AD-A183 904 A COMPACT DIGITAL COMMUNICATIONS SYSTEM PART 3 1/1
GENERALISING THE PHASE LOC..(U) ROYAL SIGNALS AND RADAR
ESTABLISHMENT MALVERN (ENGLAND) J A PRITCHARD APR 87
UNCLASSIFIED RSRE-MEMO-4027 DRIC-BR-102288 F/G 25/5 NL

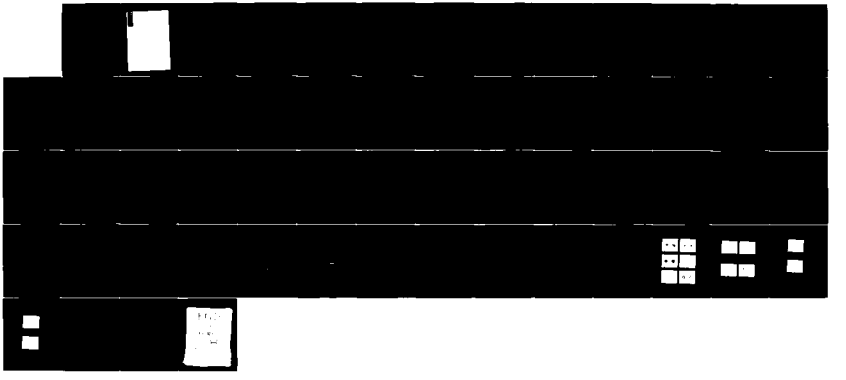

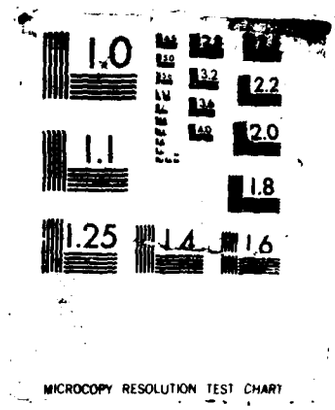
1.0
1.1
1.25
1.4
1.6
2.2
2.0
1.8
MICROCOPY RESOLUTION TEST CHART

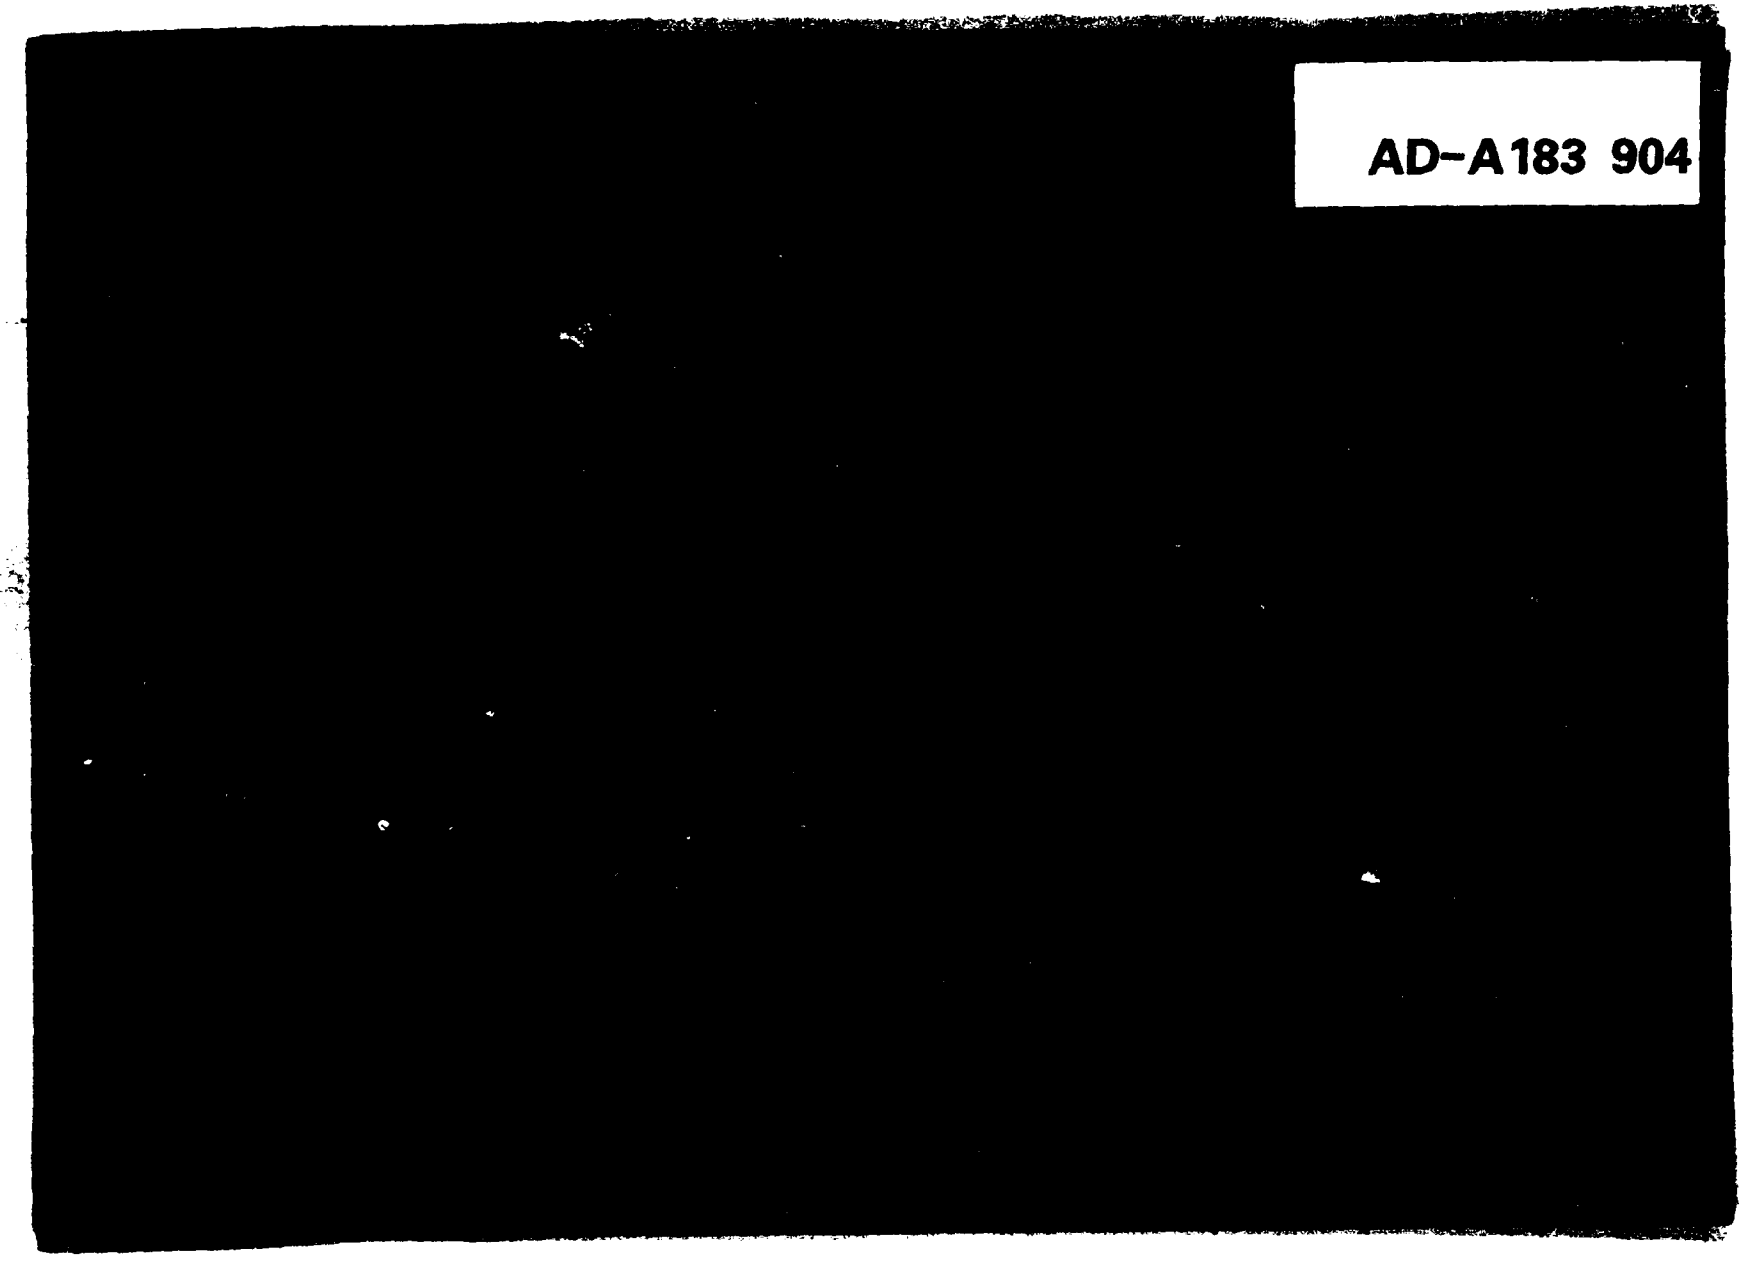
AD-A183 904

ROYAL SIGNALS AND RADAR ESTABLISHMENT

Memorandum 4027

TITLE: A COMPACT DIGITAL COMMUNICATIONS SYSTEM
PART 3: GENERALISING THE PHASE LOCKED LOOP

AUTHOR: J A S PRITCHARD

DATE: APRIL 1987

SUMMARY

We describe a generalisation of the phase locked loop principle which eliminates the problem of carrier recovery when applied to phase shift signals. The system is tolerant of signal fading and shows orders of magnitude improvements in pull in range and tracking performance over current systems used in satcoms. Pull in range exceeds 150 Hz in an absolute signal to noise ratio of worse than -10 dB. The loop will track a 50 Baud PSK modulated carrier shifting at more than 300 Hz sec^{-1}. The flexibility of this example of a new approach to communications engineering has yet to be explored.

A COMPACT DIGITAL COMMUNICATIONS SYSTEM
PART 3: GENERALISING THE PHASE LOCKED LOOP

CONTENTS

LIST OF FIGURES

LIST OF ABBREVIATIONS

PLL	Phase Locked Loop
VCO	Voltage Controlled Oscillator
SPRF	Single Pole Recursive Filter
FSK	Frequency Shift Keying
PSK	Phase Shift Keying
QPSK	Quadrature Phase Shift Keying
LCR circuit	Fundamental series resonant circuit with a capacitor, a coil and a resistor in series
PCM	Pulse code modulation
RAM	Random Access Memory
PROM	Programmable Read Only Memory
FSM	Finite State Machine
TTL	Transistor Transistor Logic
RTS	Random Telegraph Signal (constant chip rate)
DPSK	Differential Phase Shift Keying
PAL	Programmable Array Logic
S/N	Signal to Noise Ratio
Fo	Oscillator Centre Frequency
Fc	Carrier Centre Frequency

1 INTRODUCTION

1.1 Background

The purpose of a phase locked loop is to supply information about a signal which has become distorted. It does this by using prior knowledge which is stored in its dynamics. It infers the "exact" frequency or phase of the incoming signal by comparing the state of its own internal model of the signal with the distorted input signal over a period of time.

The phase locked loop (PLL) contains a model of an incoming signal, usually assumed to be a sine wave, in its voltage controlled oscillator (VCO). It infers the frequency or phase of the sine wave which it is "tracking" by continuously sampling the distorted signal presented to it. The phase locked loop also has a limited memory, so that slow variations in phase or changes in frequency (ie variations inside its "bandwidth") may be accommodated. These allowable variations define the ambiguity in the incoming signal and should exactly match our ignorance of its parameters. These systems are especially useful for tracking and demodulating noisy signals which are assumed to be quasi-sinusoidal, and where the ambiguities (data) are phase like and are on a timescale which is much longer than that of the correlation within the distortion.

Where the input signal is no longer sinusoidal, eg where it contains phase discontinuities, the phase locked loop no longer contains an accurate model of the incoming signal. Although various schemes may be used to enable a phase locked loop to "lock" to the centre frequency of a carrier, most rely on the recovery of a pure sine wave from the signal.

In this memorandum we seek to eliminate the shortcomings of the conventional phase locked loop by including in it the dynamics of the transmitted signal. A circuit is described which applies this principle to PSK modulated signals and its performance is measured. We note how the principle may be extended to different modulation schemes.

Demodulation of signals, in particular PSK signals, and the elimination of problems associated with timing recovery will be dealt with in part 4. More general schemes will be addressed in part 5.

1.2 Conventional Phase Locked Loops and their Deficiencies

Conventional phase locked loops are well documented (see for example ref 1) and the structure of a basic phase locked loop is shown in fig 1. There are three basic elements in the PLL; the voltage controlled oscillator (VCO), the phase detector, and the loop filter. For the purposes of this memorandum, the phase detector may be considered to be a multiplier and the loop filter an integrator.

The VCO outputs a model of the input signal which is assumed to be a sine wave. This is continuously multiplied with the incoming signal, and the output of the multiplier is integrated using a time constant which is long compared to the period of its inputs. The output of the integrator is an analogue voltage which is a monotonic function of the phase difference between the input and the model. Ambiguities regarding the phase or frequency of the incoming signal are expressed as points in the state space of the PLL. The system is arranged so that the input signal state becomes an attractor in that state space. Modulation schemes may be considered as a further ambiguity in the input signal, and where that modulation involves relatively slow phase or frequency excursions, may be recovered by examining the deatiled behaviour of the error voltage from the integrator. The dynamics of the conventional PLL may be tailored to suit particular requirements, principally by altering the loop filter. For example two integrators may be cascaded to store "acceleration" information; this is known as a "third order loop".

The successful operation of a phase locked loop is governed by various "trade-offs", for example that between noise and such parameters as tracking rate and pull in performance. The discussion of these matters is outside the scope of this memorandum and not strictly relevant to the discussion based on physical principles which we wish to pursue.

In this discussion, we are concerned with the failure of the conventional phase lock loop when the input signal contains components which are on a time-scale which is much less than the inverse loop bandwidth, ie those which are within the decorrelation time of the loop filter. Such signals may be for example frequency shift keyed (FSK), phase shift keyed (PSK), or wideband signals. We are not yet, however, interested in completely general classes of signal which may involve the collection of statistics using classifiers (see ref 2).

The particular example which will be used is phase shift keying, or PSK. Wider applications of the techniques outlined in this memorandum should become apparent as the discussion progresses.

A conventional phase locked loop cannot in general lock to a carrier which has "hard" or sudden phase reversals. A sudden phase reversal produces a discontinuity in the integrator input voltage, and there is a fundmental ambiguity in the correction which must be made to realign the phase of the VCO. The usual solution is to band limit the signal with a filter and then to square it by "mixing" or multiplying it by itself to recover the quasi-sinusoidal carrier (fig 2). This is unsatisfactory for a number of reasons, the two most important of which are detailed as follows:

(1) The signal to noise ratio is degraded due to the inclusion of signal X noise and noise X noise terms.

(2) Phase ambiguity remains. To demodulate the signal it is necessary to divide the PLL frequency by two and compare it with the input signal using an integrate and dump circuit (see fig 2). The phase of the PLL output may be considered to lie on a Riemann surface (ie completing a rotation in the Argand plane every 720 degrees) resulting in an ambiguity of π after division by two, ie after taking the square root.

1.3 Proposed Solution

We wish to use physical insight to eliminate some of the undesirable properties of conventional communications circuits. This is achieved by factorising the spectrum of a modulated carrier into its fundamental components and examining those components on a time-scale which is commensurate with each time constant in the signal. The three fundamental time constants present in the signal are:

(S1) a sinusoidal carrier with period τ_s

(S2) a modulation "chip" period τ_c

(S3) a carrier drift time τ_d

This breakdown of signal structure was discussed in part 1, which dealt with the identification of S1. This memorandum will now deal with S3, the longest time constant, as its determination is a prerequisite for using S2. We will then, in a later publication, attempt the complete reconstruction of the signal.

