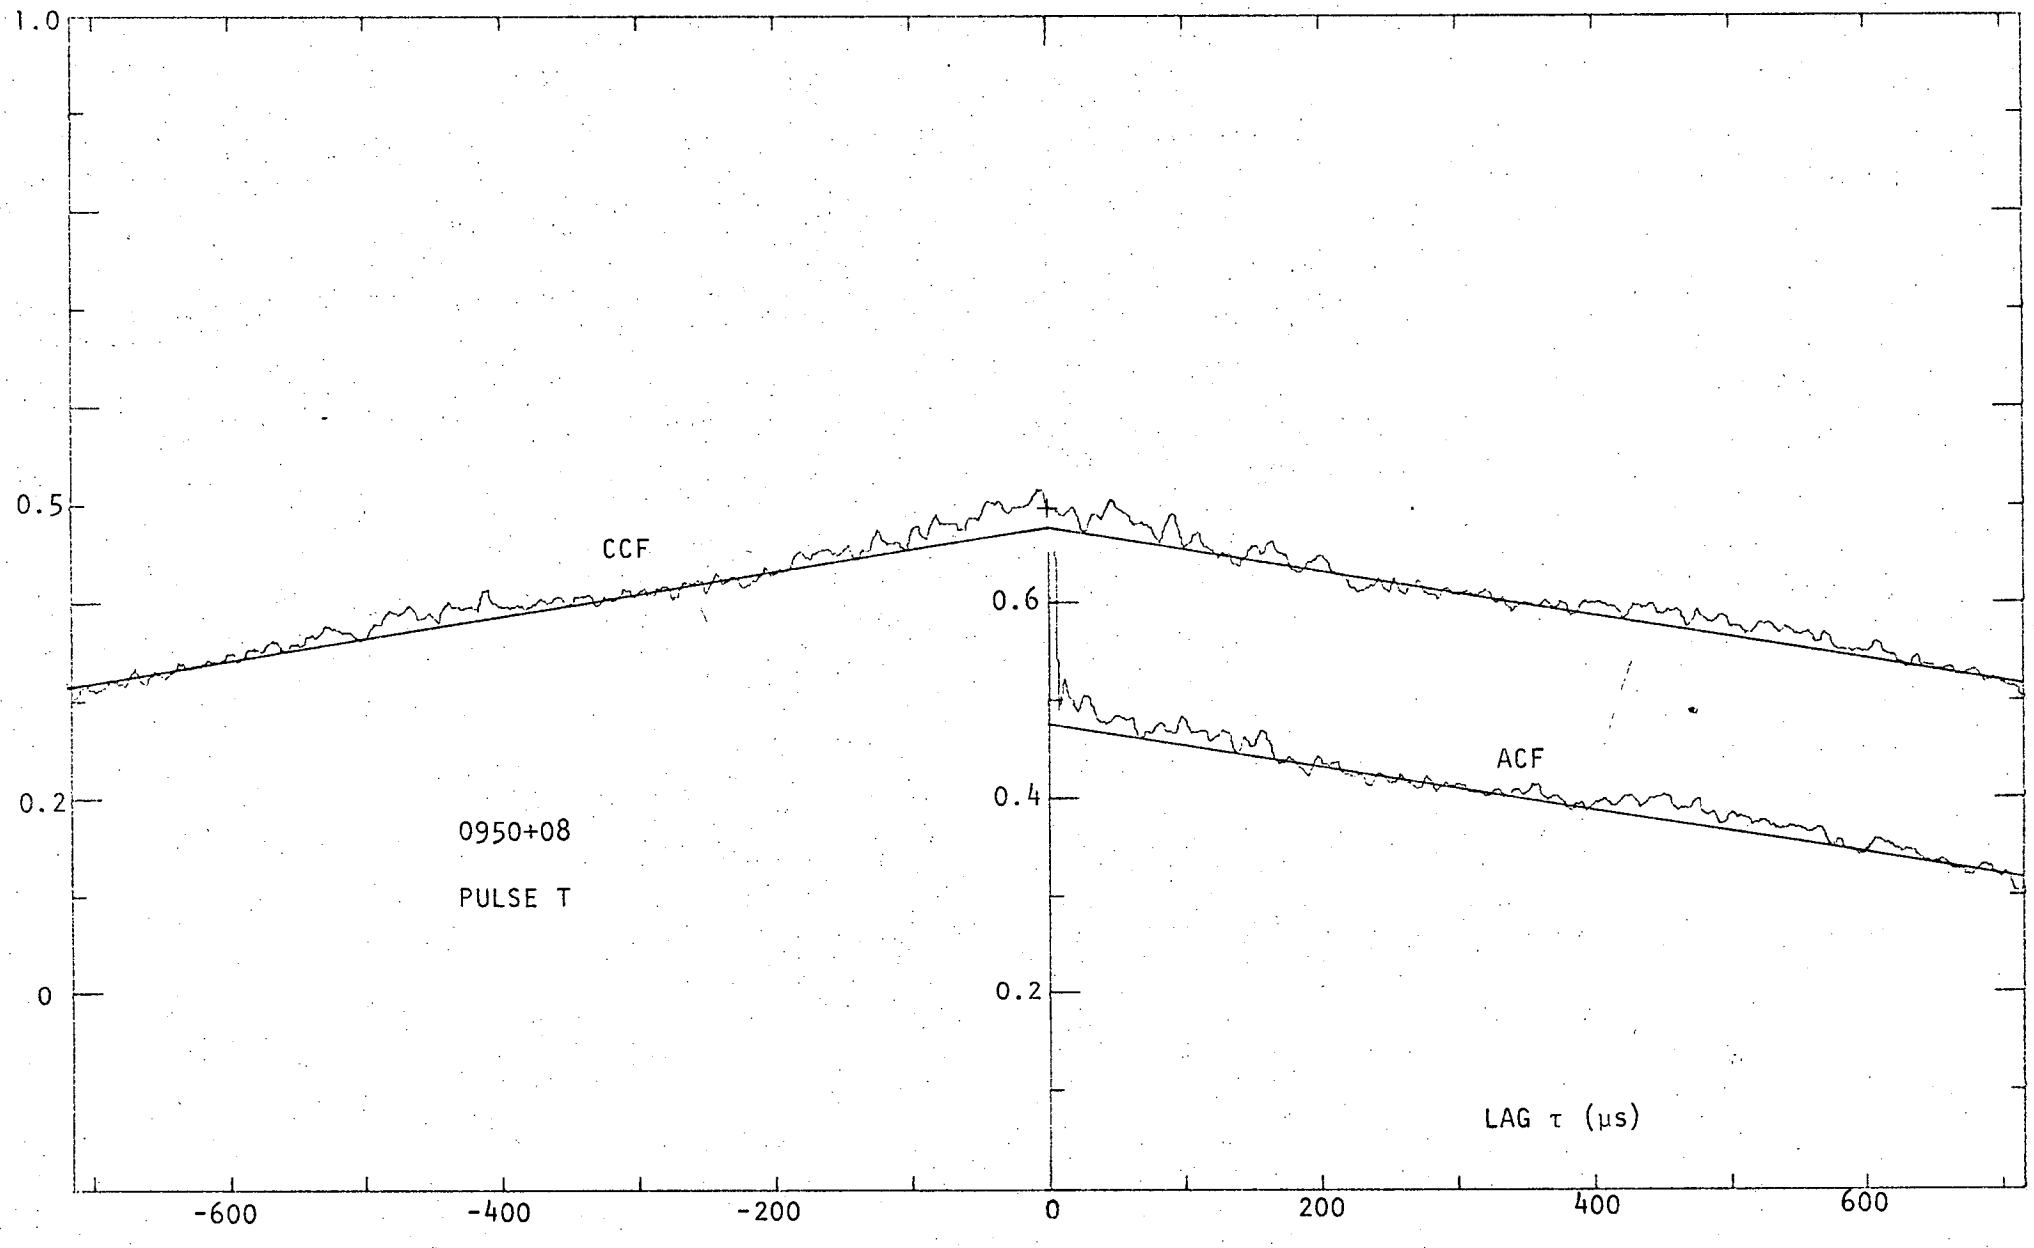
1.0
0.5
0.2
0
CCF
0950+08
PULSE T
0.6
0.4
0.2
ACF
LAG τ (μs)
-600
-400
-200
0
200
400
600

