

AD-A176 534

A COMPACT DIGITAL COMMUNICATIONS SYSTEM PART 1 RAPID
CARRIER ACQUISITION (U) ROYAL SIGNALS AND RADAR
ESTABLISHMENT MALVERN (ENGLAND) S P LITTRELL ET AL

1/1

UNCLASSIFIED


JUL 86 RSRE-MEMO-3925 DRIC-BR-100872

F/G 1772

ML



END
100
100
100



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963-A

AD-A176 534

ROYAL SIGNALS AND RADAR ESTABLISHMENT

Memorandum 3925

TITLE: A COMPACT DIGITAL COMMUNICATIONS SYSTEM
PART 1: RAPID CARRIER ACQUISITION

AUTHORS: S P Luttrell and J A S Pritchard

DATE: July 1986

SUMMARY

The problem of reducing the carrier centre frequency ambiguity in low data rate satellite communications is well known. Prior solutions to this problem have been slow, bulky or expensive. Our proposed solution does not have these drawbacks, and it is easily implemented in readily available digital hardware. This is achieved by using a 1 bit digital emulation of a bank of forced LCR oscillators to achieve an approximate Fourier decomposition.

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Avail and/or	
Dist	Special
A-1	



Copyright
C
Controller HMSO London
1986

RSRE MEMORANDUM NO 3925

A COMPACT DIGITAL COMMUNICATIONS SYSTEM
PART 1: RAPID CARRIER ACQUISITION

S P Luttrell and J A S Pritchard

LIST OF CONTENTS

- 1 Introduction
 - 1.1 Statement of the problem
 - 1.2 Prior art and its deficiencies
 - 1.3 Proposed solution
- 2 Rationale
 - 2.1 Introduction
 - 2.2 The signal structure
 - 2.3 Further constraints on the signal processing scheme
 - 2.4 A physical analogy
 - 2.5 Hardware emulation of a forced LCR oscillator
- 3 Hardware
 - 3.1 Introduction
 - 3.2 Detailed description of main circuit operation
 - 3.3 Structure of the phase tables
 - 3.4 Reading the registers with a microprocessor
 - 3.5 Register structure and size
 - 3.6 Phase and decrement table generation and oscillator emulation
- 4 Results and Discussion
 - 4.1 Introduction
 - 4.2 Computer simulation of the circuit
 - 4.3 Response of hardware to a sine wave
 - 4.4 Response of hardware to a sine wave in noise
 - 4.5 The effect of phase reversals of the noisy carrier
- 5 Conclusions

Acknowledgement
References

LIST OF APPENDICES

- 1 Derivation of Processor Structure using General Principles
- 2 Some Useful Mathematical Results
- 3 Response to a Noisy Sine Wave and Saturation Effects

LIST OF CONTENTS (Continued)

LIST OF FIGURES

- 1 Swept filter and integrator
- 2 Block diagram of rapid acquisition circuit
- 3 General register structure
- 4 Method of generating sine & cosine functions automatically
- 5 Computer simulation of power spectrum generated (sine wave at 30dBHz)
- 6 Same as Fig 5 but with hard phase reversals every 20mS
- 7 Response of oscillator to varying frequency, no noise - Argand diagrams
- 8 Response of oscillator to varying noise - Argand diagrams
- 9 Response of oscillator to 50 hard phase reversals per sec
- 10 Theoretical predictions (see text)

1 INTRODUCTION

1.1 STATEMENT OF THE PROBLEM

The object of a rapid carrier acquisition circuit is to reduce the centre frequency ambiguity of a modulated carrier. In conjunction with such a circuit, a phase locked loop may then achieve lock with the carrier in a timescale which is much smaller than that possible with a swept filter and integrator combination. This type of system is most useful in circumstances where:

- a. The signal has its energy concentrated in a small spectral region within a much larger band of ambiguity.
- b. The frequency drift of the signal is not so great as to introduce significant ambiguity into the interpretation of the data.
- c. The signal to noise ratio is close to limiting the error free data rate.

Previous realisations of such systems, although practicable, have one or more undesirable features such as high cost, complexity or bulk. In any case they are generally unsatisfactory in terms of the reduction of acquisition time. We have eliminated these undesirable features¹.

A typical and important application of such a system is in a narrow satellite communications channel where low data rate transmissions are used². The signals are often phase modulated, and there are centre frequency ambiguities introduced by local oscillator error and doppler shift due to satellite drift. In this Memorandum we consider this as a particular example of an implementation of the rapid acquisition system.

1.2 PRIOR ART AND ITS DEFICIENCIES

The original method of determining the centre frequency of a low data rate signal was to use what was effectively a narrow band swept filter and an integrator (see Figure 1). This was realised as a fixed filter feeding a rectifier and the integrator, the input signal being mixed with a swept voltage controlled oscillator. The output was tested with a threshold detector. The technique is a slow serial search and has the added disadvantage in practice of locking to the first signal encountered which is of sufficient power to exceed the set threshold. This would not necessarily be the largest or the required signal, and for example it may have been produced via nonlinearities in a satellite transponder. Typical serial search times for a 50 bit per second signal in a 12 kHz bandwidth at a signal to noise ratio of 30 dB Hz is between 0 and 40 seconds depending on the location of the signal in the band.

The latest method implemented in practice uses state-of-the-art chirp techniques. There are several successive frequency estimates made to reduce the frequency ambiguity to within the phase locked loop capture range, the first and widest search uses a chirp correlator connected to a VCO/mixer combination as before. This allows a rapid initial search. The system reduces the total acquisition time to 7 seconds regardless of the centre frequency of the signal and has the added advantage of choosing the largest signal as the carrier. The solution is complicated, and the 1984 price for the system excluding the modem was £12,000.

Other methods which have been proposed are the analogue bank of filters method and the fast Fourier Transform (FFT) method. The analogue filter bank is an attractive alternative as it is a parallel system, but has been rejected in the past as it is complex, bulky and expensive. The FFT method requires much computing power and it is expensive. Although it is the more compact and efficient of the two methods, it is not compact enough to install in small, remote terminals.

