Performance Comparison of RAKE and Hypothesis Feedback Direct Sequence Spread Spectrum Techniques for Underwater Communication Applications

F. Blackmon,
Naval Undersea Warfare Center
Division Newport
1176 Howell St.
Newport, RI 02841,
USA
blackmonfa@npt.nuwc.navy.mil

E. Sozer,
Department of Electrical &
Computer Engineering
Northeastern University
Boston, MA 02115,
USA
esozer@ece.neu.edu

M. Stojanovic
MIT, E38-376
Cambridge, MA 02139
millitsa@mit.edu
USA

J. Proakis,
Department of Electrical
& Computer Engineering
Northeastern University
Boston, MA 02115, USA
USA
Proakis@neu.edu

Abstract- Direct sequence spread spectrum (DSSS) techniques for low probability of intercept (LPI) and multi-user communications applications for underwater acoustic communications are investigated. Two promising receivers, a RAKE based receiver and a hypothesis-feedback equalization based architecture are considered. The performance of the proposed receiver structures are compared based on simulations and also actual field test data.

I. INTRODUCTION

The need for applications that involve multiple users within the same channel as well as low probability of intercept (LPI) has drawn attention to direct sequence spread spectrum (DSSS) techniques. Towards this goal, two promising variants of this technology, a RAKE based receiver architecture [2] and a hypothesis-feedback equalization based architecture [1]are investigated.

The RAKE receiver is a matched filter that uses a tap delay line to combine signals arriving over multiple propagation paths. When used with differential encoding, the RAKE receiver provides LPI and multi-user with minimum complexity. The hypothesis- feedback equalizer utilizes a chip-spaced decision feedback that is based on hypothesized values of information bits.

These techniques have been used for medium-rate data telemetry at the Atlantic Undersea Test and Evaluation Center (AUTEC). The systems were tested on the downlink from a submarine to a field of bottom mounted hydrophones. In addition, the performance of the RAKE and hypothesis-feedback equalization algorithms were compared through simulation of a number of realistic underwater channels. Section II describes the receiver algorithms. Section III provides the description of the test. The test results are discussed in Section IV. Finally, conclusions are presented in Section V.

II. SYSTEM DESCRIPTION

In this section, the DSSS transmitter and two receiver algorithms, the RAKE receiver and the hypothesis-feedback equalizer receiver, are described.

Figure 1 shows the block diagram of the DSSS transmitter. The information bits are first encoded by a rate ½ convolutional code of constraint length 7 (generator polynomials in octal are 133, 171). The interleaved encoded bits are spread using a complex spreading sequence, i.e. the same encoded bit is spread in both quadrature and in-phase branches of the acoustic channel to fully utilize the channel. After spreading, the signal is pulse shaped using a square root raised cosine filter with roll-off factor 0.25. A 5kHz passband bandwidth is used to provide a chip rate of 4000 chips/sec. The symbol rate of the system can be varied as 50bps, 100bps, 200bps, and 400bps with corresponding spreading gains of 80, 40, 20, and 10 respectively. The synchronization portion of the signal packet employs three repetitions of a quadrature modulated Gold sequence of length x.

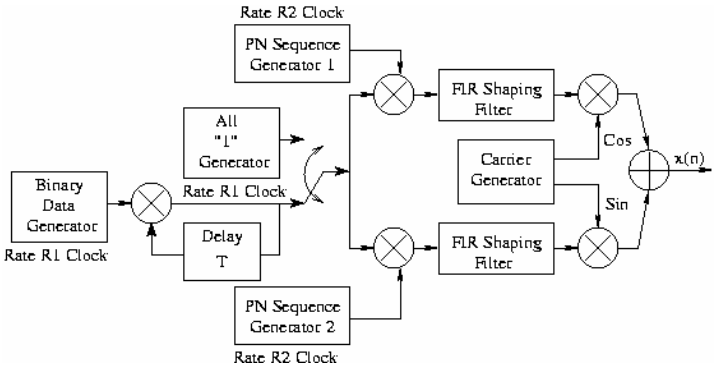

Figure 1. DSSS Transmitter Block Diagram

A. Hypothesis-Feedback Equalization DSSS

The technique of hypothesis-feedback adaptive filtering for DSSS signal detection is based on hypothesizing the value of the data symbol as +1 or –1 in the case of binary modulation. For each hypothesis, an adaptive chip-rate decision feedback equalization and synchronization are then performed. In such a manner, it is possible to update the receiver parameters at the chip rate, which may be necessary for rapidly time-varying channels. Data detection is performed by choosing the hypothesis with lower mean squared error at the end of each bit interval. At the same time, the receiver parameters corresponding to the winning hypothesis are retained for use in the next bit interval, which consists of a number of chip intervals equal to the processing gain. The detailed description of the hypothesis-feedback equalization method for DSSS is presented in [1]. The equalizer used for processing the experimental data was implemented using a standard recursive least squares (RLS) algorithm. The hypothesis-feedback equalization structure is shown in figure 2. The synchronization of the data packet was performed in two steps: coarse initial synchronization

and fine tracking. Coarse synchronization is performed on the received signal using a known replica of the synchronization signal. It provides the packet starting time and a coarse estimate of the Doppler frequency offset. These estimates are used to generate a coarsely synchronized baseband signal, sampled at twice the chip rate. Fine synchronization is required to track the residual Doppler distortion. A resampling method, based on filtering the one step phase difference within the equalizer's phase lock loop (PLL), was used . The phase difference is used by a delay locked loop (DLL) which performs linear interpolation between two adjacent samples to obtain the signal at the desired sampling time within every chip interval.

