N91-11972

A Review of Fade Detection Techniques

F. J. Pergal

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

ABSTRACT Several proposed propagation fade detection techniques are reviewed in light of general requirements presented for beacon fade characterization. The discussion includes an analysis of phase lock versus frequency lock beacon tracking loops and of excess noise injection type radiometers. The ACTS beacon fade detection schemes proposed by COMSAT and JPL are examined along with the fade detection technique used by Harris in their ACTS LBR terminal.

1. Introduction

This paper is intended as a review and comparison of several proposed fade detection techniques as applied to the ACTS experiment. The discussion of specific implementations is preceded by a general analysis of beacon tracking receivers in which the ACTS propagation beacon parameters are referenced.

In the event that a fade is detected at a transmitting terminal, compensation procedures would be initiated. The eventual compensation algorithm will depend strongly on fade statistics which will only be compiled during the ACTS experiment itself. Fade compensation algorithms will not be discussed here.

The performance of phase and frequency locked receivers is compared for the ACTS beacon dynamics. Threshold tracking performance is displayed as a function of beacon phase noise and receiver C/N_0. A brief treatment of the associated radiometer design requirements is also included, anticipating the discussion of the proposed JPL beacon receiver.

COMSAT has designed a receiver for the ACTS Beacon Measurement System at the NASA Lewis Research Center which is based on the Hewlett Packard 3586A Selective Level Meter (SLM). The SLM is controlled via PC based hardware with algorithms of COMSAT's design. Reported performance of the system is excellent (Zaks, 1989). The suitability of the COMSAT approach for a VSAT architecture, as is proposed for the ACTS Propagation Experimenter's Terminal, is investigated in detail below.

A novel technique for fade detection has been developed by the Harris Corporation for use in their ACTS LBR terminal (Manning, 1989). The Harris technique uses the LBR signal to estimate the received SNR. The algorithm operates on relatively coarse (6 bit) samples of the LBR symbol matched filter output. Differences between the averaged outputs of absolute value and square law detectors applied to the matched filter output are exploited to generate very accurate estimates of the channel SNR. Fade compensation is based on this SNR measurement without the need for a separate beacon receiver.

In the proposed JPL receiver the beacon signal is digitized at IF in a relatively wide bandwidth of 2.5 MHz. IF sampling produces complex I/Q sample pairs which are phase rotated and then filtered. Coarse and then fine filtering is accomplished using a cascade of two windowed FFTs. A frequency discriminator is implemented using outputs of the second FFT and a beacon tracking loop is then closed on the phase rotator. Radiometer measurements can be made easily using existing signals within the digital receiver as discussed below.

2.0 General Beacon Receiver Design Requirements

Referring to the Proceedings of the First ACTS Propagation Workshop (Rogers, 1989), the fundamental requirement imposed on a propagation experimenter's terminal is the measurement of fade amplitude with +/-0.1 dB accuracy. Implicit in this requirement is the capability of determining the clear sky baseline path attenuation. Presumably, satellite EIRP variations will be resolvable by means other than measurements at isolated terminals.

For a VSAT type terminal the clear sky beacon C/N_o will be on the order of 40 to 45 dB Hz (Davarian et al., 1990). Signal dynamics include phase noise of -20 dBc/Hz at 5 Hz offset and diurnal frequency drift of up to +/-1.5 ppM (Davarian, loc. cit.). Assuming that the diurnal frequency drift profile is sinusoidal, for the uplink fade beacon at 27.5 GHz, maximum drift rates of 180 Hz/minute could be experienced.

Since only beacon amplitude measurements are required, some frequency tracking error is allowable. The amount of frequency error that can be tolerated will be determined by the beacon pre-detection bandwidth. The minimum predetection bandwidth is determined by the bandwidth of the fade process. The maximum bandwidth, on the other hand, will be determined by the measurement accuracy requirement and the minimum beacon signal level for which that accuracy is to be maintained.

Assume a generic receiver as shown in Figure 1. Referring to the figure, down conversion to baseband is optional; COMSAT, for example, detects the beacon signal at IF. With baseband detection, as shown in the figure, I and Q are low pass filtered and then accurately squared to produce a harmonic free I^2+Q^2 detection statistic. For a clear sky C/N_o of 42 dB Hz and a dynamic range of 15 dB, the minimum predetection C/N_o will be 27 dB Hz. Neglecting squaring loss, the input SNR corresponding to +/-0.1 dB detector accuracy at the bottom of the dynamic range will occur with a predetection lowpass bandwidth, B_{pd}, of 7.7 dB Hz.

