SPCATS

(Sound Programme Circuit Automatic Test-Set)

A. J. GALLOWAY

Submitted in partial fulfillment of the requirements for the degree of Master of Science in the Department of Electrical and Electronic Engineering at the University of Cape Town, April 1987

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

University of Cape Town

ACKNOWLEDGEMENTS

My sincerest thanks and appreciation to Mr. J.R. Greene for supervising this thesis. His ideas and knowledge proved an inspiration throughout.

Thanks to Claude Jamotte for his suggestions and to all those that contributed to this work in some or another way.

C.S.I.R. F.R.D funding has been used towards this project.

For Tilla and Jacqueline

LIST OF FIGURES

LIST OF TABLES

CONTENTS

ABSTRACT

SPCATS has been designed and built, in prototype form to perform quality measurements on a national audio distribution network. The test-set is computer controlled and can be set to do tests at various times of the day. Tests have been implemented which are well known in the audio community making it easy to correlate results with the vast amount of literature on the subject.

The system was completed in two stages. An experimental test-set was initially built to check the overall concept and circuit ideas. With this being satisfactory implementation of the prototype was relatively straightforward.

The fundamental requirement of the system is that it performs tests on live channels. Therefore it is essential that the system be as non-intrusive as possible. The need to exercise the channel over its full frequency and dynamic ranges made the use of short tone-bursts as the test signals one of the few feasible approaches. Listening tests have shown these to cause minimal disturbance.

It was decided to use software techniques wherever possible due to their advantages over hardware methods. The transmitter section therefore consists of an insertion unit plus its computer controller. This unit is responsible for gating, scaling and generation of the test signals. The receiver section uses a data acquisition unit to feed sampled data to the computer where decimation filtering, windowing, a fast-fourier transform and various search algorithms have been implemented. Due to the high distortion purity required from the test system considerable attention had to be given to areas prone to various distortion mechanisms. This was particularly true of the analog-to-digital and digital-to-analog interfaces. The prototype system was built using most of the software and analog modules developed in the experimental test-set.

The system was tested on a broadcast quality audio channel made available for this purposed courtesy of the South African Broadcasting Corporation. The results obtained met the initial specification set and with the system proving extremely reliable.

CHAPTER 1

AN AUDIO TEST SET

1.1 INTRODUCTION

A prototype system has been developed to fill a gap in the wide variety of audio test equipment on the market. The majority of such equipment, to date, has been for doing 'bench' tests on small audio component systems, such as tape recorders and amplifiers [1,2]. Although there is a system on the market specifically designed for use on broadcast transmission installations, it is not able to quantify the different distortion types comprehensively [3]. The developed system does just this and can be used over any audio transmission link. It caters for tests that are implemented by the broadcast system mentioned above as well as the majority of tests carried out by bench-top audio testing equipment.

This system has been named 'SPCATS' (Sound Programme Circuit Automatic Test Set) and will be referred to as such in this text.

1.2 PRECEDENT

1.2.1 The Audio-dat

The Audio-dat [3] is a commercial system which non-intrusively determines characteristics of the channel it is monitoring. It uses the actual program material sent along the channel, comparing transmitted segments to the corresponding segments received. It uses this technique to determine level-faults (i.e. differences in level of programme transmitted and received), linear distortions (variations in frequency / amplitude response) and nonlinear distortions. The system cannot, however, characterise the non-linear distortions and simply indicates the gross magnitude of such distortion.

The fundamental concept of this system is that it relies on the probability that within a certain time, all possible

frequencies and amplitude levels are contained in the actual programme material sent [4].

To detect general level-faults, which includes degradation in signal-to-noise ratio and gross programme signal discrepancy due to switching errors, the system compares the signal r.m.s. values at the two ends.

The channel's frequency response causes the relative spectral content of the programme material at the two ends to vary. Subtracting these spectra gives the frequency response.

Finally, all forms of non-linear distortions are indicated by measuring the amplitude probability density functions at both ends and giving their difference as an indicator of non-linearity.

The problems with these measurement techniques are:

i) They rely on the programme having a white noise characteristic over time: One of the problems this leads to is in determining level errors. The r.m.s. reading is taken over approximately 200 milli-seconds and should the programme level be low at the time, the noise received with the program will suggest differences in level at the two ends when in fact it is only indicating a finite signal-to-noise ratio. Hence there is no guarantee that the full spectral and dynamic range of the circuit is being adequately exercised as no check is made on the amplitude of the programme material to be used for a test. In addition it is doubtful that the required accuracy for the new system could have been obtained with these techniques even if long averaging times had been used.
ii) Very little insight can be gained into the cause, or nature, of the channel defects from the results computed. The non-linear test result does not even give any idea as to whether the non-linearity is static or dynamic.

iii) It is necessary to relay information about the material transmitted to the receiver end. This is done by modulating a high frequency carrier (at 14.5 KHz) with the required information. The system is therefore not totally non-intrusive due to this carrier which has to be within the audio band due to the sharp roll-off of the channel above 15 KHz. The carrier amplitude is adjusted so that maximum reliability is maintained with a minimum of disturbance to listeners. This carrier is always present.

Emphasis has been placed on the Audio-dat system as it is in use at the South African Broadcasting Corporation. The new system will be used in conjunction with it. At present, however, the Audio-dat is only being used to monitor general level-faults. It has been in operation for over a year and from the results obtained it has been able to show linear and non-linear variations in channel gain [5]. Non-linear variations describe the dependence of gain on program level.

To avoid ambiguity SPCATS relies on commonly used, standardised tests which are well defined in the audio community and therefore meaningfull to the person analysing them.

1.2.2 Manual/Automatic Testing

Before the advent of computer-controlled instrumentation, testing of equipment was laborious and usually required a skilled person to do the actual measurement. By automating measuring procedures results have become more reliable, due to human error being removed, and the time to do tests considerably reduced, results being immediately available. The controllers for automated test instruments can either be hardware or software intensive. A software solution has the extremely attractive advantage of being easily modified as the need for changes and additions to the instrument arise [1]. However, further advantage can be gained by extending the controller to perform some of the functions implemented in hardware. These functions can then readily be changed. This

has therefore led to the modern trend of highly intensive software solutions with a minimum of analog modules [2].

For the above reasons the test system developed is software intensive with all processing of signals done by the computers. Conversion to the digital domain is done wherever feasible.

1.2.3 Use of Commercial Systems

The major difficulty with using a commercial test unit is that there are none known of that exactly satisfy the requirement. The Audio-dat has been mentioned, this being the closest system in concept, but it has has been pointed out above that it is limited in its range of tests as well as their interpretation.

An alternate approach is to use commercially available audio test equipment and build the required system from such modules. The major difficulties are:

i) Such modules are typically designed for bench-type testing.
ii) They are not designed to do analysis with non-intrusive signals.
iii) High quality audio equipment is extremely costly and this makes a specialist system built from more general modules unnecessarily costly.
iv) It is difficult to modify and add specific user requirements to a system that does not cater for this.

As an example, Tektronix, produce a programmable audio measurement system for testing high quality studio and other equipment [1]. The system satisfies the necessary requirement of high accuracy needed here but only contains a subset of the desired tests. The major problem with this system is that the signal generator and the analyser cannot easily be synchronised in time.

Practically, it would be possible to synchronise the Tektronix system and add the additional test tones needed, but this

would require considerable prototype hardware. For this reason it was considered far more viable, economically and practically, to design and produce a prototype system from circuit level.

A system, which is not commercially available, has been developed by the CBS Technology Center specifically to do audio quality tests on their thousands of tape machines. Even though they use the system to do audio quality tests, the fact that they wish to test tape machines has meant that they have adopted different tests for implementation to those of SPCATS. The point here is that although two different systems may both specifically be for audio testing, their applications may determine very different audio tests. This is to emphasize the need for a system where its properties, such as the actual tests implemented, can readily be changed. The CBS system is indeed easily modifiable due to its being software intensive. Its general design philosophy of using a limited number of specifically designed analog cards together with commercial computers has been used in SPCATS.

1.3 THE REQUIREMENT

The main stimulus for the SPCATS system has come from complaints by the general public as to the quality of the audio service provided by the South African Broadcasting Corporation. An effective system is therefore required to monitor and take measurements of the distribution network over a length of time. From these measurements, and others which may be implemented as a need for them arises, it is hoped to be able to pinpoint errors in the network and take appropriate action to correct them. For this reason it was decided to implement tests well known in the audio world so that use can be made of the wide coverage these tests have received in the literature over the years. Use can therefore be made of the vast available knowledge on the defects, and their causes, that these tests quantify. Once the reasons for the degradation in audio quality of the network are known, corrective action can be taken.

1.3.1 Summary of Desired Tests

There are two basic test types, passive and active, each with a set of type specific tests. In active testing the sound channel is driven by special stimulus signals whereas in the passive case this does not occur. In the former case, the original test signals are analysed at the receiver end, whereas in the latter case the receiver looks for spurious signals and measures noise levels.

Active tests can be further divided into three groups:

i) Linear distortion tests:
Theoretically these tests should only be done on a linear system and therefore, in practice, it is assumed that the circuit under test has very little non-linearity. This group includes measurements of insertion gain, amplitude vs. frequency and group delay.

ii) Non-linear distortion tests:
These all measure various types of non-linearity in a system. They can be categorised into static and dynamic type distortions. Static nonlinear distortion is evidenced by gain compression, and static, harmonic and intermodulation distortions. Dynamic distortion is evidenced by transient intermodulation distortion.

iii) Digital distortion tests:
This group consists of distortion modes specific to digital transmission channels. Here very little consensus has been reached by the international audio community as to techniques for characterising these distortions. Tests in this group are thus not clearly defined or well known as in the above two cases. However, the flexibility of the SPCATS system caters for investigating these tests as they can be added, modified and removed as more insight into their ability to quantify these types of distortions is gained.

Properties that are desirable to quantify here are bit-rate error, quantisation noise and 'expanded noise'.'Expanded' noise refers to noise due to the instantaneous or near-instantaneous companding used on some high quality digital sound channels.

The various audio tests, grouped in distortion type, are summarised below:

linear distortions:	insertion gain
	amplitude/frequency response
	stereo gain balance vs. frequency
	group delay
	stereo phase
nonlinear distortions:	gain compression
	harmonic distortion
	static intermodulation
	dynamic intermodulation
digital distortions:	quantisation distortion
	bit error distortion
	spurious in-band signals
noise:	idle channel noise
	'expanded' noise
	single tone interference
	hum
	spurious out-of-band signals
	crosstalk

Due to time restrictions and the desire to put the system to use at as early a stage as possible, only a subset of the above tests is presently available on SPCATS. These tests will be discussed in detail later in this text.

1.3.2 Intrusiveness

A fundamental design requirement of the system is that testing be done on live channels. This is indeed a novel feature in this equipment as commercial equipment usually tests passive equipment. The advantage in testing live channels is that they need not be taken out of service simply to measure their

performance. One possible approach is to use tone bursts that are as brief as possible, thereby causing no, or very little, disturbance to listeners. The only efficient way of implementing this is therefore to have computer controlled instruments to insert and extract the test tones at the respective transmit and receive ends. Besides satisfying the above requirement, the computer control alleviates the inefficiency and tedium of the previous manual testing. This type of testing necessitated taking the channel out of service.

The testing of television channels can be done non-intrusively as it is possible to insert test waveforms in the vertical blanking interval. Hence there is no disturbance to the viewer and it is possible to test all the important parameters of the channel [3]. The problem with testing audio circuits in a similar way is that no redundancy exists in either the time or the frequency domains. At the ends of the audio band the human ear manifests certain limitations, but this is not of much help as it is necessary to test the channel at frequencies throughout the band. In addition the full dynamic range of the channel should be exercised leaving no degree of freedom for totally non-intrusive tests with defined test signals. Besides using the actual programme material as the 'test signal' the only other alternative is to make the disturbance due to the inserted test signals as small as possible. Hence the use of short tone bursts is one solution. The length of these is fundamentally determined by the required resolution of the analysis at the receiver end. This was the method adopted, and will be dealt with, in detail, later.

1.3.3 Software Intensive System

A software intensive instrument has many advantages, some of which have been indicated above. These and others are summarised below:

i) Economy and ease of manufacture -
   Here several functions previously implemented in hardware can be done in software, thereby saving on

hardware costs. This hardware can consist, to a large extent, of readily available modules such as, analog-to-digital and digital-to-analog converters, and computer interfaces. Commercially available micro-computers would be used wherever possible.

ii) Flexibility -
Realisation of measurement and filtering functions in software makes it relatively easy to modify the tests in the light of growing experience of the kinds of imperfections encountered, and in conformity with evolving standards and new measurement techniques.

iii) Possibility of relatively non-intrusiveness testing -
Because of the requirement for short tone bursts (typically 25 milli-seconds) the use of most conventional analog techniques is excluded on account of the high Q sections that would be needed, these requiring long settling times. It is therefore essential to record the incoming test signals and the best way of doing this would be digitally seeing the analysis is to be done numerically.

CHAPTER 2

DEVELOPMENT OF A SPECIFICATION

2.1 INTRODUCTION

SPCATS (Sound Programme Automatic Test Set) is intended for unattended automatic routine monitoring of broadcast audio circuits, but it is desirable that it should be able to be used manually as well for fault-finding or development work. In normal use it consists of a centrally situated source of test signals and a number of remotely situated receivers, all under computer control.

The test sequence is automatically activated a specified number of times daily, and consists of active tests, at present, in which the sound channel is driven by special stimulus signals, which are analysed on arrival to characterise the circuit. The fundamenatal requirement of the system is that the tests be carried out on live channels and that audible disturbance of the channel be minimal. If necessary, these tests can be done late at night.

The transmitter unit is therefore responsible for activating test sequences at set times, generating test signals and gating these signals into the live channel. The receiver unit in turn monitors the channel, extracting these tones when they are received, doing the appropriate analysis on each, and printing out a report at the end of each test sequence.

2.2 DESIGN ASPECTS

2.2.1 Test-tone Duration

The fundamental property determining the minimum length of each stimulus tone is the resolution required by the analysis at the receiver end. Each tone is sampled and stored as a record for later analysis, after all the tones have been captured. As the first step in the analysis of a tone is fast

fourier transformation, the frequency resolution is set by [6]

Frequency Resolution = 1 / time record length

where the time record length is the time over which a test tone is sampled. This sets the minimum duration for tones above 125 Hz at 23 milli-seconds. These comprise 79% of the test tones. The tone lengths, and reasons for the required resolutions, are discussed in detail later in this text.

2.2.2 Synchronisation

Each test tone is accompanied by a 'sync-burst' to trigger the receiver and initiate capture. Immediately preceding each test tone a low amplitude 13 KHz sinusoid of 30 milli-seconds duration is added to the programme material. Because of its high frequency and its being added to the program material, it is hardly noticeable. The length of the sync-burst is therefore not critical, from the point of view of non-intrusiveness, in comparison to the actual test tones where the programme material is removed and replaced with the tone. These tones must therefore be kept as short as possible to minimise disturbance.

2.2.3 Test Signals

The initial set of tests for implementation was chosen in accordance with established internationally accepted norms due to their being well defined and their distortion mechanisms clearly understood in the literature [7,..,11]. More subtle distortion types requiring novel or speculative test methodology have been left for later inclusion. The tests presently implemented are those for linear and nonlinear distortions as classified in chapter 1. It is proposed that tests for digital distortions and noise, by using the versatility of SPCATS to do further investigations, be added at a later date.

2.2.4 System Stability

The SPCATS specification required a frequency and amplitude stability of ± 1 Hz and ± 0.1 dB respectively. To achieve this, a calibration procedure for each unit has been implemented, whereby all errors and offsets throughout the system are removed before putting the units into service. Due to the system stability the units should not have to be re-calibrated. Calibration is done by the receiver unit on a test sequence received from the transmitter unit. This is totally under software control.

2.2.5 Pseudo-random Noise

Initially, the idea of using a single burst of pseudo-random noise to determine the linear distortions seemed attractive. However, there are various problems with this:

i) To avoid intermodulation products invalidating results, a totally linear system is assumed. By using single tones to determine the linear distortions, this requirement is avoided as only fundamentals are measured and not their harmonics.

ii) To reduce the variance of results, either a long noise burst must be used or, alternatively, several short bursts. This, however, defeats the aim of using noise which is to decrease the total number of test tones required for each test sequence. Even with several bursts it is doubtful whether the required accuracy of ± 0.1 dB would be obtained.

iii) The extra computer storage, and additional complexity of doing the averaging of several noise bursts, was felt undesirable.

As a compromise between single tones and pseudo-random noise, a multi-frequency tone has been included. This consists of a sum of all the required frequencies at which linear distortion measurements are to be made. It therefore avoids the problems of variance and spurious harmonic products while saving on the overhead due to single frequency tones. The frequencies have been selected so that the harmonic products that are formed

due to the system nonlinearity do not fall on top of the fundamentals.

At present SPCATS makes use of single tones as well as the multi-frequency burst for comparison purposes.

2.2.6 Test Tone Levels

The test tones are not all sent along the programme channel at the same amplitudes. The tones involved in the linear tests are sent at lower levels so as not to excite any nonlinearities in the system which are amplitude dependent. The nonlinear tests on the other hand are large in amplitude so as to compare with the maximum amplitude of programme material that can be sent along the channel. This is so that any nonlinearity such programme material could excite will be shown up. Test tones can be sent at one of five levels:

+9, +6, +3, -6 and -12 dBm.

2.2.7 System Controllers

As mentioned before it was decided to use commercial micro-computers to drive SPCATS. As the functions performed by the transmitter and receiver controllers are different, appropriate models have been used which are optimal for the particular usage.

The transmitter micro-computer is responsible for:

i) functioning as a fairly accurate timer to initiate test sequences at specified times.
ii) Loading test waveforms from disk into the signal source.
iii) Interfacing to the hardware that puts the unit on line, adds the synchronisation bursts, sets the amplitude of the transmitted tones and gates in the tone bursts.

To implement these functions a Hewlett Packard series 80, model 85, microcomputer system was used due to its highly efficient capability in interfacing to external hardware [12].

The receiver computer performs the above tasks as well but includes several more complex ones:

i) needs to be capable of fairly accurate timing so as to know when a test sequence is about to start. In this way the synchronisation detection circuitry is only active for the duration of each test sequence.

ii) The appropriate gain must be set so as to compensate for the various amplitudes of the arriving test tones.

iii) The tones read by the ADC are stored either on disk, or in computer memory if enough is available, for later analysis.

iv) The computer must be capable of performing advanced signal processing.

v) The results of each test seqence must be printed as a hard copy.

Because the amount of hardware control by the computer is minimal and the need for computational power is high, the receiver has been implemented on a Hewlett Packard series 9000, model 220, computer [13].

Code for both computers has been written in the BASIC language of each due to the efficiency with which the languages handle interfacing to external hardware.

The functional specification below lays out the above formally and includes further requirements.

2.3 FUNCTIONAL SPECIFICATION

2.3.1 Introduction

This functional specification defines the performance parameters of the sound programme circuit automatic test set (SPCATS) to

- produce overall programme circuit performance data
- provide operational staff with daily statements of circuit performance

- be used as a manually operated measuring instrument when required.

2.3.2 Configuration and Operation

The test set comprises a distributed system operable under either manual or automatic computer control. It consists of a generator at the programme source (eg. the Johannesburg radio studio output) and an arbitrary number of receivers located at regional studio centres, FM rebroadcast stations and the output of precision radio receivers at transmitter control centres. The primary mode of operation consists of computer based analysis of tone-burst test signals inserted at one end of the circuit and digitally captured at the other.

a) The Signal Source

The test signals are inserted into the channel under test using an Insertion Unit which, when tests are not being performed, will be totally transparent to the normal programme material. When a test sequence is initiated, either automatically or on manual demand, control is assumed by the insertion unit via relay action. Relays are used so that total isolation of the unit from the channel is possible when a test sequence is not in progress. Normal programme material is routed through the insertion unit without audible degradation of quality. In this mode, however, synchronisation bursts are added to, while the actual test tones replace segments of the programme material. On termination of the test sequence, or in the case of unlikely failure of the test set generator, control is relinquished by the insertion unit, which automatically and in a fail-safe manner, returns to transparent mode. Great care has been taken in the design of the fail-safe circuitry so that even if power is removed from the unit, it will have no noticeable effect on the programme.

Audible disturbance of the programme circuit during a test sequence is minimal. The test signals consist mainly of short bursts of periodic signals (typically 30 milli-seconds in length).

The test tones are digitally generated and fed into the channel after digital-to-analog conversion and gain scaling. Each test tone is stored as a file on the computer's disk and after loading into a special purpose RAM, rapidly run out through the D/A conversion circuitry. A seperate oscillator produces the synchronisation signal. Each overall test tone then consists of a synchronisation burst added to, and the actual test tone replacing, the programme material temporarily, this being done by the insertion unit control circuitry.

b) The Receiver Units

The receivers are continuously connected across their respective audio channels. Loading of the channel by the receiver is negligible and there is no audible effect or impairment of signal-to-noise ratio.

After receipt of an initial synchronisation burst, with no test tone appended, the receiver and transmitter are in step, the receiver having a priori information about each tone to follow. It therefore sets the appropriate gain factor, before receiving each tone burst, which the computer in turn samples through a 16 bit analog-to-digital converter. The receiver unit, therefore, consists of the synchronisation circuitry, and an anti-aliasing filter followed by a precision gain-switched amplifier, under computer control, to preserve dynamic range to the analog-to-digital converter immediately following. The ADC runs at 44.4 kHz. The digitised incoming signals are stored in a transfer buffer for subsequent digital signal processing by the receiver computer. To maintain modularity, and due to the time taken to analyse each tone, the analysis of each test sequence is done after all the tones have been captured. All the analysis is done in software by the receiver computer.

c) Calibration

Due to the system's stability (i.e. ± 1 Hz, ± 0.1 dB) it can be calibrated before being put into service. This is done by connecting a receiver unit to the signal insertion unit and running the calibration option of the software. The calibration values are kept on a file on the receiver's disk drive. Using this procedure the overall system is calibrated.

d) Synchronisation

The system, at present, consists of nineteen independent test tones and twenty synchronisation bursts. The first burst is not accompanied by a test tone and is used to synchronise the receivers and transmitter. This is necessary so that the receiver can confirm the start of a test sequence and set the appropriate gain factor for each ensuing test tone. Should a tone not be received within a specified interval, the receiver will configure itself for the next one in the sequence. This means that a whole test sequence is not wasted simply due to a few tones being missed. (This could be caused by brief programme switching, for example, which occurs very seldom, therefore).

2.4 TESTS AND MEASUREMENTS

The tests implemented on SPCATS at present are the active tests for linear and nonlinear distortions. For clarity the tests are described seperately below, although some use data gained from the same test tone. These tests are directly based on the original project specification [14]. Note that although the exact frequencies specified in this document have not been met, the ones implemented are close enough to make no practical difference. The worst case deviation is 2.5%, 79% of the tones being within 0.5% of the original specification. The exact values of the tones is not important, only their frequency stability. The reason for not implementing the specified frequencies exactly is that the digital generator has a finite frequency resolution.

A fast fourier transform is performed on each tone recieved and the pertinent information then extracted.

2.4.1 Linear Distortions

a) Insertion Gain

This is a measure of overall insertion gain. Only the fundamental component is measured here.

Source:

frequency	996 Hz
level	-12 dBm
accuracy	± 0.05 dB

Measurement:

units	dB
accuracy	± 0.1 dB
range	+10 to -20 dB relative to nominal.

b) Amplitude / Frequency Response

Two methods are implemented. The first uses a series of single frequnecy tone bursts whereas the second uses one multi-frequency burst consisting of all the single frequencies.

Source:

frequencies	39, 127, 400, 5615, 7177, 10009, 14989 Hz
level	-12 dBm
accuracy	± 0.05 dB

Measurement:

units	dB
accuracy	± 0.1 dB
range	+10 to -20 dB relative to nominal.

2.4.2 Nonlinear Distortions

a) Gain Compression

Here the difference between a 1 KHz tone burst at -6 dBm and +6 dBm is measured.

Source:

frequency	996 Hz
level	+6 dBm, -6 dBm
accuracy	± 0.1 dB

Measurement:

units	dB
accuracy	± 0.1 dB
range	20 dB

b) Harmonic Distortion

A fundamental is inserted into the system under test and the received fundamental and second and third harmonics, up to 15 KHz, are measured.

Source:

frequencies	1) 39, 127, 400, 996 Hz
	2) 5615, 7177 Hz
level	1) +9 dBm
	2) +6 dBm
accuracy	± 0.1 dB

Measurement:

units	%
accuracy	± 0.05%
range	0.1% to 50%

c) Static Intermodulation

The channel is driven by a dual-frequency tone burst and the sum and difference frequencies extracted.

Source:

frequencies	5615 added to 7177 Hz
level	+3 dBm
accuracy	± 0.1 dB

Measurement:

units	%
accuracy	± 0.05%
range	0.1% to 50%

d) Dynamic Intermodulation

The I.E.C. adopted standard for this test has been implemented [15]. A 15 KHz sinusoid is added to a 3.15 KHz low-pass filtered square-wave. The filter is a single pole, 30 KHz low-pass. The peak-to-peak voltage ratios of the two tones is 1:4 respectively. Specific sidebands generated by the channel are measured.

Source:

frequencies	3.15 KHz square-wave and 15.00 KHz sine-wave.
level	+6 dBm
accuracy	± 0.1 dB

Measurement:

units	%
accuracy	± 0.05%
range	0.1% to 50%

2.5 ELECTRICAL CHARACTERISTICS OF THE TEST UNIT

2.5.1 Insertion Unit

The transmitter insertion unit operates in two modes, 'insertion' and 'transparent'. Its normal, or quiescent, state is transparent mode in which it is isolated by relay contact from the sound circuit and has no measurable effect on it. In the event of failure, reversion to transparent mode is automatic.

In insertion mode, the electrical performance of the unit is compatible with the measurement criteria, and causes no audible, or measureable, degradation to the programme material apart from the actual test signals inserted. The exception to this is that, only during a test sequence, a transmission notch is inserted between 13 and 14 KHz. The synchronisation burst, at 12.75 KHz, is added after the notch.

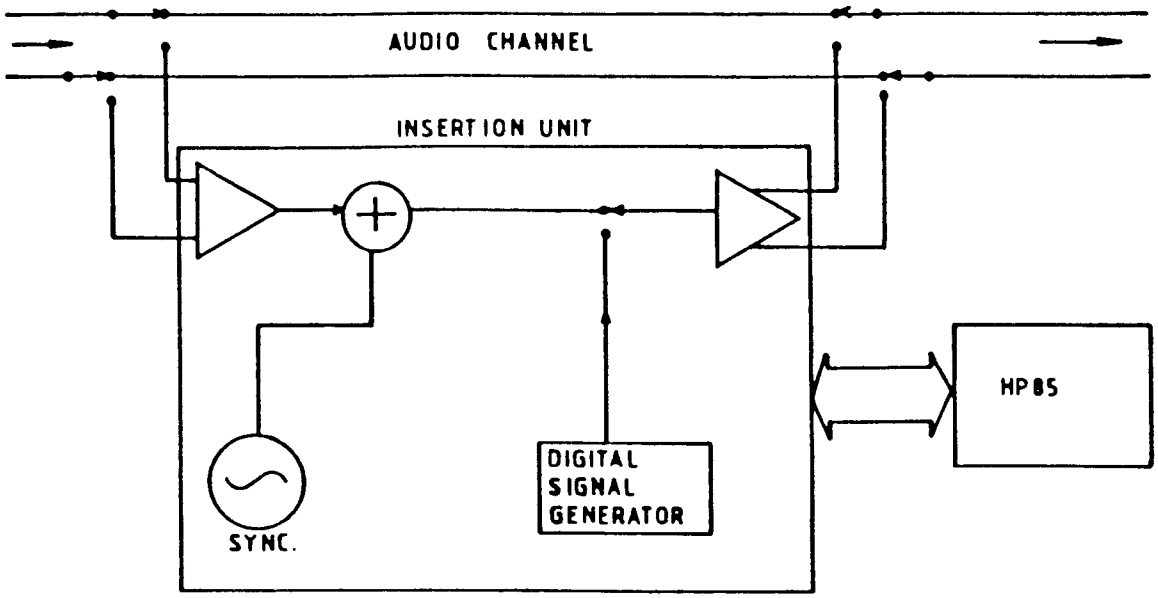

FIGURE 2.1 Insertion unit

SPECIFICATIONS:

Input impedance	40 KΩ balanced.
Output impedance	100 Ω balanced.
Common-mode rejection	> 40 dB
Distortion:	
Static	< 0.01%
Dynamic	< 0.1%
Frequency response	± 0.05 dB- 40 Hz to 13 KHz

2.5.2 Receiver Unit

This consists of a differential line receiver, as all audio lines are balanced, driving a tone decoder and gain-controlled amplifier. The signal to the decoder, however, is first high-pass filtered and that to the amplifier, passed through an anti-alias filter. The output of the gain-controlled section is passed to a 16 bit analog-to-digital converter which is read directly by the computer (fig. 2.2).

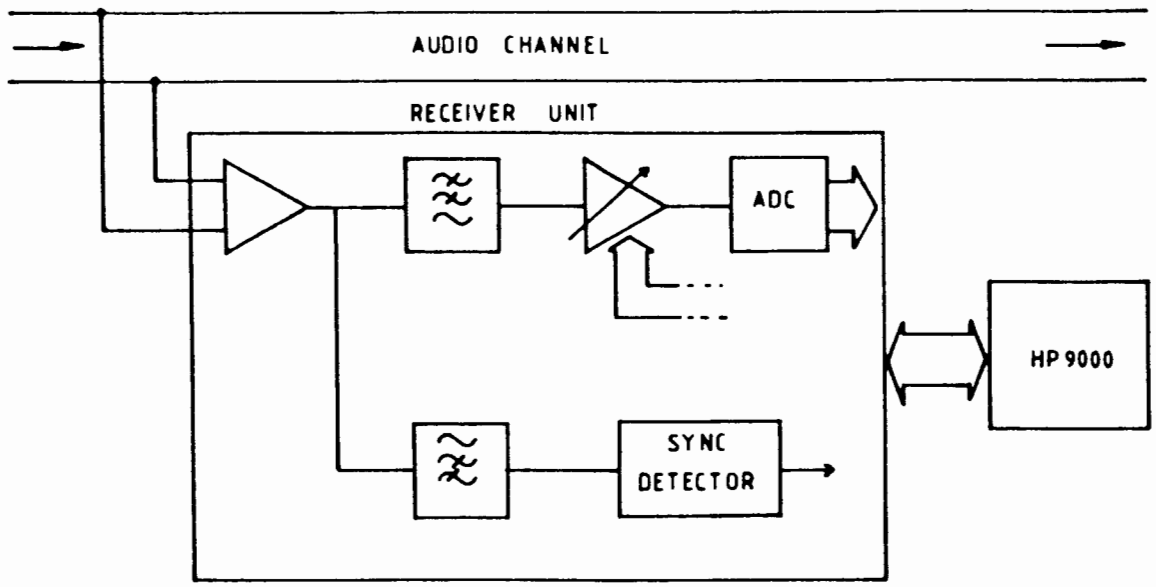

FIGURE 2.2 Receiver unit

SPECIFICATIONS:

Input impedance > 40 KΩ

The specifications in respect of frequency response, distortion and common-mode rejection match those of the insertion unit.

2.5.3 Environmental Considerations

The test set can operate in close proximity to a variety of broadcasting equipment without undesirable interaction.

Because the unit functions in-doors it only needs to operate to specification over the temperature range from 15°C to 35°C.

CHAPTER 3

SYSTEM CONCEPTS

3.1 INTRODUCTION

The principles and ideas used in the design of SPCATS will be discussed here. Some are novel while others are applications of well established concepts in electrical engineering.

3.2 SIGNAL GENERATION

The 'insertion unit' at the transmitter end consists of a signal generator, gating circuitry and the controller. The gating circuitry simply routes the programme material, synchronisation bursts and tone bursts as well as setting the appropriate gain for each tone. The generator, together with the controller is more complex and requires more discussion.

3.2.1 Generation Strategy

The functioning of the digital signal generator will be considered here and the technique used to obtain an effective 14 bit digital-to-analog converter using two 8 bit DACs. The errors inherent in DACs will also be discussed and the effect these have on the accuracy obtained from using this technique.

The generator only outputs full scale waveforms, these having previously been calculated in software and stored as files on disk. The generator has two modes of operation. When a particular waveform is desired it is read from disk and loaded into the generator. This occurs during the 'Load' mode. In the 'Run' mode the continuous-time waveform is available at the output of the generator. It is this waveform which is gated into the audio channel as a tone burst.

The generator has a dynamic range of 14 bits and due to the software error and temperature stabilisation applied, an accuracy of 13 bits. Accuracy is defined here to mean the maximum deviation of the output from its designed value, over

the full range of operating conditions [16]. With 14 bits a resolution of -84 dB is obtained whereas a resolution of -80 dB is required. Note that although a resolution of -80 dB is needed, the accuracy is specified to -70 dB [14]. This caters for cumulative system errors which make it unlikely to obtain an accuracy equal to the resolution.

a) Configuration

Due to a commercial 14 bit digital-to-analog converter (DAC) not being readily available at the time as well as the cost of such a converter, an alternative approach was adopted. A software correction technique was developed which allows the generator to consist of 2 eight bit DACs (DAC 08's), one providing the most significant byte and the other the least significant byte of a 16 bit input word. To compensate for the errors in the most significant DAC a look-up table of its precise transfer characteristic is used in calculating output waveforms. For this technique to work an overlap of 2 bits has been allowed. Thus the generator input of 16 bits maps to a 14 bit output. The word into the DACs comes from 2 eight bit, 2 K RAMs, one holding the most significant byte (MSB) and the other the least significant byte (LSB). A 12 bit counter steps through the RAMs to present successive words to the converters. The counter counts cyclically, resetting at the top of its count and then repeating. In the 'Load' mode the computer steps the counter while loading the RAMs with the new waveform whereas in 'Run' mode a crystal controlled clock is fed to it. This clock can run at 100 KHz or 20 KHz thereby catering for the high (15 KHz) and low frequency (40 Hz) waveforms.

To prevent spurious signals appearing at the outputs of the converters, latches are placed between the RAMs and the converters. Only when the RAMs have settled is the new word latched. The settling transients at the outputs of the DACs contain only high frequency energy and can be ignored. A sample/hold amplifier could have been placed after the DACs but this introduces errors of its own [17].

The outputs of the two DACs are summed, with the gain for the lower byte being one-hundredth of that for the higher. It is this weighting which makes it possible to have two identical RAMs plus DAC subsystems implementing different sections of the 14 bit waveform. The overall system layout is shown below.

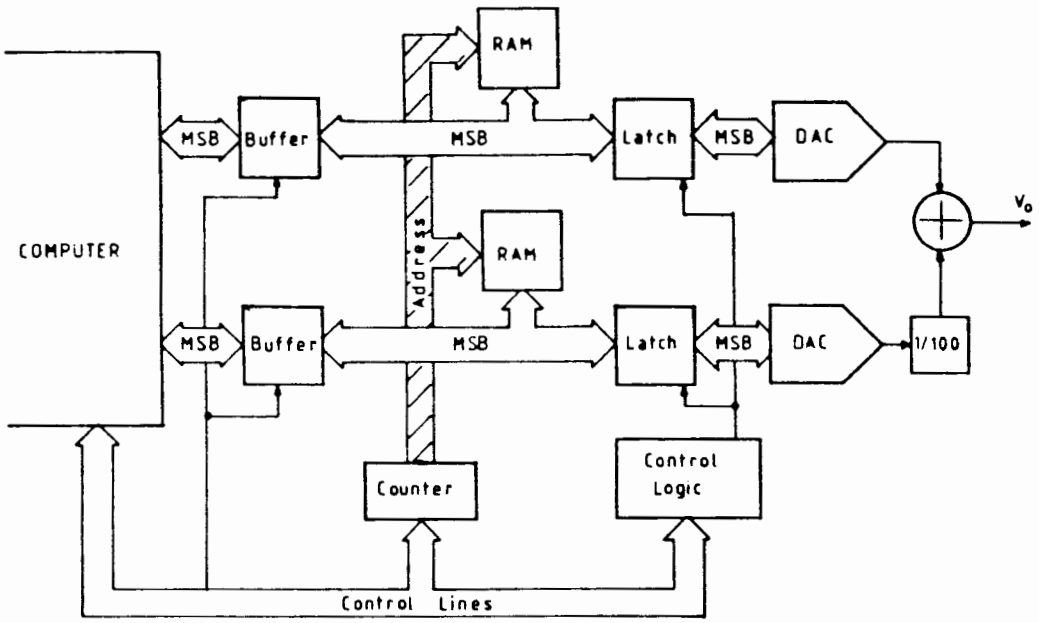

FIGURE 3.1 Digital signal generator configuration

The buffers are tri-state and decouple the computer's data bus from that of the generator once a waveform has been loaded. The DACs are used in offset binary, continuous conversion mode. The word map for each is therefore:

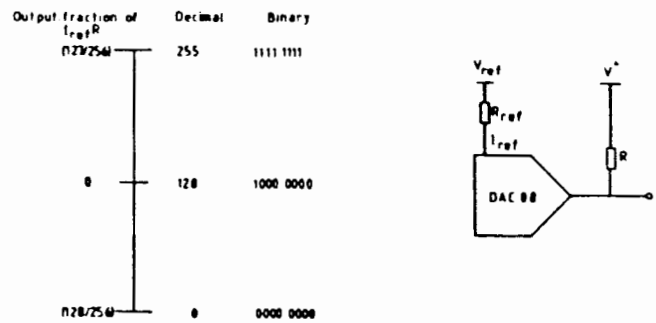

FIGURE 3.2 Generator word map

b) Digital-to-Analog Converter Errors

To show why an overlap of 2 bits was chosen it is necessary to discuss the main errors in practical DACs. These are zero, gain and linearity errors.

i) Zero (offset) error: An ideal DAC would produce zero volts output when the byte representing it is input (in our case 128). Any analog output other than zero under these input conditions is called the offset or zero error of the DAC. As this is simply an offset, in that the DAC transfer function does not pass through the origin, it can be trimmed to zero.

ii) Gain (scale factor) error: The slope of a converter's transfer function determines its gain i.e. the scale factor which establishes the relationship between the analog values and the digital codes. The deviation of the converter's actual slope from the ideal is termed the gain error. This is an error which can also be trimmed to zero.

iii) Linearity error: The zero and gain errors are linear effects and are therefore not significant in the digital signal generator. It will be shown later that they can be ignored due to the software error correction applied to the generator. The non-linearity of the DAC does require attention as these errors cannot be ignored, the function of the software correction technique being to compensate for these.

The converter's linearity error is specified as the maximum deviation of its transfer function from a straight line passing through the zero and full-scale points under the conditions in which the converter's offset and full-scale gain errors have been nulled. This definition is sometimes referred to as the 'integral linearity error'.

A more useful indicator of non-linearity is the 'differential linearity error'. This is the maximum deviation from the ideal least significant bit (LSB) analog difference between any two adjacent code values over the complete range of the DAC. For example, a converter with a differential linearity error of

± ½ LSB has all the quantum steps in its transfer function never less than ½ LSB and never more than 1½ LSBs. In the DAC08, and most commercial DACs for that matter, this error arises due to errors in the weighting of the DAC's bit current increments.

Monotonicity is an aspect of differential linearity. A monotonic DAC is one whose analog output continuously increases as its digital input code is incremented. A non-monotonic DAC has a differential error greater than ± 1 LSB.

The overall accuracy of a DAC is therefore determined by combining these and its other errors as one figure.

c) Error Analysis

Consider the accuracy of a DAC0800LC over the temperature range 15°C to 40°C and with gain and offset errors trimmed to zero at 20°C. The temperature range has been chosen as the unit is to be kept indoors. Please refer to figure 3.2 in the discussion below:

i) Offset error with temperature: This error has a drift with temperature. Assume a variation of 1 uA in the zero scale output current over the temperature range as this is not specified in the device's data sheet [18].

Total offset error ≈ 1 uA x 5 k (Where R of fig. 3.2 is nominally 5 K).

≈ 0.00 % and can be ignored.

ii) Initial gain error with temperature: As both DACs use the same V_{ref} its variation as regards the bit overlap can be ignored and only that of R_{ref} and the pull-up resistors at the outputs need be considered. Assuming 1% resistors and that they have been matched on a bridge, let their combined temperature variation be ± 0.5% over the range. The temperature variation of the DAC full-scale currents must also be taken into account (± 50 ppm/°C).

Overall gain error = ± 0.5% ± 50 ppm/°C.
= ± 0.5% ± 0.12%
= ± 0.62%

iii) Linearity error with temperature: This is specified as ± 0.19% over a 70°C range. Assuming a linear variation this can be taken as ± 0.07% over the range of interest.

Consider figure 3.3 below where the offset error has been nulled:

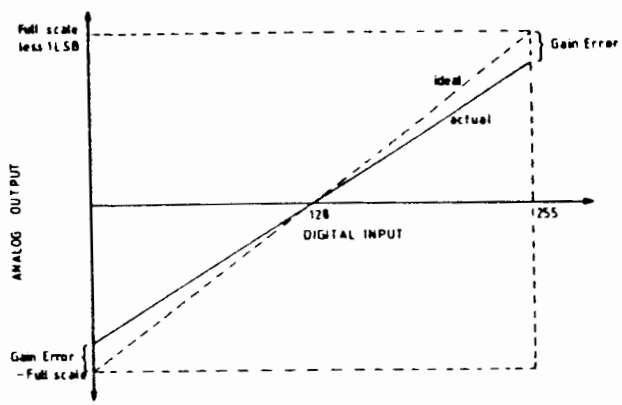

FIGURE 3.3 DAC transfer characteristic (zero offset error)

We wish to obtain effective overlap of the least significant bit (LSB) of the high order DAC (X) with the top of the low order DAC (Y). Hence in the former case we are looking at the errors in the LSB and in the latter at those in the most significant bit (MSB). The LSB of an 8-bit converter is 0.39% of the full-scale volts and we have seen that the DAC gain and offset errors can be nulled. As the maximum gain variation in the nulled DAC occurs at full-scale, the worst case error occurs at this point. This error is therefore due to the variation of the gain (set by the external components) and the DAC linearity. Hence the device's accuracy is

accuracy ≈ ± 0.62% ± 0.07%
≈ ± 0.69%

If we were using two perfect DACs we could implement a 16-bit DAC where the weighting of the MSB of the low byte would be $1/2^{n+1}$ = 1/512 due to the LSB of the high byte being precise. In this case there would be no overlap required. However, due

to the accuracy of ± 0.69% in the high byte DAC an overlap of 2 bits is required as it could be in error to approximately 1½ least significant bits (0.39% + 0.39% = .78%). The accuracy of the low byte DAC can be ignored as it is decreased by the weighting factor (1/128 x 0.69%) relative to the high byte DAC. As the computation above is for the worst case the actual overlap requirement may not be as stringent. However, a resolution of 14 bits is well within the system specification. The accuracy of ± 0.69% of the low byte DAC gives the required overall accuracy to 13 bits.

3.2.2 The Computer Interface

The generator is driven by a Hewlett-Packard HP-85 micro-computer which is interfaced via its 'general purpose input/output' (GPIO) interface. This unit has its own processor to handle such tasks as handshaking and transferring data to the computer. The interface is driven by simple commands which form part of the User's program.

The interface can provide 16-bit or dual 8-bit data exchange. The 16-bit format has been used as it takes half as long to do one such transfer compared to the time taken to do two 8-bit transfers. The data-exchange timing and logic can be configured by switches or dynamically under program control to meet a wide variety of peripheral requirements. It is this flexibility which makes this interface extremely attractive for interfacing to almost any type of custom built hardware device.

The interface has two 8-bit, low-power, bi-directional ports (A and B), and two 8-bit, output only ports (C and D). Each port uses a two wire handshake for I/O data. In 'word' mode the two ports of the same type work in parallel to implement the 16-bit word.

There is a control port of eight lines; four of these are control lines,two flag events and two indicate the status of devices. The control lines are driven by the computer and can be independently set and reset from software. The flag and

status lines are driven by the peripheral and can be interrogated from software.

Interface configuaration is done by writing an address to the interface to select one of the fifteen possible modes. Each mode selects one control line and either a flag or status line with which to do handshaking of data. The User is then free to customise the remainder of the control port to the particular application.

There are also various handshake methods possible. These are full, partial, strobe and no handshake. The full and partial modes use both handshake lines whereas strobe handshaking uses only the control line, which is pulsed whenever valid data is on the bus.

The configuration used to drive the insertion unit plus its generator is as follows:

i) Ports C and D, and A and B, used together in word mode.
ii) Control line B (CTL B) and flag line B (FLG B) used for handshaking.
iii) Pulse handshake used.

3.2.3 Computation of the Set of Test Signals

Computation of the numeric values making up a digital waveform will be discussed here for the specific case of storing these in a finite memory from where they are clocked out. The corrections applied to these values to accommodate the imperfect digital-to-analog conversion system used will be discussed later.

Any waveform clocked out of a finite memory must have its samples repeat once the end of the memory buffer has been reached. This makes it possible to keep cycling through the buffer without any discontinuities appearing in the output waveform.