The locking to and tracking of the signal will be done by a circuit which we call a generalised phase lock loop. There are certain requirements of such a circuit, these are:

(1) The feedback must be invariant under modulation.

(2) There must be no fundamental signal to noise penalty in implementation.

(3) For synchronous communications, the phase of the carrier should be obtainable.

In implementing this circuit we would like to eliminate frequency drift without building in any unnecessary structure, such as that necessary to recover the carrier, which is not actually present in the original signal. In PSK the pure, unmodulated carrier is not present in the received signal.

To achieve these objectives we will replace the VCO in the conventional phase locked loop with a more accurate representation of the signal. We will sample the signal in the frequency domain not with conventional filters, whose top-hat like response has heretofore been arbitrarily imposed on systems, but with the more natural responses of the LCR finite-element building block. We will then exploit one of the most powerful characteristics of a PSK spectrum, its mirror symmetry around the carrier frequency, to determine frequency offset by what is known as "symmetry breaking". In doing so, we will retain desired information such as the modulation symbols and the carrier phase without reconstructing the carrier, and will remove the out of band noise without using an explicit filter.

These circuits are based on physical principles and despite an inherent simplicity are sufficiently different from conventional circuits to make understanding difficult. We therefore review the basic principles and then introduce new concepts step by step, with the minimum of mathematics, to facilitate understanding.

2 RATIONALE

2.1 Introduction

In this section we will introduce the principles that are to be used in the synthesis of a generalised phase locking device. We start by reviewing the basic tool which will be used to sample the frequency domain in the vicinity of the signal; this is nothing more than a "digital analogue" of an ordinary LCR circuit as was used in part 1. We are going to explore the power of digital electronics to create an abstract entity where the parameters may be altered at will within a physical model. We will see how this can be done without the usual inconveniences such as energy dissipation, the movement of parallel plates or ferrites, crosstalk, thermal drift, noise etc.

Having given a general overview of the synthetic oscillator, we will describe in 2.3 how the circuit can be used to sample the frequency domain in an interval and return a complex number after each inverse-interval time period. This summarises the activity of that part of the spectrum during that period. We use this technique to create a phase lock loop which is substantially the same in concept as the conventional type

In section 2.4 we investigate the power spectrum of various signals and make several observations. We note the shortcomings of the phase locking scheme described in section 2.3

We then show how it is possible to sample the frequency domain more densely, and how to use the symmetry of PSK signals to achieve phase lock without recovering the carrier.

2.2 Review of the Synthetic Digital Oscillator

Figure 3 shows the layout of a "synthetic digital oscillator". Fundamentally, it works by beating the incoming signal down to zero frequency.

This circuit emulates an LCR circuit to an arbitrary degree of accuracy using the minimum number of components. It has some important advantages over conventional "analogue" or "real" LCR oscillators. For example, the Q factor and centre frequency may be soft or variable, and the state at any time is available from its registers in the form of a complex number describing the real and imaginary (cosinusoidal and sinusoidal) amplitudes. The "energy" content may be altered at will without dissipation, and there is no restriction to an exponential decay; a large range of possible dynamics may be written into the circuit through the decrement tables or by using different recursive filter structures. The circuit may be made non-physical by removing the necessity of using sine and cosine tables, in principal one may choose any functions (for example Walsh functions).

The circuit operates as follows:

A digitised, PCM encoded input signal n bits deep is multiplied separately with n bit sine and cosine functions from "phase tables" which are scanned as the input samples are taken. This emulates the oscillatory part of the LCR oscillator. The outputs from the multipliers are each input to a single pole recursive filter, and the contents of the registers in the SPRF define the state of the oscillator. During each clock cycle, the contents of the registers are decremented using tables. These form the analogy of a resistor in the conventional LCR circuit, and are realised as values stored in a PROM. The decremented values are then added to the multiplier outputs (which may be positive or negative) to form the next register contents. The registers give a direct reading of the complex amplitude of the oscillator at any particular time slice. If it is desirable to use a complex input, then the two multipliers must be replaced with a unit which can perform a complex multiplication. In the single bit case this can easily be implemented using a PAL or lookup tables. We are not yet concerned with complex input signals.

The flexibility inherent in this approach to communications engineering has yet to be fully explored. We list below some of the generalisations which may be employed to extend the capabilities of the circuit for particular applications.

(1) Phase table. Lookup tables may be used for simple cases where the lists of values are not extensive. These tables may be soft and could be changed "on the fly" by a microprocessor; most elegantly using a dual port RAM.

Where the stored functions are sinusoidal this represents an arbitrary degree of frequency agility. In this memorandum we give an example where the functions are generated automatically and adaptively according to an input signal.

(2) The multipliers. In the general case it would be wrong to restrict this function to multiplication, any required function for specific applications to more general processing could be stored in a PROM. As we are as yet only concerned with correlation and anti-correlation we shall use multipliers.

(3) The decrement tables. These define the "Q" or quality factor of the circuit in the most convenient form for our purposes which is in terms of a bandwidth. As in the case of the "phase tables" these may be soft and also updated "on the fly" by a microprocessor. This endows the circuit with an agile "Q" or bandwidth which may for example be used to accommodate different data rates in a modulated carrier. In this memorandum we will use oscillators with a constant bandwidth, but will note how a variable bandwidth may be used to improve further some of the properties of the phase locked loop. Where the next state of the "oscillator" depends on a more complicated function of its own past history or that of its "neighbours", the restriction to single pole recursive filters may be dropped.

(4) Outputs. The outputs are in cartesian form and represent real and imaginary ampltidues. Where a transform of this output is required, the required function may be performed by a PROM. For example, for the rapid acquisition circuit, power may be computed using a PROM (sum-of-squares) or where the phase is required to greater than single bit accuracy an $(x, y) \rightarrow (R, \theta)$ transformation (PROM) may be used. In the case of the phase lock loop, no more than one bit is required from each register.

The main disadvantages of the circuit are when narrow bandwidths are required or where the input word is wide. The width of the registers required (m) depends on the maximum amplitude to be accommodated and the finesse of the decrements which are needed. If $m > 16$, lookup tables are not convenient and a hardware multiplier may be needed. The circuit is designed to be efficient in its use of linear word widths however, and in most cases (especially where the signal is very noisy) a one bit input is sufficient and only the top few bit of the SPRF

registers are required for output. One exception to this situation for example might be where it is required to steer a null over a strong interfering signal in the band by computing weighted sums of the registers in a parallel bank of oscillators spread out in the frequency domain. This is an example where all of the finesse of the registers may be required together with a wider input word to emulate a "phased array".

There are two more general advantages of using this kind of circuit. No filter is required on the input, except where there is a danger of aliasing from out of band signals: the circuit is its own filter. Also much of the required signal processing may be done by using the circuits' own time constants.

In the next section we will demonstrate how an ordinary phase lock loop may be built using the "synthetic oscillator" by comparing figure 1 with figure 3. We will then demonstrate how this principle may be extended to more complicated input signals by modelling their structure and by imposing the stated requirement that the "error voltage" is invariant under a particular modulation scheme. We then apply this argument to the particular case of PSK and design hardware which is tested with real signals. Finally, we indicate that a data rate search is possible and introduce the idea of data recovery.

2.3 Synthesising a Conventional Phase Locked Loop

Figure 4 shows pictorially an LCR oscillator response curve in the vicinity of a sine wave. The amplitude of the response is $\sqrt{(\text{REAL})^2 + (\text{IMAGINARY})^2}$. If we denote the frequency of the oscillator at resonance as Fo and excite it with a frequency Fc which is offset from Fo, the oscillator is forced and the response will be reduced.

In figure 5, the response of the oscillator in the complex plane is shown. If the input frequency Fc is precisely the oscillator centre frequency Fo, then the oscillator response may be described by a stationary vector R with radius r. The angle θ of the vector is the phase of the incoming waveform. If the input frequency Fc begins to drift slowly away from the oscillator centre frequency Fo, the vector R in figure 5 begins to rotate around the origin with a reduced amplitude r in accordance with figure 4.

If we were to compute the values of $d\theta/dt$ and integrate the results over a period of time, the result would serve as an error voltage which may be used to alter the centre frequency of LCR circuit. Hence we can implement a phase lock loop by observing the states of the oscillator in the complex plane

Now compare figure 1 with figure 3. We may draw an analogy between the VCO of figure 1 and the phase (sine and cosine lookup) tables of figures 3. We may also draw an analogy between the multipliers in each diagram. In order to create a phase locked loop using a "synthetic oscillator", it is necessary to use the outputs of the single pole recursive filters as part of a loop filter.

The information about the state of the synthetic oscillator is available in Cartesian form from the SPRFs. A simple minded method of using this to implement phase locking would be to use a PROM to convert the cartesian form into and use this to drive a digital to analogue convertor. The output from this would then drive a VCO which would replace the "phase tables". However we make two observations:

(1) The decorrelation time of the oscillator is governed by its bandwidth (or "Q"). It is therefore not worth observing the contents of the SPRF registers as often as the oscillator samples the signal. Successive samples will only contain substantially independant information when they are separated by intervals greater than or equal to the oscillator decorrelation time. We call this "sparse sampling" (see ref 3).

(2) In a noisy environment it has been shown (see ref 1) that it is not (except where the noise has perverse characteristics) worth quantising the input signal to greater than one bit accuracy. We might guess that in a very noisy environment, the phase noise present in the oscillator will similarly remove the necessity to quantise the oscillator angle to more then one bit accuracy. (This implies two bits since there are two degrees of freedom in the oscillators.) The apparent phase noise is seen in the results section

Figure 6 shows how a phase locked loop may be constructed using a "synthetic oscillator". The most significant (sign) bit from each of the SPRF registers carries no amplitude information, but they together define which quadrant of the complex plane (represented by an Argand diagram) the oscillator state lies in. These two bits are input to a finite state machine (FSM) which computes $d\theta/dt$ (see ref 4 for a discussion of FSMs). The state machine may be constructed using only a PROM and a latch. It will examine the state of the oscillators at two sucessive time intervals: if the oscillator state has rotated one quadrant anticlockwise it will add one to the integrator. If the oscillator state has rotated one quadrant clockwise it will subtract one from the integrator. If the same quadrant has been seen twice or if the rotation is two quadrants, the value of the integrator is not changed. The structure of the FSM is shown in figure 7. We do not enumerate the memory contents since it will be shown later that this hardware does not need to exist explicity.

For the purpose of this discussion it is convenient to consider the integrator itself as a single pole recursive filter with a time constant of many symbols (see fig 8). This allows us to endow it with a finite, exponential memory characteristic and also to provide it with "hidden bits" which allow the required finesse in the decrement tables. The dynamics of this part of the circuit will form a subject in its own right, however this is subsidiary to the main argument and so will not be pursued further.

The output of the integrator is now a digital representation of the frequency which the "phase tables" are required to produce. In a later section we will describe how to construct the required sine and cosine generators with selectable frequency using only an adder, a latch, and two exclusive or gates.

This section has described how to construct a phase locked loop using a synthetic oscillator. Note that since there is only one oscillator in isolation (unlike the 128 described in ref 3) much greater bandwidths are now accessible. TTL clocked at 20 MHz will provide a PLL range of 10 MHz, and an ECL design might furnish ranges of 100 MHz or greater.

This phase lock loop should perform equally well if not better than the conventional counterpart. The investigation and characterisation of this circuit is left as an exercise for the reader.

2.4 Generalising the Phase Locked Loop

We now generalise the above principle. In section 2.3 it was shown that a noisy sine wave frequency and phase ambiguity may be reduced to a sparse series of single but complex scalars. This is an example of data compression where the entropy in the signal has been reduced and almost only the required information, that which removes the ambiguity in the input sine wave, has been retained. The signal was then reconstructed using this information by the state machine, integrator and sine and cosine generators.

We will now repeat this procedure where there is further ambiguity in the signal, known as modulation. We will effect a partial reconstruction using a sparse series of complex vectors. The complete reconstruction will be dealt with in part 4 where we show how to recover the data.

No general formulation has yet been derived which extends this principle to all modulation schemes, but we can envisage no fundamental difficulties in doing so. A more general but not complete scheme will be discussed in part 5.

Consider figure 9. This shows three examples of the power spectrum of the carrier which has the general form of $e^{i(\omega t + f(\phi))}$ for three case of $f(\phi)$. The first case is where the function $f(\phi)$ is a constant and the carrier is a sine wave. The second case is where the $f(\phi)$ is a square wave with amplitude π. The spectral line at the carrier frequency has become split into two main lobes at +/- the square wave frequency (+/- half the "data rate") from the original Fc. Other subsidiary sidelobes appear at +/- (2n + 1)X (data rate) where n is a natural number. These lobes have a separation equal to the data rate where the square wave is a special case of data: the binary symbols alternate between zero and one. The third diagram illustrates the case where there is phase ambiguity (information) in the carrier, and the sine wave is modulated with real data. The data is in the form of a "random" telegraph signal with a constant chip rate. As $t \rightarrow \infty$, the envelope of the power spectrum tends towards $\frac{S_{IN}(x)}{x}$, other phase change functions were considered in part 2. Their importance will become apparent in later discussions, they are not relevant here.

There are several observations to be made regarding figure 9.

(1) The minor sidelobes in figure 9(b) and figure 9(c) contain little power from the signal and relatively little information.

(2) The discontinuities in the phase function in 9(c) occur in the same place at those in 9(b); the only difference is that some are missing.

(3) We have seen in part 2 that in differential phase shift keying (DPSK) it is the ABSENCE of edges in this function which carry the information.

(4) It is the absence of edges in the "random" telegraph signal which produce the fine structure of the main lobe, and indeed of all the lobes, in fig 9(c).

(5) We are interested not in the power spectrum of the signal as shown in fig 9 but in the PHASE. In the phase is the information which is sought.

It is desirable to examine the fine structure of the central lobe of the modulated carrier in 9(c) on the timescale of its inverse width. We will first consider the naive approach to this problem and its deficiencies.

Consider a synthetic oscillator (section 2.2) deployed in the same region of frequency space as a PSK modulated carrier. We will choose the frequency of the oscillator Fo to be at or very near the carrier centre frequency Fc. We will also endow the oscillator with a "Q" such that the bandwidth of the oscillator approximately matches the width of the central lobe of the modulated signal in figure 9(c). This arrangement is depicted pictorially in figure 10. We will attempt to infer frequency error from the dynamic behaviour of the registers in the oscillator, in particular by using only the sign bit. The motion of the single bit quantised registers in the complex plane is depicted in figure 11.

11(a) Fc = Fo no phase reversal. State remains stationary.

11(b) Fc = Fo with phase reversal. State moves to opposite quadrant.

11(c) Positive frequency error, no phase reversal. Clockwise drift of oscillator state.

11(d) Negative frequency error, no phase reversal. Anti-clockwise drift of oscillator state.

11(e) Positive frequency error with phase reversal. Net anti-clockwise drift of oscillator state.

11(f) Negative frequency error with phase reversal. Net clockwise drift of oscillator state.

The above indicate that it is not possible to examine the top bit of the registers and discover the sign of frequency drift when (hard) phase reversals are present. An attempt to construct a phase locked loop using the method outlined in section 2.3 will require finer angle quantisation. We might predict a situation where extremes of phase noise will make finer quantisation of angle an inelegant solition; a balanced approach is required.

Physical intuition provides a solution to the problem. Since we have chosen to view figure 9(c) as a mixture of 9(a) and 9(b) it is a small step to note the mirror symmetry of the PRK spectrum about Fc and to consider deploying two oscillators, one on each side of the carrier centre frequency. This is the key to generalisation.

