

HIGH RATE KA-BAND DOWNLINK DIGITAL RECEIVER FOR MUOS

Arthur P. Helwig and Bin Hu
General Dynamics C4 Systems
Mail Drop H2232
Scottsdale, AZ 85257

ABSTRACT

The Mobile User Objective System (MUOS) will provide global SATCOM for the DOD using WCDMA technology. A key to achieving this is the "virtual" location of the radio base station (RBS) antennas on 4 geosynchronous satellites. The transport of the user-to-base (U2B) and base-to-user (B2U) signals between these antennas and the RBSs located on the ground is provided by Ka-band uplink/downlinks. Each Ka-band downlink transports samples for 32 WCDMA beam-carriers via three 8-PSK signals at a rate of 2457 Mbits/sec. This downlink is processed on the ground by an FPGA-based digital receiver to recover the samples and reconstruct the WCDMA uplink signals. The high data rate of this demodulator coupled with the constraints of cost, BER, availability and power consumption drives a challenging design solution. This paper discusses a number of features of this design including high speed processing architecture, hybrid analog/digital timing recovery loop, carrier recovery loop (CRL), equalizer, adaptive cross-pol interference canceller (AXPIC), frame synchronizer and turbo decoder.

INTRODUCTION

The MUOS U2B Ka-band downlink provides the virtual connection of the RBS receivers to their respective "antennas" located on the MUOS spacecraft. This connection transports digitized WCDMA beam-carriers that are sampled at the spacecraft's UHF receivers. The spacecraft's antennas include 16 beams, each covering 4 WCDMA uplink frequencies between 300 and 320 MHz. This results in a total of 64 beam-carrier sample streams that must be transported to the ground. The MUOS system includes 4 spacecrafts, each linked to 2 of 4 radio access facilities (RAF's) located around the world. Thus, each Ka-band downlink must transport samples for 32 beam-carriers. The MUOS system definition calls for this data to be sent across three 8-PSK carriers each operating at a symbol rate of 384 MHz. In trading the spectrum, power and spacecraft constraints an interface was defined that uses a rate 0.711 turbo code and Hadamard processing. Hadamard

processing reduces the required data rate by allowing the dynamic range of the downlink to be shared across all 32 beam-carriers [1]. To accomplish this, the 32 beam carriers are Hadamard transformed prior to encoding and modulation formatting in the spacecraft. This results in 32 Hadamard channels that are transported over the Ka-band downlink. In the ground modem the inverse Hadamard operation is applied to the decoded bit streams to recover the original beam-carrier sample streams. The resulting 8-PSK carriers are located between 20.2 to 21.2 GHz. Right and left hand polarization channels are used to keep the carriers within a 1 GHz band. The spectral locations of the carriers are illustrated in Figure 1.




Figure 1 - Spectral Location of RHCP and LHCP 8-PSK Carriers

Experience with applying turbo codes in wide band communication systems has shown they can be vulnerable to low frequency (correlated) residual phase noise in the demodulator. This is reduced by maintaining as high a bandwidth as possible in the carrier recovery loop, but the flexibility to accomplish this is bounded by other considerations such as link SNR. In order to mitigate this concern, interleaving of the turbo code blocks is included in the MUOS U2B downlink format. The 1024 symbol turbo code blocks are interleaved by 16. This results in frame that is 42.7 us long as illustrated in Figure 2. Each turbo code block includes a 5 symbol (15 bit) unique word.




Figure 2 - Turbo Code Block and Frame Format

GROUND SYSTEM ARCHITECTURE

The digital receiver discussed in this paper is one part of an entire ground receiver system. The other parts of this system are described briefly here as their performance characteristics must be considered in the design and analysis of the digital receiver. Refer to Figure 3 during the following description.