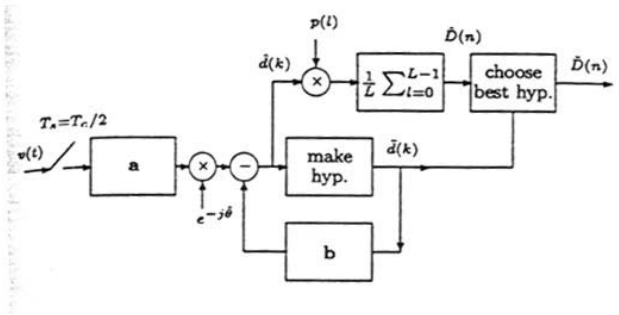

Figure 2. Hypothesis Feedback Equalizer DSSS Structure

B. RAKE-Based DSSS

The RAKE-based DSSS presented in [2] which employs differentially coherent methodology is a lower complexity alternative to the hypoithesis-feedback equalization DSSS. In the RAKE-based DSSS scheme, the symbol rate is chosen such that the channel can be assumed to be constant for two symbol durations. Then, differentially coherent modulation can be employed which does does not require explicit channel estimation. The need for an equalizer may be eliminated with a RAKE receiver if the processing gain of the system is hgigh enough to suppress intersymbol interference (ISI). In addition to suppressing the ISI, a RAKE structure can be employed to make use of the energy present in multiple propagation paths. This energy extraction process is also known as time diversity.

The first step in demodulating spread spectrum signals is synchronization, which is handled in two steps: acquisition and tracking. Acquisition is obtained by using the three synchronization pulses that precede the data packet. A DLL is employed for fine synchronization and tracking of the channel. After synchronization, the received signals are demodulated using the RAKE receiver. Figure 3 shows the RAKE filter where correlators are shown by simple multiplication. The correlator at each tap of the RAKE receiver is shown in Figure 4. Signals at the taps are first multiplied by the corresponding PN sequences. Next, the two pairs of signals are added and subtracted as shown in the figure to provide the in-phase and quadrature signals at the output of the correlator. This complex despread signal is then fed to the differential decoding block. The output signal level of each RAKE tap is compared with a threshold. The outputs that stay below the threshold are discarded, thereby reducing the noise at the combining stage. The output signals that are above the threshold are summed and the polarity of these signals is used to make a decision on the incoming bit.

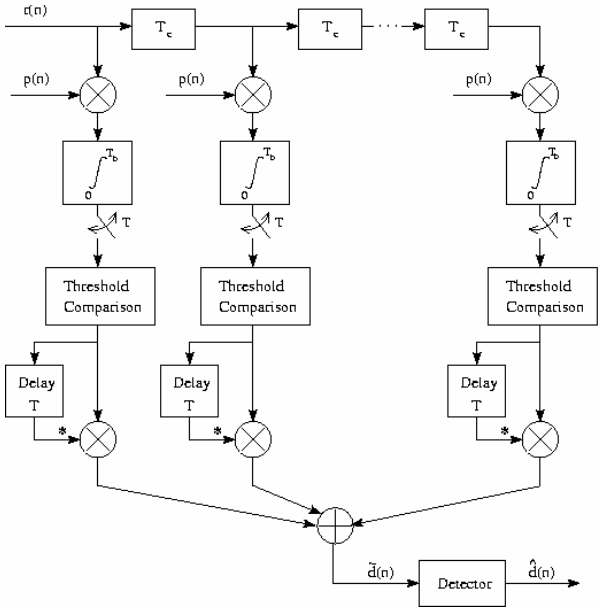

Figure 3. RAKE Based DSSS Receiver

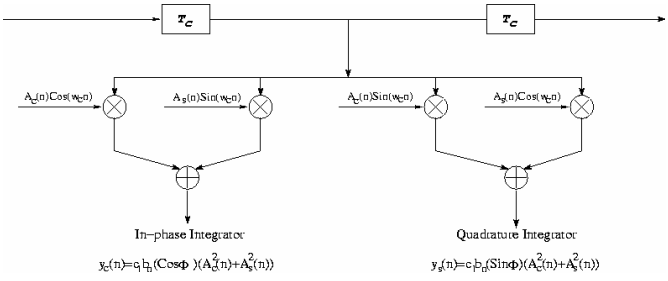

Figure 4. Correlator Used at Each Tap of RAKE Receiver

III. TEST DESCRIPTION

In February of 2002, an Underwater Range Data Communication (URDC) submarine telemetry reception test was conducted at the Atlantic Undersea Test and Evaluation Center (AUTEC) off Andros Island in the Bahamas to determine the performance of spread spectrum receiver structures as applied to acoustic data received on bottom mounted _____(AHRP) range hydrophones. A two-

channel Digital Audio Tape (DAT) recorder was used onboard the test submarine to play the signals into an onboard _____(BBLF) transducer The submarine was traveling at a speeds of up to 15 knots. On shore, the received signals from a select group of range hydrophones were recorded on an eight channel DAT recorder. The data was then post processed in Matlab to provide the results reported in this paper.

IV. DATA PROCESSING RESULTS AND DISCUSSION

The results for the receiver algorithms outlined in Section II are presented below. These results include selected field test results and simulation results for particular channel responses corresponding to four receive nodes designated as H22, H19, H18, and H16 with slant ranges varying from 1700 yards to 2200, 2000, 2200, and 2500 yards respectively from the transmitter to the receiver. The following parameters are examined: signal to noise ratio, two user energy ratio, and associated bit error rates for both receiver algorithms.

