CHANNEL IMPULSE RESPONSE ANALYSIS AND SIMULATION FOR ACOUSTIC MODEM EXPERIMENTS IN THE SINGAPORE STRAITS

Dale Green
Popoto Modem
Sandwich, Massachusetts, USA
daleg141@gmail.com

John DellaMorte
Popoto Modem
Sandwich, Massachusetts, USA
john@popotomodem.com

Jim DellaMorte
Popoto Modem *Affiliation)*
Sandwich, Massachusetts, USA
jim@popotomodem.com

Mandar Chitre
Acoustic Research Laboratory,
National University of Singapore
mandar@nus.edu

Charles Bernstein
Naval Surface Warfare Center Panama City
Panama City, Florida, USA
charles.bernstein@navy.mil

Daniel Sternlicht
Naval Surface Warfare Center
Panama City
Panama City, Florida USA
daniel.sternlicht@navy.mil

***Abstract*— A series of experiments were conducted in the Singapore Straits to test the performance of several acoustic communication waveforms and to capture time-changing Channel Impulse Responses (CIRs). The straits present a challenging shallow (20m) acoustic environment with background noise from vessels and snapping shrimp, and multipath resulting from anchored vessels, multiple small islands, and buoys in the water space. Data was collected with multiple receivers on a mooring line at fixed depths, monitored from a work boat, and a single transmitter submerged from a second work boat. Data was collected with the receivers and transmitters in static configuration, or with the transmitter drifting. A novel method for capturing the time-changing CIR over small time scales is described and the results of the experiment detailed.**

Keywords—CIR, waveform, correlator.

I. INTRODUCTION

An underwater acoustic communications (ACOMMS) experiment was conducted in July 2022 in the Singapore Straits as a combined effort between Popoto Modem, The National University of Singapore, and the United States Navy. The purpose of the experiment was to evaluate/demonstrate the performance of a variety of modulation methods and the utility of channel probes to predict and explain the signaling performance. The tests described in this paper occurred at various distances between 500 m and 2 km between a moored transmit source and a moored collection of receiving/recording devices. The transmitter was a modem provided by Popoto Modem, with a center frequency of 29 kHz, and the signal bandwidths were approximately 8 kHz. This paper deals with recordings from three Popoto modems suspended on a line at depths of 6, 10, and 12 m in 18-20 m deep water.

Among the waveforms tested were a frequency hopping (FH) variation of JANUS [1], which included a 32-chip Frequency Hopping (FH) precursor used for acquisition and temporal alignment. The individual chips were 10 ms in duration. The other waveforms included both a 100 ms x 8 kHz Hyperbolic Frequency Modulated (HFM) probe and a 1.5 sec x 8 kHz Quadrature Phase Shift Keyed (QPSK) probe. This paper discusses the uses of these probes to estimate the input SNR and to help characterize the channel.

II. ACQUISITION PROCESS

The various signaling schemes described above each support a waveform component suitable for coarse, low overhead acquisition and temporal positioning, and another suitable for high precision temporal and Range Rate (RR) alignment.

For FH signaling we rely on the 32-chip FH acquisition portion for both aspects. That is, we first search for the signal using non-coherent processing which is insensitive to RR, then follow with coherent matched filtering of the same data for the precision positioning. For the other signaling schemes, we employ, first, the HFM chirp, which is insensitive to RR, followed by our QPSK probe, which is extremely sensitive to RR. We emphasize that the initial "coarse" acquisition does not involve multi-hypothesis RR testing, which is solely reserved for the refined, matched filter processing.

In addition to the above described uses of these probes, we also use them to estimate the input Signal to Noise Ratio (SNRi) based upon the measured output SNR (SNRo) from the probe filters. For FH signaling, we measure the output SNR as:

$$SNRo = ((P-M)/\sigma) \quad (1)$$

Where P is the peak of the non-coherent output, and M and σ are the mean and standard deviation of the same output in the absence of the signal. The SNRi is estimated to be

$$SNRi = SNRo(b/W)/N^{1/2} \quad (2)$$

Where: b is the chip bandwidth, W is the operating bandwidth, and N is the number of chips.