The above may be formulated as follows:
Consider a sinusoid, $A.\sin(2\pi f_d)$ with amplitude A and frequency f_d. As a discrete signal it may be written as

$$x(k) = A.\sin(2\pi f_d kT) \quad \ldots\ldots\ldots\ldots\ldots\ldots (3.1)$$

where

$$f_s = 1/T$$

is the rate at which samples are output. The output of the DAC can therefore be regarded as a train of impulses

$$y(t) = \sum_{k=-\infty}^{\infty} y(k)\delta(t-kT) \quad \ldots\ldots\ldots\ldots\ldots\ldots (3.2)$$

where

$$y(k) = Q(x(k))$$

with Q(.) being the quantisation characteristic of the DAC and thereby converting x(k) from a continuous to discrete function.

Now to make the memory requirement finite for storing these samples, it is necessary that the samples x(k) repeat with a period N. Hence

$$x(k+N) = x(k) \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (3.3)$$

must be satisfied. This is equivalent to

$$f_d NT = L \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (3.4)$$

for some integer L. The finite set of frequencies obtainable is therefore

$$f_d = Lf_s/N \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (3.5)$$

This can be shown as follows

$$x(k) = x(k+N)$$

and for a sinusoid

$$\sin(2\pi f_d kT) = \sin[2\pi f_d (k+N)T]$$
$$= \sin[2\pi f_d kT + 2\pi f_d NT]$$

For this to be true we must have

$$2\pi f_d NT = 2\pi L$$

for L an integer. Practically L must be an integer. Hence

$$f_d = f_s L/N$$

To avoid unnecessary complication in the digital signal generator it was decided to design the control logic to cycle through the RAM in only one direction. For this reason

possible advantages in using symmetries of sine waves were not exploited [19]. In any event symmetries in the two-tone and transient intermodulation test tones are not as readily computed. Instead a large amount of RAM was used thereby increasing the set of possible output frequencies. Due to RAM being extremely cheap this solution was felt optimal.

Furthermore, the signal generator has a fixed memory size, this also saving on the hardware requirement. The above two restrictions limit the set of frequencies the generator can produce but as the aim is only to produce certain spot frequencies and not a complex frequency synthesizer, this does not matter.

One addition that was found necessary due to the large frequency span that the generator has to accomodate was to have a low (20 KHz) and a high (100 KHz) frequency clock. One or the other clock is selected from software depending on the waveform being output.

By fixing the memory size at 2K, the available frequencies are

$$f_d = f_s L/2048 \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots \quad (3.6)$$

where L is the number of waveform cycles stored. It is this requirement of an integer number of cycles that is responsible for limiting the set of realisable frequencies. Hence the output equation becomes

$$x(k) = A.\sin(2\pi Lk/2048) \quad \ldots\ldots\ldots\ldots\ldots \quad (3.7)$$

Note that the phase has been set to zero as we are free to do so.

The above formulae apply equally well to the two-tone and transient intermodulation tests. The former is simply a summation of two different frequency sinusoids and the latter also uses two summed sinusoids with the polarity of one being used to generate a square wave. The actual generation algorithm will be considered further on in this text.

3.3 SIGNAL CAPTURE

The fundamental design and analysis of the receiver unit are discussed here. The receiver unit consists of a data acquisition section interfaced to a computer. The two most important elements of the former are the analog-to-digital converter and the sample-and-hold amplifier.

3.3.1 Analog-to-digital Requirements

The system specification [14] called for a measurement resolution of -80 dB and an accuracy of -70 dB. The resolution of an analog-to-digital converter (ADC) is a fundamental property of its number of bits and so for the required resolution

$$20.\mathrm{LOG}(1/2^n) \leq 80 \text{ dB}$$

must be satisfied where n is the number of bits. With 14 bits a resolution of -84 dB is obtained.

One of the important measures of quality of a digital conversion system is the ratio of the maximum signal to the quantisation error. This ratio is a function of the number of ADC bits. The quantisation error becomes statistically random from sample to sample when the analog signal is at a high level and spectrally wide-band. To all practical intents this is true for realistic signals near the full-scale input range of the ADC. The probability density function of the error is flat with equal energy at all frequencies and therefore so is the energy density function, X(f).

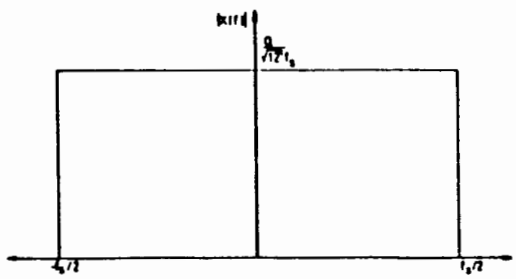

FIGURE 3.4 Quantisation-error energy density function

For a full-scale sinusoid quantised with an n-bit ideal converter, the well known expression for the signal-to-quantisation noise ratio (SNR) is

$$SNR\ [dB] = (6.02)n + 1.76$$

which gives a SNR of 98 dB for a 16 bit converter.

The accuracy of any practical ADC is less than its resolution due to errors other than that due to quantisation. As in the case of a digital-to-analog converter (DAC) these are gain (scale factor) error, offset (zero) error and non-linearity. The first two are not important as they can be cancelled by calibrating the ADC. Its non-linearity however can not be cancelled and therefore its specification must meet the -70 dB accuracy requirement.

Two further specification terms used are 'absolute' and 'relative' accuracy. Absolute accuracy is a measure which comprises gain error, zero error and non-linearity, together with noise. The relative accuracy error is the deviation of the analog value at any code from its theoretical value, after the full-scale range has been calibrated i.e. after zero and gain errors have been cancelled [20].

Calibration of the SPCATS system is done before putting it into service and it must therefore remain accurate to ± 0.05 dB over a long time. The system stability was therefore an important criterion in the design. By a system's stability we usually refer to the invariance of its characteristics with time, temperature variation, etc. This is usually quantified by the manufacturer providing stability versus temperature measurements over the device's temperature range. Parameters usually quoted include gain, zero and linearity errors versus temperature. Whereas linearity was important in choosing an ADC to satisfy the accuracy requirement rather than the gain and zero error, all three are relevant when deciding on a device's stability in this application. The zero and offset error temperature coefficients must be sufficiently small so that measurements made by the system are stable and within the ± 0.05 dB tolerance over the 15°C to 40°C operating range. The

linearity/temperature specification must guarantee the device's accuracy to better than -70 dB over the same range. Actually it is necessary that the performance be somewhat better than the limits set above so that notwithstanding cumulative errors due to the ADC and all other components of the system the overall accuracy requirement is still satisfied.

A further parameter which usually only requires checking as it is satisfied by any reasonable supply is the power-supply stability. Since high performance ADCs have on-board stabilised reference voltages they do not require extremely stable supplies.

In addition to actual performance details, attention must also be given to protocols and interfacing requirements so as to choose suitable logic and design the correct timing sequence to drive the device.

Successive approximation converters are typical in the range of speed and accuracy required and a Burr-Brown PCM75 was chosen for the system.

3.3.2 Sample and Hold Requirements

The requirements for the sample-and-hold are often assumed trivial by designers of high resolution ADC equipment when in fact this component is crucial to system performance [21]. This is borne out by high performance sample/holds and ADCs being in the same price range. The principle application of sample/hold amplifiers, as in the present case, is to maintain an ADC's input constant during conversion at a value representing the analog input at a specific time. An ideal sample/hold exactly retains the last value it had when the command to 'hold' was given, and it retains that value until the logic input dictates 'sample', at which point the output ideally jumps to the input value and follows the input until the next 'hold' command is given. Real sample-and-holds deviate from this behavior in several ways which leads to various errors. Because the sample/hold's performance in

various errors. Because the sample/hold's performance in sample mode is similar to that of a closed-loop operational amplifier it has similar specifications. Note, however, that it has a far smaller gain-bandwidth product and slewing rate and a much longer settling time due to the relatively large storage capacitor it needs to charge. Op-amps use much smaller capacitor values for compensation hence their improved performance as regards the above parameters. In addition, during the sample-to-hold, hold and hold-to-sample states, the dynamic nature of the mode-switching introduces a number of properties that are peculiar to sample/holds. These are shown in figure 3.5 and defined below.

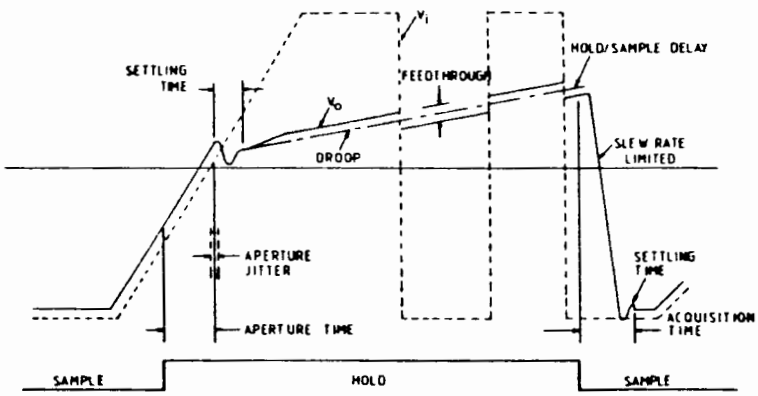

FIGURE 3.5 Sample/hold dynamic behaviour

i) Acquisition time: This is the time taken by the output of the sample/hold to reach its final value after the 'sample' command is given. The final value is specified to within a certain error band as well as the magnitude of the input step used as the input signal.

ii) Aperture time (delay): The time required after the 'hold' command for the switch to open fully. As this is simply a fixed time delay it need only be considered in a few applications where the sample must be taken precisely at the 'hold' transition.

iii) Aperture jitter: This is probably the most important sample/hold parameter. It is the variation in the aperture delay from sample to sample and due to it being random it cannot be cancelled as in the case of the delay where the hold command may be advanced by an equivalent amount.

iv) Droop: The change of the output voltage during 'hold' as a result of leakage or bias currents flowing through the storage capacitor. Increasing the storage capacitor improves the droop but increases the acquisition time and therefore sample-and-holds with on-board capacitors are specified for certain applications whereas those with external capacitors have a value chosen for the application by the User. Droop is not simply an overall d.c. offset when the input signal is near full-scale. Here a signal near the edges of the ADC input range may go out of range due to the droop of the sample/hold. This results in the sampled waveform effectively being clipped at its peaks thereby being confined to the ADC input range. This generates distortion products in the sampled waveform. Samples still within the ADC input range after the droop has occurred will simply have a d.c. offset as the rate of discharge is constant and independent of the signal's amplitude.

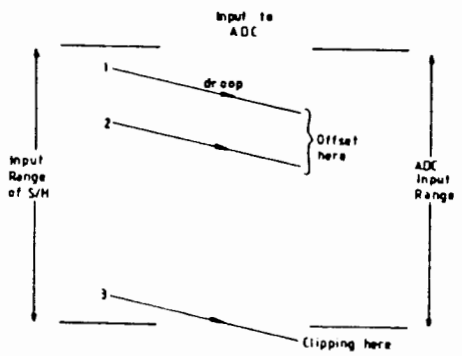

FIGURE 3.6 Effect of sample/hold droop on ADC input range

The distortion can be avoided by either making the droop sufficiently small or making the range of input signals smaller so that the range of the signals plus droop is within that of the ADC. The problem with the latter solution is that there is a loss in dynamic range equal to the number of bits by which the sample/hold input range is reduced assuming the two ranges initially equal.

v) Sample-to-hold offset: This is the shift in level between the last value in 'sample' mode and the value settled to in 'hold' mode. It consists of a linear component, the 'charge transfer' (or 'offset step'), and the non-linear 'offset non-linearity'. Charge transfer is the charge

transferred to the storage capacitor via stray capacitance, particularly that of the switch, when switching to hold mode.

vi) Feedthrough: The fraction of the input signal variation that appears at the output in 'hold'. This is caused by stray capacitance coupling from the input to the storage capacitor, principally across the open switch [20].

3.3.3 Aperture Jitter

A sample-and-hold amplifier serves two independent functions: it samples discrete values of the input signal at a fixed rate, and it holds those values constant during the interval for digitisation [17]. Until recently the only parameter specifying the former function has been the 'aperture time' and the uncertainty in this value has seldom been specified. Only in recent years with the advent of high performance sample/holds has 'aperture uncertainty' ('jitter') become a standard specification with such devices [20,21,22].

There are two mechanisms which can cause errors in the effective sampling period once the aperture time of the sample/hold has been taken into account. These are the aperture uncertainty of the device itself and phase jitter in the main digital clock signal driving it. These generate different error types in the resulting sampled waveform.

Aperture jitter is caused by the time to turn off the FET switch of the sample/hold being a function of the residual charge in the circuit and the source-gate shut-off voltage. This voltage can also be a function of the analog input signal. Because this non-linearity is primarily an asymetric effect (sign dependent), the error components will tend to be mostly second harmonic [21]. The error voltage at each sample is proportional to the signal derivative and the jitter time. The maximum signal frequency for which a specified maximum voltage error will occur can therefore be computed.

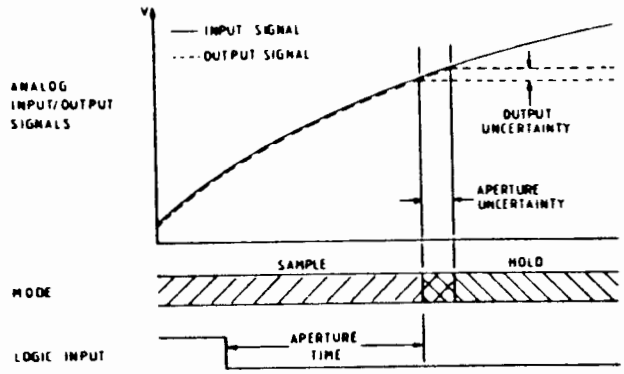

FIGURE 3.7 Aperture delay and jitter

For sine-wave inputs

$$\delta E = (f_{max}/E_{FS})/(2\pi\delta t)$$

or

$$f_{max} = (\delta E/E_{FS})/(2\pi\delta t)$$

where: δE = the allowable voltage uncertainty
E_{FS} = the sin-wave magnitude
δt = aperture jitter

Because the error is proportional to the rate-of-change of the input signal, it becomes worse as the input frequency is increased. Hence it is worst at the top of the input frequency band. To reduce the worst case peak error to one least-significant-bit for a 16 bit converter requires an aperture jitter of the order of 200 pica-seconds for a full-scale maximum-frequency audio-bandwidth signal (15 KHz). It is the requirement of small aperture jitter in high precision sample/holds that makes them so costly.

Clock jitter in the logic signal driving the sample-and-hold produces additive noise rather than distortion products. This is due to the mechanisms that cause it. The major one being noise on the signal into the comparator that generates the digital clock, this noise being random. This is applicable to any gate, flip-flop, etc. where a logic threshold is determined [23]. The error generated by the sample and hold is thus random even though it is proportional to the signal

amplitude and frequency. The general expression for the broad-band noise produced by gaussian time jitter is [17]

$$SNR = -20.LOG(2\pi f \delta t)$$

where: f = signal frequency
δt = standard deviation of the time jitter

To avoid problems due to time jitter it is essential to design the clock circuitry so that the edges of a crystal clock are used to drive the sample/hold via as few logic devices as possible as these devices each contribute a small amount of jitter.

3.3.4 Data Recording

The sampling rate of a data acquisition system is chosen based on two important trade-offs. The higher the frequency is above the baseband's maximum frequency the less severe the specification of the anti-alias filter needs to be. However, with an increase in frequency comes the more severe problem of finding compatible devices to run at the higher frequency.

The major criterion in deciding on the sampling rate was for it to be compatible with the analog-to-digital converter. As a high quality ADC was extremely difficult to obtain, this eventually having to be imported, the system design had to be done around this device.

To avoid additional complexity in the SPCATS receiver's data acquisition unit it was desirable to have the computer reading the ADC directly as each sample is converted. This avoids the need of having RAM at the acquisition unit which stores a buffer of samples to be read later by the general purpose input/output interface of the computer.

Finally, an anti-alias filter had to be designed which had a pass-band up to 15 KHz and was flat to within ± 0.05 dB, although it was found that this specification could be eased somewhat due to the calibration procedure that was implemented

in the unit. The stop-band in turn had to attenuate by 80 dB above the foldover frequency to avoid aliasing.

The criteria in choosing the sampling rate were therefore:

i) compatibility with the available ADC.

ii) Low enough so that the computer could read the ADC directly.

iii) As high as possible to make the anti-alias filter specification less stringent so that its implementation would not be unduly complex.

An investigation of the currently standardised PCM sampling frequencies showed that, as with most standards, no single standard exists. The AES recommends a 48 KHz sampling frequency for all applications although it recognises the use of 44.1 KHz in certain consumer products and 32 KHz for broadcast and transmission-related applications [24]. The Japanese have adopted 44.056 KHz as industry standard, this being specified by the 'Electronic Industries Association of Japan' (EIAJ) [25]. This is in use at the Victor Company of Japan (JVC), for example [26].

A sampling rate close to 44.056 KHz (i.e. 44.4 KHz) was chosen although any frequency satisfying the three criteria above would have been suitable. This is ideal for the ADC used in the data acquisition unit as it is specifically designed for audio PCM. In addition this frequency is well within the speed limit at which the computer can transfer data from an external device to memory. The EIAJ has specified a 20 KHz audio bandwidth for its 44 KHz sampling frequency whereas SPCATS works over 15 KHz This suggested that the anti-alias filter requirement was readily realisable as that used by Japanese industry is more severe and such filters are mass produced.

With a sample rate of 44.4 KHz and a baseband of 15 KHz the specification for the anti-alias filter is set: (figure 3.8)

i) Pass-band: 0 -> 15 KHz with ± 0.05 dB ripple.

ii) Transition-band: 15 -> 22.2 KHz

iii) Stop-band: 22.2 KHz -> ∞ with 80 dB attenuation throughout this band.

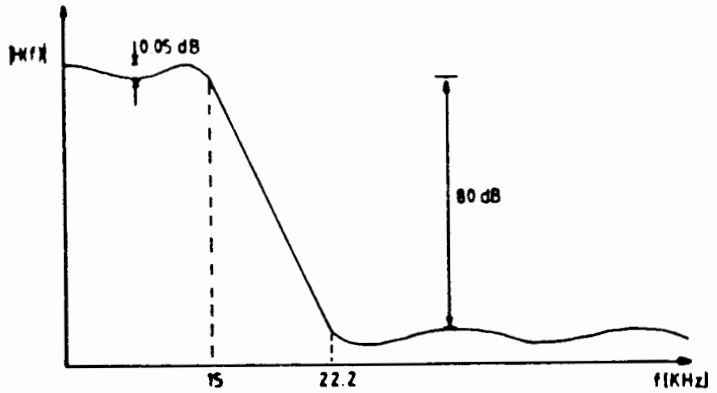

FIGURE 3.8 Anti-alias filter requirement

3.3.5 The Computer Interface

The Hewlett-Packard model 220 'General Purpose Input/Output' Interface (GPIO) is used to read data from the data acquisition hardware into computer memory as well as to select the appropriate gain on the programmable gain controller.

The interface can perform a 16-bit bi-directional data exchange between the computer and peripheral. Data can be transferred in full-duplex mode (i.e. it can receive and send data at the same time). External interrupt and User-definable signal lines are provided for additional flexibility. There are three lines to perform handshaking, one driven by the interface, the other by the peripheral and the third to indicate the direction of data flow. The interface can also be interrupted by the peripheral device. Four general purpose lines are available, two driven by the interface and the others by the peripheral. These can be configured by the User for any application specific function.

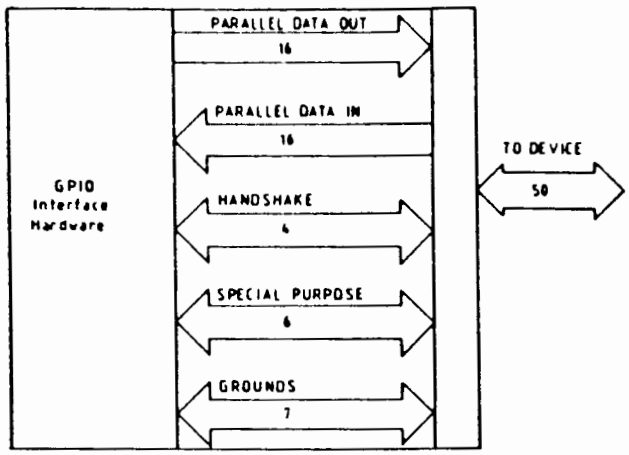

FIGURE 3.9 GPIO interface

The interface has two modes of handshaking data. Each uses all three handshake lines and the only difference between them is the relative timing between the control signals. In addition the edges of the control signals used to clock data into the interface's input registers can be hardware set.

The interface is controlled by simple program statements from where data transfers can be initiated, registers read and written to, and external interrupts initiated. There are three types of data transfer; interrupt, fast-handshake and DMA.

With an interrupt transfer concurrent processing occurs while a data transfer is in progress. Whenever the peripheral signals that it is ready to receive or send data the computer's processor is interrupted to transfer the word before continuing execution at the point it left off. The maximum obtainable speed here is 65K transfers/second. for a burst of data. Data arriving slightly slower decreases the transfer rate considerably.

The fast-handshake transfer (FHS) can move data at a maximum rate of 115K transfers/second. Here the processor remains interrupted for the duration of the transfer. Therefore, if the transfer is not completed for some reason the computer will 'hang up'. With this method the handshaking is not as critical as for the interrupt mode and all samples received by the interface will be read provided the maximum transfer rate

is not exceeded. This method has been used to read the data from the ADC.

The DMA transfer uses a DMA controller to transfer data directly between the computer's memory and the peripheral and is therefore the fastest with a maximum transfer rate of 540K transfers/second. It was not necessary to use this method as the FHS is quite adequate.

Although this interface does have a reasonable degree of flexibility it is far surpassed by that of the series 80 interface used to control the SPCATS insertion unit.

3.4 ANALYSIS

The software of the receiver computer implements a digital filter, fast-fourier transform and various search algorithms for measuring harmonic components.

3.4.1 The Fast Fourier Transform

The Fast Fourier Transform (FFT) is the basic starting point for almost all modern digital signal processing. It is so significant that special purpose FFT chips have come on to the market in recent years to drastically reduce the time taken to compute the transform using conventional software techniques.

The fast fourier transform is simply an efficient algorithm for the computation of the discrete fourier transform (DFT). The DFT is, in turn, the discrete approximation to the continuous fourier transform. Because of this there are some important factors which need to be borne in mind when applying the DFT to a sampled, continuous waveform, this being the major application of the transform.

The DFT can be expressed as

$$X(mF) = \sum_{n=0}^{N-1} x(nT)\, e^{-j(2\pi/N)mn}$$

where: T = time domain sampling period
= $1/f_s$
F = frequency resolution of spectrum
m = m'th frequency sample
n = n'th waveform time sample
N = total number of samples taken
= 1/FT

This being an approximation to

$$X(f) = \int_{-\infty}^{\infty} x(t)\, e^{-j2\pi ft} dt \qquad \text{for } x(t) \text{ non-periodic}$$

and

$$X(mF) = 1/t_p \int_{t_p} x(t)\, e^{-j2\pi Ft} dt \qquad \text{for } x(t) \text{ periodic}$$

where: t_p = period of x(t)
= 1/F

It can be shown that the sampling of a function in one domain implies periodicity in the other; either inherent or forced periodicity [27]. Hence, as a continuous waveform is sampled in practice and the numerical DFT applied there is sampling in both domains and hence each domain assumes periodicity in the other. The period in the time domain is the time t_p over which the N samples are taken and in the frequency domain, f_s, the sampling frequency. In each domain there are N sample points in one period. In the frequency domain this periodicity does not cause difficulties and simply means that the spectrum repeats along the frequency axis.

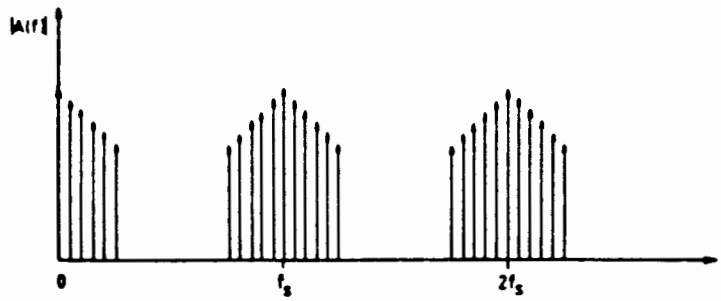

FIGURE 3.10 Image-spectra due to sampling

Over one period in the frequency domain the DFT can be thought of as a parallel bank of N filters with equal bandwidths spanning the range 0 to f_S Hz, where the output of each filter represents a spectral line at its corresponding point in the discrete DFT frequency axis. Each filter has a finite bandwidth and hence finite resolution. For a particular filter, any frequency falling within its band will be said to occur at the filter's center frequency. This is precisely what occurs at each of the N frequency points of the discrete spectrum. As will be shown later, the shape of each of these N 'filters' is determined by a window function applied to the N real-time samples. Hence the -3 dB point of such a filter is important for accurately determining the amplitude of a frequency component that does not happen to fall precisely on the center frequency.

The periodicity requirement in the time domain does however cause problems which require careful attention. As a finite time record of samples is assumed periodic by the DFT problems arise due to discontinuities at the time record ends giving rise to smearing of the spectrum. Consider the continuous sinusoid, sampled over a finite time record shown below

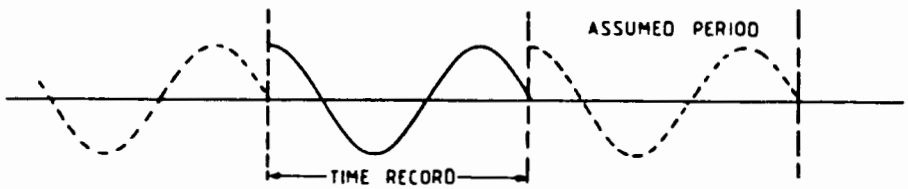

FIGURE 3.11 Assumed periodicity of time record

To alleviate this problem a window function is applied to the time record so that samples gradually taper to zero at the record ends and hence discontinuities are not generated [28].

FIGURE 3.12 Windowed time record

a) Considerations in using the DFT

The DFT is used in the SPCATS system to approximate the Fourier transforms of the continuous-time tone-bursts. The three properties that must therefore be borne in mind in its application are as follows.

i) Aliasing:

This is well known and results from a signal sampling frequency not satisfying the sampling theorem, which states that the sampling frequency of a system must be greater than twice the input signal bandwidth if such a signal is to be accurately reproduced. If this is not satisfied overlap of the image-spectra (or 'aliasing') occurs.

ii) Leakage:

This has been discussed and is the term used for the distortion of the spectrum due to the finite time over which the signal is sampled. Time-gating is effectively the same as multiplying the continuous signal by a rectangular pulse of unity amplitude. Note that there is a special case where leakage does not occur. This is where the sampled signal is exactly periodic in the time record. Cascading M (where $M \rightarrow \infty$) such records will correctly give the original signal over these M records. In other words, the signal has no discontinuity at the interface between any two records and as the DFT assumes the signal to have a period equal to the record length, no leakage will occur. However, in practice this condition is virtually impossible to enforce as the system sampling

rate is usually fixed and the input signals one wishes to measure are not necessarily harmonically related to this sampling frequency.

iii) Picket-fence effect:

This effect is produced by the inability of the DFT to observe the spectrum as a continuous function, since computation of the spectrum is limited to integer multiples of the DFT resolution, F. The exact behavior of the fourier transform can therefore only be seen at discrete points. There are two ways of reducing this effect:

a) Padding - Here zeroes are added to the end of the time record, thereby effectively making it longer and increasing the DFT resolution (i.e. F). This therefore changes the frequencies at which the discrete spectrum points occur making it possible to align frequency components with these points. The problem with this is that the DFT now computes more points, thereby non-linearly increasing the computation time. In addition an FFT routine is required that can cater for variable length input sequences while the lengths of sequences which can be handled by efficient algorithms are limited to integral powers of two.

b) Windowing - Use a window function which has a flat response over one frequency bin. Now a component falling between any two bins will have its amplitude read correctly at the expense of accurate frequency determination.

The latter approach has been used in the SPCATS system as a window function must be used in any case to solve the leakage problem and as the exact frequency of an incoming tone does not need to be accurately determined, since it is already known. The effect of windows is discussed in detail further on in the text.

b) Algorithms for the FFT

The efficient use of the DFT in numerical machines was made possible by the fast fourier transform algorithm presented by Cooley and Tukey in 1965 [29]. Since then many variants of this have been proposed [30]. Of these one of the most important is that of Sande and Tukey.

All FFT algorithms use special groupings of the DFT terms to reduce the number of complex additions and multiplications, the number required by the DFT being N^2, where N is the number of samples. For the Cooley-Tukey algorithm this is reduced to $N.LOG_2N$. Computation and a memory reduction is possible due to a process called decimation in time (Cooley-Tukey), or decimation in frequency (Sande-Tukey) [27].

An important simplification occurs where the time record only consists of real-valued samples (not complex as in the general case). Here a two-to-one reduction in computation and storage can be obtained by putting the N-term real record in the form of two N/2 separate records, forming an effective N/2 complex record to be used in either of the algorithms above [28,31]. The exact mechanisms of these algorithms can be found in most conventional digital signal processing texts and will therefore not be discussed here [27].

3.4.2 Windowing

The fundamental problem with a finite length record of a sampled signal is the discontinuities at the record ends. The finite time, or 'window', during which the waveform is sampled affects the resulting spectrum. As mentioned before, where the sampled signal is not periodic in the time record smearing of the spectrum occurs. For this reason window functions are used together with almost all FFTs. The time record is multiplied by the window function thereby making the discrete spectrum agree more closely with the continuous one.

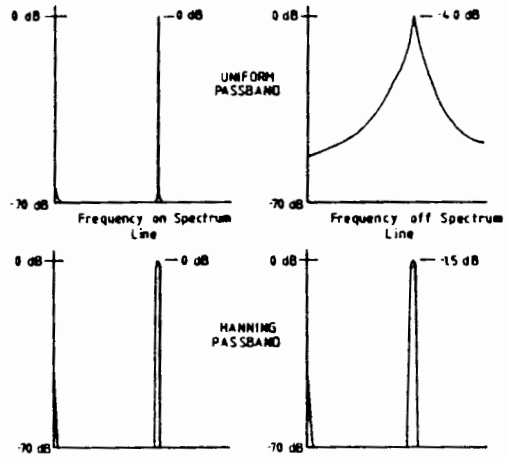

FIGURE 3.13 Effect of windows on resultant spectrum

It is the choice of window function being critical in many applications which has led to much investigation by researchers.

In the frequency domain the window's fourier transform is convolved with the spectrum of the input signal. Hence, instead of obtaining dirac impulses at the frequency positions in the spectrum where the signal's harmonic components occur, there is a widening of these impulses due to the convolution. This leads to the problem of resolving two frequencies close together, but of greatly differing amplitudes, as the smaller one is concealed by the larger due to the window function.

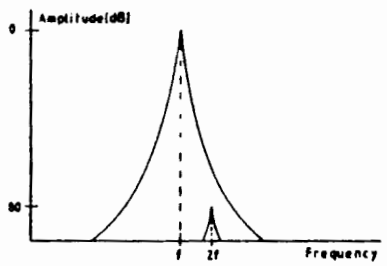

FIGURE 3.14 Effect of window on nearby harmonic of fundamental

This limited ability of most window functions to resolve a small harmonic close to the fundamental was the major problem with the analysis section of the SPCATS system. We were not simply at liberty to increase the resolution and thereby separate the harmonic components in the spectrum of the received signal as this would necessitate a longer tone-burst

due to a longer time record being required. Hence a window function which rapidly drops off to its stop-band with no side-lobes above the -80 dB specification had to be found. In addition the window had to be maximally flat over one FFT bin.

To illustrate the problem consider the following:
The time record that has no window function applied to it has an implicit rectangular window. As the fourier transform of a rectangular pulse is a sin(x)/x function we get smeared harmonic components, with the shape of the window's transform, when the spectra are convolved.

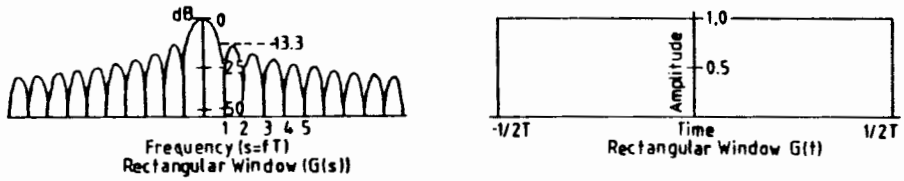

FIGURE 3.15 Rectangular window

In this case, to read a harmonic which is adjacent to a fundamental, and 80 dB smaller, the harmonic must be a considerable number of FFT frequency bins away. To effect this it is necessary to increase the FFT resolution and hence the time record length until the smearing does not conceal the smaller component.

Now consider the same problem above, but using a window function which is implemented on the majority of commercial spectrum analysers. This is the Hanning window which has improved sidelobes while maintaining considerable frequency accuracy.

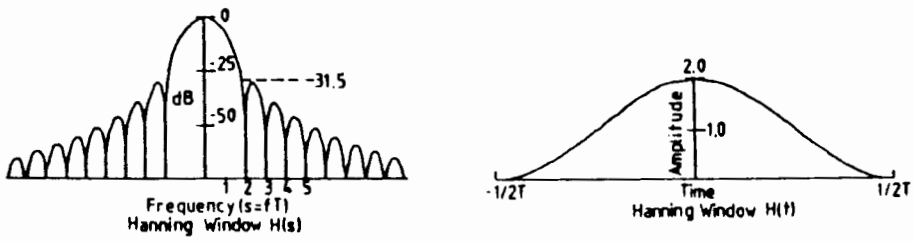

FIGURE 3.16 Hanning window

By frequency accuracy is meant the ability to accurately determine in which FFT bin a particular component falls. For

this window the first sidelobe is -31.5 dB down on the center frequency, whereas in the case of the rectangular window it is -13.3 dB. Hence our -80 dB harmonic can be measured closer in to the fundamental with the Hanning window. For this window a shorter test-tone can therefore be used as less frequency resolution is required. This window is not used in the test-set analysis section as the tone lengths would still be unacceptably long. As amplitude accuracy is important in the system and not that of frequency, a window function is used that is based on that implemented in a commercial spectrum analyser. Here the -80 dB harmonic need only be six FFT bins away to be read [31].

Note that in decreasing the peak sidelobe amplitude, the width of the center lobe increases and this is why it becomes more difficult to resolve fine frequency detail. It is this increase in the main lobe width that leads to the flatness over one FFT bin. The actual window used on the SPCATS will be discussed in detail later.

3.4.3 Digital Filtering

A digital filter is a computational process by which a sequence of numbers, usually samples of an analog signal, is transformed into a second sequence of numbers. This computational process may correspond to high-pass, low-pass, band-pass, or band-stop filtering, integration, differentiation, or any other function.

Digital filters have many advantages over their analog counterparts:

- Changes resulting from variations in component values are non-existent.
- Periodic calibration is not required.
- The performance from unit to unit is stable and repeatable.
- Flexibility is available since filter responses can be altered by simply changing arithmetic coefficients.
- Arbitrarily high precision can be achieved, limited only by the word-length used by the algorithm.

a) The requirement

In the analysis section of the SPCATS system it is necessary to vary the FFT resolution in accordance with the frequencies of the test tones received. The low frequency tones (40 Hz and 125 Hz) require greater resolution to extract their harmonics than the higher frequency tones. The FFT's resolution is fundamentally determined by the time record length (T), the sampling frequency (f_s) and the number of samples in the record where $f_s = N/T$. The frequency resolution is then

$$F = 1/T = f_s/N$$

Hence there are various ways of increasing the resolution:

i) Lengthen the time record (take more samples thereby increasing N). This has been shown to be an undesirable solution.

ii) Reduce the sampling frequency. This improves the resolution in the same ratio as the frequency decrease and would be an ideal solution were it not for the cut-off frequency of the anti-aliasing filter also having to decrease with the sampling rate so as to avoid aliasing.

iii) The solution to the problem is to lower the sample rate in software. Here the sampled waveform is digitally low-pass filtered and then samples discarded at regular intervals. This process is known as sample rate decimation.

Decimation allows the hardware sampling rate and the anti-alias filter cutoff frequency to remain constant. This is the method adopted in the majority of modern spectrum analysers [32,33]. An additional factor here is that advantage is taken of the benefits of digital over analog filters mentioned above.

As a example of decimation to increase the resolution by a factor of five, we would pass the output of the ADC through the digital filter, filtering the data to one-fifth the

original bandwidth, and then discard four out of every five samples.

Digital filters suffer from problems to maximise noise performance and dynamic range in a similar way to their analog counterparts. The problems in the former stem from rounding errors due to finite word lengths. The most important are:

i) Overflow - This is analogous to overload in an analog filter. Here a sample value in the filter becomes too large for the finite word length and thereby causes distortion in the output. In addition, an overflow can cause a digital filter with feedback to go into an oscillatory mode that is sustained by repeated overflows. This is called 'overflow oscillation'.

ii) 'Self noise' - Rounding or truncation of intermediate results to conform to the system's word length generates noise within the filter wherever this occurs.

iii) 'Limit cycle' oscillations - This phenomenon occurs in many digital signal processing systems where there is a rounding operation within a feedback loop. A non-zero output occurs even though the input is zero due to truncation and rounding of arithmetic results. These non-zero outputs can either be a constant value (termed a 'deadband') or oscillatory, this determined by the number of feedback terms in the filter's 'difference equation' [32]. One approach to solving this problem is to increase the word length so that the limit cycles are reduced to an acceptable level. However, this is not attractive in applications where the word length needs to be kept to a minimum. The standard approach is to add a small amount of 'dither' signal into the filter nodes where rounding or truncation occurs. The dither variance is sufficient to break up the limit cycles but its amplitude low enough to increase the noise floor only slightly. This is analogous to dither added to ADC systems to uncorrelate noise generated by low-level, narrow-band signals.

The above anomalies are important design considerations in systems requiring short word lengths. With the SPCATS receiver's computer these problems do not occur due to real arithmetic (53 bit mantissa) being used for the filter. Where a system accuracy of -80 dB is required a word length of about 20 bits or less is unacceptable [32] (for IIR type filters, these being used in the majority of commercial spectrum analysers [32, 33]). For this reason it was not possible to use integer arithmetic for the filter although this would have the attraction of speeding up the algorithm considerably.

b) Filter types

At this point a brief review of the possible digital filter types will be made with emphasis on that adopted in the system's decimation filter.

A digital filter is designed by determining the coefficients of the input-output algorithm by some approximation technique. The input-output function of the filter is known as its 'difference equation' and can be expressed as

$$y(n) = \sum_{i=0}^{k} a_i x(n-i) - \sum_{i=1}^{k} b_i y(n-i)$$

This leads to the discrete transfer function in the z-domain of

$$H(z) = N(z)/D(z) = \sum_{i=0}^{k} a_i z^{-i} / (1 + \sum_{i=1}^{k} b_i z^{-i})$$

Digital filters are classified by the duration of their impulse responses and according to the type of realisation adopted. The former leads to two important classes:

i) Infinite impulse response (IIR): Here the impulse response, h(n), has an infinite number of samples.
ii) Finite impulse response (FIR): The FIR impulse response has a finite number of samples.

The possible filter realisations are:

The possible filter realisations are:

i) Recursive realisation: With this realisation the present output depends on the input as well as previous values of the input and output. There is therefore feedback present. The difference equation has both a_i and b_i terms which are non-zero.

ii) Nonrecursive realisation: Here the output depends only on the present and previous input values. The b_i terms of the difference equation are therefore all zero.

iii) Fast fourier transform realisation: The input signal is transformed by the FFT, filtered in the frequency domain and then transformed back to the time domain.

Theoretically the two filter classes can be implemented with any of the filter realisations. However, in practice an IIR filter is easier to implement recursively and a FIR as a nonrecursive or FFT realisation. We therefore associate these filter types with such realisations. Some of the important differences between IIR and FIR filters are discussed below.

i) Errors due to quantisation, rounding, and coefficient inaccuracies are far more problematic in IIR filters. This is due to the feedback present in an IIR realisation. This necessitates long word lengths for the coefficients of narrow band IIR filters.

ii) A FIR filter is always stable as all its poles are at the origin.

iii) Linear phase is easily obtainable in the FIR case but not so for the IIR.

iv) The FIR filter requires a considerably higher order (number of coefficients) than the IIR filter for the same roll-off and Q.

v) Determining the coefficients of the FIR filter is considerably more difficult than for the IIR realisation where relatively straightforward design procedures are available. The only practical way of designing FIR filters is by means of specialised computer programs [34].

vi) An advantage that FIR decimation filters have over their IIR counterparts is that output samples need only be computed at the lower output frequency whereas in the IIR

case an output sample must be computed for each input even though several of these samples are discarded every output cycle [35].

vii) A FIR filter requires a lot more computer storage due to its large number of coefficients in comparison to an equivalent IIR implementation.

Due to there being no word length problems (53 bits) with the SPCATS filter we were free to choose either an IIR or FIR type. Noise performance and stability were therefore not criteria used in making this decision due to the more than adequate number of bits available. Linear phase was not a requirement either.

Possibly the most important criterion was the availability of techniques to compute the coefficients. The direct procedures available for IIR implementations were felt more desirable than the computer techniques usually used for FIR filters, software for the latter not being readily available. In addition, FIR filters require considerably more multiplications and additions per output sample due to the large number of coefficients required to get the same response as an IIR realisation. This is due to the advantage of feedback being lost. The computation time saved in only having to calculate output samples in a FIR realisation at the final decimation output rate, instead of at the input sample-rate, did not outweigh the disadvantages mentioned above. Therefore an IIR realisation has been used for the SPCATS decimation filter.

c) Determining filter coefficients

There are three basic techniques for determining IIR filter coefficients. These are the bilinear transform, the step-invariance method and the impulse-invariance method. Of all the techniques available for implementing a desired filter digitally, the bilinear transform is the most well formulated and widely used.

The bilinear transform is a direct mapping from the analog s-plane to the digital z-plane where

$$s = \frac{C(1 - z^{-1})}{(1 + z^{-1})}$$

and C is the mapping constant between the two planes. It can be shown that the behavior of the analog function over the infinite frequency range is mapped into the range from D.C. to f_o in the digital filter. This effect, known as 'warping', causes noticeable compression of the original analog transfer function in the digital implementation which is an advantage in a low-pass filter due to the improved roll-off obtained.

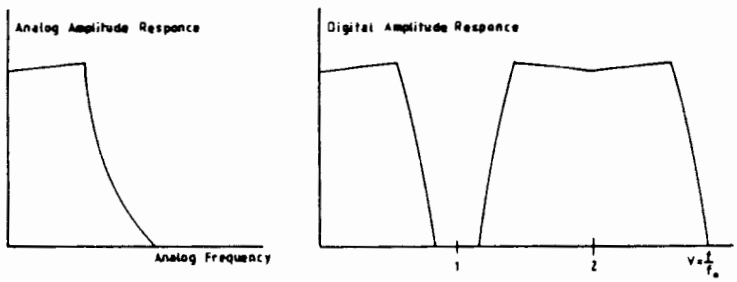

FIGURE 3.17 equivalent response in analog and digital domains

This is due to the mapping constant C being expressed as

$$\tau = C.TAN(\pi/2(f/f_o))$$

where: τ = frequency in the analog domain
f = frequency in the digital domain
f_o = foldover frequency

Plotting τ versus v, where $v=f/f_o$, shows clearly why warping occurs.

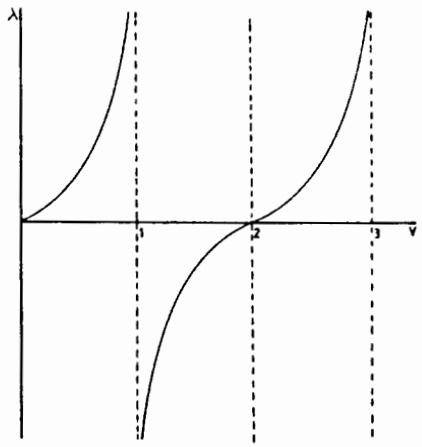

FIGURE 3.18 Relationship between frequency variables of the analog and digital functions

In practice it is not necessary to go through the whole process of computing the digital coefficients via the transform as tables are available expressing the digital coefficients in terms of the analog coefficients where the respective transfer functions are of the form

$$G(s) = \frac{A_0 + A_1 s + A_2 s^2 + \ldots\ldots + A_k s^k}{B_0 + B_1 s + B_2 s^2 + \ldots\ldots + B_k s^k}$$

for the reference analog function, and

$$H(z) = \frac{a_0 + a_1 z^{-1} + a_2 z^{-2} + \ldots\ldots + a_k z^{-k}}{1 + b_1 z^{-1} + b_2 z^{-2} + \ldots\ldots + b_k z^{-k}}$$

being the corresponding digital domain function.

The precise details of this technique and those mentioned briefly below can be found in any digital signal processing text [27].

The impulse-invariance method uses the criterion that the impulse response of the digital filter be the same as the corresponding impulse response of the reference analog filter at the sampling points. The step-invariance method on the other hand sets the step response of the digital filter equal to that of the corresponding step response of the reference analog filter at the sampling points.