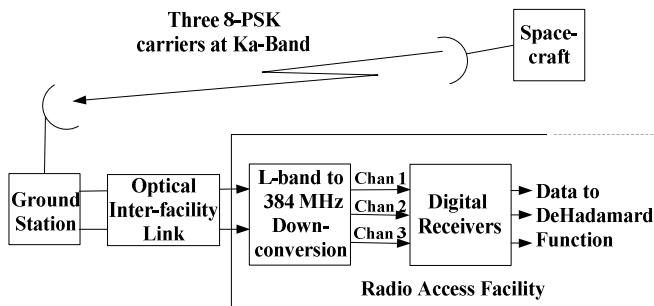


Figure 3 - Ground System Architecture

As previously mentioned, samples of 32 beam carriers taken at the spacecraft UHF receivers are transported to the ground via three 8-PSK carriers. Two of these carriers are contained in the right hand polarization (RHCP) and the remaining one is carried in the LHCP. The ground antenna feed separates these polarizations into two separate signals and downconverts them to L-band. At L-band the RHCP contains one carrier at 1.226 GHz and one carrier at 1.774 GHz; the LHCP contains one carrier at 1.774 GHz. Due to the long distance (up to 1000 feet) from the ground station to the RAF building containing the digital receivers, an optical interface is used to transport the LHCP and RHCP signal between the two locations. At the RAF building each carrier is bandpass filtered and Downconverted to a center frequency of 384 MHz. The resulting Channel 1, 2 and 3 signals are then passed to the digital receivers.

DIGITAL RECEIVER IMPLEMENTATION

Given the previous background and architecture discussion, some key digital receiver features and requirements are established. The receiver is actually

three identical receivers whose outputs must be synchronized in order to accomplish the de-Hadamard operation that reconstitutes the beam carrier samples taken on the spacecraft. This synchronization is in addition to the other usual synchronization operations (symbol sync, carrier recovery, and frame sync) required for demodulation. Since two of the three 8-PSK carriers are transported on the same carrier frequency at K-band but on different polarizations, adaptive cross-pol interference cancellation (AXPIC) is a required function. This is driven by cross-pol interference that occurs during fade events in the k-band channel. From a cost perspective, it is highly desirable to use "commercial-off-the-shelf" (COTS) FPGA boards. For the downlink receiver a VME-based FPGA board product has been selected. While this approach offers flexibility, the power dissipation and thus the number of FPGA boards per VME chassis is limited. This requires trade-offs against the level of digital signal processing that can be attained and the benefits this affords.

In view of the above considerations, the digital receiver architecture, shown in Figure 4 was established.

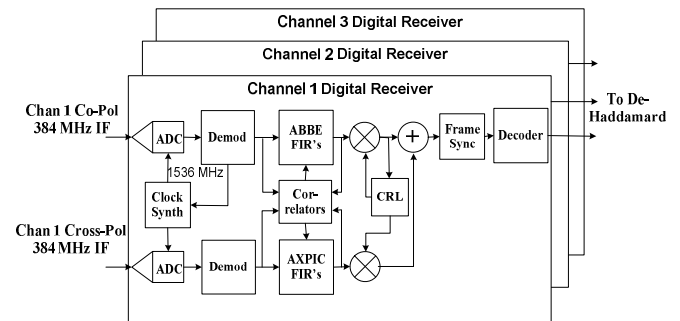


Figure 4 - Digital Receiver Architecture

The receiver consists of co-pol and cross-pol ADC's and demodulators, followed by an adaptive baseband equalizer (ABBE), adaptive cross-pol interference canceller (AXPIC), carrier recovery loop (CRL), frame synchronizer and decoder. Each receiver output feeds the de-Hadamard function. The COTS FPGA boards are able to operate at clock rates as high as about 150 MHz. However, at this rate design compilations become very difficult. The signal processing is, therefore, parallelized by a factor of four. This results in an average processing rate of 1/4 the symbol rate or 96 MHz.

Demodulator

A block diagram of the demodulator is shown in Figure 5.

