

Channel Effects on Direct-Sequence Spread Spectrum Rake Receiver During the KauaiEx Experiment

Paul Hursky*, Vincent K. McDonald[†], and the KauaiEx Group

* *Center for Ocean Research, SAIC, 10260 Campus Point Drive, San Diego, CA 92121*

[†] *Space and Naval Warfare Systems Center, San Diego, San Diego, CA 92152*

Abstract. Acoustic communications using a direct-sequence spread spectrum modulation with a RAKE receiver was tested during the KauaiEx experiment (in Hawaii). The oceanography was measured along the communications path during this test, so that the propagation of the communications transmissions could be subsequently modeled. Such simultaneous measurements provide a continuous picture of how the ocean waveguide evolved during the transmissions, enabling us to isolate what was influencing communications performance. This paper presents an analysis of DSSS Rake receiver performance at this shallow water coastal site.

INTRODUCTION

The SignalEx experiments (see [1]) have been conducted to gain an understanding of how various oceanographic phenomena influence the “channel” seen by underwater acoustic communications systems, and to further measure how these phenomena impact the performance of these systems. The ocean “channel” is a challenging one, because it is a fluctuating waveguide, with a rough surface in motion, an irregular bottom, and a volume in which wave speed can vary with depth and range. The severity of these effects and the relatively limited bandwidth in the underwater channel renders many of the techniques developed for terrestrial wireless or wire channels less effective.

A large repertoire of modulation schemes and receiver designs can be deployed to combat these channel impairments, all with parameters that potentially need to be tuned. This flexibility is also a curse, for it is difficult to predict in advance what particular setting suits a given environment. A goal of the SignalEx work is to develop the ability to predict communications performance, given a description of the oceanography at the targeted site and using acoustic propagation models, so that acoustic communications systems can be deployed at their optimal settings, turn-key, and with reasonably predictable performance.

Acoustic propagation modeling at low frequencies and for a “frozen” ocean (i.e. ignoring time-varying phenomena, such as surface wave motion), even with arbitrary range and depth dependencies, is well developed and successful in predicting complex propagation environments (e.g. witness the many successful demonstrations of

matched field processing). However, at higher frequencies, where ocean dynamics due to surface motion and water column sound speed fluctuations, boundary roughness, and air bubbles all have a more dramatic effect, modeling is not as mature. In fact, it is not clear what needs to be modeled to predict acoustic communications performance.

After briefly reviewing the DSSS-Rake receiver, we present results of testing this receiver in three different shallow water ocean sites. We conclude by demonstrating a channel simulator that demonstrates the sort of time-varying phenomena that must be included in a communications performance prediction, especially for receivers having tracking loops (the DSSS-Rake receiver has a delay-locked loop for tracking the symbol timing).

DIRECT SEQUENCE SPREAD SPECTRUM (DSSS) AND RAKE RECEIVER

Spread spectrum modulation techniques, as their name implies, use more bandwidth than strictly necessary to transmit a particular information bit rate, using their redundant bandwidth to either provide gain or to allow multiple users to share the same band. Conceptually, using extra bandwidth provides more “slots” for information than is needed for a given information bit rate. Frequency shift keying over a sparse “frequency hopping” sequence (with different users assigned different hopping patterns) is one form of spread spectrum modulation. Phase shift keying according to a pseudo-random noise (PRN) sequence (with each user assigned a different sequence) at a higher rate than the information rate is another form, called direct sequence spread spectrum (DSSS), or code division multiple access (CDMA). In this paper, we will focus on the DSSS/CDMA receiver in a single user, point-to-point configuration, testing a particular receiver design (described in [2]). Improvements to this basic receiver have also been described in [3].

In DSSS/CDMA, a single carrier is continuously transmitted, phase modulated by a bi-polar “spreading” sequence (of 1’s and -1’s, called chips). If L is the “spreading factor”, L chips are transmitted for every information bit. An information bit is conveyed by setting the sign of the block of L chips used to transmit it. In the next section, we will be presenting results of transmitting 4000 chips/second in an 8-16 kHz band with information bit rates of 50 bps (80 chips/bit), 100 bps (40 chips/bit), and 200 bps (20 chips/bit).