A. Field Test Data

In the case of nodes H22, H19, and H18, the RAKE DSSS receiver followed by the Viterbi decoder successfully demodulated the received signals. Occasionally, at 200 and 400 bps, there were 20 or less errors at the output of the RAKE receiver prior to decoding. In each case, the multipath spread was between 0.5ms to 10ms which corresponds to 2 to 40 chips. The RAKE receiver was set with a RAKE tap delay line length equal to 10 taps. Table 1 shows the typical performance of the RAKE receiver for Node H22 demonstrating the excellent results that can be obtained. These results are typical of the nodes H19 and H18.

Node H16 presented a significant multipath arrival around the 50^{th} chip. Therefore, the RAKE filter length was increased to 80 chips in some cases and 100 chips in other cases [WHY SO MANY?]. Table 2 shows the performance results for packet 13 received on node H16. With the extended filter, the amount of noise that entered the decision device increased despite the thresholding in the RAKE. Therefore, the threshold value was increased to 1.0 for the 400 bps case. This decreased the number of errors but not significantly (marked with asterisks in table 1.). As a second approach, the RAKE filter length was limited to the first group of multipath, i.e. the filter length was reduced to 10 taps. In all other packets for this node, the number of errors at the output of the RAKE DSSS prior to decoding was less than 15 by varying the number of RAKE taps using either 10 or 80 taps and the number of errors was reduced to 0 after decoding. When the energy in the RAKE taps was checked for packet 13, user 1, 400 bps, it was observed that the main arrival was lost which may occur when the DLL loses the lock on the main arrival.

Table 1.Bit Error Results of RAKE Receiver (H22 Packet 1)

Chan. #	User	Rate	# Taps	Rake Errors	Decoder Errors
22	1	50	10	0	0
22	1	100	10	0	0
22	1	200	10	0	0
22	1	400	10	2	0
22	2	50	10	0	0
22	2	100	10	0	0
22	2	200	10	0	0
22	2	400	10	4	0

Table 2.Bit Error Results of RAKE Receiver (H16 Packet 13)

Chan. #	User	Rate	# Taps	Rake Errors	Decoder Errors
16	1	50	100	0	0
16	1	100	100	1	0
16	1	200	100	12	0
16	1	200	10	193	101
16	1	400	100	223	183
16	1	400	10	133	101
16	1	400	100	218	215
16	2	50	100	0	0
16	2	100	100	0	0
16	2	200	100	11	0
16	2	200	10	0	0
16	2	400	100	125	117
16	2	400	10	12	0
16	2	400	100	117	63

The hypothesis-based feedback equalization DSSS receiver prior to decoding produced no errors in all cases except for one case. Typical processing results for the node H22 channel response is shown in Figure 5 and the chip scatter plot with associated figures of merit are shown in figures 6a through 6d for data rates of 50bps, 100bps, 200bps, and 400bps respectively.

The channel response and the data processing results for node H16 [WASN"T THIS 17?]packet 1, user 1, for the 400bps case, the only case where there was residual error at the output of the hypothesis equalizer, are shown in figures 7 and 8. Figure 7 characterizes the performance of pre-processing to obtain coarse synchronization. The channel response magnitude is shown, which is the cross-correlation between the Doppler-corrected received preamble and the known preamble. The indicated Doppler frequency offset is the coarse initial estimate. Figure 8 characterizes the performance of the adaptive hypothesis feedback receiver prior to decoding. This figure shows the chip scatter plot, the data scatter plot, and the estimated phase and delay. Indicated in the figure are the receiver parameters used for processing as well as the estimated residual Doppler frequency offset, bit error probability, and the output SNR.

The number of errors in this packet is 2 producing a probability of bit error of 0.0025.

