Underwater Acoustic Single- and Multi-User Differential Frequency Hopping Communications

Dianne Egnor*, Luca Cazzanti†, Julia Hsieh† and Geoffrey S. Edelson*
*BAE Systems
P.O. Box 868
Nashua, NH 03061-0868 USA
†University of Washington - Applied Physics Lab
Box 355640
Seattle, WA 98105 USA

***Abstract*— Differential frequency hopping (DFH) is a fast frequency hopping, digital signaling technology that achieves the desirable performance features of non-interfering spread spectrum operation, spectral re-use, fading mitigation, and interference resistance. Therefore, DFH coding provides the critical capability for multiple users to seamlessly communicate in the bandwidth-limited acoustic channel. In previous work, DFH coding has been shown to be superior to other coding schemes in additive Gaussian white noise and Rayleigh-fading environments when considering the joint constraints of multiple user access, detectability mitigation, and the presence of jamming.**

In this paper, we describe the auto-synchronizing single-user DFH decoder we have developed for a single hydrophone receiver. We present the performance of this decoder on multi-user simulated data and on multi-user data collected at sea during the Rescheduled Acoustic Communications Experiment (RACE08). We use the Sonar Simulation Toolset (SST) to produce the simulated data for soft through hard bottom compositions to provide a range of multipath severity to gain insight into DFH performance across environments. Based on these initial results, the DFH waveform shows considerable promise for computationally minimal, high reliability communications among uncoordinated users in an underwater acoustic channel.

I. Introduction

Differential frequency hopping (DFH) is a frequency hopping digital signaling technology that achieves the desirable performance features of noninterfering spread spectrum operation, spectral reuse, multipath fading mitigation, and interference resistance [3], [4], [7].

For DFH waveforms, the frequency of the transmitted tone depends on both the current data symbol and the previous transmitted tone. That is, given a data symbol X_n and the frequency of the previous hop F_{n-1}, the frequency of the next hop is determined as $F_n = G(F_{n-1}, X_n)$ where the function G can be viewed as a directed graph that has nodes corresponding to frequencies and vertices labeled with input data.

Trellis models, often used in depicting and analyzing convolutional codes, are easily applied to a differential frequency-hopped signal, as shown in Fig. 1. The vertical axis of the trellis corresponds to frequency, while the horizontal axis corresponds to time intervals. The set of states at any given time corresponds to the set of all possible frequencies that may be transmitted by the system. For a hop set of size M, there

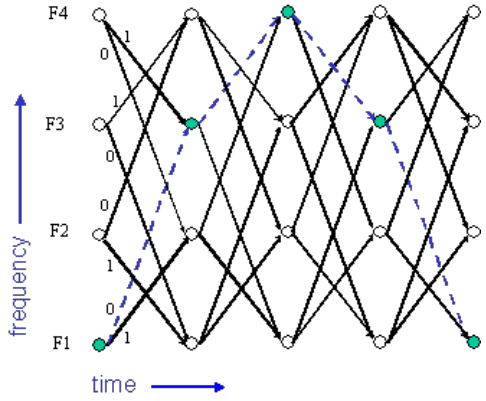

Figure 1. Example DFH trellis for a hop set of size four

are M possible states at each stage in the trellis. The branches leaving each state terminate at the frequencies that are possible at the next hop, given the current frequency state. A label on each branch indicates the encoded bits that correspond to the transition from the current transmitted frequency to the next transmitted frequency.

A trellis-based DFH receiver can reconstruct transmissions that are missing due to a fading channel or collisions with other users. The trellising also allows for the simultaneous demodulation of multiple users by assigning unique trellises to individual users. For the trellis in Fig. 1, the bit rate is one bit per hop, the hop set size M is 4, and the data sequence shown by the dotted line is 0110. Note that the first detection at frequency F_3, corresponds to a 0 data bit and the second detection at F_3 corresponds to a 1 data bit. This illustrates the DFH feature that the sequence of detections, not the detections themselves, carry the information.