FIGURE 9.3.1C

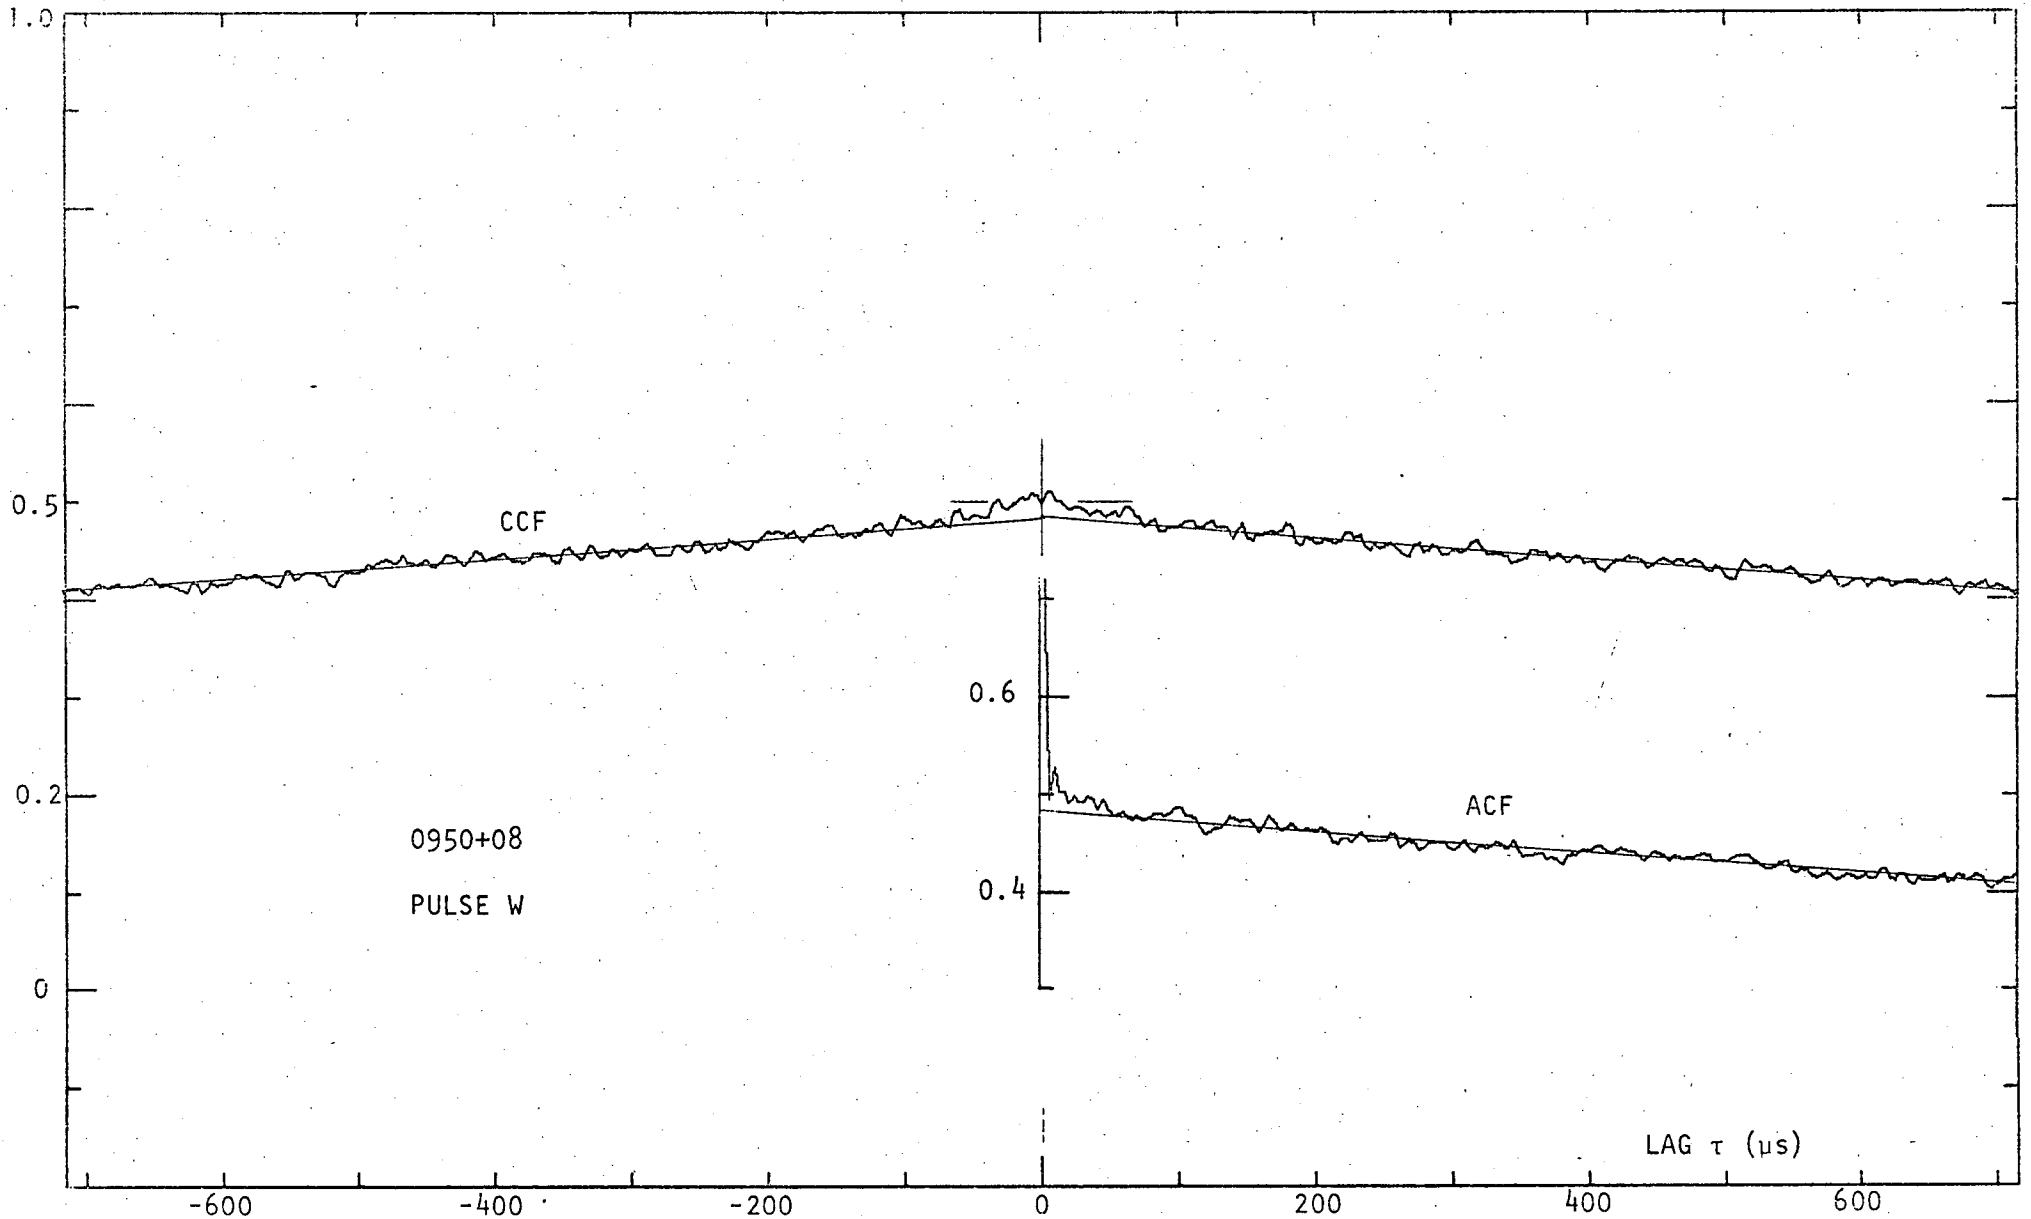
1.0
0.5
0.2
0
CCF
0950+08
PULSE W
0.6
0.4
ACF
LAG τ (μs)
-600
-400
-200
0
200
400
600