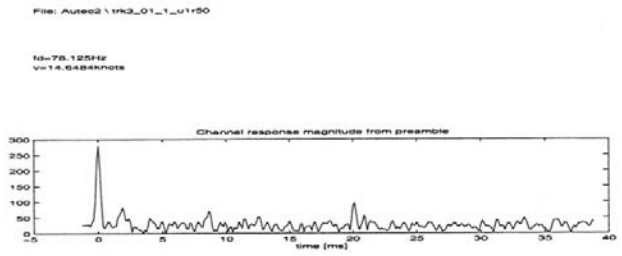

Figure 5. H22 Packet 1 User 1 Preamble Processing Results

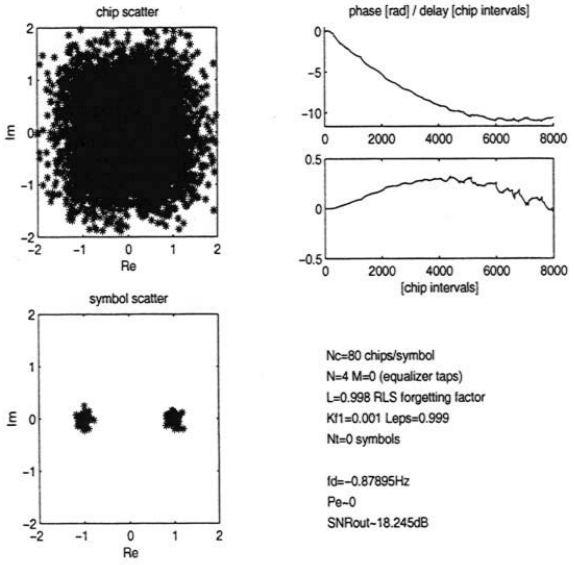

Figure 6a. H22 Packet 1, User 1, 50bps Data Processing Results

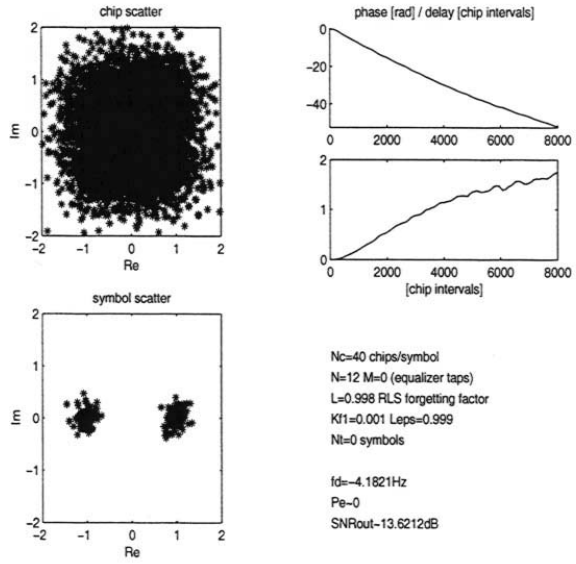

Figure 6b. H22 Packet 1, User 1, 100bps Data Processing Results

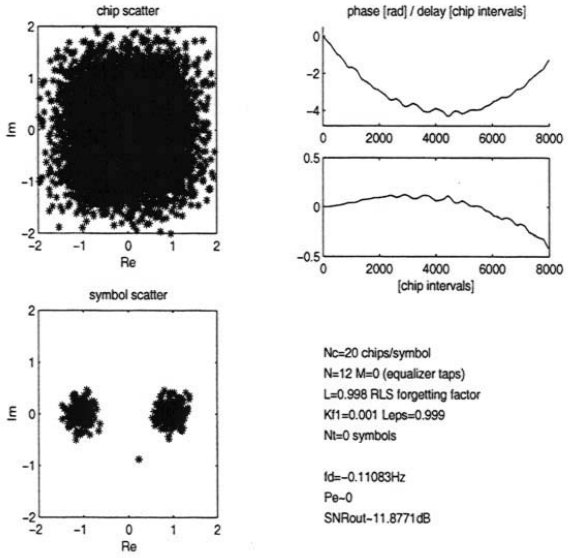

Figure 6c. H22 Packet 1, User 1, 200bps Data Processing Results

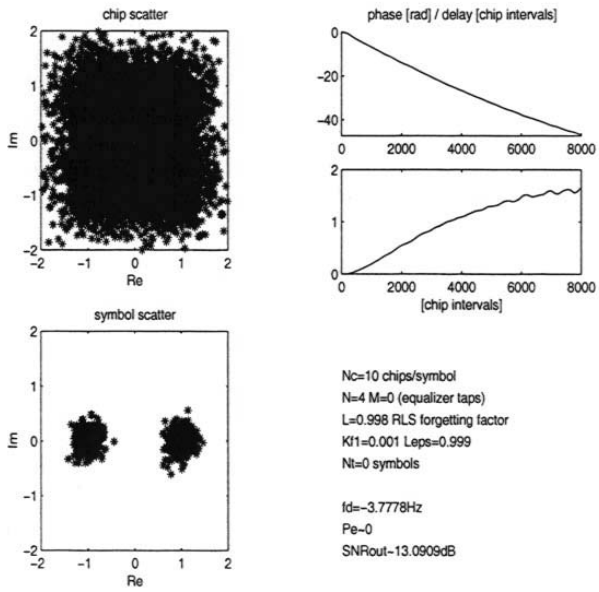

Figure 6d. H22 Packet 1, User 1, 400bps Data Processing Results

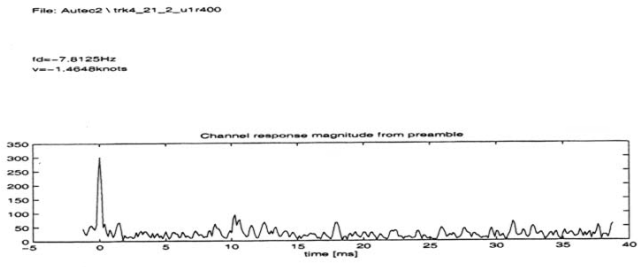

Figure 7. H16 Packet 1 User 1 400bps Preamble Processing Results

Due to the low number of residual errors, the Viterbi decoder that follows corrects these errors. It should be noted that training was not required and that 10 feedforward and feedback equalizer taps were used. The total Doppler frequency offset is estimated to be between 13 and 14Hz which differs from that obtained with the RAKE receiver. The fact that the hypothesis feedback receiver does not exploit the distant multipath, yet provides good performance, indicates that the RAKE receiver too is capable of suppressing the interference from the distant multipath. Thus, it is possible that the reason for the RAKE receiver's failure to produce better results on these types of packets is not in multipath suppression but in synchronization.