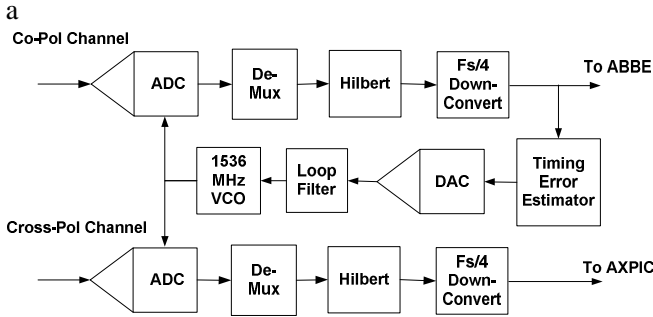


Figure 5 - MUOS Downlink Demodulator

A key feature of this design is the use of coherent sampling of the 384 MHz IF input signal. In a digital receiver, the sampling frequency can be independent of the symbol rate, and the optimum sampling time can be found using digital interpolation techniques [2] such as Farrow filtering [3]. However, the parallelized processing required for this design makes decimation more complicated and drives up the FPGA resource requirements. In this design, the signal is sampled coherently at 4 times the symbol rate, or 1536 MHz. The ADC output is de-multiplexed by a factor of 16 which reduces the average processing rate to 96 MHz. The Hilbert transform that follows creates a complex signal and suppresses the negative frequency component of the input signal by at 40 dB across the band of interest. This requires a 19 tap FIR filter at 4Rs (4x Symbol Rate), but due to the odd symmetry and the alternate nulls of the Hilbert transform impulse response, only 5 multipliers are required for implementation.

In the co-pol path, 6 Hilbert transform outputs are computed for each block of 16 samples out of the demux. Four of the 6 samples are used for the subsequent one sample per symbol processing. The two additional outputs are mid-symbol values that are required for the symbol timing estimator and associated lock detector. In the cross-pol path, which runs at two times symbol rate, 8 Hilbert transform outputs are computed for each block of 16 samples from the demux. The total number of multipliers required for Hilbert transform operations is $(6+8)*5 = 70$.

Following the Hilbert the signal is downconverted to baseband using $F_s/4$ downconversion. Since the downconversion LO is 1/4 of the 1536 MHz sample rate, the required operations involve simple switching and inversion operations and involve no multipliers. The downconverter output is then passed on to the ABBE/AXPIC and to the timing error estimator. This estimator uses mid symbol and on-time samples to develop its output and is described in reference [4].

A feature of this design is the hybrid analog/digital clock recovery loop. In this design all components from the ADC through the timing error estimator [4] are digital. The timing error estimator output then drives a DAC which provides an analog drive to the loop filter and 1536 MHz VCO. The loop design is 2nd order type 2 and tracks symbol timing with minimum static phase error in the ADC sample clock. The expected effective timing jitter has been analyzed using 2nd order loop analysis that accounts for jitter contributions from the transmitter clock, the VCO and the variance of the timing error estimator. The predicted jitter is less than 0.0056 cycles rms (2 degrees rms). This will have a negligible impact on the decoder performance which is specified to be 10^{-5} BER at 9.2 dB SNR.

ABBE and AXPIC

The demodulator output drives the ABBE and AXPIC functions. Typically equalizers are used to mitigate the effect of over-the-air channel impairments, but for the MUOS downlink receiver, the equalizer serves other needs. For this system, the equalizer in conjunction with the bandpass filter in the L-band to 384 MHz downconverter provides the matched filter for the demodulator. This approach is due to the fact that the equalizer is operating at one sample per symbol (T-spaced). Because of this, it is not able to limit the noise bandwidth of the demodulator path. Limiting noise bandwidth, therefore, is the purpose of the bandpass filter noted above. However, the equalizer is able to shape the frequency response inside of the $\pm R_s/2$. (R_s = symbol rate). In this offset frequency range it is able to compensate for undesired filtering effects in the ground and space system hardware.

