Low Cost Coherent Demodulation for Mobile Satellite Terminals

Santanu Dutta and Steven J. Henely

Mobile Communication Satellite System
Rockwell International
400 Collins Road, NE
Cedar Rapids, Iowa 52498
(319) 395 8257

ABSTRACT

This paper describes some low cost approaches to coherent BPSK demodulation for mobile satellite receivers. The specific application is an Inmarsat-C Land Mobile Earth Station (LMES), but the techniques are applicable to any PSK demodulator. The techniques discussed include combined sampling and quadrature downconversion with a single A/D, and novel DSP algorithms for carrier acquisition offering both superior performance and economy of DSP resources. The DSP algorithms run at 5.7 MIPS and the entire DSP subsystem, built with commercially available parts, costs under $60 at quantity-10,000.

INTRODUCTION

Low cost mobile terminals are essential to the commercial success of many of the mobile satellite services (MSS) being launched today. This is because the services are based on the premise of a mass market, whose penetration will be critically dependent on the terminal cost. Driven by the desire to also minimize the space segment cost, some mobile satellite service providers, such as Inmarsat in its "C" standard, have opted for coherent rather than differential PSK modulation. Coherent demodulation is more complex than differential detection, the more popular approach, and tends to increase the terminal cost.

Modern PSK demodulators are almost invariably implemented by DSP. However, it is not widely recognized that, per today's pricing of programmable DSP chips, the cost of the DSP solution goes up quite rapidly with the processing speed (millions of instructions per second, or MIPS) and the on-chip memory. Floating point chips also extract a premium over fixed point devices. Table 1 shows the MIPS, internal memory and unit cost at quantity-10,000 for some popular programmable DSP chips.

It is clear from Table 1 that there is considerable incentive to engineer the DSP alogrithms for minimum MIPS and on-chip memory, while keeping the performance *acceptable*. In the present case, "acceptable performance" was that defined in the Inmarsat-C System Definition Manual (SDM).

Table 1. Comparison of Low Cost DSP Chips

Vendor	Product	Type	Cycle Speed (MHz)	On-Chip Prog. Mem.	On-Chip Data Mem.	Unit Cost @ qty.10K ($)
Analog Devices	ADSP-2105	16-bit Fixed Pt.	10.0	1K x 24 (RAM)	512 x 16 (RAM)	20.00
	ADSP-2101	16-bit Fixed Pt.	12.5	2K x 24 (RAM)	1K x 16 (RAM)	38.00
	ADSP-21010	32-bit Fl. Pt.	12.5	external	external	49.88
Texas Instr.	TMS320C25	16-bit Fixed Pt.	12.8	4K x 16 (ROM)	544 x 16 (RAM)	14.86
	TMS320C51	16-bit Fixed Pt.	28.5	8K x 16 (ROM) 1K x 16 (dual RAM)	1K x 16 (RAM)	28.89
	TMS320 C31-27	32-bit Fl. Pt.	13.5	4K x 32 (dual ROM) 2K x 32 (dual ROM)		40.00

Prices are for the Industrial temperature range (-25 C to +85 C),
except for TMS320C51 and TMS320C31-27, which are currently
not available in that temperature range.
(The Industrial temperature range is preferred for land mobile services)

In this paper, the techniques discussed are (1) simultaneous bandpass sampling and quadrature downconversion using a single A/D, and (2) a novel DSP algorithm for carrier acquisition, with features suitable for MSS. Other DSP innovations are also featured in the Rockwell Inmarsat-C terminal but are not discussed here for lack of space and proprietary reasons.

Demodulator Hardware Architecture

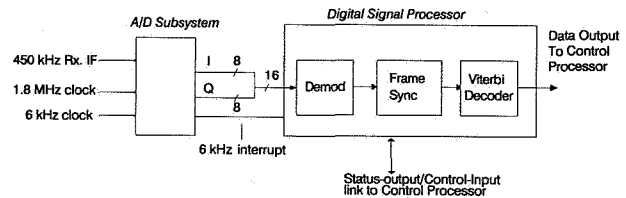Figure 1 shows functions performed in DSP in the Rockwell LMES.