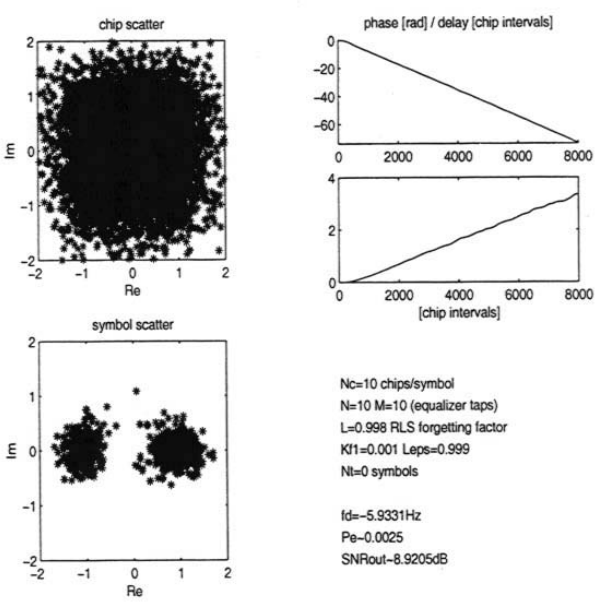

Figure 8. H16 Packet 1, User 1, 400bps Data Processing Results

The multi-user capabilities of each receiver structure were tested with the field test data from packet 21 node H16. The signals from two users were combined in the passband asynchronously, simulating a multiple access interference environment. The combined packet was processed for both users. No knowledge of the interfering user's presence or the type of interfering signal is assumed. Table 3 and 4 show the results of various signal-to-interferer ratios (SIRs) of 5dB, 0dB, and –5dB for the RAKE and hypothesis feedback receivers respectively. User 2 is the intended user.

Table 3. RAKE DSSS Receiver Multi-User Results

Intended User	SIR (dB)	Data Rate (bps)	# Taps	Rake Errors	Decoder Errors
2	5	50	10	0	0
2	5	100	10	0	0
2	5	200	10	3	0
2	5	400	10	10	0
2	0	50	10	0	0
2	0	100	10	2	0
2	0	200	10	15	3
2	0	400	10	90	24
2	-5	50	10	0	0
2	-5	100	10	15	0
2	-5	200	10	69	63
2	-5	400	10	261	192

Table 4. Hypothesis Feedback Equalizer DSSS Receiver Multi-User Results

Intended User	SIR (dB)	Data Rate (bps)	# Taps	Equalizer Errors	Decoder Errors
2	5	50	10	0	0
2	5	100	10	0	0
2	5	200	10	0	0
2	5	400	10	0	0
2	0	50	10	0	0
2	0	100	10	0	0
2	0	200	10	0	0
2	0	400	10	2	0
2	-5	50	10	0	0
2	-5	100	10	0	0
2	-5	200	10	0	0
2	-5	400	10	-	-

The RAKE multi-user results shown in table 3 demonstrate that at high processing gain of 80, at a data rate of 50bps, no errors are detected at the receiver output. Errorless performance at a data rate of 100bps is possible with the use of the decoder. At other SIR values and data rates there are residual errors at the output of the RAKE and the decoder. The worst case is at an SIR of –5dB at a data rate of 400bps where a third of the packet is in error.

The hypothesis feedback multi-user results shown in table 4 demonstrates excellent multi-user performance at almost all combinations of SIR values and data rates. At an SIR of 0dB, there are a couple of residual errors at the output of the equalizer; but these are easily corrected by the decoder. At a data rate of 400bps and a SIR of –5dB, the receiver cannot stay in convergence. Additional simulation, not presented here, reveals that at a SIR of –5dB and data rate of 400bps, successful performance is possible only at a very high input SNR.

Next, a specific case with data rates of 100bps and 400bps will be examined. The energy ratio of user 2 to user 1 was approximately two, i.e. an SIR of –6dB, where user 1 is the intended user. In the 100bps case, the packets from each user are synchronous to within one chip interval. There are no errors for the hypothesis feedback receiver in the 100bps case, although the receiver performance is better for the stronger user 2. However, the RAKE receiver has no errors for user 2 yet has 12 errors for user 1 prior to decoding and no errors after decoding as shown in the RAKE receiver multi-user results listed in table 5. In the 400bps case, the signals are offset in time by 3.5 chip intervals. Table 5 shows that there are errors at the output of the RAKE receiver for both users at this data rate. However, there are many more errors for user 1 as expected. A stronger error correction code may have been able to remove all the errors at the output of the decoder for user 1.

Table 5. RAKE DSSS Receiver Multi-User Results (SIR=-6dB)

Intended User	Rate (bps)	# Taps	Rake Errors	Decoder Errors
1	100	10	12	0
2	100	10	0	0
1	400	10	184	204
2	400	10	33	4

For the same case using the hypothesis feedback receiver, the stronger user's signal is detected easily but the weaker user suffers because the reduced processing gain for the 400bps case is not sufficient to overcome the combined multiple access and multipath interference. With sufficiently long training the algorithm converges, and more importantly, synchronization is maintained. It should also be noted that cross correlation between the users spreading sequences can be very large over the short interval of 10 chips since the processing gain decreases from 80 to 10 as the data rate changes from 50bps to 400bps. High cross correlation reduces the ability to suppress multiple-access interference when the propagation channels of the two users is similar. The user 1 and user 2 equalization results are shown respectively in figures 9 and 10. There are no errors for user 2 after equalization as shown in figure 10 as compared to the residual 2 errors remaining after decoding the RAKE receiver's output. However, user 1 whose results are shown in figure 9 can not separate the BPSK symbols appropriately as seen in the symbol scatter plot. There are 18 errors at the equalizer output in this case. This is a factor of 10 less than the number of errors present at the output of the RAKE receiver for the same case. The convolutional decoder has no problem decoding for this user to provide no errors. In the RAKE case as seen from table 3, there was still 204 errors after decoding. The comparative results show that the hypothesis feedback DSSS receiver outperforms the RAKE DSSS receiver under these multi-user conditions due to the equalization method's ability to distinguish the two users by applying different training sequences as well as being able to synchronize effectively.