Differential frequency hopping waveforms were originally proposed for operation in HF (High Frequency) bands [4]. In more recent work [7], the DFH concept was generalized to any frequency range, and trellis concepts were applied. Bounds have been developed to characterize performance in additive white Gaussian noise (AWGN) and Rayleigh fading channels. It was shown in [7] that generalized DFH waveforms demonstrate excellent single-user performance, and are

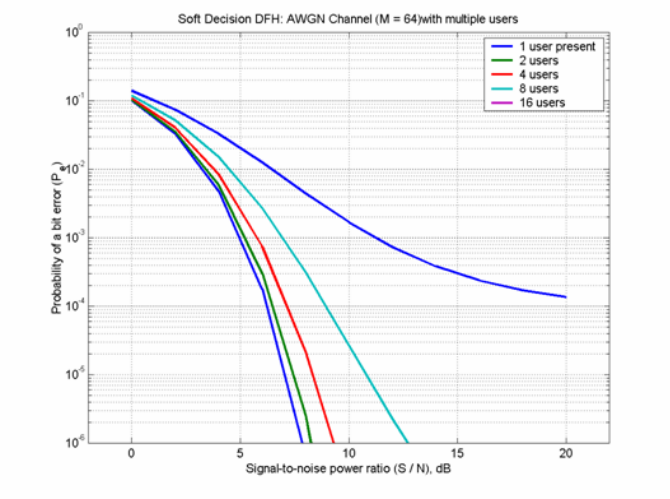

Figure 2. Multiuser SD-DFH probability of bit error (one bit per hop, hopset size of 64 frequencies)

tolerant of co-channel (multi-user) interference. DFH has been favorably compared to conventional FSK, Fast Frequency Hopped MFSK (FFH/MFSK), and Direct Sequence Spread Spectrum (DSSS) [7]–[9].

Because the DFH waveform is tolerant of interfering signals, it is well suited for multiple access environments. In a multi-user DFH system, each transmitter uses a unique trellis pattern. At each receiver, the decoder follows the trellises that correspond to the users of interest. The compatibility of two or more trellises can be measured by the distance between the assigned trellises. The trellis depth directly reflects the actual tolerance to multi-user interference (MUI), as it is a measure of how well a particular trellis structure accepts worst-case hit sequences from interferers. Judicious trellis design can maximize the trellis depth of assigned trellises.

The multiuser SD-DFH probability of bit error is plotted for $M = 64$ frequencies in Fig. 2, with a bit rate of one bit per hop. It can be seen that a moderate number of additional (interfering) users does not significantly affect the bit error rate performance bound.

DFH modulation is also self-synchronizing. Because the data are encoded in the intervals between successive hops, both bulk frequency and time offsets can be determined in the decoder from the waveform itself, without the use of training symbols. Furthermore, the described approach does not rely on centralized controllers, and requires no orchestration between users for conferencing and bandwidth packing, beyond the assignment of each user's trellis G.

This paper presents the performance results for an efficient, self-synchronizing, single-user DFH receiver approach that provides an improved anti-jam, MUI-tolerant communications capability appropriate for the underwater acoustic channel.

II. DFH in the Underwater Acoustic Channel

Acoustic communications waveforms propagating in the underwater channel can be severely distorted in both time and frequency, causing corruption of the received data [5]. Variations in water depth, bottom type, sound speed profile, and source and receiver location can cause a wide range of multipath interference, which results in time spreading of the transmitted signals and consequent intersymbol interference (ISI) in the received bit stream. Different wind and surface wave conditions, in combination with platform motion, can result in varying degrees of Doppler shift and spreading.

Multipath interference in underwater acoustic propagation is the prime cause of ISI in communication signals. In frequency hopping modulation schemes, this causes energy transmitted in one time-frequency bin to extend further in time than the intended duration of that bin. Doppler spreading has a related effect on frequency-hopped signaling, in that energy from one time-frequency bin can leak into neighboring frequency bins at a given time instant. Furthermore, the frequency fading characteristics of the underwater channel may cause the energy in some time-frequency bins to be drastically attenuated relative to other bins [1], [11], [12].

Our first assessment of the capabilities of DFH modulation in the underwater channel was conducted by simulating the underwater propagation of DFH signals with the Sonar Simulation Toolset (SST) [2]. SST is a set of software modeling tools for producing high-fidelity acoustic timeseries that realistically take into account the effects of the underwater environment on the propagating signals. SST flexibly accommodates both arbitrary acoustic signals generated by external signal generators and digitized experimental signals recorded during at-sea tests. In SST, underwater environments are defined by parameters relevant to acoustic signal propagation and reception: sound speed profile, bathymetry, surface and bottom characteristics, ambient noise levels, etc. These parameters are particularly important to underwater DFH signal analysis because they control the degree of multipath interference in the waveforms.