Much effort has gone into discovering PRN sequences that have good auto and cross correlation properties (see [4]), so that delayed versions of the same sequence (e.g. in the presence of multipath) or multiple different spreading sequences (e.g. transmitted by multiple users with distinct spreading sequences) will not interfere with each other. We will use Gold sequences, although the receiver we are using does not reap the full benefit of these sequences because each information bit is spread over only a subset of the entire sequence (ultimately, a consequence of the limited bandwidth we are using, relative to the information bit rates we have chosen to test).




Figure 1. Diagram of DSSS-Rake receiver.

Rake receivers (see Figure 1) were designed to exploit the above-mentioned good correlation properties of spread spectrum sequences by combining the copies of the transmitted signal, which arrive at different multipath arrival times. Because the signal has been “spread” (by a sequence with “good” auto-correlation properties), the overlapping arrivals only interfere with each other to a limited extent.

The Rake receiver consists of a bank of correlators (called “fingers” of the Rake receiver), each matched to the known PRN sequence. Each correlator is applied at a different delay in order to recombine or despread the multipath arrivals before “detecting” which information bit value was sent. The RAKE fingers are spaced to cover the time spread of the channel, or at least the subset of it that we seek to recombine.

Because each symbol period (i.e. information bit) contains only a subset of the sequence, and we cycle through the entire sequence as we move through multiple bits, we must initially synchronize our copy of the sequence with the received sequence. This is done by detecting three repetitions of a subset of the sequence. After this initial acquisition, a delay-locked loop (DLL) is used to track any subsequent drift, by adjusting the timing to minimize the amplitude difference between early and late versions of a dominant Rake finger output.

The receiver we are testing uses differentially coded information bits. This means we do not have to track carrier phase, but only detect how the phase changes from one symbol to the next (i.e. whether the current RAKE output is plus or minus the previous RAKE output, where outputs are produced at the bit rate). The taps of the RAKE receiver will constructively add to the detection output if their corresponding multipath arrivals have the same phase across two bit periods.

Note the recombination of the RAKE taps is only possible because overlapping delayed copies of the spreading sequence do not interfere with each other. Overlapping copies of a single carrier system that has not been spread would interfere with each other.

To summarize our expectations of the receiver being tested: we expect that the DLL will track a slowly varying time offset between the transmitter and receiver, and we expect that the multiple taps of the RAKE receiver will provide some gain in a multipath environment, if the multipath arrival structure is stable over a symbol period. We have put these expectations to the test using a receiver array in shallow water spanning most of the water column, so that we get a comprehensive picture of how performance varies over depth, time, and the accompanying oceanographic changes.

PERFORMANCE AND ANALYSIS OF EXPERIMENT DATA

Figure 2 shows the experiment configuration off the coast of Kauai and the time-varying impulse response at three depths, measured by applying a matched filter to LFM chirps (sweeping from 8-16 kHz in 50 ms, repeated every 250 ms). The strongest arrivals were observed near the bottom, with the second group of arrivals showing progressively more compact structure as we move from the surface to the bottom.

Figure 3 shows the array geometry, the sound speed profile, and bit error rate versus depth and time for three bit rates (50, 100 and 200 bps). The performance was measured over two diurnal cycles and nearly the entire water column.

The performance was uniformly better near the bottom. As Figure 2 shows, most of the energy is in the second group of arrivals, and the concentration only gets stronger as we move from the ocean surface to the ocean bottom. As seen from looking at the sound speed profile in Figure 3 (and as described in [5]), this concentration of energy is due to a thermocline that extends almost to the bottom, where a strong duct is formed, focusing these later arrivals.

The concentration of errors peaking at hour 30 seems to be correlated with the changes in sound speed profile, in particular the degree of ducting near the bottom. When the duct is weaker, due to mixing in the water column, the second group of arrivals is no longer so concentrated, and the receiver seems to have less of a dominant signal to lock onto.