B. Simulation Data

Two separate simulation channels were created based on the experimental data processed in the previous section in order to demonstrate the LPI capabilities of both DSSS receiver structures. The particular figure of merit is the lower limit on the SNR at which acceptable receiver performance is still available. The SNR per bit is equal to the SNR per chip multiplied by the number of chips per bit. Channel 1

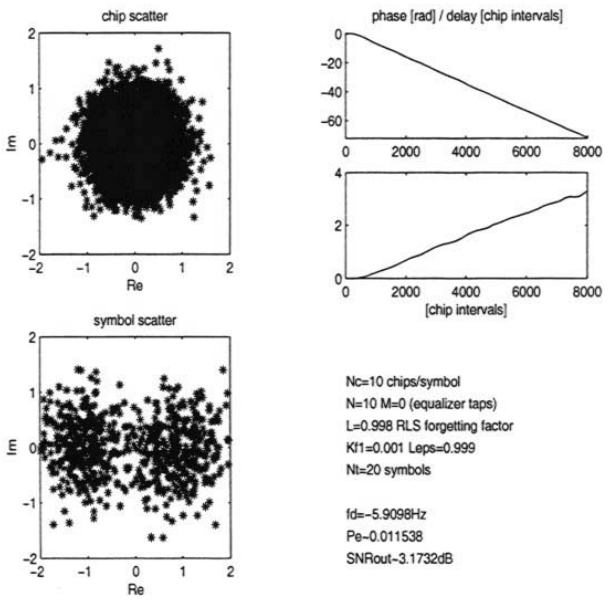

Figure 9. Hypothesis Feedback User 1, H16, Packet 21 Data Processing Results

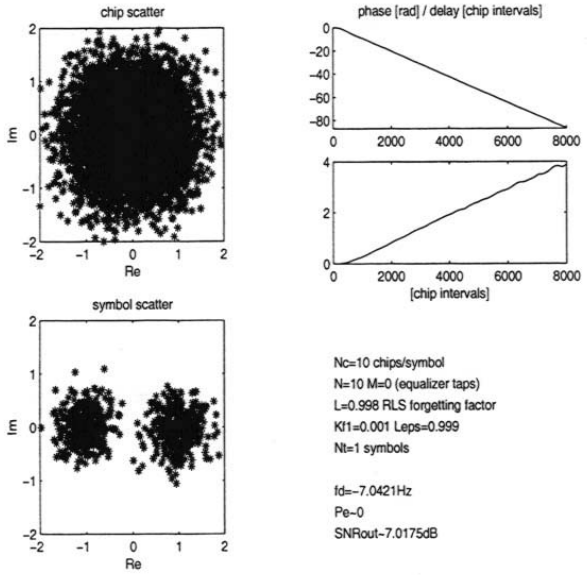

Figure 10. Hypothesis Feedback User 2, H16, Packet 21 Data Processing Results

represents the mutlipath responses seen on nodes 22 and 19 where the multipath arrivals are within a couple of chips. Channel 2 represents the multipath responses seen on node 16 where a strong second group of multipath arrivals are present around the 50th chip. Channel 1 coefficients were set to w = [1.0 0.5 0.4]. Channel 2 coefficients are the same for the main arrival and scaled version of these coefficients is used for the second group of arrivals as shown in figure 11.

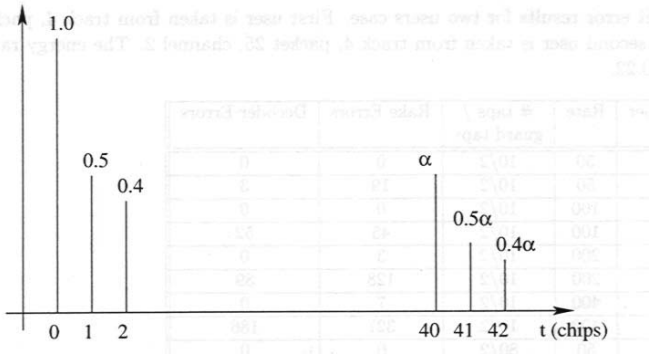

Figure 11. Multipath Profile of the Simulated Channels

Figure 12 shows the RAKE DSSS receiver bit error rate versus SNR per chip for channel 1 for the case of 50 bps and 400 bps. Both the output of the RAKE receiver before and after decoding are shown. A difference of about 9dB between the two data rate cases is seen since the energy per bit for the 50 bps case is 8 times more than the 400 bps case. The Viterbi decoder provides an additional 4 dB gain around BER = 10^{-2} . The error rate after decoding drops below 10^{-4} at –11dB for 50 bps which is well below the noise level. In the case of 400 bps, the SNR must be increased to –1dB to achieve the same performance.