Assuming a fade process bandwidth on the order of 5 Hz, this indicates that, in a sense, the ACTS transmitted EIRP is nearly optimal for a propagation experiment using the VSAT architecture. It should also be noted that the post detection noise spectrum will be triangular extending out to $2*B_{pd}$. To avoid noise aliasing either post detection filtering or an output sample rate of at least $4*B_{pd}$ should be used.

If a larger measurement tolerance is possible at the lower end of the dynamic range, or a bank of detectors could be used, then the predetection bandwidth could be opened up proportionately. In any case, the beacon tracking system must maintain the beacon within the predetection bandwidth over the dynamic range of interest. In addition, the tracking system should allow for easy initial acquisition or reacquisition after a long fade. The acquisition process should be fully automatic allowing for unattended operation and, in the case of the 20.2 GHz receiver, should be highly resistant to false locking on modulation sidelobes.

2.0.1 Phase Locked vs. Frequency Locked Tracking

Additive bandpass Gaussian noise and beacon phase noise will introduce angle jitter on the beacon tracking oscillator. In the case of additive noise, the narrower the loop bandwidth, the more accurately the receiver will track the beacon. To track out beacon phase noise, however, the loop should be as wide as possible. For the ACTS beacons, phase noise is significant enough that for a VSAT receiver using either phase or frequency locked tracking, loop bandwidth has to be carefully chosen. The best choice minimizes the total angle tracking error variance.

The contribution to tracking phase jitter variance due to additive noise is $[C/N_oB_L]^{-1}$ radians^2 for the phase locked loop. For the frequency locked loop the situation is more complicated (Natali, 1984) and depends on the predetection SNR and on details of the discriminator design. For a predetection SNR >> 1 and B_L << B_I, where B_I is the equivalent predetection IF bandwidth, assuming a cross product discriminator with integrate and dump filtering, the frequency tracking error variance is given by:

$$[0.051(B_L/B_I)B_I^2]/SNR_I$$

As an example, the VPI Olympus 12.5 GHz beacon receiver frequency locked loop has a noise bandwidth of approximately 10 Hz with a predetection bandwidth of 200 Hz. For a C/N_o of 30 dB Hz, the above expression gives an RMS frequency tracking error of roughly 4.5 Hz.

In the case of beacon phase noise the tracking angle error variance is given by the integrated product of the linearized loop error response and the beacon angle noise power spectral density (Gardner, 1979). For the range of B_L in question, the beacon phase noise power spectral density in dBc/Hz has close to a pure flicker of frequency characteristic. The phase noise power spectral density, $S_{phase}(f)$, will equal h_3/f^3, where f is the frequency offset from the beacon carrier frequency and h_3 is a proportionality constant. Expressed as frequency noise the beacon angle modulation PSD will simply be h_3/f, i.e. 1/f or "flicker" noise. For the ACTS beacons h_3 is approximately equal to 1.25.

The total angle error variance will be the sum of the two component variances due to additive and phase noise. For the phase locked loop with a sinusoidal phase detector the commonly accepted threshold corresponds to an RMS phase error of .316 radians. For the frequency locked loop the situation is again slightly more complicated. The frequency error that can be tolerated depends on details of the discriminator design (Natali, 1984).

Generally, the discriminator range will be from $+/-0.2B_I$ to $+/-0.5B_I$, where B_I is the effective prediscriminator IF bandwidth. Allowing for three sigma excursions, the RMS frequency error that can be tolerated will be approximately one third of the discriminator range. The larger the range the greater the frequency error that can be tolerated before thresholding occurs.

There are of course constraints on the discriminator range. Usually, discriminator action is obtained by comparing the phase of a signal with that of a delayed version of itself, i.e. by delay and multiply operations or by differencing the power measured in two frequency offset filters. The multiplication and power measurement operations introduce squaring loss. In order to avoid significant squaring loss in the loop, the effective prediscriminator bandwidth cannot be too large.

For both the phase locked and frequency locked loops there are also practical constraints on the predetection bandwidth. These arise because of component tolerances and voltage offsets. About the best that can be said is that the optimal tracking strategy and parameter settings have to be determined on a case by case basis.

2.0.2 Radiometer Design Considerations

To meet the fade measurement accuracy requirement, for the propagation terminal, a baseline clear sky path attenuation must be established. An excess noise injection radiometer such as used in VPI&SU's Olympus terminal suite (Stutzman, 1990) and shown in block diagram form in Figure 3. is one means to this end. Referring to the figure, a directional coupler at the receiver front end is alternately switched at a low rate (dwell times of two to five seconds would be typical) between a diode excess noise source and a termination.