All the bit error rates shown are without error correction (which would probably recover fully from bit error rates up to 12 percent or so). Arguably, the 100 and 200 bps tests yielded marginal performance (too many instances of bit error rates beyond any error correction repair). This is a consequence of the relatively small spreading factors (40 for 100 bps and 20 for 200 bps), which determine the matched filter gain. The amount of multipath in the channel in which the receiver was tested requires more gain to overcome the multipath interference. Increasing the bandwidth would allow longer PRN sequences (with more gain) to be used, without sacrificing the information bit rates. Adding an equalizer and a chip-rate timing loop, as suggested in [3], are other potential remedies, although they are more computationally demanding.




Figure 2. Kauai experiment configuration (upper left), and measured impulse response functions at three depths: near surface (upper right), middle of water column (lower left), and near the bottom (lower right). Only a small portion of the impulse response, around the earliest arrivals, is shown so that the fluctuations in arrival time can be observed.




Figure 3. Upper left figure shows array element placement relative to sound speed in the water column. Bit error rates versus array element and time (45 h total) are also shown: at 50 bps (upper right), 100 bps (lower left), and 200 bps (lower right) in Kauai.

CHANNEL SYNTHESIZER FOR PERFORMANCE PREDICTION

Our approach for a channel synthesizer is to start with a frozen ocean model, using the best available sound speed profile for the targeted site, and to calculate arrival times and complex amplitudes using the Bellhop gaussian beam propagation model. We have had good success modeling the high frequency channel with such an approach, demonstrating matched field processing in the 8-16 kHz band (see [6]). Besides time of travel, amplitude and phase information, ray (and gaussian beam) models typically provide information about how many times individual arrivals have interacted with the boundaries. This additional information enables us to add fluctuations the otherwise static paths, for example striving to reproduce fluctuations caused by interactions with a moving surface. As the measured impulse response functions in Figure 2 show, sometimes this is as simple as providing a sinusoidal variation whose amplitude and frequency match the wave motion of the ocean surface (as in the impulse response near the bottom). Other times, such surface interactions are clusters of arrivals with no easily identifiable track (as can be seen in the impulse response near the surface).




Figure 4. Channel simulation: “frozen ocean” impulse response function (upper left), simulated and estimated time-of-arrival showing Doppler shift caused by source and surface motion (upper right), differential output of Rake fingers (lower left), DSSS constellation color coded to indicate known values of information bits (lower right). At higher Doppler rates, when surface motion reinforces source motion, we see fading on the Rake finger outputs and the resulting closing of the eye pattern in the constellation as the high phase rate causes time spreading of the matched filter outputs.

Figure 4 shows an example of testing the DSSS Rake receiver against a synthetic channel, in which the transmitter drifted toward the receiver (this constant Doppler produces a mild ramp in the arrival times), and the path being tracked by the timing

acquisition loop (i.e. the DLL) corresponded to a surface path having a sinusoidally varying arrival time. The superposition of these two motions resulted in arrival times having a downward trajectory with a sinusoidal oscillation. Note how the constellation (shown in the lower right subplot) expands and contracts, following the Doppler acceleration shown in the tracking loop. This causes bit errors when the SNR drops, due to the Rake matched filters not being able to “follow” the Doppler accelerations.

This is a simple example of why realizations (and not just averaged quantities, such as spectra) must be synthesized for the multipath fluctuations, especially for simulations of receivers that have tracking loops. Furthermore, each path may have a distinct Doppler process, depending on which part of the water column it has traveled through.

Our channel synthesizer relies upon resampling the transmitted waveform, realizing a true time-varying time dilation (and not just a phase shift: our 8-16 kHz band does not qualify for a narrowband approximation) corresponding to a Doppler trajectory. Furthermore, each path arrival is given a potentially distinct Doppler trajectory, so that arrivals that have interacted with the ocean surface at different locations have suitably independent trajectories, and arrivals that have interacted with the surface more than once have greater “spread” than paths that have interacted only once.