FIGURE 9.3.1D

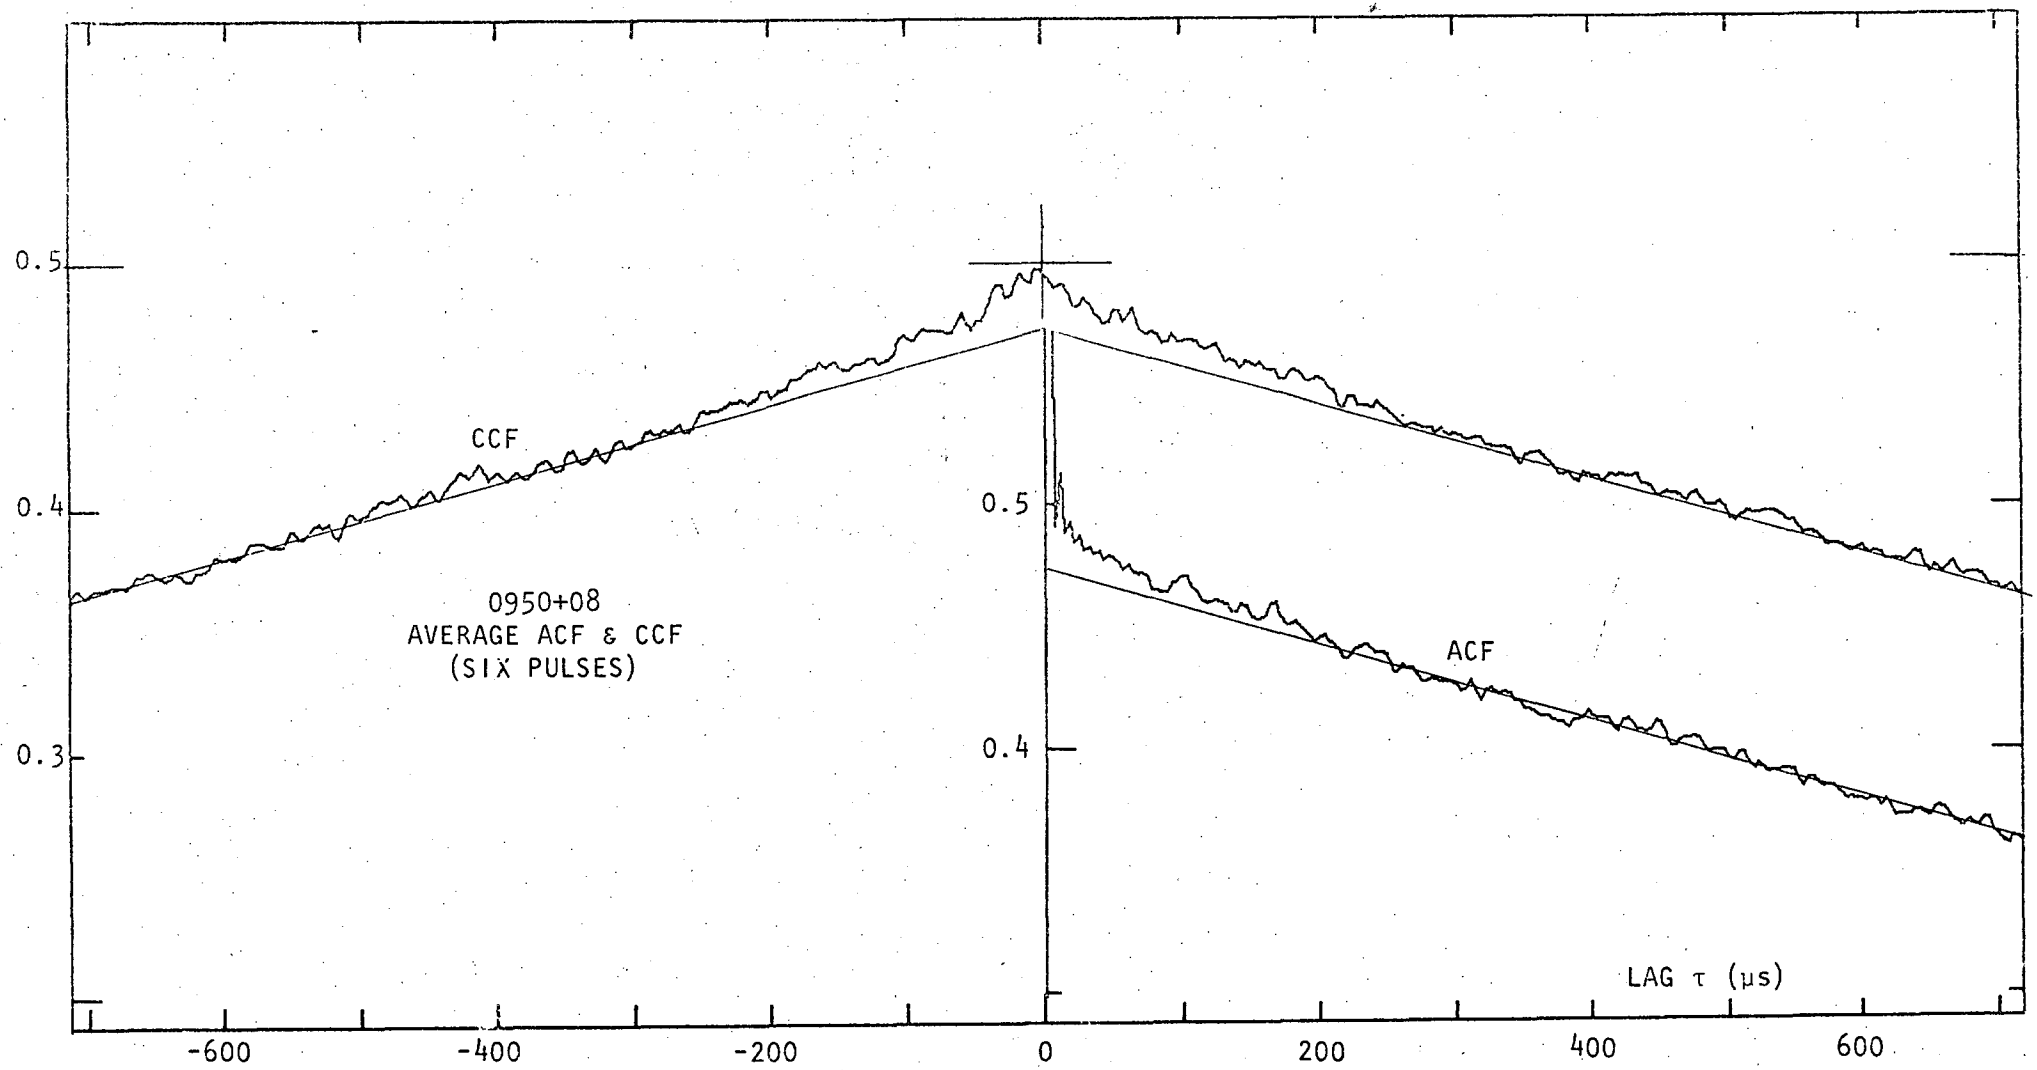
0.5
0.4
0.3
CCF
0950+08
AVERAGE ACF & CCF
(SIX PULSES)
0.5
0.4
ACF
LAG τ (μs)
-600
-400
-200
0
200
400
600

FIGURE 9.3.1E

Pulse	S/N ξ	intercept $r_I(0+)$	Figure	comments
A	~1		B	break at 65 ± 20 μs
H	1			weak break at 0.3 ± .1 ms
K	~1			straight line
P	0.7			straight line
T	0.9		C	break at 170 ± 30 μs (see text)
W	2.5	0.505 ± .02	D	weak break at 70 ± 20 μs
sum		0.494 ± .005	E,F	structure extending to 250 μs
OFF	0	0.495 ± .005	A	
C	1.8	0.49 ± .08		straight line, background corrected
D	1.05	0.52 ± .07		straight line, background corrected

TABLE 1 : Summary of PSR 0950+08 intensity ACFs. Maximum lag is 720 μs. Background subtraction was only performed for pulses C and D and so only the intercepts for these two are sensitive tests of the AMN model. However the errors quoted for these contain a nominal ±0.05 to account for possible systematic errors in the background correction.

CCF "cap" is narrower than expected. Pulse T has a microstructure feature extending to about 0.17 ms and an enhancement near 0.5 ms. This enhancement which indicates weak periodicity corresponds to the spacing of the three broad micropulses that make up pulse T (Figure 7.2C). Similarly the ACF of pulse W has an enhancement near 1.3 ms (not shown in Figure 9.3.1D) corresponding to the spacing of the two main parts of this pulse. Another analysis of pulse W in which dispersion was not removed produced a CCF in which the "cap" has shifted from zero-lag to be centred at 60 ± 10 μs. This agrees well with the expected shift, namely the bandwidth sweep time of 51 μs.

The correlation functions for the six pulses were averaged, each being weighted very approximately by the S/N of the pulse. (Had the ACVs been averaged as implied by the definition (5) of an ACF then pulse W would have been dominant. This is because an ACV is roughly proportional to intensity squared.) The mean ACF shows a weak microstructure feature extending to 250 ± 50 μs. In spite of the small number of pulses this is consistent with the values Δt_μ = 175 μs obtained by Hankins (1972) and Δt_μ = 200 μs obtained by Kardashev *et al* (1978). The CCF also has a cap about 50 μs wide. The corresponding ACF feature is visible under the side-lobes of the zero-lag spike (Figure 9.3.1F).