Assume that the diode excess noise ratio (ENR) and the physical temperature of the switch, coupler and any other connecting hardware are well regulated. This scheme then will provide constant differential noise power against which to measure changes in the system noise temperature (and receiver gain).

Loss mechanisms in the excess noise path will make a contribution to system noise temperature; but, it is possible to keep these losses in the range of a few dB. Components are mounted on a temperature regulated plate in an effort to stabilize the differential noise temperature. In addition, the system would also be periodically calibrated against hot/cold loads.

Elektronik Centrallen has cited (1988) the results of a two year life test of 35 GHz noise diodes with ENRs on the order of 23 dB which shows ENR variations of roughly +/-.03 dB. From the data records it appears that much of this small variation may even be due to seasonal changes in the test environment.

The radiometer is used to measure differences in system noise temperature, the varying sky noise temperature being the component of system noise that is of interest. At any particular system noise temperature, the accuracy of an absolute system noise temperature measurement is given by $T_{sys}/(BT)^{1/2}$, where B is the equivalent predetection IF noise bandwidth and T is the measurement integration time. Differences in absolute noise measurements will have an accuracy of $(2)^{1/2}$ times this number.

The following heuristic argument is intended as motivation for choosing the excess noise injection radiometer IF bandwidth. The detected excess injected noise provides the "yardstick" against which changes in absolute system noise temperature are measured. The excess noise should therefore be as large, and as accurately determined as practical. Along these lines, if changes in receiver gain are very slow, long term averaging of the detected excess noise differences would remove most of the random measurement inaccuracy. For a system noise temperature, T_{sys}, of 1500 ^{O}K, and an integration time of 3 seconds, a 1 ^{O}K accuracy could be produced with an IF bandwidth of 2.5 MHz.

2.1 The Harris Algorithm and Implementation

The Harris fade detection technique operates on the coded LBR signal and makes use of the output of the receiver's coherent amplitude detector. The input signal is MSK with additive bandpass noise. The receiver signal which is processed is the output of a baseband matched filter. The density of the filter output at the sampling instants is a zero mean Gaussian density convolved with delta functions at +/-A, where A is the signal amplitude at the matched filter output.

The quantized matched filter output is operated on by two parallel detector stages. One stage is a square law detector, the other an absolute value detector. Following each detector operation are two identical accumulation and scaling stages. The processing is illustrated in Figure 3.

For sufficiently high input SNRs, the two Gaussian components (centered at +/-A) of the matched filter output density will have very little overlap. Under these conditions the outputs of the absolute value and square law detector chains are proportional to A and $(A^2 + s^2)$, respectively, where s^2 is the variance of the Gaussian components of the matched filter output density.

The accumulated absolute value detector output is squared to produce an estimate of A^2. The ratio, $A^2/(A^2 + s^2)$, is then formed and used to drive a lookup table. Since this ratio is a monotonic function of SNR (i.e. of A^2/s^2), it easy to compute true SNR via the lookup. At the same time by modifying the lookup table the overlap effect mentioned above can be compensated. The result is an SNR estimate that is accurate over a wide dynamic range.

The number of baseband matched filter samples accumulated for an SNR estimate is 640. Further a minimum of 900 of these ratios (conditioned by the lookup) are averaged to generate the final SNR estimate used by the fade compensator. The noise densities at the outputs of the square law and absolute value detectors will be similar (Blachman, 1966). For input SNRs (in the matched filter bandwidth) of 6 dB or better there will be very little squaring loss so that the variance of the resulting SNR estimate can be easily calculated.

The main advantage of the Harris technique from a receiver design viewpoint is that it is independent of AGC effects. There is also the potential for excellent dynamic range. This will to some extent depend on Bit Synchronizer performance degradations with low SNRs and the resulting inaccuracies in estimated symbol amplitude. Finally the Harris algorithm as implemented does not require an independent beacon receiver which can amount to a large cost savings.

2.2 The COMSAT Fade Detection Receiver

The COMSAT fade detection receiver operates on the ACTS uplink and/or downlink fade beacons. In the fully configured Beacon Measurement System three Hewlett Packard 3586A Selective Level Meters (SLM) are used to track and measure three independent beacons simultaneously. For the BMS, the clear sky C/N_o is 62.1 dB Hz for the 20.2 GHz beacons and 58.5 dB Hz for the 27.5 GHz beacon.