What remains to be done is constructing more realistic models for the surface, volume and bottom interactions (i.e. changing bathymetry as the communicating platforms move around the site) that can be driven by measurements of the prevailing oceanography and weather (e.g. wind speed, sea surface temperature, etc).

CONCLUSIONS

We have tested the DSSS-Rake receiver at depths spanning the water column for several tidal cycles while simultaneously measuring the prevailing oceanography along the propagation path from source to receiver. Not surprisingly, we can correlate bit error rates with many of the oceanographic phenomena at both sites. We have discussed how using PRN spreading sequences having the touted “good cross-correlation” properties in the DSSS-Rake receiver does not realize the gains implied by these properties, because only subsets of the full sequences are used for any single information bit. As a result, given the typical channel spreads in underwater channels, only modest rates can be reliably achieved with the DSSS-Rake receiver, unless the bandwidth is dramatically increased.

We have demonstrated a channel synthesizer that accurately reproduces the prevailing “frozen ocean” multipath and also captures a variety of time-varying channel effects, including continuously and rapidly varying Doppler at surface-interacting paths and discontinuous times of arrival due to volume and rough boundary scattering. As an example of why it is important to use actual realizations of time-varying channels (and not just spectral averages that do not capture the detailed phase trajectories of the various multipath components), we have shown how the Doppler due to surface motion can cause sporadic outages in the DSSS-Rake receiver. Note that a “frozen ocean” channel would not stress the various tracking loops that are

typically used in receivers for various phase-modulated signaling schemes, especially higher-rate phase-coherent schemes.

ACKNOWLEDGMENTS

This work was supported by ONR through the ONR High Frequency Initiative and the SignalEx project.

The original code for the DSSS-Rake receiver was obtained from Ethan Sozer and John Proakis with whom we had many fruitful discussions on acoustic communications. We also benefited from later work on this receiver by Michael Porter.

The data used to test this receiver was collected during the KauaiEx experiment, with Michael Porter as chief scientist, and the KauaiEx group, all contributing ideas and resources. The KauaiEx Group is: Michael B. Porter, Paul Hursky, Martin Siderius (SAIC), Mohsen Badiy (Univ. Delaware), Jerald Caruthers (Univ. Southern Miss.), Daniel Rouseff, Warren Fox (Univ. Washington), Christian de Moustier, Brian Calder, Barbara J. Kraft (Univ. New Hampshire), Keyko McDonald (SPAWARSYSCEN), Peter Stein, James K. Lewis, and Subramaniam Rajan (Scientific Solutions).

REFERENCES

1. M. B. Porter, V. K. McDonald, P. A. Baxley, and J. A. Rice, "SignalEx: Linking environmental acoustics with the signaling schemes," *Proceedings of MTS/IEEE OCEANS'00 Conference*, 2000, pp. 595–600.
2. E. M. Sozer, J. G. Proakis, M. Stojanovic, J. A. Rice, A. Benson, and M. Hatch, "Direct sequence spread spectrum based modem for under water acoustic communication and channel measurements," *Proceedings of MTS/IEEE OCEANS'99 Conference*, pp. 228 - 233, September 1999.
3. M. Stojanovic and L. Freitag, "Hypothesis-feedback equalization for direct-sequence spread-spectrum underwater communications," *Proceedings of MTS/IEEE OCEANS'00 Conference*, pp. 123 - 129, September 2000.
4. D. V. Sarwate and M. B. Pursley, "Cross-correlation properties of pseudorandom and related sequences," *Proc. IEEE*, vol. 68, no. 5, pp. 593-619, 1980.
5. Martin Siderius, Michael Porter and the KauaiEx Group, "Impact of thermocline variability on underwater acoustic communications: results from KauaiEx," *Proceedings of HF Acoustics Conference*, La Jolla, California, March 1-5, 2004.
6. Paul Hursky, Martin Siderius, Michael B. Porter, and Vincent K. McDonald, "High-frequency (8-16 kHz) model-based source localization", *J. Acoust. Soc. Am.*, vol. 115, no. 6, pp. 3021-3032, June 2004.