1.3 PROPOSED SOLUTION

There are two key elements which point to the most efficient solution and its hardware realisation. These are:

a. We consider that a system is more likely to be optimum if it emulates the physics of the behaviour of the object under study. The determination of a frequency implies that the system would have "eigenstates" which are analogues of different narrow frequency bands.

b. If we are to employ digital hardware most efficiently and the bandwidth of the digital hardware is much greater than the width of the satellite band under consideration, we would then wish to use the digital bandwidth to develop the physical analogue in parallel.

A simple physical system which is responsive in a precisely definable manner to a particular band of frequencies is the LCR oscillator. We shall discuss the rationale for using this model and its particular implementation in hardware in the next section.

2 RATIONALE

2.1 INTRODUCTION

The rationale in the design of any good signal processing hardware derives from one's prior knowledge of the signal structure. Broadly speaking, the complexity of the design reflects the richness of structure present in the signal. Obviously other considerations must enter into the design, and typically these involve trading off the complexity of the hardware against the quality of the signal analysis being performed. In our application these additional considerations will become clear in the following sections.

2.2 THE SIGNAL STRUCTURE

In order to specify our prior knowledge of the signal structure we must examine how the signal was generated, and the type of noise processes that are present. The signal consists of components with different characteristic timescales:

(S1) A sinusoidal carrier with period τ_s

(S2) A code chip period τ_c

(S3) A carrier and/or code drift time τ_d

which we shall assume satisfy the inequality

$$\tau_s \ll \tau_c \ll \tau_d \quad (2.1)$$

The fact that these characteristic timescales are well separated enables us to factorise the signal processing into three stages:

(P1) Hardware to lock on to the sinusoidal carrier S1

(P2) Software to demodulate the signal encoded in S2

(P3) Software to track (ie retain lock with) S1 and/or S2

We shall be concerned only with P1 in this Memorandum; P2 and P3 will be dealt with in future reports.

In order to design P1 we must introduce prior knowledge of the corresponding signal structure. This prior knowledge is:

(PK1) S1 is a pure sine wave of frequency ω , plus

(PK2) additive white gaussian noise

where PK2 models the known properties of the noise processes in satellite communication channels. Furthermore we may impose a higher level of prior knowledge which expresses the fact that the carrier frequency ω is to be found in a limited band of frequencies

(PK3) $0 \leq \omega \leq \omega_{\max}$

PK1-3 completely specify our prior knowledge of S1 plus its associated noise; we shall use this information together with some additional constraints to design P1.

2.3 FURTHER CONSTRAINTS ON THE SIGNAL PROCESSING SCHEME

The optimal linear processing scheme for threshold detection of S1 corresponding to PK1 and PK2 is the Fourier Transform (FT). A FT is essentially a set of filters that are matched to an equivalent set of pure sine wave signals. The signal to noise ratio (SNR) is maximised in the FT channel which precisely matches the frequency of the incoming carrier. Thus we could perform P1 by seeking the FT channel with the maximum response. In order to achieve this aim we must introduce some additional prior knowledge in order to constrain the precise means whereby the FT is realised.

An obvious constraint that we must place on any FT method is:

(C1) The FT "window" must have a length of order τ_c .

This ensures that the width of the FT channels is of order of the modulation superimposed on the carrier, which in turn ensures that the energy contained in the carrier and its modulation (ie S1 with S2) does not wander outside the range of the FT channel that it happens to fall in originally. It also ensures that each FT channel is as narrow as it could be subject to this condition. Another obvious constraint is:

(C2) The set of FT channels must have substantially non-overlapping windows.

This will ensure that S1 and S2 will only fall into one or two FT channels simultaneously. The final constraint on the general properties of the FT channels is:

(C3) The entire frequency range $0 \leq \omega \leq \omega_{\max}$ must be "covered" by FT channels.

Together the constraints C1-3 specify the type of FT (up to factors of order of unity) which is required to perform P1. However it is by no means obvious that a standard FT algorithm will be the optimal solution within our constraints; we have only needed to specify the properties of the transformation up to factors of order of unity. This is evident in the use of loose language in defining the constraints C1-3 above. We shall now capitalise on this residual freedom in order to design a P1 which looks similar to a standard FT, but which possesses certain features that permit a simple hardware realisation.

2.4 A PHYSICAL ANALOGY

We strongly believe that physical intuition is a powerful means of seeking short-cut solutions to problems, therefore we shall seek a physical analogue of the processing operations that P1 must perform. Then we shall design processing hardware that emulates as simply as possible the behaviour of the physical analogue. This type of approach may not be desirable to everybody, so we include a formal derivation of the equivalent mathematical representation in Appendix 1.

The FT operation arrived at after imposing the set of constraints C1-3 merely emulates a bank of filters. We shall therefore introduce a physical model where each filter is an LCR oscillator. The centre frequency and Q-factor of each LCR oscillator is adjusted to give the desired FT channel response. The incoming signal then corresponds to an external source which stimulates the LCR oscillator to respond accordingly.

The response of a particular FT channel then corresponds to the response of the equivalent LCR oscillator. We are not interested here in the phase of the response, so we shall use the rate of energy dissipation in the resistor R as a measure of the response. This will fluctuate in a manner which depends upon the centre frequency and Q-factor of the LCR oscillator; we may average out these fluctuations by integrating the energy dissipated over a period of time. Let us denote the choice of an LCR oscillator physical analogy as constraint number 4:

(C4) Choose that particular FT realisation which corresponds to the response of a forced LCR oscillator.