Referring to Figure 4., the down converted beacon signal is input to the SLM. The signal is frequency converted to a 15.625 KHz IF and band pass filtered to a selectable IF bandwidth. The filtered signal is RMS/log detected with a detector time constant of .011 seconds and the conditioned, detected signal is sampled at a 6&2/3 Hz rate. A self calibration feature is also available in the instrument under bus control.

Receiver tuning information is derived from the bandpass limited 15.625 KHz IF signal. The limited signal drives a second order phase locked loop with B_L of approximately 600Hz (Hewlett Packard, 1990). The phase locked loop VCO runs at 20 times the IF signal frequency. The VCO frequency is counted over a very accurate 0.5 second time interval to obtain the average signal frequency reading. Frequency tracking is essentially open loop with the first LO being set via bus commands from an algorithm of COMSAT's design. The 0.1 Hz frequency resolution and long term stability available with the external clock option allow extremely accurate tuning under bus control.

For a VSAT architecture C/N_o is approximately 15 dB less than what COMSAT is working with in the BMS. To track a VSAT beacon signal the IF bandwidth would probably have to be lowered to 20 Hz. At that IF bandwidth the detector time constant might become a problem. Furthermore, beacon tracking in the open loop mode may require developing a sophisticated frequency estimation algorithm to assure quick reacquisition after fades. In summary, it is probably safe to say that fairly extensive testing would be required to ascertain the suitability of the COMSAT approach in the VSAT application.

2.3 The JPL Digital Beacon Receiver

A block diagram of the JPL Beacon Receiver is shown in Figure 5. Relatively coarse quantization is used to sample the signal at a nominal IF of 20 MHz. Note that the ratio of the nominal IF center frequency to the sampling clock frequency is of the form:(1 + 1/4).

The choice of IF and sampling frequencies allows I and Q relative to a reference sinusoid at the nominal IF frequency to be reconstructed from the IF samples. In the general case, the same will be true if $f_{IF}=f_s*(N +/- 1/4)$, where f_{IF} is the IF frequency, f_s is the sampling clock frequency and N is any integer. The only restriction on f_s is that it be at least twice the band width of the sampled IF signal.

It has been known for some time that a high degree of measurement accuracy is possible with coarse quantization in the presence of a dithering signal (Widrow, 1955). It is also well known that, in general, the quantizing levels have to be extremely accurate. Blachman (1984) has shown that for a memoryless quantizer operating on a noisy signal, the effective quantizer transfer function will be the noise free transfer function convolved with the noise density. The noise essentially smooths over A/D nonlinearities. By operating at a low enough SNR (<-20 dB), the A/D step accuracy problem is avoided in the JPL receiver. The quantization operation is also effectively memoryless because of the independence of the signal and noise.

The A/D output sequence is sorted into an I/Q sample pair stream, then phase rotated using a complex mixer (a look up table) and formatted into blocks. A (complex) FFT is then performed on blocks of I/Q sample pairs. The order of the FFT is determined by acquisition time requirements as discussed below. The FFT operation is equivalent to passing the signal through a filter bank. The time sequences produced by the FFT bins are then detected and filtered. Following which the output from the FFT bin containing the carrier component is identified by the controlling processor.

As part of the acquisition process a fine grained filter bank operation locates the carrier more precisely. This filtering operation is performed by a second or fine FFT taken on the sample stream corresponding to that bin from the first or coarse FFT operation that contains the carrier component.

Carrier tracking is accomplished using an AFC loop closed on the complex mixer through a phase accumulator. The accumulator increment is controlled by an FLL where the discriminator function is implemented using the outputs of the fine FFT (Natali, 1984). Updates to the loop take place at a relatively low rate corresponding to roughly ten times the loop bandwidth.

During acquisition the AFC loop is disabled. In the tracking mode the carrier is downconverted to a frequency close to DC by the complex mixing operation. The loop will keep the carrier centered in a relatively narrow spectral window. The downconverted signal is amplitude detected and filtered in the controlling processor and passed to the terminal computer. The processor will be designed to support a 10 Hz carrier amplitude sampling rate.

The radiometer and signal processing channels share a common IF. The coarse FFT bins <u>not</u> containing beacon, telemetry or ranging components are detected, summed and filtered to estimate the channel noise variance. Currently simulations are being run to test the effectiveness of FFT windowing using coarse A/D quantization.