We deploy the oscillators so that the centre frequencies are separated by the inverse data rate and the oscillator pair lie symmetrically over the carrier centre frequency. This is shown pictorially in fig 12. It means that each oscillator covers one of the main lobes of the power spectrum shown in figure 9(b). The bandwidth of the oscillators is chosen to be approximately equal to the inverse symbol time. The state of the oscillator

pair is now deemed to be part of the state of the model which describes the incoming signal and will now replace the phase detector in the conventional phase locked loop. The VCO will be replaced by a pair of complex $e^{i\omega_1 t}$ and $e^{i\omega_2 t}$ waveform generators. These will form the "phase tables" for the oscillator pair where ω_1 and ω_2 are separated by the inverse data rate. We will now investigate the state transitions of the oscillator pair for various classes of segment in the input signal which are one symbol long, by observing trajectories in the complex plane as before. Note that we are not yet interested in absolute phase, and also that the details of the paths are profoundly affected by both the nature of the phase transitions and the bandwidth of the oscillator.

The approximate behaviour of each oscillator is enumerated in figure 13 for each case that appears in figure 11, and is largely self explanatory. Additive, Gaussian noise will make the end points of the vectors uncertain and in the limit where the noise power in the band of the oscillators exceeds the signal power this will affect the sign bits.

For the purpose of constructing a phase locked loop it is desirable to find an objective function which will return a frequency error given a series of single bit complex vectors from the SPRF registers. This is easily done by inspection. The real scalar which returns the frequency error is

$$\frac{d\phi_{01}}{dt} + \frac{d\phi_{02}}{dt}$$

where ϕ is the phase of the registers and the derivatives are obtained from successive samples of the SPRF registers.

Figure 14 shows, in block digram form, the structure of the generalised phase locked loop.

It is now instructive and interesting to note for the purpose of synchronous communications, the synchronised oscillator in the receiver no longer exists explicity. However the state of the oscillator which should be there may be inferred from the states of the oscillator pair 01 and 02. We have eliminated the need for removing the phase reversals by frequency doubling.

3 COMPUTER SIMULATIONS

3.1 Introduction

Before we can construct hardware, the principles which we have described have to be tested by simulation.

To validate the assumptions which were made to produce figure 13, a computer program was developed in November 1985 to simulate the circuit shown in figure 14. This was written in Fortran 77 and executed with a Data General MV6000 computer; the Argand diagrams were plotted using a Calcomp 81 graph plotter. A selection of the results are reproduced in this memorandum.

Section 3.2 illustrates the conditions which are depicted in figure 13. Section 3.3 shows the results of simulations where noise was added to the synthetic signal which was built up inside the computer. This was to test the validity of the assumption that noise should only perturb the behaviour of the oscillators and not radically alter their behaviour.

Section 3.4 shows the behaviour of the centre frequency and integrator in the complete system when the feedback loop is closed and the system is "acquiring" the signal. This is done under differing noise conditions.

3.2 Argand Diagram Tests

Each oscillator had the initial condition of zero energy content, so all of the parametric plots start from the origin. The units on the axes are linear and represent amplitude; they need to be scaled to be interpreted physically. For the purpose of this discussion such interpretation is not necessary.

Figure 15 shows the behaviour of the simulated register contents of O1 and O2 where the input is a pure sine wave and Fo = Fc. The curves are plotted over 1000 input samples which in the program was just over two symbols long. The plot does not appear to make a complete revolution of 2π, however this is an artifact of the initial conditions, and plots where 2000 samples have been taken show a complete extra revolution.

Figure 16 shows similar plots except that hard (abrupt) phase discontinuities are now made every 480 samples. The symbol length was defined to be 480 samples.

Figure 17 is similar to figure 15 except that there is now a frequency error of 0.25 times the inverse data rate = 0.5 of the oscillator separation.

Figure 18 is similar to figure 17 except that now there are abrupt phase reversals every 480 samples.

3.3 Behaviour of the Oscillators Stimulated with a Noisy Signal

Simulations were performed where the input signals were noisy. Gaussian noise was added over the entire frequency band which was used as an input to the oscillators. This band was approximately 240 times the bandwidth of each oscillator.

Figure 19(a) shows a similar plot to that in figure 15(a) but in such a way that the signal to noise ratio has now been reduced from infinity to 30 dB Hz. In the 0 to 12 kHz input band used this represents an absolute signal to noise ratio of about -10.8 dB. Note that the "Q" of the oscillator does much of the averaging and that 100 points are not needed to describe a close approximation to the plotted trajectory. This is a graphic illustration of why it is only necessary to sample the registers sparsely. 30 dB Hz is a typical noise level that is found in 50 bit per second satellite communications using DPSK. The fundamental limit for an error rate of one bit in 10^5 in such a system is about 27.5 dB Hz.

Figure 19b shows a parametric plot of the registers of oscillator 2 over 8 symbols at a signal to noise ratio of 21 dB Hz, or approximately -16 dB if the entire input band is counted. Phase reversals were present.

3.4 Phase Locking Simulations

We now turn our attention to the behaviour of the simulated system of figure 14. Three sets of two graphs are reproduced here, figure 20 is at a signal to noise ration (S/N) of infinity, figure 21 is at an S/N of 30 dB Hz, and figure 24 is at a S/N of 24 dB Hz.

Part (a) of each figure shows the centre frequency of the oscillator pair vs the number of samples. The vertical scale is in units of π per sample.

Part (b) of each figure shows the value stored in the integrator in arbitrary units vs time. The time constant of the integrator was chosen to be five data symbols long. This was because it was considered that the maximum frequency drift which could occur without producing ambiguities in the data (where we are relying on the principle of separating the components in the signal using time constants) was about one symbol in five, or π error in 5π of modulated signal.

Note that the "hidden bits" in the integrator allow changes in the register which are smooth compared to the input ("visible") bits.

4 HARDWARE

4.1 Introduction

We now derive the structure of the hardware which will be used to test a particular example of the generalisation of the phase locked loop. The synthesis will be broken down into a number of discrete sections.

The first will deal with the single bit sine and cosine generators and how the frequency of the lower oscillator is input as a 16 bit word. The frequency of the second oscillator is found by adding another word (the data rate word) to the number representing the frequency of the lower oscillator. For the purposes of this particular implementation, the data rate is constant.

The second section will show how the oscillators are realised, and how all four components are multiplexed together to save space and hardware.

The third section will deal with the state machine and the integrator which are shown in figure 14. The reason why a dotted line was drawn around these components was that they do not exist explicitly in the circuit; instead they are emulated by a microprocessor. We shall see that the oscillators compress the data so efficiently that a relatively slow microprocessor oversampling the registers by a factor of 8 has about 2500 machine cycles available to analyse the output between each sparse sample. The following section will then outline the microprocessor software.

The clock generator and input circuits will then be described, showing how the detailed circuit timing, interrupts and outputs are derived using the microprocessor clock.

Apart from the microprocessor, the digital to analogue convertors, and a PAL, the circuit was constructed from standard TTL logic chips. The chip count of the prototype circuit was 32 including the input circuits and the Argand diagram display facility. It was built on a double eurocard with space to spare, as with other circuits we desire eventual incorporation in small terminals. After amalgamation with other circuits and hybridisation, an appreciable reduction in size can be envisaged.

4.2 Sine and Cosine Generators

Figure 23 shows the basic structure of the sine and cosine generators. The latch is a store upon which modulo 2^m addition is performed. The frequency input to the adder determines the rate at which the system cycles completely, and we call each such cycle equivalent to 2π radians. The outputs are the most and second most significant bits of the latch. For those who wish to generate sine and cosine tables to more than one bit accuracy, we suggest that more bits are used from the latch and either the AMD29526-29529 sine and cosine generator chips or PROMs are used to generate the waveforms. We note that the two most significant bits follow the sequence 00, 01, 10, 11 and that if these refer to quadrants, then the signs of (sin, cos) progress as (+,+), (+,-), (-,-), (-,+). The inset in figure 23 shows how a section of the PAL (programmable logic chip discussed later) is used to make the conversion to sine and cosine, albeit with a phase shift. Of course the integrator does not "know" about this phase shift.

Figure 24 show the structure of the sine and cosine generators which we shall use. There are two latches to accommodate the two different frequencies used; the circuit multiplexes between the two. The frequency input is that of the lower oscillator (of course the integrator does not "know" this either), and the upper frequency is calculated from both the lower frequency and the frequency difference (data rate) in the lower adder. We are not yet interested in variable or adaptive data rates. Therefore the generator timing is chosen so that one line in the adder may be tied to one of the system clocks to provide the oscillator separation (see section on timing etc). We do not include the circuit diagram, this should be obvious from the layout.

4.3 SPRF and Multiplier Structure

In our example the bandwidth which we are interested in is only 12 kHz. Therefore it would be wasteful to build the circuit such that all SPRFs are distinct. The prototype circuit multiplexes all four parts together.

The full circuit diagram of this section of the circuit is shown in figure 25 to illustrate the hardware compression achieved by multiplexing. The PAL is configured to act as the sine/cosine conversion logic, the multipliers, and the real/imaginery multiplexer. Note that the multiplier is an exclusive OR and not an exclusive NOR. This arrangement inverts the complex plane through the origin. This is of no consequence since the adders have been wired to perform two's complement arithmetic on the complement of the multipliers. Hence the complex plane is inverted twice.

The number of bits required in the registers may be derived from the bandwidth. We will derive the necessary values here as the same basic argument applies to all circuits of this type. The form of the decrement tables (or function) is value → (1-a)X (value) where a is related to n such that (1-a) n = 1/e n is the logarithmic decrement period measured in cycles. For a sampling rate of 48000 samples per second 1/50 second (for 50 baud data) corresponds to 960 samples. (1-a) is then given by:

$$(1-a) = \left(\frac{1}{e}\right)^{\frac{1}{960}}$$

$$= 0.99896$$

The number of bits m required in the registers is given by the maximum excursion which can occur before the decrement equals the input. This excursion is:

$$\frac{1}{1-0.99896}$$

$$= 961$$

Since there are both positive and negative excursions, the registers have to be able to accommodate a range of 1922, this requires 11 bits. These 11 bits are on the LEFT of the decimal point. As stated in part 1 (ref 3), we will require as many "hidden bits" on the right of the decimal point to allow sensible decrements when there are very small numbers in the registers. Therefore to accommodate the full dynamic range we require 22 bits.

There are two observations to be made here. Firstly, 50 bits per second is a relatively low data rate requiring an unusually high "Q" or low bandwidth, and this bandwidth is independent of the total range over which the phase lock loop is to operate. Wide words are a small price to pay here since as they become wider with decreasing data rate more work is done by the SRPRFs, and the sparse sampling becomes sparser pro rata. Secondly, any hard limiting causing a drop in the effective "Q" will have a profound effect upon the motion of the oscillators in the Argand plane caused by modulation symbols. It may also effect the demodulation efficiency. (Those who like to simplify circuits might like to try using the identity function in the decrement tables and let the hard limiting define the "Q").

Since we are not yet interested in demodulation but only wish to test the principle of generalised phase locking using these circuits we have chosen to use 16 bit registers and to accept a certain amount of hard limiting. 8 bits are used for decrement purposes, and the extremes of the decrement tables are altered to prevent overflow and the resulting undesirable (in this case) toroidal topology of the Argand plane.

Although we could conceive simple hardware to perform decrements, PROMs were considered to be more versatile for development. That is why they were chosen. A program was written on a Data General MV20000 computer to produce the contents of the tables with user selectable "Q" and automatic overflow prevention. The resulting bit patterns were stored in a file and down loaded to a Stag PPZ chip programmer. Two 64K x 8 AM27512 PROMs were used, one for the "hidden bits" and one for the "visible bits".

Four 4 x 4 register files were used to store the real and imaginary contents of both frequencies. The eight hidden (least significant) bits are not affected by the adder, therefore only two four bit adders were required. No latches were used in the loop; the timing was such that the register files could operate properly.

4.4 The Microprocessor

The microprocessor used to emulate the state machine was a Hitachi HD63P01M1, chosen because it would accept a 2764 8K x 8 erasable PROM as a piggyback. This was useful for developing the machine code. The microprocessor has four ports. Port 1 was used to input the (top eight) bits of each register in turn. Ports 2 and 3 were used to output the frequency required to the sine and cosine generators. This left port two, which contains a serial interface, for experiments involving data input/output.

The processor was interrupted by a pulse generated from its own clock which had been divided down. Each interrupt represented a sparse sample where the processor reads the registers, performs the necessary calculations, and outputs a new frequency. After each interrupt the processor goes into "sleep" mode and awaits the next interrupt.

Each symbol in the prototype system was 20 mS long. The interrupts were arranged to over sample the contents of the registers by about a factor of 4, ie the registers are sampled approximately every 2.5 mS. This was done to avoid the "wagon wheel effect", similar to the appearance of a wagon wheel going the wrong way on a cinema screen because of the sampling rate of the camera shutter. This was to enhance the pull in range of the phase lock loop.

The microprocessor code is not trivial as parts are time critical. It is also tedious, so rather than reproduce it in full we explain the main functions step by step

(1) The contents of the four registers are mirrored twice in the processor internal RAM. One set is for the current sparse sample, another for the previous sample. When interrupted, the processor reads the four registers each at the correct multiplexer time slot from the SFRFs and puts the contents into the current sample registers.

(2) The most significant bit from each of the registers in the current and previous sample are arranged and loaded into the index register. This is used to point to an address which loads another register, the phase error register, with the value of the total phase error. This has the range [-2,2]. This method of calculating the error sacrifices a small amount of the memory for speed.

(3) The phase error register is added to a sixteen bit signed integrator register. The integrator register is then duplicated and right shifted in the appropriate manner and subtracted from the original. This emulates a single pole recursive filter. The number of right shifts will determine the time constant.

(4) The integrator register is duplicated again and once more right shifted in an appropriate manner and added to another sixteen bit register in which the current frequency is stored. In this case the number of right shifts determine the gain.

(5) The current sample registers are written over the previous sample registers and the processor puts itself to sleep. Note that it is no good putting the micro into a wait loop since it needs to exit synchronously in order to read the oscillator states at the correct time.

4.5 Input and Clock Circuits, Circuit Timing

In the input circuit, a Lin-Cmos operational amplifier was used to keep the power supply to a single +5V rail. Care was taken to minimise the DC component entering the oscillators. A D type latch was used to hold the input steady during a sample period. Edges occurring during a sample period can propagate through

the very fast PAL and upset the adder, making the register values take in the wrong value. Curiously, since the basic circuit is stable in the sense of the physical (LCR) system which it emulates, this fault tends to degrade rather than ruin its performance.

The clock generator was a simple set of dividers and invertors, with differentiators to provide the NMI interrupt to the processor and the write pulse to the register files. The circuit timing diagram is shown in figure 26 for a duration of just over two input samples.

The circuit which is used to display the Argand diagrams is not shown but is very simple. It consists of four AD7545 12 bit digital to analogue convertors and an address decoder to enable each in turn from the clocks.

The sampling rate was chosen as 48 kHz rather than 24 kHz in this particular circuit. This was in order facilitate correct operation of the sine and cosine generators and to avoid possible aliasing near the top of the band which was predicted by theory.

The sixteen bit representation of frequency was chosen so that one line could be used to introduce the frequency dither required to define the data rate; this allowed one line to the data rate adder (in figure 24) to be tied to clock 2. The frequency resolution was chosen to be just less than 1 Hz; more accurate frequencies were approximated by dither.

5 RESULTS AND DISCUSSION

5.1 Introduction

The circuit was built to test the feasibility of wider use of the synthetic oscillator circuit in communications systems, and in particular to find out if the prototype PLL could compete with more traditional designs. With this end in mind, there was no attempt to optimise the circuit or to try alternative configurations.

One value of integrator time constant, gain and oscillator bandwidth was used throughout. It was noted that the required "Q" could not be attained due to the hard limiting.

The time constant was chosen to be approximately three symbols. We have said that the feedback in the phase lock loop must be independent of the modulation. A time constant of three symbols smooths any fluctuations caused by the range of phase change functions which may be used (see ref 5). The gain was chosen to be about 3/80 Hz per π error in a symbol.