Intensity ACF and CCF calculations were also performed upon pulses C and D. Windows of 4.8 ms were used and the background was subtracted according to equation (8). The resulting ACF intercepts can therefore be treated as more sensitive tests of the AMN model than those where the subtraction is not done. Neither intercept is significantly different from 0.5 whether account is taken of possible errors in the background correction or not. Nor does either set of curves show any microstructure feature above the "noise"; they are well-fitted by straight lines.

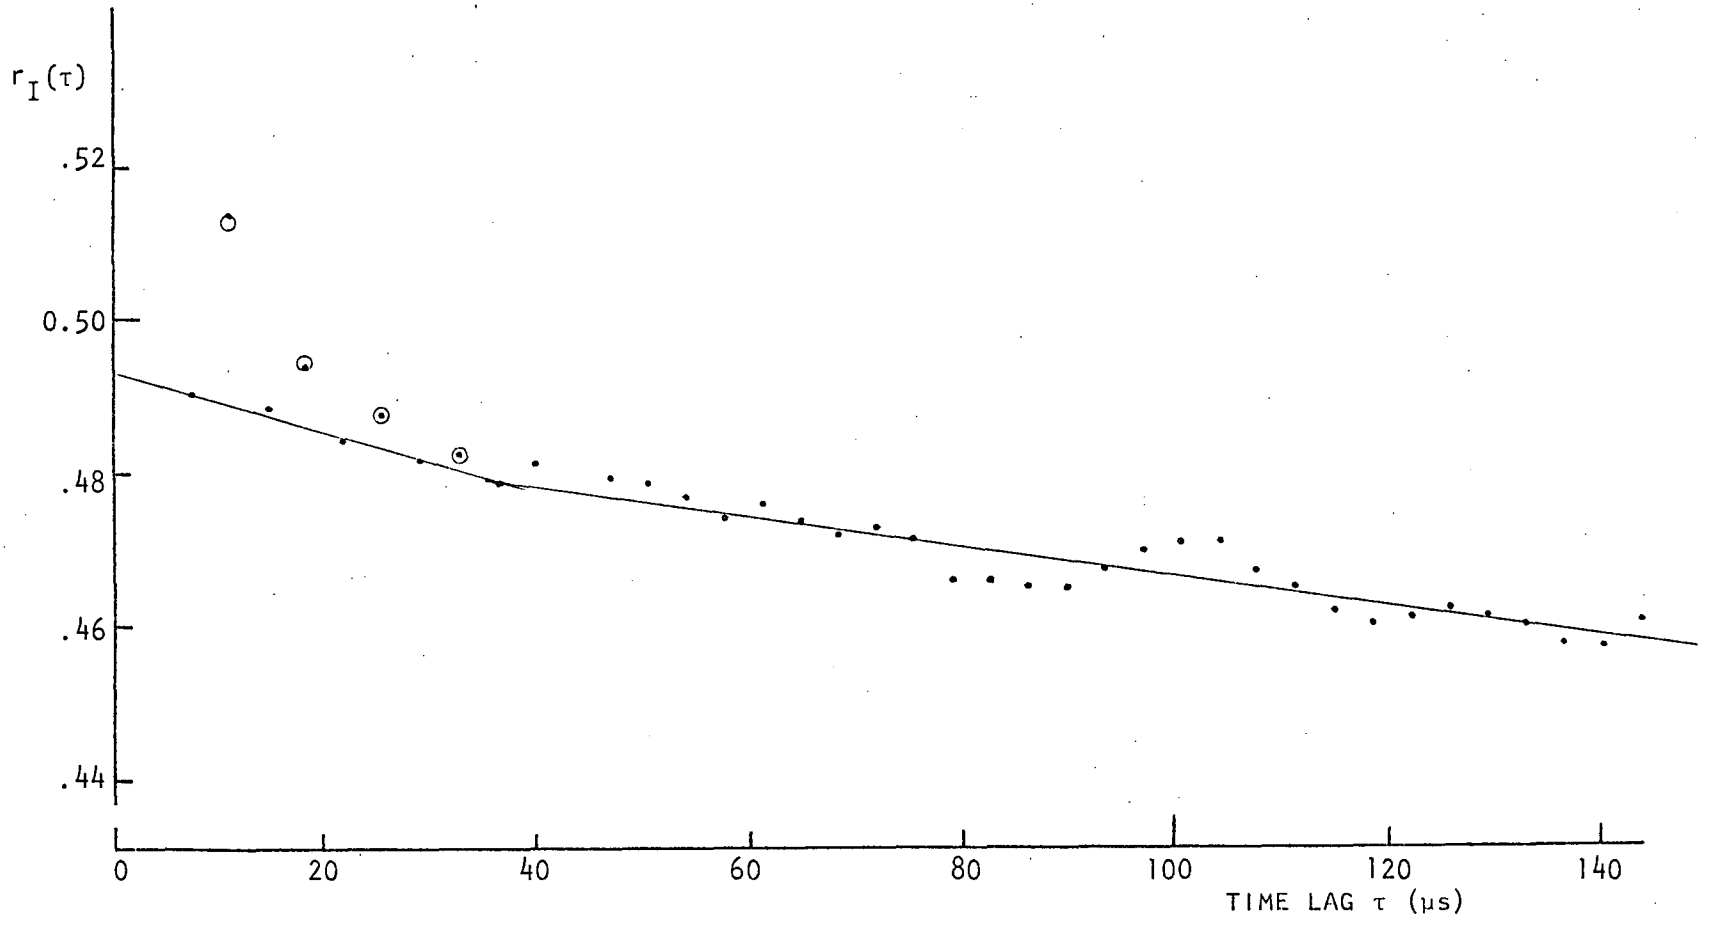

FIGURE 9.3.1F . The first 40 points of the mean ACF shown in Figure 9.3.1E. Every second point is statistically independent. The circles indicate the sinc^2 side-lobe values expected with the straight line fit to $r_I(0+) = 0.492$. Note that the off-pulse background has not been subtracted.

9.3.2 Power Spectrum Auto-correlations

ACF analyses of the power spectra of five pulses from PSR 0950+08 were performed. In each case the discrete power spectrum was calculated from 7 ms of data (31K samples plus 1024 zeroes) and the passband was flattened. Each ACF was calculated by the FFT method to a maximum lag of 140 kHz (1024 points) using a frequency window of 1.8 MHz (13K points). Background subtraction was applied to four of the pulses according to the "flat passband" method. Some additional analyses were performed on pulse T. The results of all these ACF analyses are summarized in Table 2.

All of the ACF curves are well-fitted by straight lines and the fluctuations about each line are similar to analyses of off-pulse data. After background correction each intercept, $r_p(0+)$, agrees with the prediction of the AMN model, which is $r_p(0+) = 0.5$. However the result for pulse D requires account to be taken of errors in the background correction. The corrected intercept is $0.46 \pm .02$ where the quoted standard error is based on the fluctuations in the curve. From equation (13) with $\xi = 0.65$ and $\Gamma = 0.96$ the bias in the intercept which would be caused by a fractional over-estimate, ε, in the background level is -0.7ε. Therefore a value $\varepsilon = 0.03$ would suffice to reconcile the intercept with the AMN prediction and the uncertainty in the background of pulse D is of this order.

Pulse	time window (ms)	signal-to-noise ξ	corrected intercept $r_p(0+)$	comments
P	7.0	~0.7	-	fits straight line
T	3.7	0.68	.52 ± .02	fits straight line
T	7.0	0.38	.49 ± .13	fits straight line
W	7.0	1.20	.506 ± .05	fits straight line
C	7.0	1.25	.503 ± .02	fits straight line
D	7.0	0.65	.46 ± .02	fits straight line

TABLE 2 : Summary of power spectrum ACF analyses for PSR 0950+08. In each case the frequency window was 1.8 MHz. Quoted errors for $r_p(0+)$ do not include possible systematic errors in the background correction which are about ±0.05.