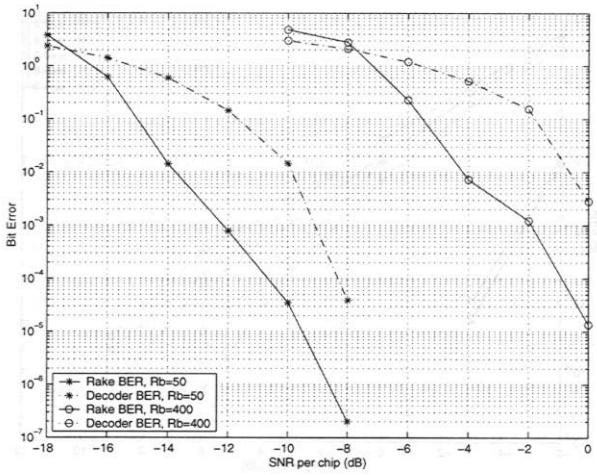

Figure 12. RAKE Receiver Channel 1 BER vs. SNR/chip

The bite error rate performance of the RAKE receiver for the second channel is presented in figures 13 and 14 for data rates of 50 bps and 400 bps respectively for various scaling values of the coefficients representing the second group of multipath arrivals. The dashed lines in the figures represent the output of the receiver prior to decoding and the solid lines represent the output of the decoder. The 400 bps case suffered from inter-symbol-interference and the bit error rate curve converged to an error floor. However, the 50 bps case does not show any sign of error floor. The processing gain of

this data rate/processing gain case is enough to suppress the ISI and provide LPI communication capability. It should be noted, although not shown that the 100 bps case does not show an error rate floor; but the 200 bps case tends towards an error rate floor for larger amplitude coefficients of the second group of multipath arrivals.

For both channel models the same hypothesis feedback equalization structure was used for comparison to the RAKE receiver performance curves. Six fractionally spaced feedforward taps and no feedback taps were used. Once again the 50 bps (processing gain of 80 chips per bit) and 400 bps (processing gain of 10 chips per bit) were used. The results are summarized in figures 15 and 16. The following observations can be made from the simulation results. First, there is a consistent difference of 9 dB between the 400bps and the 50bps bit error rate curves, corresponding to the difference in the respective processing gains. The receiver performance is slightly worse on channel 2 compared to channel 1. In channel 2, the receiver does not exploit the energy of the distant multipath cluster, but only that of the first cluster. The useful energy is thus slightly less in channel 2 than in channel 1.

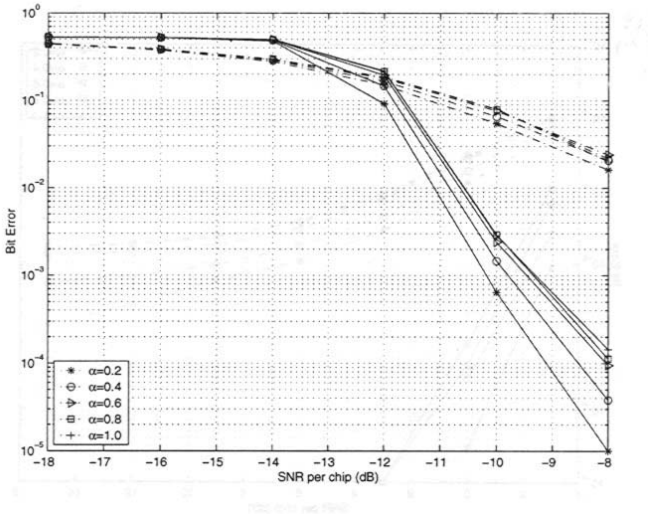

Figure 13. RAKE Receiver (50bps) Channel 2 BER vs. SNR/chip

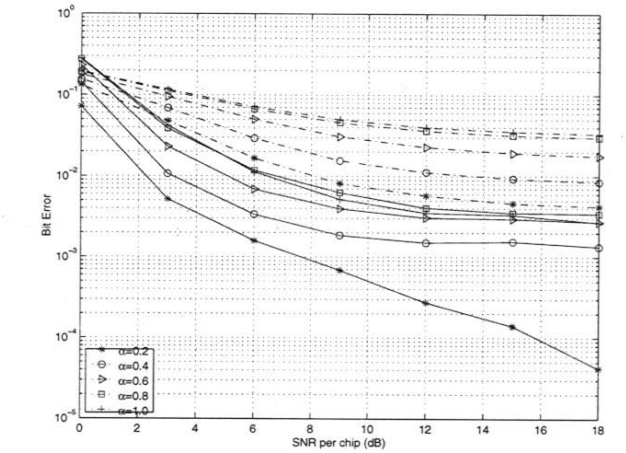

Figure 14. RAKE Receiver (400bps) Channel 2 BER vs. SNR/chip

The results shown for the hypothesis feedback equalization receiver structure are those of an uncoded system and "acceptable performance" can be defined as that which gives a relatively high bit error rate. For a bit error rate of 10^{-2}, a bit SNR of 3dB is needed in channel 1, while a bit SNR of 5dB is needed in channel 2. For a bit error rate of 10^{-3}, a bit SNR of 6dB is needed in channel 1, while a bit SNR of 8dB is needed in channel 2. The coding gain for the code used in the simulation is approximately 3dB which provides the same BER performance but at 3dB lower bit SNR. It should be noted that the hypothesis feedback receiver performance is better by 1 to 2dB than the RAKE receiver for channel 1 for a given BER. The difference in performance increases for channel 2. In this case, the hypothesis feedback equalization method provides 3.5dB improvement over the RAKE receiver prior to decoding for the 50bps case.

An interesting observation was made by comparing the simulation results of the hypothesis feedback receiver to those of the differentially coherent RAKE receiver. In the case of channel 2 transmission at 400 bps, the differentially coherent RAKE receiver exhibited an error floor. The error floor is due to residual ISI that prevents the differentially coherent detection from functioning normally. No error floor was observed using the hypothesis feedback receiver within the same SNR range.