Figure 1. Demodulator Hardware Architecture

The received signal at the IF of 450 kHz is simultaneously sampled and quadrature downconverted to the nominal frequency of 0 Hz by a single 8-bit A/D. Thereafter, the complex samples are processed in the DSP chip to perform the functions of carrier frequency and phase estimation, symbol synchronization, BPSK matched filtering, frame deinterleaving, UW detection and tracking, erasure sensing and Viterbi decoding. Of these, only carrier frequency and phase estimation are discussed here.

INPUT SAMPLING/ DOWNCONVERSION

Subharmonic quadrature sampling was the technique used. This combined the processes of input sampling and quadrature downconversion, leading to a significant reduction or simplification of the signal conditioning circuitry preceding the A/D. We first discuss two conventional approaches and then describe the approach used.

Conventional Approach A

Probably the most conventional approach is to multiply (mix) the IF signal with two quadrature-phase local oscillator (LO) signals at IF. The mixer outputs are lowpass filtered and digitized by "slow" A/D converters. Figure 2 shows a block diagram of this approach.
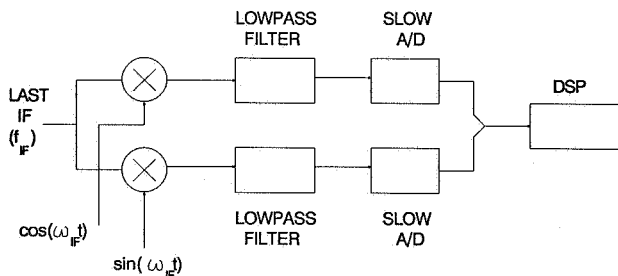

Figure 2. Conventional I/Q Downconversion and Input Sampling

The required operating bandwidth of these A/D's is the sampling rate, which, according to the Nyquist theorem, is lower bounded by the width of the receiver's stopband. In the present design, the stopband width is approximately 6 kHz.

For ideal circuit components, this yields mathematically perfect I and Q samples. However, because of the separate analog paths of the I and Q signals, the possibility of gain and phase mismatch exists. Selecting components with good match increases the terminal cost. There is also the problem of DC bias if active mixers are used. The two-path sampling approach also requires many more discrete components than a single-path approach, as used in Conventional Approach B and the present implementation.

Conventional Approach B

Another approach common in today's digital wireless modems is single-path sub-Nyquist sampling with the complex downconversion in DSP. In this approach, the input sampling rate has to be at least 2-times the IF stopband width. After digital Hilbert transformation, the sampling rate can be decimated to a frequency equal to the stopband width.

Implemented Approach

The bandpass IF signal at 450 kHz was sampled by *pairs of pulses*, with an intra-pair time separation of $1/(4.\mathrm{IF}) = 1/(1.8 \text{ MHz})$ and inter-pair separation of $1/(6 \text{ kHz})$. Figure 3 explains the concept.
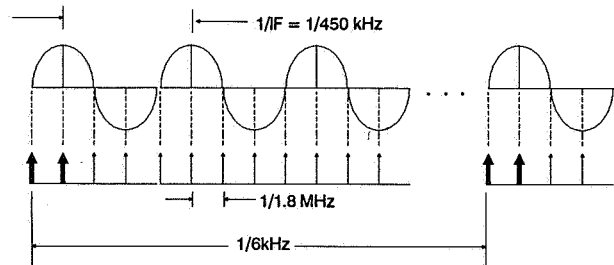

Figure 3. Concept: Combined I/Q Downconversion and Input Sampling

The process may be thought of as sampling with the complex sampling function, $(1+j1)\delta(t-nT)$. The 90-degree phase shift between the sampling pulses in each pair is achieved by time staggering. As in Conventional Approach B, only a single A/D is required. Moreover, the input rate is 6 kHz, unlike Conventional Approach B, in which the rate would be 12 kHz. Figure 4 shows the hardware block diagram.
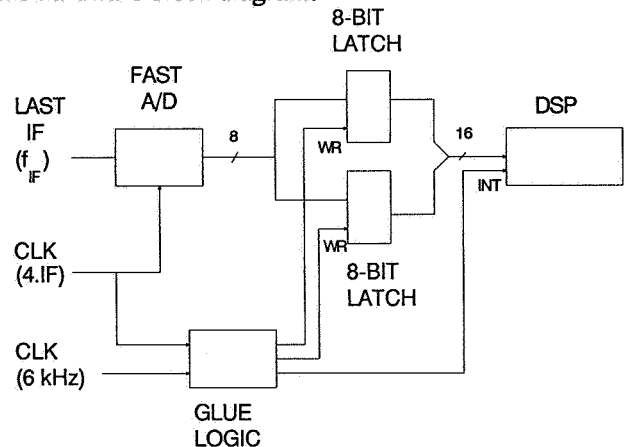