2.5 HARDWARE EMULATION OF A FORCED LCR OSCILLATOR

The following results are derived in Appendix 1. We shall assume throughout the following analysis that the raw signal has been translated to baseband, that it has been band limited to remove out-of-band noise, and that it has been sampled at no less than the Nyquist rate. The in-phase and in-quadrature amplitude responses ($A_{I,k}(n)$ and $A_{Q,k}(n)$) of forced oscillator k at sample n have the form

$$\begin{aligned} A_{I,k}(n) &= \sum_{r=0}^{\infty} \alpha^r \cos[(n-r)\phi_k] x_{n-r} \\ A_{Q,k}(n) &= \sum_{r=0}^{\infty} \alpha^r \sin[(n-r)\phi_k] x_{n-r} \end{aligned} \quad (2.2)$$

where x_n is sample n of the forcing term (ie incoming signal), α is the decrement (or decay) factor of the oscillator, and ϕ_k is the phase advance per sample of the oscillator. The squared response is given by

$$E_k(n) = [A_{I,k}(n)]^2 + [A_{Q,k}(n)]^2 \quad (2.3)$$

The goal of P1 is to emulate digitally the behaviour of a set of oscillator amplitude responses, and to make suitable measurements on the corresponding $E_k(n)$ in order to lock on to S1.

The first hardware consideration is what number of bits to retain from each signal sample x_n . This question can be addressed by considering the relative noise amplitudes due to Gaussian noise on the communication channel, and noise associated with the quantisation of the x_n . We must ensure that the quantisation noise never dominates the Gaussian channel noise if we are substantially to retain the signal information that is present. For our purposes 1 bit quantisation accuracy is sufficient to ensure that these conditions are met in almost all signal and noise scenarios. The consequent simplicity of the hardware design is highly desirable.

An immediate consequence of the decision to use 1 bit signal quantisation is that we need only represent the amplitudes of the $\cos(n\phi_k)$ and $\sin(n\phi_k)$ terms to 1 bit accuracy. This approximation may not be extended to the decrement factor α because then only two α values (say 0 and 1) could be represented, which would be an unacceptable constraint on the form of the LCR oscillator which could be emulated.

Let us represent these two constraints as:

(C5) Quantise the signal to 1 bit accuracy.

(C6) Quantise the $\cos(n\phi_k)$ and $\sin(\phi_k)$ terms to 1 bit accuracy. These representations are stored as "phase look-up tables".

Further constraints on the detailed form of the hardware will be defined to the next section, which deals in depth with the various hardware compromises that must be made in order to arrive at an economical design.

3 HARDWARE

3.1 INTRODUCTION

The hardware implementation will be discussed in the specific context of the application for which it was originally designed. This was the narrowing of the frequency ambiguity of a 50 baud DPSK telegraph signal in a 12 kHz bandwidth. The width of the relevant channels was 10 kHz, but some margin is always allowed to ensure overlap of adjacent channels. The signal to noise ratios commonly encountered in such channels are often enough to cause errors, and the frequency of the carrier is correspondingly difficult to determine. The signal to noise ratio over the whole band is in practice often as poor as -15 dB.

The basic layout of the system is shown in Figure 2. The non-oscillatory part of the dynamics of each LCR oscillator is realised as a pair of single pole recursive filters (SPRFs). Each pair consists of a real part and an imaginary part, which respond to "cosinusoids" and "sinusoids" respectively. Since the digital bandwidth is chosen to be a factor of 128 greater than the input sampling rate, the same circuitry can service 128 oscillators, the state of each of which is recorded in a register in the recursive filter loop. These memories are arranged so that any particular one may be interrogated by the microprocessor system which is to perform the necessary squaring, summing and accumulating functions which are needed to obtain a power spectrum. Alternatively the "sum of squares" operation may be performed by a read only memory (eg AMD 27512) configured as a look-up table. Computer simulation has shown that although there is some small penalty in acquisition time, less than 8 bits are needed from each register.

The input signal is multiplied by the 128 different sets of terms in the phase tables which correspond to the centre frequencies of each LCR analogue. Thus the phase tables are chosen for each "oscillator" in such a way that an input frequency corresponding to the centre frequency of a particular oscillator is converted to dc at the input to the corresponding SPRF.

The criteria (C5) and (C6) discussed in the rationale specify that the signal and the phase tables be quantised to one bit accuracy. In principle this means that only the sign bit is retained, and the multiplier is reduced to an exclusive NOR gate. Moreover, if the exclusive NOR is replaced by an exclusive OR, the effect is to invert the complex planes represented by the contents of the registers through the origin. This will not affect the operation of the rapid lock circuit since we are interested at this stage only in signal powers. Note that the exclusive NOR and the adder may be absorbed as two extra address lines in the decrement tables. Although the contents of the "decrement tables" becomes more complicated, this will reduce the component count.

The input is sampled at the Nyquist frequency of the satellite band and held by the latch shown in Figure 2. Since the system is designed to operate at baseband and the total bandwidth is 12 kHz, a new sample is taken every $1/24000$ sec, or every 41.67 microseconds. In this period, 128 SPRFs must be serviced using their individual internal states, the input sample, and their individual phase table values for that particular sample time to define a new state. This implies a system clock of frequency $24000 \times 128 = 3.072$ MHz. This is easily realisable using a "lock step" system, TTL, and fast but readily obtainable memories.

128 "parallel" LCR circuits were chosen for the following reasons:

- a. 2^N (N positive integer) was a highly convenient number of frequency channels to use from a hardware point of view. No special considerations were needed for counters and registers for example.
- b. 128 oscillators were the maximum number that could easily be implemented in hardware without running into problems with speed and (from the end user's point of view) power consumption.
- c. 128 oscillators satisfies the criteria (C1-C3) mentioned in the rationale regarding the ratio of the linewidth of the modulated signal to the width of the oscillator response; $12000/128 = 93.75$ Hz which is of the order of the linewidth of the modulated signal.

In the next section the detailed operation of the circuit will be discussed, followed by an introduction and explanation of how such parameters as the register size and sparse sampling interval was arrived at.

3.2 DETAILED DESCRIPTION OF MAIN CIRCUIT OPERATION

The block diagram of the circuit is shown in Figure 2. The circuit diagram of the realisation will vary according to the implementation of the individual engineer. The basic system works using a "lock step" principle driven by a clock with a fundamental frequency of 3.072 MHz. The clock may best be implemented as an oscillator running at some harmonic of this frequency, using dividers and gates to obtain the correct phases for write enable pulses, etc.