With an IF filtered to a noise bandwidth of 2.5 MHz as above, the (complex) I/Q process will have a bandwidth (assuming 4 MHz wide filter skirts) of about 2 MHz. With a 16 MHz sampling clock, the process will be oversampled by a factor of 2. With the above choice of IF and sampling clock frequencies, the sample sequence forms a pattern: I, Q, -I, -Q, I, Q, etc. Therefore, samples 2 and 3 modulo 4 can be discarded.

Using a pair of United Technologies UT69532 IQMAC Processor chips clocked at 16 MHz, a 1024 point complex FFT will run in 0.320

milliseconds; 4 chips will do the FFT in .160 milliseconds. Using 1024 points for the coarse FFT channelizes the I/Q stream into bins with noise bandwidth of roughly 4 KHz.

For a C/N_0 of 43 dB Hz this provides a predetection SNR of about 7 dB even without the fine FFT processing. Since acquisition consists in measuring the relatively slowly changing power of the bin outputs, subsampling the outputs is possible. Needless to say very fast acquisition and reacquisition (in the sub-second range) following deep fades becomes possible.

Experience with Olympus indicates that similar tracking loop parameters and carrier amplitude processing will be appropriate for ACTS signal processing. Loop noise bandwidths on the order of 2 to 3 Hz would enable tracking into the 15 dB C/N_0 region. Note that at peak carrier frequency drift rates (180 Hz/minute) some type of loop aiding would be required for such narrow loop bandwidths.

3.0 Conclusion

Three specific fade detection techniques developed for the ACTS program have been examined. A comparison of the techniques points out the advantages of each in its particular application. The COMSAT design is a very clean implementation for the case where C/N_0 is relatively high. Its applicability for a VSAT propagation terminal architecture, however, would have to be carefully investigated.

The Harris receiver only requires a nominal AGC to keep the signal level reasonably positioned in the A/D converter input range. In this sense the Harris approach has a real advantage. The Harris technique could also be adapted to a beacon receiver but then radiometer information is given up. With this the ability to sense clear sky attenuation independent of satellite EIRP variations is also lost.

The JPL receiver has the advantage of simultaneously measuring channel noise along with beacon amplitude. This minimizes the effects of receiver gain variations on the beacon amplitude measurement. All tuning is done digitally after A/D conversion, so that the need for a selective IF along with a tuned LO is eliminated. The all digital implementation is also attractive from the standpoint of trouble free operation in unattended terminals.

Acknowledgement: The research described in this paper was performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

References

Blachman, N. M., Noise and its Effects on Communication, p. 93, New York, McGraw Hill, 1966.

Blachman, N. M., "The Effect of a Nonlinearity upon Signals in the Presence of Noise", IEEE Trans. COM-21, pp. 152-154, February 1973.

Davarian, F., F. Pergal, D. Chakraborty and W. Stutzman, "ACTS Propagation Terminal Prototype Planning and Design", Proceedings of NAPEX XIV, May 1990.

Elektronik Centrallen, "20/30 GHz Multi-Frequency Radiometer Feasibility Study", Horsholm, Denmark, May 1988.

Gardner, F. M., Phaselock Techniques, New York, McGraw Hill, 2nd ed., 1979.

Hewlett Packard, Private communications with personnel at the HP Lake Stevens Instrument Division.

Manning, R. M., "The Harris 'Mean Squared over Variance' Technique", Proceedings of the First ACTS Propagation Studies Workshop, December 1989.

Natali, F. D., "AFC Tracking Algorithms", IEEE Trans. COM-32, pp. 935-947, August 1984.

Rogers, D. V., "Report of the Systems Group", Proceedings of the First ACTS Propagation Studies Workshop, December 1989.

Stutzman, W. L., "Olympus Propagation Terminal Hardware and Experiments", Proceedings of NAPEX XIV, May 1990.

Widrow, B., "A Study of Rough Amplitude Quantization by Means of Nyquist Sampling Theory", IRE Trans. CT-5, pp. 266-276, December 1956.

Zaks, C., "NASA Ground Station Beacon Measurement Subsystem", Proceedings of the First ACTS Propagation Studies Workshop, December 1989.