Figure 4. Hardware Block Diagram: Combined I/Q Downconversion and Input Sampling

The economy of this approach, both in sampling cir-

cuitry and DSP code is obvious. Compared to Conventional Approach A, the advantages are in circuit matching requirements and low parts count. Compared to Conventional Approach B, it avoids the digital Hilbert transformer and runs at a lower input sampling rate.

Although economical in hardware and immune to mismatch errors, this approach does have its own imperfection. However, the imperfection is acceptable here because of the noisy input signal. The input samples are at perfect phase quadrature only at the nominal IF (450 kHz). If the band of interest (stopband) is small compared to the nominal IF, the phase error relative to quadrature, even at the band edge, is small. For example, in the present design, the stopband is approximately +/-3 kHz. The phase error for this frequency is $\theta = +/-0.6$ degrees. It can be shown that this creates a cochannel self-interference term that is at $20\log_{10}\{\sin(\theta)\}$ relative to the desired signal, i.e approximately -40 dBc. The fact that the modem operates in E_s/N_0 of typically 3.7 dB makes this level of self-interference quite acceptable.

The present approach also requires a faster A/D than the conventional approaches. In the present design, an operational bandwidth of 4.IF = 1.8 MHz is required, as opposed to a bandwidth of 6 kHz in Conventional Approach A and 12 kHz in Conventional Approach B respectively. However, low cost flash A/D's in the low MHz range are now available, making this approach a better choice from a cost standpoint. The cost of the A/D used, in quantity-10K, was $6.75.

CARRIER ACQUISITION ALGORITHMS

Demodulator Requirements

The detailed performance requirements are given in the SDM [1] and are not repeated here. However, the key challenges are highlighted.

Transmit Signal Characteristics
Modulation:	Unfiltered BPSK
Coding:	Rate-1/2 Convolutional
Symbol rate:	1200 bps

Fading Channel Characterisitics
Unfaded C/N_0:	34.0 dBHz
Fading Type:	Rician, C/M = 7 dB
Fading bw.:	0.7 Hz

Blocked Channel Characteristics
Unblocked C/N_0:	35.0 dBHz
Duration:	2.7 s
Period:	8.9 s

Doppler Characteristics
Max. Shift:	+/- 850 Hz
Variation Rate:	+/- 10 Hz/s

In [1], the performance requirements are specified in terms of the Packet Error Rate (PER) for the fading channel and the blocked channel separately. Performance specifications for different packet sizes and input

C/N_0 are provided; here we cite only those for the 128-byte packets and the above C/N_0 values for illustration.

Typical SDM Performance Requirements
PER(128) fading ch.:	8%
PER(128) blocked ch.:	10%

When the demodulator performance requirements are translated into carrier acquisition requirements, the following facts emerge:

(1) The low prevailing C/N_0, while making the demodulation task difficult, makes it possible to use non-ideal processing techniques. This was exploited in the input sampling scheme.

(2) Conventional phase locked loop techniques for carrier acquisition will not work because of the conflicting requirements of *large capture bandwidth* (+/- 850 Hz) and *rapid acquisition* on the one hand, and *low phase noise* on the other. The capture range was actually set even higher, at +/- 1300 Hz, to enable rapid frequency search on power up. During the latter phase, the receiver is hopped in 2.5-kHz steps. Rapid carrier acquisition is required so that (a) the initial frequency search time is short, and (b) not many bits are lost when the LMES emerges from a blockage or transmit mode (the communication is half-duplex).

Review of Carrier Recovery Techniques for BPSK Demodulation

The two major problems in BPSK demodulation are recovery of the carrier phase and symbol clock phase. In this paper, we discuss only the former.

Carrier recovery may be performed by either open loop or closed loop techniques. One popular open loop technique is to continuously operate an FFT in the background and use it to obtain a coarse frequency estimate; this is used to aid a closed loop carrier synchronizer. An alternative approach is described by Viterbi [2]. Both of these techniques are much more complex and demanding of DSP resources than the closed loop techniques. It was therefore decided to implement the present demodulator based on closed loop techniques alone.