Techniques for implementing FIR filters are all based on the fourier series method. This method uses the fact that the digital filter amplitude response A(f) is a periodic function of frequency and can therefore be expressed as a fourier series in the frequency domain. Via the impulse response the FIR transfer function can then be obtained. There is an approximation that arises here in that only a finite number of the fourier series terms are used for the filter. This leads to poor filter responses which do not agree with that desired. This problem is similar to that of the fast fourier transform due to finite time records where the input signal is effectively multiplied by a rectangular window. In the filter case, the fourier series is effectively multiplied by such a window. The various window functions used in FFT analysis are therefore also used with FIR filters to compensate for their finite number of coefficients. Note that with the FFT the window function is applied in the time domain whereas with the filter it is used in the frequency domain.

3.4.4 Data Presentation

It is desirable to present concise, complete reports after each test sequence. As several tests may be run each day unnecessary detail would prove tedious to the User. Results are printed at the end of each test and the following information included:

i) Channel over which test performed.

ii) Date and time of test - If meaningful conclusions are to be drawn from the data this information is essential. It could, for example, pinpoint errors on a certain section of the network which is only switched in as part of the channel at certain times of the day.

iii) All results of a particular test type are grouped together - For example, all frequency response measurements are printed in succession.

iv) The frequency and test level are included with each result.

v) Each test result is displayed in conventional units - Frequency response results are therefore in decibels, static intermodulation distortion as a percentage, and so on. This makes the results more meaningful to the User without conversion to familiar units being necessary.

CHAPTER 4

AUDIO-QUALITY TESTS

4.1 INTRODUCTION

An overview is given here of the main analog distortion types and the tests implemented on the SPCATS system to measure these.

Distortion types are generally classified as being linear or non-linear. Various tests have appeared over the years to quantify these and these have been used so as to provide results which can easily be compared with those of other audio systems.

4.2 LINEAR DISTORTION

Linear distortion causes changes in the original phase and magnitude relationships between individual frequency components. Note that linear distortion does not occur if the amplitude of each component is changed by a fixed amount, or the whole signal is shifted in time by a fixed amount. These two variations may occur and do not comprise distortion as the shape of the input signal is still maintained.

There are two fundamental consequences of linearity:

i) Superposition applies - This means that a composite signal is always the sum of the individual, but simultaneously processed signals.

ii) The magnitude scale factor between the input and output is preserved - Thus, if the input signal is doubled, so will the output.

The standard way of quantifying the linearity of a system is by means of magnitude and phase plots, against frequency. However, to obtain more information about a system's performance, tests for actual non-linearity are also carried out.

In practice it is only possible to design a system to be linear over a certain range of input amplitudes and frequencies. It is for this reason that the linear tests mentioned above should be performed within the correct bandwidth and amplitude range if meaningful results are to be obtained. Knowing that a system is linear over a certain range, the effect of non-linearity outside this range can be determined by means of non-linear tests at specific relatively large signal levels. By performing the same tests in the linear region of the system a measure can be made of its deviation from an ideal linear system.

4.3 NON-LINEAR DISTORTION

This type of distortion is responsible for the generation, by the audio system, of harmonic, intermodulation and other 'new' frequency components when a signal is applied.

These new frequency components combine with the original input to produce an output that is not a true reproduction of the input. Whereas non-linear distortion generates unwanted frequencies, linear distortion does not.

Non-linear distortion may be static or dynamic in character:

- One of the most important forms of amplitude non-linearity is due to clipping of the input signal. This generates harmonics which are derived from the spectrum of the clipped portion of the waveform. The harmonic content of the output spectrum then gives an indication of the type of clipping involved [36]. With all non-linear distortion a knowledge of the harmonic content of the distortion will give an idea of its cause [36,37]. For example, in the case of clipping the following deductions can be made:
  i) If symmetrical clipping occurs, the resulting nonlinear distortion consists of only odd harmonics.
  ii) Unsymmetrical clipping gives rise to a d.c. term as well as even and odd harmonics.

- The rise time of an input signal being greater than that of the system gives rise to frequency dependent non-linearity. Here dynamic intermodulation products are generated.

4.4 MEASUREMENT OF DISTORTION

The tests used by SPCATS to quantify the various distortion types are laid out below. Additional non-linear tests have been implemented as redundant testing enables the particular distortion mechanisms present in a system to be identified.

4.4.1 Linear Tests

'Frequency response' and 'insertion gain' measurements are made. The frequency response is done at spot frequencies across the audio band whereas insertion gain is a measure of the difference between the amplitudes of a 1 KHz tone at the transmit and receive ends of the channel.

4.4.2 Non-linear Tests

The three generally accepted non-linear measurements, 'total harmonic', 'static intermodulation' and 'transient intermodulation' distortion are included in the SPCATS set of tests. In addition a measure of 'gain compression' is also made.

a) Gain Compression

Here a 1 KHz tone is sent at a low amplitude and thereafter at a high amplitude. The amplitude differences between the two tones are measured at the transmit and receive ends. The difference between these two values then gives the gain compression due to static non-linearity of the system. This test is therefore equivalent to that of total harmonic distortion except that the result is expressed in a different form.

b) Total Harmonic Distortion

Here a single tone of frequency 'f' and amplitude 'A' is applied as the test signal. Depending on where the system gain is zero the output signal will contain a d.c. term and terms at multiple values of the input frequency. The percentage harmonic distortion is then given as:

$$THD(\%) = [(E_2^2 + E_3^2 + ...) / E_1]^{\frac{1}{2}} \times 100$$

where: E_1 = amplitude of the fundamental
E_2, E_3 ... = amplitudes of the harmonics

Because of the nature of this test, a signal source having an extremely low distortion content must be used. However, the measurement accuracy specification for the system has been met with the digital signal generator. Note that this test only measures static distortion.

One of the disadvantages of THD measurements is that they can only be considered meaningful when the testing frequency is less than about one-third of the upper limit of the bandwidth of the system being measured. Measurements outside this band will be false as the generated harmonics will have been filtered out.

c) Static Intermodulation Distortion

The I.E.C. specification of the original CCIFF intermodulation test has been implemented [15,38].

The method uses two closely spaced frequencies, f_1 and f_2, of equal amplitudes and measures their difference frequency products. Depending on the test frequencies this measures varying degrees of static as well as dynamic intermodulation. At high frequencies the rise time of the test signals may become significant and thus generate slew-rate limited products.

The I.E.C. refers to static intermodulation as 'difference frequency distortion'. There are three types defined; second

and third order as well as total difference frequency distortion. These are computed from:

Second-order:

$$D_2 = \frac{V(f_2-f_1) + V(f_2+f_1)}{V(f_1) + V(f_2)}$$

Third-order:

$$D_3 = \frac{V(2f_2-f_1) + V(2f_1-f_2)}{V(f_1) + V(f_2)}$$

Total:

$$D = \frac{(V_2^2 + V_3^2)^{\frac{1}{2}}}{V(f_1) + V(f_2)}$$

where: f_1, f_2 = the respective frequencies of the tones
$V(...)$ = amplitude of specified frequency component
$V_2 = V(f_2-f_1) + V(f_2+f_1)$
$V_3 = V(2f_2-f_1) + V(2f_1-f_2)$

The I.E.C. has specified f_1 and f_2 as 8 KHz and 11.95 KHz respectively. However, the SPCATS specification called for 5.6 KHz and 7.2 KHz [14]. Using these frequencies, results predominantly indicate static non-linearity.

The I.E.C. specify second and third order difference frequency distortion, as well as total, as these give an indication of the order of the curvature of the system. When the non-linearity is due to even-order curvature, the major intermodulation products are $f_1 \pm f_2$. With odd-order curvature there are none of these with the main products being at $2f_1 \pm f_2$ and $2f_2 \pm f_1$. Depending on the positions of the fundamentals, not all the above products will fall within the band. For example; using the I.E.C. frequencies the product $f_1 + f_2$ does not do so whereas it does with the actual frequencies used. This product has therefore been included in the second-order difference frequency calculation used.

The primary advantage of the difference frequency test is that its two tones do not need to have a high harmonic-free content, compared with that of the THD test tones. In addition the test-signal better approximates real program material.

This test is most useful at the upper end of the system band where the THD measurement cannot be used due to the roll-off.

d) Transient Intermodulation Distortion

The test for dynamic distortion was originally devised to explain the disagreement between conventional tests, which would indicate a high quality audio system when actual listening tests revealed unacceptable sound reproduction. It was felt that dynamic distortion was responsible for this as it depends on the frequency characteristics of the system, whereas the conventional tests are primarily amplitude related.

An indicator of dynamic distortion is transient intermodulation (TIM). TIM is caused by the rise time of an input signal exceeding that of the amplifier. The large-signal bandwidth of a system determines how fast it can respond to a step function, or in practice, a transient with a fast rise time. The response time of a system to a unit step input will shorten as the cut-off frequency is raised. In the limit, a system with infinite bandwidth will respond to a dirac impulse [36]. Dynamic distortion testing is a method of quantifying this property of a system.

It is therefore good engineering practice, in any type of distortion test, to use signals whose rise times are consistent with the system's bandwidth. For the audio bandwidth, where the -3 dB frequency is 15 KHz, the maximum theoretical rise time which can be reproduced linearly is 30 volts/micro-second [36].

The test signal defined for transient intermodulation distortion measurements consists of a high frequency sinusoid, 15 KHz, and a square-wave of four times the sine-wave's

amplitude. The specific frequency values have been chosen so that the optimum separation between the sine-wave, square-wave harmonics and the intermodulation products is obtained [37].

A 30 KHz single pole, low-pass filter is used to limit the rise time of the square-wave. As was discussed above, this is required so as not to to drive the system with frequencies above its bandwidth. Were the square-wave limited to 15 KHz there would theoretically be no TIM present. Hence a value of 30 KHz is chosen to obtain a standardised measure. For high quality equipment with a bandwidth well above 15 KHz, a 100 KHz filter is the standard used [15].

The amount of transient intermodulation is defined as:

$$TIM(\%) = [(V_1^2 + ... + V_9^2)^{1/2}/V_s] \times 100$$

where the V_i are the voltage levels corresponding to the following frequencies:

Vi	Component	Frequency[KHz]
V_5	$f_s - 5f_q$	0.90
V_4	$f_s - 4f_q$	2.28
V_6	$f_s - 6f_q$	4.08
V_3	$f_s - 3f_q$	5.46
V_2	$f_s - 2f_q$	8.64
V_8	$f_s - 8f_q$	10.44
V_1	$f_s - f_q$	11.82
V_9	$f_s - 9f_q$	13.62
V_s	f_s	15.00

and f_q and f_s are the frequencies of the square and sine waves respectively.

The filtered square-wave exercises the system over a large dynamic range. If during the transition the system is slew-rate limited, transient gain compression will occur. The resulting modulation of incremental gain will modulate the smaller sinusiodal signal present, causing sidebands to appear about the harmonics of the larger signal. The magnitude of these harmonics is an indicator of the presence of TIM.

It is important to realise that the test as described measures total intermodulation distortion i.e. dynamic as well as static. By replacing the square-wave with a triangular-wave, of equal amplitude and frequency, the static intermodulation alone can be measured [39].

The magnitude of the intermodulation products of the square-sine test can be used to identify the limiting mechanism of the test circuit. It can be concluded as to whether symmetrical, unsymmetrical, abrupt or gradual slew-rate limiting is occurring [37].

CHAPTER 5

INITIAL INVESTIGATIONS

5.1 INTRODUCTION

An experimental test-set was designed and built before the final prototype. This was to test the system's concept and solve the problems that are bound to arise with any prototype equipment. Much of this initial experimental circuitry was designed in such a form that it could be incorporated with minimum modification into the final prototype, in order to minimise unnecessary redesign. The insertion unit did not change whereas the receiver section was redesigned using circuit modules that had been tried and tested in the experimental set. The receiver software also recquired modification for use with the final system.

The availability of an HP3561A Dynamic Signal Analyser made the initial investigations possible. Without such an instrument it is doubtful whether the project would have been possible at all. By means of the experimental test-set it was possible to iron out most of the problems before actually building the prototype thereby keeping changes to this minimal.

Testing of the experimental set was done in a manner as close as possible to the circumstances under which the final set would be used. In this way it was possible to pinpoint any errors particular to the modus operandi of the system.

5.2 DIGITAL SIGNAL GENERATOR

The insertion unit, consisting of the signal generator together with its gating and timing circuitry, was built in its final form. As this has been covered in some detail in earlier parts of this text the software error correction technique used will be discussed here together with the spectral results obtained. In addition the precise signal-to-

quantisation noise ratio (SNR) will be calculated for each test-tone type.

5.2.1 Software Error Correction

The digital signal generator was covered in chapter 3. By ensuring that the voltage to the DACs is stable it is possible to correct the most significant byte's DAC and only be left with the inaccuracy of that due to the least significant DAC. The automatic software calibration feature is implemented as follows:

The most significant byte (MSB) DAC is stepped through all its possible states (0 -> 255) and the output read with a precision voltmeter. The least significant byte (LSB) DAC is set at 128 and held constant, this being its theoretical zero volts output. The deviation of this value from zero simply gives a constant d.c. offset which is less than the required 13 bit accuracy and can therefore be ignored. With the above procedure a file of actual output voltages of the MSB DAC is obtained with an accuracy equal to that of the voltmeter. This must therefore be equal to or greater than the desired 13 bits of the signal generator. The inaccuracies of the LSB DAC are ignored and it is assumed to have a perfect linear characteristic. This assumption would not be possible were a higher accuracy required. Readings of this DAC's output for a zero and full-scale input are only taken and values between these two found by linear interpolation. In this way any desired output voltage can be obtained, accurate to 13 bits, as the sum of a MSB voltage, obtained from the file look-up table, and an interpolated LSB voltage.

The above procedure was automated by using the HP85 which controls the insertion unit to step the MSB while keeping the LSB constant. By connecting a HP precision voltmeter (HP3497A) to the computer as well it was easy to read its value after each increment into a file on the computer. The general purpose I/O interface (GPIO) was connected to the generator and the Hewlett Packard instrumentation bus (HPIB) to the voltmeter. The HPIB is HP's implementation of the IEEE-488

general-purpose instrumentation bus (GPIB). The calibration set-up is shown below.

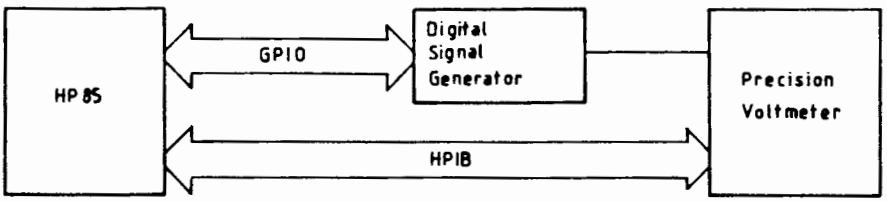

FIGURE 5.1 Error calibration system

Reading output voltages of the MSB DAC and using these together with the LSB to obtain the closest approximation to a continuous analog voltage does not mean that the overlap of 2 bits becomes redundant. This is still necessary in order to ensure that all the required states can be reached.

5.2.2 Signal to Noise Ratios

The signal to quantisation noise ratio (SNR) usually quoted for data converters is that for a full-scale sine-wave. However, as the test tones are not all single sinusoids, the SNRs for the other tone types will be derived here so that these can be compared with actual results obtained. It was shown in chapter 4 that the tones consist of single sine-wave, dual sine-wave and square- plus sine-wave types. The following calculations are only valid for full-scale signals. In this case it can be shown that the RMS quantisation noise is

$$N_Q = \delta v/(12)^{\frac{1}{2}} \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots \quad (5.1)$$

where: δv = the quantisation voltage step

i) Sine-wave: For a full-scale sine-wave

$$y = A.SIN(wt)$$
$$= (2^n \delta v/2)SIN(wt)$$
$$= (2^{n-1}\delta v)SIN(wt)$$

where: n = number of bits

A = half full-scale range of converter

The RMS value is

$$V_{RMS} = 2^{n-1}\delta v/(2)^{\frac{1}{2}} \quad \ldots\ldots\ldots\ldots\ldots\ldots \quad (5.2)$$

Hence the signal to quantisation noise power is

$$SNR = \frac{V_{RMS}}{N_Q} = \frac{2^{n-1}\delta v/(2)^{\frac{1}{2}}}{\delta v/(12)^{\frac{1}{2}}}$$

and

$$SNR_{[dB]} = 6.02n + 1.76 \quad \ldots\ldots\ldots\ldots\ldots \quad (5.3)$$

ii) Dual Sine-waves: Where a composite signal is derived by summing two sinusoids of different frequencies, each has an amplitude of A/2 so that the full-scale range of the converter is never exceeded. The composite signal is therefore

$$y = A/2.SIN(w_1t) + A/2.SIN(w_2t)$$

Although the total power of the signal is the sum of the individual component powers, the RMS value is the RMS sum of the RMS values [40].

$$\begin{aligned} V_{RMS} &= [1/2(A/2)^2 + 1/2(A/2)^2]^{\frac{1}{2}} \\ &= A/2 \\ &= 2^{n-1}\delta v/2 \quad \ldots\ldots\ldots\ldots\ldots\ldots \quad (5.4) \end{aligned}$$

This gives a SNR of

$$SNR = \frac{V_{RMS}}{N_Q} = \frac{2^{n-1}\delta v/2}{\delta v/(12)^{\frac{1}{2}}}$$

and

$$SNR_{[dB]} = 6.02n - 1.25 \quad \ldots\ldots\ldots\ldots\ldots \quad (5.5)$$

iii) Square- plus sine-waves: For a square-wave of (3/4)A and a sinusoid of A/4, the two waveforms

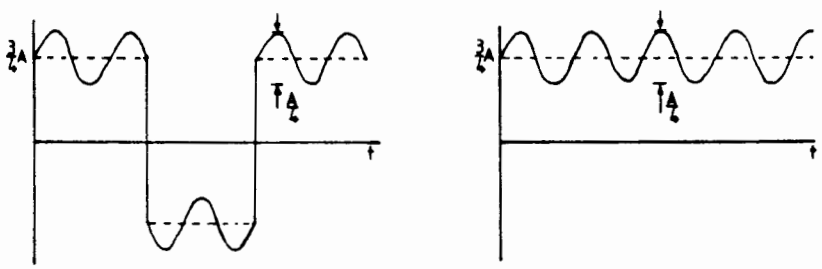

FIGURE 5.2 RMS equivalence of sine/square waveform

have the same RMS value

$$V_{RMS} = [(3/4)A^2 + (1/2)(1/4)A^2]$$
$$= 2^{n-1}\delta v(19/32)^{1/2} \quad \ldots\ldots\ldots\ldots\ldots\ldots \quad (5.6)$$

therefore

$$SNR = \frac{V_{RMS}}{N_Q} = \frac{2^{n-1}\delta v(19/32)^{1/2}}{\delta v/(12)^{1/2}}$$

and

$$SNR_{[dB]} = 6.02n + 1.33$$

Using the above the SNRs for the three tone types for the 14-bit case can be calculated as:

- single sinusoid: 86.0 dB.
- dual sinusoid: 83.0 dB.
- square plus sinusoid: 84.0 dB.

5.2.3 Tests of Signal Generator Performance

The HP3561A dynamic signal analyser made it possible to accurately determine the performance of the signal generator. The analyser has an 80 dB measurement range which is precisely that needed for the system.

The signal generator has, as expected, a noise floor below -80 dB with harmonic distortion below the -70 dB accuracy requirement. Tests were done at all the sine-wave frequencies as well as for the other tone types with all satisfying the above specification.

One result shown up by the tests, which had initially been overlooked, was the inherent SIN(X)/X roll-off of the generator's output characteristic. This results in the high frequency sinusoids having smaller amplitudes than those at the low frequencies even though their numerical peak amplitudes are identical. This effect can be explained as follows:

The numerical values fed into a DAC can be represented as a series of dirac impulses. The DAC however keeps its output constant until it receives the following word. The output

therefore looks identical to that which would be obtained from a sample-and-hold placed at this point. The only difference is that the sample/hold removes the switching transients caused by the bits of the input word changing. However, to avoid the additional errors created by an output sample/hold, latches were placed between the RAMs and the DACs. In this way a new value is presented to the DACs only once the RAMs have settled. The small transients due to the latch outputs changing were found not to be objectionable.

The ideal impulses into the DAC together with its output are shown below. In this discussion we can ignore the transients.

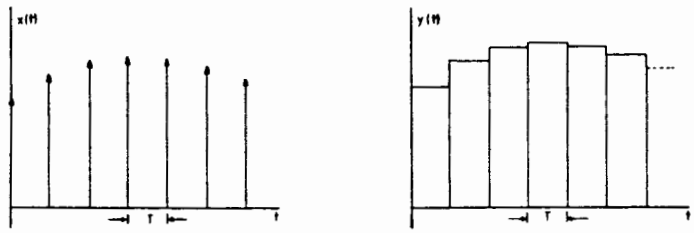

FIGURE 5.3 Operation of zero-order hold

Holding the impulse constant over one sampling period is termed a 'zero-order' hold. The output of the DAC can be considered as the convolution of its input, x(t), with a unity rectangular pulse, p(t).

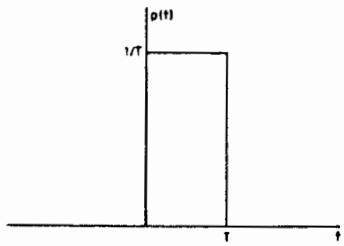

FIGURE 5.4 Unity rectangular pulse

As convolution in the time domain corresponds to multiplication in the frequency domain, the signals spectrum is multiplied by that of p(t). The amplitude spectrum of p(t) can be expressed as

$$A(f) = SIN(\pi fT)/(\pi fT)$$

where: T = sampling period
f = frequency

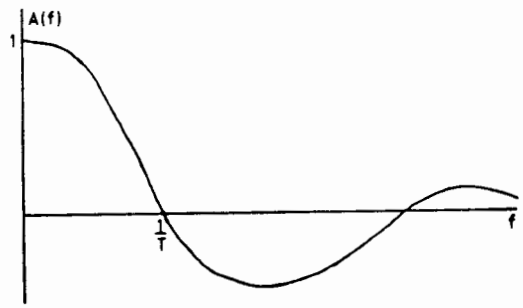

FIGURE 5.5 Pulse spectrum

Consider the case of the SPCATS signal generator where the sampling frequency is 100 KHz and the frequency of interest is 15 KHz.

Now $A(f) = SIN(\pi fT)/(\pi fT)$

$= SIN[15\pi(1/100)]/[15\pi(1/100)]$

$= 0.963$

and

$A(f)_{[dB]} = -0.32$

Relative to d.c. therefore the response at 15 KHz is down by 0.3 dB. In applications where the roll-off is unacceptable the output low-pass filter has a rising characteristic over its passband to compensate for this. This is known as 'aperture correction' [17]. With SPCATS the roll-off can simply be ignored due to the overall system calibration reading the peak amplitude of each test-tone and storing these in a file.

5.3 THE DYNAMIC SIGNAL ANALYSER

Besides its use for doing measurements on the system the HP3561A dynamic signal analyser (DSA) was also used as a data acquisition unit for the experimental investigations. The analyser's buffer of real-time samples can easily be transferred via the HPIB to a stand alone computer for processing. However, in this application we were only concerned with the time-domain samples. The only processing done on these samples is hardware decimation filtering to bring the sampling frequency down from the fixed 256 KHz hardware rate to that specified by the User. In the experimental receiver it was therefore not necessary to implement a digital filter for this purpose.

The ADC used in the analyser will be considered here as this was the element used for data acquisition and as attempts were made, using discrete components, to build a similar device for the final prototype system.

5.3.1 Description and Specification

The HP3561A DSA is one of the latest members of HP's family of FFT analysers. Due to the VLSI circuit technology used the instrument can compute phase and magnitude spectra over the ranges of 125 uHz to 100 KHz which makes it ideal for audio use.

The attraction of dynamic signal analysers in comparison to their swept-frequency counterparts is that the former captures a buffer of samples before actually analysing the data. With the later the signal must remain stationary during the sweep time which limits its applications. It is this characteristic of DSAs which makes it possible to download the time-domain samples.

The instrument has a dynamic range of 80 dB and an amplitude accuracy of ± 0.15 dB over its entire temperature and humidity operating range [41]. A feature that was found extremely useful is its ability to compute the total harmonic distortion of a specified fundamental. This was used extensively in testing the performance of the signal generator. In addition the instrument can be instructed to automatically select the largest spectral component on its display. This is a quick and effective way of selecting the fundamental of a set of harmonics.

The instrument performs an FFT on a record of 1024 samples. After each FFT computation the display is updated and a new record taken and the process repeated. Due to the custom hardware and the TMS320 digital signal processing chip used, a nearly continuously varying display is obtained. This is attractive for using the DSA to tune equipment. In addition to the continuous update feature a time capture mode is available. Here a buffer of 1 to 40 records is taken and each

record can be scrutinised in turn or analysis done on the whole buffer.

Due to the phase characteristics of the anti-alias filter near its roll-off spectra are only displayed up to 0.4 times the sample rate. To be precise, 400 of the 1024 frequency-domain sample points are used.

The requirement that test-tones be as short as possible is limited by the resolution recquired to measure adjacent harmonics. The 'flat-top' window available on the analyser allows detection and measurement of closely spaced spectral components. Using this window the sampling frequency was set at 51.2 KHz (20 KHz display span) for the middle to high frequency tones and 8.0 KHz (3.125 KHz span) for the low frequencies (40, 125 Hz). This set the minimum tone lengths at 20 msecs. and 128 msecs. respectively.

The analyser has a hardware external TTL trigger. On receipt of a trigger pulse the time capture buffer is filled. In our case the buffer consisted of 1 record. The trigger pulse was generated using external hardware which detected the synchronising burst preceding each tone.

5.3.2 Interfacing

The DSA was linked to the receiver's computer (series 9000, model 220) via the HPIB. As it is possible to to set the analyser up with commands sent to it via the HPIB the computer was used to automate the whole system.

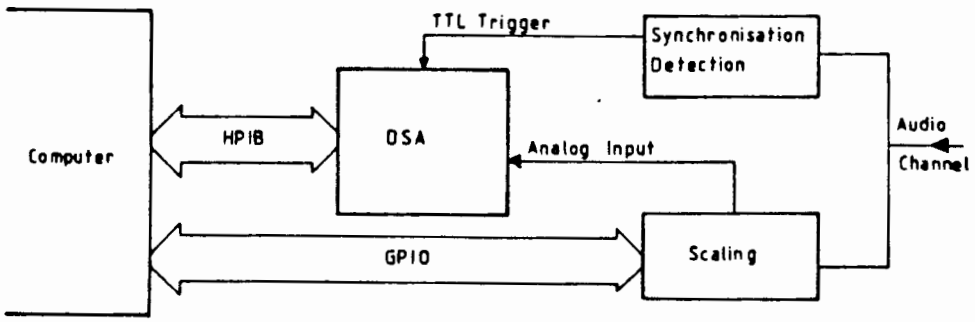

FIGURE 5.6 Investigation configuration

Due to variation in level of the tones being sent from the transmitter end, it was necessary to apply inverse scaling to equalise the amplitudes of the tones into the analyser. This scaling was done in external hardware although the range setting of the DSA could simply have been changed for each tone. This was rather done in hardware, however, as the final system could not take advantage of this variable range. The GPIO (general purpose I/O) interface was used to select the gains. The software for this experimental system configured the DSA before each tone was sent and also selected the correct inverse scaling factor. Then the system would pause, the computer polling the DSA, until the trigger pulse had been received and the buffer of samples taken. The computer would then read the 1K samples and store these on disk for later analysis.

Reading the DSA buffer is not a simple matter of reading 1024 samples as these are preceded by 175 words (16 bit word-length) of internal data [42]. This data contains information about the configuration and status of the instrument. As this information was not relevant in this application it was discarded and only the actual 1024 samples used. The samples of each tone were specifically written as a file to disk to improve modularity. This strategy was also used with the prototype and therefore the software for the analysis did not have to be changed. The second reason for using files is that the analysis of each tone is very slow due to the large amount of signal processing involved. By transmitting all the tones first and then doing the analysis, disturbance of the channel can be confined to about 6 minutes. This makes it possible to do test runs during particular types of program material, e.g. a news bulletin.

The capture of a test-tone sequence can therefore be summarised as:

```
configure DSA for next tone.
poll DSA to detect when trigger has occurred.
read the time buffer discarding the first 175 words.
store the 1024 samples of the time record to disk.
repeat until all the tones received.
```

5.3.3 Analog to Digital Conversion

Due to the front-end of the DSA being used as the analog to digital converter for the experimental test-set it will be considered in more detail here. The analog to digital conversion section consists of a range selector at the input, followed by a nine-pole passive anti-alias filter and finally the ADC. The range selector varies the gain of the input signal so that the input range of the ADC is constant.

The concept used in the design of HP's ADC is similar to that used in the SPCATS signal generator where a converter of lower resolution than that of the composite converter is used.

The technique uses an 8-bit ADC and a 13-bit DAC, performing two successive conversion passes but with the scaling of the input signal changed between passes.

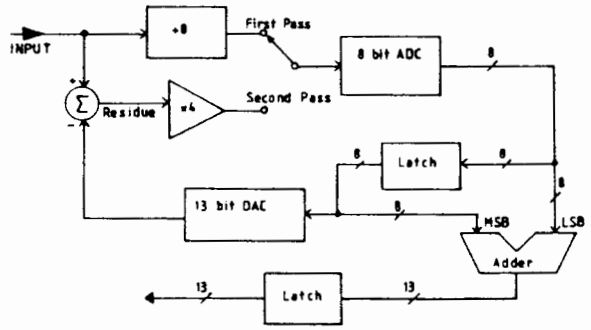

FIGURE 5.7 Residue converter

This residue converter has an overlap of 3 bits as two 8-bit conversions are made, added, and a 13-bit result obtained. As for the digital signal generator the overlap is recquired due to the errors of the 8-bit ADC not only being that due to the quantisation uncertainty (cf. section 3.2.1).

The effect of low amplitude, narrow band input signals on a converter's noise floor was considered for the DACs of the signal generator. Due to such signals causing the quantisation noise to become correlated and harmonics to appear instead of a flat noise floor, dither is applied to uncorrelate the signal and quantisation noise. HP engineers have found however

that the dither is only really effective if it is of the same magnitude as the input signal [43]. However, it is then necessary to subtract the dither signal from the final ADC output so that the noise floor is not increased. In addition the dither signal must be as broadband as that required from the quantisation noise. These recquirements are most effectively met by using a pseudo-random noise source where its analog and digital representations are available and the bandwidth can be set by the sequence length.

Attempts were made to construct a residue converter for the SPCATS system using discrete components. However, insurmountable problems were found in suppressing noise and this idea was therefore shelved. This is probably one of the reasons for HP engineers designing their ADC on a semi-custom integrated circuit.

5.4 SYNCHRONISATION

Synchronisation of the transmitter and receiver systems provided some difficulty due to the synchronisation information arriving at the receiver added to the program material. A method that works no matter what program is present at the synchronisation instant was therefore required. The synchronisation methodology and circuitry developed for the experimental set remained unchanged in the final prototype.

5.4.1 The Problem

Some technique is required to indicate to the receiver that a test sequence has commenced and when to sample the channel so as to input each test-tone. The synchronisation method has to generate a trigger pulse which is accurately defined in time relative to the tone-burst. If the time of occurrence of this pulse varies the length of the test-tone must be increased so that it falls within the capture window. The capture window is the time between the trigger occurring and the 1024 samples having been taken. Consider figure 5.8

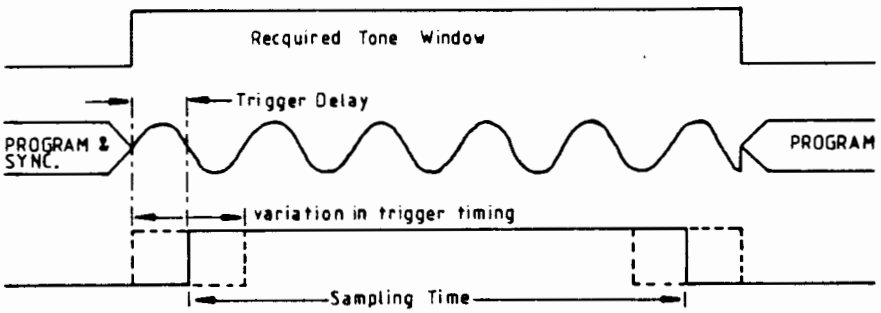

FIGURE 5.8 Synchronisation window

The test-tone must be present over the time taken to sample it as well as for the maximum variation in the trigger timing. A factor which further complicates the problem is the delay between the synchronisation information arriving at the receiver and the time taken for the detection circuitry to generate the trigger.

The other important consideration here is the intrusiveness of the synchronisation technique itself. This has been considered above for the actual test-tone and it has been shown how the technique adopted could increase the required tone length. It is essential to blank out the program material for the duration of the test-tone, this factor being largely responsible for the disturbance caused by the test tones. Hence it is crucial that the tone length be as short as possible, a few milli-seconds making a significant difference. However, where signals are added to the program material the disturbance is markedly reduced, particularly if this occurs near the top end of the audio band. This was confirmed by listening tests done in the laboratory. The length of the disturbance is also not as critical as in the case of the test-tones.

Various techniques could have been adopted for the synchronisation system:

i) Modulation of a low level high frequency carrier:
This technique is used in the Audio-dat system with the carrier at 14 KHz. The advantage of this is that the carrier can be used for synchronisation as well as for data transmission. The problems are that the carrier's

amplitude must be low enough and the data rate low so that the bandwidth is kept small relative to that used for audio (15 KHz).

ii) Pseudo-random spread-spectrum technique:

By synchronising two identical pseudo-random noise sources at the two ends and multiplying the data signal at the transmitter and the composite signal at the receiver with the source, the data signal can be sent across the channel.The data signal can then be a pulse sent immediately ahead of the test-tone. Note that the pseudo-random noise 'carrier' must be continuously present so that the two ends remain synchronised. A short burst of the carrier just before each tone would not allow sufficient time for the two ends to synchronise.

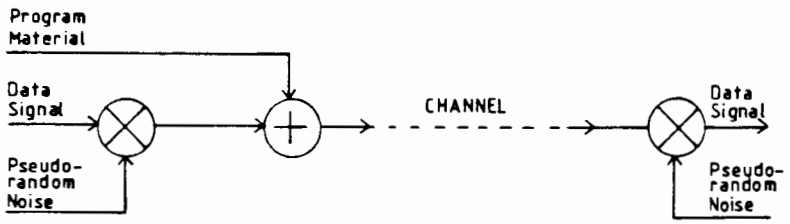

FIGURE 5.9 Spread spectrum synchronisation

The pseudo-random noise source can be designed to have a white noise spectrum over the audio band. Note however that for a narrow-band data signal this additional noise is not objectionable but that as the signal bandwidth is increased so this noise floor rises. To obtain a well defined pulse at the receiver would require a relatively large carrier bandwidth. In addition, getting the two noise sources to synchronise initially significantly increases the circuit complexity required.

iii) Add short signal bursts:

Adding short bursts of a specific signal directly in front of each test-tone helps with the requirement of keeping the trigger relative to the tone accurate with respect to time. The sync-burst can be at a lower level than, and added to, the program material. In addition it can be placed high up in the audio band. In addition its

disturbance is masked by that of the test-tone. This was confirmed by listening tests where it was found that a high frequency burst two to three times longer than the average test-tone was hardly noticeable.

5.4.2 The Solution

A 12.75 KHz sync-burst was decided on, placed directly in front of each test-tone. This frequency was chosen to be as high as possible, but sufficiently far away from the low-level carrier of the Audio-dat (14.5 KHz). Detection was facilitated by the availability of a special purpose chip (XR-2211).

Each test-tone is preceded by a 40 msec. sync-burst and an initial calibration burst is also sent so that the system can configure itself for the first tone. The insertion unit uses a 12.75 KHz oscillator and gating circuitry to:

- add the sync-burst to the program material by gating in he oscillator.
- the oscillator and program material are gated out and the test-tone fed into the channel.
- the test-tone is gated out and the program material transmitted as before.

The synchronisation section of the receiver circuitry consists of the tone decoder together with a 'discriminator' which checks the time for which the trigger pulse is high and low and thereby removes spurious trigger pulses.

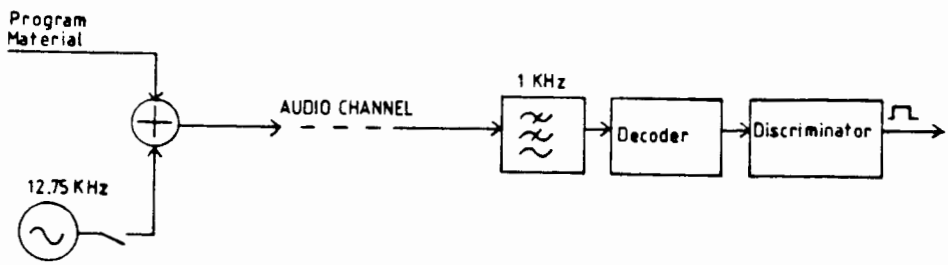

FIGURE 5.10 Synchronisation layout

The high-pass filter simply removes any d.c. from the tone detector's input. Short pulses (less than 5 msecs.) generated at the detector's output due to high energy program material

around 12.75 KHz are removed by the discriminator. Only when its input has been held high for 10 msecs. does its output follow and only after its input has returned low for more than 1 msec. does the discriminator output do likewise. Consider the following diagram:

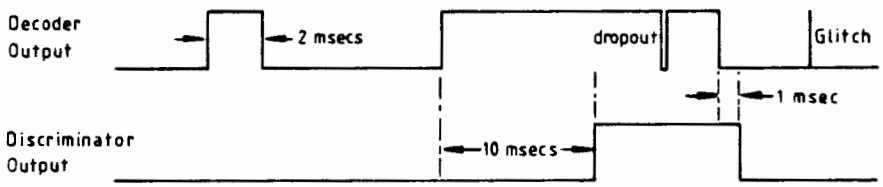

FIGURE 5.11 Discriminator timing

The 1 msec. high-to-low delay had to be kept short as the test-tone length had to be increased by an equivalent amount to compensate for this. This delay was found adequate however as any drop-outs while the decoder output is high are very brief.

The XR-2211 consists of a phase-locked loop and associated logic circuitry to generate pulses when the loop is in lock. The basic functioning of the tone detector is as follows:

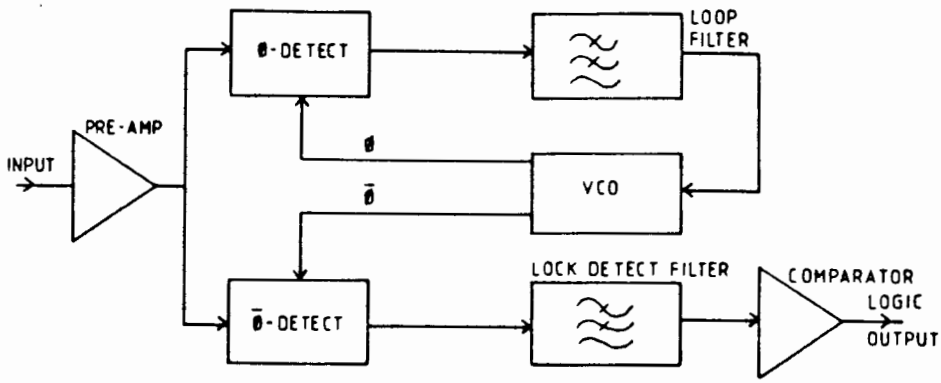

FIGURE 5.12 Tone detector chip

The top phase detector, low-pass filter and VCO form a standard PLL, the phase detector being a standard EXOR type. The bottom phase detector checks for the VCO and input signal being in phase and causes the logic output to change state. The lock detect filter limits the logic response time so that 'chatter' does not occur at the output due to noise and the PLL not actually being in lock.

Using external components the following parameters of the device can be set [44]:

- VCO center frequency: This was set to 12.75 KHz.
- The detection bandwidth ±δf: Here detection bandwidth refers to the pulling range of the VCO. δf = 170 Hz was chosen.
- The loop damping factor and the lock detect time constant: Increasing the former improves the out-of-band signal rejection, but increases the PLL capture time. The latter controls the response time of the lock detection output logic.

The detailed tone detection design is given in appendix E.

5.5 SYSTEM TESTS AND MODIFICATIONS

Extensive tests were done on the experimental system to ensure accuracy and total reliability. A high-quality re-broadcast receiver was used to obtain program material over the full audio bandwidth. This was therefore used as the program source into the insertion unit. This unit and the receiver were then connected with a balanced line.

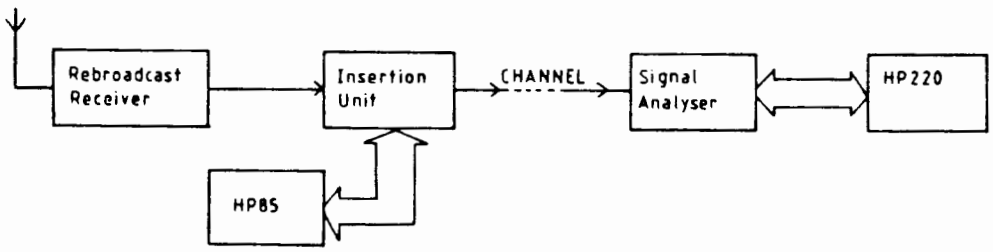

FIGURE 5.13 Test configuration

This configuration was therefore identical to that which would be used in practice except that the channel between the two ends was only a few meters long and not hundreds of kilometers.

5.5.1 Synchronisation

Due to the utmost reliability required from the synchronisation circuitry software was written to extensively test this. A routine was written for the transmitter to send a

sync-burst every 2 seconds and a routine at the receiver checked that each burst was received, keeping running totals of the number recognised and of those missed. This test was usually run overnight.

Initial results were totally unsatisfactory and the whole synchronisation system had to be thoroughly checked. The problem was found to be caused by the input amplifier of the tone detector. This is not in fact a linear amplifier but a limiter so that the loop can lock to input tones of any amplitude greater than 2 mV. RMS. This works fine if only single frequency tones are input. However, the application here has the sync-bursts in the presence of program material with the amplitude of the bursts being less than the average level of the latter. The limiter therefore generates harmonic and intermodulation products, from the program material present, at the synchronising frequency. This either causes spurious triggers to occur or genuine sync-bursts to be ignored due to the harmonics, 180° out of phase, being present and therefore cancelling the sync-burst. The problem with spurious triggers was overcome with the discriminator but not the missed triggers. As the broadcast audio channel itself has a sharp roll-off at 15 KHz, interfering products can only be generated from frequencies at or below 6.375 KHz. Instead of discarding the tone detector and building a discrete PLL/logic system it was considered more practical to place a high-pass filter in front of the decoder chip. To make the most use of the filter's roll-off a fourth-order state-variable filter was designed and built with a cut-off frequency of approximately 18.5 KHz. Hence the drop from 12.75 KHz to 6.4 KHz is 40 dB/decade with the poor roll-off in the transition band being well above the synchronisation frequency.

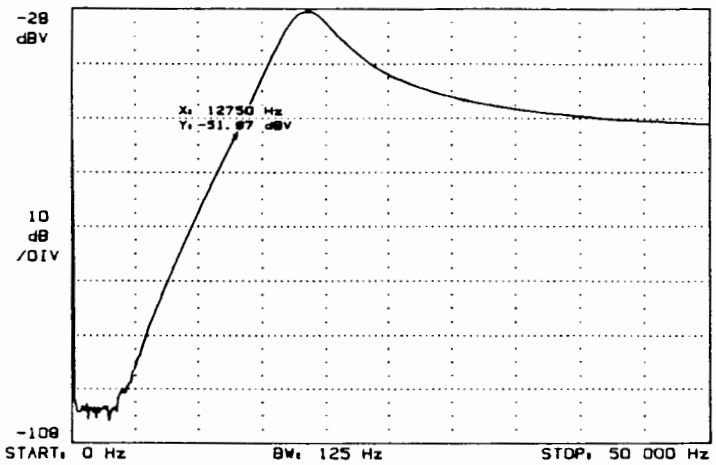

FIGURE 5.14 Sync. high-pass filter

In conjunction with the high-pass filter a 12.75 KHz notch filter was placed in the path of the program material at the transmitter. This removes program energy at this frequency as it was found that this caused occasional spurious triggers to occur. Listening tests done with the notch in and out of circuit showed the loss in highs to be hardly noticeable, this being acceptable for the duration of a test sequence.

A further factor adding to the synchronisation problem was drift of the VCO center frequency with temperature. This was corrected by replacing the VCO timing capacitor with a high quality polystyrene type. The variation of the VCO center frequency with temperature was then measured in an environmental chamber. The results of this are shown in appendix E.

With these modifications the test runs were repeated and in a total period of 24 hours no synchronisation bursts were missed (hence 0 out of 43200).

5.5.2 Additional Modifications

Although a voltage regulator was used for the signal generator's DACs the tests showed that the system stability did not meet the specification. A high accuracy, high stability zener was therefore added to correct this problem.

CHAPTER 6

SYSTEM DESIGN

6.1 INTRODUCTION

The prototype test-set was designed and built using ideas and principles developed in the experimental set. In addition several of the hardware modules constructed for the latter were used in the prototype. In addition the transmitter software was not changed and limited additions and modifications done to that of the receiver.

Important system features not previously discussed will be presented here. The function of each hardware module will be explained as well as the main features of the software. This chapter will therefore build on material presented in earlier ones.

6.2 OPERATIONAL CONSIDERATIONS

One of the main requirements of the system is that it be completely fail-safe or as near to this as possible. Further points which have not yet been considered are line-matching of the unit and the dynamic headroom needed to cater for the variable channel gain.