The results are divided into two sections. The first illustrates quantitative results showing the behaviour of the Argand diagrams. The second set of results are preliminary measurements where the tracking performance and pull in range are tested, and the effects of noise are measured.

5.2 Qualitative Results

The results in this section (apart from the spectra) are shown as pairs of Argand diagrams. The lower frequency is on the left and the upper frequency is on the right. The real (cosinusoidal) response is plotted horizontally and the imaginary response vertically. All exposures were 1/15 sec (except for fig 27(c) - averaged noise). The Argand diagrams were displayed on a Tektronix 7704A oscilloscope with a 7A26 plug-ins. The display circuits were built in to the phase lock loop board. The Argand diagrams are state diagrams of the oscillators, and the brightness of each point is proportional to the time which the oscillator spends in that state. Note that, as in part 1 (ref 3), the order in which the points are accessed is not obvious from the photograph. Some of the displays are limited by sharp borders which appears to confine the state space within a square. This is due to the finite register size which hard limits the maximum oscillator amplitude. Methods of reducing the register size while retaining the dynamic range of the system is outside the scope of this memorandum and will be discussed in future publications.

Figure 27 shows the basic responses of the circuit. 27(a) shows the response of the system to wideband noise, ie Gaussian noise which has an approximately flat spectrum from 0 to 12 kHz. This was the circuit's input frequency span.

Figure 27(b) shows the result of the same input with the addition of a large sinusoidal signal well outside the band of both oscillators. This illustrates an interesting consequence of using single bit quantisation: the principle of superposition no longer applies and the amount of "energy" which remains within the bands of the oscillators is reduced. The large signal has "stolen" some of the energy from the oscillator bands. If the circuit is to be operated where large, unwanted signals are present in the band, greater than single bit quantisation will be required on the input.

Figure 27(c) shows the same display as 27(a) except that the exposure was made for ten seconds. Although some perturbations in the probability distribution are still visible, the structure which was apparent in the shorter exposure has mostly been swamped. Compare this with the equivalent illustration in ref 3: the finer quantisation has resulted in a smoother representation.

Figure 27(d) shows the response of the oscillators to a pure sine wave which was tuned to be at the centre frequency of the upper oscillator, ie 25 Hz above the centre frequency of the oscillator pair. The upper oscillator has become stationary and occupies very few points in its state space. The lower oscillator still responds strongly but rotates around the Argand plane at the rate of 100 π per second. Note that the occupied states form a square: this is because the registers are saturating. As we have discussed before, the direction of rotation depends on the sign of the frequency error.

Figure 27(e) shows the effect of a frequency error of 200 Hz from the centre of the system. The oscillator with the greater frequency error is rarely saturating, and the motion in the Argand plane is now nearer to a circle. In figure 27(f) the frequency error is now 400 Hz. The occupied states now approximate circles and have contracted considerably. The motion in state space averages to an annulus became of the complicated approximations present in the sine and cosine generators and the decrement tables.

Figure 28 shows the system phase locked to a signal which is in varying amounts of noise. Although there are indications that the system remains in lock below a noise level of 20 dB Hz (an absolute value of less than -20 dB) this has not been shown. The display is not easily distinguished from noise at signal levels of less than 25 dB Hz, and the signal to noise ratio cannot be measured using our method (detailed later) to within +/- 1 dB at the lower signal levels. In such regimes the definition of lock itself becomes ill defined (ref 1).

As the signal recedes into the noise, the phase lock loop begins to show frequency errors which involve ever greater excursions from the signal frequency. In the limit where the signal has disappeared completely, the loop executes a "drunkard's walk" in the frequency domain. The exact characteristics of this walk have not been investigated, but suffice to say that in the close approximation to Gaussian noise which we use the signal can disappear for tens of seconds: on its reappearance at the same or near the same frequency the loop will rapidly reacquire. This behaviour is particularly useful where fading is encountered, as it often is in satellite channels. We have not yet considered lock detection, and indeed this may not be necessary. In the case where the signal is drifting in frequency and is too weak to lock to, loss of signal may be detected by the rapid acquisition circuit (ref 3). If the phase lock loop frequency does not lie close enough to the frequency indicated by the rapid acquisition circuit, the PLL frequency may be reset accordingly, ie to a value where it will reacquire the signal.

Figure 29 shows the effect of a modulated carrier on the state sequence of the oscillators. There are 50 "hard" or quasi-instantaneous phase reversals per second. 29(a) shows the Argand diagrams where the signal to noise ratio is infinity, and 29(b) is the same where the signal to noise ratio is 40 dB Hz. We show only the effect of "hard" phase reversals for the following reasons:

(1) Oscillator saturation strongly affects the paths in state space.

(2) The elucidation of oscillator behaviour due to phase modulation is the subject of another publication and is not relevant to phase locking.

Different phase change functions result in various paths such as "banana" or epicyclic type motions depending on various parameters.

Figure 30(a) shows the noise spectrum which was added to the signal. The display which was used was a Tektronix 7633 oscilloscope with a 7L5 spectrum analyser plug in. The display covers 20 kHz, and the spectrum is flat to within about +/- 1 dB up to 12 kHz, where it falls off at a rate of about 10 dB per kHz. The noise source was a VHF tuner with a shorted input followed by a filter to shape the output spectrum to the required characteristics.

Figure 30(b) shows how the signal to noise measurements were made. The noise was centered between two horizontal line on the graticule, and the distance in centimetres between the middle of the trace and the peak to be measured was multiplied by 10 and added to the filter bandwidth (in this case 10 Hz). The figure shows a measured noise figure of 36 dB Hz. As the vertical scale was logarithmic (10 dB per cm) there was little error on the peak unless the measurement was less than 30 dB Hz, however the error on the noise "floor" was about +/- 1 dB.

5.3 Quantitative Results

We have concentrated on the parameters which most concern us, ie pull-in range in noise and tracking capability.

The requirements for the specific application of low data rate satellite communications are that the pull in range should match the resolution of the rapid acquisition circuit, and that the tracking rate should be at least 0.38 Hz sec^{-1}. The former requires a very substantial improvement over existing designs.

We have said in the hardware section that only one value of integrator time constant and feedback gain was used. Changing these values clearly does affect the acquisition and tracking performance, although measurements to enumerate these interrelationships remain to be done. We have chosen the values stated in the hardware section to illustrate one example of how the device will perform.

Signal acquisition appears to be reliable up to over +/- 150 Hz from the carrier. At +/- 180 Hz the acquisition is unreliable and has a probability of about 0.5, in any case acquisition in these circumstances may take from 2 to 10 seconds. At +/- 100 Hz acquisition takes less than 0.5 sec.

Figure 31 shows the behaviour of the loop centre frequency vs time for different initial conditions. These results were obtained by decoding the multiplexed digital values from a Tektronix 7D01 logic analyser. These graphs are examples only of the behaviour of the centre frequency of the circuit as lock is approached. In the nomenclature of single harmonic motion (which is not quite appropriate in this nonlinear system) the behaviour may appear to be over, under or critically damped when the same starting conditions were applied. We have found, however, that increasing the time constant does eventually lead to consistantly "underdamped" behaviour and eventually loss of lock - as physical intuition would lead one to believe.

The acquisition range fulfils the requirement comfortably and shows an improvement of between one and two orders of magnitude in pull in range when compared with conventional systems used for low data rate satcoms.

Figure 32 shows tracking performance in Hz sec^{-1} for various signal to noise ratios. The requirement is exceeded by more than three orders of magnitude. Note the inaccuracy where low signal to noise ratios are plotted: this is due both to the difficulty in measurement at these signal levels and to the difficulty of deciding if the circuit is still in lock. These measurements were taken using a sudden start to the oscillator sweep and with no attempt to optimise the oscillator structure or parameters. Hence the optimised, dynamic performance is expected to be considerably better.

The intercept with the x axis could be interpreted as the signal to noise ratio below which the circuit will not track at all. This is a dangerous assumption, as the definition of lock has become hazy in these conditions and the behaviour is not (insofar as the noise is random) deterministic.

The circuit appears to interpret the phase reversals as added noise. We are not yet concerned with the absolute phase of the oscillators and we have not yet synchronised the oscillator separation to the data rate. There will therefore be random errors in the error register within the processor, but this does not destroy the invariance criterion we have stated for the feedback. The phase lock loop will acquire and maintain lock when the data rate is smoothly varied from 0 Baud to well over 500 Baud. This indicates the possibility of variable or adaptive data rates and we shall see in a future publication how a data rate search could be made by varying the oscillator separation.

6 CONCLUSIONS

The performance of the generalised phase locked loop applied to PRK has exceeded expectations. By using the power of the pure representations available in digital electronics in conjunction with physical insight we are developing a fundamentally new approach to communications engineering.

The undesirable frequency doubling necessary for carrier recovery from PSK signals has been eliminated. For synchronous demodulation the frequency and phase of a "ghost" carrier may be inferred from the combined state histories of a pair of synthetic digital oscillators.

There is an improvement in acquisition range of between one and two orders of magnitude over systems which are used at present. The tracking performance exceeds the requirement by more than three orders of magnitude. This removes some of the problems associated with low data rate satellite communciations to and from moving vehicles. The system is robust when signal fading into Gaussian noise occurs.

This memorandum has described only one basic system which uses the synthetic oscillator approach. A few of the possible extentions to the PLL idea are listed below.

(1) The order of the loop may be changed to suit varying requirements. The loop which has been described is analogous to a conventional second order loop. Clearly a third order loop, for example, may be built by cascading two (software) integrators.

(2) The range of the loop can easily be extended to at least 10 MHz using TTL or 100 MHz if ECL is used.

(3) There is the possibility of rejecting unwanted signals by increasing the number of quantisation levels on the input and providing a better approximation to superposition.

(4) The scheme can be extended to different modulation schemes by varying the number of oscillators and their method of deployment, ie "Q" and centre frequency, to retain invariance of the feedback during modulation.

In the following publications we will attend to demodulation, and will address some of the problems of the current scheme. For example, the word lengths in the registers are considered unnecessarily long and we seek a better representation and method of defining the "Q".

We shall also explore some more general schemes for the communication of information using (mainly) the synthetic digital oscillator approach. For example, the circuit envisaged to synthesise spectra described at the end of part 2 (ref 5) is considered capable of simulating a signal with the same sidelobe structure as common modulated carriers but with the "wrong" phase behaviour. We wish to make use of this capability.

7 ACKNOWLEDGEMENT

The author woild like to thank Dr S P Luttrell for much help with the theory, and many excellent ideas which have helped to make this whole project become a reality.

8 REFERENCES

1 Floyd M Gardener: "Phaselock techniques", John Wiley, 1966.

2 J A S Pritchard: Real time demodulation of signals using statistical methods", Electronics letters vol 21 No 25/26 5 December 1985.

3 RSRE Memorandum 3925 "A compact digital communications system part 1: Rapid acquisition" S P Luttrell and J A S Pritchard.

4 RSRE Memorandum 3992 "Compact bit slice writable control store for experimental information processors" J A S Pritchard.

5 RSRE Memorandum 3994 "A compact digital communications system part 2: a stored waveform PSK modulator" J A S Pritchard.

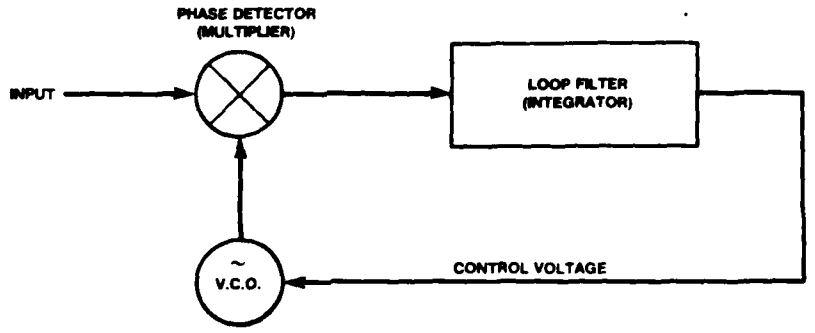

FIG 1 BASIC PHASE LOCK LOOP

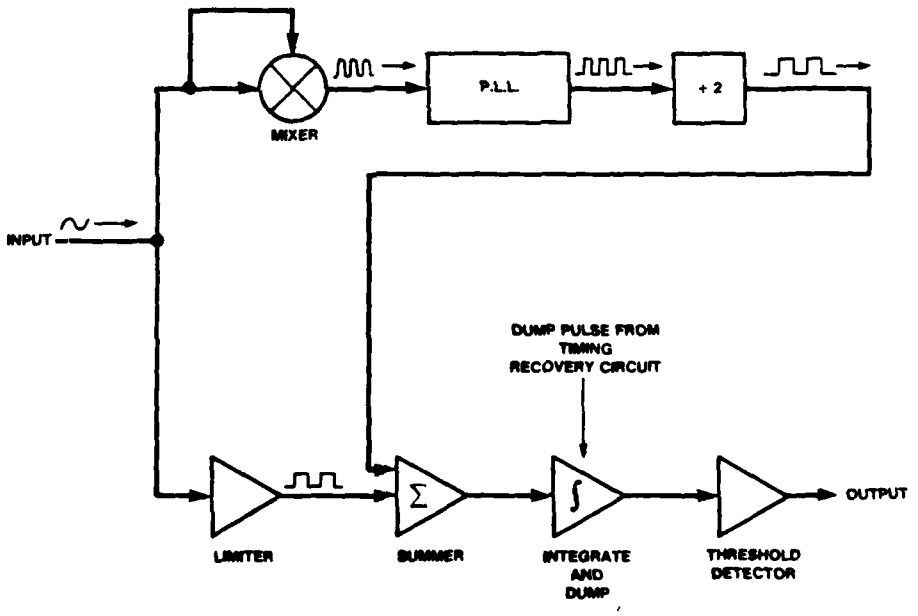

FIG 2 BASIC RELEVANT COMPONENTS OF PSK DEMODULATOR (FILTERS NOT SHOWN)

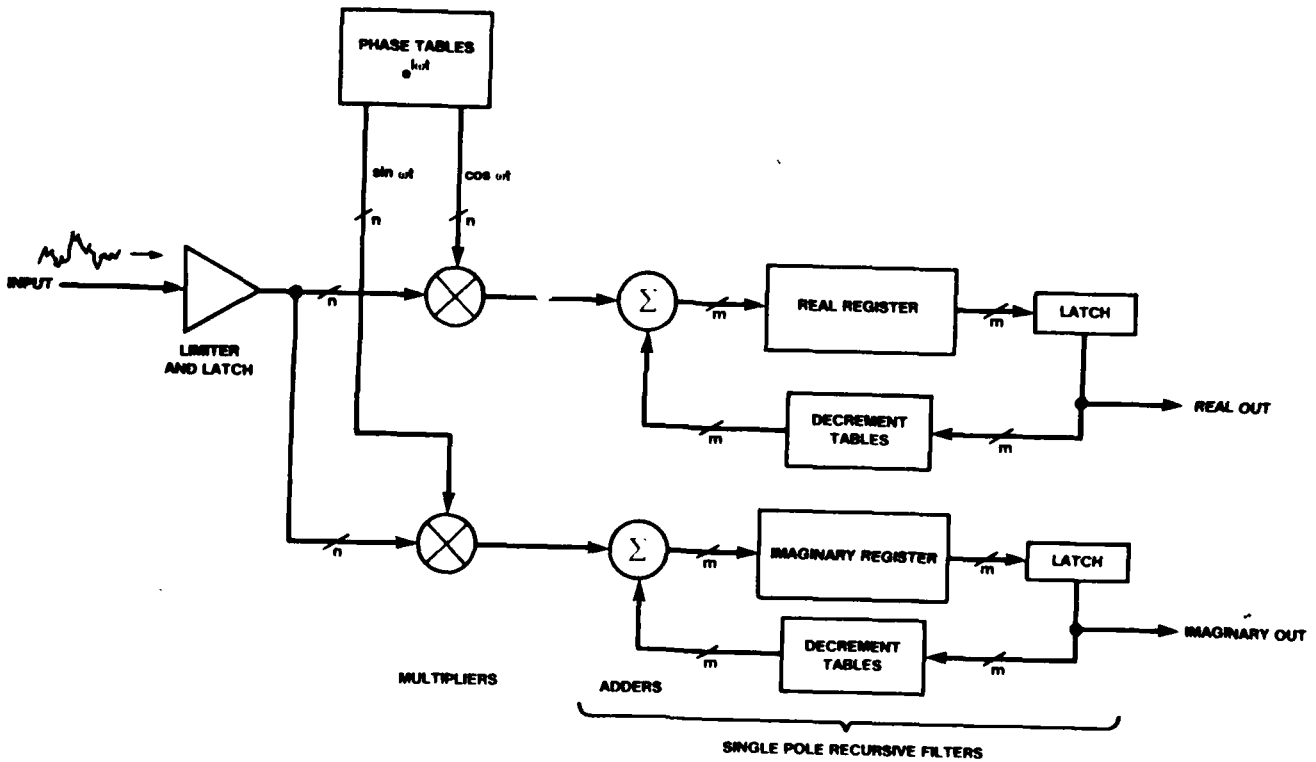

FIG 3 BLOCK DIAGRAM OF "SYNTHETIC OSCILLATOR"