A block diagram of the ABBE and AXPIC are shown in Figure 6. In order to compensate for asymmetric frequency response effects, the equalizer has a complex FIR structure that includes cross I/Q branch filters. Due to FPGA resource limits, each of the four FIR filters is 13 taps long. Because of the parallel processing of 4 samples wide at a 96 MHz rate, the ABBE FIR filters use a polyphase implementation and require a total of 256 multipliers. The purpose of the AXPIC FIR's is to develop an estimate of the cross-pol signal at the sampling instant in the co-pol channel. A 3-tap FIR filter is sufficient for application as it is not necessary to estimate ISI in the cross-pol channel. In order to allow the AXPIC to adapt to timing differences between the co-pol and cross-pol channels, T/2 spacing is used for the FIR filters. The AXPIC FIR's also use a polyphase implementation and require a total of 48 multipliers.

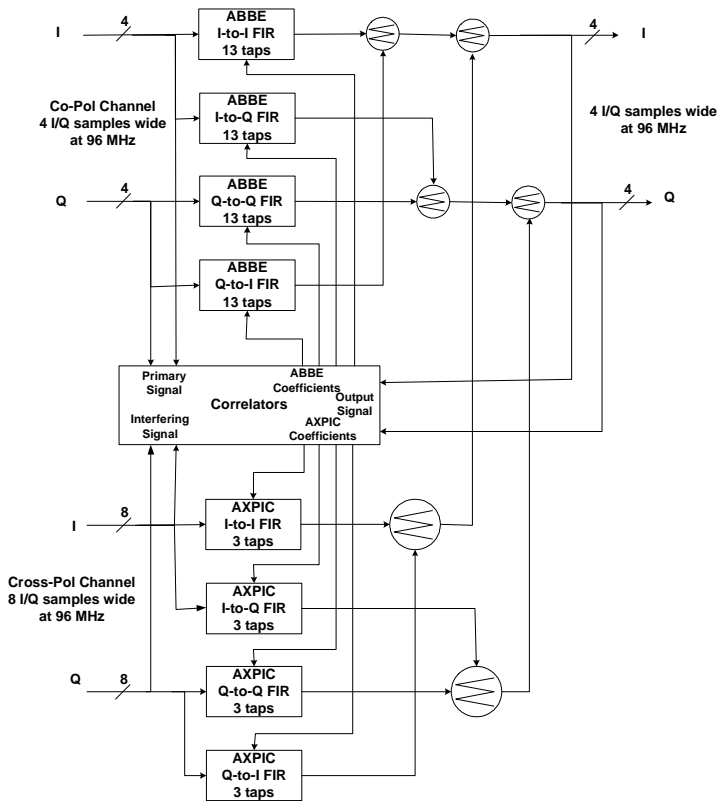


Figure 6 Block Diagram of ABBE and AXPIC

The FIR filter coefficients for the ABBE are adapted using a least mean squared (LMS) update algorithm which is illustrated in Figure 7. A hard decision is made on each symbol, and an error signal, e , is formed by subtracting the actual symbol from the symbol decision. The FIR coefficient vector, c , is then updated by correlating the input signal, x , with the error signal as follows:

$$c(n+1) = c(n) + \mu * x(n) * e(n) \quad (1)$$

μ is referred to as the convergence factor and is set to optimize convergence time and final mean squared error (MSE) of the output constellation. The coefficient adaptation for the AXPIC also uses LMS.

Carrier Recovery Loop (CRL)

Following the ABBE/AXPIC, the CRL tracks out any residual frequency error and phase variations. Carrier recovery can be accomplished by open loop phase estimation on blocks of data or pilot symbols or by a continuously tracking phase locked loop. Block phase estimation is not feasible with the limited FPGA resources. A very short unique word (5 symbols out of 1024 symbol turbo code block) has been allocated; so phase estimation using the unique word as a pilot

symbol is not deemed feasible. Since the MUOS k-band downlink signals are continuous, a continuously tracking phase locked loop is the chosen approach.