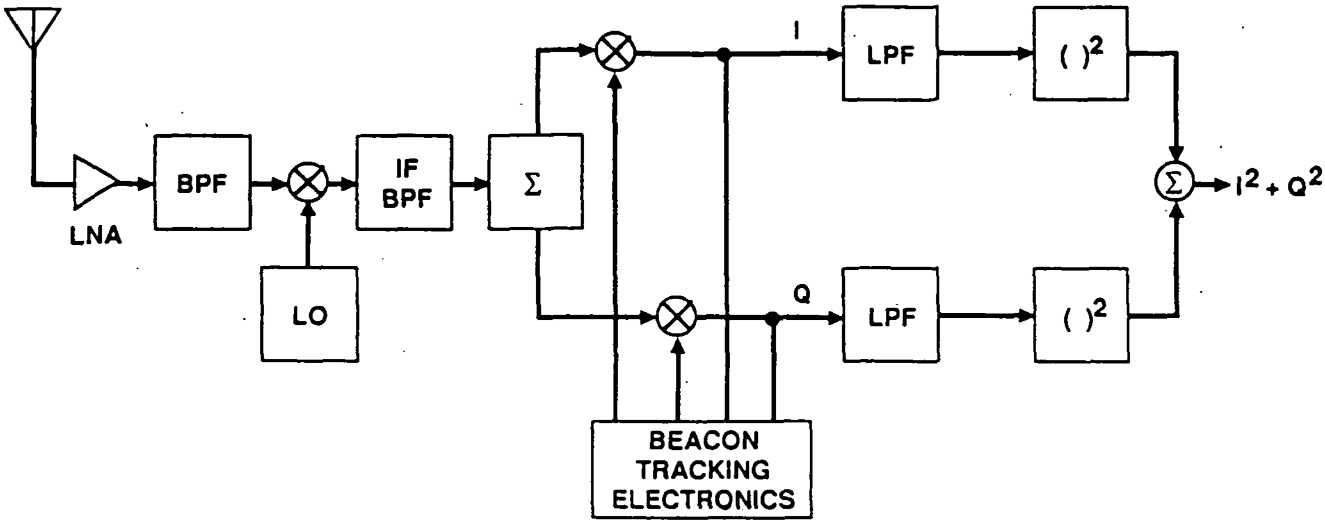

Figure 1. Generic Beacon Receiver Block Diagram

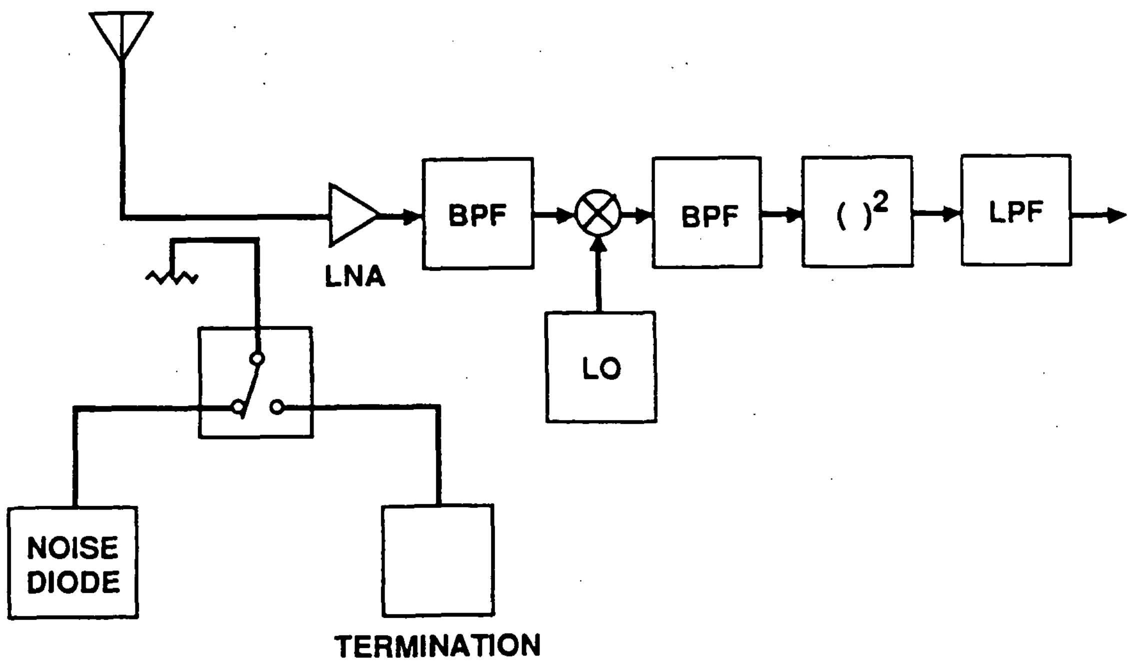

Figure 2. Excess Noise Injection Radiometer Block Diagram.