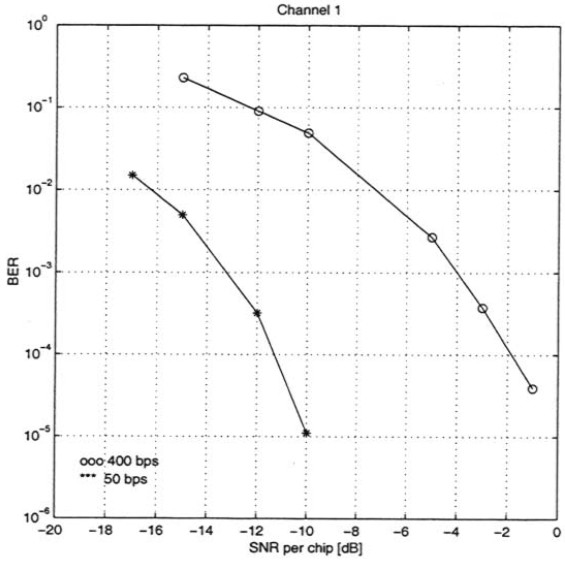

Figure 15. Hypothesis Feedback Receiver Channel 1 BER vs. SNR/chip

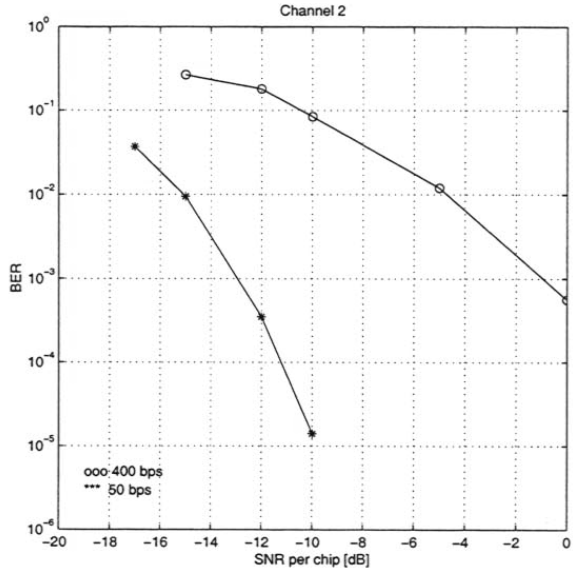

Figure 16. Hypothesis Feedback Receiver Channel 2 BER vs. SNR/chip

V. CONCLUSIONS

There are a number of conclusions that can be made based on the field test and simulation performance results. In the field test data case, both DSSS receivers, the differentially coherent RAKE receiver and the hypothesis-based feedback adaptive receiver, have been successfully applied to the AUTEC data set. This data set is characterized by varying Doppler spread, corresponding to transmitter velocity of up to approximately 15 knots. The multipath properties of the channel are also time-varying. The channel usually contained a cluster of arrivals associated with the main path and in some instances a distant multipath cluster. The delay spread within the main cluster varied between two and ten chip intervals.

In the large majority of the received signal packets, the hypothesis-based feedback equalizer employed a feedforward filter spanning five chip intervals which was sufficient to provide excellent quality of detected symbols. It was observed that acquiring and maintaining correct synchronization, in both phase and delay, was crucial to proper receiver operation. Apart from the time-varying Doppler, caused by transmitter motion, there appeared to be little time-variation in the multipath strength.

From the comparative study of the two receivers, both receiver structures are applicable to the channels that were studied. The major difference between the two approaches is in their ability to track the time-variation in the signal. For channels where it is necessary to track the signal phase and delay on a chip-by-chip basis, the hypothesis-based feedback equalizer is recommended. However, if Doppler correction can be performed with sufficient accuracy prior to data processing, then the RAKE receiver offers good performance as well as simplicity of implementation.

The comparative results show that the hypothesis feedback DSSS receiver outperforms the RAKE DSSS receiver under these multi-user conditions due to the ability to simultaneously distinguish the two users and mitigate the effect of multipath, while maintaining synchronization.

The SNR performance of the algorithms indicates that the hypothesis feedback equalizer is able to provide BER improvement over the RAKE receiver for the same chip SNR. Also, there is an error floor that is seen for the RAKE receiver which does not exist for the hypothesis feedback equalization receiver.

Regarding the future use of the hypothesis feedback receiver, the following issues should be addressed: multi-channel operation as a way of improving performance in all conditions (severe multipath as well as multiple-access interference), channel-estimation based sparse decision-feedback equalization as an efficient way of suppressing distant multipath, and design of multi-user signaling techniques that maximally exploit the available bandwidth expansion through code selection.

VI. ACKNOWLEDGEMENTS

This work was sponsored through the Underwater Range Data Communications (URDC) program at the Naval Undersea Warfare Center Division Newport. The program manager for this effort is Wilfred Canto.

VII. REFERENCES

[1] M. Stojanovic and L. Freitag, "Hypothesis-Feedback Equalization for Direct-Sequence Spread-Spectrum Underwater Communications," IEEE Oceans'00 Conference, Providence, RI, Sept. 2000.

[2] E. M. Sozer, J. G. Proakis, M. Stojanovic, M. Hatch, J. A. Rice, and A. Benson,, "Direct Sequence Spread Spectrum Based Modem for Under Water Communication and Channel Measurement," IEEE Oceans'99 Conference, Seattle, WA, Sept. 1999.