For matched filter processing, such as used with the chirp waveform, (1) applies, with P, M, S evaluated from the magnitude squared matched filter output. The input SNRi is estimated according to (3):

$$SNRi = SNRo/(TW) \quad (3)$$

III. Channel Estimation and Modeling

The FH signals include a 32-chip acquisition portion which is used for acquisition, estimation and correction for RR, and Vernier arrival estimation,. A by-product of these functions is to produce an estimate of the CIR. Fig. 1 shows the CIR estimates for 20 transmissions from day 1 of the experiment. The x-axis shows 0-6 ms of time, while the y-axis is a linear scale of the magnitude-squared of the matched filter output (power). We observe that the duration of the CIR is, on average, approximately 3 ms. The use of the acquisition portion for CIR estimation is not ideal because of significant sidelobes in the autocorrelation function, but it is "free" in the sense that no other probe is necessary.

A broadband, pseudo-random probe signal was used in support of the Multi-Frequency Shift Keyed (MFSK) and QPSK signaling. This 1.5 second waveform was built to serve two purposes: as a channel probe, and as the training vector for our QPSK modulated message. To fit the latter requirement, it was constructed in a manner identical to that used for the modulation. Each baud period consisted of exactly 4 samples of a single complex number (phase) at a baseband sample rate of 25600 (complex) samples/sec. The baud duration therefore was 0.15625 ms, and the signal had a nominal bandwidth of 6.4 kHz. The waveform was passed through a loose low pass filter to reduce sidelobes. No other shaping of the chips was used (e.g., raised cosine). The Peak-to-Average Power Ratio (PAPR) was approximately 1 dB.

For use as a probe, we generated 57 subcorrelators from the probe by extracting overlapping segments, each 100 ms in duration, each overlapped with its neighbor by 25 ms, as the design only allows for 1.5 s minus 100 ms for building the probes. The purpose is to correlate each subcorrelator with the received probe to reveal phase and power (or amplitude)

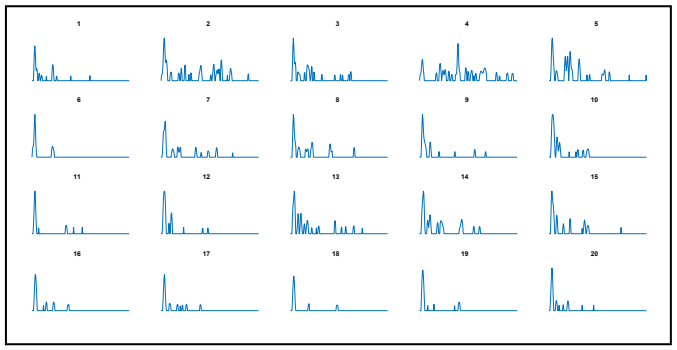

Fig. 1. CIR estimates made from matched filter processing of the 32-chip acquisition portion of the FH waveform. The x-axis reflects 0-6 ms, and the vertical scale is a linear scale of power.

changes each 25 ms. The lag and time indices δt are the same: If fs is the baseband sample rate, then δt = 1/fs. Figure 2 shows an Average White Gaussian Noise (AWGN) example, with the upper plot showing the first 5 subcorrelator outputs for the same data stream. The lower plot shows all 57 subcorrelator outputs, now aligned in time and truncated to 4 ms lag time following the peak.

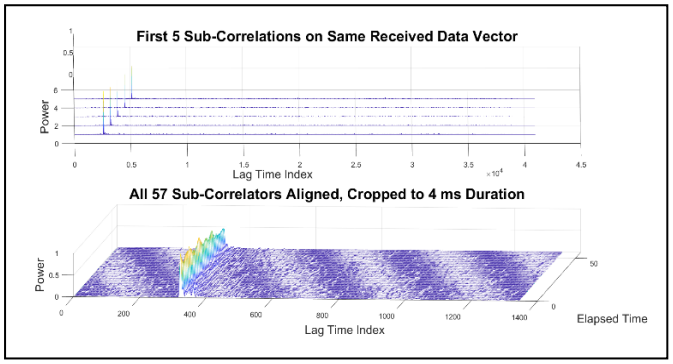

Fig. 2. Use of correlators.