SST also allows specification of arbitrary locations and trajectories of acoustic sources and receivers within an environment. SST uses acoustic propagation models and standard signal processing techniques to produce properly calibrated, realistic digital timeseries of the signals as they would appear at the receivers in the specified environments. These timeseries can then be operated on by the same signal processing algorithms that operate on acoustic data measured at sea.

In our SST-based study, we simulated an ocean environment characterized by a downward-refracting sound speed profile (SSP), and a 100m flat-bathymetry water depth. Omni-directional acoustic sources (representing different DFH users) were placed at a mean distance of 5km away from a single-hydrophone omni-directional receiver, and all sources and the receiver were placed at a water depth of 50m. We varied the properties of the ocean bottom between a soft bottom (designated Sandy Clay) and a harder bottom (designated Medium Sand). Typically, harder bottom types reflect more acoustic energy than softer bottom types. The consequent stronger multipath arrivals at the receiver lengthen the duration of the channel impulse response (CIR) of environments characterized by harder bottom types. Longer CIRs lead to more severe ISI,

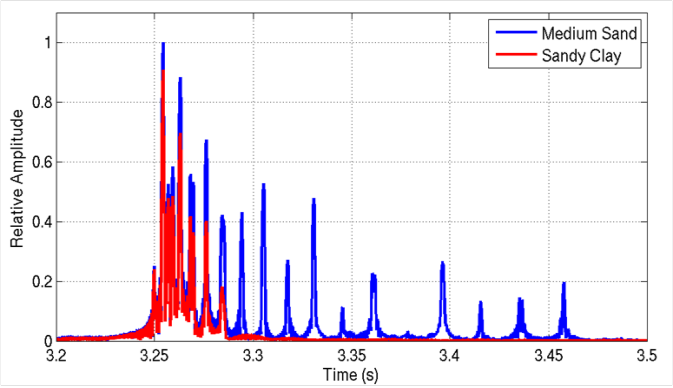

Figure 3. The CIRs for two different ocean environments characterized by a hard bottom (Medium Sand) and a softer bottom (Sandy Clay).

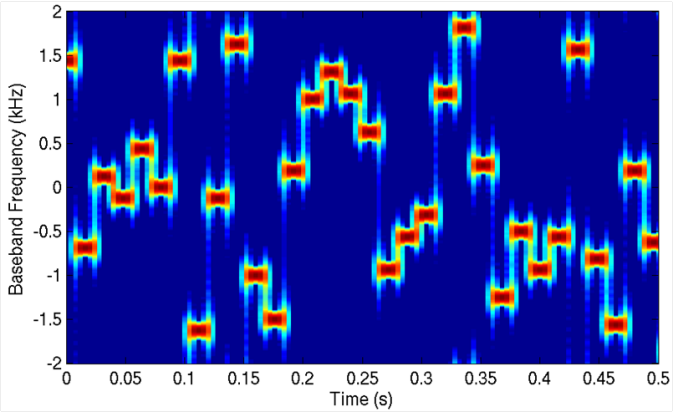

Figure 4. Spectrogram of the beginning portion of the transmitted DFH signal for one of the users.

and can limit effective communications [1], [11], [12].

An effective way to illustrate the differences in CIRs between two environments using SST is to transmit a simulated linear FM sweep, and to match-filter the simulated reception with that same transmitted FM sweep [10]. Fig. 3 shows the results obtained using this procedure (using a 4kHz FM sweep centered at 14kHz) for the two environments simulated in this study. Note how the blue curve (for the Medium Sand harder bottom) has a much longer effective CIR, which can cause severe ISI.

The simulations involve transmitting simulated signals from eight different users (using eight different trellises) offset slightly in range from each other. Each signal is 60 seconds long at a symbol rate of 62.5 bits per second. Fig. 4 shows a spectrogram of the beginning portion of the transmitted DFH signal for one of the users.