The highest frequency to be analysed in the system is 12 kHz, and the system samples at the Nyquist rate for the highest frequency in the band. The input passes through a hard limiter and the positive or negative value is input to an edge triggered latch as a 1 or 0 respectively.

While each input sample is held in the latch, the following sequence of operations takes place:

- a. The clock advances the counter by 1. The lower 7 bits of the counter select (i) the real and imaginary values of particular oscillator registers, and (ii) the correct pair of entries in the phase tables relating to that oscillator and the current sample.
- b. The values of the oscillator registers are latched; these form the addresses of the decrement factor look-up tables.
- c. The exclusive NOR of the current sample and the phase table real and imaginary parts is added to the outputs of the decrement tables. A 1 output from an exclusive NOR increments the values by 16 input to a register, a 0 decrements the value by 16.
- d. The new values are latched into the registers, and the counter advances by one.

This sequence is repeated until all 128 registers have been serviced, after which the next input signal sample is latched.

The contents of the decrement tables are chosen so that the behaviour of the SPRF satisfies criterion (C4): the response of the circuit is similar to the response of a forced LCR oscillator. The tables are generated by a computer program which emulates the system, and are downloaded to an EPROM programmer.

Since the clock is running at 3.072 MHz the cycle of operations a to d above must take place in 325 nanoseconds. There are two critical paths in the circuit around which a signal has to propagate in this time:

- a. The path from the counter through the phase look-up table, the exclusive NOR and the adder to the registers.
- b. The loop from the registers, through the latch, decrement table and adder, and back to the registers again.

The system does allow 150 ns access time memories to be used for the phase look-up and decrement tables, but the registers to be used are preferably 35 ns RAM with separate input and output terminals.

3.3 STRUCTURE OF THE PHASE TABLES

The phase look-up tables in the case of a 128 oscillator, single bit system are arranged as a 32K x 2 memory. The lowest 7 address lines select the outputs for a particular oscillator (register pair). The upper 8 address lines select the relevant phase of each basis frequency.

Since, in this particular implementation of the system, the lowest frequency signal used has a period of 256 samples, the entire set of tables can be considered to be recorded modulo 256. When the tables are exhausted the counter automatically restarts at zero and the tables repeat.

The contents of the tables may be generated by carefully hard limiting the sine and cosine functions of a computer. The system emulator is used, and the contents of tables used in simulations are downloaded to an EPROM programmer.

3.4 READING THE REGISTERS WITH A MICROPROCESSOR

The system clock may of course be made synchronous with that of a microprocessor, but this is not necessary.

A 7 bit wide comparator may be used to request the contents of a particular register pair. The microprocessor puts the number of the oscillator which it is required to read on the lines to the comparator. When the counter next addresses the appropriate register pair, the comparator outputs a pulse which writes the register contents into separate latches. The pulse from the comparator may be used to interrupt the processor. Hence, the processor can observe the states of the oscillators at its leisure.

Because the decorrelation time of the oscillator amplitudes is substantially longer than the sample to sample time interval, the microprocessor may be arranged to process all of the information available from the oscillator registers. This is achieved by using the microprocessor to observe the contents of each register pair at times which are separated by an order of the decorrelation time (see Appendix 2 for details). This is called "sparse sampling".

The order in which the oscillators are observed may be optimised to reduce the time for which the processor has to wait for the signal from the comparator.

3.5 REGISTER STRUCTURE AND SIZE

The register structure and size needs to be given some consideration if the best performance for any particular application is to be obtained from the system. Clearly, if the registers are too small then hard limiting will take place. Conversely, if the registers are too wide, they will always remain far from saturation and there will be bits which are never used.

The general scheme of register construction for this type of system is shown in Figure 3, for which there are two quantities to be decided. Firstly there is the number of "hidden bits" (ie the bits which are not observed by the microprocessor), which are used to supply sufficient finesse to the system ensuring correct accuracy of the decrements. Secondly there is the "increment bit" where increments and decrements resulting from correlation between the input samples and the phase tables are fed in.

Only the bits which are of equal or greater significance than the "increment bit" need be fed into the adder, since other bits are not affected by the addition or subtraction and may be fed directly into the registers. The maximum value that may appear in the registers (if not hard limited) is decided by the decrement value.

The bits which are fed into the adder from the decrement tables are arranged so that the binary value of those bits is 1 less than it should be in terms of the register value after decrement. The input to the adder from the exclusive NOR gate is to the second least significant bit. Hence, if the output from the exclusive NOR is a 0, the overall effect is to decrement the registers. If the output is a 1, the registers are incremented. By using the fact that the decrement table is implemented as a memory, no explicit subtraction needs to be performed.

There are cases where small registers which do hard limit may be used. Where the system size must be small such a trade-off with performance is acceptable. The rapid carrier acquisition may be delayed by a fraction of a second in such a system, but this will have no operational significance.

The rapid acquisition system has real and imaginary registers each of which are 8 bits wide. The number of hidden bits is chosen to be 4, and the increment bit is chosen to be the 5th (counted inclusively from the least significant bit). The microprocessor has access to all 16 bits of each register pair, but computer simulation has shown that this accuracy is not necessary, and that the less significant bits need not be observed. There is a corresponding graceful degradation in performance as more bits are omitted.

3.6 PHASE AND DECREMENT TABLE GENERATION AND OSCILLATOR EMULATION

Two aspects of circuit design will be mentioned in passing which will become important when more general circuits using the "synthetic oscillator" approach are contemplated.

The LCR circuits which are emulated by the system may easily be made agile in centre frequency and bandwidth. This could be achieved by making both the phase and decrement tables soft, ie different sets of values may quickly be written in by a microprocessor if the appropriate circuitry were added. If a separate decrement table were used for each SPRF by using a larger memory, each oscillator could have a different bandwidth.