A simple phase locked loop cannot be used because a BPSK signal has a suppressed carrier. However, modified phase locked loops are usable, such as the squaring loop and the Costas loop. Many text books, e.g. [3], provide extensive coverage of both techniques. The Costas loop has the advantage over the squaring loop that it is capable of wider bandwidth operation [ibid p.304]. Therefore, the Costas loop was chosen.

Cahn has analyzed the performance of the Costas loop and shown that, in most receive applications, there is a conflict between the required lock-time/capture-range and the acceptable level of phase noise [4]. This is explained below.

The capture range of a basic Type I phase locked loop (without a perfect integrator in the loop filter) is directly proportional to the loop resonance frequency, and hence also the loop bandwidth [3, p.364]. As the loop bandwidth determines the amount of phase noise in the loop's voltage controlled oscillator (VCO), it is clear that *large capture bandwidth and low phase noise are conflicting requirements for a Type I loop*.

In Type II loops (loop filter has a perfect integrator), the capture range is theoretically unbounded. In practical systems, it is bounded by the loop's dynamic range. Thus, *for Type II loops, the capture range and phase noise (loop bandwidth) are unrelated*.

We now examine the lock time. This is given by Gardner for Type II loops as [5, p. 76]

$$T_{acq} = 4.2(\Delta f)^2/B_n^3 \qquad (1)$$

where,

T_{acq}:   acquisition time
Δf: frequency offset
B_n: loop noise bandwidth

The expression for Type I loops is very similar [Spilker, p.364] and differs only in the multiplying constant. Irrespective of the type of loop, note that the lock time is inversely proportional to the cube of the loop bandwidth. This makes it difficult to simultaneously achieve rapid phase lock during carrier search, and low phase noise during carrier tracking.

Cahn proposed to overcome this problem by creating an outer frequency locked loop around the inner phase locked loop, as shown in Figure 5 (excluding the adaptive AFC gain control, which is the contribution of the present work).
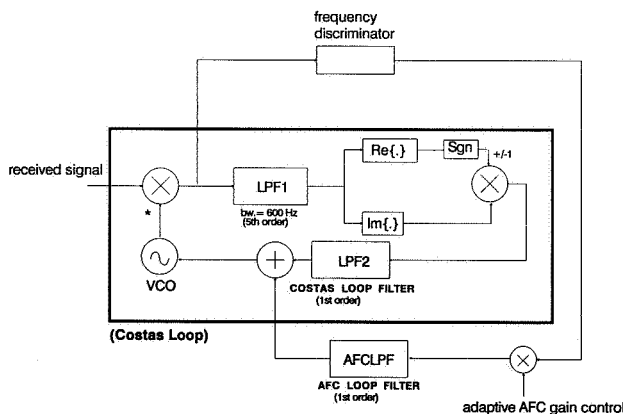

Figure 5. AFC-aided Costas Loop

The outer loop acts as an automatic frequency control (AFC) loop; this configuration is known as the "aided" phase locked (or Costas) loop. Both the outer

and the inner loops provide error signals to the VCO input. However, the AFC loop's contribution is proportional to the *frequency difference*, and not the *phase difference*, relative to the input carrier. This configuration increases the capture range and reduces the lock time because an AFC loop can perform the task of pulling in carriers with large offset much better than a phase locked loop. However, the AFC loop also contributes noise to the VCO's driving function. Therefore, the AFC loop bandwidth has to be limited so that its contribution to the VCO's phase noise is small compared to that of the phase locked loop. Cahn found an AFC loop bandwidth of 0.1 times the bandwidth of the phase locked loop to be a suitable choice.

New Carrier Recovery Scheme

Although Cahn's AFC loop solves the capture range problem and provides some help in reducing the lock time, the latter is still unacceptable per the present design goals, given below.

Carrier Recovery Design Goals
Loop Bandwidth: 60 Hz
(determine by phase noise
and Doppler rate tracking)
Capture Range: +/- 1300 Hz
Mean Acquisition Time: 1 s

In the present demodulator, the lock time was further reduced over Cahn by making the AFC loop gain adaptive. The inner loop was a Type II Costas loop with 60-Hz loop bandwidth, a capture range of over 100 Hz and an acquisition time (for 100-Hz offset carriers) of approximately 0.8 s. The outer loop had a capture range of +/- 1300 Hz and an acquisition time of approximately 1.0 s. The adaptive AFC gain control scheme is described below.