Fig. 5 shows a spectrogram of the simulated reception in the Sandy Clay environment. The time spread is not a noticeable feature in this figure, but frequency-selective fading is evident. Some frequency bins are noticeably attenuated compared to others due to closely spaced multipath arrivals that destructively interfere at some frequencies and constructively interfere at others. Fig. 6 shows a spectrogram of the simulated

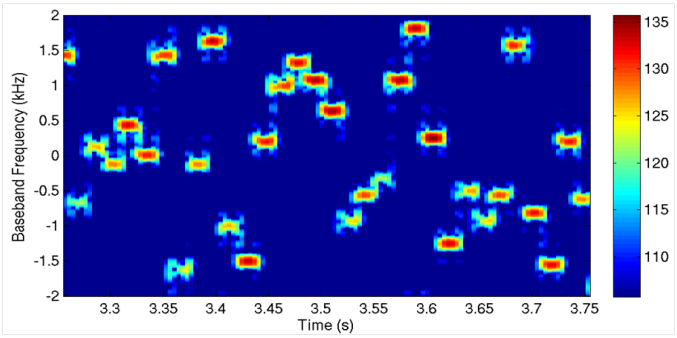

Figure 5. Spectrogram of the DFH signal from Fig. 4 after it propagates through the Sandy Clay environment.

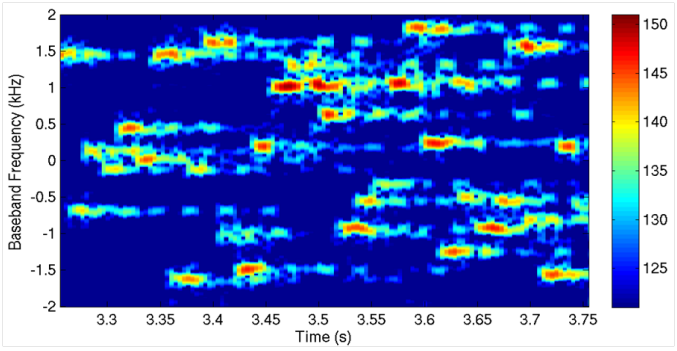

Figure 6. Spectrogram of the DFH signal from Fig. 4 after it propagates through the Medium Sand environment.

reception in the Medium Sand environment. The increased time spread is very evident in comparison with that seen in Fig. 5, as is the more pronounced frequency-selective fading.

In order to quantify the performance of DFH modulation in these two environments, we measured the bit error rate (BER) of the received signals for the two environments. Fig. 7 shows the average BERs for the two environments as a function of the number of simultaneously transmitting users. As expected, the BERs for Sandy Clay are lower than the BERs for Medium Sand, confirming the intuition that harder bottom types are more likely to cause ISI and consequently higher BERs. Also note that as the number of simultaneous users increases, the BERs increase in both environments. This is because in this multi-user simulated scenario we assume single-user demodulation: the receiver demodulates each user independently, without knowledge of other active users, which it effectively treats as interferers. Thus, the higher the number of interferers, the higher the BER. For the harder bottom type this effect is particularly evident, because a higher number of multipath arrivals from each user combine with the multipath arrivals from the interfering users, compounding the negative effect on the BER.

This simulation shows that the single-user DFH modulation scheme can be effective for underwater multi-user acoustic communications in relatively benign environments, as the BERs for Sandy Clay are below 10^{-3} for up to four simultaneous users, and well below 10^{-2} for five to eight users.

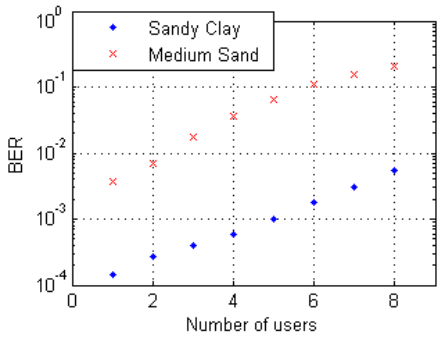

Figure 7. BERs as a function of the number of simultaneous users for the two environments simulated in SST.

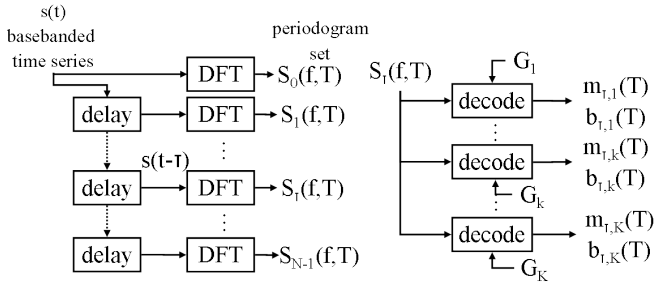

Figure 8. The DFH acquisition process. The first step is shown on the left: form periodograms at a family of hypothesized sample delays. The second step is shown on the right: attempt to decode each periodogram with each possible user code G_k

III. Demodulation

Demodulation of the DFH signal requires signal acquisition. Signal acquisition determines which users are active, as well as their respective sample delays and Doppler shifts. Unlike most other waveforms, DFH signal acquisition does not require a preamble or training symbols. Therefore, acquisition can be repeated as often as the variability of the channel demands, with no coordination required with the transmitter. Signal acquisition is accomplished in two steps (shown in Fig. 8): (i) form periodograms at a family of hypothesized sample delays and (ii) attempt to decode each periodogram with each possible user code G_k. Each decoder outputs a metric sequence as well as a bit sequence. From the metric sequence, active users are detected at each user's sample delay. This acquisition process can lag real-time, as long as the lagged data is buffered for demodulation once active users are detected at their corresponding delays.