It should be recognised that the use of memories as look-up tables for "phase" and "decrement" values allows complete flexibility. For example, one is not restricted to encoding sine and cosine functions into the phase tables, and neither is one restricted to using a "window" which represents exponential decay.

Where very accurate centre frequencies or (correspondingly) narrow bandwidths are required, some sacrifice of the versatility of look-up tables must be traded with a reduction in size gained by using special purpose circuitry. Hence, a special purpose multiplier may be used in place of the decrement tables; the multiplicand determining the bandwidth. Similarly, special purpose hardware may be used to generate the phase tables. An example is shown in Figure 4. Here, the required period divided by four is supplied to a presetable down counter; the required centre frequency precision is supplied by the counting rate. Note that (using this circuit) separate generators are needed for each "oscillator".

Where the sampling rate is sufficiently low and only a small number of oscillators are required, the circuit may be emulated using a micro-processor. This would mean that the special purpose registers and tables would reside in the memory space of the processor, and that (comparatively) long word lengths and correspondingly narrow bandwidths would be available.

4 RESULTS AND DISCUSSION

4.1 INTRODUCTION

The system which has been described was simulated using a Data General MV6000 computer. The software was written in such a way that the important parameters such as the number of oscillators and the register structure could be altered easily. The contents of the look-up tables could be downloaded directly into a PROM (Programmable Read Only Memory) programmer so that they could be used in a hardware realisation of the system. The software was designed to emulate the real circuit operation as closely as possible.

The hardware was constructed using TTL gates, 150 ns EPROM for the look-up tables and 35 ns registers for the oscillator contents. A double Eurocard was used, but some further care in layout and the choice of memory components would reduce the size by about 50%. For example, if a 1024x8 were used as a "decrement table", the adder and multiplier could become part of the look-up table by using two extra address lines. Hybridisation could further reduce the size to about a half of a Eurocard. The real and imaginary output latches were each temporarily connected to an 8 bit digital to analogue converter (DAC) so that the contents could be viewed directly on an oscilloscope as an Argand diagram. The registers corresponding to a single oscillator could be observed by setting up a 7 bit switch connected to the input of the comparator.

Some results from the computer simulation are presented, together with results obtained from the hardware realisation.

4.2 COMPUTER SIMULATION OF THE CIRCUIT

The response of the system to a sine wave and a sine wave with hard phase reversals occurring every 20 milliseconds was investigated. The amount of noise in the 12 kHz band was varied during the experiments. A typical value has been chosen to illustrate the emulator performance. The output is shown as a power spectrum in each case.

Note that the spectrum shown is a result of sparse sampling the contents of each of the oscillators and averaging the results. Normally this operation would be done by a microprocessor connected to the system. This would place a relatively light burden on the microprocessor because of the sparse sampling scheme.

A typical communications signal of the type described working at 50 bits per second would be expected to have at worst a signal to noise ratio of 30 dB Hz. For synchronous phase reversed keying (PRK), the fundamental error rate is about 1 in 10^5 at 29 to 30 dB Hz. Hence, a link using PRK is designed to operate with a lowest signal to noise ratio of 30 dB Hz. As current modems are not adaptive, a margin of at least 6 dB is usually added by increasing transmitter power to account for unknowns such as fading.

Figure 5 shows the results of using the emulator to produce a power spectrum where the input is a pure sine wave at 0.3π per sample, in 0.1 second (simulated) steps. Similarly Figure 6 shows a similar set of plots except that the input is a sine wave with phase reversals at 20 ms intervals. The signal to noise ratio in each case is 30 dB Hz. The horizontal axis is the oscillator number arranged in order of increasing frequency, and the vertical axis is the power (in arbitrary units). The input frequency of 0.3π per sample is not special and lies between the centre frequencies of two oscillators.

4.3 RESPONSE OF HARDWARE TO A SINE WAVE

Sections 4.3 and 4.4. refer to signals with only the basic time-scale τ_s of the signal ((S1) in the rationale).

Figure 7 shows the contents of a pair of oscillator registers displayed on an oscilloscope where the input is a pure sine wave. The persistence of the screen and the shutter speed of the camera (1/30 sec) has averaged the display. No noise is present, and pictures are shown for different offset frequencies of the input from the oscillator centre frequency.

As in all of the photographs shown in section 4, the format is an Argand diagram with the real part on the horizontal axis and the imaginary part on the vertical axis. The real part is the "cosinusoidal" or In-phase channel, and the imaginary part is the "sinusoidal" or Quadrature channel. The picture represents a state diagram for the individual oscillator which is being observed. The brightness of each point is proportional to the amount of time which the oscillator spends in the corresponding state, but the order in which the states are accessed is not obvious from the picture. The outer boundary of the Argand diagram forms a square; since the registers are limited to eight bits the machine cannot access the states outside this square which it would otherwise occupy. This leads to some interesting behaviour as will be discussed later. If the frequency of the incoming sine wave is within a few Hertz of the centre frequency of the oscillator whose behaviour is being observed, the bright points on the image move around the periphery of the square as expected.

In the first photograph, the oscillator spends nearly all of its time on the periphery of the square, most often at the corners. This indicates that the registers are nearly always saturated and that the value of the power ($\text{Real}^2 + \text{Imaginary}^2$) is the maximum possible.

As the frequency moves away from the oscillator centre frequency, the oscillator spends less time in the corners and begins to occupy states near the centre more frequently. This is indicated by a brightening of the spots which are nearer the middle of the screen.

Eventually the pattern begins to collapse towards the centre as the input signals cease to be correlated with the phase tables. Finally the response is similar to that when pure noise is input, except that the pattern is more compact. This is because the input energy lies completely outside the oscillator's band.

4.4 RESPONSE OF HARDWARE TO A SINE WAVE IN NOISE

Figure 8 shows the response of an oscillator to a sine wave at the oscillator centre frequency but with varying amounts of noise present. The results are calibrated in the terms of reference which would be applied to a real signal (ie in dB Hz).