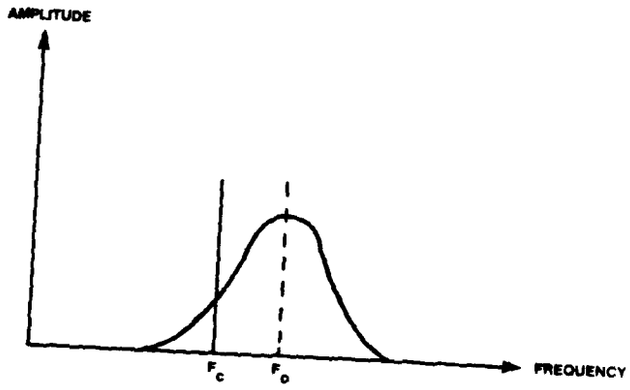

FIG 4 OSCILLATOR RESPONSE IN VICINITY OF SINE WAVE

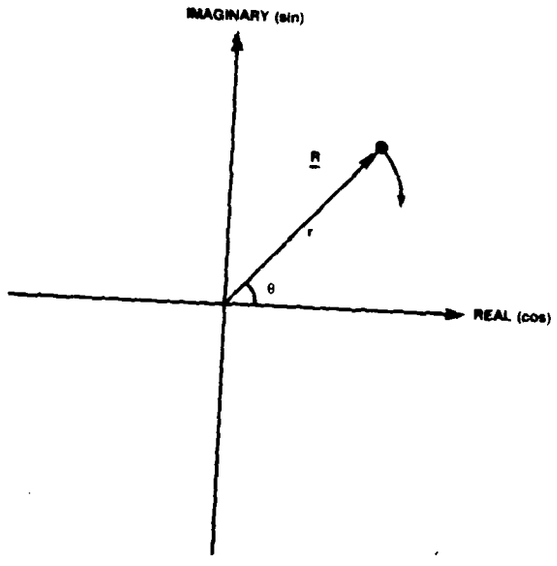

FIG 5 COMPLEX OSCILLATOR RESPONSE TO SINEWAVE OFFSET FROM CENTRE FREQUENCY

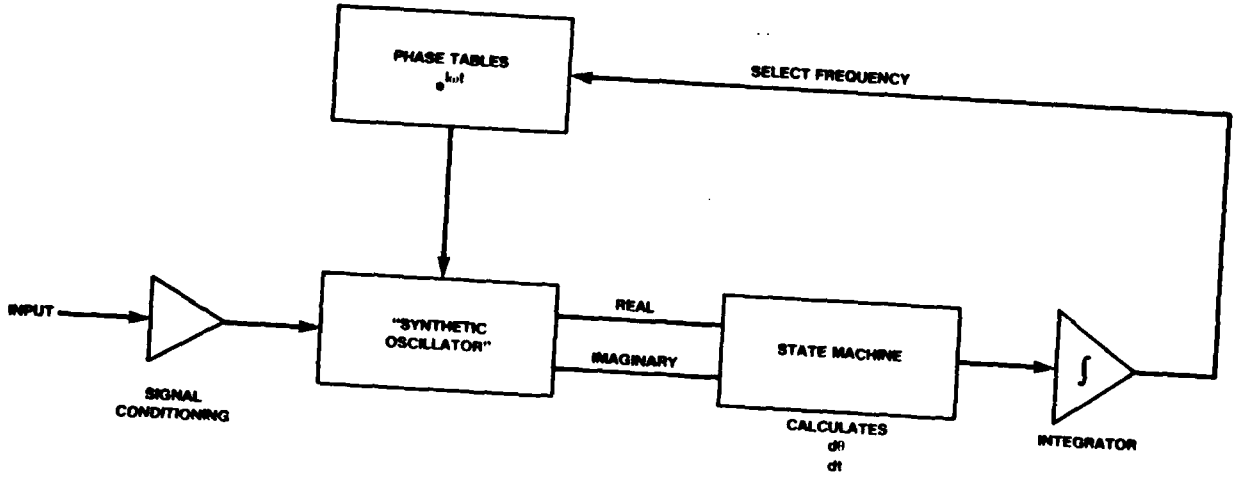

FIG 6 SIMPLE PHASE LOCK LOOP USING "SYNTHETIC OSCILLATOR"

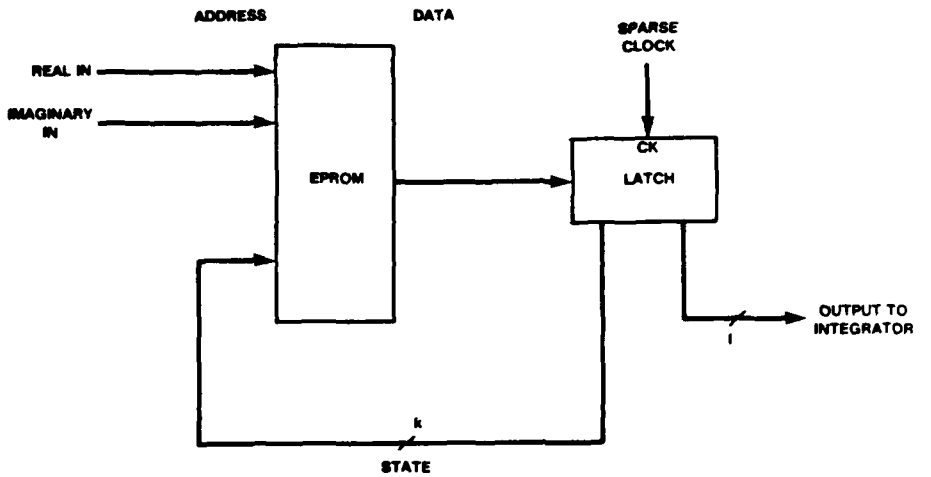

FIG 7 STATE MACHINE TO MONITOR PHASE DRIFT

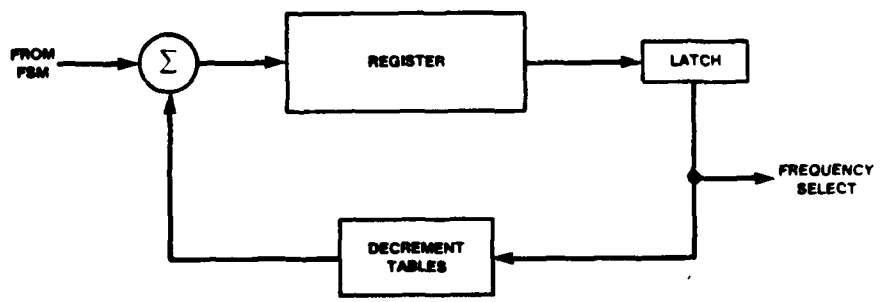

FIG 8 FREQUENCY ERROR INTEGRATOR (SPRF)

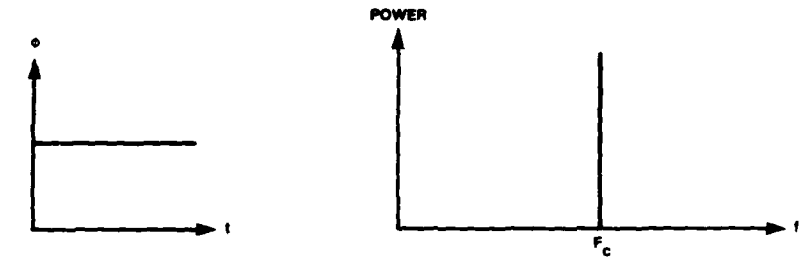

9(a) f(o) = c

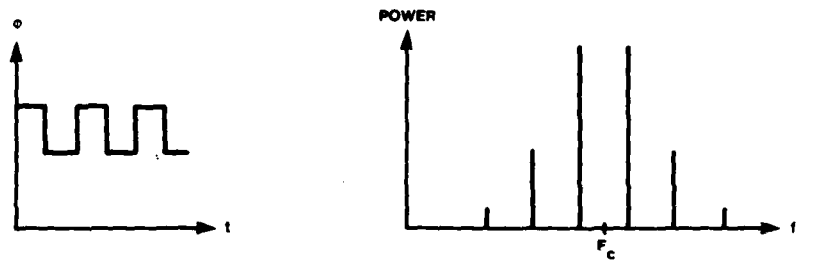

9(b) f(o) = SQUARE WAVE

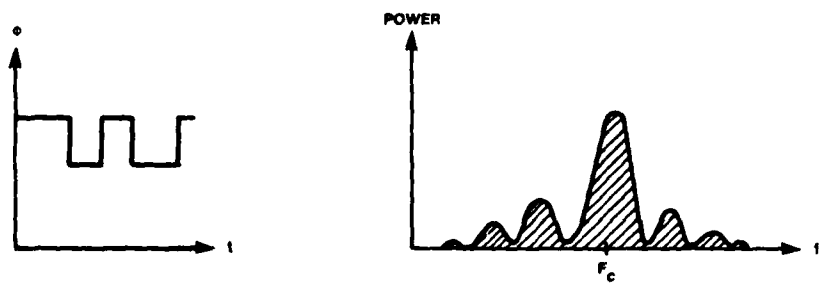

9(c) f(o) = RTS

F(o) SPECTRUM

FIG 9 POWER SPECTRA OF VARIOUS WAVEFORMS (SEE TEXT)

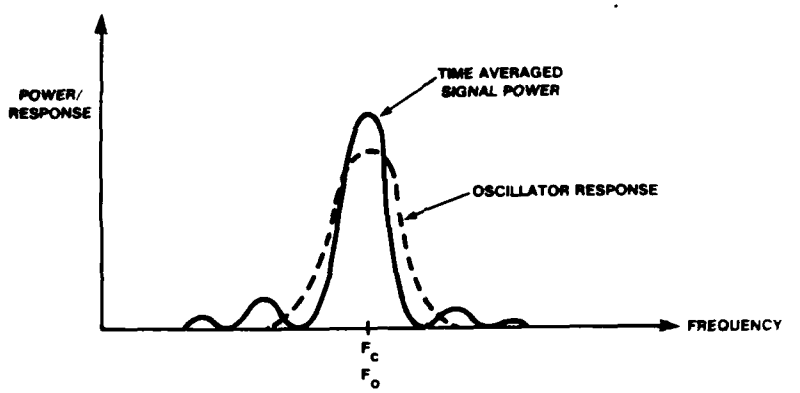

FIG 10 PICTORIAL REPRESENTATION OF OSCILLATOR DEPLOYMENT

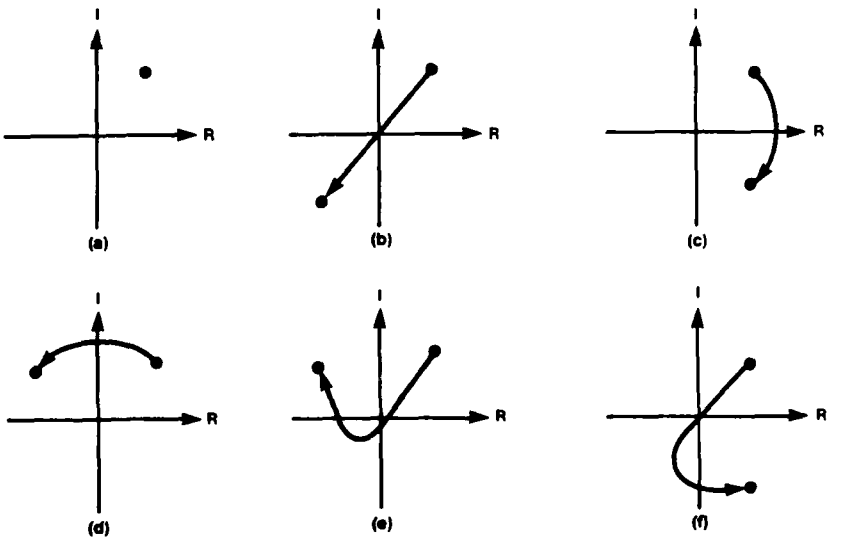

FIG 11 SINGLE OSCILLATOR TRAJECTORIES IN COMPLEX PLANE (SEE TEXT)