Once users have been detected at their corresponding delays, the demodulation process consists of the following steps for each user (shown in Fig. 9): delay by the user's delay, form the periodogram using a DFT, and decode using the user's code. This process results in an efficient, auto-synchronizing, single-user demodulator. For the doubly-spread channel, Doppler effects can be handled in the same way as delay is handled in this simple implementation: hypothesize a family of Doppler effects and compensate for them before forming the periodogram.

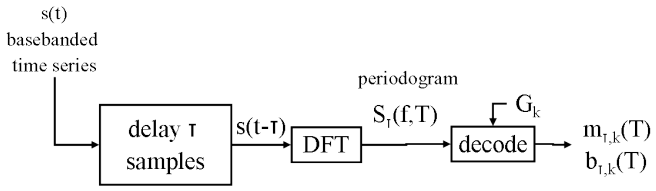

Figure 9. The single-user demodulator

To describe the decoder, we first define the following notation:

- $S_\tau(f,T)$ is the value at frequency bin f at hop interval T for the periodogram S calculated at sample delay τ;
- $f_T = G_k(f_{T-1}, b_T)$ shows the operation of the code G for the user k: the frequency transmitted at the current hop f_T is a function of the frequency transmitted at the previous hop f_{T-1} and the bit in the current hop interval b_T;
- $m_k(T)$ is the output metric sequence for the user k as a function of hop interval T;
- $m_k(T,f)$ is the intermediate metric state for the user k as a function of hop interval T over the vector of trellis states f.

For each destination state f, $m_k(T,f) = S_\tau(f,T)$ + the maximum of the following two candidates: $m_k(T-1,f_0)$ and $m_k(T-1,f_1)$; where f_0 and f_1 are defined by $f = G_k(f_1,1)$ and $f = G_k(f_0,0)$. This decision is lossless: if the two hypotheses were carried forward through the transmission, any metric value downstream of the smaller of these two options could never exceed the corresponding metric value downstream of the greater one. The output metric sequence $m_k(T)$ is selected as the intermediate metric state with the highest metric value over the states f. This selection can be made without increasing error d hops after the current hop interval T, where d is the depth of the user code G.

IV. Results

A. RACE08 Experiment Setup

The Rescheduled Acoustic Communications Experiment 2008 (RACE08) experiment was conducted 1-17 March 2008 in Narragansett Bay at the University of Rhode Island's Narragansett Bay campus, in water depths ranging from 9 to 14 meters. The surface conditions were primarily windblown chop, and the sound speed profile was approximately isovelocity, varying with the tides (primarily due to salinity changes) between 1450 and 1470m/s.

The experiment layout consisted of a four-element source vertical line array, a reference receiver element, and three receiver vertical line arrays. We present results using a single hydrophone from each of two of the receiver arrays: one 400m north of the sources and one 1km north of the sources. The four transmitter elements were treated independently, as separate users (or as a jammer). Six DFH user configurations were collected, as shown in Table I.

TABLE I

DFH USER CONFIGURATIONS TESTED IN RACE08

source type	config1	config2	config3	config4	config5	config6
ITC-1007	silent	user 1	user 1	silent	user 1	user 1
AT-12ET	silent	user 2	user 2	silent	user 2	user 2
AT-12ET	silent	silent	user 3	jammer	jammer	user 3
AT-12ET	user 1	silent	user 4	user 1	silent	jammer

One transmitter element is of a different type than the others, and transmitted 6dB lower than the other elements within the signal band. Users transmitting from the same element type transmit at the same power. The transmitter elements are separated by 60cm. The users' signals are delayed, in order to arrive asynchronously at the single-hydrophone receiver. The jammer consists of a very slow chirp, which crosses the band over the entire duration of the signal (60s), so that it acts like a tonal jammer within a particular hop interval. The users k are distinguished by distinct codes G_k. They occupy the same band (same hopset) and are uncoordinated in transmit time. They all have a bit rate of 68bps, transmitting a maximum of 3593 bits per 60s transmission interval. In addition, after the users' transmissions, a short LFM chirp (1s) across the 4kHz band (9-13kHz) was transmitted from one source element, for channel characterization. Note that information from this channel characterization is not used by the demodulator, nor in signal acquisition.