Figure 7 ABBE Coefficient Update Algorithm

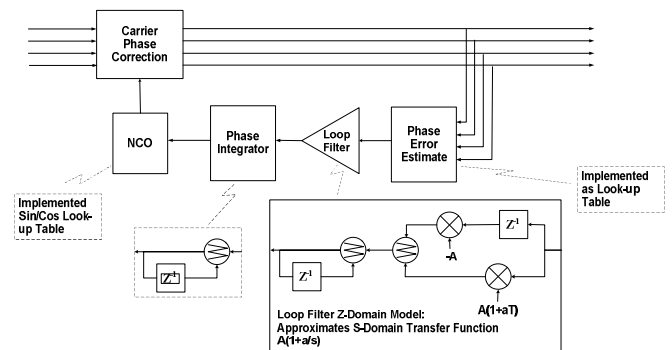


Figure 8 - Block Diagram of Carrier Recovery Loop

The CRL is illustrated in Figure 8. It uses an eighth order phase error estimator of the following form:

$$PD(n) = \text{IMAG} \left[\frac{(I(n) + jQ(n))^8}{(I(n)^2 + Q(n)^2)^4} \right] \quad (2)$$

This phase error estimator is discussed in reference [5]. The estimator is implemented as a look-up table (LUT). Using the LUT, a phase error estimate is computed for each symbol. Since the design is parallelized by 4, 4 outputs are computed at a 96 MHz rate. Each group of four outputs are summed together and then passed on to the loop filter. The rate at which this sum is being computed is much greater than the CRL loop bandwidth (96 MHz versus ~100 KHz), so the effect of the summing operation on the loop response is negligible. The CRL loop filter is of the form $A(1+a/s)$ in the s -domain which results in a 2nd order type 2 loop. This produces negligible static phase error due to Doppler frequency error which is less than one part in 10^{-7} or about 2 KHz. The loop filter output is used to index into a sin/cos lookup table which then provides the phase correction for the input signal.

The residual phase noise in the demodulator has been of particular interest in this design, and a summary of the

results of its analysis are presented here. As mentioned previously, one purpose of the CRL is to track (within its bandwidth) the phase jitter of the demodulator input signal. This jitter is due to phase noise contributions from different sources including the transmitter LO, the receiver LO's, and the receiver ADC sample clock¹. The use of an optical interface between the ground station antenna and RAF building is another opportunity to add phase noise, but in the MUOS system this is a negligible contributor. Within its loop bandwidth, the CRL tracks the composite phase jitter imparted by the various contributors. Outside the loop bandwidth, the jitter is not tracked and contributes to residual phase noise in the demodulator. For a CRL closed loop response of $H(f)$, and a phase noise power spectrum on the demodulator input signal of $S_{LO}(f)$, the power spectrum of the resulting residual phase noise can be written as:

$$S_{RESd}(f) = S_{LO}(f)[1 - H(f)] \quad (3)$$

The expression $[1-H(f)]$ is sometimes referred to as the error response of the loop. The predicted SSB residual phase noise contribution due to the composite LO phase noise is shown in Figure 9 for the case of a 100 KHz 3dB bandwidth in the closed loop response and a loop damping factor of 1. The error response is high-pass, and the input phase noise power spectrum decreases with increasing frequency; thus, the resulting residual phase noise spectrum increases, peaks, and decreases with increasing frequency.

The CRL is required to operate down to an effective E_s/N_0 of about 9.2 dB². Figure 10 shows the predicted spectral contribution of the phase estimator variance for this case assuming a 100 KHz 3dB bandwidth in the closed loop response and a loop damping factor of 1.

Given the contributions of the input phase jitter and the phase estimator variance, it is desirable to find the optimum CRL loop bandwidth. Figure 11 plots the integrated phase noise versus loop bandwidth for the composite as well as the individual contributors. It shows that minimum phase noise of about 3.8 degrees occurs for a loop bandwidth of 40 to 100 KHz. Because the phase estimator gain is a function of the E_s/N_0 , the CRL loop gain will be compensated versus E_s/N_0 in

order to maintain the CRL loop bandwidth (and thus residual phase noise) in an acceptable range.