The first diagram shows the response of an oscillator to a signal of 50 dB Hz. This is a signal where the signal power in the whole 12 kHz band is about 10 dB greater than that of the noise. More oscillator states are now accessed because of the noise component.

As the signal to noise ratio is decreased, the region of commonly occupied states becomes larger, until at 20 dB Hz the appearance of the display is similar to that of noise, except that the centre is offset. At 30 dB Hz an offset is readily apparent, but it must be remembered that the oscilloscope screen is integrating the signal.

The striations which are apparent in the photographs arise from the hard limiting due to the limited register size and quantisation effects. The input signal increments or decrements each register by 16 every time that a sample is taken by the machine. It is an artefact of the decrement table organisation that if a register approaches its maximum or minimum value, it is reset to a value which is 16 or 48 (depending on the next sample value) units away from the extremum.

The quantisation of the decrement tables then allows the registers to change their value by only 15, 16 or 17 after each subsequent sample is taken. If the next hard limiting occurs after only a few samples have been taken, then the effect of the above restricted register changes limits the number of register states that can be occupied. This clearly manifests itself as striations on the oscilloscope screen.

4.5 THE EFFECT OF PHASE REVERSALS OF THE NOISY CARRIER

In this experiment, the timescale (S2) is introduced into the signal. A realistic value of $\tau_c = 20$ ms is used.

Figure 9 shows the oscillator Argand diagram which results when there are 50 hard phase reversals per second of the carrier, and a telegraph signal at that rate with varying amounts of noise present.

As the noise level is increased, the distribution in the display tends towards a "dumbbell" shape. The poorer noise tolerance is to be expected since the energy of the modulation signal is spread over a greater bandwidth than before, and consequently less is "absorbed" by an oscillator centred on the carrier frequency.

Note that hard phase reversals repeated 50 times a second split the line in the carrier spectrum into two main lobes at ± 25 Hz from the original frequency. If the frequency source is retuned by 25 Hz in either direction, a stationary display similar to those in Figure 9 is obtained. The system does not however have sufficient resolving power to separate the two components completely. The implications of a system which can do this will be discussed in another publication.

5 CONCLUSIONS

Using general physical arguments based on factorisation of length/time-scales, and on an LCR analogy, we have identified an irreducible description of the structure of a low data rate carrier that is hidden in white noise. This description involves a 1 bit representation of the incoming (noisy) signal which is then correlated with a set of 1 bit templates in order to achieve a set of matched filter operations. The outputs from this are added to registers which have a limited memory time (ie the register bank of forced LCR oscillators).

The implementation is such that 128 oscillators operating at baseband (ie 0-12 kHz) may be emulated by a single serial digital circuit operating at high speed (ie 3 MHz). Although the input characteristic frequency is 12 kHz, the characteristic frequency with which the register contents vary is 94 Hz ($= 12 \text{ kHz}/128$) which is the inverse register decay time. This allows interpretation of a restricted number of the registers by a relatively slow serial device such as a microprocessor.

The circuit was constructed and its performance was measured; all the results were in agreement with the theoretical predictions. Preliminary results indicate that a 50 bit/s phase reverse keyed telegraph signal at 30 dB Hz in a 12 kHz bandwidth may be located and phase locked to in $\frac{1}{2}$ second or less.

The flexibility inherent in this approach to communications signal processing has not been fully explored in this Memorandum. By careful design of the initial filtering operation and/or by simultaneous interpretation of the behaviour of several of the registers more sophisticated use of available bandwidth is possible. One of us (JP) is researching into these possibilities and in the future will publish extensions of the basic ideas outlined here.

ACKNOWLEDGEMENT

The authors would like to thank Mr Paul Wells for constructing the circuit.

REFERENCES

- 1 "Signal Transforming Device", D/PD Pats/P.0288, Patent Application No 8611,180.
- 2 "A Manpack Satellite Communications Earth Station", C H Jones, The Radio and Electronic Engineer, Vol 51, No 6, p259, June 1981.

APPENDIX 1: DERIVATION OF PROCESSOR STRUCTURE USING GENERAL PRINCIPLES

We shall derive the hardware structure from a set of general principles:

ASSUMPTION 1

The hardware will measure a finite number N of statistics $s_k(n)$ ($k = 1, 2, \dots, N$) of the data samples $x(n)$.

ASSUMPTION 2

In order that the hardware structure be simple, we shall demand that the $s_k(n)$ be linearly self-recursively definable, so

$$s_k(n) \equiv \alpha_k(n) s_k(n-1) + \beta_k(n) x(n) \quad k = 1, 2, \dots, N \text{ and } n \in \mathbb{Z} \quad (A1.1)$$

where $\alpha_k(n)$ and $\beta_k(n)$ are both real. $\alpha_k(n)$ is interpreted as a "memory" term, and $\beta_k(n)$ is interpreted as "weighting" the forcing term $x(n)$.

ASSUMPTION 3

To further simplify the hardware we shall restrict the $\alpha_k(n)$ to have the form

$$\begin{aligned} \alpha_k(n) &\equiv \alpha \in \mathbb{R} \\ 0 &\leq |\alpha| < 1 \end{aligned} \quad (A1.2)$$

This restriction on $\alpha_k(n)$ still permits us to specify a "memory length", but it can now neither depend on the statistic nor the time.