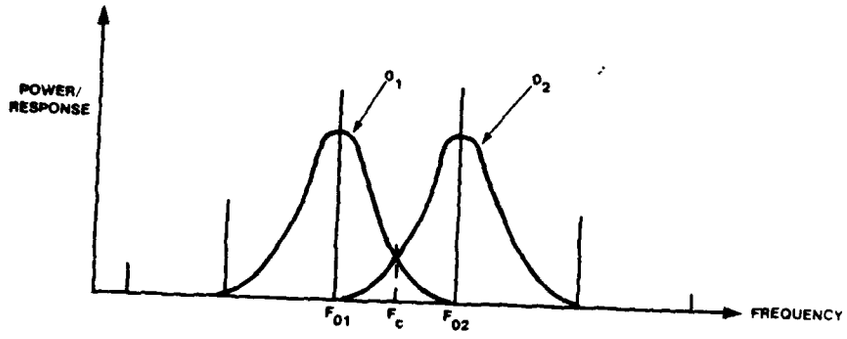

FIG 12 PICTORIAL REPRESENTATION OF OSCILLATOR PAIR DEPLOYMENT

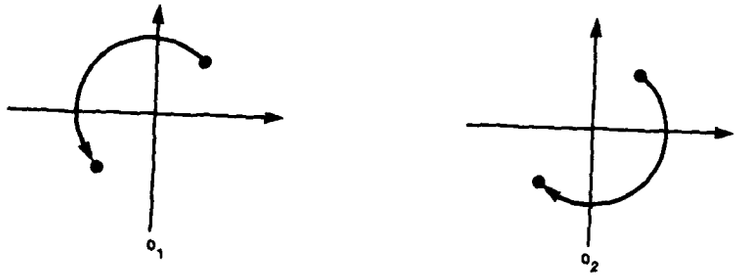

13(a) PURE CARRIER - NO FREQUENCY ERROR $\sum \Delta o = 0$

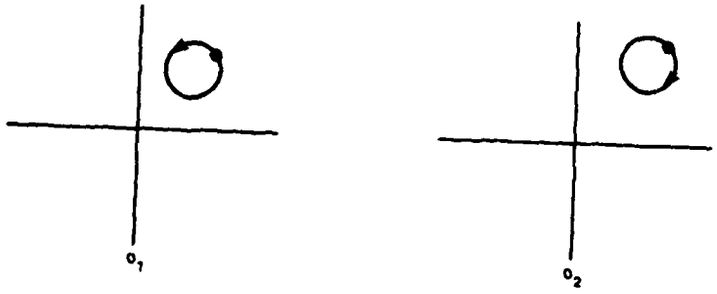

$\sum \Delta o = 0$

13(b) CARRIER WITH PHASE REVERSAL - NO FREQUENCY ERROR

O_1 O_2

FIG 13 COMPLEX MOTIONS OF OSCILLATOR PAIR

13(c) PURE CARRIER, + FREQUENCY ERROR

13(d) CARRIER + PHASE REVERSAL, + FREQUENCY ERROR

13(e) PURE CARRIER, - FREQUENCY ERROR

13(f) CARRIER + PHASE REVERSAL, - FREQUENCY ERROR

O_1 O_2

FIG 13 CONTINUED

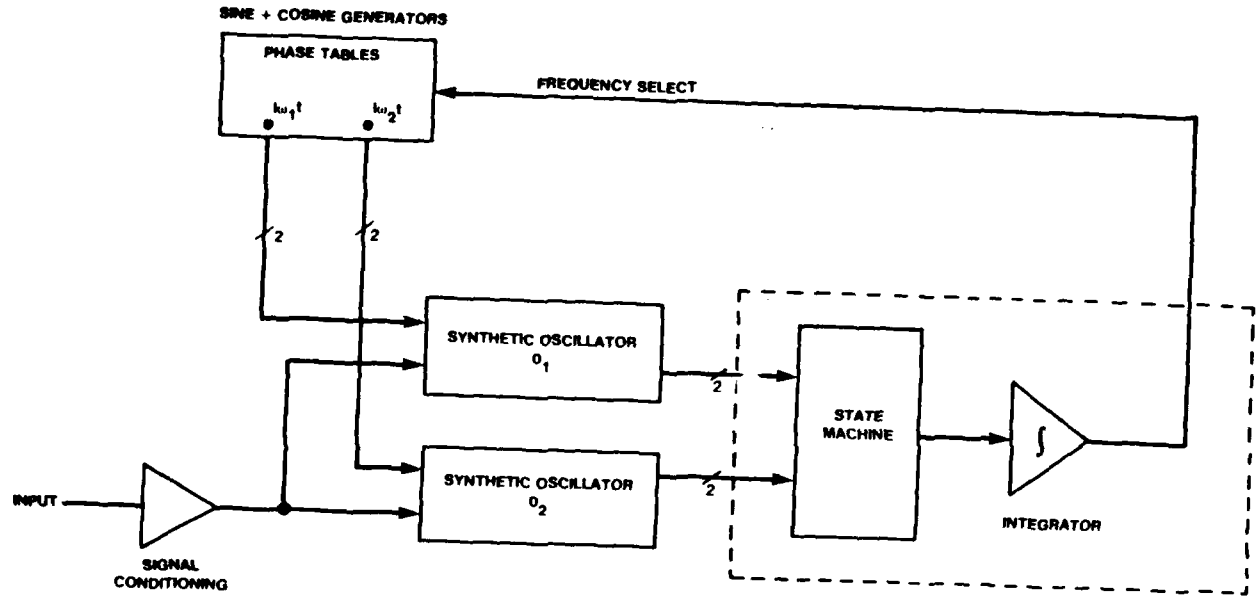

FIG 14 BLOCK DIAGRAM OF PHASE LOCK LOOP GENERALISED FOR PSK

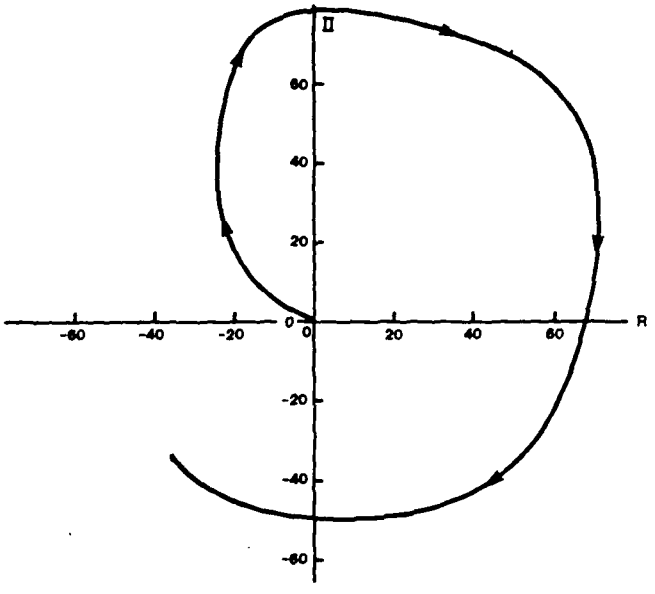

FIG 15(a) PARAMETRIC PLOT OF CONTENTS OF BIN 1 (O_1)

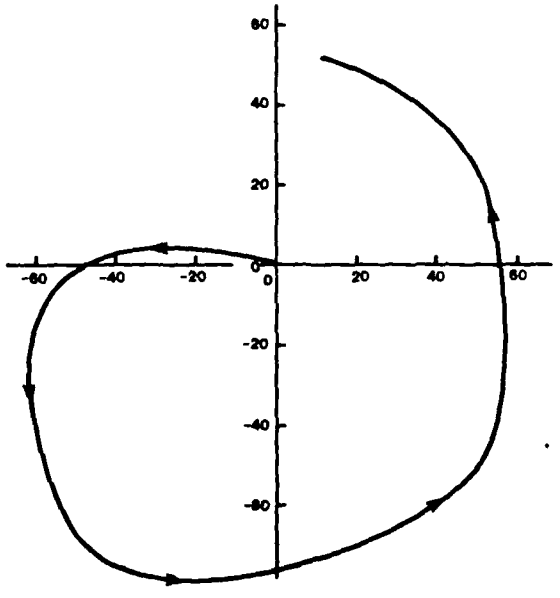

FIG 15(b) PARAMETRIC PLOT OF CONTENTS OF BIN 3 (O_2)

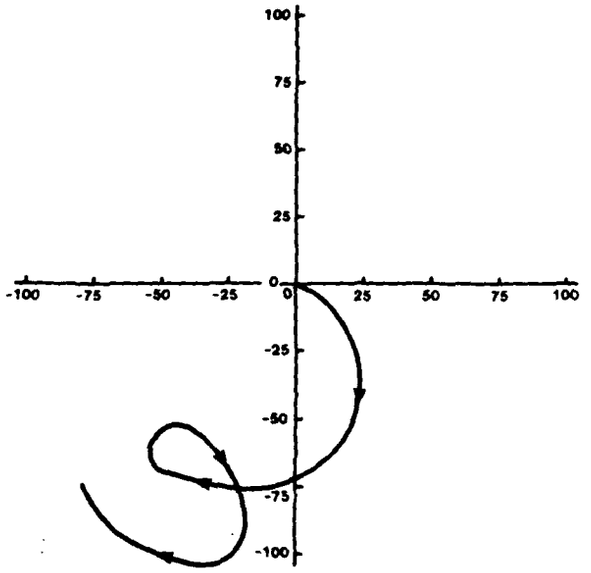

FIG 16(a) PARAMETRIC PLOT OF CONTENTS OF BIN 1 (O_1)

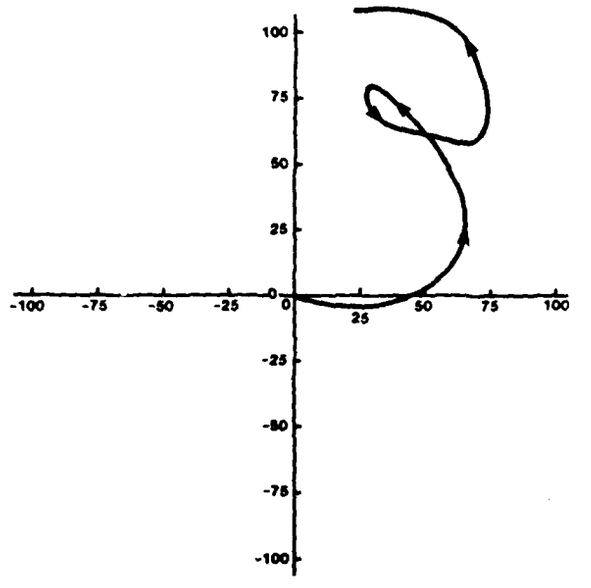

FIG 16(b) PARAMETRIC PLOT OF CONTENTS OF BIN 3 (O_2)

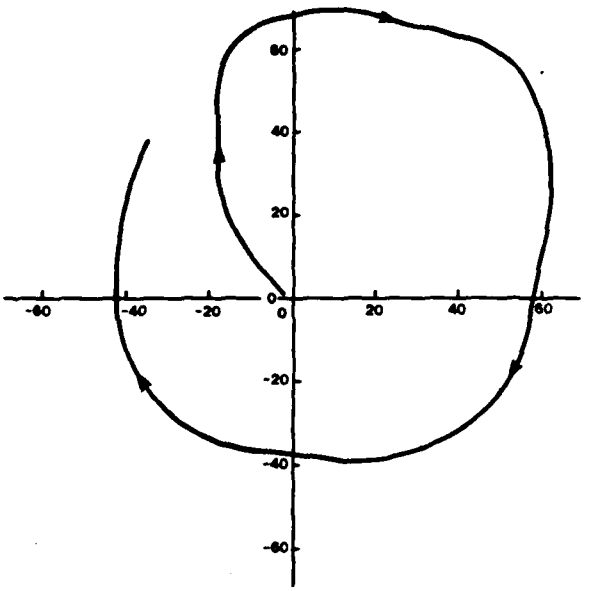

FIG 17(a) PARAMETRIC PLOT OF CONTENTS OF BIN 1 (O_1)

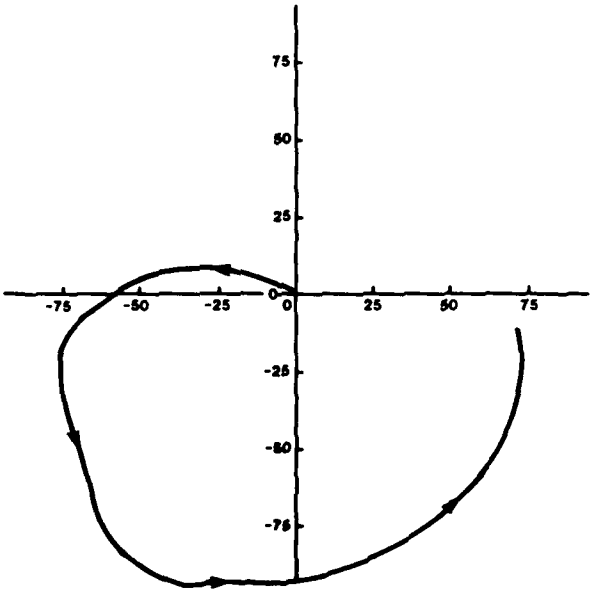

FIG 17(b) PARAMETRIC PLOT OF CONTENTS OF BIN 3 (O_2)

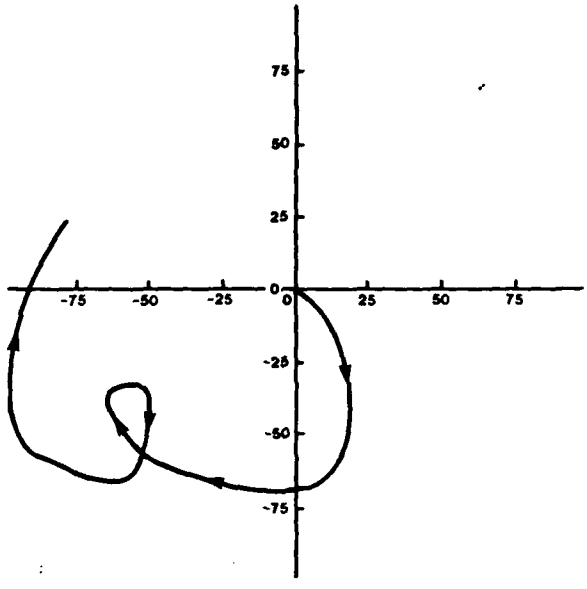

FIG 18(a) PARAMETRIC PLOT OF CONTENTS OF BIN 1 (O_1)

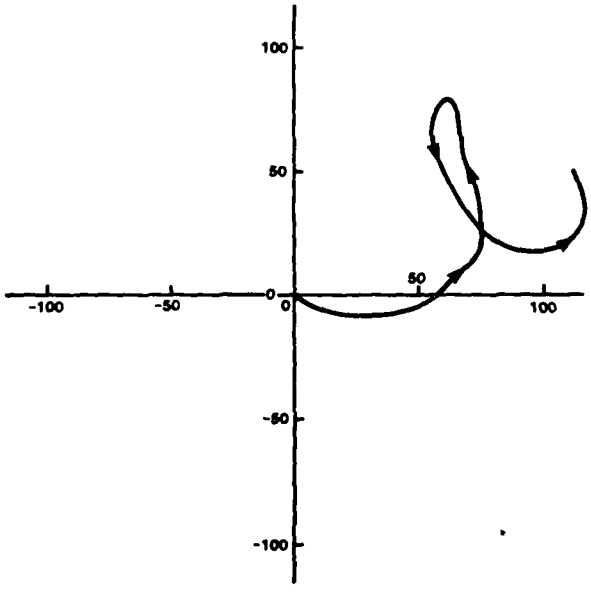

FIG 18(b) PARAMETRIC PLOT OF CONTENTS OF BIN 3 (O_2)

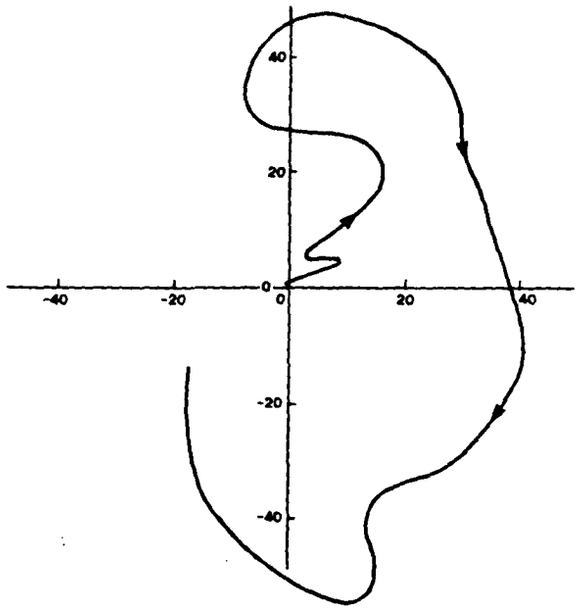

FIG 19(a) PARAMETRIC PLOT OF CONTENTS OF BIN 1 (O_1)

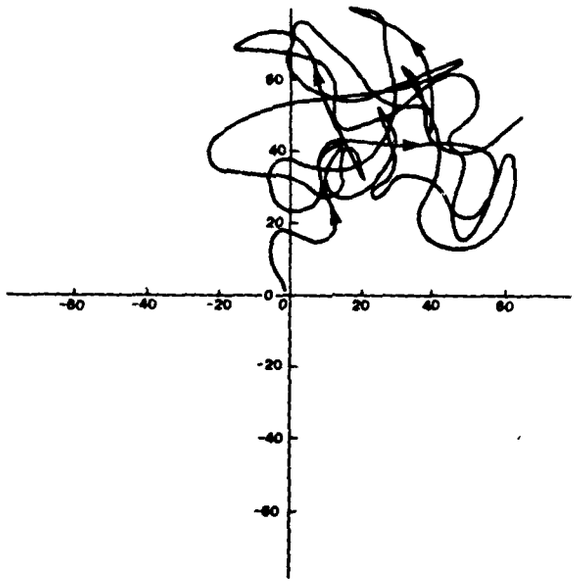

FIG 19(b) PARAMETRIC PLOT OF CONTENTS OF BIN 3 (O_2)