In an ideal channel, the output of the subcorrelators, at the peak power locations, will be identical in both phase and power. We sampled the received data for all of the 20 transmissions of probe data to measure the phase at each subcorrelator. Fig. 3 shows how the unwrapped phase angle of the dominant path changes across the 1.4 seconds of the probe, sampled every 25 ms, for the pings marked at the end of each curve.

Fig. 3 shows that the phase angles are not constant, as they would be from an ideal channel, although the phase rotations are approximately linear and quite different among pings. We surmise that the observed phase rotations are caused either by the channel or by platform motion.

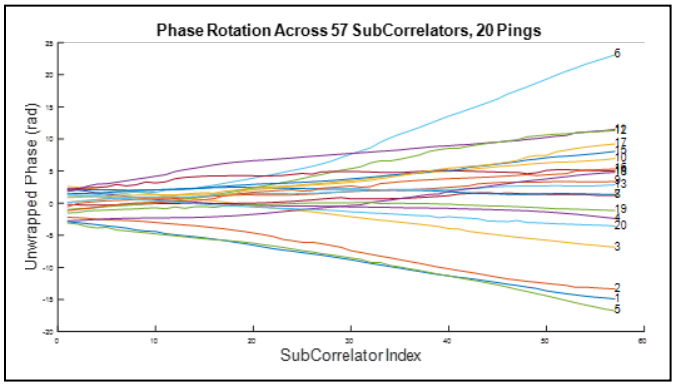

Fig. 3. 57 measurements of phase angle sampled every 25 ms for each of 20 different pings.

As an example of probe processing applied to at-sea data, we show in Fig. 4 the waterfall displays for all 20 transmissions during one test period using the probe for the QPSK signal. The vertical axis is power on a decibel scale, while the x-axis covers approximately 4 ms of lag time. The remaining axis is time, expressed as 25 ms intervals across 1.4 seconds.

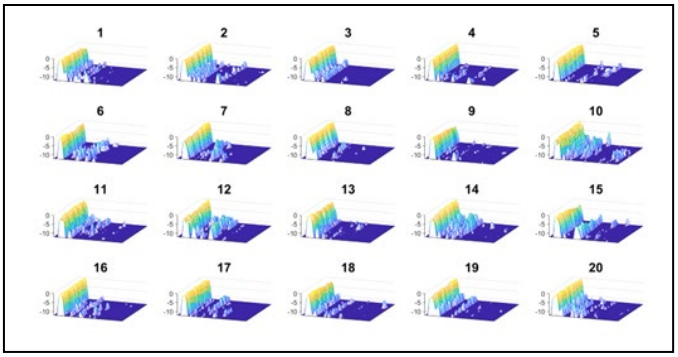

Fig. 4. 20 Examples of CIRs from one test opportunity.

We next explored one of the pings in detail, using the data to construct a time-varying model of the channel and applying that model in a simulation mode for both the original transmitted probe, and a much shorter Linear Frequency Modulated (LFM) signal. Fig. 5 shows the analysis process for ping 10 from Fig. 4. The upper left plot is the same as shown in Fig. 4, with more details provided.

Our process finds and follows across time and lag each distinct peak in the display. The upper right plot of Fig. 5 shows the result of such "peak-picking". We then collected the data into specific tracks, with amplitude data shown in the lower left-hand plot, shown here for 8 tracks. In the lower right-hand plot we show the phase angles for each of strongest tracks, with measured speed (m/s) for each track written on the plot.

Speed V is measured by fitting a straight line to the lag indices for a given path, measuring the lag difference dL at the start and end of the line, then solving the equation:

$$V = dL \cdot C_S / T \tag{4}$$

where C_S is the sound speed (~1500 m/s) and T is the pulse duration (maximum of 1.4 seconds).

The track and phase angle data are measured at 25 ms intervals. We next perform a spline interpolation of each, for all paths. Fig. 6 shows the amplitude and phase angle interpolations, respectively.