Assumptions 1, 2 and 3 allow us to solve the defining equation for $s_k(n)$ to give

$$s_k(n) = \sum_{r=0}^{\infty} \alpha^r \beta_k(n-r) x(n-r) \quad (A1.3)$$

ASSUMPTION 4

To further simplify the hardware we shall assume that the $\beta_k(n)$ are periodic in n . The natural choice for the period is N . This is because there are exactly N statistics s_k , so perfect information retention by the s_k can only ever be achieved for at most N samples of $x(n)$. It is therefore pointless to permit the $\beta_k(n)$ (for fixed k and variable n) to depend on more than N parameters. This periodicity restriction may be written as

$$\beta_k(n+N) \equiv \beta_k(n) \quad (A1.4)$$

For the same reason the memory length introduced by α must be of order of N , so typically

$$\alpha \equiv 1 - \frac{1}{N} \quad (A1.5)$$

to ensure that

$$\alpha^N \ll 1 \quad . \quad (A1.6)$$

ASSUMPTION 5

It remains to determine only the detailed form of the $\beta_k(n)$. This can be decided only when the structure of the signal which we wish to detect in the $x(n)$ is specified. Clearly it is desirable for a single statistic s_k to respond strongly to a given signal, for then a simple threshold operation may be used to determine the signal from the information collated by the statistics s_k . We may ensure this behaviour by defining

$$\beta_k(n) \equiv x_k(n) \quad (A1.7)$$

where $x_k(n)$ is the time series obtained from a noiseless version of desired signal k . This specification is more complicated when either the noise is correlated, or the signals $x_k(n)$ are not orthogonal.

For sinusoidal $x_k(n)$ the output takes the form

$$s_k(n) = \sum_{r=0}^{\infty} \alpha^r \begin{matrix} \cos[(n-r)\phi_k] \\ \sin[(n-r)\phi_k] \end{matrix} x(n-r) \quad (A1.8)$$

where ϕ_k is the phase advance of $x_k(n)$ per sample, and the cosine or sine are chosen according to whether the in-phase or in-quadrature channel is being considered.

APPENDIX 2: SOME USEFUL MATHEMATICAL RESULTS

The results that are quoted in this Appendix ignore the effects of quantisation. This approximation is best when there is no hard limiting and the signal to noise ratio at the input is less than of order of 1 (ie quantisation effects are minimal). These requirements are not independent of each other, and they are usually met in practice.

Let us define some notation:

- τ $\equiv$ raw sample the separation
- $m\tau$ $\equiv$ sparse sample time separation
- $a_k^{\cos}(n)$ $\equiv$ amplitude of cosine oscillator k at $t = n\tau$
- $a_k^{\sin}(n)$ $\equiv$ amplitude of sine oscillator k at $t = n\tau$
- $E_k(n) \equiv [a_k^{\cos}(n)]^2 + [a_k^{\sin}(n)]^2$
- ω_k $\equiv$ angular frequency of oscillator k
- $\phi_k \equiv \omega_k \tau$ $\equiv$ phase per time τ for oscillator k
- $\cos(n\phi_k)$ $\equiv$ cosine phase table entry for oscillator k at $t = n\tau$
- $\sin(n\phi_k)$ $\equiv$ sine phase table entry for oscillator k at $t = n\tau$
- α $\equiv$ decrement factor
- $x(n)$ or x_n $\equiv$ raw data value at $t = n\tau$

Oscillator response equation:

$$\begin{aligned} a_k^{\cos}(n) &= \sum_{r=0}^{\infty} \alpha^r \cos[(n-r)\phi_k] x_{n-r} \\ a_k^{\sin}(n) &= \sum_{r=0}^{\infty} \alpha^r \sin[(n-r)\phi_k] x_{n-r} \end{aligned} \quad (A2.1)$$

$$E_k(n) = \sum_{r,s=0}^{\infty} \alpha^{r+s} \exp[-i(r-s)\phi_k] x_{n-r} x_{n-s} \quad (A2.2)$$

Average response to a cosine wave input $x_n = \cos(n\phi)$:

$$\langle E_k(n) \rangle \simeq \frac{1}{4} \frac{1}{1 + \alpha^2 - 2\alpha \cos[\phi - \phi_k]} \quad (A2.3)$$

where we have used the approximation $|\phi - \phi_k| \ll |\phi + \phi_k|$.

Average response to white noise input $\langle x_n \rangle = 0$, $\langle x_m x_n \rangle = N\delta_{mn}$:

$$\langle E_k(n) \rangle = \frac{N}{1 - \alpha^2} \quad (A2.4)$$

$$\langle E_k(m) E_k(n) \rangle = \left(\frac{N}{1-\alpha^2} \right)^2 \left[1 + \alpha^{2(m-n)} \left(1 + \frac{(1-\alpha^2)^2}{|1 - \alpha^2 e^{-2i\phi_k}|^2} \right) \right] \quad (A2.5)$$

Comparison of (A2.4) and (A2.5) reveals that $E_k(m)$ and $E_k(n)$ are substantially decorrelated when $\alpha^{2|m-n|} \ll 1$ (this result is related to (A1.6)). (A2.4) and (A2.5) provide a theoretical justification for the sparse sampling scheme that was adopted.

APPENDIX 3: RESPONSE TO A NOISY SINE WAVE AND SATURATION EFFECTS

The effect of 1 bit quantisation of this input signal may be substantially accounted for by ensuring that $\langle x_n^2 \rangle = 1$. The result in (A2.3) is quoted for $x_n = \cos(n\phi)$ where $\langle x_n^2 \rangle = \frac{1}{2}$. The result in (A2.4) is quoted for $\langle x_n^2 \rangle = N$. Combining these two results in such a way that the SNR (per sample on average) is R yields

$$\langle E_k(n) \rangle \simeq \frac{R}{1+R} \cdot \frac{1}{2} \cdot \frac{1}{1 + \alpha^2 - 2\alpha \cos[\phi - \phi_k]} + \frac{1}{1+R} \cdot \frac{1}{1-\alpha^2} \quad (A3.1)$$

For the particular implementation that we have described in this Memorandum, namely satellite communications, the relevant parameters are

$$f_{\max} = 12000 \text{ Hz} \quad ; \quad f_{\text{bin}} = 93.75 \text{ Hz}$$

$$\tau = 41.7 \text{ } \mu\text{s}$$

$$\alpha = 0.98773$$

The phase per sample per bin width is

$$\begin{aligned} \Delta\phi_{\text{bin}} &= 2\pi f_{\text{bin}}\tau \\ &= 2.456 \times 10^{-2} \text{ radian} \end{aligned}$$

A conversion table between SNR R and the equivalent measure R' in dB Hz is

$R'/\text{dB Hz}$	R
20	0.0083
30	0.083
40	0.83
50	8.3

where the expression

$$R' = 10 \log_{10} [12000 R]$$

has been used.