9.4 CORRELATION ANALYSES FOR PSR 1749-28

9.4.1 Intensity Correlations

The intensity ACF was evaluated for seven pulses from PSR 1749-28. The pulses were treated in two groups : pulses A, B, C and pulses F, H, J and K. Because of the strong scintillation fringes in the observed spectrum of this pulsar, I selected data from 70 kHz channels whose centre frequencies fell on the fringe maxima. Twenty-seven intensity sequences (traces) were chosen from the first group of pulses and sixteen from the second. For each trace a window of 5.9 ms (824 points) was centred on the observed emission. Background subtraction was performed and this was based on a total of five off-pulse stripes. The ACF of each trace was calculated out to a lag of 1.43 ms (200 points) and these ACFs were averaged within each group of pulses. For the first group of pulses the averaging was weighted by the strength of the pulse; ACVs were averaged in the second group.

For both groups of pulses the mean ACF curve is approximately fitted by a straight line. The curve for the second group is shown in Figure 9.4.1A. Individually the two ACF curves appear to vary randomly about the fitted straight line. However when smoothed to a resolution of about 0.1 ms the residuals show the same pattern. In both curves there is a dip between 0.1 and 0.3 ms and a broad hump between 0.4 and 0.7 ms. Since the two sets of pulses are disjoint the similarity suggests that these features are real. In particular, the hump suggests periodicity in the pulsar emission with periods in the range 0.55 ± .15 ms. Both features are consistent with the power spectrum ACF of the second group of pulses, discussed in the next Section.

Background correction was applied to the mean ACF for the second group of pulses according to equation (8). The corrected intercept is $r_I(0+) = 0.52 \pm .025$ where the error contains a component ±0.015 representing the effect of a nominal 5% uncertainty in the background. This was

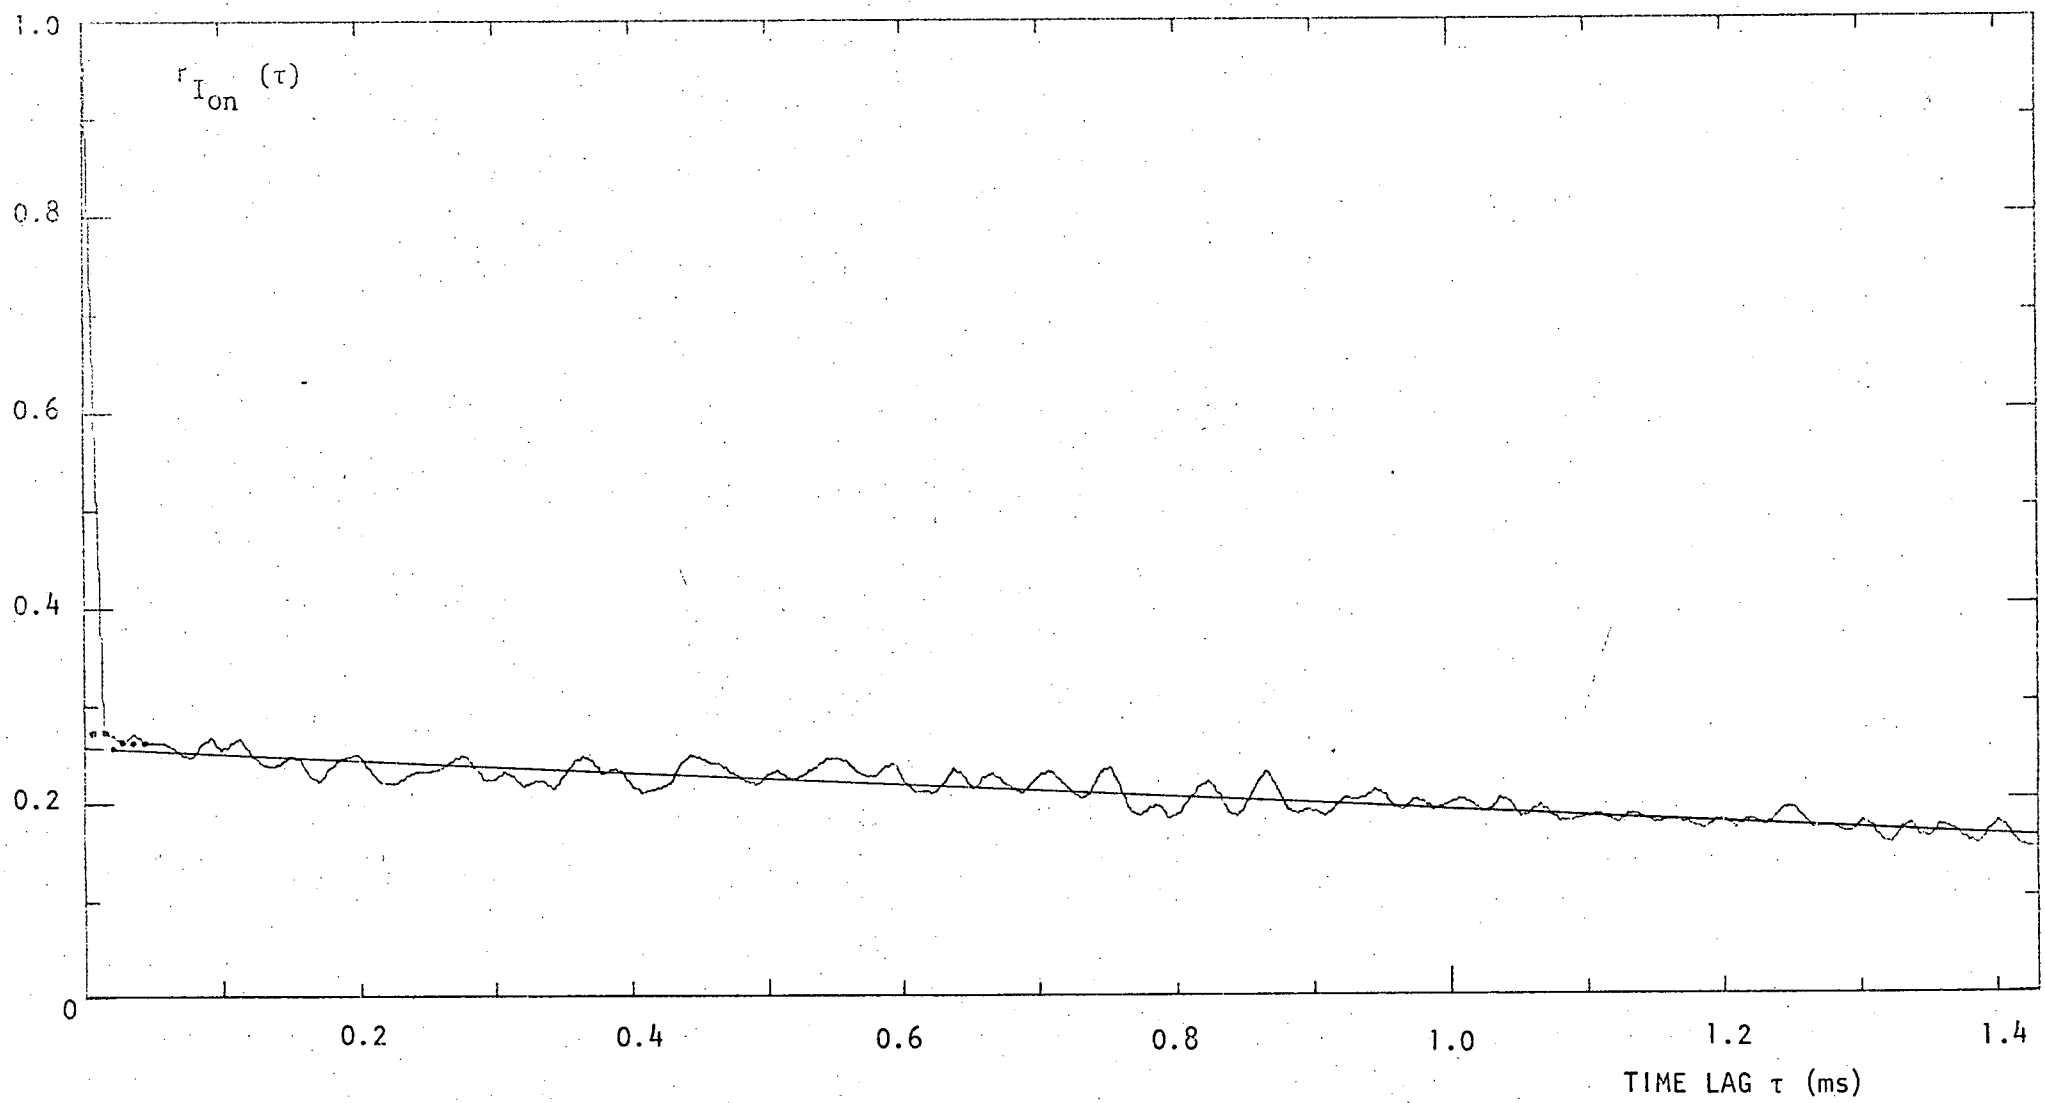

FIGURE 9.4.1A Intensity ACF averaged over four pulses (F, H, J, K) from PSR 1749-28 at 408 MHz. For each pulse four channels were used each with bandwidth 70 kHz and centred on the strongest scintillation fringes in the pulsar spectrum. The time window was 5.9 ms; the time resolution is 14.3 μs. The dots show the ACF values after subtracting the zero-lag spike with its sidelobes. The plot shows the ACF before background correction.

calculated using (13) with $\Gamma = 2.8$ and mean S/N ratio $\xi = 0.85$.