6.2.1 Fail-safe Features

It was decided that the insertion unit, this being the only section that can disrupt program broadcasts, should be in circuit only during a test sequence and that for the rest of the time it be totally isolated from the channel. This has been accomplished by using relays to insert the unit into the channel under test. This method was used for the following reasons:

i) The isolation and ON-resistance and linearity obtained are better than that of semiconductor switches.

ii) Relays are more robust than their semiconductor counterparts and can be configured so that the effect of failure is highly predictable.

iii) The relays which have been used remove the insertion unit from the channel should the unit loose power. That one of the two states of a relay occurs when it is de-energised is attractive for fail-safe applications.

iv) Although it was initially thought that the switching of the relays would generate loud 'clicks' on the channel this was found not to be the case. In fact this switching is hardly noticeable. In addition, the relays are only switched twice during a test sequence; at the beginning to bypass the channel and at the end to return the channel to its normal configuration.

The relays are switched by the computer but instead of using it to keep the relays energised for the duration of a test sequence a logic circuit in the insertion unit does so. This will only energise the relays on receipt of fifteen pulses from the computer all within 60 milliseconds. Therefore, any spurious transients received from the computer, such as those generated when applying power to it, will be ignored. Tests done on the relay driver verified that this is indeed so. In addition the circuit must be 'refreshed' with a set of pulses at least every 30 seconds in order for the relays not to drop out. This feature caters for the computer being switched off while power to the insertion unit remains on, as well as the unlikely case of the program 'hanging'. The following situations have therefore been catered for:

i) Power removed from only the insertion unit => the relays drop out as they are not driven from the computer.

ii) Power removed from only the computer => channel returns to its normal configuration after 30 seconds.

iii) Power removed from both the computer and insertion unit => relays immediately drop out.

iv) Insertion unit, computer, or both switched on => relays not energised.

v) Program 'hangs' => relays drop out after 30 seconds. This situation is hardly likely as the software has been extensively tested.

At the receiver, the test-unit consists only of a high-impedance input across the channel which cannot disrupt program transmission.

Other features have been incorporated to allow the software to recover from any test-tones not being captured. As pointed out before this was a problem with the experimental test-set which was only solved later in the prototype. The fail-safe features were however kept. These include

i) Calibration trigger: This is about 20 milliseconds longer than normal sync-bursts but hardly noticeable as it it not accompanied by a test-tone. With the calibration trigger received the software can work out the position of a test-tone in the sequence where tones have been missed. Therefore to make sure that the calibration trigger is identified by the detection circuitry it has been lengthened.
ii) Trigger detection: Except for the calibration trigger, its arrival not being precisely known, the trigger detection circuitry is only activated two seconds before a test-tone is due to be received. It is immediately deactivated after receipt of the tone. This was added to decrease the likelihood of spurious triggers.
iii) System clocks: The starting times of all test sequences are set on the two computers so that the trigger detection circuitry at the receiver is only activated a few minutes before the calibration trigger is due. An overlap has been allowed to cater for any mismatch in the clocks at the two ends.

6.2.2 Line-matching and Dynamic Headroom

The system specification [14] required an input impedance of 40K for the insertion unit and as high an impedance as possible for the data acquisition unit. The actual broadcast

channels have 600Ω balanced lines. The above units therefore have balanced inputs with the insertion unit having a balanced output as well. The 40K requirement has been met by placing a resistor in front of the differential input amplifier and the high impedance by using a high input impedance op-amp. The balanced output of the insertion unit uses the low output impedance of two op-amps. The above is shown in figures 6.1 and 6.2.

The insertion gain specification [14] states a -20 dB to +10 dB dynamic range. As the full-scale voltage of the ADC is ± 5 volts the 0 dB voltage level has been set at 1.5 volts. The programmable gain section of the data acquisition unit therefore scales all test-tones from their required transmission amplitudes (i.e. +9, +6, +3, -6 and -12 dBm) to 1.5 V. nominal. Hence a +10 dB gain through the channel will give a full-scale input to the ADC.

6.3 ANALOGUE SUBSYSTEMS

The various analog modules of SPCATS will be considered so that the reader can analyse the circuit diagrams given in appendix D. The function of each module will therefore be considered here and not a detailed analysis given of the circuit diagrams.

6.3.1 Insertion Unit

The insertion unit interfaces to the HP85 and is responsible for tone generation, setting test-tone durations and the gating of program material, sync-bursts and the tones.

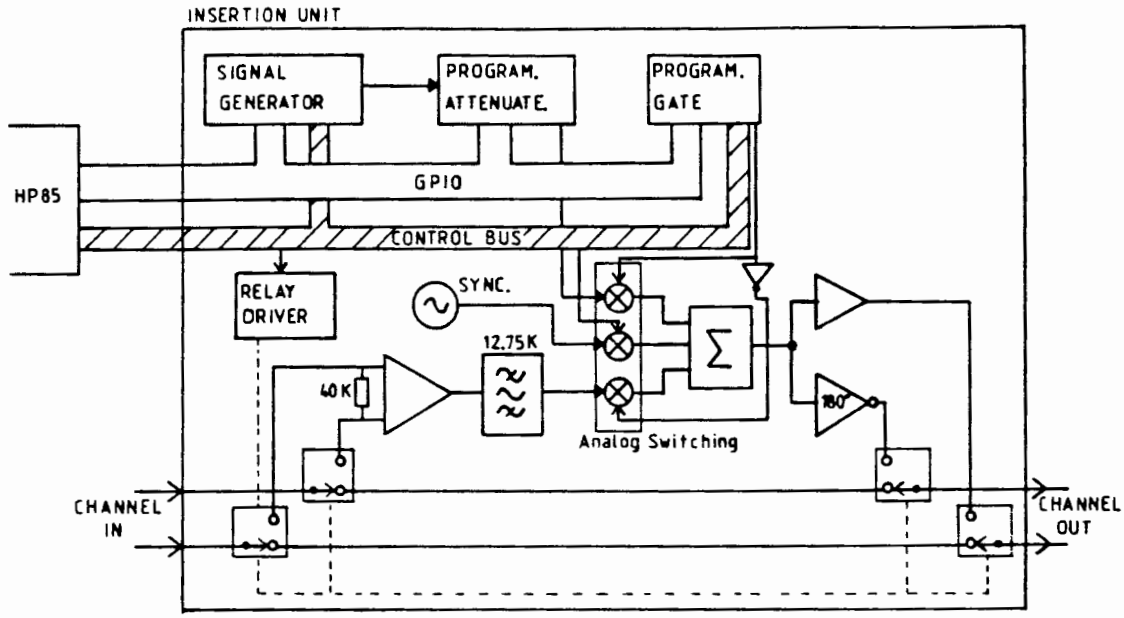

FIGURE 6.1 Insertion unit

Note that the insertion unit uses single-ended signals internally, the conversion to and from differential being done at its input and output.

i) Digital signal generator: This generates the test waveforms and makes them available as a continuous signal at its output.

ii) Programmable attenuator: The gain of each test-tone is set from the HP85.

iii) Analog switch and summer: Here the sync-burst is added to the program material, its duration being set by the width of the pulse on the gate control line from the computer. Directly after the sync-burst the test-tone is inserted while gating out the program, the length of this being set by the programmable gate.

iv) Programmable gate: Four bits of the GPIO set the pulse duration, the pulse being output on the negative edge of the sync. gating control line.

v) Relay driver: Detects the fifteen pulses from the computer required for it to energise the relays and lets them drop out after 30 seconds unless pulsed again by the computer within this time interval.

vi) Sync. oscillator: Generates a 12.75 KHz triangular wave. This arrives at the receiver as a sinusoid due to the 15 KHz roll-off of the channel.

6.3.2 Data Acquisition Unit

This unit adjusts the gains of the arriving test-tones, controls the PCM75 ADC and the sample/hold, detects the synchronisation tone-bursts and generates the gate to strobe the data into the HP220 computer.

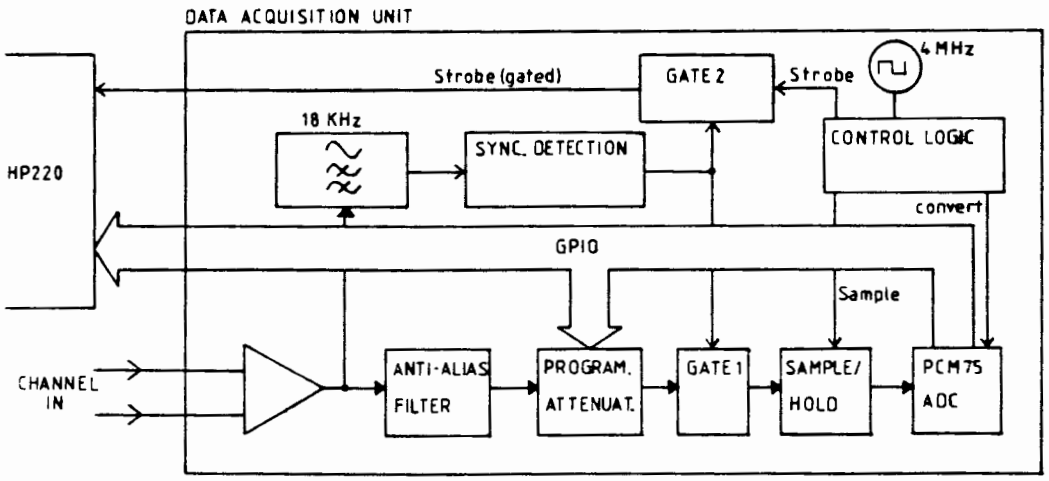

FIGURE 6.2 Data acquisition unit

i) 18 KHz low-pass filter: Removes frequency components below 6.5 KHz to avoid the generation of harmonics and intermodulation products in the tone-decoder doing the synchronisation detection.

ii) Sync. detection and trigger detection: An XR2211 tone-decoder has been used to detect the 12.75 KHz sync.-burst and generate a short trigger pulse.

iii) Anti-alias filter: In the pass-band (0 -> 15 KHz) this has ripple of ± 0.05 dB and stop-band attenuation above 22 KHz of -80 dB.

iv) Programmable attenuator: The incoming test-tones are scaled to 1.5 V. here to provide a constant nominal input level to the ADC.

v) Gate 1: The input signal is only fed to the ADC while the test-tone is present to avoid overloading the device due to the programmable attenuator scaling the program material as well.

vi) Gate 2: This only allows strobe pulses to be received by the computer while a test-tone is being sampled.

vii) Control logic: This generates the 'sample' command of the sample/hold, at 44.4 KHz, and the corresponding 'convert' pulse to initiate an ADC conversion at the end of the 'sample' pulse. In addition continuous 'strobe' pulses are also generated.

6.4 ANALOG-TO-DIGITAL CONVERSION

The experimental test-set used the dynamic signal analyser (HP3561A) as its data acquisition module. The prototype set therefore required a complete analog-to-digital (ADC) system to replace this. Such systems consist in the main of an anti-alias filter and an analog-to-digital converter with its associated control logic. These elements will be considered with respect to the requirements for SPCATS.

A choice existed between using a commercial ADC card specifically designed for the receiver's model 220 computer or a custom built system. The former was considered preferable and tests done on the card. This showed that the card did not perform as well as suggested by the manufacturer's specifications and a custom system had therefore to be used. An ADC specifically for audio PCM has been used and together with careful design of the control logic a high performance system built. Obtaining a satisfactory sample-and-hold amplifier proved difficult due to cost and availability and so the system uses a discrete sample-and-hold built for the application.

6.4.1 The Anti-alias Filter

An anti-alias filter with ± 0.05 dB ripple in the pass-band and 80 dB attenuation above the foldover frequency of 22.2 KHz was required (cf. section 3.3.4). This was satisfied with a passive tenth-order elliptic filter. This filter type has the smallest transition-band, for a specified attenuation, of the standard types. In comparison the dynamic signal analyser uses a similar filter, except of ninth-order [43]. The test-set

filter consists of two cascaded fifth order sections due to the required component accuracy being less for two such filters than a single tenth order filter. The filter was designed using standard published tables to satisfy the above parameters. The circuit diagram for the complete filter is shown in appendix D. Note that due to a 3 dB loss in each filter section an amplifier is used at the filter output to restore unity gain. As only magnitude spectra are required from the test-set, the phase of the filter is irrelevant and therefore this was not a consideration in choosing the filter type. The overall filter response is shown below.

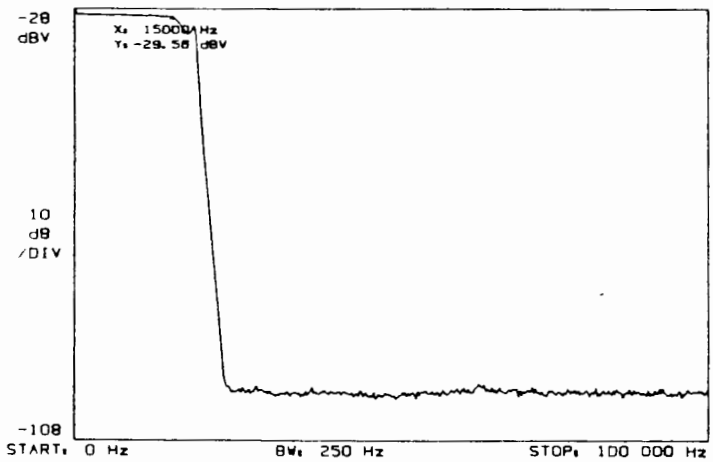

FIGURE 6.3 Anti-alias filter amplitude response

Although the modern trend is to use active filters, a passive implementation was chosen for the following reasons:

i) Noise floor: With active filters it is difficult to get the noise floor below -65 dB to -70 dB The anti-alias filter required this at -80 dB.

ii) Dynamic range: The input signal range is 30 dB and this aggravates the noise problem even more. The problem is caused by the signal amplitude extremes within the filter. A high order active filter is usually made from cascaded second-order sections. Unless the poles of the overall transfer function are grouped correctly the signal is extremely large at the outputs of some sections and very small at others. Where the signal is large distortion and even clipping can occur due to the signal driving the op-

amp at its limits and where small, the signal plus noises is amplified in the following stage.

iii) Component count: Due to the accurate tuning of each filter section and the -80 dB spectral purity required, state-variable sections would probably be used. As each section is second-order and uses four active elements the overall filter would have a very high component count. With the passive implementation only four inductors were used, these being the only bulky elements.

iv) Interfacing: Adding an active module to a system must be done with care so that it does not couple noise through its supplies to the rest of the system. Worse still, it could cause the whole system to become unstable. Due to the passive filter being extremely modular the above problems are virtually non-existent.

v) Filter stability: An active filter has the potential of becoming unstable whereas the passive implementation can never oscillate independently of the system.

Investigations were done on a switched-capacitor filter chip requiring a minimum of external components for its operation. The National Semiconductor MF6 is a 6th order Butterworth low-pass filter [45]. Measurements done on the device proved it unacceptable for the application here:

i) Stop-band attenuation - The noise floor is only slightly better than -60 dB.

ii) Harmonic distortion - The performance as regards distortion is poor. Tests of harmonic distortion in the middle of the audio band (7.2 KHz) gave results of approximately -50 dB THD. In the plot below the top spectrum shows the test signal and that below the output of the filter.

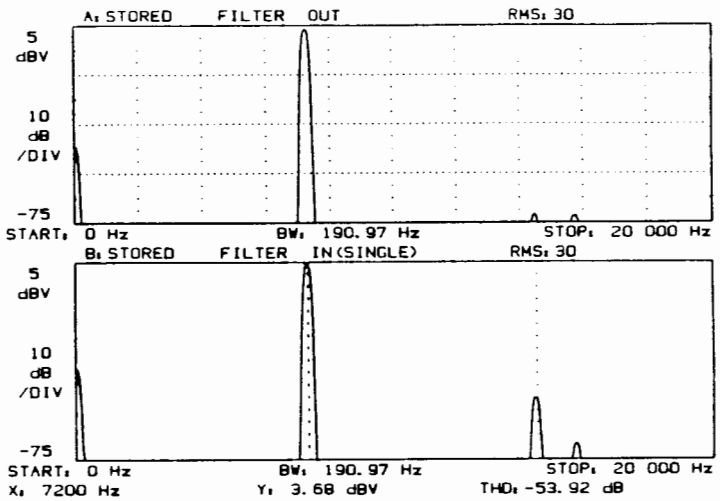

FIGURE 6.4 Switched capacitor filter THD

6.4.2 Hewlett-Packard ADC Interface

The HP98640A 7-channel Analog Input Interface is an analog-to-digital conversion system for Hewlett-Packard 9000 series 200 computers. The system has a resolution of 13 bits and its input channels can be scanned so that readings can be taken of several independent events. In the application here only one channel was required. The interface has four input ranges to cater for a wide variety of applications. Its sample rate is easily set from software. Besides its normal operation from software the card can be set up to take readings on an external TTL compatible hardware trigger being received. By using the synchronising circuit's trigger pulse the interface was instructed to initiate a sampling sequence in an identical way to that used for the DSA. Only a small amount of hardware change was thus necessary to replace the DSA with this ADC card. There is a library of subroutines supplied with the interface with respective sets for HP's Pascal or BASIC language systems. However, the external interrupt mode cannot be used from Pascal and therefore BASIC had to be used. In all the I/O situations encountered in the design of the test-set HP BASIC was found the most suitable of the available programming languages for the model 220 computer.

In an application using one channel the maximum sampling rate is specified at 55 KHz [46]. With the 13-bit resolution and a

linearity of 0.02% it initially seemed that the interface would satisfy the -70 dB accuracy requirement of the test-set. However, investigations showed the ADC to be totally unsatisfactory.

Fast-fourier transforms done on the buffers of samples taken by the card of a high quality sine-wave (THD < -80 dB) [47] showed considerable 'comb' distortion. An example, for an input signal of 5.6 KHz being sampled at 40 KHz, is shown below.

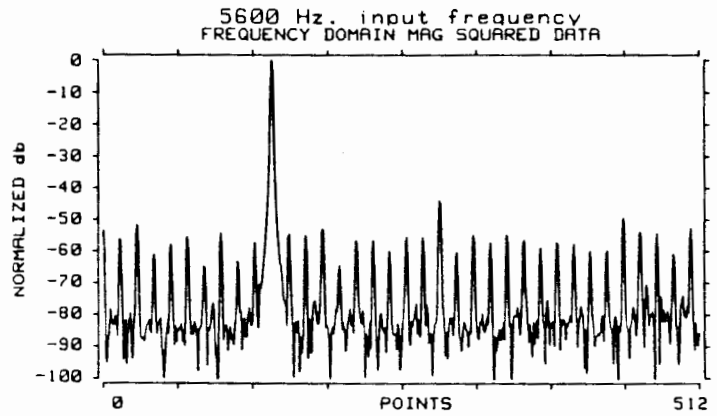

FIGURE 6.5 HP ADC comb spectrum

This effect was seen to occur only at high frequencies within the audio-band. By viewing the time-domain samples for such a waveform periodic errors in the signs of samples near the zero-crossings were found.

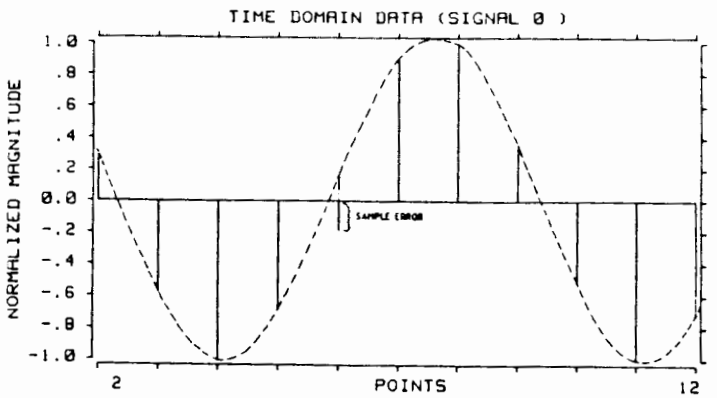

FIGURE 6.6 HP ADC sample sign errors

The specifications give a maximum input slew-rate of only 400 mV/usec for the ± 10 volt input range. This sets the

maximum input frequency at

$V = A.SIN(wt)$

and $dV/dt = Aw.COS(wt)$

so that $dV/dt_{max.} = 2\pi Af_{max.}$

i.e. $f_{max.} = dV/dt_{max.}.(2\pi A)^{-1}$

and so for the case above

$f_{max.} = 6\ 366$ KHz

where: $dV/dt_{max.} = 400$ mV./usec.

$A = 10$ Volts

So although the interface can sample at 55 KHz it can theoretically only input full-scale waveforms at one-ninth this frequency. This is rather surprising and is the reason for this factor initially being overlooked. It seemed reasonable to assume that the ADC would be able to sample waveforms up to the foldover frequency to within the specifications quoted. Note that the 5.6 KHz test frequency is below this theoretical maximum value but close enough in practice to support the reasoning given above.

Although it would have been possible to do modifications to the card it was decided not to due to the cost of the interface and the fact that there was no guarantee that the problem could be solved without extensive tests and modifications. It was therefore decided to build a custom analog-to-digital conversion system.

6.4.3 Custom ADC system

The analog-to-digital conversion system used in the final prototype consists of a Burr-Brown PCM75 audio PCM converter, a sample-and-hold and associated control logic. A National Semiconductor LF398 sample and hold was initially used and tested with surprisingly good results when considering its cost in relation to high quality sample/holds. However, its harmonic distortion was excessive and a discrete sample/hold was used in the final system.

a) Interfacing to the ADC

The PCM75 is 16-bit audio quality converter for use in digital recording studios, for example. It has more than 90 dB dynamic range (96 dB) and THD of 0.004% for a full-scale input. The maximum conversion time to 16 bits is 17 usecs. and the device is compatible with the EIAJ STC007 specification [25] (see section 3.3.4). This compatibility therefore made the ADC ideal for use in the test-set. All the device's specifications are well within the test-set requirements.

The maximum 17 usec. conversion time can be decreased at the expense of the THD [48]. To accommodate the sample/hold and the sampling frequency of 44.4 KHz, a conversion time of 16.5 usecs. has been used. This slight decrease had no noticeable effect on the THD measurements of the test-set. Due to the on-board reference voltage, power-supplies with 1% ripple are adequate.

Although the device's gain and offset errors are fairly large at ± 0.1% these are, as pointed out before, not important and can be trimmed to zero by means of external potentiometers. The differential linearity error, however, is and here it is ± 0.0015%, easily satisfying the SPCATS accuracy requirement of better than -70 dB. Note that with precision modules of this nature a warm-up time is specified. Here it is given as five minutes. The gain and offset drifts with temperature are specified and well within the test-set requirements. Although this device might be considered an 'overkill' it should be borne in mind that the excess performance means that the requirements of other parts of the system can be relaxed somewhat as the ADC contribution to the cumulative error can be ignored. In addition, an ADC to satisfy the 14-bit requirement would still have been costly with no guarantee of correct distortion behaviour unless it were specifically designed for audio. Because of availability and knowing that the PCM75 would satisfy the requirement, without needing to test it, it was the safest ADC to import.

The device can be used with different data formats; straight binary for unipolar input signals and offset binary for bipolar input signals. By inverting the most-significant bit, two's complement can be obtained. All data is complementary code. Due to the HP220 using two's complement notation and having 16-bit integers, the ADC was easily configured to be read directly by the computer's GPIO interface (see section 3.3.5). It has been pointed out before that the GPIO uses a full two-line handshake for all data. For this reason the digital control logic of the ADC system puts a pulse on its handshake line at the end of each conversion. The acknowledgement sent by the computer is ignored as the GPIO can read the data fast enough and it is not necessary for the ADC to wait for the GPIO to become ready. In addition the ADC circuitry does not need to check that the GPIO is ready to receive each burst of samples as each data-block transfer is initiated well before the test-tone is received (see the system software discussion below).

The ADC is set to do continual conversions but is read only while a test-tone is present. At first the hardware interrupt line of the GPIO was used, this being pulsed by the trigger pulse from the synchronisation circuit. However, measurements of the waveforms and timing signals showed a delay of approximately 3 msecs. from the time the trigger was received to the first sample being read. As pointed out before such a delay necessitates making the test-tone longer by 3 msecs. which then infringes on the non-intrusiveness requirement. An alternative method was therefore adopted. By gating the strobe signals to the interface in hardware it is possible to output the command to fill a buffer with samples and the computer then effectively to 'hang' until it has received the correct number of strobe pulses to fill the buffer with data. Here there is no measurable delay between the trigger and the first sample being read. This was checked with a storage-oscilloscope using this trigger and comparing the samples in the buffer to those on the 'scope'. The logic timing for filling the buffer is shown below:

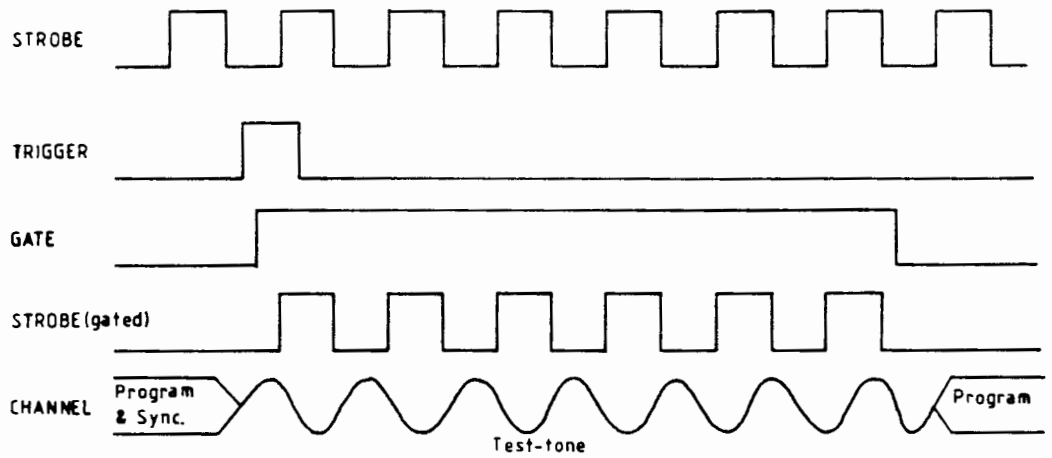

FIGURE 6.7 Logic to transfer a buffer of samples to computer memory

The complete circuit diagram for the ADC system can be found in appendix D.

b) Sample and Hold

A sample-and-hold is used to hold the input voltage at the instant of sampling so that it is constant while the ADC does the conversion. The possible errors that can arise in this 'hold' mode have been discussed before (see section 3.3).

Due to the unavailability of a precision sample/hold to match the performance of the ADC tests were done on a LF398. With a 1000 pF. storage capacitor a sample time of 5 usecs. was chosen and a settling time of 1 usec, these being compatible with the specifications of the device [18]. The sample time is necessary for the output voltage to reach the input voltage after the transition from 'hold' to 'sample' and the settling time for the output voltage transient to die away after changing to 'hold'. With the sampling frequency of 44.4 KHz (period of 22.523 usec.) this leaves 16.5 usec. for the ADC conversion. As the ADC starts conversion on the negative edge of its 'convert' command the 1 usec. delay was incorporated in this pulse.

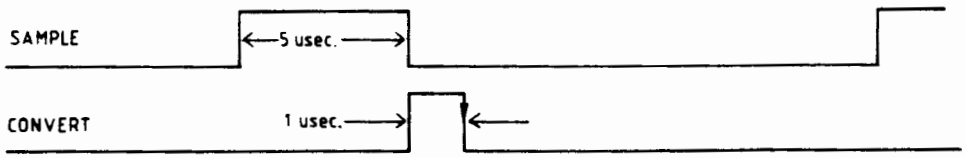

FIGURE 6.8 ADC logic timing

Great care was taken in the design of the data acquisition logic so as to keep clock jitter on the negative going edge of the 'sample' command to a minimum. This is the critical edge in the whole system as it fixes the instant at which the sample is taken. Any uncertainty in this edge will generate gaussian noise causing the overall noise floor to increase (see section 3.3.3).

A 4 MHz crystal, divided by 90, has been used to set the basic 22.523 usec. period. By using two BCD counters and a collection of logic gates it is possible to decode various states of the counters and produce the 'sample' and 'convert' commands of the recquired duration and relative timing. The complexity here is caused by using the clock to toggle flip-flops, with their inputs in the correct states, to generate the recquired commands. In this way their edges are only subject to the jitter in the clock and the flip-flops. The jitter of the crystal clock is very small and by using high speed logic for the flip-flops it has been kept to a minimum. From the diagram below it can be seen that the noise floor is as expected thereby showing that the clock jitter has indeed been kept sufficiently small with the design philosophy adopted. However, the harmonic performance (approximately -68 dB) is just unacceptable. Again, testing was done using a pure sinusoid (THD < -80 dB).

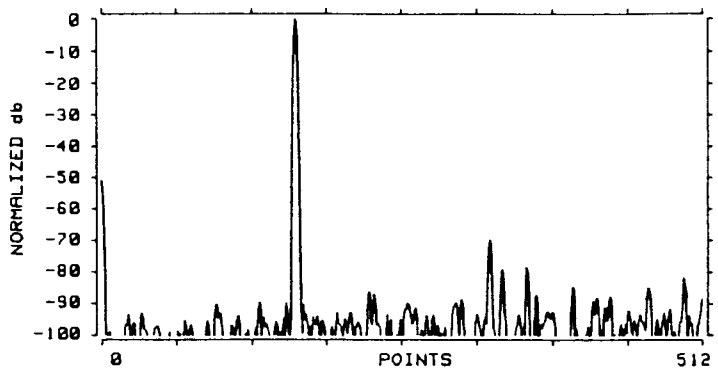

FIGURE 6.9 LF398 sinusoid response

The LF398 is poorly specified by the manufacturer without an aperture time even being quoted. However, due to the harmonic distortion of the device being primarily second-order it seems reasonable to assume that it is due to aperture jitter, this being the distortion caused by such jitter (cf. section 3.3.3).

A study of sample-and-holds produced by various semiconductor manufacturers revealed several devices suitable for use with the ADC. However, none of these were readily available and would have had to be imported. It was therefore decided to build a discrete sample/hold. A data sheet was obtained for Sony's CX20018 16-bit audio PCM ADC where the discrete sample/hold used by them with the ADC is shown [50]. It was therefore decided to use this as the basis for design.

The sample/hold uses a well-known configuration with the storage capacitor in the feedback of the second op-amp.

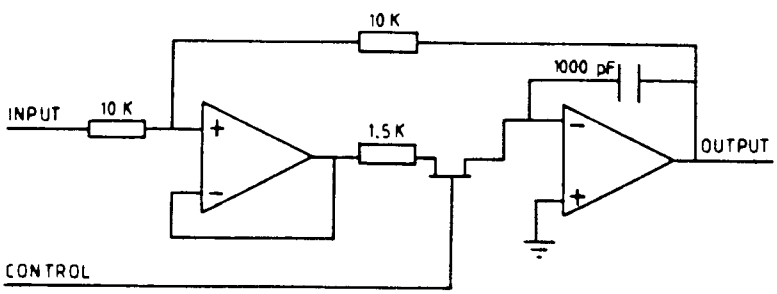

FIGURE 6.10 SPCATS discrete sample-and-hold

A high precision analog switch (Siliconix DG308) has been used as the control element as this is responsible for the majority of sample/hold errors. The switch has an ON-resistance of 70Ω

(typically) [54]. The important property of a switch is that the ON-resistance does not vary with changing signal voltage. This switch has low r_{DS} and the input signal does not significantly modulate the ON-resistance. High speed LF356 op-amps have been used. A detailed circuit diagram can be found in appendix D.

Tests on the system showed an increased noise floor, as expected of a discrete circuit, as well as improved THD. performance. The transient intermodulation distortion is however about the same for the discrete device as for the LF398. By removing the analog switch and replacing it with a simple FET the second harmonic shot up to -40 dB due to the ON-resistance modulation.

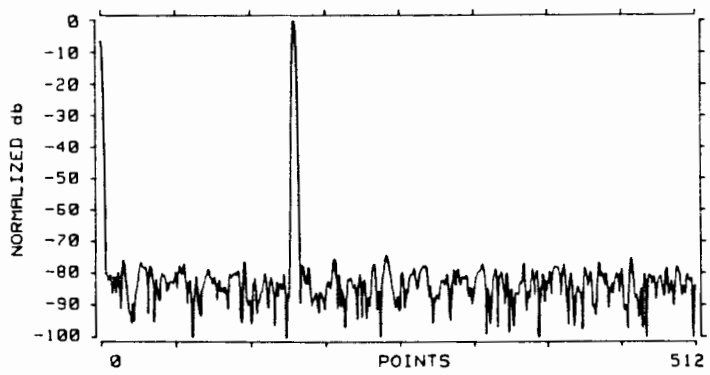

FIGURE 6.11 Discrete sample/hold sinusoid response

6.4.4 System Testing

All tests on the system components were done with a high quality audio test unit designed and built by the author in the course of his undergraduate studies [47]. The unit generates pure sine-waves (THD. < -80 dB), a twin-tone for static intermodulation measurements (with < 0.05% static distortion of its own) and a TIM. test-tone (< 0.1% TIM). It is computer controlled and uses the HP3561A dynamic signal analyser to take readings at the output of the circuit under test.

Analysis of data files on the HP220 was done with an HP9827A Waveform Analysis Pack. In addition to FFT's of data and displaying the resulting spectra for investigation, the package was also used to verify the performance of the

software decimation filter. Due to its versatility the package was even used to check the effect of various windows, the final choice being based heavily on these results. However, the major application of the pack was in checking the performance of the data acquisition system developed.

6.5 SYSTEM SOFTWARE

the SPCATS software consists of a program written for the transmitter unit and one for the receiver unit. As the transmitter's HP85 only functions as the insertion unit's controller its program is largely occupied with interfacing to the hardware via the GPIO. The receiver's HP220 on the other hand is predominantly involved with the analysis of the received test-tones. It uses its GPIO to interface to the data acquisition unit from where the sampled test-tones are read. The various techniques used in these two programs will be discussed and where appropriate pseudo-code presented. A detailed description of the programs will not be given as this is considered no more informative than reading the programs themselves.

6.5.1 Signal Generation

The transmitter unit consists of the insertion unit with its HP85 controller. The HP85 is a single unit consisting of a keyboard, display, printer and tape-drive. To reduce wear and improve data transfer rate a 3½" floppy-disk drive has been substituted for the tape-drive. This speeds up the transfer time for programs and data files considerably.

All control of the insertion unit is done via the GPIO. The test-tones are not computed each time they are output, a separate program having been written to calculate these and store each tone on disk. To load a test-tone into the signal generator its file is read and transferred to the generator. Due to the considerable amount of time taken to compute a test-tone it is not practical to calculate each tone at transmission time.

The transmitter program uses the timer of the computer to initiate test sequences at times set by the User. The User therefore sets the system clock, the time at which tests are to commence and the repetition period of the sequences. Test-tones are transmitted at 20 second intervals to allow sufficient time for each tone to be read from disk and loaded into the generator. A second timer is used for this purpose to generate an interrupt every 20 seconds. Before each tone is transmitted the word to set the correct gain and duration is output through the GPIO to the insertion unit. At the completion of a test sequence the program waits until the repetition time has elapsed before starting the next sequence. This can be summarised as

```
Repeat for each test sequence:
        Wait until set time is reached
        Enable 20 second interrupts
        FOR I= 1 TO number_of_tones
                OUTPUT gain,duration word for I'th tone
                Read I'th tone from disk and load into signal
                generator
                Wait for interrupt
        NEXT I
```

At present the system uses nineteen test-tones to cover the tests discussed in chapter 4.

a) Controlling the Interface

The BASIC of the HP85 is resident and the only Hewlett-Packard supported programming language available for the computer. It is not nearly as structured as that of the HP model 220 machines although their I/O command sets are the same. The BASIC, moreover, does not cater for procedures and only allows two digit variables.

Words (16-bit) are output to the GPIO using 'OUTPUT' commands

in the application program. For example

```
OUTPUT 411 USING "#,W" ; 128
```

outputs the ASCII word ('W') for 128 to the interface at address 4 (the GPIO) with no statement terminator character being used ('#'). The last two digits of the three digit address ('11') set the mode of operation of the GPIO. Ports A and B have been used (hence 16 bits) to drive the signal generator and ports C and D for the gain settings, test-tone duration and channel selection. Ports A and B are used to load the calculated waveform integers into the generator's RAM. Ports C and D however have their individual bits assigned to particular functions of the insertion unit. The bit map is

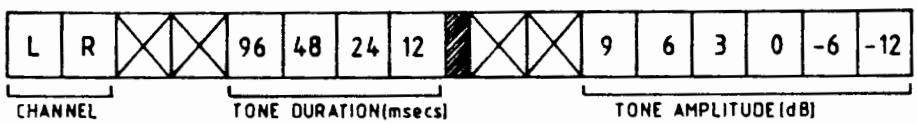

FIGURE 6.12 Insertion unit bit map

By setting the appropriate bits the required amplitudes and tone durations can be selected. With the channel selection, test-tones can be sent on the left or right stereo channels. At present only one channel has been implemented for mono use. Once the required bits have been chosen for the word the corresponding integer is output to the interface. The set of integers required for the test-tones has been stored as data in the program. The BASIC allows values to be stored as data in program lines from where they can be read by the program. This is similar to 'constants' used in Pascal.

b) Signal Calculation Algorithm

It has been mentioned before that test-tones are calculated and stored on disk as files for later use by the transmitter program. Special purpose software has been written to generate these files.

Two programs have been written for the computation of the test waveforms. The first determines the exact voltages of all the

most-significant byte DAC's possible output states and creates a file of these values. It also approximates the quantisation step of the least-significant byte DAC (see section 5.2.1). Using this voltage and the file of voltages the samples of any desired waveform can be calculated.

The second program calculates files for single sinusoids, dual sinusoids and the TIM test waveform. Using the equations for these waveforms the exact sample values are calculated. The closest value that the signal generator can output is then found and used as the sample. This is done as follows:

With the precise value of a sample known, a linear interpolation is done between the zero-scale and full-scale voltages of the MSB DAC to locate the theoretically closest voltage. A search is then made of the voltages on either side of this value to locate the closest real voltage. This may not be the one found by interpolation due to the errors of the DAC. Once this value has been found a LSB DAC voltage, to be added to that of the MSB DAC, is calculated which makes the overall voltage the best approximation to the exact value. The LSB voltage is found by using linear interpolation between the two end point voltages of this DAC. The differential linearity of this DAC is therefore assumed to be zero. This results in a 14 bit DAC with an accuracy of 13 bits.

Although the voltages from the two DACs are simply added, forming the integer representing this total voltage is not as straightforward. This is due to the respective bytes of each DAC's voltage having the same binary weighting whereas the top byte of the integer, which corresponds to the MSB DAC, is 2^8 times greater than the lower byte. One way of overcoming this problem is to multiply the MSB DAC byte by 2^8 and add the least significant byte. This will give an integer with the bytes in the correct positions. An alternative technique was used in the program in which the two bytes are converted to 8-bit strings, these strings concatenated to give a 16-bit string which is converted in turn to an integer. Once the MSB and LSB values have been determined, the integer is formed as

follows:

```
Top_byte$ = MSB converted to a string (8 bits long)
Lower_byte$ = LSB converted to  string (8 bits long)
Total$ = Top_byte$ & Lower_byte string
Integer = Total$ converted to an integer
```

where: $ = a string variable
& = concatenation of two strings

The algorithms used to obtain the theoretical sample values are considered next. Note that for each waveform 2048 (2K) samples are determined and the resultant array written to a file. All waveforms are calculated with an initial amplitude of unity and the scaling done thereafter.

i) single sine-wave:

$$L = INT[(2048f_i/f_s) + 0.5]$$

where: L = number of cycles in the file
f_i = desired sine-wave frequency
f_s = sampling frequency
INT = truncates a real number to give an integer

then

$$y(n) = SIN(2\pi nL/2048)$$

where: n = n'th sample point from 1 to 2048

ii) Dual sine-wave:

This is simply the summation of two single sine-waves of half amplitude each to maintain an overall unity amplitude.

$$y(n) = \tfrac{1}{2}SIN(2\pi nL_1/2048) + \tfrac{1}{2}SIN(2\pi nL_2/2048)$$

iii) TIM waveform:

This waveform consists of a square-wave at 3.18 KHz and a sinusoid at 15.00 KHz. The amplitude ratio of these is 4:1 respectively [37]. The square-wave frequency selected above by M. Otala, the test's designer, does not agree with that specified by the IEC. They have specified this

to be 3.15 KHz. The finite resolution and fixed clock of the signal generator allows only an approximation of these frequencies. The sine-wave frequency is set by

$$L_S = INT[(2048f_i/f_S) + 0.5)]$$
$$= 307 \text{ cycles}$$

now

$$f_i = L_S f_S/2048$$
$$= 14990 \text{ Hz}$$

is the closest frequency. For the square-wave the sign of a 3 KHz sinusoid is used to generate the waveform. The closest frequency is determined by

$$L_{SQ} = 65 \text{ cycles}$$

and $f_i = 3.17$ KHz

For reasons to be discussed below $L_{SQ} = 64$ cycles has been used. this sets the square-wave frequency at 3.125 KHz. This is a 0.8% shift from the IEC. frequency which is itself displaced by 0.7% from the original frequency. The discrepancy between the specified two frequencies suggests that their exact values are not critical and so the generator's approximation is acceptable. The TIM waveform equation is therefore

$$y(n) = (4/5)SGN[SIN(2\pi nL_{SQ}/2048)] + (1/5)SIN(2\pi nL_S/2048)$$

and has a peak amplitude of unity.

After each sample is calculated it is scaled by the MSB DAC's full-scale voltage before being passed to the algorithm responsible for finding the closest generator output voltage. the form of the algorithm, using a sinusoid as an example, is as follows:

```
y(I) = array of 2048 samples
L=64
FOR I= 1 TO 2048
      y(I) = SIN(2πLI/2048)
      y(I) = y(I)*(MSB DAC's peak output voltage)
NEXT I
OPEN a file for y(I)
OUTPUT y(I) to disk
CLOSE the file
```

The reasons for changing L_{sq} to 64 instead of using 65 will now be given. Consider the equation

$$x(k) = A.SIN(2\pi f_d kT) \quad \ldots\ldots\ldots\ldots\ldots\ldots \quad (6.1)$$

Now it was shown in section 3.2.3 that

$$f_d = f_s L/N \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots \quad (6.2)$$

and substituting in (6.1)

$$x(k) = A.SIN(2\pi Lk/N) \quad \ldots\ldots\ldots\ldots\ldots\ldots \quad (6.3)$$

where: f_d = frequency of sinusoid to be generated
f_s = generator sampling frequency
= 1/T
N = number of samples in buffer
= 2048 in this case
L = number of cycles in buffer

It can be shown that where L and N are 'relatively prime' (i.e. their highest common denominator is 1) the values of the sequence of samples only repeats after N samples [19]. In the signal generator N is the buffer size of 2048 samples. However, where L and N are not relatively prime the sequence repeats within the buffer after N/d samples where 'd' is the highest common denominator of L and N. This can be shown as

$$\begin{aligned} x(k+N/d) &= A.SIN(2\pi L(k+N/d)/N) \\ &= A.SIN(2\pi Lk/N + 2\pi L/d) \\ &= A.SIN(2\pi Lk/N) \\ &= x(k) \end{aligned}$$

and from (6.3)

Hence the sequence repeats after N/d samples.

It was found that unless the samples of each period of the sinusoid used to generate the square-wave are identical, the period of the resulting square-wave may vary through the buffer. This is due to certain cycles of the sine-wave having more samples of one sign than others. Hence the period variation occurs when the signs of the samples are used. One way of ensuring that this problem does not occur is for the sample sequence to repeat after each cycle and not only at the

end of the buffer. Hence L and N must not be relatively prime but instead N/d must equal L, the number of cycles in the buffer, where d is the number of samples per cycle. For the square-wave L_{sq} was calculated as 65 but 64 has been used. Hence

$$N/d = L$$

or $$N/L = D$$

In this case

$$N/L = 2048/64$$

and so $$d = 32$$

The number of samples in each cycle is therefore 32. Hence if N/L is an integer the sample sequence will repeat every cycle.

The above condition forces each cycle of a sine-wave to have identical samples. It does not mean however that the number of positive and negative samples per cycle are equal. This is necessary if a square-wave with an equal mark/space ratio is to be generated as in the case of the TIM waveform. More formally, the criterion

$$x(k) = -x(k + d/2)$$

ensures that the numbers of opposite sign samples are equal per cycle. For 'd' even this criterion is satisfied [19]. The requirements for an equal mark/space ratio constant period square-wave generated from a sinusoid are therefore:

i) The number of samples in the buffer (N) must be integer divisible by the number of cycles (L).

ii) The number of samples per cycle must be even.

An interesting result was observed when varying the number of cycles with the generation of a sine-wave. Where the buffer length, N, (here 2048 samples) and the number of cycles, L, are relatively prime there is a slight improvement in the THD of the tone over the case where they are not. In particular where the sample sequence repeats every cycle the THD. is at its worst. This is due to the averaging effect the samples of a relatively prime sequence have on the quantisation error over several cycles. As pointed out before the quantisation noise is not gaussian unless the signal is broadband and of large amplitude. This observation is only mentioned for

interest as it has no significant effect on the digital signal generator's performance.