The AFC gain scheme first investigated was:

IF(NOT.(inner loop lock)) THEN
     AFC GAIN = HIGH
ELSE
     AFC GAIN = LOW

Practical implementation of this scheme revealed a number of problems. It was found that, in order to achieve the target lock time of 1 s, the AFC loop gain had to be increased to a very high value. At this high AFC gain, the phase noise contributed by the AFC loop often prevented the inner loop from locking. Moreover, the AFC loop gain would undergo damped oscillation around its steady state value for an unacceptable length of time before the frequency uncertainty settled down to within the 100-Hz capture range of the inner loop. This meant that the decision to switch the AFC gain to a low value could not be based on a lock detector operating on the inner loop. The AFC loop would have to autonomously switch gain, based on some measurement of its own state.

Autonomous AFC Gain Switching

The key requirement is for the AFC loop to determine that it is sufficiently close to its steady state value. The time response of the AFC loop's error signal, e_{AFC}, to a step change in frequency, for large AFC gains, has the characteristic underdamped shape shown in Figure 6.
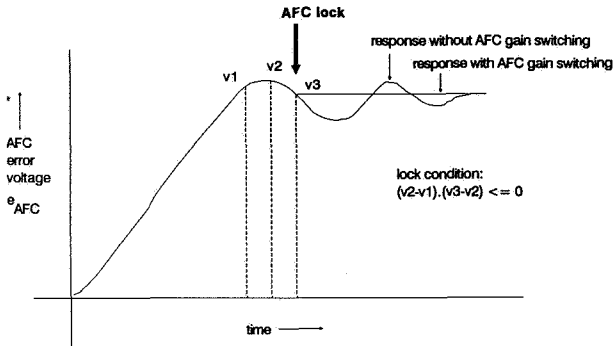

Figure 6. Typical AFC error signal response without noise (artist's impression)

It is clear that a change in sign of the derivative of e_{AFC} (noise assumed to be absent) indicates that e_{AFC} has just traversed its first peak. Usually, this point is sufficiently close to the steady state value. Thus, a change is sign of the first derivative of e_{AFC}, say e_{AFC}', may be taken as the signal to clamp down the AFC gain. Figure 6 shows this conceptually.

When noise is present, this approach is not fool-proof as noise can cause premature sign changes in e_{AFC}'. The following remedy was applied to this problem. e_{AFC} was filtered by a 1.6 Hz bandwidth filter before its derivative was taken. However, this measure, by itself, could not eliminate all occurrences of spurious AFC lock indication. Thus a waiting time of 0.8 s was introduced on each occurrence of AFC lock. During this time, the AFC gain would be clamped down to its LOW value. If, at the end of this period, the inner loop still indicated no lock, the AFC gain would be returned to its HIGH value. The waiting time was selected to be 0.8 s as this was the acquisition time of the inner loop.

Special Accommodation for Fading and Blocked Channels

Some customization of the above concepts were incorporated to further improve performance in fading and blocked channels. Instead of two AFC gains, three gain values were used -- HIGH, MEDIUM and LOW.

High AFC gain was used on "initial search" for the carrier. The "initial search" condition was defined to exist on power up and on changing the receiver's tuned frequency. If, after once acquiring the carrier, it was lost (presumably due to fading or blockage) then the MEDIUM gain was applied. The use of a medium gain

ensured rapid resynchronization when the carrier returned from a fade or blockage. As it had been gone only for a short period, the frequency could not have changed very much. If the inner loop remained continuously out of lock for more than 25 s, the "initial search" condition was declared to exist, on the assumption that significant frequency change might have occurred in the intervening period (due to Doppler or oscillator drift). The LOW gain was applied only when the inner loop was locked.