Other ACF analyses were performed on the second group of pulses. These used time windows of double the duration described above and all of the spectrum including the scintillation troughs. As would therefore be expected the S/N ratio was reduced by a factor ~4 and no new information about the pulsar was gained.

9.4.2 Power Spectrum Auto-correlations

ACF analyses of the power spectra of eight pulses from PSR 1749-28 were performed. Because of the high dispersion of this pulsar the spectra were calculated from time windows of 14 ms (62K samples, as a matter of course in the dispersion removal). Since after dispersion removal all of the pulses are less than 7 ms wide, the S/N ratios of the spectra are about half what could have been achieved by a much more expensive analysis (two more 64K FFTs per pulse). In each case the ACF was calculated to a maximum lag of 70 kHz (1024 points) using a frequency window of 1.8 MHz (26K points). Background correction was applied to all of the ACFs using the "flat passband" method and this was based upon four stripes of off-pulse data. The ACVs of seven pulses were averaged; the eighth, that of pulse D, was not included because its spectrum covers only 0.9 MHz.

All of the ACF curves show breaks near 3 kHz or 30 kHz. The details are given in Table 3 while Figure 9.4.2A shows the mean ACF.

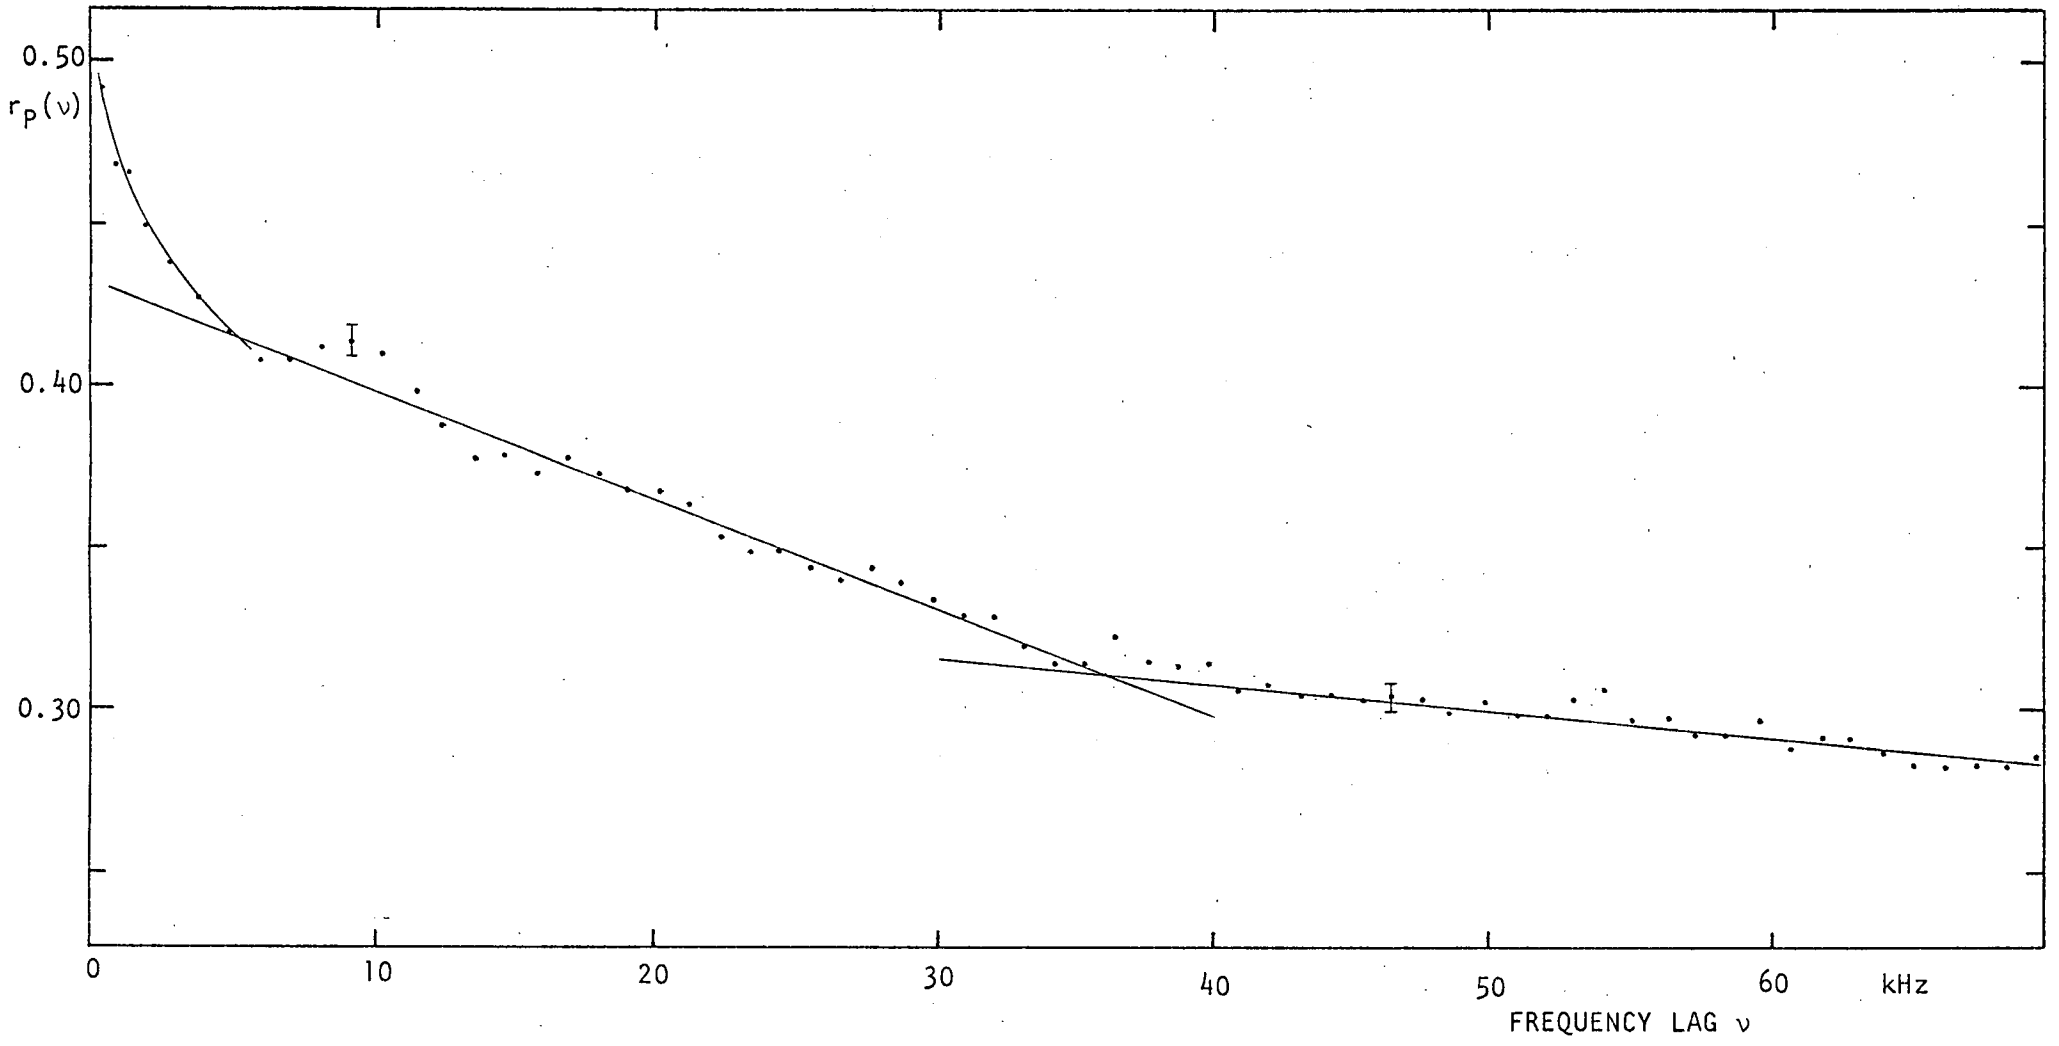

FIGURE 9.4.2A ACF of power spectrum averaged over seven pulses from PSR 0950+08 at 408 MHz. Background correction has been performed. Each point is independent and typical 2σ error bars are shown. The fitted straight line and curve are discussed in the text.