Figure 9 Residual Phase Noise Contribution Due to Phase Jitter at the Demodulator Input




Figure 10 - Residual Phase Noise Contribution Due to Phase Estimator Variance

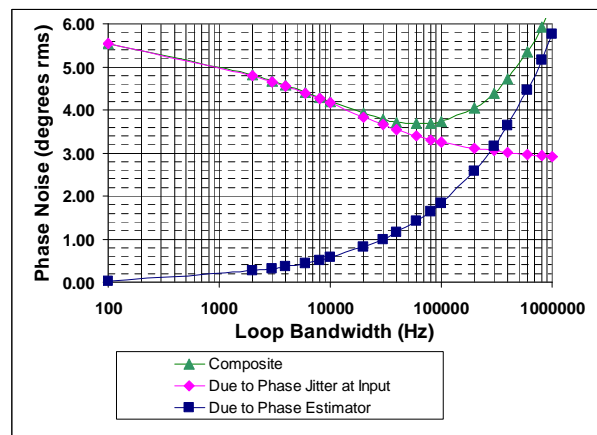


Figure 11 Residual Phase Noise versus CRL Loop Bandwidth

¹ The phase noise contribution of the ADC sample clock must be considered due to the use of IF sampling.

² Effective E_s/N_0 includes the effect of distortion and un-cancelled cross-pol interference in addition to thermal noise.

Frame Sync

A maximum likelihood approach is used to provide frame synchronization and achieve carrier phase ambiguity recovery. It is implemented by correlating received hard decision bits against the known unique word bit pattern. To increase the detection performance and phase ambiguity recovery accuracy, all the 16 TPC blocks within in one U2B Ka-Band frame are required for the correlation.

Specifically, three correlations are used to provide frame synchronization and detect the carrier phase ambiguity.

1. *Declare the location of the TPC block:* received noisy 8-PSK symbols are converted to bits by first choosing the 8-psk symbol closest in distance to the received symbol. The demodulated symbol is mapped into the bit pattern assigned to the symbol by the 8-psk constellation definition. The resulting bit stream is correlated against the fixed unique word bit pattern by computing the number of mismatches between the bit stream and the bit unique word over the length of the unique word. The number of bit mismatches is compared to a pre-determined threshold λ_0 . If the number of bit mismatches is less than or equal to λ_0 , then a unique word detection is declared and the location of the TPC block is found.

2. *Determine the block count within the frame:* The block count information is obtained to determine the specific location of a TPC block within the TPC frame. Because the block count is randomized and interleaved, these operations must be removed in order to recover the information. Also, since this processing is performed before decoding, a best fit correlation algorithm is required to reduce the probability of error on the block count.

3. *Detect phase ambiguity:* 8 parallel correlations are chosen to determine which one of the 8 phase corrections should be applied to the signal.

Frame synchronization across three demodulator paths is required to perform the De-Hadamard function due to how the Hadamard transform is applied in the spacecraft. The timing recovery loops for all three channels must be locked before data is sent to the frame sync hardware. This architecture will minimize the frame acquisition processing time across all three channels.

In order to minimize the use of FPGA hardware resources, the host computer is used to find the beginning of the frame with respect to a counter value. Hardware then uses the counter value computed by

software that identifies the beginning of the frame to commence data flow to the decoder. Data flow is not started until frame sync is detected in all three demodulator paths.

As a summary, the frame synchronization and carrier phase ambiguity recovery design achieves the following results at 2dB Eb/N0,

- Probability of missed detection $< 10^{-10}$.
- Probability of false alarm $< 10^{-23}$.
- Probability of frame synchronization error $< 10^{-9}$
- Probability of demodulation phase error $< 10^{-17}$

Turbo Decoder

The turbo product decoder is based on turbo product code block. The selected TPC code is (64,56) x (48,39) where the (64,56) component is shortened from a (128,120) extended-Hamming code and (48,39) is shortened from a (256,247) extended-Hamming code.