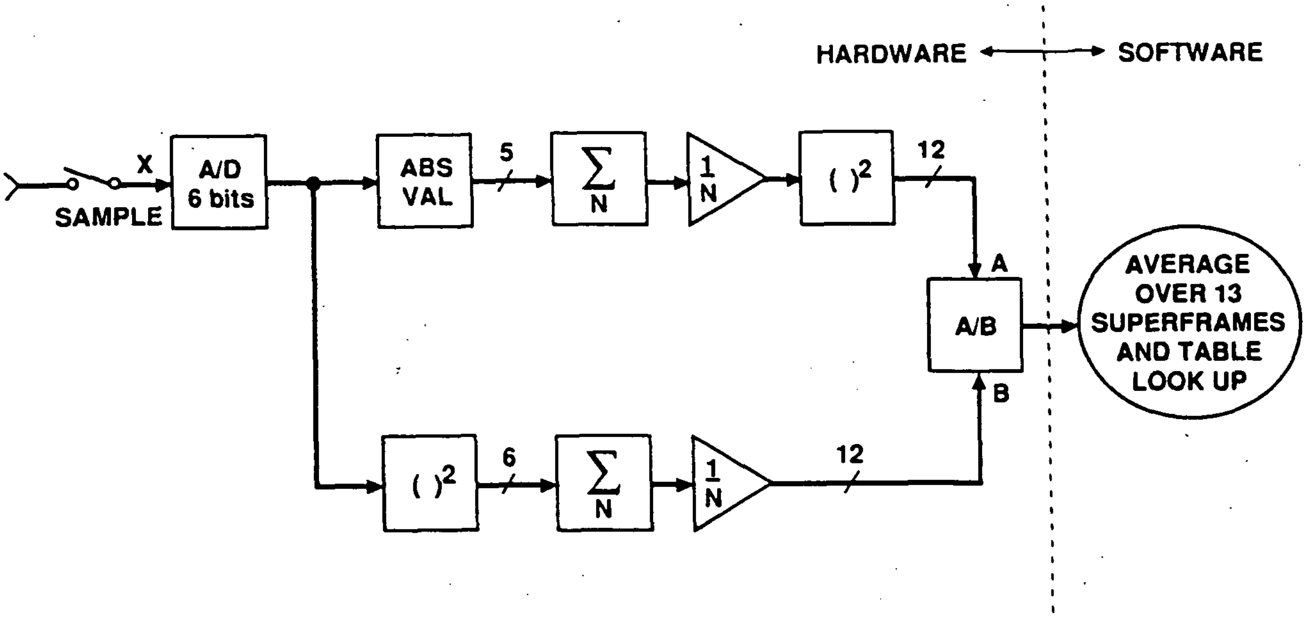

Figure 3. The Harris LBR Fade Detection Technique.