pulse	S/N (T = 14 ms)	Breakpoints (kHz)
A	0.41	3.5 , 27
B	.20	3.3 , noisy
C	.45	2.6 , 40
F	.39	exponential to 33
H	.17	5 only
J	.15	exponential to 25
K	.25	23
sum	0.29	3.5-5, 35
D	~0.3	3.3

TABLE 3 : Summary of power spectrum ACFs for PSR 1749-28

Before discussing the shape of the mean ACF curve I will deal briefly with the intercept, $r_p(0+)$. This is 0.50 ± .055 where the uncertainty is made up of ±0.015 due to the observed fluctuations in the curve near lag v=0, and ±0.04 representing the effect of a nominal 4% uncertainty in the background level (using equation (13) with $\Gamma \simeq \xi \simeq 0.29$). Whether the possible background error is included or not the observed intercept agrees with the AMN model.

The mean ACF curve has a well-defined break at 35 kHz which is explained by interstellar scintillation. The decorrelation frequency for this pulsar at 408 MHz was estimated as 16 kHz (Table 3 of Section 6.4.2), which is good agreement with a breakpoint at 35 kHz for a particular pattern of scintillation fringes. This pattern was shown in Figure 7.3D. Since it has a quasi-period ~ 160 kHz, 35 kHz is a reasonable value for the half-width of the average fringe.

The exact frequency of the break near 4 kHz is less well defined although the feature is clear. There is an approximately exponential curve meeting at 5 kHz the straight line that extends to 35 kHz. Note that there is no excess correlation at small lags due to sidelobes of the zero-lag spike (confirmed by the off-pulse power spectrum ACFs) so that the 4 kHz feature is not an artefact. The corresponding timescale is 250 ± 50 μs and I identify this as the microstructure timescale, Δt_μ, of PSR 1749-28.

There is also a significant peak in the ACF curve at 9.5 ± 1.5 kHz. This indicates periodicity in the time domain with period P_μ = 105 ± 15 μs. There are some peaks in the corresponding intensity ACF curves with quasi-periods in this range but the signal-to-noise ratio of these curves is insufficient for a definite confirmation of the periodicity.

These first measurements of Δt_μ and microstructure period, P_μ, for PSR 1749-28 are shown in Table 1 of Section 1.3.4 with published data for 12 other pulsars. The value for Δt_μ fits well the linear relationship between Δt_μ and pulsar period P found in the Arecibo results (Cordes 1978b). However the period, P_μ = 105 μs, is shorter than any other published value.

CHAPTER 10

CONCLUSION

10. INTRODUCTION

The videotape system for dispersion removal has been proven by the observation of four pulsars. The signals from two of these were sufficiently strong for some conclusions to be drawn. The observations of PSR 0950+08 at time resolutions down to 0.5 μs generally agree with those of other workers. PSR 1749-28 had not been observed before with microsecond time resolution. The identification of a 250 μs characteristic timescale for this pulsar needs confirmation by further observations preferably with a better signal-to-noise ratio. The AMN model was assumed in obtaining this result from ACFs of power spectra.

No significant deviations from the AMN model have been found, either in the qualitative analyses of dynamic spectra, or in the running mean analyses of intensity series, or by measurement of the modulation index of the noise through auto- and cross-correlations. The power of the latter test was reduced by uncertainties in the off-pulse background which relates to the performance of the videotape system. The observations of Hankins and Boriakoff (1978) apparently did not suffer such uncertainties and they reported agreement of their ACFs with the AMN model. This was with a 0.8μ s resolution. Is it worth-while attempting to increase the bandwidth, say by a factor of two, for measurements of the modulation index? Apparently not since there are no reports of deviations of marginal statistical significance (e.g. 1 σ) at the present resolution. Therefore the bandwidth and time resolution would need changing by a larger factor. However interstellar scattering seriously restricts the best time resolution that can be obtained at frequencies below 1 GHz for most pulsars. Since this smearing effect is approximately proportional to f^{-4}, better time resolution may require modest increases in observing frequency, f. However for two pulsars, PSR 1919+21 and PSR 2016+28, the fraction of the emission that is microstructure decreases with frequency.

I suggest that more simultaneous observations of microstructure at widely-spaced frequencies below 1 GHz be made. Time resolutions of 1μs should be sufficient. These observations may lead to some understanding of the factors which decrease the correlation of the microstructure between different observing frequencies. Also the time lag of best cross-correlation as a function of pulsar longitude may suggest the location of the emission. For example, is a single micropulse emitted at various points (at different frequencies) along a curved field line?

Finally I shall discuss briefly a development of the videotape system in which the recording medium is replaced by a video-disk. Such a system is presently being developed by Drs. P.A. Hamilton and P.M. McCulloch at the University of Tasmania. The advantages include: an increase in bandwidth (by a factor of two); greater stability of system gain and noise (as the vararies of tape tension, especially during stop-action replay, are replaced by a well-controlled head-to-disk separation); the ability to record only within a time window (equal to the disk rotation period of 40 ms) centred on the integrated pulse profile. This last feature has obvious advantages but may not be so useful for simultaneous observations at different frequencies where the dispersion delay is more than a few milliseconds. Such a situation would require at least a dual-head video-disk instead of the less expensive solution of recording the two data channels by splitting the video band with a sub-carrier related to the reference tone. The videotape system records continuously and so there is no difficulty with time windows. (This assumes that the segmenting of the data into stripes is countered by using two recorders as described in Section 7.6.) Furthermore the storage capacity of a video-disk is much less than that of a reel of videotape, even with the inclusion of the off-pulse data. Therefore an observing session using a videodisk must be interrupted periodically to digitize the data while the digitizing of a videotape can be deferred until convenient.

REFERENCES

Backer,D.C. 1973 Astrophys.J. 182,245

---. 1974 Astrophys.J. 190,667

Backus,P.R. and J.M.Cordes 1978 Bull.A.A.S. 10,446

Bartel,N. 1978 Astron.Astrophys. 62,393

Benford,G. 1977 Mon.Not.R.astr.Soc. 179,311

Benford,G. and R.Buschauer 1977 Mon.Not.R.astr.Soc. 179,189

Bergland,G.D. 1969 IEEE Spectrum 6,41 (reprinted in "Digital Signal Processing", ed. Rabiner and Rader, IEEE Press)

Boriakoff,V. 1976 Astrophys.J.(Letters) 208,L43

Bracewell,R.N. 1965 "The Fourier Transform and its Applications" McGraw-Hill, New York

Buschauer,R. and G.Benford 1976 Mon.Not.R.astr.Soc. 177,109

Chen,H-H.,M.A.Ruderman and P.G.Sutherland 1974 Astrophys.J. 191,473

Cheng,A.F. and M.A.Ruderman 1977a Astrophys.J. 212,800

---. 1977b Astrophys.J. 214,598

---. 1979 Astrophys.J. 229,348

Cheng,A.F.,M.A.Ruderman P.G.Sutherland 1976 Astrophys.J. 203,209

Cooley,J.W. and J.W.Tukey 1965 Math.Comput. 19,297 (reprinted in "Digital Signal Processing", ed. Rabiner and Rader, IEEE Press)

Cooley,J.W.,P.A.Lewis and P.D.Welch 1969 IEEE Trans.Audio.Electroac. June, (reprinted op cit)

---. 1970 J.Sound.Vib. 12,315 (reprinted op cit)

Cordes,J.M. 1975 Astrophys.J. 195,193

---. 1976a Astrophys.J. 208,944

---. 1976b Astrophys.J. 210,780

---. 1978a Astrophys.J. 222,1006

---. 1978b paper presented at the "Symposium on Pulsars and Supernova Remnants", Sydney, April 1978, and published

in Aust.J.Phys. 32,9 (March 1979)

---. 1978c paper presented at the "JPL/NASA Workshop on Magnetospheres", August 1978