6.5.2 Acquisition and Analysis

The software used for the receiver unit will be considered here. Besides functioning as the controller for the unit a large amount of the required signal processing has been done in software. The receiver hardware therefore simply functions as the data acquisition unit for the system.

The overall program structure will be discussed with particular emphasis on the signal processing done, I/O interfacing and any unusual features.

a) Overview

A Hewlett-Packard series 9000 model 220 has been used for the receiver computer. It consists of 640K RAM, one floppy and one 4.8 Megabyte hard disk, GPIO and HPIB interfaces, and keyboard, monitor and printer. The GPIO is used to select gains and read the ADC of the data acquisition unit. The HPIB interfaces to all the computer's peripherals (disk drives and printer) and was used in the experimental test-set to control and read data from the dynamic signal analyser. The computer is an industrial version making the addition and removal of component cards easy with a large number of slots available for accessories (8).

The computer uses a 68000 processor making it extremely fast and can even have a DMA controller added to speed up memory and external transfers.

The BASIC programming language for the model 220 has been used for all the code. Features of this BASIC are:

i) I/O is extremely well-handled providing sufficient flexibility at the byte level to avoid the need for assembler routines. Of all the high-level languages available for this computer the BASIC is the best at

handling I/O. This was an important consideration due to the hardware interfacing required with individual bits having to be set and reset. The I/O capabilities of Pascal, including HP's implementations, are virtually non-existent.

ii) The language allows modularity identical to that of Pascal. the most important requirement here is the availability of 'procedures'. As with standard Pascal this BASIC uses procedures with local variables and parameter passing. In fact there is no difference between the languages except that the BASIC uses line numbers and is interpreted. It does however use labels for marking lines and so the User never needs to become involved with the line numbers. In fact the language has certain advantages over standard Pascal. Subroutines are allowed within a procedure having access to all its local variables. A 'common area' can be declared where variables are stored which can be accessed from anywhere within the program. In other words these global variables can be used from within the local environment of a procedure. This avoids the need of tedious repetition of the same variables being passed to several procedures. Therefore, as regards modularity there was no reason found to choose Pascal above this version of BASIC.

iii) Not having a compiler for the language did prove a handicap as the routines for the FFT and digital filter - in fact the whole analysis section of the software - could have been speeded up considerably if these routines were compiled.

iv) The BASIC has three types of variables; 'Reals' with a 53-bit mantissa and 11-bit exponent, 16-bit 'integers' and 8-bit 'bytes'. Although Pascal can define many 'types' the lack of this feature was not found to be a drawback at all. Integers have been used wherever possible to increase speed and only where precision is required, such as in the FFT and Digital filter, have 'reals' been used.

For the above reasons BASIC was chosen above other programming languages, in particular Pascal, in this particular instance.

b) Structure and Modularity

The software has been written in two sections; the first is responsible for interfacing to the hardware, timing the test sequence and storing samples to disk, and consists of two procedures. One initialises the hardware for each tone to be received and the other reads data from the ADC and stores it on disk. The second software section does the analysis of the test-tones and prints the results for each. Here two main procedures are also used; one to do windowing, the FFT and various search algorithms, and the other to output results to a printer. The software configuration is shown below in hierarchical order.

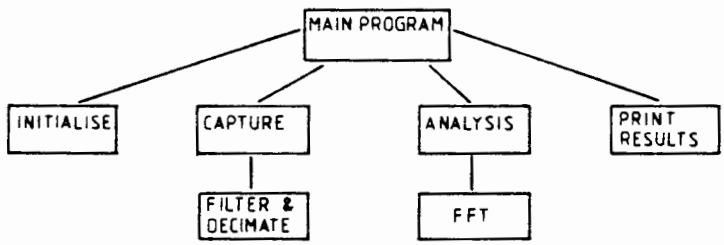

FIGURE 6.13 Receiver software structure

The main program only consists of calls to each level 1 procedure in turn, the same calls being repeated for each test sequence. Using pseudo-code the calling program is essentially

```
REPEAT forever:
        CALL Initialise
        CALL Capture

        CALL Analysis
        CALL Print_results
```

At present the software only works as a mono system but provision has been made for stereo use later involving slight modifications to the program.

c) Sequencing and Housekeeping

This is handled by the 'Initialise' routine. Here the User sets the system clock, the time at which tests are to start and the repetition time of the test sequences. At the set start-time the hardware trigger circuitry and the interface to the acquisition unit are enabled. Hereafter the program exits this routine and the procedure to do the actual capture of the test-tones is called. The transmitter starts transmitting test-tones five minutes after the set-time, this delay being considered sufficient to cater for slight errors in the system clocks or their settings.

At the end of a test sequence the program returns to this routine to wait for the start of the next sequence. The 'Initialise' routine is only called once to set up the system, this not being required thereafter.

d) Test-tone Capture

Capture and storage to disc of test-tones is done by the 'Capture' procedure.

The procedure waits for the initial calibration trigger before setting the correct hardware gains, for subsequent tones, and starting a timer to check that each tone in the sequence is received. As discussed above the calibration trigger was adopted with the experimental test-set as a fail-safe feature and has been kept in the final prototype as it is felt that the more such features a system has the more reliable it should be. There is a limit to this in that the fail-safe features themselves become so numerous that they themselves may lead to system malfunction. However the software here does not even approach this situation.

After the calibration trigger is received the program uses the timer to check that each tone it captures is the correct one in the sequence. As mentioned in Chapter 5 the problem with the sync-detection circuitry was only solved after the

software had been written and it was decided to keep all fail-safe features.

The basic operation of the capture procedure is therefore:

```
WAIT for calibration trigger
REPEAT until all tones received:
        Set correct hardware gain
        Enable data TRANSFER 2 secs. ahead of tone.
        WAIT for transfer of tone to complete
        IF correct tone not received THEN
                Use clock to decide which tone in fact received.
        END IF
        Attach label to tone
        SAVE tone to disk
```

The technique used to read data from the ADC will be discussed as it uses the TRANSFER mode of the computer. Normal I/O is handled by OUTPUT and ENTER commands. However they are too slow to transfer data at 44.4 KHz. The transfer mode can enter data at up to 540 K words/second (see section 3.3.5).

TRANSFER uses a buffer, the size of which is declared in the program, and assigns a 'path-name' to the buffer. The interface forming the other end of the communication is also assigned a path-name. So , for example, the buffer may be assigned the path-name @Buffer and the GPIO interface connected to the hardware, @Gpio. The TRANSFER command moves data between the two path-names, the direction of the transfer being determined by the order of these in the command. Consider the case of data being entered, therefore from @Gpio to @Buffer. Once the data is in the buffer it can only be manipulated as a block. To obtain access to individual values it is necessary to read the buffer into an equal sized array. The above can be written as:

```
ASSIGN @Buffer to a BUFFER[size]
ASSIGN @Gpio to GPIO interface
TRANSFER @Gpio to @Buffer; WAIT
ENTER @Buffer; Samples(*)
```

The transfer can use interrupt, fast-handshake of DMA mode (see section 3.3.5). Here fast-handshake has been used, specified by the WAIT in the transfer line above, giving a maximum transfer rate of 115 K words/second.

Due to the fixed 44.4 KHz hardware sampling frequency, sample-rate decimation is required for the low frequency tones (40, 125 Hz) to increase their FFT resolution (see section 3.4.3). Due to software decimation being used in this software all the ADC samples required must be read into memory before they can be put through the digital filter and decimated. To avoid having to write the un-decimated data to disk the filtering and decimation routine is called from the 'Capture' procedure and is not part of the analysis software. This is the only signal processing not included in this software. The 40 Hz tones are decimated by 6 and the 125 Hz by 2. This sets the buffer sizes at 12K and 4K bytes respectively, the samples being two-byte integers. Because these buffers, in particular the larger one, require a large amount of memory they are only declared for the duration of the transfer. The memory used is released immediately thereafter for other use. With hardware decimation , such as in the HP3561A dynamic signal analyser, the filtering and decimation is done on the samples as they are taken thereby avoiding the need to store redundant samples. This is made possible by using a TMS320 to implement the filter, it being fast enough to handle the 256 KHz hardware sampling frequency. The use of this chip, or a similar device, was considered unnecessary in the test-set as only four tones need to be decimated, there being no speed requirement here and memory costing considerably less than a TMS320.

e) Signal Processing

The main signal processing algorithms used in the system will be considered next. Where relevant reasons for their choices will be given. Of interest then is the window function used, the decimation filter and the different decimation ratios

chosen. The FFT has been discussed previously (section 3.4.1) and the algorithm need not be presented, this being a standard implementation which can be found in almost all digital signal processing texts. It will suffice to say that the Sande-Tukey algorithm for real-valued input sequences has been used to perform a 1024 point FFT.

The 'Analysis' routine proceeds as follows:

```
REPEAT for each test-tone:
        ENTER data file from disk
        Window data
        FFT data to obtain magnitude response
        Use search algorithms to read required harmonics
```

Although the decimation filtering is done at the time of capture it will be discussed here with the analysis software. Finally, the routine to print the results at the end of each test will be mentioned.

f) Windowing

The problems arising due to finite length sample-records have been shown previously (section 3.4.2). This will be considered more formally here.

N uniformly spaced samples of a time-domain signal give rise to N uniformly spaced samples of the associated periodic spectrum under the DFT. The DFT of finite length data is effectively the projection of the observed signal onto a basis set spanning the observation interval. This basis set consists of the sine and cosine terms of the fourier transform. Only signal frequencies which coincide with those of the basis set will project onto a single basis vector. The basis vectors correspond to the N DFT points calculated, each spaced by f_S/N Hz where f_S is the sampling frequency. Frequencies not corresponding to one of these points will project onto the entire basis set [51]. It is this projection which is minimised by the use of window functions. Note that the

problem is caused by the finite length data records and not the sampling process.

As a means of comparing the performance of different windows various parameters have been defined in the literature. Those of particular relevance in identifying a small spectral line close to a large one will be considered here. A detailed mathematical treatment will not be given and can be found elsewhere [52].

i) Equivalent Noise Bandwidth:
Multiplication of a record of samples by a window function corresponds to convolution in the frequency-domain i.e. the time domain signal

$$f_W(nT) = w(nT).f(nT)$$

has a transform

$$F_W(w) = W(w)*F(w)$$

The amplitude of the harmonic estimate at a given frequency is therefore biased by the accumulated broad-band noise included in the bandwidth of the window.

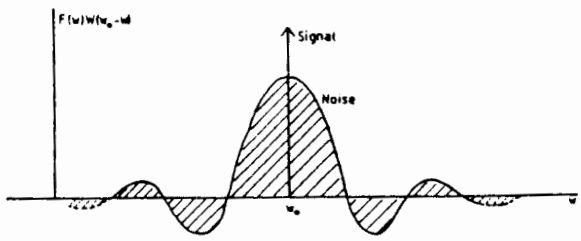

FIGURE 6.14 Broad-band noise accumulated by window adds to harmonic estimate

We therefore want small bandwidth windows to minimise this accumulated noise. A measure of this bandwidth is the equivalent noise bandwidth (ENBW) of the window. This has the usual definition of a rectangular filter with the same peak power and accumulating the same noise power as the original filter (in this case, window). ENBW figures are given for the windows tabulated in table 6.1, but normalised by the noise power obtained with no window present. A buffer of samples with no window function applied is effectively multiplied by a rectangular window.

Hence the normalisation above is relative to a rectangular window. Note that a quantity called 'Processing Gain' (PG) can be shown to be the reciprocal of ENBW.

ii) Coherent Gain:

The DFT can be considered as a bank of matched filters with each having the same fixed gain. This gain is proportional to the window applied and is a maximum for the case where no window is used (hence rectangular). For any other window the gain is reduced due to the window tapering to zero at the boundaries. This reduction in coherent gain represents a known bias on spectral amplitude. Coherent gain, normalised by that of the rectangular window is shown in table 6.1. Note that this gain is computed at a DFT output point.

iii) Scalloping Loss (Picket-fence Effect):

This is a measure of the additional loss in coherent gain due to a tone with a frequency mid-way between two bins and not at a DFT output point as for the coherent gain. The scalloping loss is therefore the ratio of coherent gain for a tone located half a bin from a DFT sample point to the coherent gain for a tone located at a DFT sample point.

iv) Sidelobe level:

Consider the diagram below

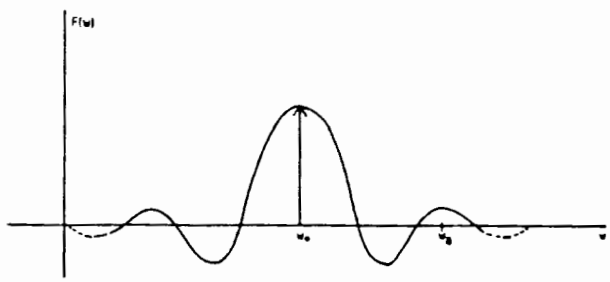

FIGURE 6.15 Spectral leakage effect of windows

For a tone at w_0 and a reading taken at w_a, the value at w_a will not be zero due to energy in the sidelobe at this frequency. In addition, were a small component present at w_a it would be in error by the amount contributed by the sidelobe at this point. Note that this effect is different

to the error caused by the window response centered at the particular frequency being measured due to the broadband noise collected. This is a measure of the area enclosed by the window whereas the effect being considered here is concerned with the peak amplitude of the sidelobes. This is of major concern, and far more important than the ENBW, when trying to measure a small harmonic close to a large one.

v) Minimum Resolution Bandwidth:
This is the bandwidth required to detect two adjacent, equal amplitude, components. If the signals are too close together they will exhibit a single spectral peak. However, where the window main lobes intersect at less than their -6 dB values the two tones can be detected. The -6 dB points are necessary as the adjacent lobes sum coherently and must have dropped by more than half at their crossover point if the peaks are to be detected. The -6 dB bandwidth is usually expressed in bins (1 bin = f_s/N) and is a measure of the width of the window's main lobe.

The above parameters are given below for the more well known windows as well as that used in the analysis routine, the Blackman-Harris window.

Window	Highest side-lobe level (dB)	Coherent gain	Equiv. noise BW (bins)	Scallop loss (dB)	6.0-dB BW (bins)
Rect-angle	-13	1.00	1.00	3.92	1.21
Hanning	-32	0.50	1.50	1.42	2.00
Hamming	-43	0.54	1.36	1.78	1.81

4-sample Kaiser-Bessel (a=3.0)	-69	0.40	1.80	1.02	2.44
Minimum 4-sample Blackman-Harris.	-92	0.36	2.00	0.83	2.72

TABLE 6.1 Relative window performance figures

It has been mentioned before that decreasing window side-lobe amplitudes increases the main-lobe width. This can be seen from the table where the 6.0-dB BW increases as the side-lobe levels decrease. What is desirable here is that this increase in the main-lobe width is kept to a minimum. This relative trade-off is a performance criterion of the particular window type. The decrease in coherent gain with a decrease in side-lobe level is due to the increased tapering required towards the window boundaries.

Two main criteria were used in choosing the SPCATS window; side-lobe level with as narrow a 6.0-dB BW as possible, and the scallop loss. As the test-tones being analysed do not fall on DFT points the effect of this on their amplitudes had also to be considered. A window with a low scallop loss as well was therefore required. One of the windows satisfying these criteria most effectively is the minimum 4-sample Blackman-Harris window. This window is defined by

$$w(n) = a_0 - a_1\cos(2\pi n/N) + a_2\cos(4\pi n/N) - a_3\cos(6\pi n/N) \qquad n = 0,1,2,\ldots,N-1 \qquad \ldots\ldots\ldots\ldots (6.4)$$

where: $a_0 = 0.35875$
$a_1 = 0.48829$
$a_2 = 0.14128$
$a_3 = 0.01168$
n = n'th time domain sample

To obtain the -92 dB side-lobe levels it was found that the coefficients had to be used to five decimal places. Changing the last digit of even one coefficient was found to have a dramatic effect on the side-lobe performance.

To show the difference in detection of a small harmonic close in to a large one the plots below have been included. The first is for the Hanning window and the second for the Blackman-Harris.

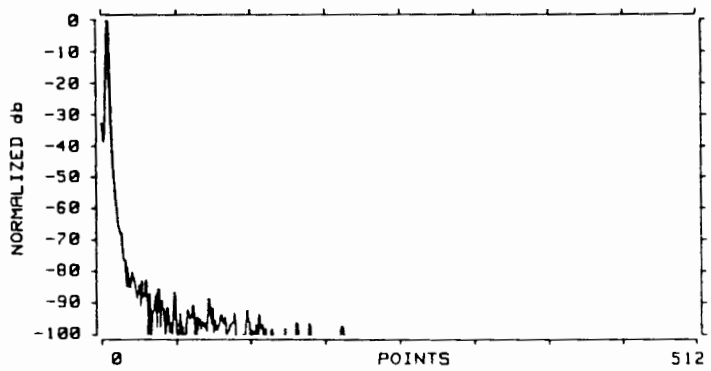

FIGURE 6.16 Hanning window harmonic detection

With this window the second harmonic of the 40 Hz test-tone at -68 dB relative to the fundamental is hardly noticeable (sampling frequency of 7.2 KHz).

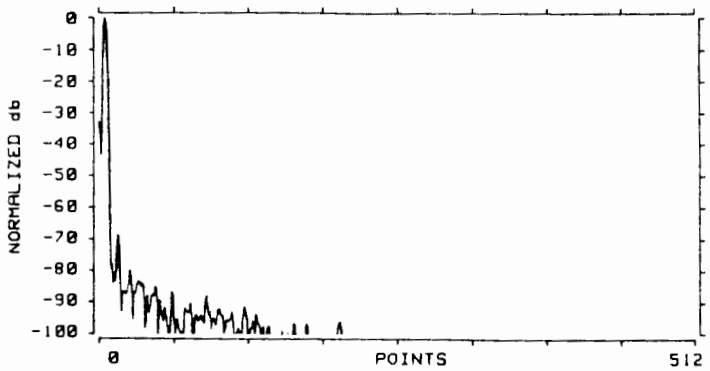

FIGURE 6.17 Blackman-Harris window harmonic detection

The vast improvement of the Blackman-Harris window, due to its low side-lobe level, is seen above.

Application of equation (6.4) to an array of 1K samples would

be implemented as

```
FOR n = 0 to 1023
      Window_sample=a0-a1COS(2πn/1024)+a2COS(4πn/1024)-a3COS(6πn/1024)
      Samples(n) = Samples(n)*Window_sample
NEXT n
```

where: Samples(*) = array of samples
Window_sample = n'th window sample

g) Decimation Filtering

The need for decimation of the low frequency tones (40, 125 Hz) has been considered previously (see section 3.4.3). With a sampling frequency of 44.4 KHz and a 1024 point FFT, the spectral resolution is 43 Hz. For a 40 Hz fundamental, the harmonics are therefore spaced by less than 1 FFT bin. In the discussion above the 6.0-dB bandwidth of the Blackman-Harris is 2.72 bins. As pointed out this is the required bin spacing to detect two adjacent equal amplitude components. In the case of detecting an adjacent harmonic 80 dB down on the fundamental, the bin spacing must be increased and hence the FFT resolution. Tests done on the window showed that a bin spacing of six or more allows effective detection of the harmonic. It is interesting to note that similar tests done on the Hanning window required a spacing of 15 bins at least in detecting the harmonic.

For the 40 Hz test-tones, the bin width needs to be approximately 7 Hz or less to satisfy the above requirement. Hence the required decimation ratio is

$$d_{40} = 43/7 \approx 6$$

With the 125 Hz test-tones a bin width of

frequency resolution = tone frequency/6
= 21 Hz

is required. This is satisfied by a decimation ratio of

$$d_{125} = 43/21 \approx 2$$

With the hardware sampling frequency of 44.4 KHz and the anti-alias filter which has 80 dB attenuation above 22 KHz no

aliasing occurs. To reduce the sampling rate in software it is also necessary to avoid aliasing. Therefore to decimate by 6, giving a sampling rate of 7.4 KHz, the spectrum must have no components greater than -80 dB above the 3.7 KHz foldover frequency. The decimation by 2 has a less stringent requirement, the new foldover frequency being at 11.1 KHz. Consider the spectrum of a 40 Hz test-tone shown below with no decimation or digital filtering having been done.

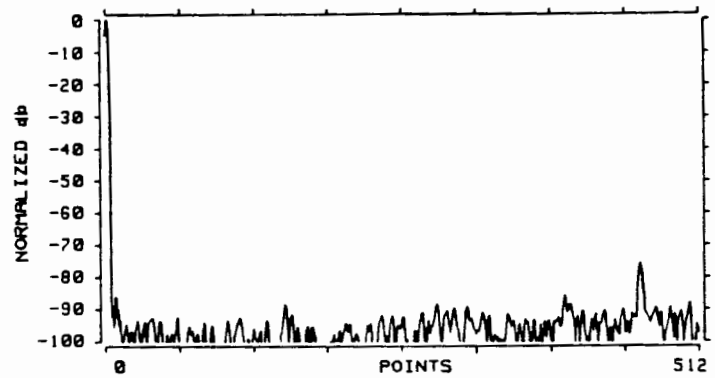

FIGURE 6.18 Unfiltered 40 Hz spectrum (f_s=44.4 KHz)

This shows the noise floor to be well below the -80 dB requirement except for a spurious harmonic at 15.6 KHz. This is generated by the transmitter computer. The filtering requirements for both the 40 Hz and the 125 Hz are the same as both need their un-decimated spectra clear to 11 KHz and the noise floor extends up to 15 KHz. Once filtered the decimation can be done. The detailed design of the filter is given in appendix B. An example of a 40 Hz filtered and decimated test-tone is given below.

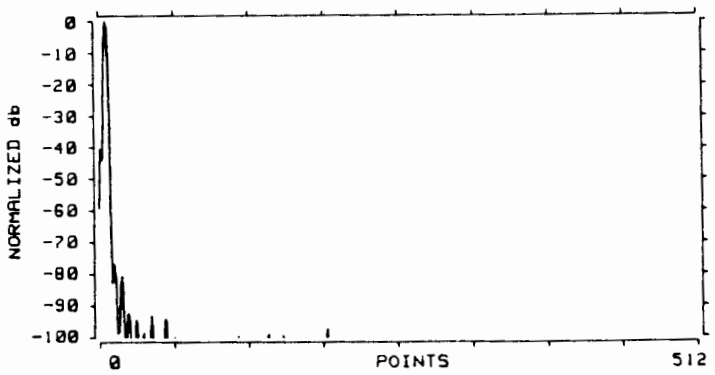

FIGURE 6.19 40 Hz tone, filtered and decimated by 6 (original sampling frequency of 44.4 KHz)

Here the harmonics can all be detected, this not having been the case in figure 6.18. The results for the 125 Hz test-tones are similar.

The decimation algorithm is straight forward with 5 out of every 6 samples being discarded in the 40 Hz case and every alternate sample of the 125 Hz test-tone. Buffers of 6K and 2K samples respectively are read from the ADC into two arrays for processing. After filtering the decimation is implemented as

```
J=Ratio
FOR I= 1 TO 1024
      Samples(I) = Samples(J)
      J = J + Ratio
NEXT I
```

After the completion of the above algorithm, the array of samples, Samples(*), contains the decimated samples as the first 1024 array elements. These samples are windowed and fast-fourier transformed hereafter.

h) Search Algorithms

Once the FFT has been computed for each test-tone, the magnitude spectral points are stored in an array for use in obtaining the distortion test results discussed in Chapter 4. Knowing the spectral resolution and the precise fundamental frequency, or frequencies, of each test-tone the closest frequency bin can be calculated and the harmonic read.

Initially the search algorithms used did not calculate the precise frequency bin of a particular harmonic component using the exact fundamental frequency, this being a priori information. Instead the approximate fundamental frequency was used to calculate the harmonic to an accuracy of two to three bins with the algorithm then reading the amplitudes of the components on either side of the selected bin to find the largest one. This method works well where the fundamental's frequency is not very stable and the harmonics being measured are larger than the surrounding noise. It was found however

that with the small harmonics to be read the algorithm would select a larger adjacent noise spike instead of the harmonic. An algorithm calculating the precise location of each component to be read has therefore been used instead as the test-tone frequencies are controlled by a crystal clock and stability is not a problem.

6.6 CALIBRATION

Calibration of the test-set is done before putting the unit into service. The receiver and transmitter sections are connected by means of a short balanced line and the calibration option of the receiver software run. This is identical to the normal receiver run except that the amplitude of each test-tone is stored in a file on disk. The transmitter section does its normal test run here. For this calibration procedure to be meaningful, the two sections need to have a high degree of temperature and drift stability.

The advantage of such a calibration procedure is that linear errors throughout the system can be taken account of. For example, the roll-off of the test-tone amplitudes with an increase in frequency due to the digital generation is accounted for with this technique. Accurate trimming of individual elements of the test-set is also eliminated, the overall 'trimming' values being stored on disk for use with each test run.

Only those test-tones measuring absolute levels need a calibration value. Readings such as THD, static intermodulation distortion (SIM) and TIM obtain their results from considering the relative spectral content of the received test signal. Tests such as frequency response, insertion gain and gain compression subtract their calibration values from the received test-tone to obtain absolute results.

CHAPTER 7

OPERATIONAL EVALUATION

7.1 DESCRIPTION OF TRIALS

Before putting the final prototype into service it was tested thoroughly to ensure that it is reliable and meets performance expectations.

Two series of tests were done. The first was carried out in the laboratory in the same way as those done for the experimental set. Program material was fed from a rebroadcast receiver into the insertion unit and the output of the unit joined to the data acquisition unit by a twisted pair of wires.

The second series of tests was done by courtesy of the S.A.B.C. The insertion unit, with its HP85, was set up at their Seapoint studios in Cape Town. The studio output-signal from the station 'Radio Good Hope' was fed through the insertion unit and on to the transmitters at 'Constantiaberg'. This transmitter caters for most of the Cape Flats and covers the University as well. By connecting the SPCATS receiver to the rebroadcast receiver a transmission link was established. By setting the HP85's clock and the test sequence start and repetition times it was possible to leave the unit unattended at the studios indefinitely with the time of each test run being precisely known. At the receiver the corresponding times were set. Note that the system was calibrated in the laboratory before leaving the insertion unit at the S.A.B.C. Results of a sample calibration run are shown in figure 7a.

7.2 EVALUATION

The laboratory tests will be considered first. As for he experimental test-set the performance of the synchronisation circuitry was checked using the two programs written specifically for this, one for the transmitter and the other for the receiver. Sync-bursts were transmitted every 2 seconds

and the receiver kept running totals of those acknowledged and those missed. This test was run for several hours and confirmed the results obtained for the experimental set with no triggers being missed.

A calibration run was done in the laboratory and several test runs thereafter to check the stability of the results for a series of tests. Theoretically these should have all given the same results as the calibration and test runs were done across the same twisted pair. The set of test runs showed the system stability to satisfy the specification [14].

The purpose of the second series of tests was to do a final check on the synchronisation circuitry and to get an idea of the distortion present on the high quality channels of the S.A.B.C. No hitches were found with the former and the quality of the channel tested was found to be very good. Here it was naturally not possible to use the programs written to check the synchronisation in the laboratory due to the disturbance it would have caused to listeners of Radio Good Hope.

With test sequences being repeated every half hour and on the hour it is interesting to note that only after four days of this did listeners start telephoning the S.A.B.C. to query the 'blips' that they kept hearing. This was done to check public response and suggests that with the tests done as intended, namely once or twice a day, late at night or early in the morning, the majority of listeners will never know that tests are being done. As pointed out before the test-tones were not intended to be totally non-intrusive but rather to cause minimal disturbance to listeners. That the queries only started after four days seems to agree with the listening tests done in the laboratory in that the tests are hardly noticeable.

A copy of the results obtained for one of the tests between Seapoint and the University is also shown below. The results are copies of the computer's printouts.

CALIBRATION RUN (BUILT 0&H)
9 Oct 1986

Test time (hours:mins.)= 0: 0

PRINTOUT AT 13:49:05

FREQUENCY RESPONSE:

FREQ.(Hz)	LEFT(dB)	RIGHT(dB)
37	0.00	
130	0.00	
391	0.00	
998	0.00	
5598	0.00	
7160	0.00	
10023	0.00	
14970	0.00	

Multi-tone response:

391	0.00	
998	0.00	
5598	0.00	
7160	0.00	
10023	0.00	
14970	0.00	

INSERTION GAIN

1 KHz @ -12 dBm:	LEFT(dB)	RIGHT(dB)
	0.00	

GAIN COMPRESSION

1 KHz @ +6 & -6 dBm:	LEFT(dB)	RIGHT(dB)
	0.00	

TOTAL HARMONIC DISTORTION:

FREQ.(Hz) @ +9 dBm:	LEFT(dB)	RIGHT(dB)
37	-76.72	
130	-98.83	
391	-89.91	
998	-90.89	
FREQ.(Hz) @ +6 dBm:	LEFT(dB)	RIGHT(dB)
5598	-78.32	
7160	-79.93	

STATIC INTERMODULATION DISTORTION:

Test tone @ +3 dBm:	LEFT(%)	RIGHT(%)
	.31	

TRANSIENT INTERMODULATION:

Test tone @ +8 dBm:	LEFT(%)	RIGHT(%)
	.17	

(a) Calibration

SEAPOINT U.C.T.
10 Oct 1986

Test time (hours:mins.)= 7:30

PRINTOUT AT 07:51:33

FREQUENCY RESPONSE:

FREQ.(Hz)	LEFT(dB)	RIGHT(dB)
37	.53	
130	.06	
391	.02	
998	0.00	
5598	.09	
7160	.30	
10023	.06	
14970	.11	

Multi-tone response:

391	-.03	
998	0.00	
5598	.11	
7160	.29	
10023	-.04	
14970	.14	

INSERTION GAIN

1 KHz @ -12 dBm:	LEFT(dB)	RIGHT(dB)
	-2.75	

GAIN COMPRESSION

1 KHz @ +6 & -6 dBm:	LEFT(dB)	RIGHT(dB)
	.02	

TOTAL HARMONIC DISTORTION:

FREQ.(Hz) @ +9 dBm:	LEFT(dB)	RIGHT(dB)
37	-43.36	
130	-62.01	
391	-56.09	
998	-57.14	
FREQ.(Hz) @ +6 dBm:	LEFT(dB)	RIGHT(dB)
5598	-53.68	
7160	-47.38	

STATIC INTERMODULATION DISTORTION:

Test tone @ +3 dBm:	LEFT(%)	RIGHT(%)
	.57	

TRANSIENT INTERMODULATION:

Test tone @ +8 dBm:	LEFT(%)	RIGHT(%)
	2.99	

(b) Test

FIGURE 7.2 Sample SPCATS results

For the calibration run the THD values are all well within the -70 dB requirement, as is that for static intermodulation distortion (SIM). The TIM marginally satisfies its specification.

The channel tested was relatively short which might have contributed to the good performance observed. The THD at aproximately -50 dB over the frequency range, the SIM at .6%, and the TIM of 3% are probably adequate for voice and popular music.

However, on a station such as 'Radio Allegro' , which airs predominantly classical music, this might be the reason for listeners not being satisfied with the service, especially over longer links. It is anticipated that the results for a long link such as that from Cape Town to Johannesburg will not be nearly as good as the above due to the large amount of Post Office equipment involved between the two ends.

CHAPTER 8

FUTURE DEVELOPMENT

8.1 INTRODUCTION

Possible extensions and improvements to the SPCATS prototype will be considered here as well as a different approach which may result in a system in which testing is totally non-intrusive. Due to the time schedule of the project it was not possible to try any of these.

8.2 SPCATS EXTENSIONS AND IMPROVEMENTS

The problem with the existing calibration technique is that the transmitter and receiver units need to be at the same location and continued accuracy depends on their long-term stability. A preferable technique would be for these units to be calibrated independently of each other. For example, each unit could measure the voltage across an internal high accuracy zener-diode and use this as a calibration reference. The problem here is that this can only be done at one frequency, say 1 KHz. If the drift in the system is only localised (around 1 KHz. in this case) then the remainder of the system will be in error. To do the calibration at various frequencies throughout the audio-band becomes complex as a variable amplitude frequency response amplifier is then required to deal with the localised drift. This adjustment could be done under software control but would be different to the software compensation used at present. Due to the independence of the two ends it would be necessary for the software to adjust the hardware so that the output of the insertion unit is always at a predefined level. It would therefore not be possible to implement such a calibration technique purely in software as is done for the present system. Were a data link available the variation in the insertion unit's hardware could be relayed to the receiver where it could be compensated for.

Although the technique used to implement the 14-bit DAC in the digital signal generator works perfectly it is not a simple process to calibrate this DAC and generate the waveform files. Therefore, to make copies of this system is extremely laborious as each unit must be calibrated and have its own set of waveforms calculated. A better approach in this regard therefore would be to use a commercial 14-bit DAC as there is no calibration required in this case and the same test-tone files can be used for any number of units.

An improvement to the existing signal generator, however, would be to use a look-up table for the least-significant byte (LSB) DAC as well. This would not increase the resolution as the two bit overlap fixes this at 14 bits (see section 3.2.1). However, the accuracy at present is only 13 bits due to the linear interpolation used to obtain the LSB value. This therefore assumes a perfectly linear DAC and ignores the differential linearity error. the use of a second look-up table would take this into account and provide a 14-bit DAC with a 14-bit accuracy. Hence the only errors would be those due to the quantisation uncertainty and drift.

8.3 AN ALTERNATIVE APPROACH

Some consideration has been given to the possibility of a truly non-intrusive testing technique based on natural program material. The basic philosophy of such a technique would be to identify the linear and non-linear distortions of the channel by a detailed comparison of the program material at each end. Suitable material might be music recorded on a high-fidelity compact-disc player. A segment of this material could then be used as a 'test' signal. This could be transmitted at a known time and captured at the remote end. Careful synchronisation would be required to ensure that the segments compared at each end were indeed identical. In order to compare the segments an accurate copy of the original material would be kept at the remote end. Subject to an adequately wide information bandwidth in the original source which fully exercises the channel there should be no special problems in inferring linear distortions in this way seeing this is the principle

used in the Audio-dat system. Here a statistically based comparison of the energy at each end in various frequency bins is made. Detailed calculations would however be required in deciding whether sufficient accuracy could be attained with this technique.

The non-linear distortions, however, present problems which would require novel work. An example could be the use of adaptive waveform equalisation using a transversal filter-type equaliser to perform as close a match as possible of the two waveforms. These could then be statistically compared in order to quantify the non-linear distortions. This technique is believed to be potentially viable. The work could draw on a great deal of recent research into adaptive equalisation but nevertheless aspects of it are , novel. It would also be necessary to quantify the distortions in terms of the standard audio-quality tests so that results can easily be interpreted.

Unless the 'test' signal is a 'jingle' or similar material, it would have to be played at appropriate times to fit in with the rest of the material of a particular programme. The SPCATS system also has this limitation as it is preferable to run tests late at night or early in the morning. In both situations the user is therefore not at liberty to run tests at random.

CHAPTER 9

CONCLUSION

Important ideas, investigations and results from this text will be highlighted here. In particular novel features of SPCATS, and areas where advanced electronic engineering procedures have been applied, will be mentioned.

The text opened with the requirements for the test-set pointing out that no suitable commercial equipment is available. The Audio-dat is the closest system in concept but does not provide the same results as SPCATS. It is a requirement of the system to perform tests on live channels and this testing must be as non-intrusive as possible. The technique that has been adopted is to use short tone-bursts. A software intensive system was shown to be preferable and this strategy has been adopted.

The second chapter developed a specification for the system showing what tests, system stability and tone durations were required. The basic functioning of the system was laid out together with its configuration and required electrical characteristics.

In Chapter Three the important concepts used in the design and implementation of the test-set are presented. The decision to make the system software intensive meant that a fairly significant amount of digital signal processing algorithms had to be used at the receiver end. The features of digital filters are discussed and the reasons for the choice of an IIR implementation for the decimation filter given. The use of digital sample-rate reduction is presented and it is shown how it allows an effective variable sampling frequency system to be designed using a fixed hardware rate and constant anti-alias filter. The dramatic effect that the choice of FFT window has on the resolution of the FFT is demonstrated, and the choice of the most appropriate window is motivated.

A novel 14-bit programmable digital signal generator has been constructed using two 8-bit DACs and a software error correction technique to give a final accuracy of better than 13 bits. To aid in this discussion the errors and irregularities of real-world data conversion systems are presented. An error analysis for the DAC08 is given, this being the type of DAC used in the generator.

A discussion of the sampling frequencies used in the literature and in industry is given to show why the data acquisition unit's sampling frequency of 44.4 KHz. was chosen. Interfacing of the data acquisition and insertion units to their respective microcomputer controllers is presented.

Chapter Four gives a detailed specification of the tests implemented by SPCATS categorising these into those that measure linear distortions and those for non-linear distortions. Particular attention is given to transient intermodulation distortion and the mechanisms causing it.

Before designing and constructing the final prototype test-set an experimental system was made to test the concepts that were to be used in the final set. This is presented in Chapter Five with particular attention being given to the synchronisation system developed. Various problems were encountered which were only solved after detailed consideration of the tone-decoder at the heart of the system. A testing procedure was adopted which was considered a good simulation of the real conditions under which the system would operate. Special-purpose software was written for these tests and used to iron out the synchronisation problems. The majority of hardware built and software written for the experimental system was used in the final prototype. Hence the prototype was constructed with a reasonable amount of confidence that it would work.

The sixth chapter presents a detailed discussion of the hardware and software used in the prototype. Important parts of the software have been highlighted and pseudo-code given to show their implementation. Where appropriate, plots of actual

data obtained from analysis of hardware and software modules have been presented.

A detailed discussion of the requirements for the digital generation of waveforms is given with a mathematical treatment for the number of cycles that can be stored in a buffer of finite length and the effect these have on the symmetry of sinusoids.

Extensive tests were done on the final prototype to ensure reliability and that it meets its specification. Besides simulated tests, identical to those done on the experimental test-set, the system was used on a real audio broadcast channel. This is discussed in Chapter Seven with sample test results being given.

Chapter Eight considers possible improvements to the present system to improve the accuracy of the signal generation to 14 bits and overcome the difficulty with the present calibration system. In addition an alternative technique for testing live channels is proposed where this testing is totally non-intrusive.

The SPCATS prototype that has been developed has been found to be reliable and to meet the initial specification set. Use has been made of digital signal processing techniques to ease hardware requirements and provide flexibility for the addition of new test procedures or the modification of existing ones. Although the system is not totally non-intrusive it is felt that the disturbance of the test-tones is minimal. It has not been possible to test several audio broadcast channels due to the limited time for the project. It is therefore not possible at this stage to comment on the quality of the audio broadcast network in this country.

REFERENCES

[1] Cabott, R.C."Design Factors in a Programmable Distortion Measurement System." J. Audio Eng. Soc., Vol. 32,No. 12,pp. 964-983, Dec. 1984.

[2] Finger, R.A. "Computer-Based Signal Processing for Audio Electronic Measurements." J. Audio Eng. Soc., Vol.31, No.10, pp. 755-759, Oct. 1983.

[3] Hempel, J. "Automatic Quality Monitoring and Data Transmission for Radio Broadcasting Installations During Ongoing Program." Nachrichtentechn. Z., Vol. 27, pp. 403-406, 1974.

[4] Jakubowski, H. "Analyse des Programmmaterials des Hörrundfunks." Rundfunktechn. Mitt., Vol. 15, pp. 275-284, 1971.

[5] Mills, R.M. "Results of Monitoring Insertion Gain and Programme Levels with the Audio-dat System." Private Communication, Nov. 1986.

[6] Hewlett-Packard Application Note "Understanding the HP3582A Spectrum Analyser." April 1978.

[7] Preis, D. "Linear Distortion." J. Audio Eng. Soc., Vol. 24, No. 5, pp. 346-367, June 1976.

[8] Waddington, D.E. "Intermodulation Distortion Measurement." J. Audio Eng. Soc., Vol. 12, No. 3, pp. 221-225, July 1964.

[9] Otala, M., Leinonn, E. "The theory of Transient Intermodulation Distortion." IEEE Trans. ASSP, Vol. ASSP-25, No. 1, Feb. 1977.

[10] Bloch, A. "Measurement of Nonlinear Distortion." J. Audio Eng. Soc., Vol. 1, No. 1, pp. 622-67, Jan. 1953.

[11] Leinonen, E., Otala, M. "Correlation Audio Distortion Measurements." J. Audio Eng. Soc., Vol.26, No. 12, pp. 12-19, Jan./Feb. 1978.

[12] Hewlett-Packard, HP-85: I/O Programming Guide, Jan. 1981.

[13] Hewlett-Packard, Basic Programming Techniques, with extensions 2.0, for the HP series 200 computers, Dec. 1982.

[14] Project 1984/1. "Development of a sound programme circuit automatic test set." South African Broadcasting Corporation.

[15] Supplement to IEC publication 268-3, Feb. 1983

[16] Clayton, G.B. Data Converters, Hong Kong: Macmillan Press, 1983.

[17] Blesser, B.A. "Digitisation of Audio: A Comprehensive Examination of Theory, Implementation, and Current Practice." J. Audio Eng. Soc., Vol. 26, No. 10, pp. 739-770, Oct. 1978.

[18] National Semiconductor, Data Acquisition Handbook, 1980

[19] Duttweiler, D.L., Messerschmitt, D.G. "Analysis of digitally Generated Sinusoids with Application to A/D and D/A Testing." IEEE Trans. Commun., Vol. COM-26, No. 5, May 1978.

[20] Analog Devices, Data Acquisition Products Catalog, 1978.

[21] Van de Plassche, R.J., Schouwenaars, H.J. "A Monolithic High-Speed Sample-and-Hold Amplifier for Digital Audio." IEEE J. of Solid-State Circuits, Vol. SC-18, No. 6, Dec. 1983.

[22] Stafford, K.R., Gray, P.R., Blanchard, R.A. "A Complete Monolithic Sample/Hold Amplifier." IEEE J. Solid-state Circuits, Vol. SC-9, No.6, Dec. 1974.

[23] Kuleshov, V., Nelson, D. "A First-Passage Statistics of a First-Order Markov Process with a Linear Ramp." Res. Lab. Elec. Quart. Progress Report, Vol. 97, pp. 181-184, April 1970.

[24] AES Recommended Practice for Professional Digital Audio Applications Employing Pulse-Code Modulation - Preferred Sampling Frequencies, AES5-1984, J. Audio Eng. Soc., Vol. 32, No. 10, Oct. 1984.

[25] Electronic Industries Association of Japan: STC - 007 "Consumer Use PCM Encoder-Decoder 1979-6.

[26] Maraoka, T., Yamada, Y., Yamazaki, M. "Sampling Frequency Considerations in Digital Audio." J. Audio Eng. Soc., Vol. 25, No. 4, pp. 260-264, April 1980.

[27] Stanley, W.D., Dougherty, G.R., Dougherty, R. Digital Signal Processing, 2nd ed. Reston: Prentice Hall. 1984.

[28] Bergland, G.D. "A Guided Tour of the Fast Fourier Transform." IEEE Spectrum, pp. 41-52, July 1969.

[29] Cooley, J.W., Tukey, J.W. "An algorithm for the machine calculation of complex Fourier series." Math. Comput., pp. 297-301, April 1965.

[30] Bergland, G.D. "A Fast Fourier Transform Algorithm for Real-Valued Series." Communications of the ACM, Vol. 11,No. 10, pp. 703-710, Oct. 1968.

[31] Cox, R.G. "Window Functions for Spectrum Analysis." HP Journal, pp. 2-13, June 1985.

[32] Panek, C.R., Kator, S.F. "Custom Digital Filters For Dynamic Signal Analysis." HP Journal, pp. 28-36, Dec. 1984.

[33] Schmidt, L.A. "Designing Programmable Digital Filters for LSI Implementation." HP Journal, pp. 15-23, Sept. 1978.

[34] Digital Signal Processing Committee, Programs for Digital Signal Processing, New York: IEEE Press, 1979.

[35] Crochiere, R.E., Rabiner, L.R. "Optimum FIR Digital Filter Implementations for Decimation, Interpolation, and Narrow-Band Filtering." IEEE Trans. Acoust., Speech, Signal Processing, Vol. ASSP-23, No. 5, Oct 1975.

[36] Preis, D. "Impulse Testing and Peak Clipping." J. Audio Eng. Soc., Vol. 25, pp. 2-14, Feb. 1977.

[37] Leinonen, E. Otala, M., Curl, J. "A Method for Measuring Transient Intermodulation Distortion (TIM)." J. Audio Eng. Soc., Vol. 25, pp. 170-176, April 1977.

[38] Small, R.H. "Total Difference-Frequency Distortion: Practical Measurements." J. Audio Eng. Soc., Vol. 34, No. 6, June 1986.

[39] Takahashi, S., Tanaka, S. "A new method of measuring transient intermodulation distortion: Comparison with the conventional method." J. Audio Eng. Soc., Vol. 30, pp. 10-15, Jan./Feb. 1982.

[40] Edminster, J.A. Electric Circuits, Schaums Outline Series: McGraw-Hill Inc., 1965.

[41] Epstein, J.S. "Versatlie Instrument Simplifies Dynamic Signal Analysis at Low Frequencies." HP Journal, pp. 4-11, Dec. 1984.

[42] Hewlett-Packard, 3561A Dynamic Signal Analyser Operating Manual, Nov. 1983.

[43] Epstein, J.S., et. al. "Hardware Design for a Dynamic Signal Analyser." HP Journal, pp. 12-20, Dec. 1984.