DEMODULATOR PERFORMANCE RESULTS

The "proof of the pudding" for the above techniques is in meeting the SDM PER requirements and the derived requirement of 1-s carrier acquisition time. Figures 7 and 8 show the PER performance of the Rockwell LMES demodulator in the SDM fading channel.
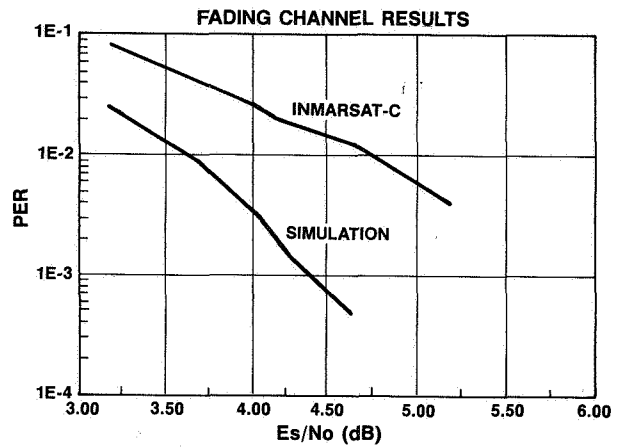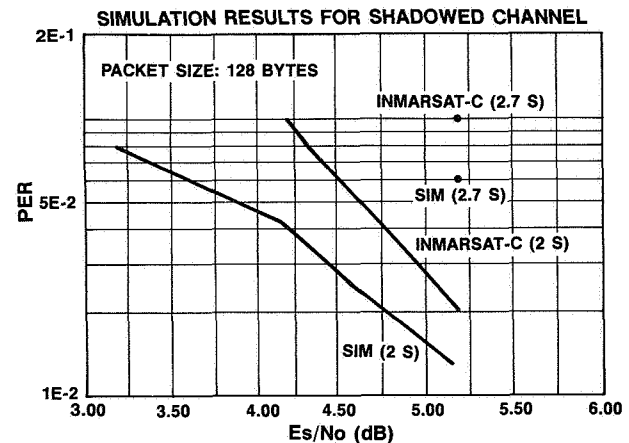

Figure 7. Fading Channel Performance

Figure 8. Blocked Channel Performance

It is clear that Inmarsat-C performance requirements are satisfied.

397

The statistics of the carrier acquisition time, for 850 Hz offset, in a fading channel with unfaded C/N_0 = 34 dBHz, is given below (simulation results).

Carrier Acquisition Statistics
Mean carrier acquisition time: 1.1 s
Median carrier acquisition time: 0.9 s
90-percentile carrier acquisition time: 1.8 s

Since the other major aim of the design was cost minimization, the outcome of that effort is noted below.

The processing speed requirement of the Demodulator part of the algorithm is 4.8 MIPS, while that for the entire DSP subsystem is 5.7 MIPS. The program code size is approximately 1.5K. The implementation was based on one Analog Devices ADSP 2101 chip, for which the quantity-10K unit price is approximately $38. If the program memory size could be reduced to under 1K, the ADSP 2105 chip could be used at the quantity-10K unit price of $20. This is considered feasible by additional innovations in code optimization and is planned for future product revisions.

The entire cost of the DSP subsystem, including external memory, A/D and other sampling circuit components, is under $60.

SUMMARY

When addressing a mass market, it is important to minimize product cost while keeping product performance above a defined level of acceptability. While the costs of DSP parts have been falling in general, a significant difference still exists between the "low-end" parts with modest MIPS and on-chip memory, and the higher end parts featuring greater DSP resources. In this paper, some novel input sampling techniques and DSP algorithms are presented which helped to realize an Inmarsat-C demodulator using low-end parts.

REFERENCES

[1] *Inmarsat-C System Definition Manual (SDM), Release 2.0*, Inmarsat, London, September 1992.

[2] Viterbi, A. J. and Viterbi, A. M., "Nonlinear Estimation of PSK-Modulated Carrier Phase with Application to Burst Digital Transmission," *IEEE Transactions on Information Theory*, Vol. IT-29, No. 4, pp. 543-551, July 1983.

[3] Spilker, J. J. Jr. *Digital Communications by Satellite*, Prentice Hall, 1977, pp. 302-305.

[4] Cahn, C. R., "Improving Frequency Acquisition of a Costas Loop," *IEEE Transactions on Communications*, Vol. COM-25, No. 12, pp. 1453-1459, December 1977.

[5] Gardner, F. M., *Phaselock Techniques*, 2nd ed., John Wiley & Sons, 1979.