Cordes,J.M. and T.H.Hankins 1973 Bull.A.A.S. 5,18

---. 1977 Astrophys.J. 218,484

Cordes,J.M.,J.Rankin and D.C.Backer 1978 Astrophys.J. 223,961

Eastlund,B.J. 1968 Nature 220,1293

Ellis,G.R.A. 1973 Nature 241,387

Elsässer,K. and J.Kirk 1976 Astron.Astrophys. 52,449

Ferguson,D.C. 1976a Astrophys.J. 205,247

---. 1976b circular letter dated Oct 24, 1976

Ferguson,D.C.,D.A.Graham,B.B.Jones,J.H.Seiradakis and R.Wielebinski 1976 Nature 260,25

Ferguson,D.C. and J.H.Seiradakis 1978 Astron.Astrophys. 64,27

Ferguson,D.C. and R.Wielebinski 1977 New Scientist (9 June 1977),586

Flowers,E.G.,J-F.Lee,M.A.Ruderman,P.G.Sutherland,W.Hillebrandt and E.Muller 1977 Astrophys.J. 215,291

Ginzburg,V.L. and V.V.Zheleznyakov 1975 Ann.Rev.Astron.Astrophys. 13,511

Ginzburg,V.L.,V.V.Zheleznyakov and V.V.Zaitsev 1969 Astrophys.Space Sci. 4,464

Goldreich,P. and W.H.Julian 1969 Astrophys.J. 157,869

Hamilton,P.A.,P.M.McCulloch,J.G.Ables and M.M.Komesaroff 1977 Mon.Not.R.astr.Soc. 180,1

Hankins,T.H. 1971 Astrophys.J. 169,487

---. 1972 Astrophys.J.(Letters) 177,L11

---. 1973 Astrophys.J.(Letters) 181,L49

---. 1974 Astron.Astrophys.Supp. 15,363

Hankins,T.H. and V.Boriakoff 1977 Bull.A.A.S. 9,433

---. 1978 Nature 276 (2 Nov 1978),45

Hankins,T.H. and B.J.Rickett 1975 Methods Comp.Phys. 14,56

Hankins,T.H.,V.Boriakoff,J.M.Cordes and B.J.Rickett
1977 Bull.A.A.S. 9,288

Hansell,G.E. 1969 "Filter Design and Evaluation", van Nostrand

Hardee,P.E. and W.K.Rose 1976 Astrophys.J. 210,533

Henricksen,R.N. and J.A.Norton 1975 Astrophys.J. 201,719

Henricksen,R.N. and D.R.Rayburn 1974 Mon.Not.R.astr.Soc. 166,409

Hillebrandt,W. and W.Müller 1976 Astrophys.J. 207,589

Hinata,S. 1979 Astrophys.J. 227,275

Hinata,S. and E.A.Jackson 1974 Astrophys.J. 192,703

Jackson,J.D. 1975 "Classical Electrodynamics", 2nd edn.,Wiley, New York

Kardashev,N.S.,A.D.Kuz'min,N.Ya.Nikolaev,A.Yu.Novikov,M.V.Popov,T.V.Smirnova,
V.A.Soglasnov,T.V.Shabanova,M.D.Shinskii and Yu.P.Shitov
1978 Sov.Astron. 22,583

Kawamura,K. and I.Suzuki 1977 Astrophys.J. 217,832

Kirk,J.G. and D.terHaar 1978 Astron.Astrophys. 66,359

Klauder,J.R.,A.C.Price,S.Darlington and W.J.Albershiem
1960 Bell System Tech.J. 39,745

Komesaroff,M.M.,P.A.Hamilton and J.G.Ables 1972 Aust.J.Phys. 25,759

Kraus,J.D. 1966 "Radio Astronomy", McGraw Hill, New York

Large,M.I. and R.H.Frater 1969 Proc.IREE Aust. 30,227

Manchester,R.N. and J.H.Taylor 1972 Astrophys.Lett. 10,67

---. 1977 "Pulsars", Freeman, San Francisco

Manchester,R.N.,E.Tademaru,J.H.Taylor and G.R.Huguenin
1973 Astrophys.J. 185,951

Michel,F.C. 1973 Astrophys.J. 180,207

Mills,B.Y.,R.E.Aitchison,A.G.Little and W.B.McAdam
1963 Proc.IRE.Aust. 24,156

Okamoto,I. 1974 Mon.Not.R.astr.Soc. 167,457

Oster,L. and W.Sieber 1976 Astrophys.J. 210,220

---. 1978 Astron.Astrophys. 65,179

Rickett,B.J. 1975 Astrophys.J. 197,185

---. 1975b Bull.A.A.S. 7,455

Rickett,B.J.,T.H.Hankins and J.M.Cordes 1975 Astrophys.J. 201,425

Ruderman,M.A. 1972 in "IAU Symposium 53. The Physics of Dense Matter", ed. Tsuruta

---. 1976 Astrophys.J. 203,206

Ruderman,M.A. and P.G.Sutherland 1975 Astrophys.J. 196,51

Smith,F.G. 1970 Mon.Not.R.astr.Soc. 149,1

---. 1977 "Pulsars", Cambridge University Press

Stone,J.M. 1960 "Radiation and Optics", McGraw Hill, New York

Sturrock,P.A. 1971 Astrophys.J. 164,529

Sutton,J.M. 1968 Aust.J.Phys. 21,221

---. 1971 Mon.Not.R.astr.Soc. 155,51

Taylor,J.H. 1974 Astron.Astrophy.Suppl.Ser. 15,367

Taylor,J.H. and R.N.Manchester 1975 Astron.J. 80,794

Taylor,J.H.,R.N.Manchester and G.R.Huguenin 1975 Astrophys.J. 195,513

Udal'tsov,V.A. and V.N.Zlobin 1974 Astron.Astrophys. 37,21

Zlobin,V.N. and V.A.Udal'tsov 1976 Sov.Astron. 19,683

Zheleznyakov,V.V. and V.E.Shaposhnikov 1972 Astrophys.Space Sci. 18,141

Zverev,A.I. 1967 "Handbook of Filter Synthesis", Wiley, New York

APPENDIX A : LIST OF COMPUTER PROGRAMS

The following is a brief summary of the main computer programs written to implement the videotape dispersion removal system.

(i) PDP11 programs in FORTRAN and assembler, with assistance from P. A. Hamilton :

VM2 -- performs stop-action digitizing with a validation algorithm that requires agreement of each sample on any two successive replays of the same stripe

VMORDR -- rearranges the samples into time sequence

VMEDIT -- interactive editing of the data

VMCOP2 -- transcription to 9-track tape

(ii) PDP8 program in assembler, using a files package implemented by P. S. Whitham :

VD5 -- stop-action digitizing with averaging over successive replays of a stripe

(iii) The following programs run on the Burroughs B6700 of the University of Tasmania Computing Centre. Unless otherwise stated each program is written in Burroughs Extended ALGOL. "Main-line" programs are marked with an asterisk.

COMP6 -- compares sets of samples between files created by VD5 which represent the same stripe. Needed to detect "gliches" in the sampling comb.

DPF -- batch detection and editing of dropouts

*LOAD -- sets up a disk file ready for dispersion removal, subtracts the mean, etc.

*TBU -- (600 lines, Trig-table generation, Big FFT, Unscramble) performs a discrete Fourier transform on large files of real data e.g. 64K

*SMP -- (480 lines) plots power spectra on log or linear scales

with smoothing over rectangular windows. Also writes 32 block-averages to a 'bpass' file.

PRINTUP -- forms weighted sums of 'bpass' files. For off-pulse data.

DIANA -- (FORTRAN) calculates cubic spline coefficients from a 'bpass' file for passband flattening

*PHASE -- corrects the complex spectrum for instrumental phase

*DYN -- (1100 lines) divides the spectrum into channels, flattens the passband, performs pre-detection dispersion removal by use of the FFT algorithm, writes the resultant intensity series to a 'tracefile' with house-keeping information in a 'traceindex', smooths and plots the traces, calculates an average across traces, writes this to an 'avfile' and plots it. There are many options including the ability to replot 'tracefiles' and 'avfiles'.

*ZPLOT -- displays the dynamic spectrum represented by a 'tracefile' (assuming many channels) as a blackness plot on a lineprinter

MELANIE -- hidden-line plotting of a dynamic spectrum

SLANT, PROFILE - calculate averages across traces, modulation indices, etc.

HIST/TRACE -- calculates histograms of relative intensity

*CORR -- (900 lines) calculate, plot and sum ACFs and CCFs of traces in a 'tracefile', and writes them to a 'dfile'. Many options

CORR/REPLOT -- replots a 'dfile'

CORR/SETUPSPECTRUM, SUMDFILES -- for use with CORR

TAS/PLOTPAK -- a library of plotting subroutines and utility programs. This has become the installation's plotting system.