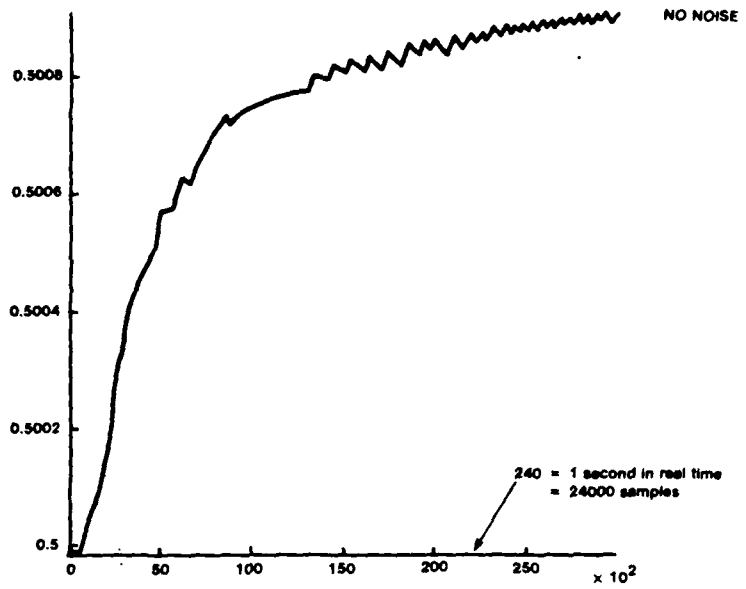

FIG 20(a) GRAPH OF CENTRE OSC. FREQUENCY

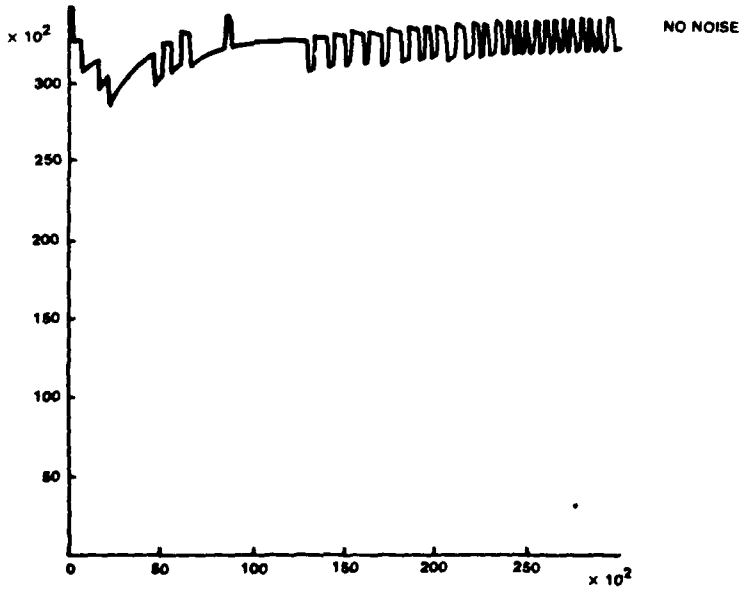

FIG 20(b) GRAPH OF TRACKING INTEGRAL

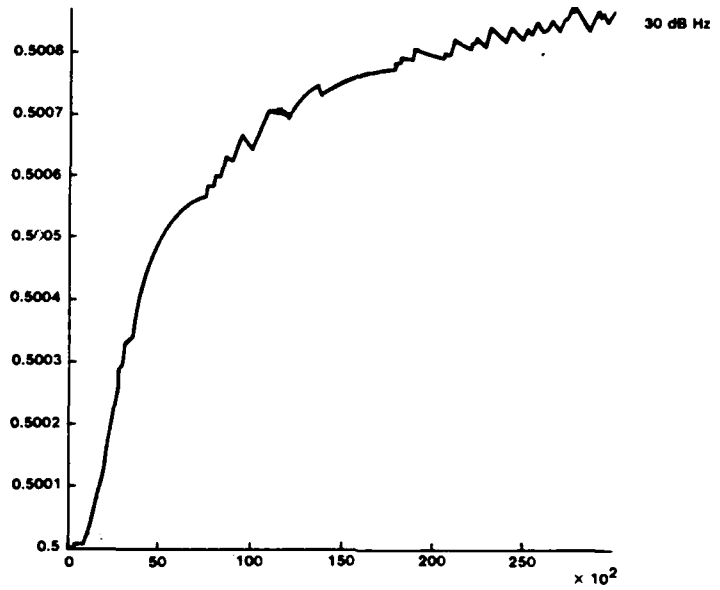

FIG 21(a) GRAPH OF CENTRE OSC. FREQUENCY

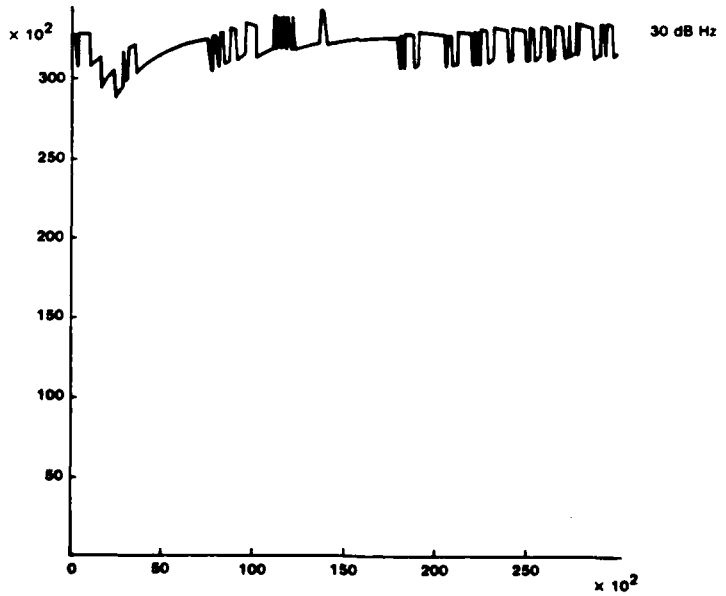

FIG 21(b) GRAPH OF TRACKING INTEGRAL

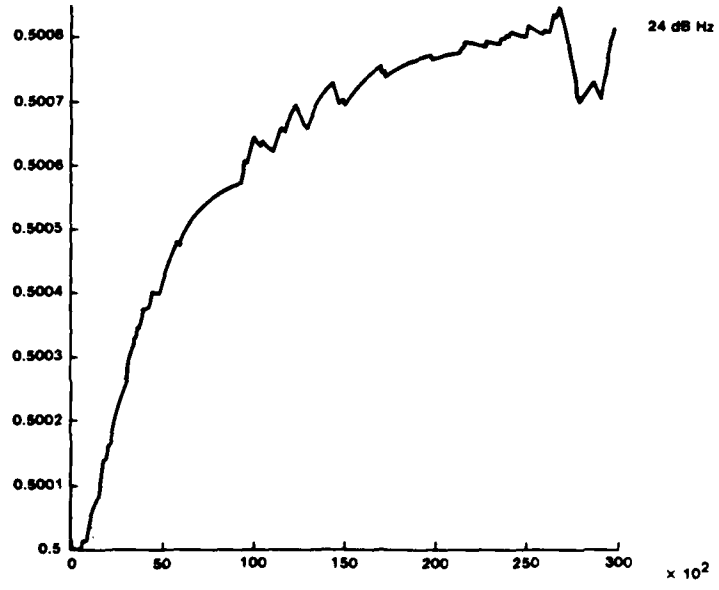

FIG 22(a) GRAPH OF CENTRE OSC. FREQUENCY

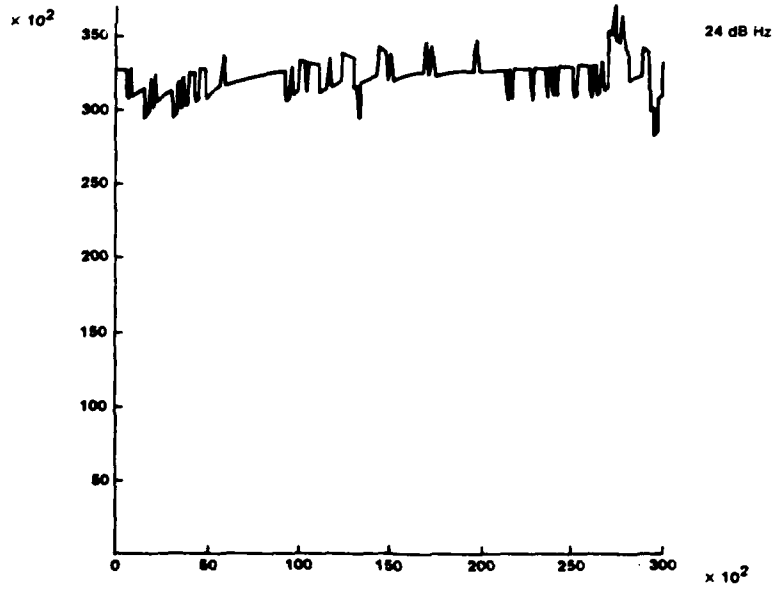

FIG 22(b) GRAPH OF TRACKING INTEGRAL

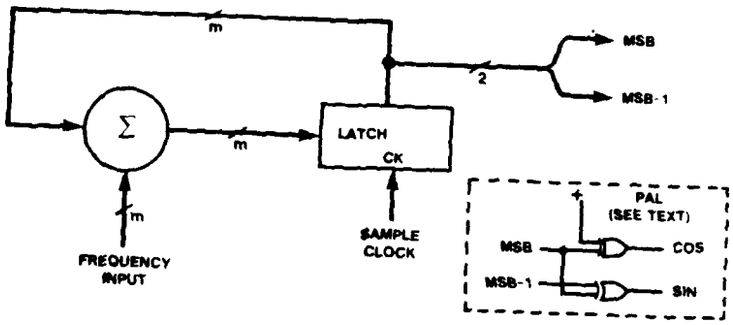

FIG 23 BASIC STRUCTURE OF SINE & COSINE GENERATORS

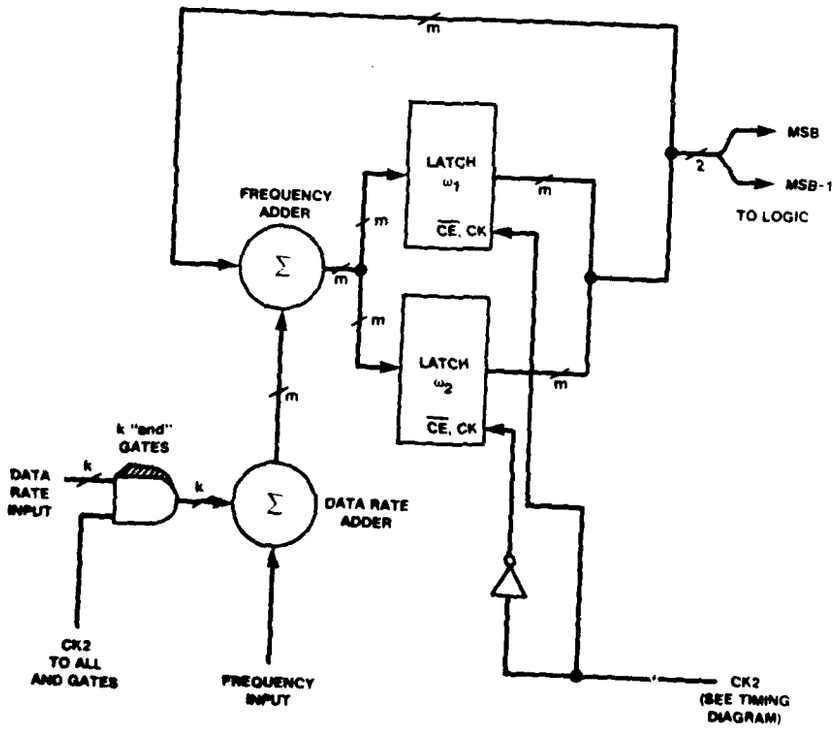

FIG 24 FULL STRUCTURE OF SINE & COSINE GENERATORS

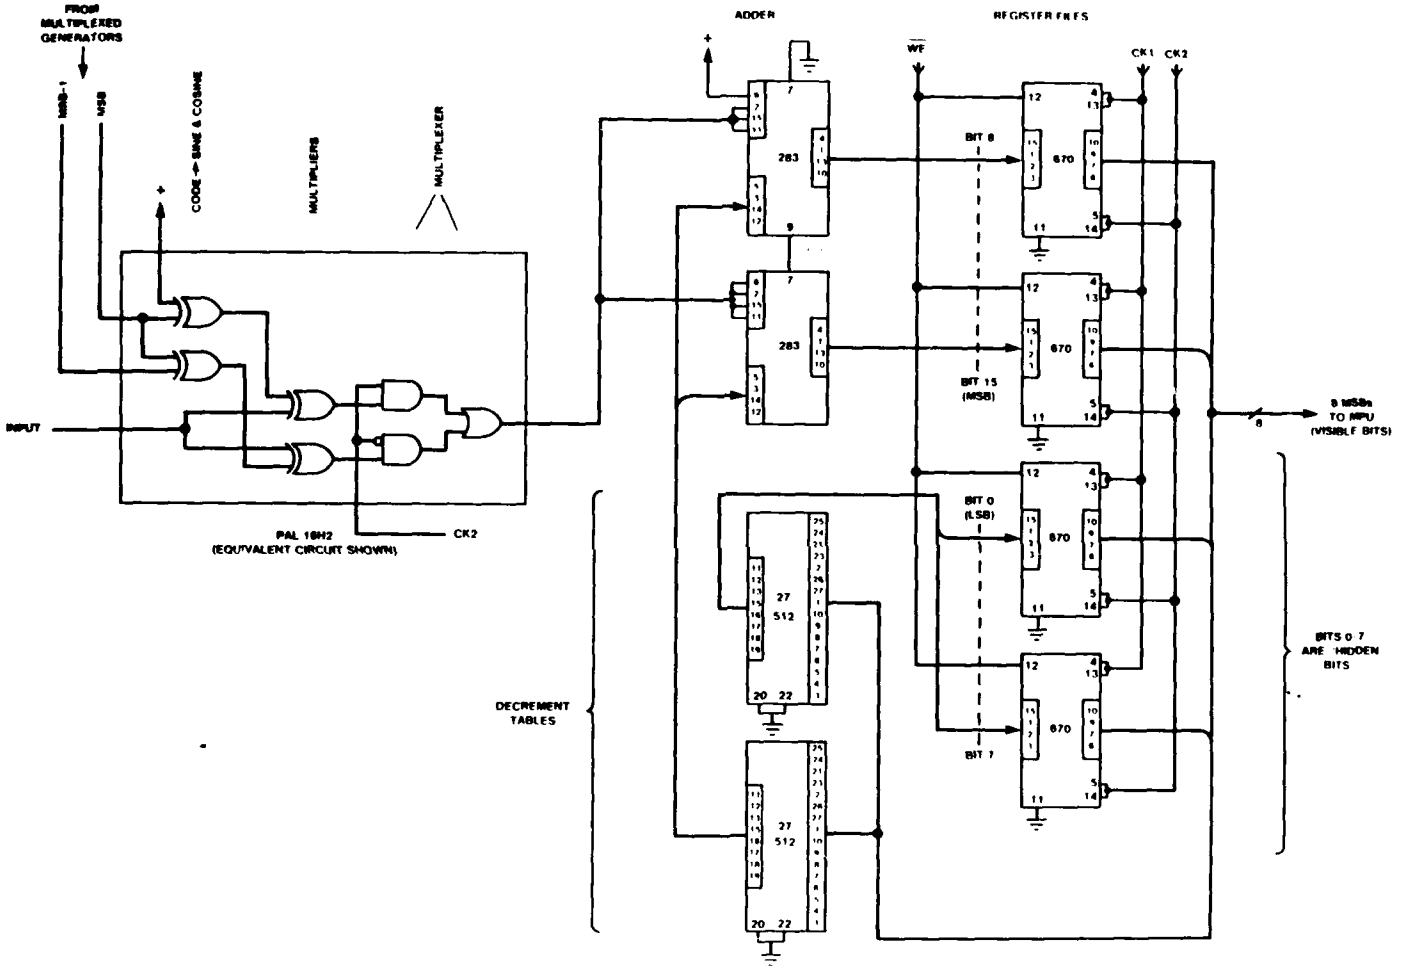

FIG 25 FULL STRUCTURE OF SPRFs AND MULTIPLIERS (POWER CONNECTIONS NOT SHOWN)