B. RACE08 Channel

A characterization of the underwater acoustic channel for RACE08 was done by computing CIRs calculated from the transmitted LFM chirps. The chirps spanned the same 4kHz bandwidth occupied by the DFH signals transmitted in the experiment, centered around 11kHz. Fig. 10 shows a typical CIR for the receiver placed 1km North of the sources (transmission file 0781355F10; from the top element of the source array to the element of the receiver array). Note the main peak at approximately 54.7 seconds; the other peaks are due to multipath. Peaks that are very closely spaced together, such as the main peak and the one at about -9dB, can cause frequency fading. Note also the peaks at approximately 55.3 and 55.9 seconds. These peaks are due to long-delay multipath caused by the geography of the experiment location. These multipath arrivals cause intersymbol interference.

C. RACE08 Results

The results from the RACE08 data (over a total of 2231 1-minute receptions) are shown in Tables II and III.

For unjammed transmissions, the maximum BER of almost 10^{-1} occurs on the low-power user while three other users are transmitting at 6dB greater power. The impact of increased range is minimal, as the results at 1km range are very similar to those at 400m range. The improvement in BER with increased range in the 4-user case can be attributed to reduced multipath at the greater range, indicating that the performance is limited

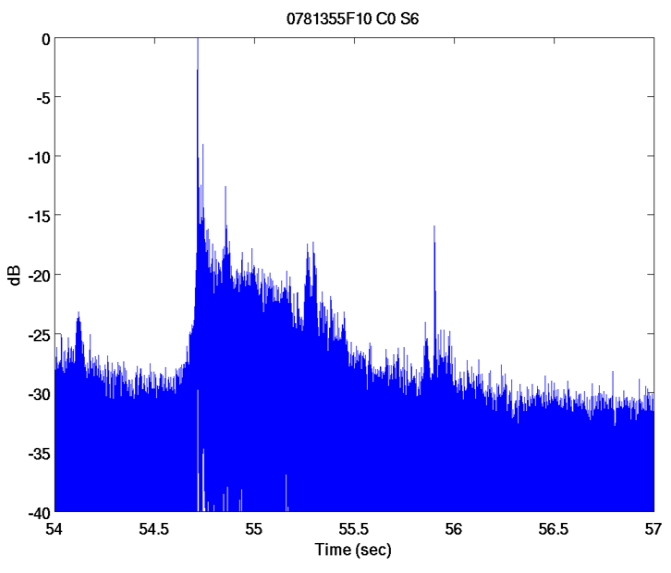

Figure 10. Example of a typical CIR for the RACE08 underwater acoustic channel, obtained by replica correlation of a LFM chirp, for reception file 0781355F10.

TABLE II

DFH RESULTS FOR UNJAMMED TRANSMISSIONS IN RACE08

source type	configuration	receiver distance from sources 0m	400m	1000m
AT-12ET	single user	0	0	0.35%
ITC-1007	user 1 of 2	0	0.95%	1.25%
AT-12ET	user 2 of 2	0.08%	0.03%	0.17%
ITC-1007	user 1 of 4	0	8.97%	8.26%
AT-12ET	user 2 of 4	0.12%	0.66%	0.36%
AT-12ET	user 3 of 4	0.44%	0.66%	0.53%
AT-12ET	user 4 of 4	3.94%	1.67%	1.09%

by multipath and MUI. The degradation in BER with increased range in the two-user case indicates that the performance is limited by attenuation, or SNR, especially as the results for user 2 of 2 compare favorably to the single-user case, with a BER near 10^{-4} at 400m and 10^{-3} at 1km. In addition, the different user configurations were collected on different days. Therefore, variation in performance could also result from varying channel conditions.

For jammed transmissions, the maximum BER of 15% occurs on the low-power user while two other users are transmitting at 6dB greater power. The impact of the jammer's presence is minimal at the reference phone, but more pronounced when offset from the source, indicating that the jammer's effect is compounded by the underwater acoustic channel. For both the single-user and two-user trials, the improvement with increased range is more pronounced than in the unjammed case, further supporting our conjecture that the performance in the presence of the jammer is primarily limited by its multipath interference, which is attenuated at greater range.