[44] EXAR data sheet, "XR-2211 FSK Demodulator/Tone Decoder."

[45] National Semiconductor data sheet, "MF6 6th Order Switched Capacitor Butterworth Lowpass Filter."

[46] Hewlett Packard, "HP Installation and Reference Manual, HP98640A 7-channel Analog Input Interface Hardware", july 1984.

[47] Galloway, A.J. "An Audio Quality Test Set." B.Sc. Eng. thesis submitted to the University of Cape Town, Nov. 1984.

[48] Burr-Brown data sheet, "PCM75 16-bit Hybrid Analog-to-Digital Converter."

[49] Kaiz, S.J. "A 16-bit A-D-A Conversion System for High-Fidelity Audio Research." IEEE Trans. ASSP, Vol. Assp-23, No. 1, Feb. 1975.

[50] Sony data sheet, "CX 20018 Dual 16-bit, 44 KHz., Multiplexed A/D."

[51] Cooley, J.W., Lewis, P.A., Welch, P.D. "The finite fourier transform." IEEE Trans.Audio Electroacoust., Vol. Au-17, pp. 77-85, June 1969.

[52] Harris, F.J. "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform." IEEE Procs., Vol. 66, No. 1, pp. 51-83, Jan. 1978.

[53] National Semiconductor, Logic Data Book, 1982.

[54] Siliconix, Analog switch and IC product data book, Jan. 1982

APPENDIX A

USER'S GUIDE

OVERVIEW

The test unit consists of the Transmitter and the Receiver sections. Tests can be run in a timer mode or a single test set off by the User directly (for use in a laboratory environment).

The unit can use a particular run to calibrate itself as well. This is simply an option in the set-up menu.

There are also two programs to do initial fine tuning of the trigger circuitry at the Receiver.

RUNNING THE SYSTEM

Transmitter:

1) Load and run the program "TX2" (disables keyboard) or "TX1" (keyboard enabled).
2) Respond to the initial two questions. Note that tests are only done in mono.
3) Select "TIMING" to send tests at specified times or "CONTIN" to immediately start sending a test sequence.

Receiver:

1) Load and run "RX2"
2) Respond to the initial set-up questions. Again note that only the mono mode is allowed.
3) If the system needs to be re-calibrated for some reason then respond in the affirmative to this question. Note that the calibration run must only be done in the 'CONTIN' mode. If done in the 'TIMING' mode every ensuing run will re-calibrate the system.
4) If a printout is desired, connect a printer to the HPIB and set to address 701.

5) Select 'TIMING' or 'CONTINUE' as for the transmitter unit.

Timing Mode:

1) At both the transmitter and receiver the computer real-time clocks must be set when in this mode.
2) Then the time for the start of the sequence of tests must be entered. Note that this must be set for a future time as tests will only start when the clock reaches this set value.
3) The User will then be asked to enter a repeat time for the tests. This must be greater than or equal to .5 of an hour i.e. tests can be sent at a maximum of one test sequence every 1/2 hour.
4) The final time asked for stops tests being sent between this time and midnight. At midnight tests start again at the repetition rate specified earlier. This is used to stop tests being sent in peak listening time.
5) Once in 'TIMER MODE' the units can be left unattended to perform tests at the desired times.

NOTE: IT IS ESSENTIAL TO SET THE SAME TIMES AT BOTH THE RECEIVER AND THE TRANSMITTER.
The units have been programmed to take account of this.

Contin mode:

1) Set the receiver to this mode before doing so at the transmitter so that the receiver does not miss the calibration trigger. This is due to the transmitter sending this trigger as soon as it enters this mode.
2) Again make sure, if a printout is desired, that a printer is connected at address 701.

Note that this mode was only added to do one off tests, in a laboratory for instance.

Calibration

If a calibration run is requested the data file 'CAL' is updated. At the same time a copy of the old version of 'CAL' is made to 'CAL_BACK'. This feature has been included so that it is possible to recover from an accidental re-calibration of the system. Should this happen it is only necessary to copy 'CAL_BACK' to 'CAL'.

TUNING THE TRIGGER CIRCUIT

This tuning is done by varying the potentiometer on Board 3 of the Receiver unit. By running 'SYNC_SET' at the transmitter and 'SYNC_CHK' at the receiver, synchronising bursts will be sent every 2 seconds and the receiver display whether or not it is receiving them. The two extremes where triggering justs starts to occur should thus be found and the potentiometer set to the mid-point of these.

APPENDIX B

LOW PASS DECIMATION FILTER DESIGN

REQUIREMENT

Due to the low frequency spectra (40 Hz, 125 Hz.) containing harmonic components between 15 KHz. and 21 KHz, it is necessary to low-pass filter before performing the decimation. The worst of these, at 15.6 KHz, is due to the scan-rate of the transmitter computer and was found to have a worst case amplitude of -50 dB. The noise floor is at -85 dB. Therefore, to satisfy the -80 dB. design requirement [14], a filter that has rolled off by at least 30 dB. at 15 KHz. is required. A second-order analog Butterworth filter was chosen as the reference for the digital filter as it satisfies this requirement. Such a filter type is very appropriate due to it being maximally flat, this being of more interest here than the large transition band. As there is only a -85 dB. noise floor between the fundamentals (40, 125 Hz.) and their harmonics, and 15 KHz, the important criterion here is sufficient attenuation above the latter frequency.

The project specification only requires measurement of the second and third harmonics in the system but the filter, being conservative, caters for the fourth as well. This therefore sets the maximum frequency of interest at 500 Hz. i.e. 125x4 Hz. By choosing a 3 dB. cut-off frequency of 2.5 Khz, the filter response is flat to < .05 dB. over the range 0 to 500 Hz. (see figure B2).

DESIGN

Assign

$f_c = 2500$: the filters' cut-off frequency.
$l_r = 1$: the normalised cut-off frequency of the reference analog filter.
v_r : the corresponding digital cut-off.

C : mapping constant from analog to digital domain

f_s = 44433 Hz : system sampling rate.

f_o = 22217 Hz : foldover frequency.

For a cut-off of 2.5 KHz the normalised digital cut-off frequency is the ratio

$$v_r = f_c/f_o = 2500/22217 = 0.112526 \quad \ldots\ldots\ldots\ldots (1)$$

This determines the mapping constant

$$C = 1_r COT(\pi/2)v_r = 5.598471 \quad \ldots\ldots\ldots\ldots\ldots\ldots (2)$$

Now at 15 KHz. the filter must have rolled off by 30 dB. at least. This corresponds to a normalised digital frequency of

$$v_a = 15000/22217 = 0.675 \quad \ldots\ldots\ldots\ldots\ldots\ldots (3)$$

and a normalised analog frequency of

$$1_a = C.TAN(\pi/2)v_a = 10.0 \quad \ldots\ldots\ldots\ldots\ldots\ldots (4)$$

As this is well out of the transition band of the analog filter, and from figure B1, a second-order transfer function with a roll-off of 40 dB./decade is required. A first-order section, with its 20 dB./decade, will not suffice.

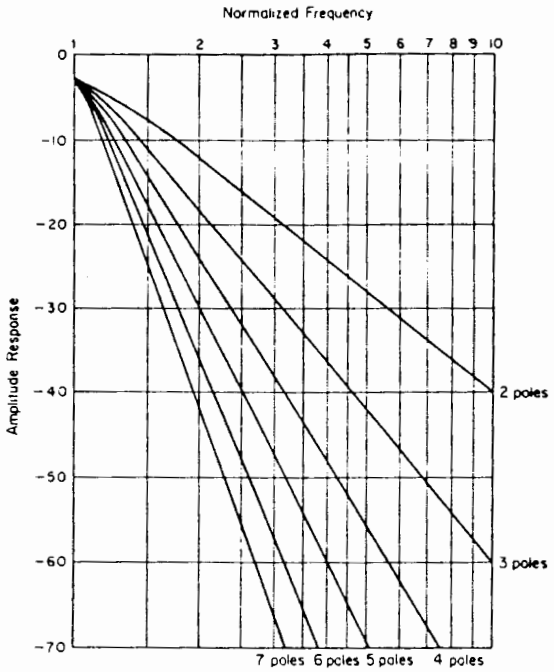

FIGURE B1 Normalised amplitude response vs. Butterworth filter order

The analog transfer function is therefore

$$G(s) = A_0/(B_0+B_1s+B_2s^2)$$

where: $A_0 = B_0 = 1$

$B_1 = 1.414214$

$B_2 = 1$

.....(5)

Using the Bilinear transformation [27]

$$H(z) = (a_0 + a_1z^{-1} + a_2z^{-2})/(1 + b_1z^{-1} + b_2z^{-2})$$

$$= a_0(1 + 2z^{-1} + z^{-2})/(1 + b_1z^{-1} + b_2z^{-2})$$

where: $A = B_0 + B_1C + B_2C^2$

$= 40.260309$

$a_0 = (A_0 + A_1C + A_2C^2)/A$

$= 0.024838$

$a_1 = 2a_0$

$a_2 = a_0$

$b_1 = (2B_0 - 2B_2C^2)/A$

$= -1.507335$

$b_2 = (B_0 - B_1C + B_2C^2)/A$

$= 0.606688$

.....(6)

Note that by substituting $z = e^{j\pi v}$ where $v = f/22217$ the magnitude of the filter response can be computed. A program was specifically written for an HP-15C to compute this and the plot is shown below.

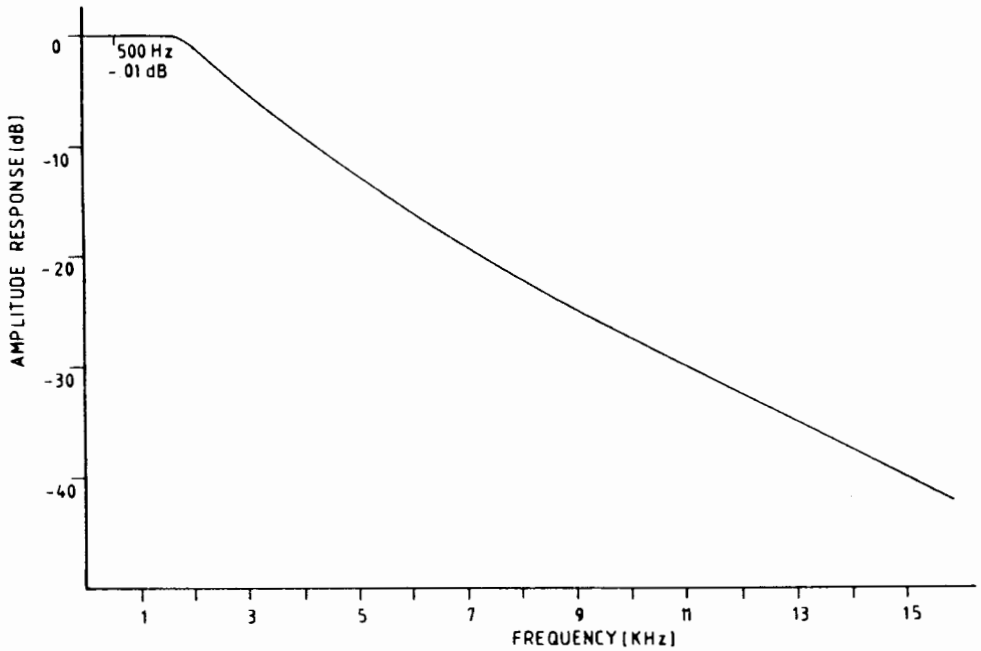

FIGURE B2 Amplitude response of digital filter

From the z-domain transfer function, H(z), the difference equation is

$$y(n) = -b_1y(n-1) - b_2y(n-2) + a_0x(n) + a_1x(n-1) + a_2x(n-2)$$

where: n = n'th sample
y(n) = output sample at n
x(n) = input sample at n
.....(7)

IMPLEMENTATION

To show actual implementation of a difference equation, pseudo-code for equation (7) is included here.

As the array of samples contains integers it is necessary to check for possible overload in the filter due to reals being used inside the filter, but the output being integers. Single variables have been used wherever possible in the routine to

avoid the extra time taken by the computer when accessing arrays.

Note that real numbers, with a 53 bit mantissa for the HP series 9000 computers, are used in the prototype and therefore word-length effects do not exist.

Because the filter has a finite buffer of samples to process, a priori data, i.e. before the first sample, is not available. All samples before time n=0 are, therefore, set to zero. The array of input samples is put through the filter with the output samples being written back into the same array.

```
REAL A0, B1, B2
A0 = .0248384
B1 = -1.5073345
B2 = 0.6066879
X(n-1) = 0  ! Sample before present one being worked on.
X(n)   = 0  ! Present input sample.
Y(n-1) = 0  ! Latest output before present one being computed.
Y(n)   = 0  ! present output being computed.

FOR I = 1 TO 1024
  X(n-2) = X(n-1)
  X(n-1) = X(n)
  X(n)   = Samples(I) ! Samples contains the 1 K
                      INTEGER samples.

  Y(n-2) = Y(n-1)
  Y(n-1) = Y(n)

  Y(n)   = A0(X(n) + 2X(n-1) + X(n-2))
  Y(n)   = Y(n) - B1Y(n-1) - B2Y(n-2)

  IF Y(n) > 32767 THEN Y(n) = 32767
  IF Y(n) < -32767 THEN Y(n) = -32767
  Samples(I) = Y(n)
NEXT I
```

RESULTS

Using the HP Waveform Analysis Pack the two plots below were obtained. They are both of a 40 Hz. fundamental sampled at 44.4 KHz. The unfiltered spectrum is shown below.

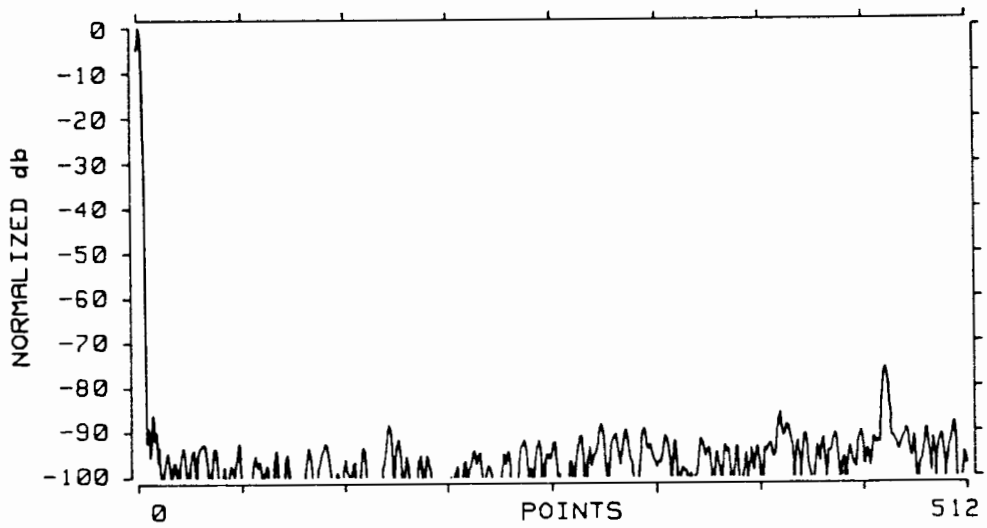

FIGURE B3 Unfiltered waveform

After passing the above sampled waveform through the filter the following response was obtained.

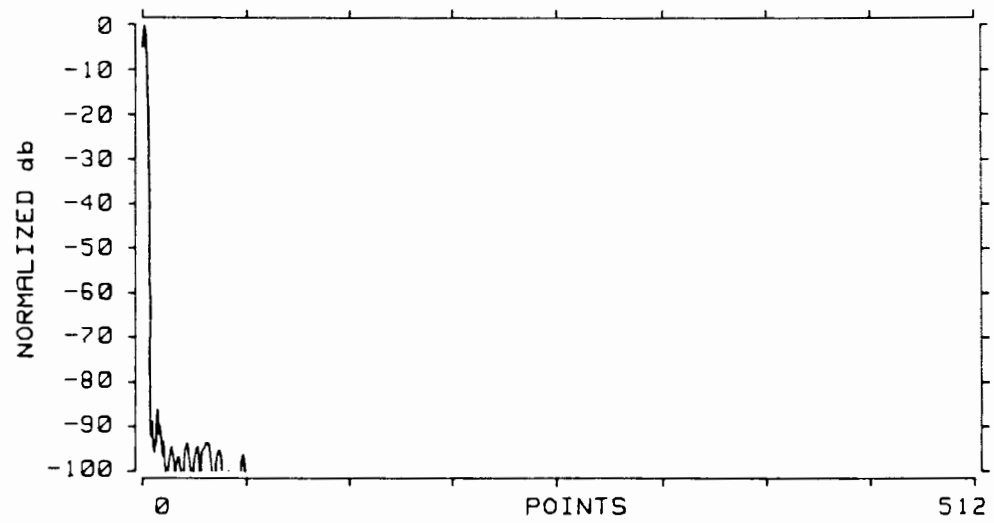

FIGURE B4 Filtered waveform

APPENDIX C

DIGITAL TIMING GENERATION

INTRODUCTION

Due to the need for minimal jitter of the sampling pulse to the sample-and-hold of the analog-to-digital converter board, the digital generation design was done so as to include the crystal clock as an input to the final gating element generating the pulse. In this way jitter is limited to that of the clock circuitry and this gating element.

By using a 4 MHz clock and dividing down by 90 the desired sampling rate of 44.4 KHz is obtained.

TIMING

In the following discussion refer to figures F1 and F2:

The 4 MHz. clock circuitry output ('CLOCK') is first divided by 10 ('U10'). and then by 9 ('U12'). The clock period is therefore 250 ns., the period of the first counter 2.5 us. with a 250 ns. pulse and the period of the second counter 22.5 usecs. The divide-by-10 counter output ('COUNTER1') is used to clock the divide-by-9 counter.

When the first counter reaches count 3 this state is detected ('COUNT3'). The 5 ('COUNT5') and 7 ('COUNT7') counts of the second counter are also detected. To limit the length of these pulses they are gated with the output of the first counter ('COUNTER1') in 'U14' to give the shortened 5 count pulse ('AND5') and shortened 7 count pulse ('AND7'). This is necessary so that the 'J' input of the 'SAMPLE' pulse flip-flop has returned to logic 0 by the time it is reset. The 'AND5' pulse is used to set the sampling pulse flip-flop ('U15'); 'AND7' resets the same flip-flop while at the same time setting the convert pulse flip-flop ('U15'). The latter flip-flop is reset by the 'COUNT3' pulse from the first counter. The-sample-and hold samples while its control line is

high and is in hold when this is low. The ADC starts conversion on the nagative edge of 'CNVRT'.

The race condition between 'COUNTER1' and 'COUNT5' at 'U14' causes a glitch on 'AND5' which is ignored by the remaining circuitry. This is due to the two counters being positive edge triggered whereas the flip-flops are negative edge-triggered. The maximum length of the glitch is approximately 52 ns., this being set by the propogation delay from the clock input of the counter to its terminal count output, 'COUNTER1' [53]. The maximum turn on and turn off delays for the AND gate is 20 ns. and 15 ns. respectively [53]. Hence, for the worst case, the glitch on 'AND5' will end approximately 70 nsecs. after the positive edge of the system clock ('CLOCK'). As the time between edges of the system clock is 125 nsecs. this leaves a set-up time for the flip-flop of 55 nsecc. before the negative edge of the clock. As the flip-flops only require a minimum set-up time of 20 ns. the glitch will not have any effect on the 'SAMPLE' line. Similarly, the glitch on 'AND7' will not be seen on 'CNVRT'.

The positive edge of the sample and hold control pulse ('SAMPLE') is used to strobe data into the computer's General Purpose I/O interface. Although the control pulses to the ADC are continuous, the strobe pulses to the computer are gated so that the ADC is only read when there is a tone present. 'GATE2' is used to control the NOR gate ('U16') to facilitate this. While 'GATE2' is high no strobe pulses are sent to the interface as 'STROBE' is kept low.

DIGITAL TIMING DIAGRAM

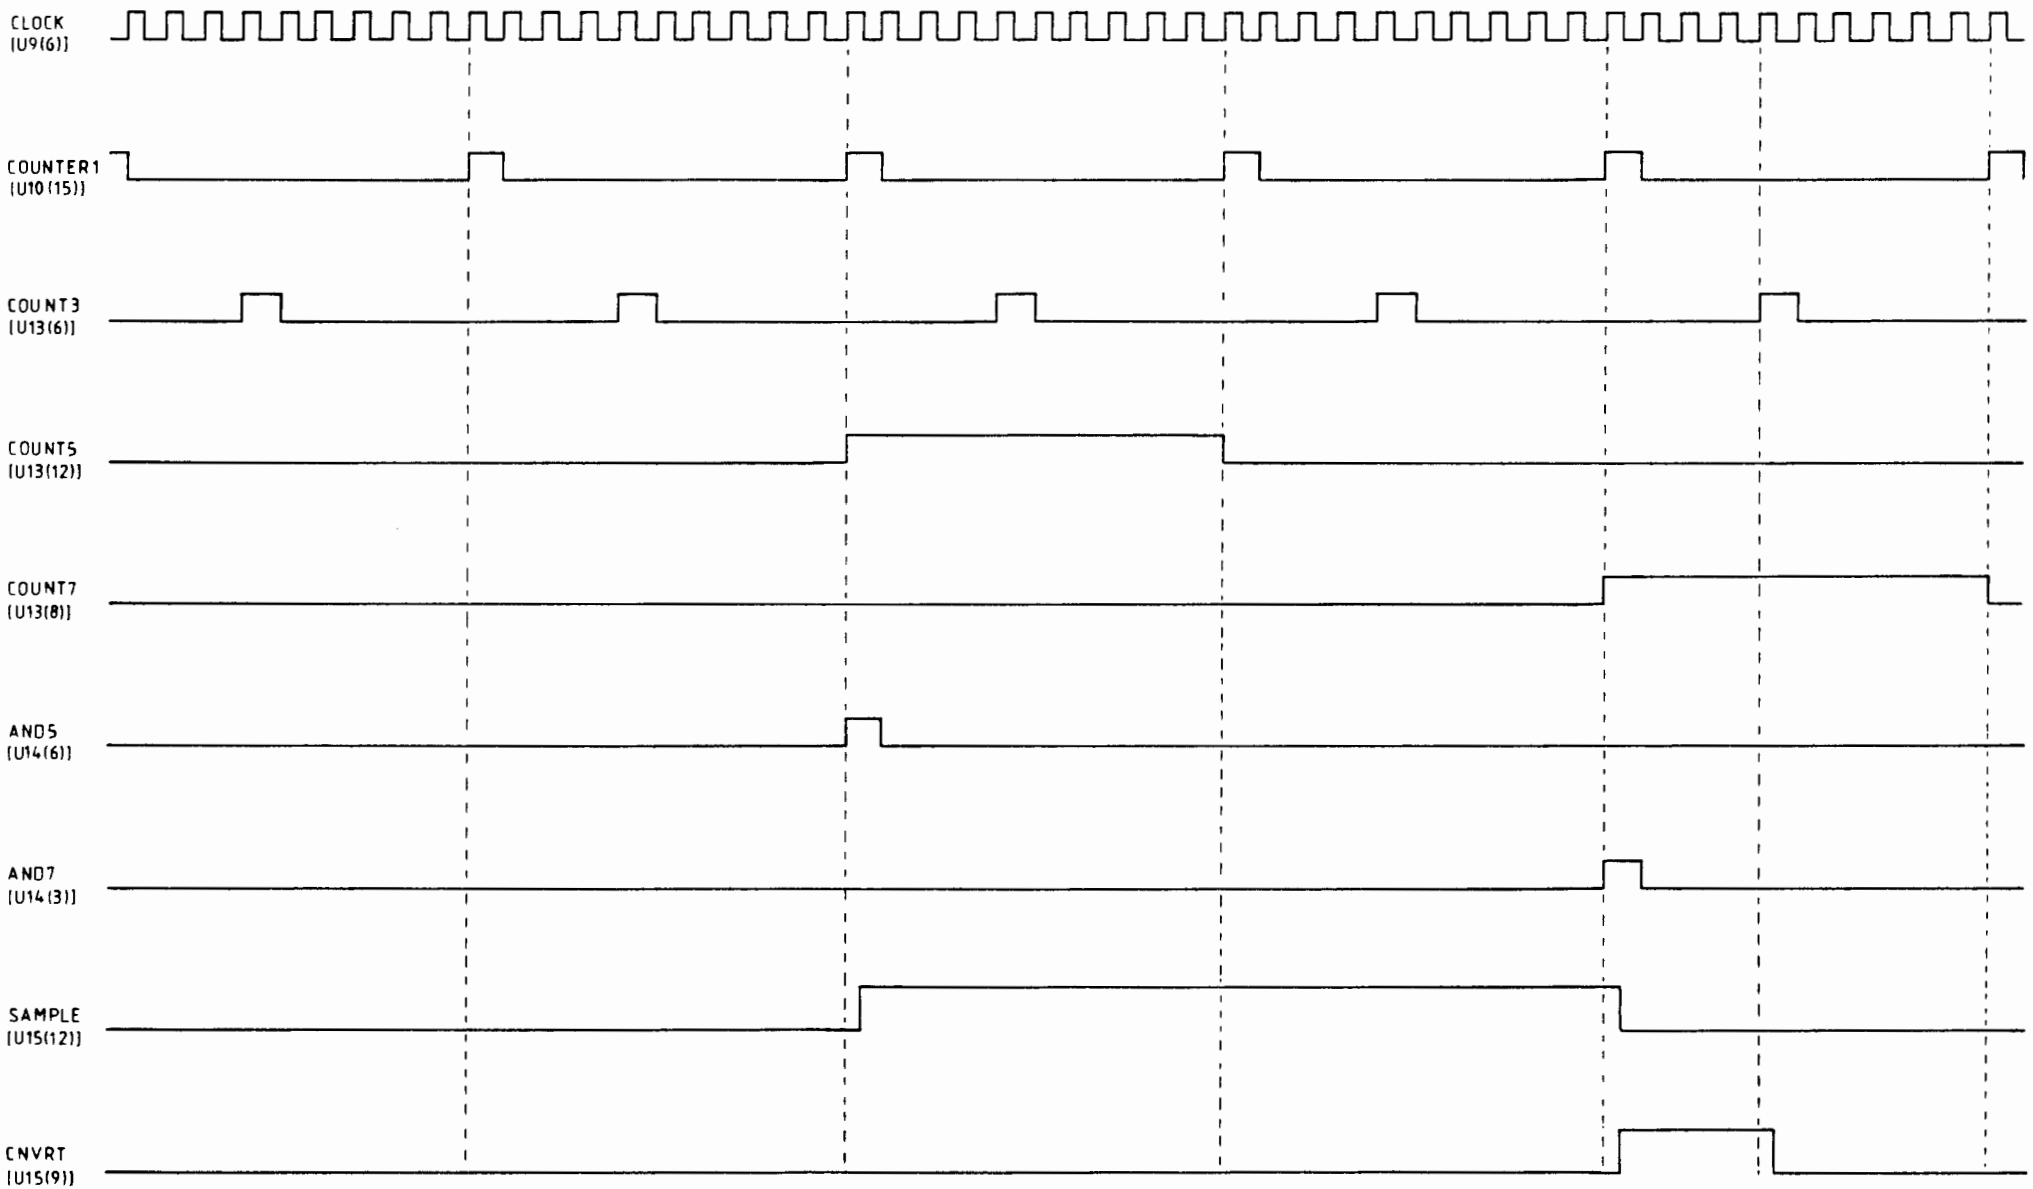

FIGURE C1 Digital timing diagram

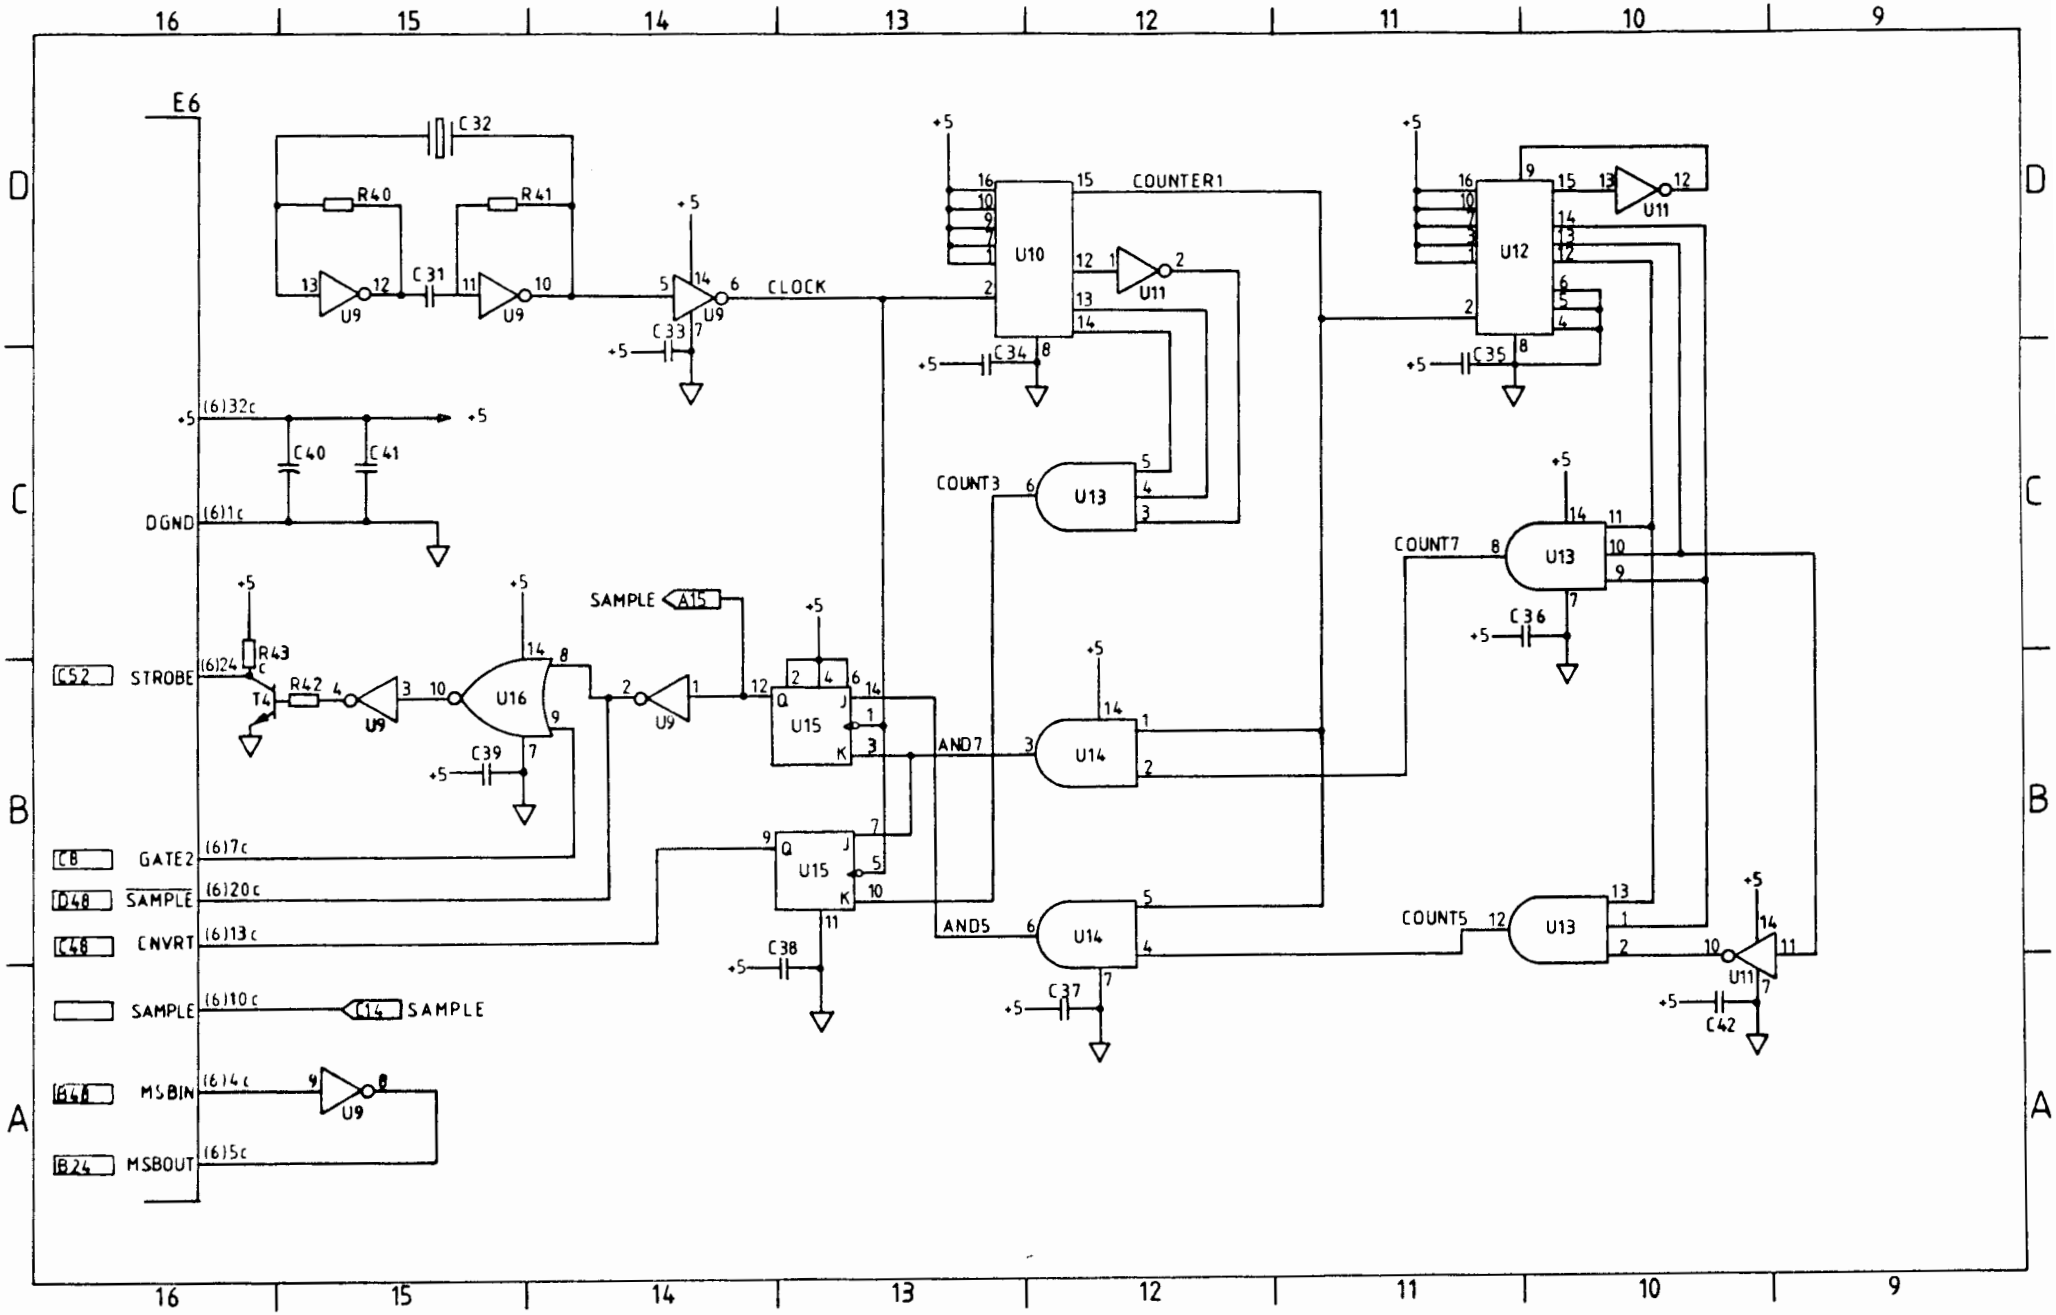

FIGURE C2 Digital generation circuitry

APPENDIX D

OPERATION OF RECEIVER CIRCUITRY

BOARD 1:

This consists of the sample-and-hold, analog-to-digital converter and latches. The latches are used to keep the data word constant after conversion by the ADC. The top of the board is covered with a ground plane connected to analog ground. This is crucial in limiting noise.

Inputs

1) Analog and digital grounds and -15, +15 and +5 Volt supplies, all being zenered to protect the ADC.
2) The gated received tone ('SIGNAL') is fed into the discrete sample-and-hold.
3) The inverse of the sampling pulse ('SAMPLE').
4) The convert command for the ADC, also used by the latches ('CNVRT').

Outputs

1) The 16 bit unipolar, complemented data word.

Sample and Hold

Due to the excellent THD. (Total Harmonic Distortion) requirement of the sample-and-hold, a high quality analog switch is required for the discrete S&H. For this reason 'U29' was used. However, the gating pulse recquired by this device is not TTL compatable. Hence level shifting was required while still maintaining sharp edges on the pulse. Hence the sampling pulse ('SAMPLE') is first shifted before being fed to the switch.
The S&H samples for 5 usecs. and is then allowed to settle for 1 usec. before A/D conversion begins.

SYSTEM LAYOUT

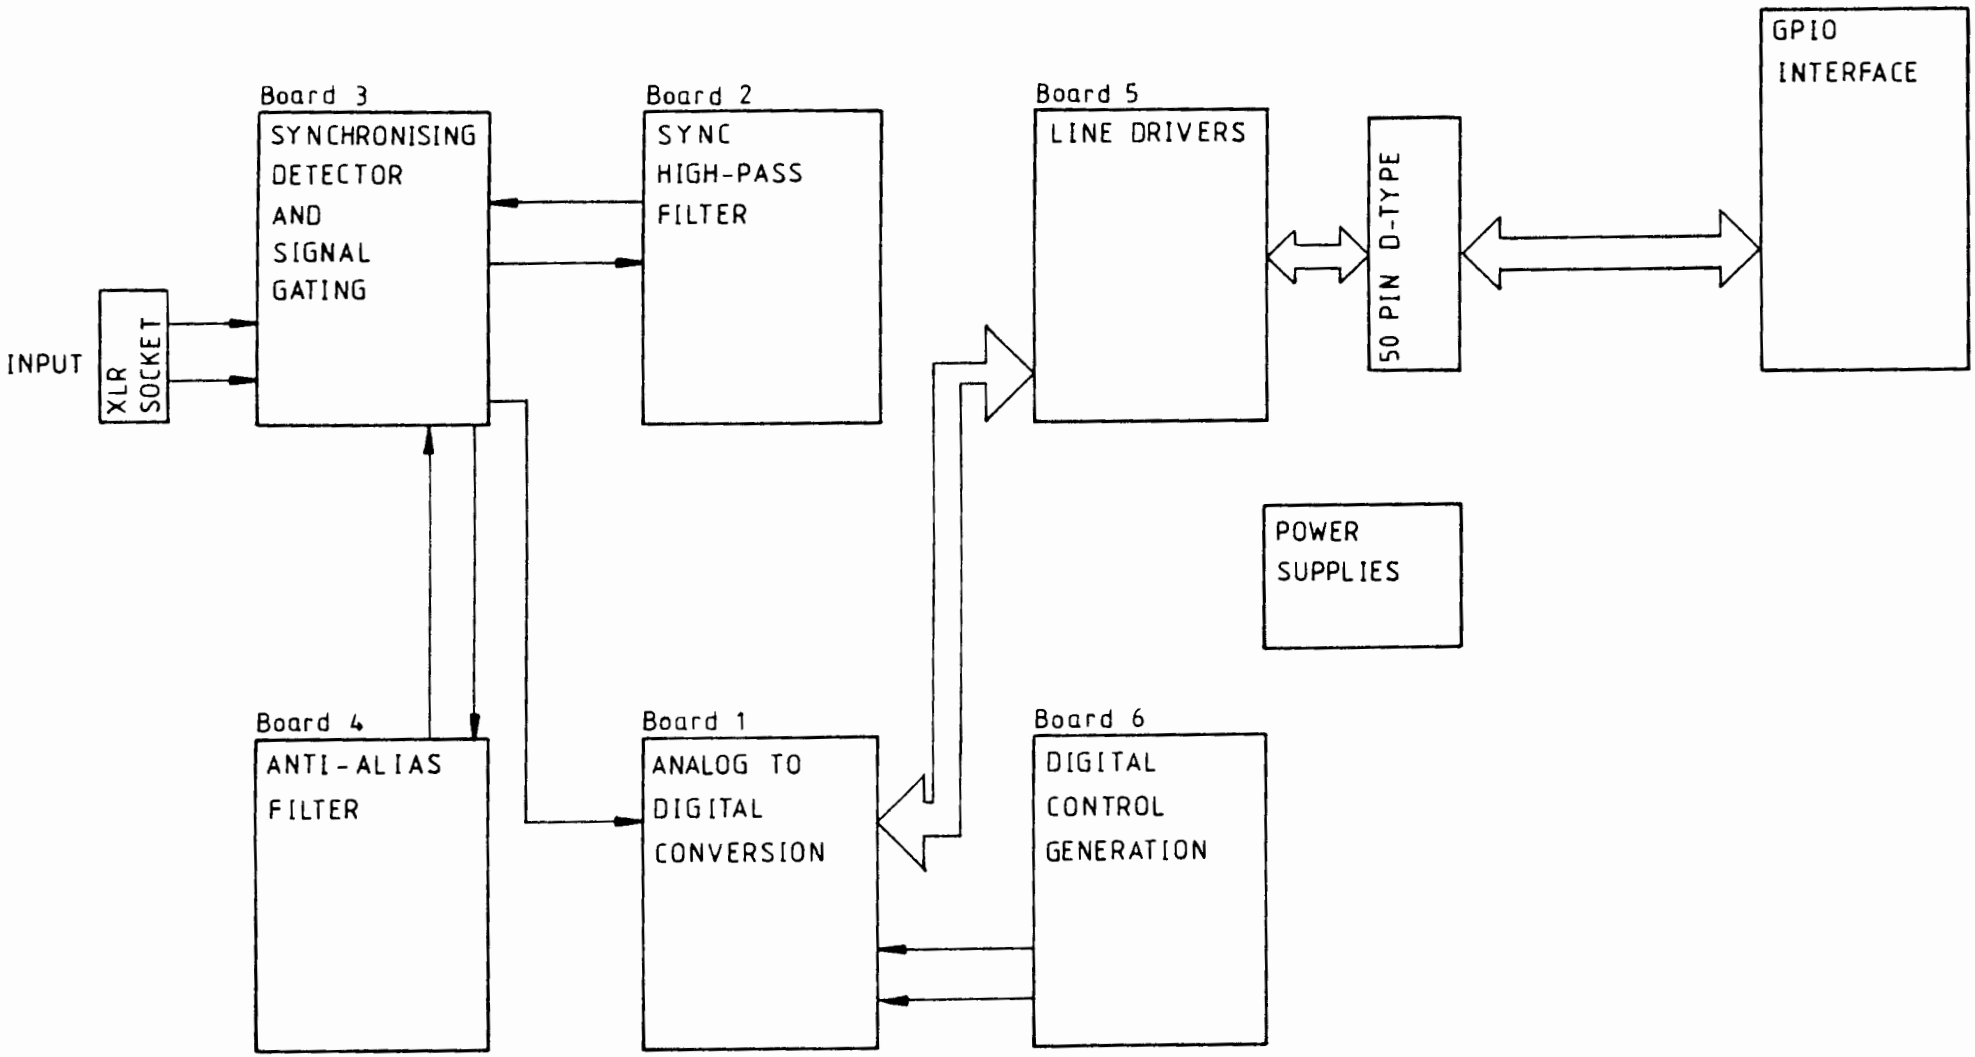

Analog to Digital Convertor

To keep the distortion of the ADC to a minimum its conversion time must be set to the maximum i.e. 17 usecs. At a sampling frequency of 44.4 Khz. it is necessary to make this slightly less (i.e. 16.2 usecs.) due to the 6 usecs used by the S&H. The potentiometer, R77, is used to set this so that the status pin (pin 18 - open collector) of the ADC returns low approximately 250 nano-secs. before the sample line ('SAMPLE') goes low for the start of a sampling period.

D.C. offset of the ADC can be zeroed by adjusting potentiometer R78, although this is not critical.

Because the layout of the board is critical if noise is to be minimised, the following should be borne in mind:

i) The analog and digital earths should only be joined at one place and as close to the ADC as possible.
ii) The decoupling should be placed as close as possible to the relevant pins of the ADC.
iii) An analog-earth guard ring must be layed around pins 26 and 27 (the comparator input).

Using these techniques the prototype has a noise floor of -85 dB. with an input signal at -10 dB. of FSR. It is most important to only connect the analog and digital earths at the chip and nowhere else in the system.

Because the data output by the ADC is unipolar and complemented it is converted to two's complement by inverting the MSB. and then setting the GPIO of the computer to read negative-true logic.

BOARD 2:

This high-pass filter is used to remove program material frequency components below 6.5 Khz. so that these will not cause harmonics to spuriously trigger the tone decoder, or

cause triggers to be missed, due to the limiter at its front end.

The cut-off frequency is at 18.5 Khz. , thereby ensuring that maximum attenuation occurs below the 12.75 Khz. of the synchronising tone as the roll-off below this frequency is at -40 dB./decade.

Inputs

i) Analog ground and +15, -15 Volt supplies.
ii) Single-ended program material ('PGOUT').

Outputs

1) Filtered program material ('DEC').