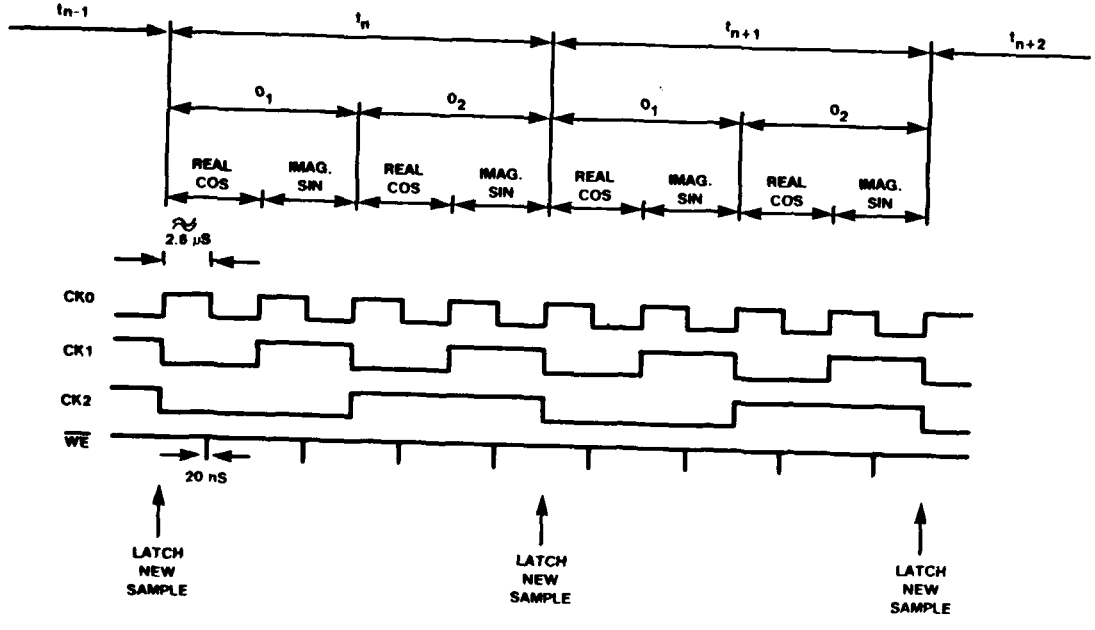

FIG 26 TIMING DIAGRAM FOR GENERALISED PLL PROTOTYPE

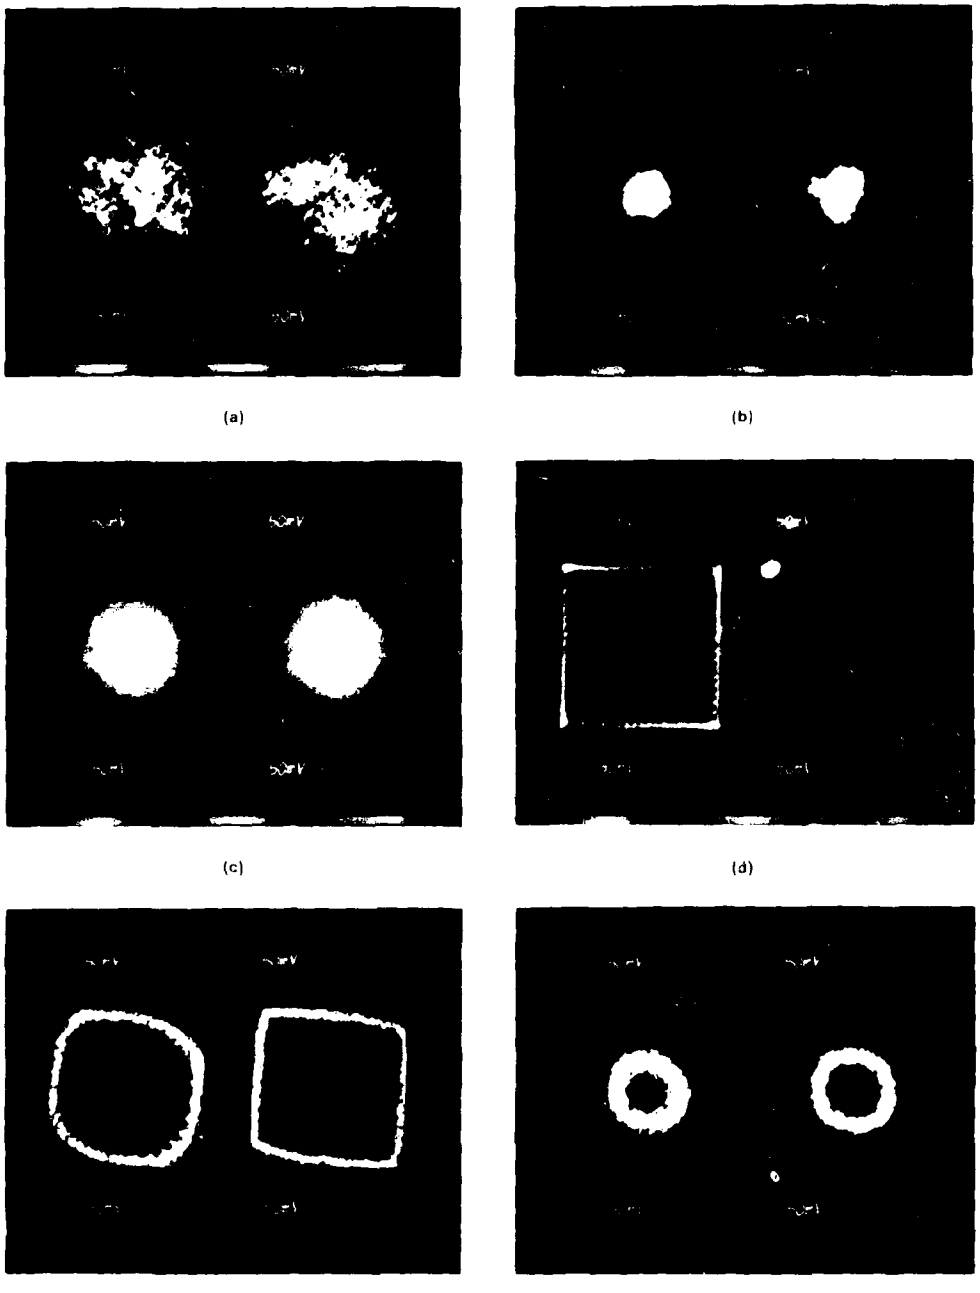

(a) (b)

(c) (d)

(e) (f)

FIG 27: BASIC RESPONSES OF PROTOTYPE

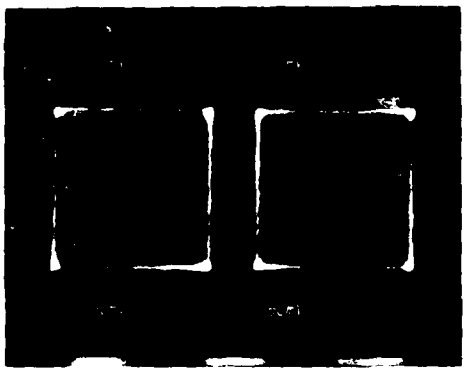

(a) S/N = ∞

(b) S/N = 40 dB Hz

(c) S/N = 30 dB Hz

(d) S/N = 25 dB Hz · 2 dB Hz

FIG 28: ARGAND DIAGRAMS WHEN PHASELOCKED (VARYING AMOUNTS OF NOISE)

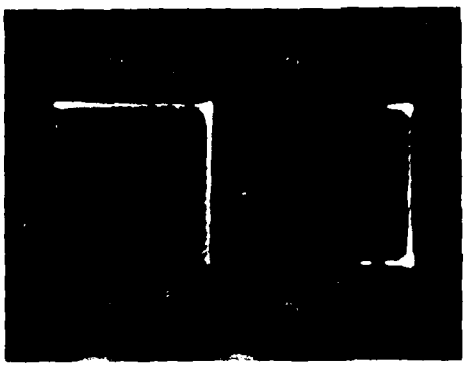

(a) S/N =

(b) S/N = 40 dB Hz

FIG 29: EFFECT OF PHASE REVERSALS ON ARGAND DIAGRAMS

(a) NOISE SPECTRUM

(b) S/N MEASUREMENTS (35 dB Hz)

FIG 30: NOISE AND S/N MEASUREMENTS

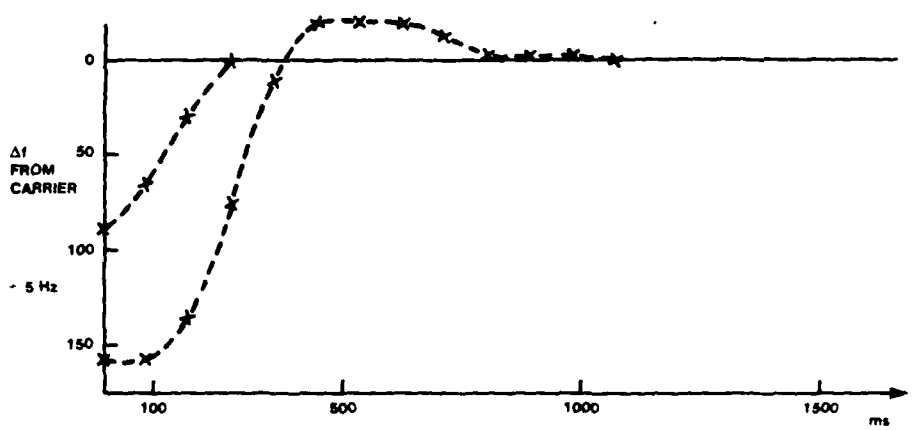

(a) PURE CARRIER

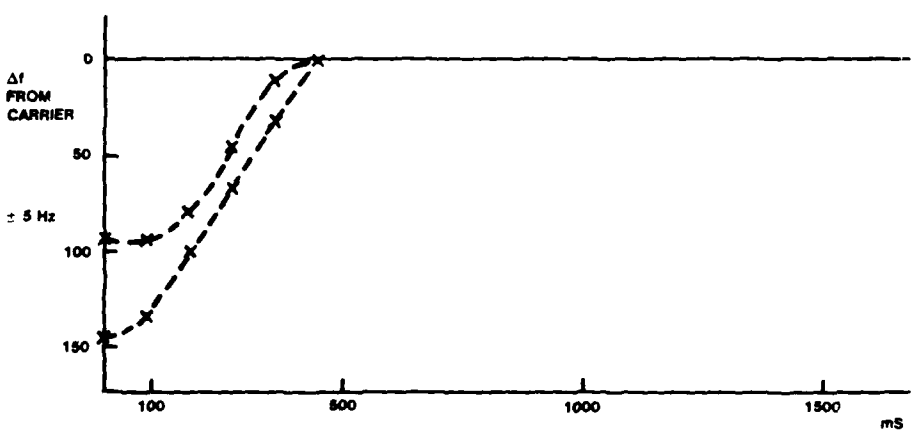

(b) WITH PHASE REVERSALS AT 50 BAUD ON CARRIER

FIG 31 PULL IN CHARACTERISTICS OF PHASE LOCK LOOP

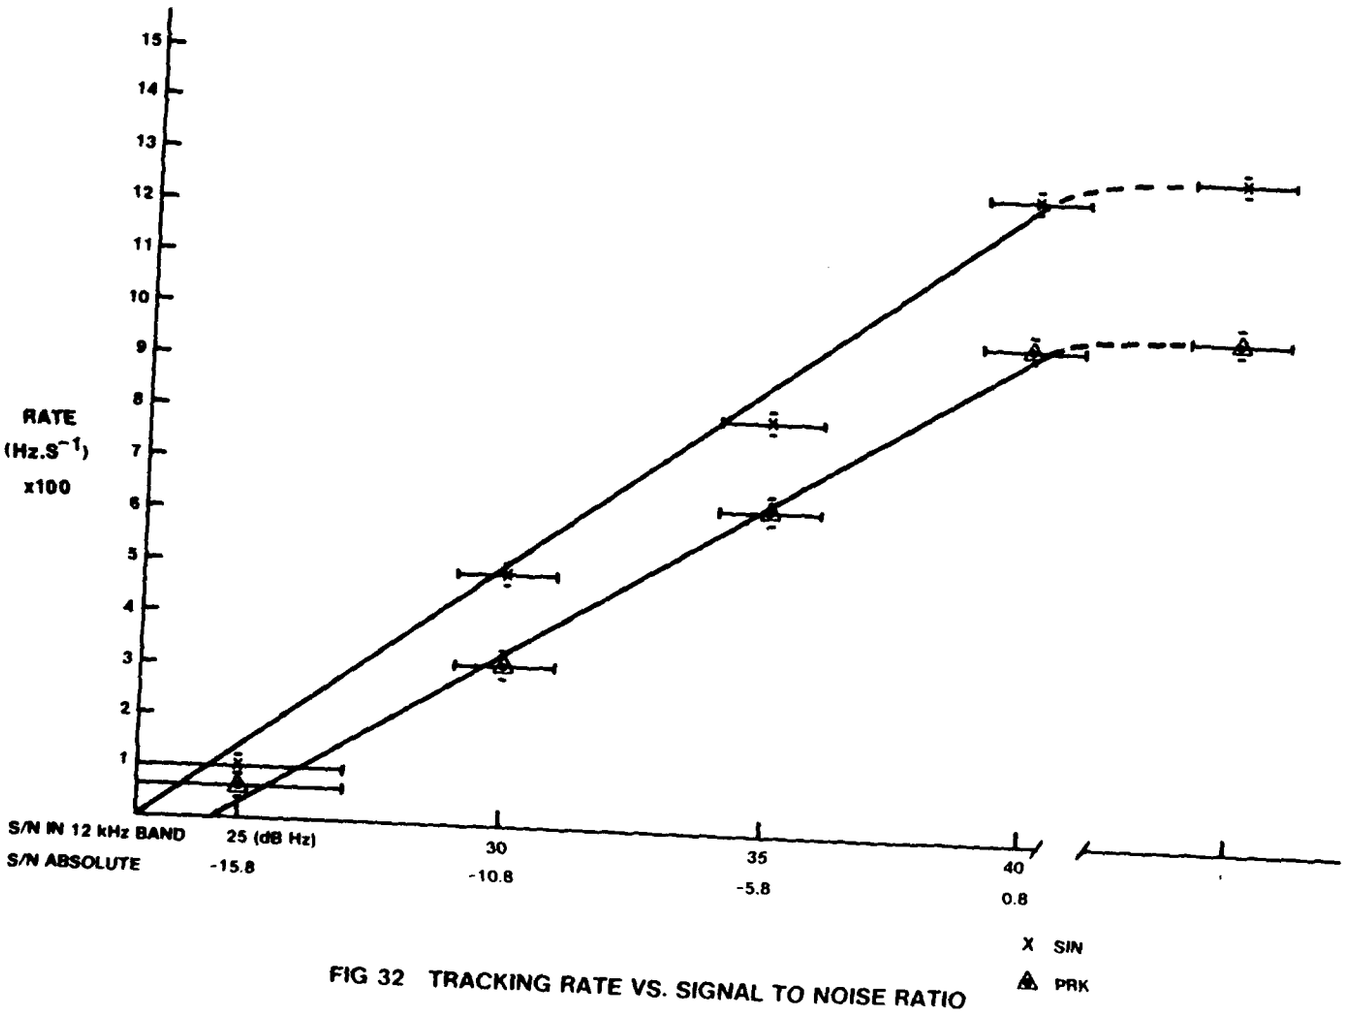

FIG 32 TRACKING RATE VS. SIGNAL TO NOISE RATIO

END

DATE
FILMED

9-87