Also note the variation in performance across the users at

TABLE III

DFH RESULTS FOR JAMMED TRANSMISSIONS IN RACE08.

source type	configuration	receiver distance from sources 0m	400m	1000m
AT-12ET	single user	0	3.81%	1.27%
ITC-1007	user 1 of 2	0	5.16%	2.40%
AT-12ET	user 2 of 2	0.08%	3.88%	1.22%
ITC-1007	user 1 of 3	0	15.04%	14.69%
AT-12ET	user 2 of 3	1.49%	4.10%	3.90%
AT-12ET	user 3 of 3	2.30%	3.77%	4.31%

the reference phone near the sources: the receiver is closest to user 1, so a similar power diversity across the users occurs due to the variety in distance between source and receiver. We observe that the power diversity due to the different source types results in worse performance in multiuser scenarios, when the single-user decoder here described is used.

V. CONCLUSION

We have demonstrated, through simulation and collected data, the capabilities of the DFH waveform in the underwater acoustic channel. The simplicity of the transmitter and receiver, together with its ease of use in a multi-user environment, make the waveform a good candidate for computationally minimal, high reliability communications among uncoordinated users in an underwater acoustic channel.

Development is continuing on demodulator enhancements that improve performance with a minimal computational impact, focusing on mitigating multiple user interference and channel effects such as fading. In addition, a demodulator that takes advantage of a full receiver array will show better performance than these results, which use a single-hydrophone receiver.

ACKNOWLEDGMENT

The authors would like to thank Brinda Ramaswamy of BAE Systems and Warren Fox of BlueView Technologies. This work was supported by the Office of Naval Research under Contract N00014-07-C-0306 and under Grant N00014-07-1-0297.

REFERENCES

[1] W. Chung, C.R. Johnson, Jr., and M.J. Ready, "Characterization of Multipath Distortion of FSK Signals," *IEEE Sig. Proc. Letters*, vol. 9, no. 1, Jan 2002, pp. 26-28.

[2] R.P. Goddard, "The Sonar Simulation Toolset, Release 4.1: Science, Mathematics, and Algorithms," APL-UW TR 0404, Applied Physics Laboratory, University of Washington, Mar 2004.

[3] K. Halford and M. Brandt-Pearce, "Multistage Multiuser Detection for FHMA," *IEEE Trans. on Comms.* , vol. 48, no. 9, Sept 2000.

[4] D.L. Herrick, and P.K. Lee, "CHESS, A New Reliable High Speed HF Radio," *Proceedings of MILCOM96*, 1996.

[5] D.B. Kilfoyle and A.B. Baggeroer, "The State of the Art in Underwater Acoustic Telemetry," *IEEE J. Oceanic Engineering*, vol. 25, no. 1, Jan 2000, pp. 4-27.

[6] R. Kozick and B. Sadler, "Maximum Likelihood Multi-User Detection for Fast Frequency Hopping/Multiple Frequency Shift Keying Systems," Proceedings of WCNC2000, 2000.

[7] D.G. Mills and G.S. Edelson, "CHESS Study Final Report," report for DARPA and ARFL, Feb 2001.

[8] D.G. Mills, G.S. Edelson , and D.E. Egnor, "A multiple access differential frequency hopping system," *Proc. MILCOM 2003*, vol. 2, pp. 1184-9, 2003.

[9] D.G. Mills, D.E. Egnor, and G.S. Edelson, "A performance comparison of differential frequency hopping and fast frequency hopping," *Proc. MILCOM 2004*, vol. 1, pp. 445-50, 2004.

[10] J.G. Proakis, *Digital Communications*, 4th edition, Boston: McGraw Hill, 2001.

[11] W-B. Yang and T.C. Yang, "High Frequency FH-FSK Underwater Acoustic Communications: The Environmental Effect and Signal Processing," *High Frequency Ocean Acoustics*, M. B. Porter, M. Siderius , and W.A. Kuperman, eds., American Institute of Physics, 2004.

[12] W-B. Yang and T.C. Yang, "M-ary Frequency Shift Keying Communications over an Underwater Acoustic Channel: Performance Comparison of Data with Models," *J. Acoust. Soc. Am.*, vol. 120, no. 5. Nov 2006, pp. 2694-2701.