The turbo product decoder uses an iterative decoding algorithm for product code built using linear block codes³. It is based on a soft-input/soft-output decoder for decoding the component codes so that near optimum performance is obtained at the output of each iteration [6]. This soft-input/soft-output decoder is a Chase decoder which delivers soft outputs instead of binary decisions. The soft output of the decoder is an estimation of the log-likelihood ratio (LLR) of the binary decisions given by the Chase decoder [7]. There are three parameters that drive the turbo product decoder performance: quantization bits, number of iterations and the complexity. The complexity is defined as the number of candidate codewords used for each Hamming component code decode. The parameters are defined as follows:

1. Quantization: 5 bits quantization with +/- 2.5 quantization range
2. Iteration: 4 iterations
3. Complexity of the Chase Algorithm: Search over 48 candidate codeword at the least reliable 6 bit positions with a maximum of 4 bit flips

³ It is a specific case of the more generic serial concatenated block code where the TPC has a specific interleaver map (row column encoding).

REFERENCES

1. David K. Lee and Randall K. Bahr, "Dynamic Range Compression using Hadamard Processing and Decorrelation Spreading," *Presented at MILCOM Oct. 2006*.
2. Floyd M. Gardner, "Interpolation in Digital Modems-Part I: Fundamentals", *IEEE Transactions on Communications*, Vol. 41 No. 3, March 1993.
3. Floyd M. Gardner and Robert Harris, "Interpolation in Digital Modems-Part II: Implementation and Performance, *IEEE Transactions on Communications*, ", Vol. 41 No. 6, June 1993.
4. Kurt Kallman et. al., United States Patent 5,774,508, "Data Synchronizer Phase Detector and Method of Operation Thereof", June 30, 1998, Assignee - Motorola, Inc.
5. Yair Linn and P. Takis Mathiopoulos, "A New Family of NDA Carrier Phase Detectors for Coherent M-PSK Receivers", Department of Electrical and Computer Engineering, University of British Columbia.
6. Charles C. Wang, "Improved Metric for Binary Turbo Decoding Using M-ary PSK Signals," *the Aerospace Corporation*.
7. Ramesh Mahendra Pyndiah, "Near-Optimum Decoding of Product Codes: Block Turbo Codes," *IEEE Tran. On Communications*, Vol. 46 No.8.August 1998

Figure 12 shows a comparison of simulated floating point, simulated fixed point, and hardware measurement of the TPDD BER performance. The hardware BER test was conducted with input symbol data stored in memory and supplied directly to the TPDD, and the information bits produced by the TPDD being checked against reference data. The fixed point simulation data was generated using the quantized version of the floating point simulation code with E_b/N_0 values of 4.2dB, 5.5dB, and 6.8dB, respectively. Each test file contained 16 different frames of data. Fixed point implementation results in about 0.9 dB of degradation at 10^{-5} BER with respect to floating point simulation.




Figure 12 Comparison of Floating Point and Fixed TPDD BER Performance

CONCLUSION

The MUOS U2B downlink digital receiver processes data at a rate of 2457 Mbits/s while using COTS FPGA boards. In accomplishing this function and meeting key requirements including BER, cost, availability and power consumption, trade-offs have been made to reduce digital hardware requirements. These include use of a T-spaced equalizer (rather than T/2 spaced), and a hybrid analog-digital coherent sampling loop. These trades have resulted in only a modest increase in hardware complexity while allowing each of the 3 demodulators to be implemented on a single FPGA board. The turbo decoder implements a rate 0.711 decode function and uses four iterations. It is implemented on 3 FPGA boards per demodulator path for a total of 9 boards. This results in a total of 12 FPGA boards in the receiver.