BOARD 3:

Here differential program material is input and converted to single-ended. This is then used by the synchronising and gain-scaling sections of the card.

Inputs

i) Balanced channel signal ('IN1' & 'IN2').
ii) Output from the high-pass filter ('DEC') into the tone decoder.
iii) Output from the anti-alias filter ('TONE') into the gain-scaling section.
iv) Analog ground and +15, -15 Volt supplies.
v) Data output lines (DI0,...,DI5) from the GPIO which select one of the available gain values.

Outputs

i) Single-ended signal ('PGOUT') to the high-pass and anti-alias filters.

ii) The gating pulse ('GATE2') which enables the strobe pulses to the computer.
iii) The gated test tone ('SIGNAL') is passed to the A/D card.

Tone Decoder

'U3' has been configured as a tone decoder to detect the synchronising pulse sent together with each tone.
The potentiometer R14 is needed to do the fine frequency adjustment of the VCO. It is important to find the two extreme positions where triggering just starts to occur and to set the potentiometer in the middle of this range. This should be done with the unit in its case and after its equilibrium temperature has been reached.

The program 'SYNC_SET' at the transmitter can be used to do this as it sends a trigger sync.-burst every two seconds.

Gain setting

Due to the test tones being sent along the channel at different levels and the need for all tones into the ADC to be at 1.5 Volts, it is necessary to apply a suitable gain factor to accomplish this. Hence the following bits of the GPIO correspond to the following transmitted gains:
+9 dBm: DI5
+6 dBm: DI4
+3 dBm: DI3
0 dBm: DI2
-6 dBm: DI1
-12 dBm: DI0

Gating

The synchronisation pulses from the tone decoder are used to generate two gating pulses, one ('GATE1') approximately 1 msec. before the other ('GATE2').
The first enables the signal ('SIGNAL') to the ADC and hence allows settling time before the second which enables strobing of the ADC by the computer.

BOARD 4:

The components of the anti-alias filter should be carefully chosen to ensure that they are to specification. The variable inductors are used to do the fine tuning, the desired response being flat to 15 Khz (within .1 dB) and ripple and noise minimised in the stopband.

The filter consists of two cascaded fifth-order elliptic sections.

Inputs

1) Analog ground and +15, -15 Volt supplies.
2) Program material from the differential to single-ended converter ('PGOUT').

Outputs

1) The output from the anti-alias filter.

BOARD 5

The line-drivers were added to eliminate bit errors at the computer due to noise and crosstalk along the GPIO cable connecting the test unit to the computer.

Inputs

1) Data word, with the MSB inverted, from the A/D card.
2) Digital ground and the +5 Volt supply.

Outputs

1) Data word to the computer via the GPIO cable.

BOARD 6:

Here the digital control signals for the A/D card are generated. The elaborate circuitry has been used to ensure that the period between control pulses is absolutely constant over a short interval. This avoids jitter problems.

Inputs

i) Digital ground and +5 Volt supply.
ii) The strobe gating pulse ('GATE2').
iii) The MSB from the A/D card.

Outputs

i) The strobe pulses ('STROBE') to the computer. The GPIO reads the data on the positive edge of this pulse (i.e. 'BSY' clock source of GPIO used).
ii) The Sample ('SAMPLE') and convert ('CNVRT') commands to the ADC.
iii) The inverted MSB ('MSBOUT') of the data word.

POWER SUPPLIES:

+5 V: 250 mA.; Regulated with 1% or less ripple.
+15, -15: 100 mA.; Regulated with 1% or less ripple.

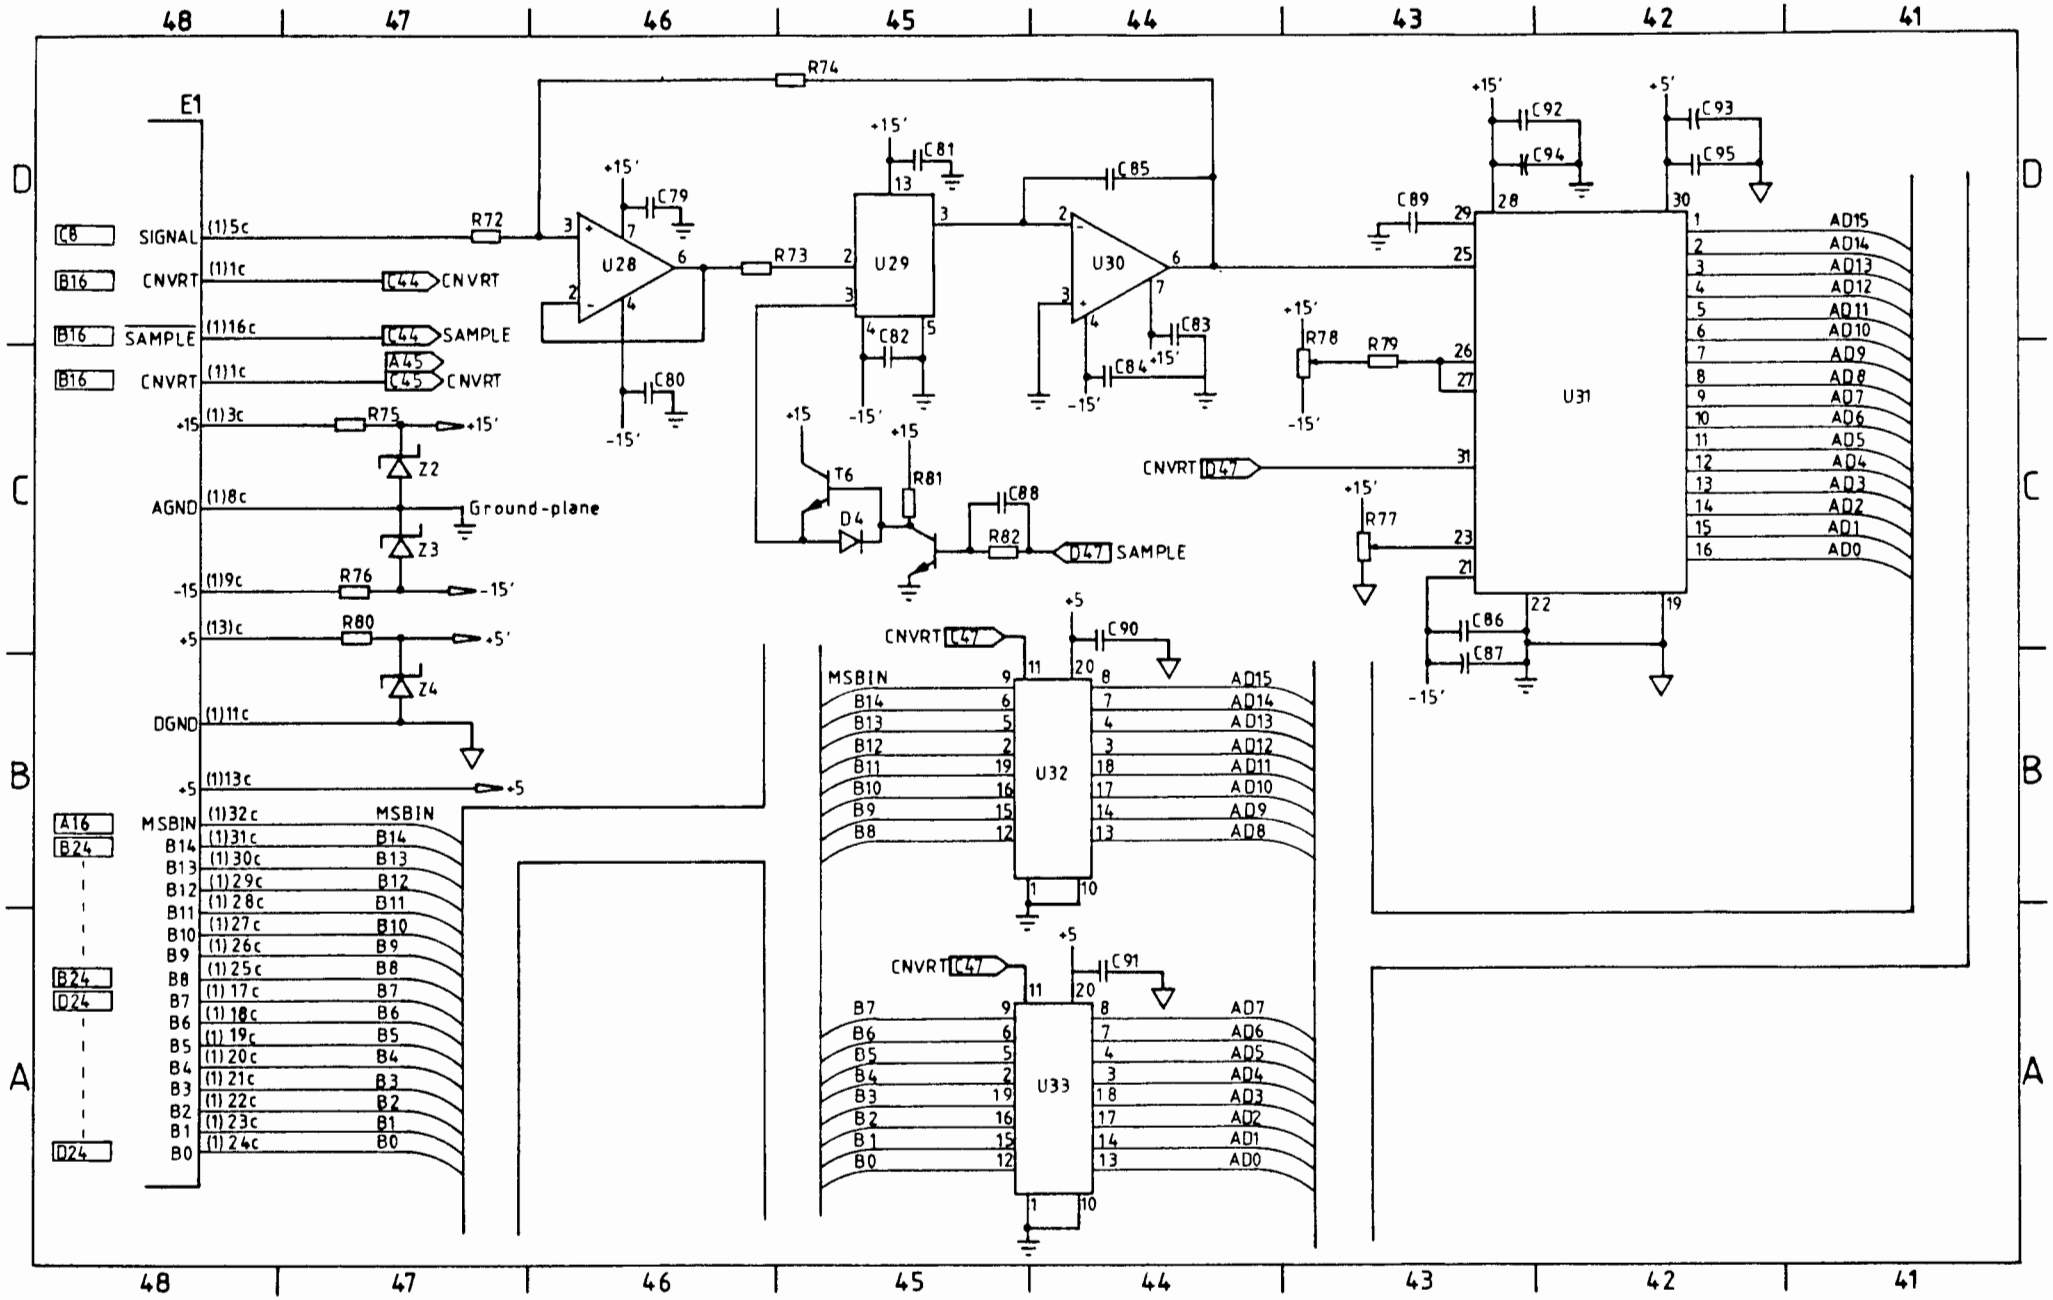
48 47 46 45 44 43 42 41
D C B A
E1
C8 SIGNAL (1)5c R72
B16 CNVRT (1)1c C44 CNVRT
B16 SAMPLE (1)16c C44 SAMPLE
B16 CNVRT (1)1c A45 C45 CNVRT
+15 (1)3c R75 +15'
AGND (1)8c Ground-plane
-15 (1)9c R76 -15'
+5 (13)c R80 +5'
DGND (1)11c
+5 (1)13c +5
A16 B24 B24 D24 D24
MSBIN (1)32c MSBIN
B14 (1)31c B14
B13 (1)30c B13
B12 (1)29c B12
B11 (1)28c B11
B10 (1)27c B10
B9 (1)26c B9
B8 (1)25c B8
B7 (1)17c B7
B6 (1)18c B6
B5 (1)19c B5
B4 (1)20c B4
B3 (1)21c B3
B2 (1)22c B2
B1 (1)23c B1
B0 (1)24c B0
Z2 Z3 Z4
R74 R73
U28 U29 U30 U31 U32 U33
C79 C80 C81 C82 C83 C84 C85
C86 C87 C88 C89 C90 C91 C92 C93 C94 C95
+15' -15' +5'
T6 D4 R81 R82
D47 SAMPLE
CNVRT D47
R77 R78 R79
MSBIN B14 B13 B12 B11 B10 B9 B8
B7 B6 B5 B4 B3 B2 B1 B0
AD15 AD14 AD13 AD12 AD11 AD10 AD9 AD8 AD7 AD6 AD5 AD4 AD3 AD2 AD1 AD0

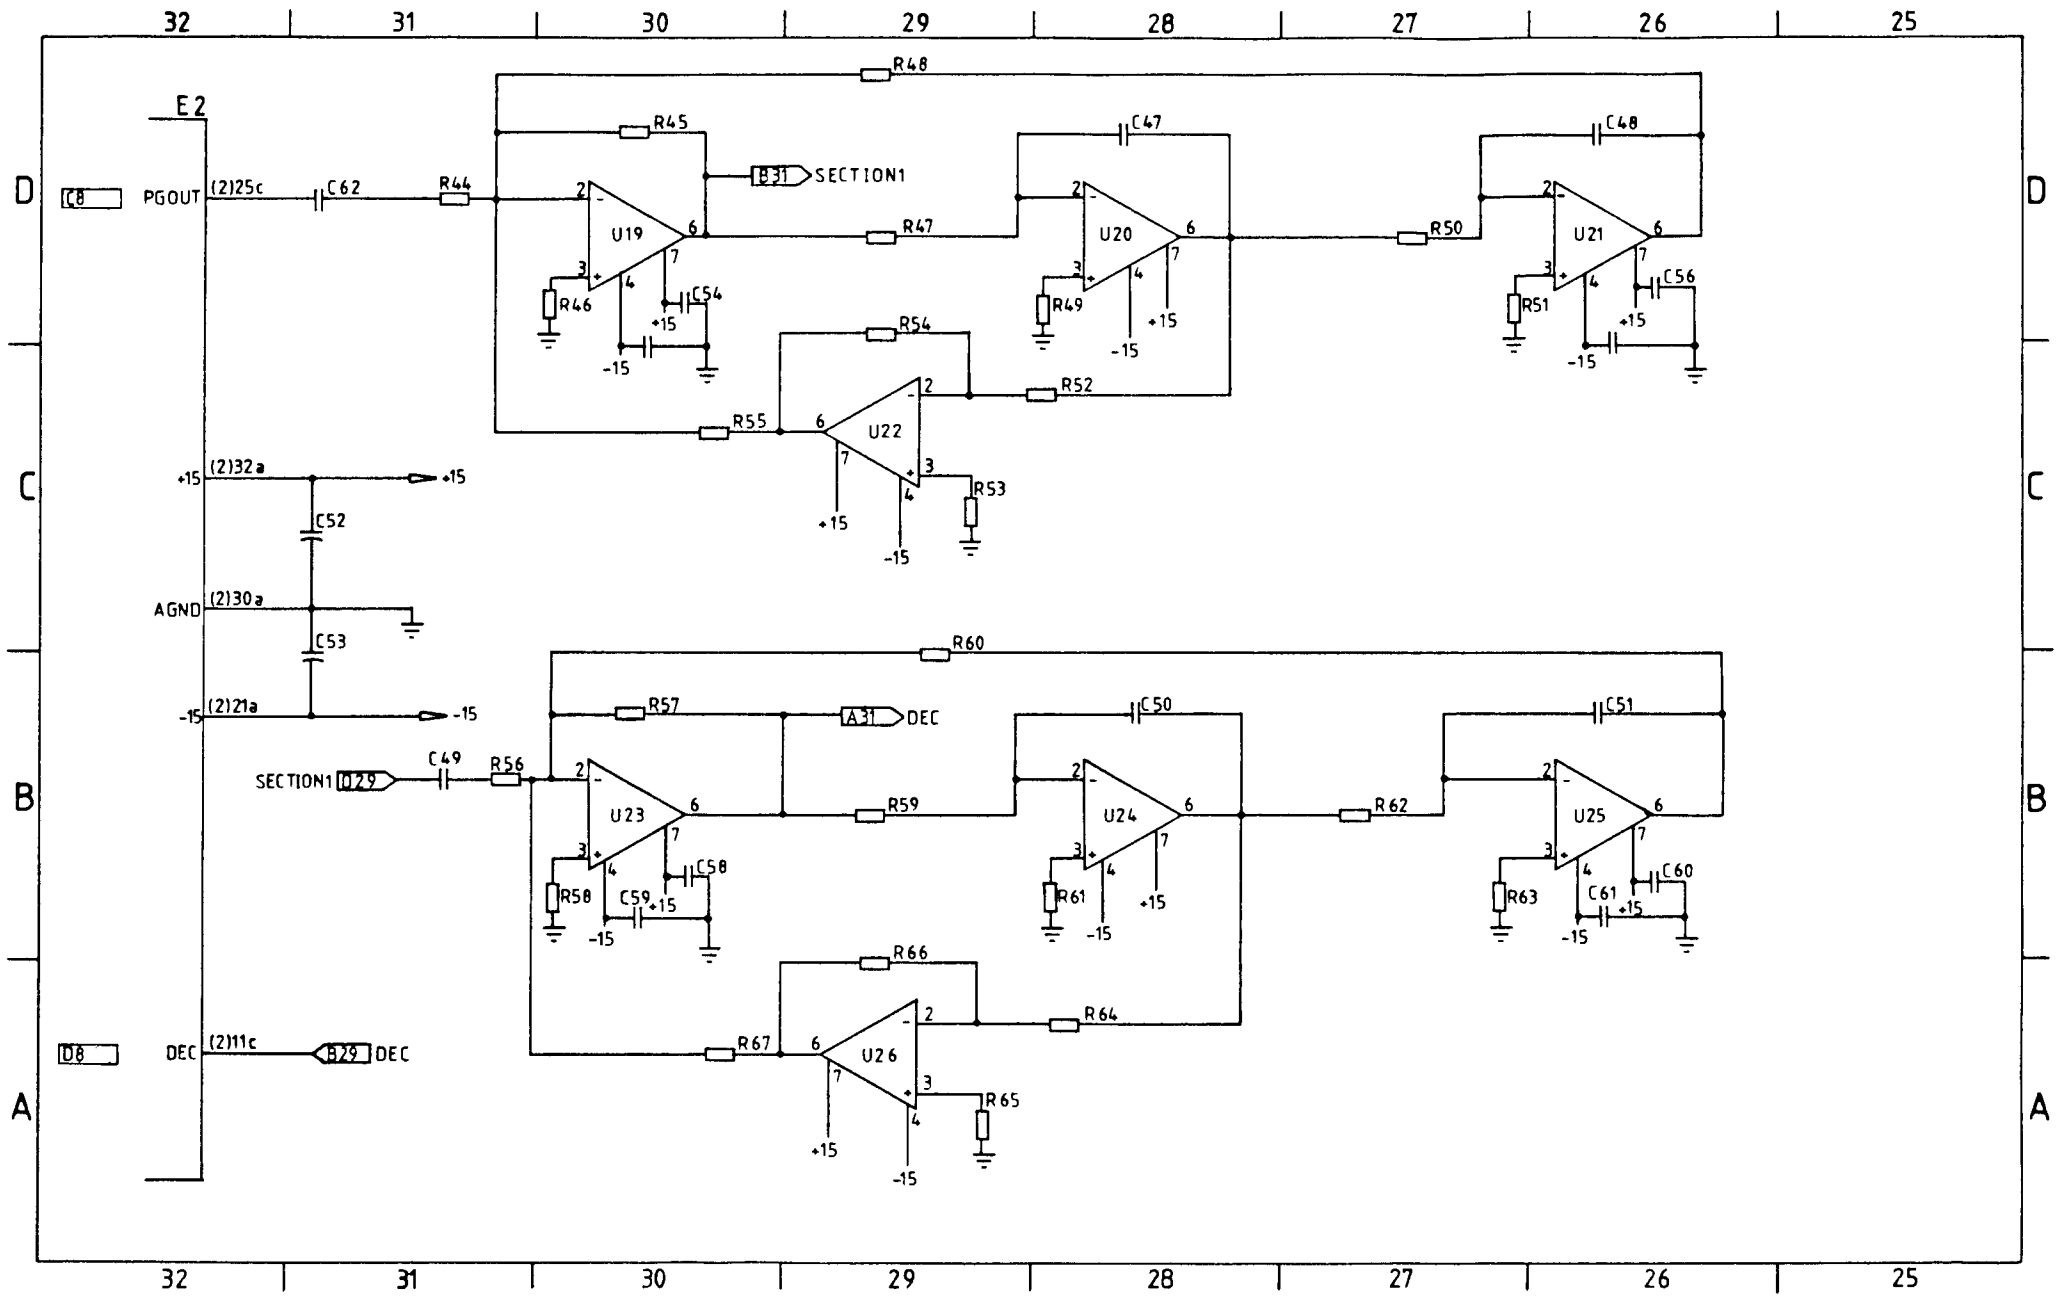
32
31
30
29
28
27
26
25
D
C
B
A
E2
C8
PGOUT
(2)25c
C62
R44
R45
R48
B31
SECTION1
U19
R46
C54
+15
-15
R47
C47
U20
R49
R50
C48
U21
R51
C56
R54
R52
R55
U22
R53
+15
(2)32a
C52
AGND
(2)30a
C53
-15
(2)21a
R60
R57
A31
DEC
SECTION1
D29
C49
R56
U23
R58
C59
C58
R59
C50
U24
R61
R62
C51
U25
R63
C61
C60
R66
R64
R67
U26
R65
D8
DEC
(2)11c
B29
DEC

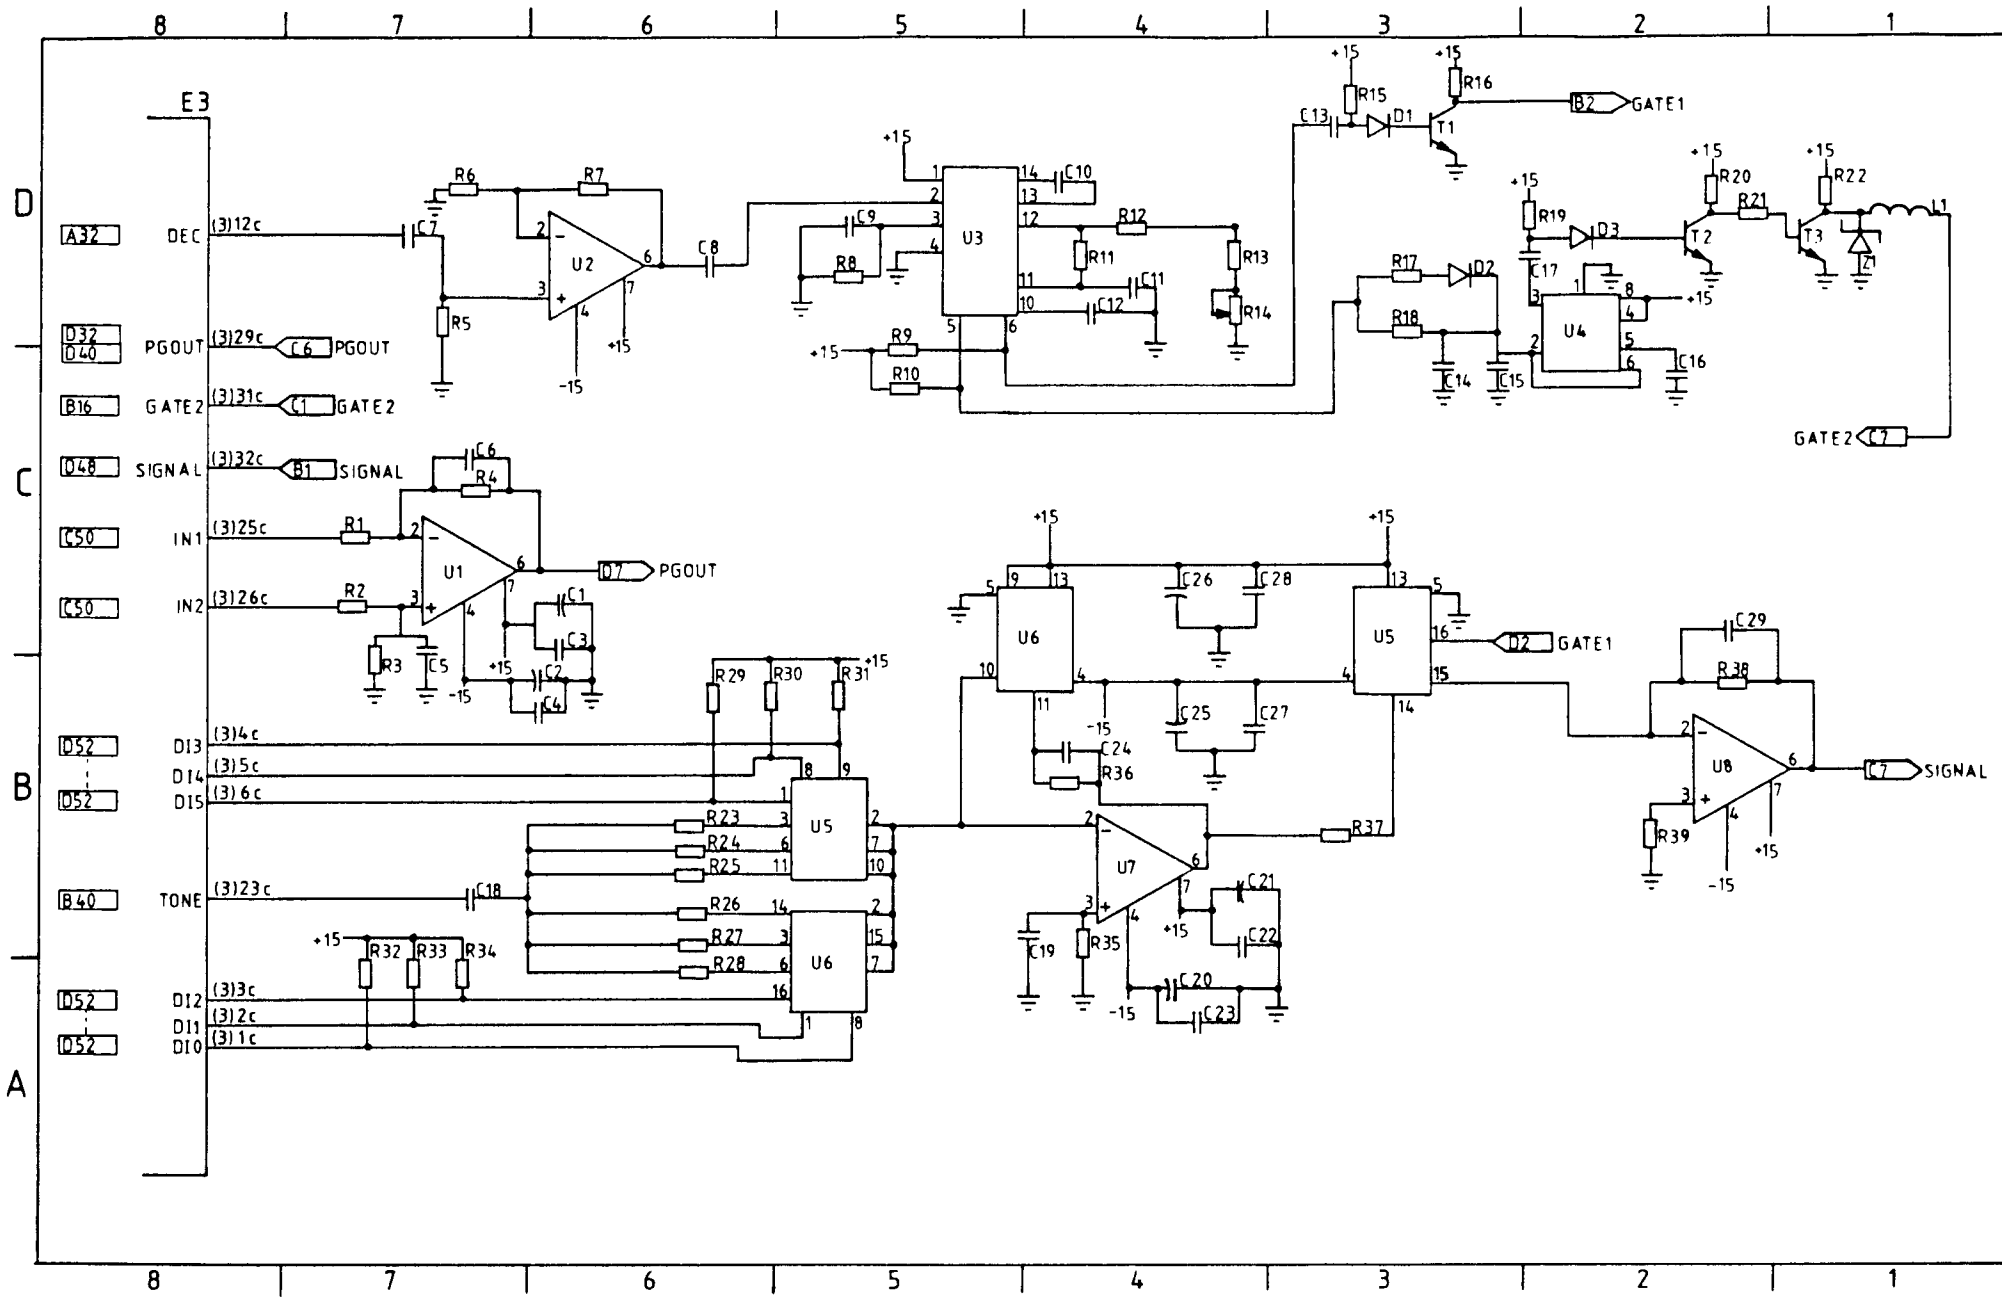
E3
A32
DEC
(3)12c
D32
D40
PGOUT
(3)29c
C6
PGOUT
B16
GATE2
(3)31c
C1
GATE2
D48
SIGNAL
(3)32c
B1
SIGNAL
C50
IN1
(3)25c
C50
IN2
(3)26c
D52
DI3
(3)4c
DI4
(3)5c
D52
DI5
(3)6c
B40
TONE
(3)23c
D52
DI2
(3)3c
DI1
(3)2c
D52
DI0
(3)1c
D7
PGOUT
B2
GATE1
D2
GATE1
C7
SIGNAL
C7
GATE2
U1
U2
U3
U4
U5
U6
U7
U8
T1
T2
T3
D1
D2
D3
L1
Z1
+15
-15

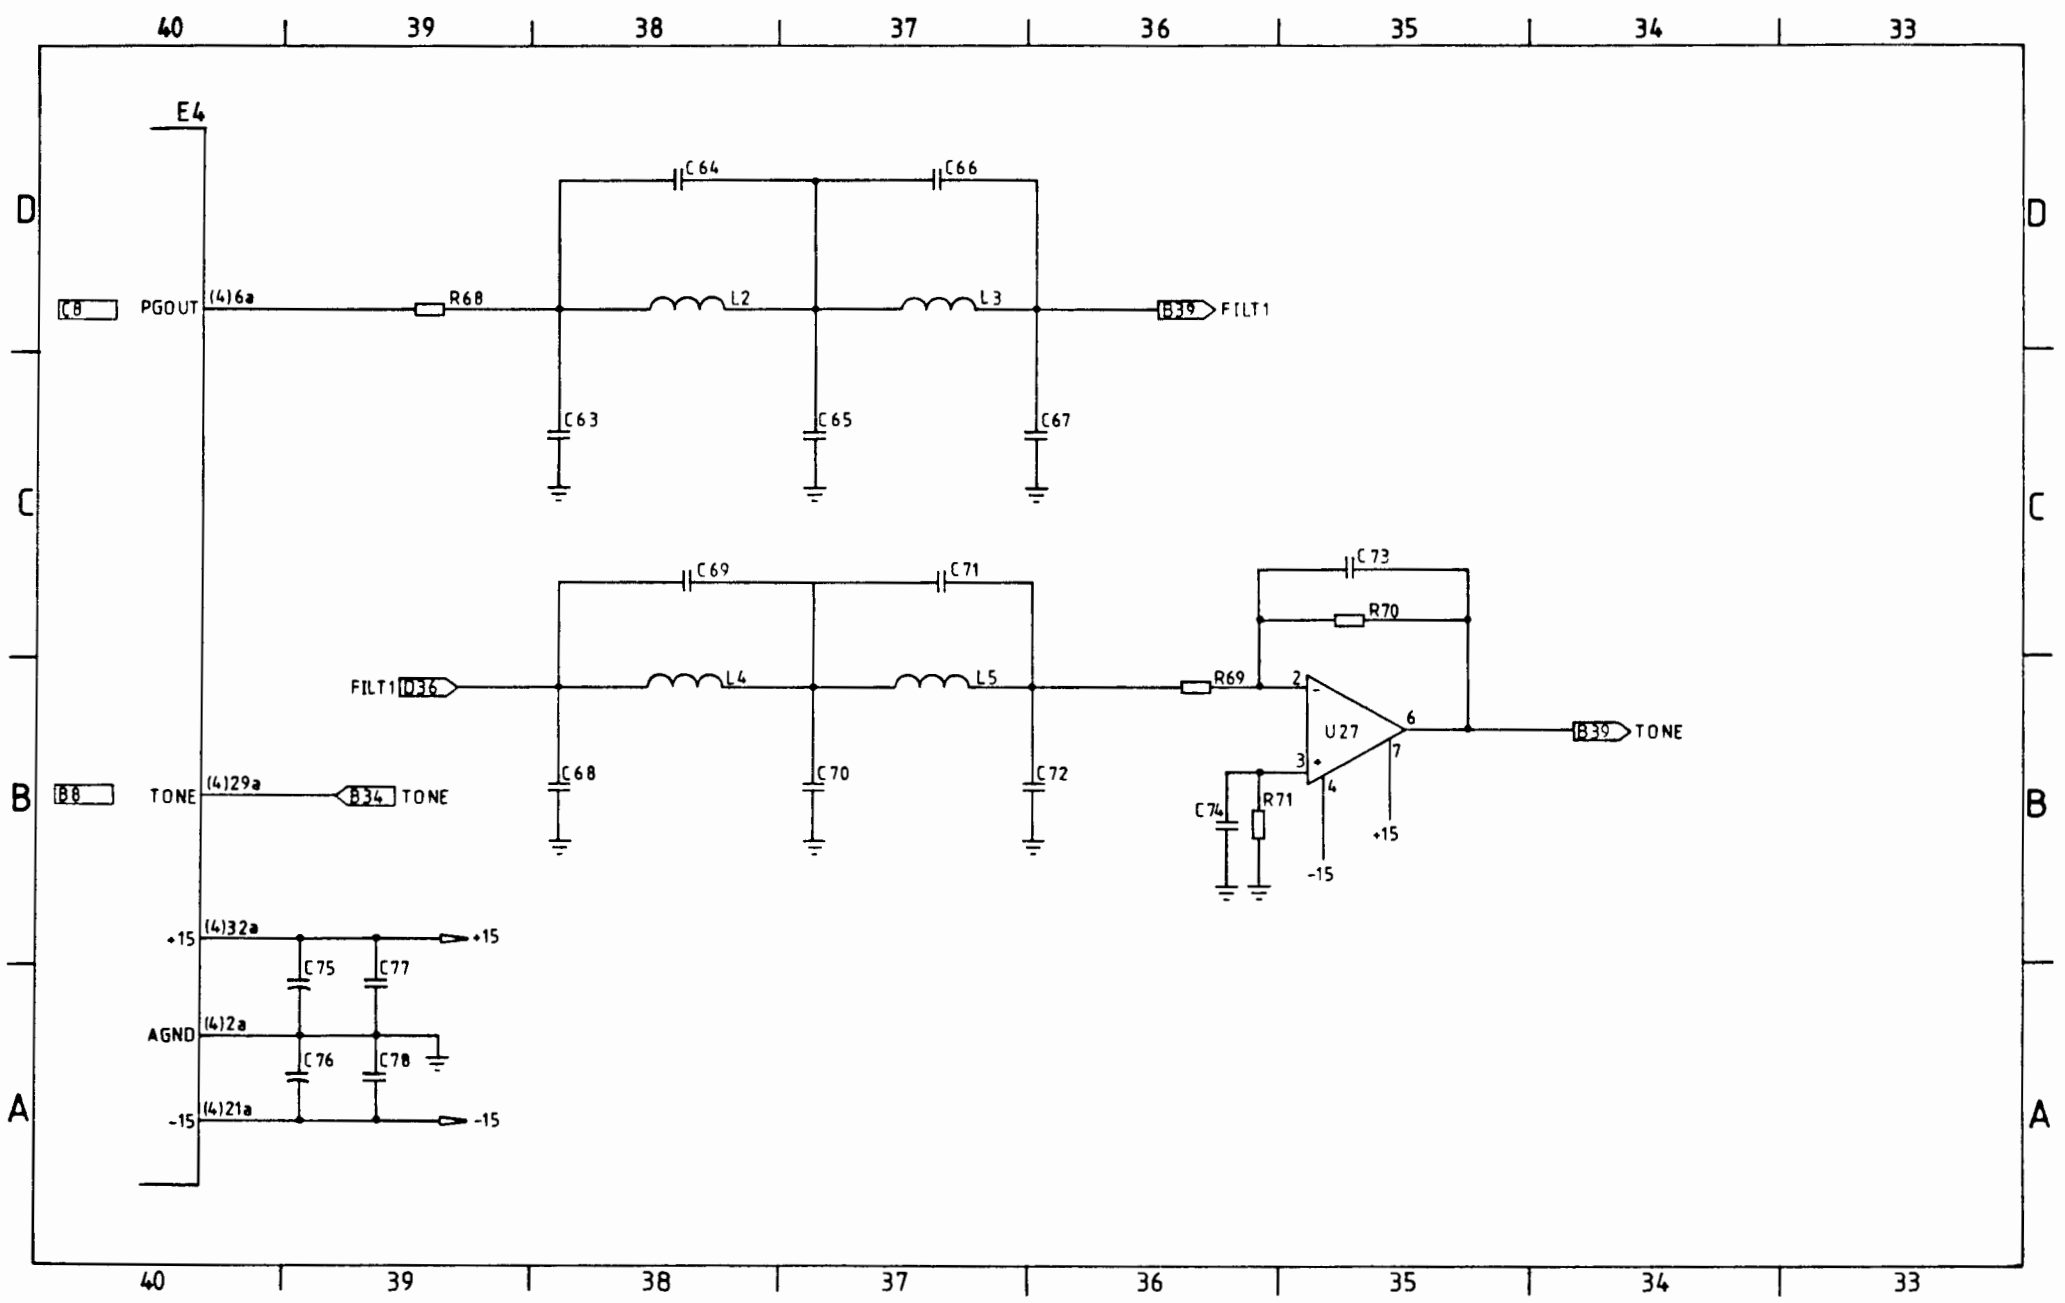
40
39
38
37
36
35
34
33
D
C
B
A
E4
C64
C66
C8
PGOUT
(4)6a
R68
L2
L3
B39
FILT1
C63
C65
C67
C69
C71
C73
R70
FILT1
D36
L4
L5
R69
2
U27
6
B39
TONE
3
7
4
C68
C70
C72
C74
R71
+15
-15
B8
TONE
(4)29a
B34
TONE
+15
(4)32a
+15
C75
C77
AGND
(4)2a
C76
C78
-15
(4)21a
-15

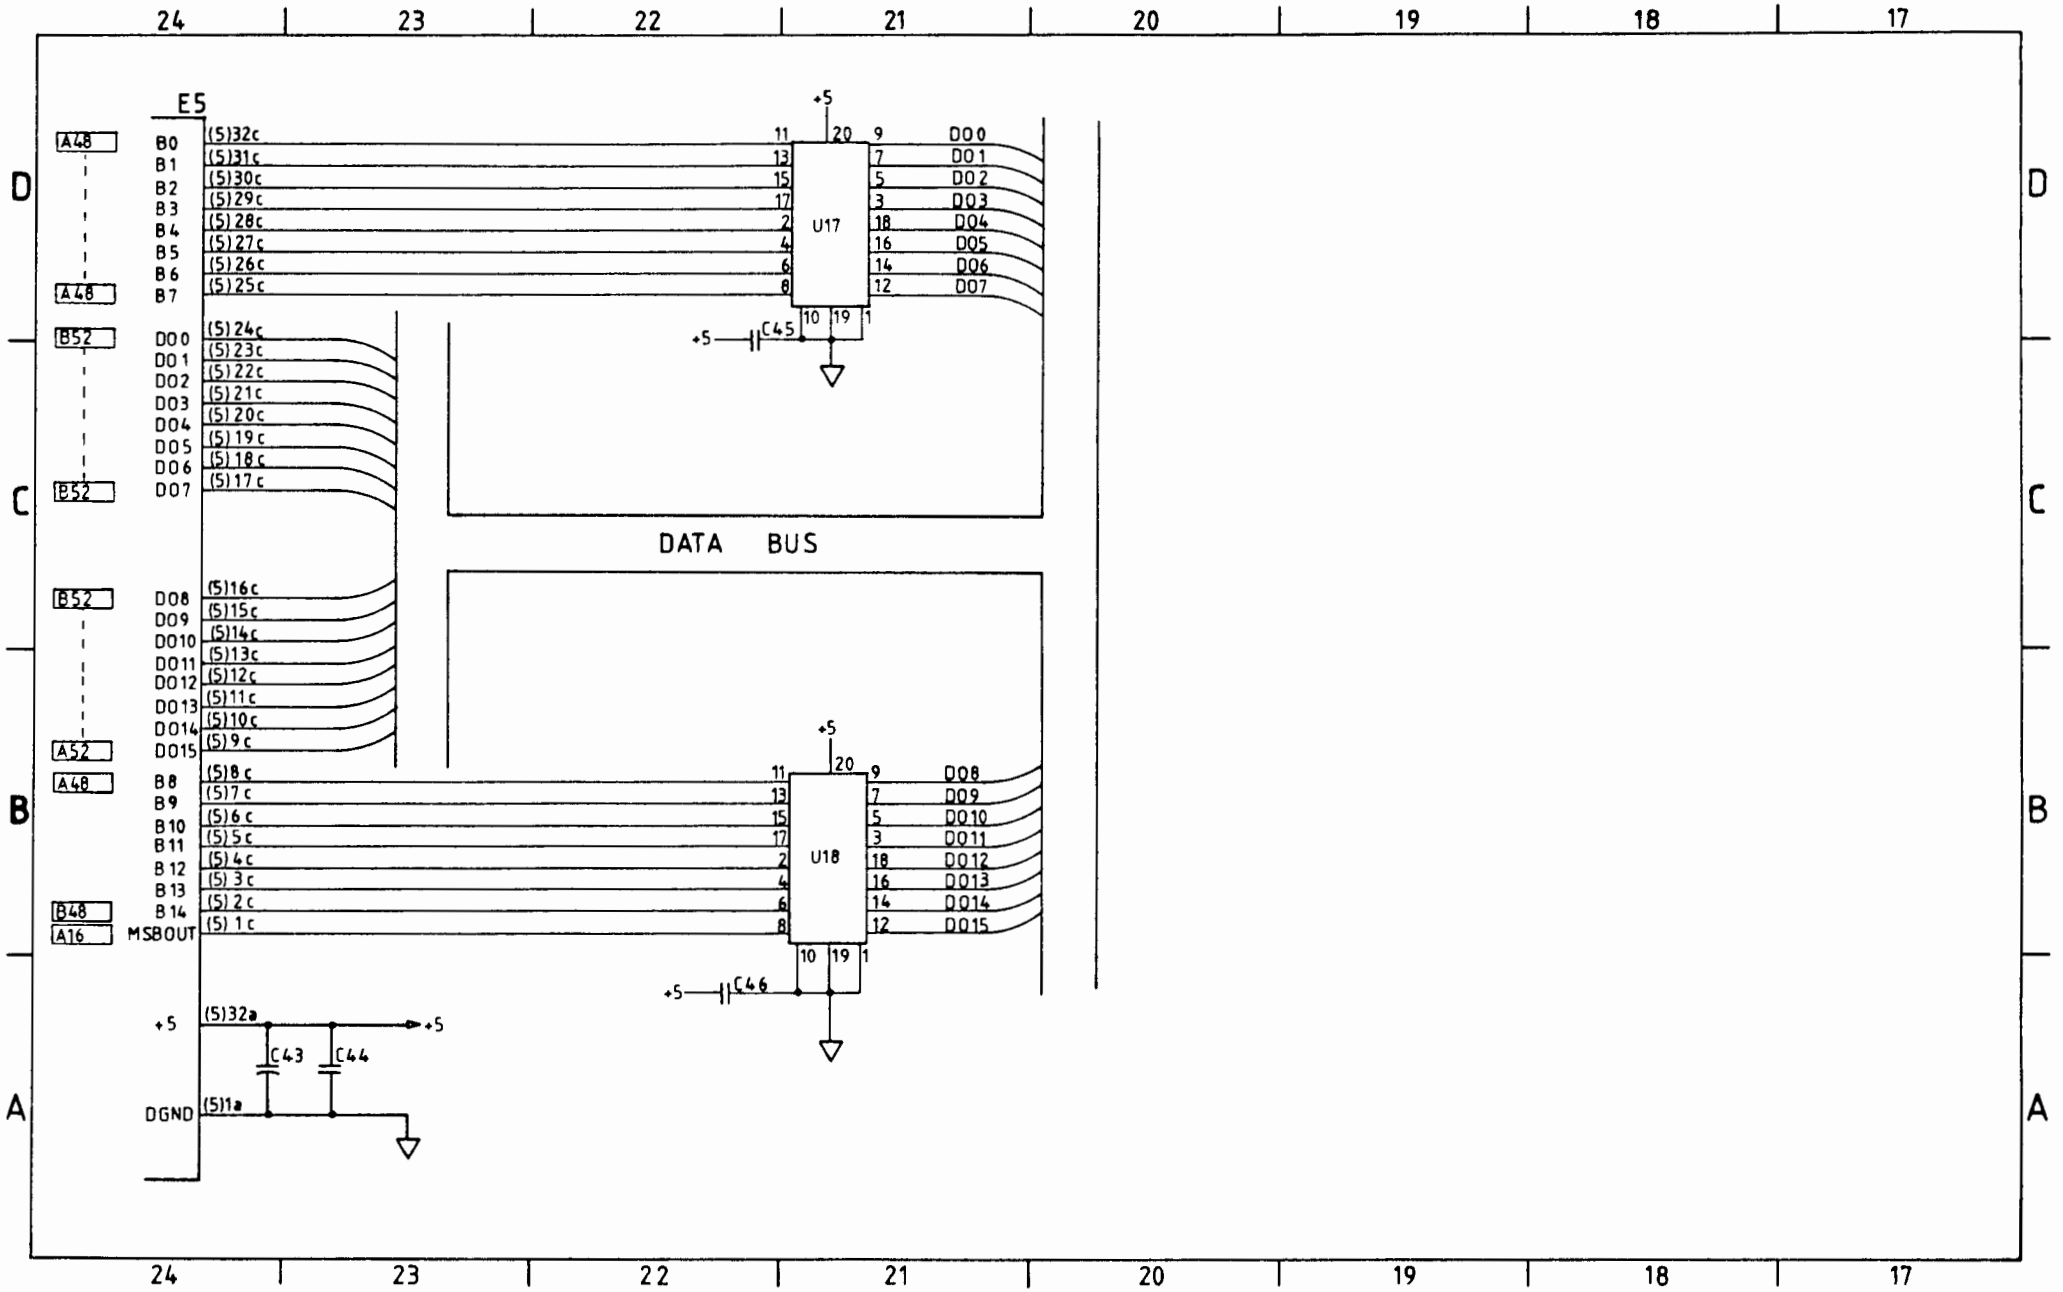
24
23
22
21
20
19
18
17
D
C
B
A
E5
A48
B0
B1
B2
B3
B4
B5
B6
A48
B7
B52
D00
D01
D02
D03
D04
D05
D06
B52
D07
(5)32c
(5)31c
(5)30c
(5)29c
(5)28c
(5)27c
(5)26c
(5)25c
(5)24c
(5)23c
(5)22c
(5)21c
(5)20c
(5)19c
(5)18c
(5)17c
+5
U17
C45
DATA BUS
B52
D08
D09
D010
D011
D012
D013
D014
A52
D015
A48
B8
B9
B10
B11
B12
B13
B48
B14
A16
MSBOUT
(5)16c
(5)15c
(5)14c
(5)13c
(5)12c
(5)11c
(5)10c
(5)9c
(5)8c
(5)7c
(5)6c
(5)5c
(5)4c
(5)3c
(5)2c
(5)1c
U18
C46
+5
(5)32a
C43
C44
DGND
(5)11a
D00
D01
D02
D03
D04
D05
D06
D07
D08
D09
D010
D011
D012
D013
D014
D015