Figure 10 shows $\langle E_k(n) \rangle$ given by (A3.1) plotted against k , where the value of k is given as the difference between the actual k and the centre frequency k . This figure shows the response (in squared units) versus displacement from centre frequency (in bin widths). Various dB Hz levels are shown.

It is instructive to compare the predictions of Figure 10 with the results displayed in Figures 7 and 8. The Argand diagrams shown have an x and y scale that extends from -8 to +8. This corresponds to the fact that we add in the incoming signal at bit number 5 of the 8 bit oscillator registers, and so the oscillator has only 4 bits to represent (the integer part of) its response. The maximum energy (squared response) that the oscillator can store is therefore 128 units ($= 8^2 + 8^2$) and this occurs in the corners of the Argand diagram. At the extremities of the x and y axes the energy is 64 units. These energy units are the same as those used in Figure 10.

Figure 11 can now be compared directly with the results in Figures 7 and 8. Clearly any value of $\langle E_k \rangle > 128$ will saturate at least one of the oscillator registers, so there is a substantial range of SNR within which saturation should be clearly observed. An important caveat to note is that Figure 11 refers to $\langle E_k \rangle$ (not E_k), so there is the possibility of intermittent saturation for values of $\langle E_k \rangle$ that seem "safe". Referring to Figure 11, the oscillator should be saturated for a pure sine wave input for $k = 0, 1, 2$. This is borne out by the results shown in Figure 7.

At $k = 5$ ($\Delta f \sim 470$ Hz), $\langle E_5 \rangle = 33$, $\sqrt{\langle E_5 \rangle} = 5.7$,
and at $k = 21$ ($\Delta f \sim 2$ kHz), $\langle E_{21} \rangle = 1.9$, $\sqrt{\langle E_{21} \rangle} = 1.4$.

Both of these results are consistent with Figure 8.

The results that are shown in Figure 8 correspond to the intercepts (on the $\langle E_k \rangle$ axis) in Figure 10. These intercepts are in tabulated form

SNR/dB Hz	$\langle E_0 \rangle$	$\sqrt{\langle E_0 \rangle}$
50	2968	54.5
40	1529	39.1
30	292	17.1
20	68	8.2
NOISE	41	6.4

Comparing these results with Figure 8 shows that the results are compatible. Clearly the effect of noise is to blur out the response so that an oscillator that is saturated (according to $\langle E_0 \rangle$) will spend some of the time unsaturated.

The overall impression is that (A3.1) provides a good estimate of $\langle E_k \rangle$. The theory could be extended to account for quantities such as $\langle (E_k)^n \rangle$ $n > 1$, which would provide a more rigorous means of analysing the effects of noise.




FIGURE 1 SWEPT FILTER AND INTEGRATOR




FIGURE 2 BLOCK - DIAGRAM OF RAPID ACQUISITION CIRCUIT




FIGURE 3 GENERAL REGISTER STRUCTURE




FIGURE 4 METHOD OF GENERATING SINE & COSINE FUNCTIONS AUTOMATICALLY







FIGURE 5 COMPUTER SIMULATION OF POWER SPECTRUM GENERATED SINE WAVE AT 300HZ







FIGURE 8 COMPUTER SIMULATION OF POWER SPECTRUM GENERATED - SINE WAVE AT 30dBHz WITH HARD PHASE REVERSALS EVERY 20ms







FIGURE 7 RESPONSE OF OSCILLATOR TO VARYING FREQUENCY, NO NOISE. ARGAND DIAGRAMS $\frac{1}{30}$ SEC EXPOSURE




50dB Hz




40dB Hz




30dB Hz



20dB Hz



NOISE



NOISE AVERAGED FOR 4 SEC (4 SEC EXPOSURE)

FIGURE . 8 RESPONSE OF OSCILLATOR TO VARYING NOISE. INPUT IS AT OSCILLATOR CENTRE FREQUENCY. INPUT BAND = 0 - 12 KHz - EXP $\frac{1}{30}$ SEC.



FIGURE .9 RESPONSE OF OSCILLATOR TO 50 HARD PHASE REVERSALS PER
SECOND




FIGURE.10 (SEE TEXT)

DOCUMENT CONTROL SHEET

Overall security classification of sheet ...UNLIMITED.....

(As far as possible this sheet should contain only unclassified information. If it is necessary to enter classified information, the box concerned must be marked to indicate the classification eg (R) (C) or (S))

1. DRIC Reference (if known)	2. Originator's Reference Memorandum 3925	3. Agency Reference	4. Report Security Classification Unlimited	
5. Originator's Code (if known)	6. Originator (Corporate Author) Name and Location Royal Signals and Radar Establishment			
5a. Sponsoring Agency's Code (if known)	6a. Sponsoring Agency (Contract Authority) Name and Location			
7. Title A COMPACT DIGITAL COMMUNICATIONS SYSTEM PART 1: RAPID CARRIER ACQUISITION				
7a. Title in Foreign Language (in the case of translations)				
7b. Presented at (for conference papers) Title, place and date of conference				
8. Author 1 Surname, initials Luttrell S P	9(a) Author 2 Pritchard J A S	9(b) Authors 3,4...	10. Date	pp. ref.
11. Contract Number	12. Period	13. Project	14. Other Reference	
15. Distribution statement				
Descriptors (or keywords)				
continue on separate piece of paper				
Abstract The problem of reducing the carrier centre frequency ambiguity in low data rate satellite communications is well known. Prior solutions to this problem have been slow, bulky or expensive. Our proposed solution does not have these drawbacks, and it is easily implemented in readily available digital hardware. This is achieved by using a 1 bit digital emulation of a bank of forced LCR oscillators to achieve an approximate Fourier decomposition.				

END

DATE
FILMED

3-87