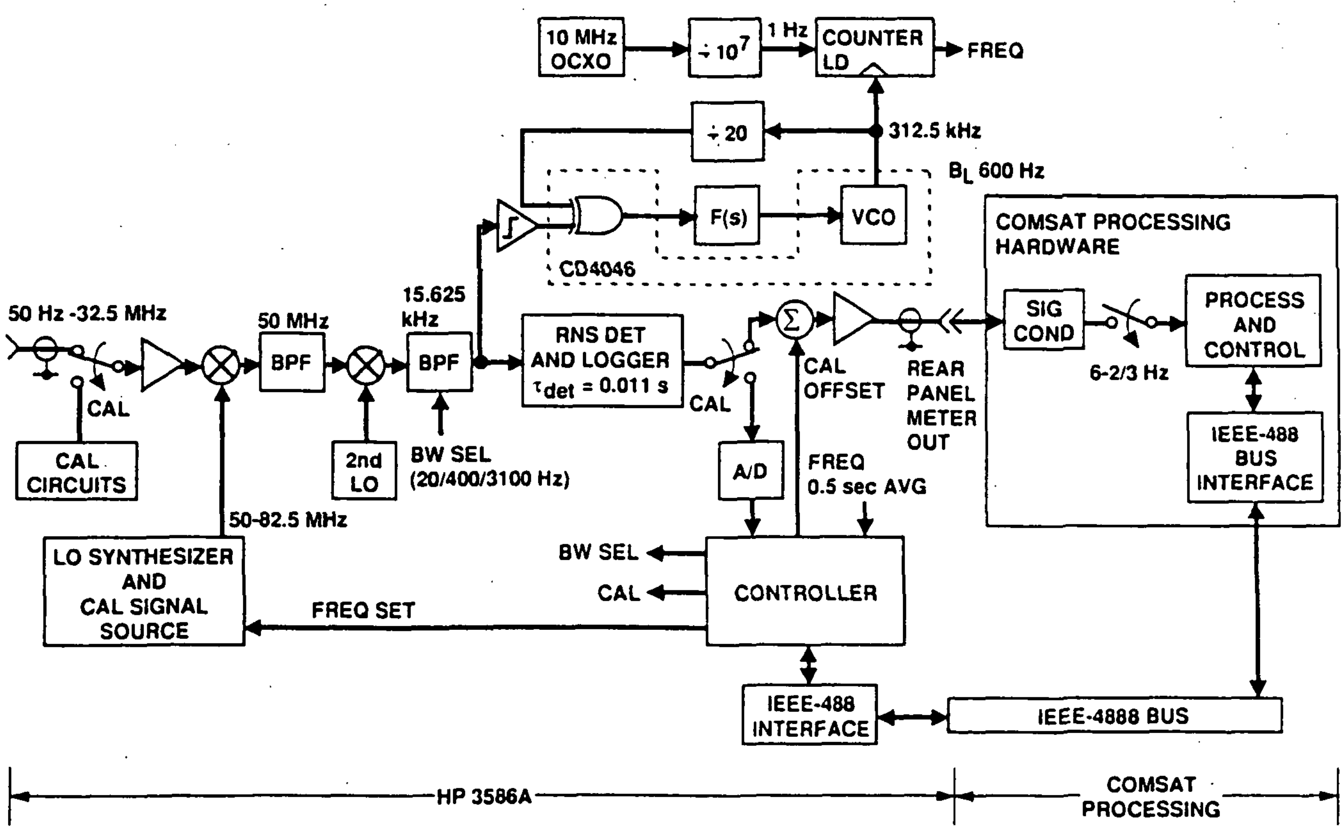

Figure 4. The COMSAT Fade Detection Technique.

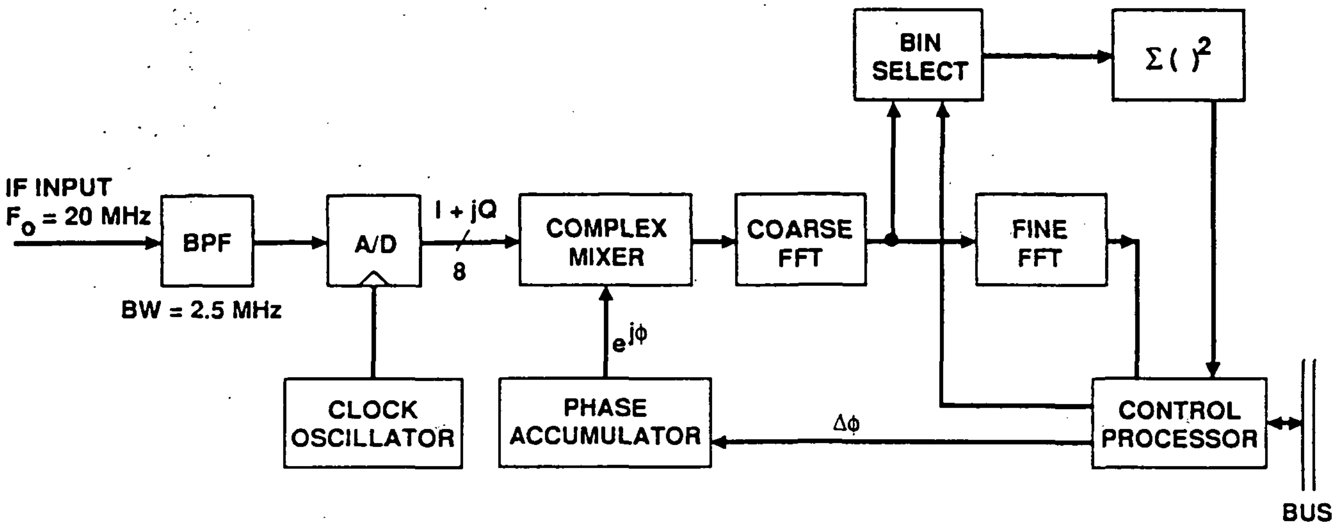

Figure 5. The JPL Digital Beacon Receiver.