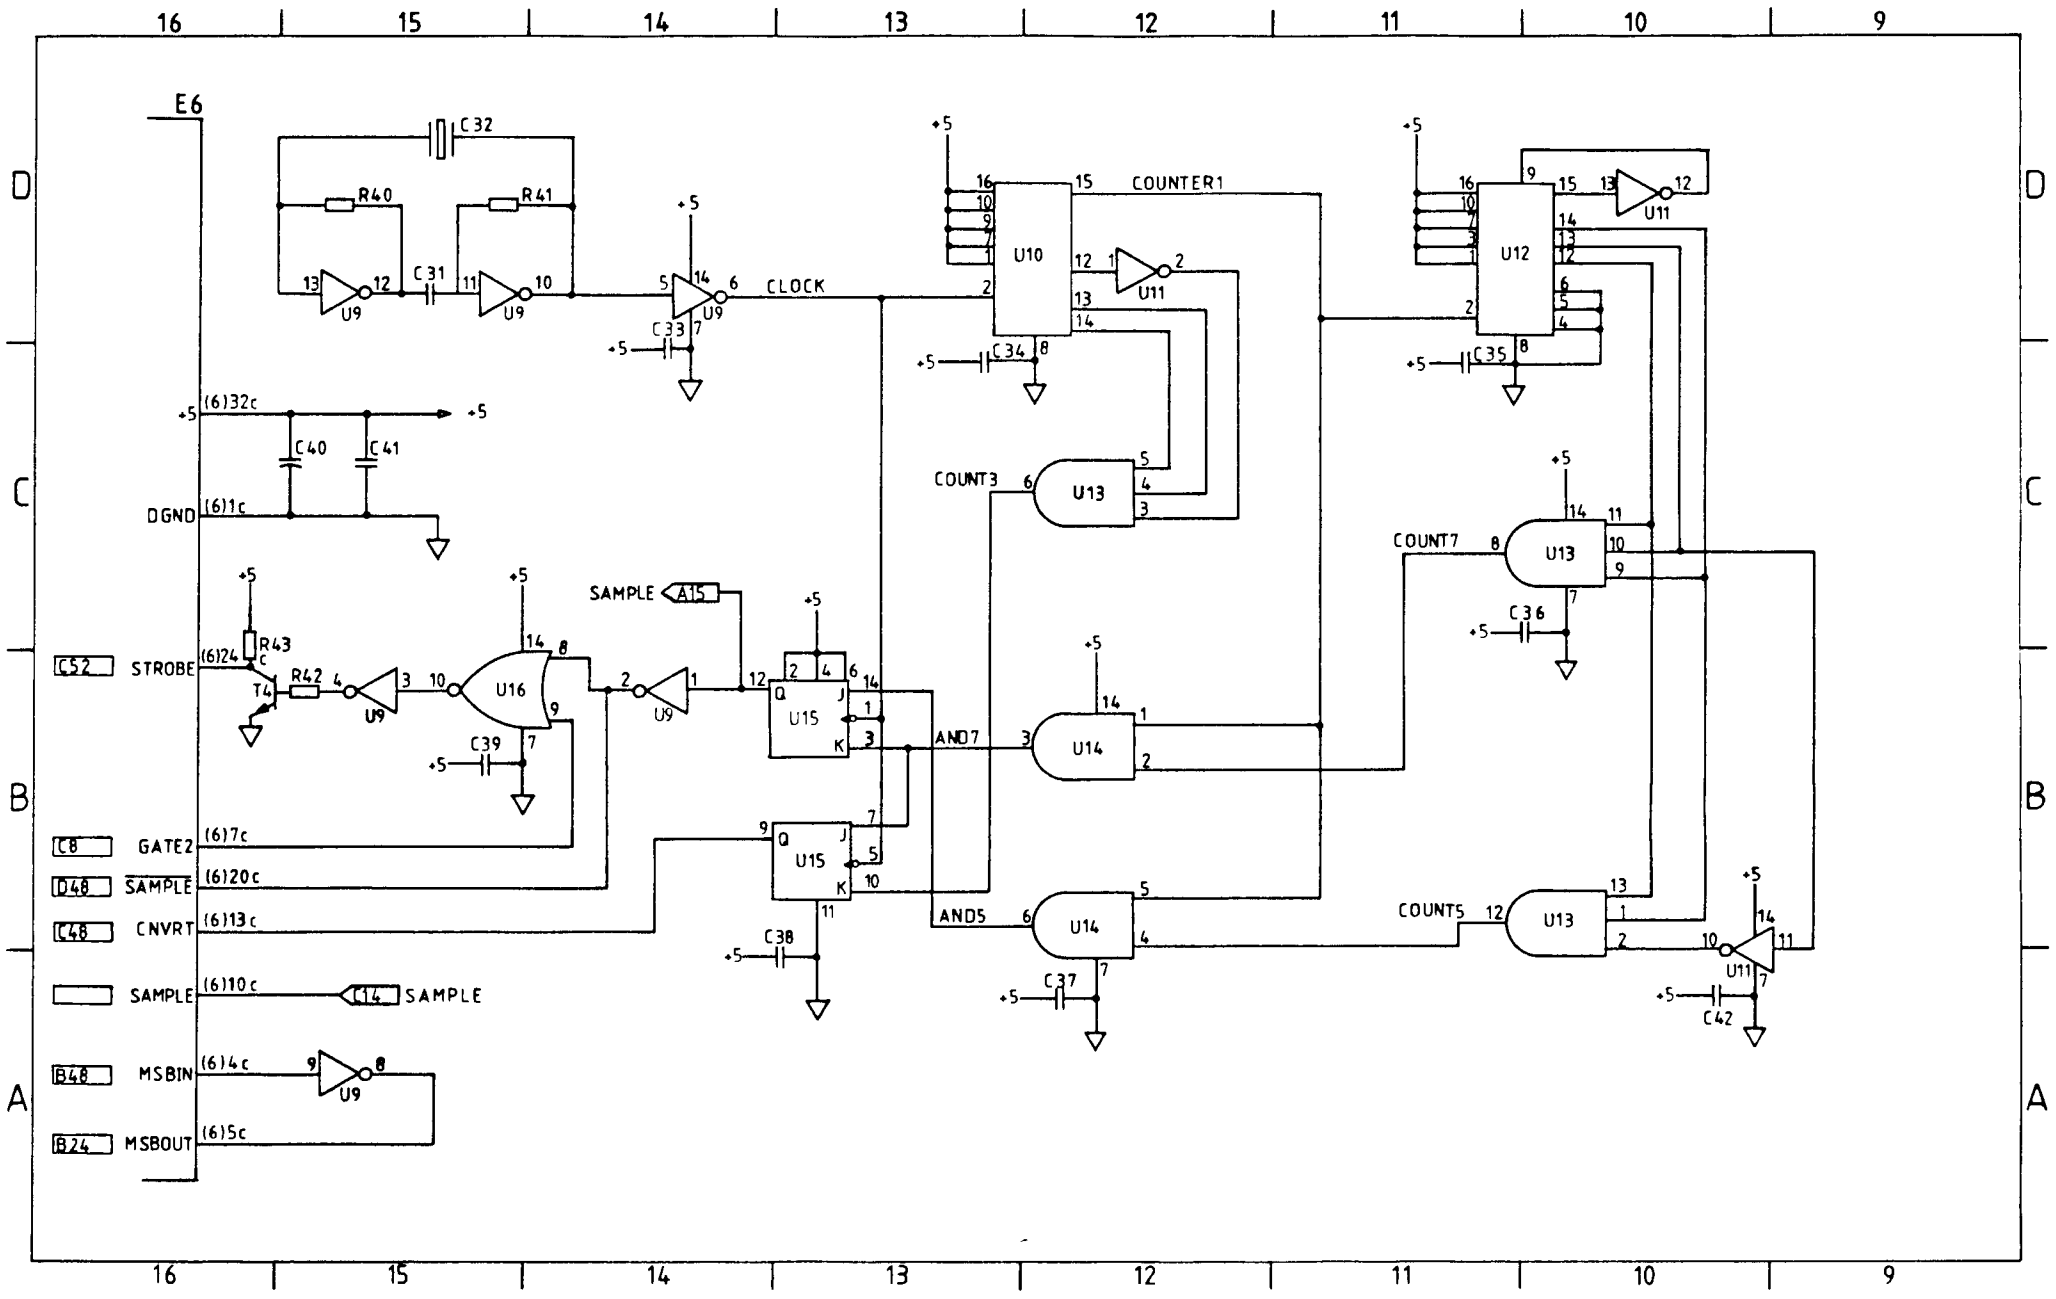
16
15
14
13
12
11
10
9
D
C
B
A
E6
C32
R40
R41
C31
U9
CLOCK
C33
+5
U10
COUNTER1
C34
U11
U12
C35
U11
(6)32c
C40
C41
DGND
(6)1c
COUNT3
U13
COUNT7
U13
C36
SAMPLE
A15
R43
(6)24
C52
STROBE
T4
R42
U16
C39
U15
AND7
U14
C8
GATE2
(6)7c
D48
SAMPLE
(6)20c
C48
CNVRT
(6)13c
U15
AND5
U14
C38
C37
COUNT5
U13
U11
C42
SAMPLE
(6)10c
C14
SAMPLE
B48
MSBIN
(6)4c
U9
B24
MSBOUT
(6)5c

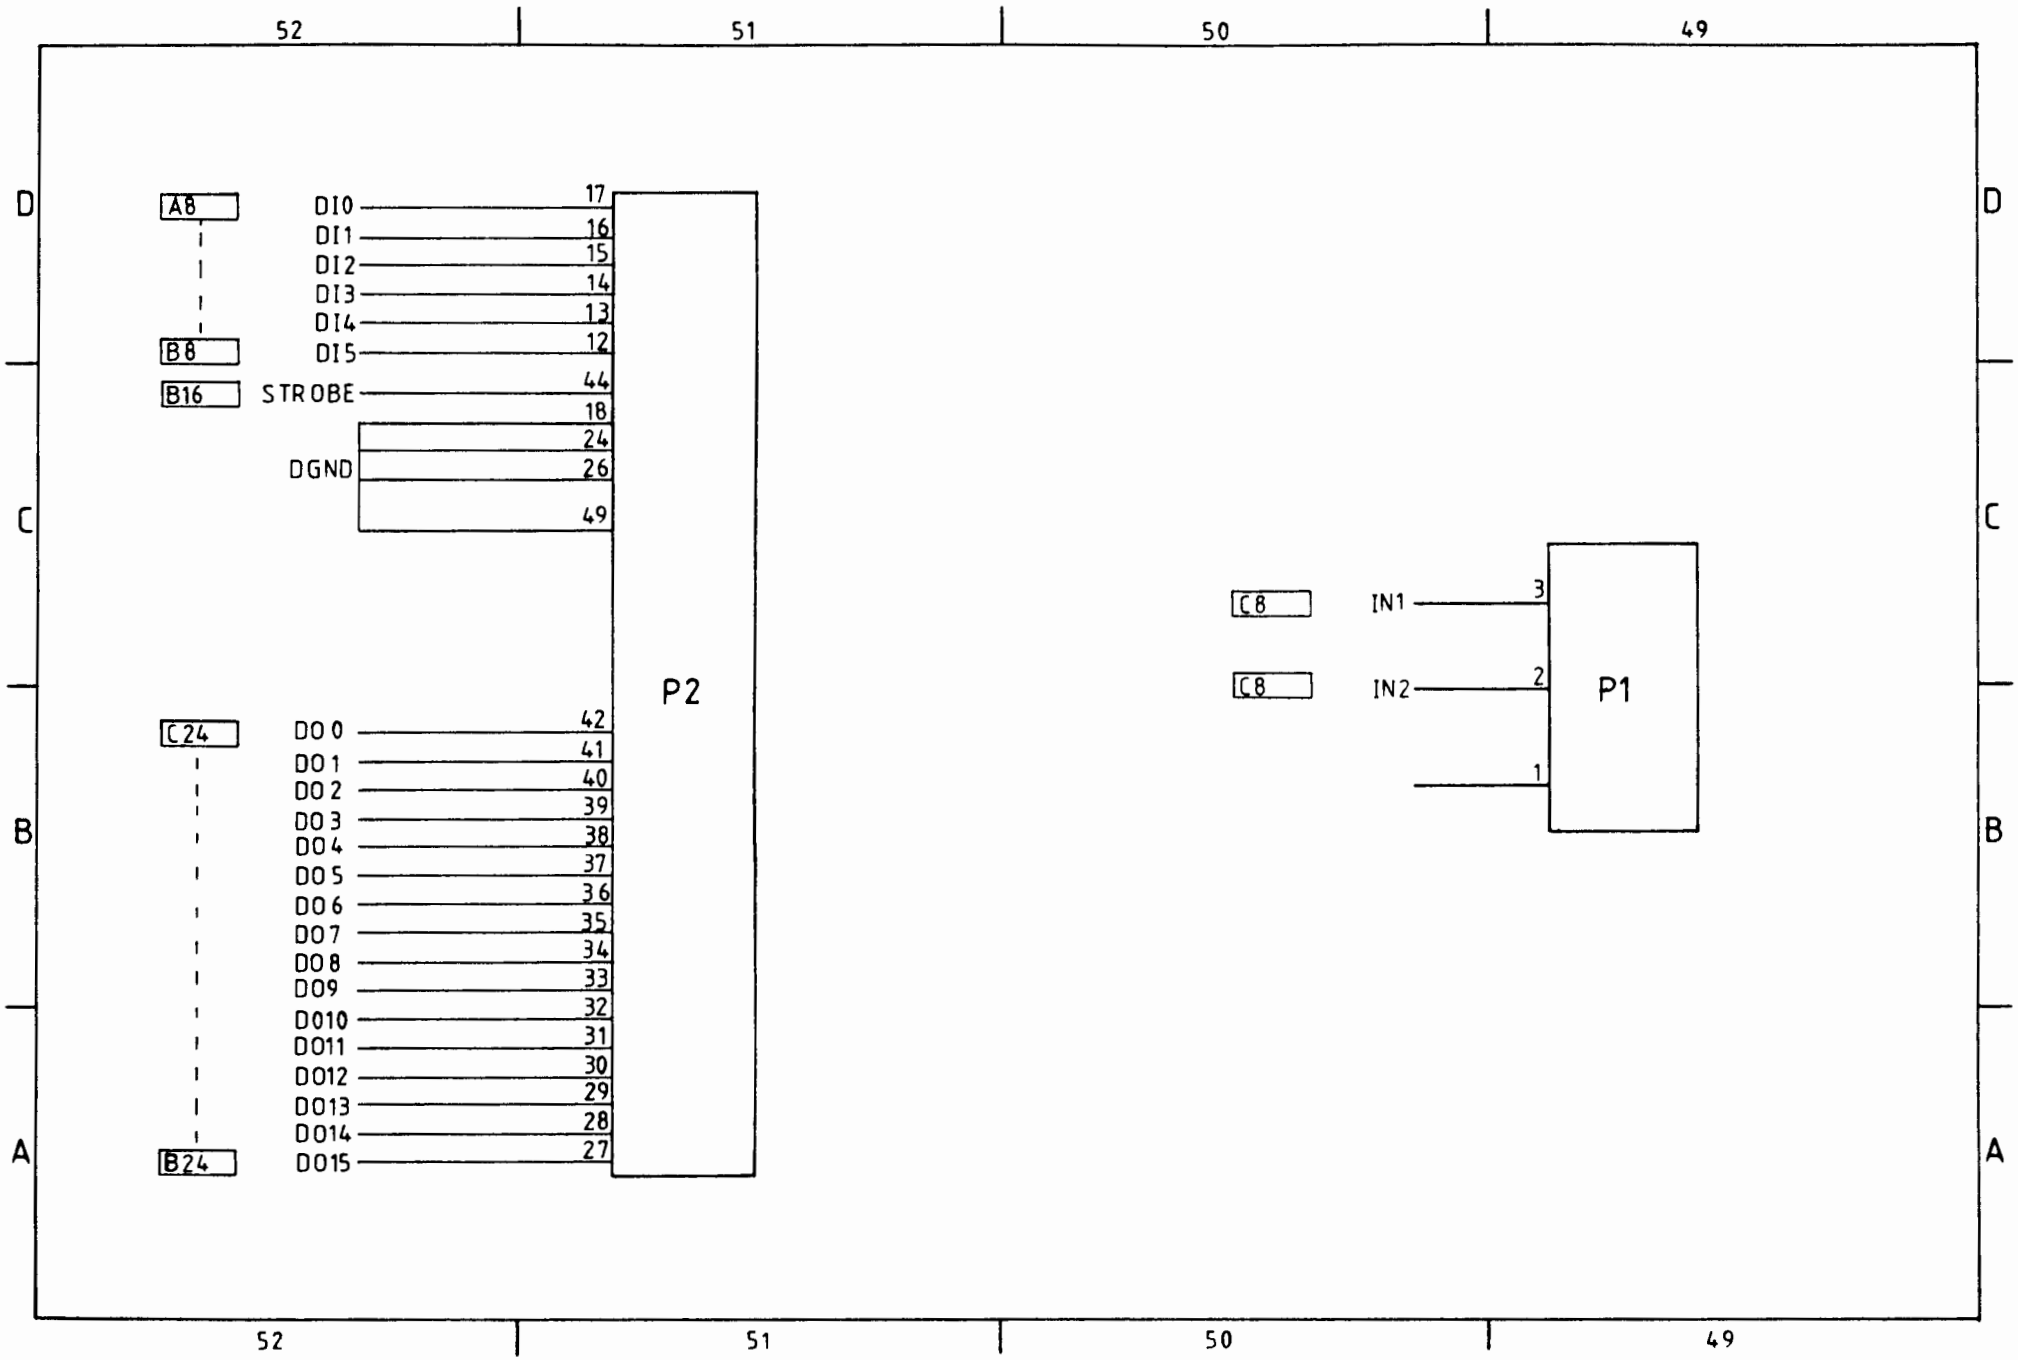
52
51
50
49
D
C
B
A
A8
B8
B16
DI0
DI1
DI2
DI3
DI4
DI5
STROBE
DGND
17
16
15
14
13
12
44
18
24
26
49
C24
B24
DO0
DO1
DO2
DO3
DO4
DO5
DO6
DO7
DO8
DO9
DO10
DO11
DO12
DO13
DO14
DO15
42
41
40
39
38
37
36
35
34
33
32
31
30
29
28
27
P2
C8
C8
IN1
IN2
3
2
1
P1

RECEIVER CIRCUIT DIAGRAM COMPONENTS

Reference Designation	Value	Description	Mfr Part Number
C1	10uF 35VDC	Tantalum	
C2	10uF 35VDC	Tantalum	
C3	63nF	Ceramic	
C4	63nF	Ceramic	
C5	1nF	Polystyrene	
C6	1nF	Polystyrene	
C7	1nF	Polystyrene	
C8	.1uF	Ceramic	
C9	40nF	Polystyrene	
C10	5.6nF	Encapsulated Polystyrene	
C11	40nF	Polystyrene	
C12	.1uF	Polyester	
C13	.68uF	Polyester	
C14	.068uF	Polyester	
C15	.068uF	Polyester	
C16	.01uF	Polyester	
C17	.82uF	Polyester	
C18	2.2uF	Polyester	
C19	1nF	Polyester	
C20	4.7uF 25VDC	Tantalum	
C21	4.7uF 25VDC	Tantalum	
C22	3n3	Ceramic	
C23	3n3	Ceramic	
C24	270pF	Ceramic	
C25	4.7uF 25VDC	Tantalum	
C26	4.7uF 25VDC	Tantalum	
C27	3n3	Ceramic	
C28	3n3	Ceramic	
C29	82pF	Polystyrene	
C31	1nF	Ceramic	
C32	4.00Mhz	CRYSTAL	
C33	1nF	Ceramic	
C34	1nF	Ceramic	
C35	1nF	Ceramic	
C36	1nF	Ceramic	

C37	1nF	Ceramic
C38	1nF	Ceramic
C39	1nF	Ceramic
C40	10uF 63VDC	Elec.
C41	33nF	Ceramic
C42	1nF	Ceramic
C43	22uF 25VDC	Elec.
C44	33nF	Ceramic
C45	4n7	Ceramic
C46	4n7	Ceramic
C47	1.5nF	Polystyrene
C48	1.5nF	Polystyrene
C49	.68uF	Polyester
C50	1.5nF	Polystyrene
C51	1.5nF	Polystyrene
C52	10uF	Elec.
C53	10uF	Elec.
C54	10nF	Ceramic
C55	10nF	Ceramic
C56	10nF	Ceramic
C57	10nF	Ceramic
C58	10nF	Ceramic
C59	10nF	Ceramic
C60	10nF	Ceramic
C61	10nF	Ceramic
C62	.068uF	Polyester
C63	7110pF 1%	Encapsulated polyester (Combination)
C64	1356pF 1%	Encapsulated polyester (Combination)
C65	13015pF 1%	Encapsulated polyester (Combination)
C66	4300pF 1%	Encapsulated polyester (Combination)
C67	5080pF 1%	Encapsulated polyester (Combination)
C68	7084pF 1%	Encapsulated polyester (Combination)
C69	1365pF 1%	Encapsulated polyester (Combination)

C70	13015pF 1% (Combination)	Encapsulated polyester
C71	4220pF 1% (Combination)	Encapsulated polyester
C72	5040pF 1% (Combination)	Encapsulated polyester
C73	1.5nF	Ceramic
C74	1.5nF	Ceramic
C75	1uF	Tantalum
C76	1uF	Tantalum
C77	22nF	Ceramic
C78	22nF	Ceramic
C79	3n3	Ceramic
C80	3n3	Ceramic
C81	3n3	Ceramic
C82	3n3	Ceramic
C83	3n3	Ceramic
C84	3n3	Ceramic
C85	1000pF	Polystyrene
C86	10uF 63VDC	Elec.
C87	33nF	Ceramic
C88	270pF	Ceramic
C89	.01uF	Polyester
C90	3n3	Ceramic
C91	3n3	Ceramic
C92	10uF	Elec.
C93	10uF	Elec.
C94	33nF	Ceramic
C95	33nF	Ceramic

D2	signal diode 1N4148
D2	signal diode 1N4148
D3	signal diode 1N4148
D4	signal diode 1N4148

E1	Eurocard Plug Connector
E2	Eurocard Plug Connector
E3	Eurocard Plug Connector

E4		Eurocard Plug Connector	
E5		Eurocard Plug Connector	
E6		Eurocard Plug Connector	
L1	2.2nH		
L2	12.26nH 5%	variable	core: 3H1, A315
L3	8.97nH 5%	variable	core: 3H1, A315
L4	12.26nH 5%	variable	core: 3H1, A315
L5	8.97nH 5%	variable	core: 3H1, A315
P1		XLR 3-POLE Socket	
P2		50 PIN D-TYPE Socket	
R1	3.32K 1% 1/4 W.		
R2	3.32K 1% 1/4 W.		
R3	3.32K 1% 1/4 W.		
R4	3.32K 1% 1/4 W.		
R5	100K 5% 1/4 W.		
R6	2.2K 5% 1/4 W.		
R7	2.2K 5% 1/4 W.		
R8	470K 5% 1/4 W.		
R9	5.6K 5% 1/4 W.		
R10	5.6K 5% 1/4 W.		
R11	1.0M 5% 1/4 W.		
R12	1.78K 1% 1/4 W.		
R13	10.0K 1% 1/4 W.		
R14	10K Wire pot 20 turn paralled with 4.64K 1% 1/4 W		
R15	560K 5% 1/4 W.		
R16	18K 5% 1/4 W.		
R17	2.2K 5% 1/4 W.		
R18	27K 5% 1/4 W.		
R19	560K 5% 1/4 W.		
R20	15K 5% 1/4 W.		
R21	56K 5% 1/4 W.		
R22	4.7K 5% 1/4 W.		

R23	26.483K 1% 1/4 W.
	(26.1K + 383)
R24	18.762K 1% 1/4 W.
	(18.2K + 562)
R25	13.274K 1% 1/4 W.
	(13.0K + 274)
R26	9.399K 1% 1/4 W.
	(9.09K + 309)
R27	4.710 1% 1/4 W.
	(4.12K + 590)
R28	2.365K 1% 1/4 W.
	(2.00K + 365)
R29	18K 5% 1/4 W.
R30	18K 5% 1/4 W.
R31	18K 5% 1/4 W.
R32	18K 5% 1/4 W.
R33	18K 5% 1/4 W.
R34	18K 5% 1/4 W.
R35	18.2K 1% 1/4 W.
R36	18.2K 1% 1/4 W.
R37	10.0K 1% 1/4 W.
R38	10.0K 1% 1/4 W.
R39	5.62K 1% 1/4 W.
R40	1K 5% 1/4 W.
R41	820 5% 1/4 W.
R42	27K 5% 1/4 W.
R43	5.6K 5% 1/4 W.
R44	33.2K 1% 1/4 W.
R45	10.0K 1% 1/4 W.
R46	2.74K 1% 1/4 W.
R47	6.81K 1% 1/4 W.
R48	10.0K 1% 1/4 W.
R49	10.0K 1% 1/4 W.
R50	5.11K 1% 1/4 W.
R51	10.0K 1% 1/4 W.
R52	33.2K 1% 1/4 W.
R53	10.0K 1% 1/4 W.
R54	10.0K 1% 1/4 W.
R55	10.0K 1% 1/4 W.
R56	15.0K 1% 1/4 W.

R57	10.0K 1% 1/4 W.	
R58	10.0K 1% 1/4 W.	
R59	5.36K 1% 1/4 W.	
R60	10.0K 1% 1/4 W.	
R61	10.0K 1% 1/4 W.	
R62	6.81K 1% 1/4 W.	
R63	10.0K 1% 1/4 W.	
R64	33.2K 1% 1/4 W.	
R65	10.0K 1% 1/4 W.	
R66	10.0K 1% 1/4 W.	
R67	10.0K 1% 1/4 W.	
R68	1.00K 1% 1/4 W.	
R69	1.00K 1% 1/4 W.	
R70	2.00K 1% 1/4 W.	
R71	.655K 1% 1/4 W.	
R72	10.0K 1% 1/4 W.	
R73	1.5K 5% 1/4 W.	
R74	10.0K 1% 1/4 W.	
R75	100 5% 1/4 W.	
R76	100 5% 1/4 W.	
R77	5K Multi-turn Trimpot	TYPE-960-20
R78	100K Multi-turn Trimpot	3006p-1-104-100K
R79	1.8M 1% 1/4 W.	
R80	33 5% 1/4 W.	
R81	2K7 5% 1/4 W.a	
R82	12K 5% 1/4 W.	

T1	transistor	BC108
T2	transistor	BC108
T3	transistor	2N2222A
T4	transistor	2N2222A
T5	transistor	BC109C
T6	transistor	BC109C

U1	FET Op-amp	LF351
U2	FET Op-amp	LF351
U3	Tone Decoder	XR2211CP
U4	timer	555
U5	switch array	DG308CJ

U6	switch array	DG308CJ (Siliconix)
U7	FET Op-amp	LF351
U8	FET Op-amp	LF351
U9	inverters	74S04
U10	Sync. 4-bit counter	74LS160
U11	inverters	74LS04
U12	Sync. 4-bit counter	74LS160
U13	Triple 3-input AND gates	74LS11
U14	Quad 2-input AND gates	74LS08
U15	Dual J-K Flip-Flops	74H73
U16	Quad 2-input NOR gates	74LS02
U17	Octal Line Drivers	74LS244
U18	Octal Line Drivers	74LS244
U19	FET Op-amp	LF351
U20	FET Op-amp	LF351
U21	FET Op-amp	LF351
U22	FET Op-amp	LF351
U23	FET Op-amp	LF351
U24	FET Op-amp	LF351
U25	FET Op-amp	LF351
U26	FET Op-amp	LF351
U27	FET Op-amp	LF351
U28	FET Op-amp	LF356
U29	switch array	DG308CJ (Siliconix)
U30	FET Op-amp	LF356
U31	ADC	PCM75JG (Burr Brown)
U32	Positive edge-triggered latches	74LS374

U33		Positive edge-triggered latches	74LS374
Z1	5V1	Zener	
Z2	18V	Zener	
Z3	18V	Zener	
Z4	5V1	Zener	
+5	250mA	.5mVRMS Ripple	PM 529
+15,-15	100mA	.5mVRMS Ripple	PM 502

APPENDIX E

SYNC DETECTION CIRCUITRY DESIGN

INTRODUCTION

Detection of the synchronising burst from the transmitter unit is done using EXAR's FSK Demodulation/Tone Decoder chip (XR-2211). The chip has been set to work in the tone decoder mode and consists of a phase-locked loop and associated circuitry. By selecting external components the VCO frequency, loop damping factor, detection bandwidth and out-of-band signal rejection are set.

DESIGN

The circuit diagram for the completed design is shown below.

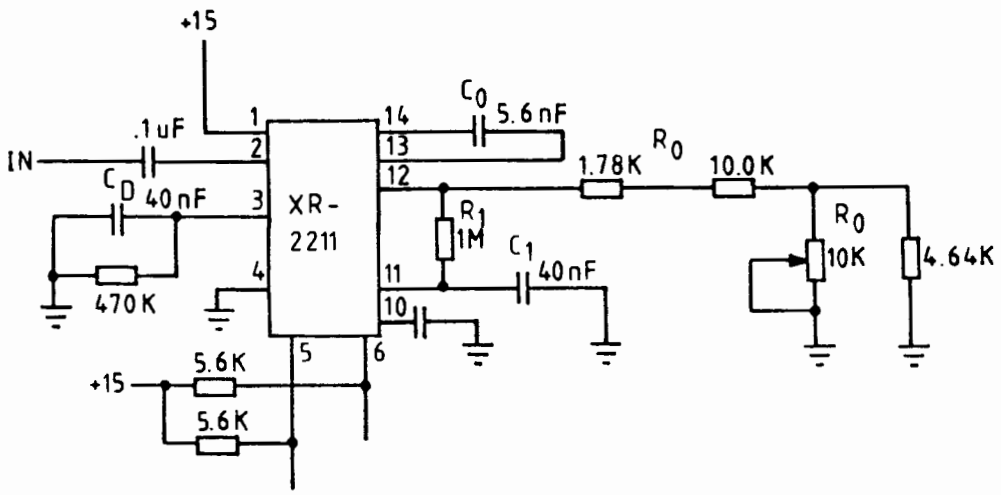

FIGURE E1 Tone decoder configuration

1) Timing Resistor (R_0)

This value is initially chosen and the rest of the design then done around it. Due to availability 11.78K was put in series with a 10K multi-turn wire trimpot paralled with 4.64K (all 1% metal film for temperature stability). With the trimpot centered this gives a design value for R_0 of 13.37K.

2) Timing Capacitor (C_0)

This value is calculated using

$$C_0 = 1/R_0 f_s$$

where f_s is the tone frequency and hence the centre frequency of the VCO (12.75 Khz). Therefore

$$C_0 = 1/(13.37E3)(12.75E3)$$
$$C_0 = 5.9 \text{ nF}.$$

Any value of C_0 in the range 5.4 nF to 6.6 nF will do. Due to the temperature stability requirements of the VCO a high quality, low temperature coefficient capacitor should be used. In the prototype an encapsulated polystyrene was used.

Tests done in an environmental chamber showed a variation of 77 Hz. over the range 0°C to 50°C for the complete decoder circuit section (fig. E2).

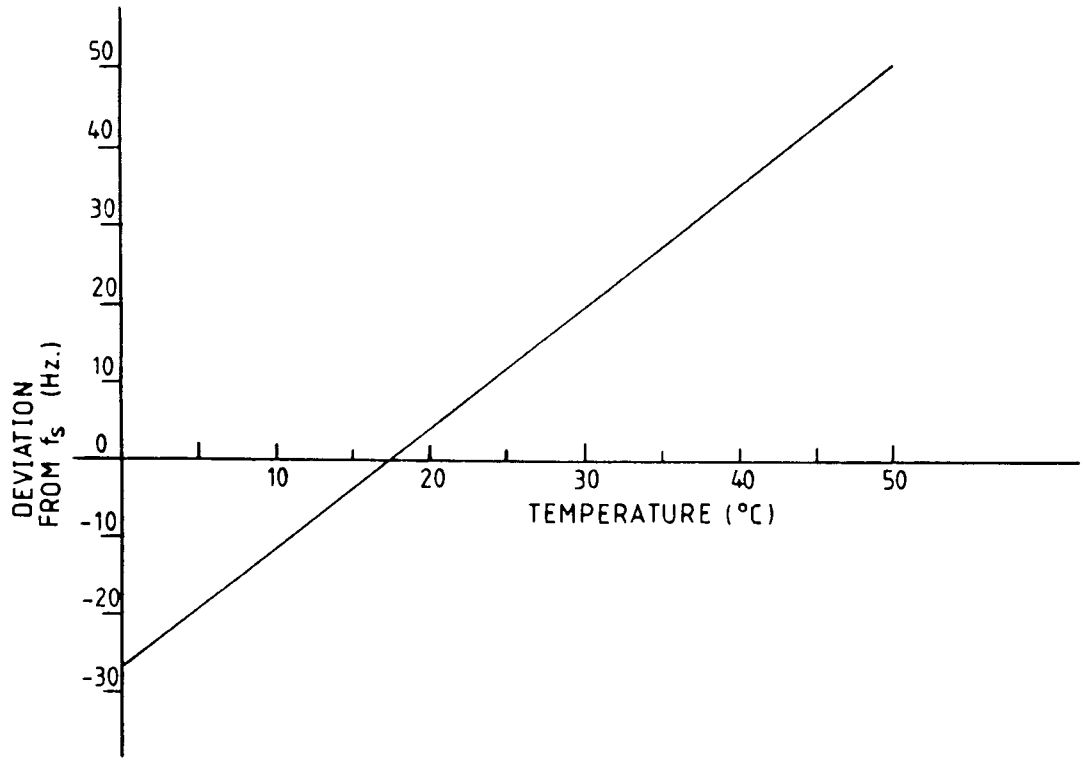

FIGURE E2 VCO center frequency vs. temperature

3) Detection Bandwidth Resistor (R_1)

The bandwidth ± δf is set at δf = 170 Hz. This value accomodates the slight drift in VCO frequency due to temperature drift. Hence

$$R_1 = R_0(f_0 / \delta f)$$
$$= (13.37E3)(12.75E3 / 170)$$
$$R_1 = 1\ M\Omega$$

4) Loop Damping Factor Capacitor (C_1)

Although the manufacturers recommend a damping factor of ½ for most applications, damping of 1/10 was used here as this showed best performance. However, this value was not very critical.

$$C_1 = C_0/16\zeta^2 \qquad \text{where: } \zeta = \text{damping}$$

For a damping factor of ½

$$C_1 = .25C_0$$
$$= 1.5\ nF$$

The design here uses C_1 = 40 nF.
Note that increasing C_1 improves the out-of-band signal rejection, but increases the PLL capture time.

5) Lock Detect Filter Capacitor (C_D)

This avoids chatter at the logic outputs due to program material near the capture bandwidth. The resistor R_D is set at 470K and so C_D is given by

$$C_D \geq (16/\text{capture range in Hz})$$

Hence

$$C_D = 16/323$$
$$C_D = .05\ uF$$

is used here.

For all the components above, capacitors are polystyrene and resistors carbon, unless otherwise indicated.

CHECKING VCO FREQUENCY

The free-running frequency of the VCO can be measured by doing the following:

1) Remove any input to pin 2 and short pins 2 and 10 together.
2) Disconnect capacitor C_D at pin 3 and measure the frequency at this pin.

This procedure is also used to set the VCO frequency.

APPENDIX F

SAMPLE-AND-HOLD CHARGING RESISTOR

An interesting effect was noticed with different values of charging resistance used for the discrete sample/hold built.

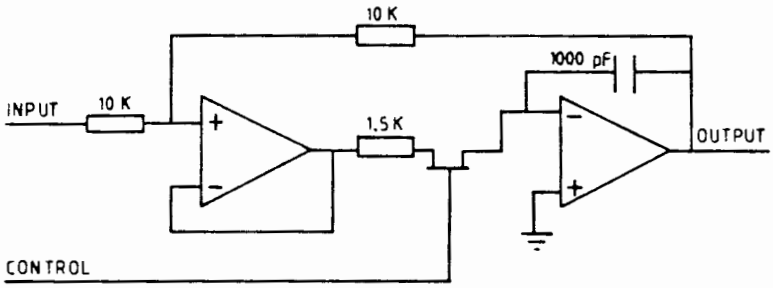

FIGURE F1 Discrete sample-and-hold

With no resistance the second harmonic had an amplitude of approximately -67 dB.

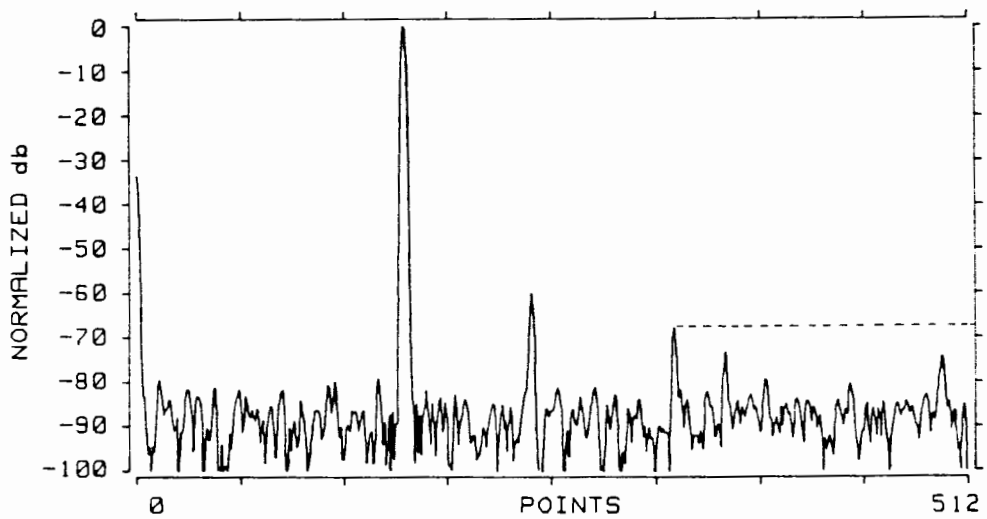

FIGURE F2 Sample/hold with zero series resistance

By increasing this to 1K and above the harmonic disappeared into the noise although an increase in the noise floor was observed for resistors significantly above 1K.

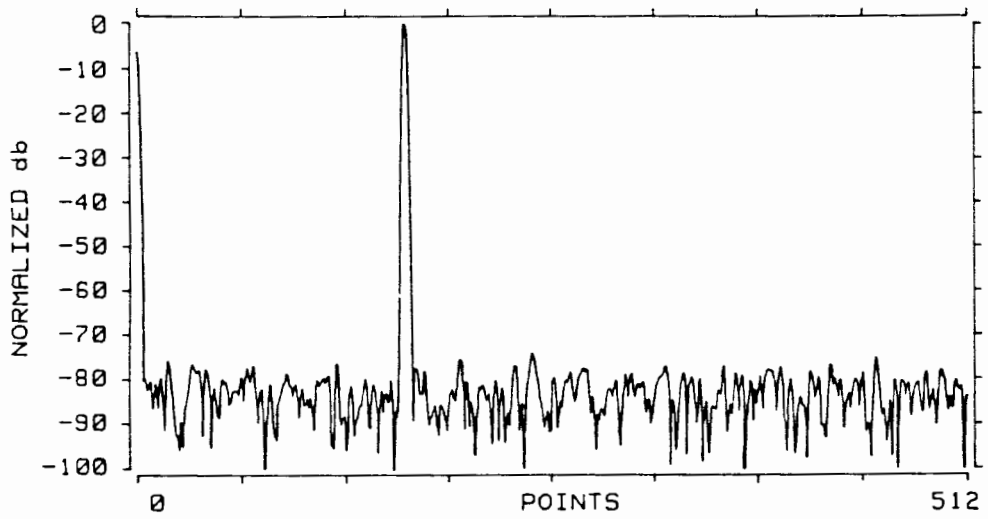

FIGURE F3 Sample/hold with 1.5K series resistance

In the case of a sample/hold at the output of a DAC an improvement in distortion performance is obtained by not allowing the op-amp driving the capacitor to slew. By adding sufficient series resistance so that the capacitor charges exponentially the non-linearity of the slewing op-amp is avoided and replaced with a linear effect [49].

The distortion is seen at the output of a digital-to-analog conversion system due to the sample/hold generating the output waveform and any non-linear behaviour during the device's operating cycle forms part of the waveform. However, in the case of an ADC system the situation is different. The ADC only starts a new conversion once the output of the sample/hold has settled in the 'hold' mode. During 'sample' its output is not considered and any distortion produced during this interval should be irrelevant. Although the reason for the variation in the observed distortion of the SPCATS data acquisition system, due to different resistance values, is not clearly understood it is thought to be due to the slew-rate recovery time of the op-amp driving the capacitor. Such an effect could be present during the initial stages of the 'hold' mode. The fact that the Sony Corporation uses a 1K resistor in the sample/hold of its CX20018 ADC suggests that the improvement in the distortion performance due to this is well founded [50].

22 JUN 1987