The interpolated data can be used to impose channel conditions on any waveform with the same duration, or less. As an example, assume the waveform is the same transmitted probe used in the Singapore experiments. We convert the phase angle vector to a complex phase vector using an exponential transformation and multiply the waveform by this phase vector. We multiply the result by the amplitude vector, and delay the product by the measured delay of the track (relative to the delay time of the first peak in Fig. 5, upper-left plot). If the measured track speed is other than zero, we time-compress/dilate the track. We sum all of the path-delayed and scaled waveforms together to build a composite received signal. We may then return the resulting signal to the analysis routine to reproduce a presentation similar to that of Fig. 5, as shown in Fig. 7 for the 10th transmission, showing the Impulse Response (IR) over time and lag.

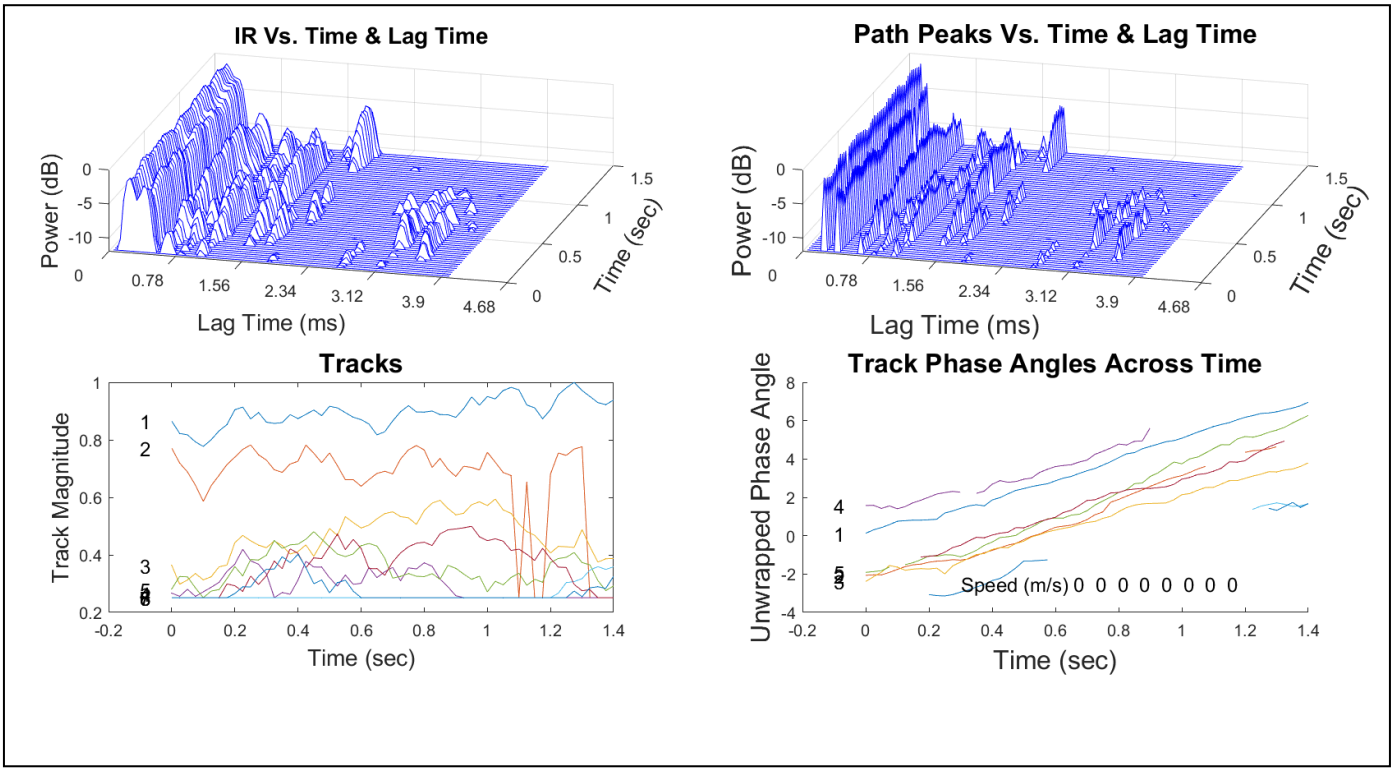

Fig. 5. Steps in analyzing the measured time-varying CIR.

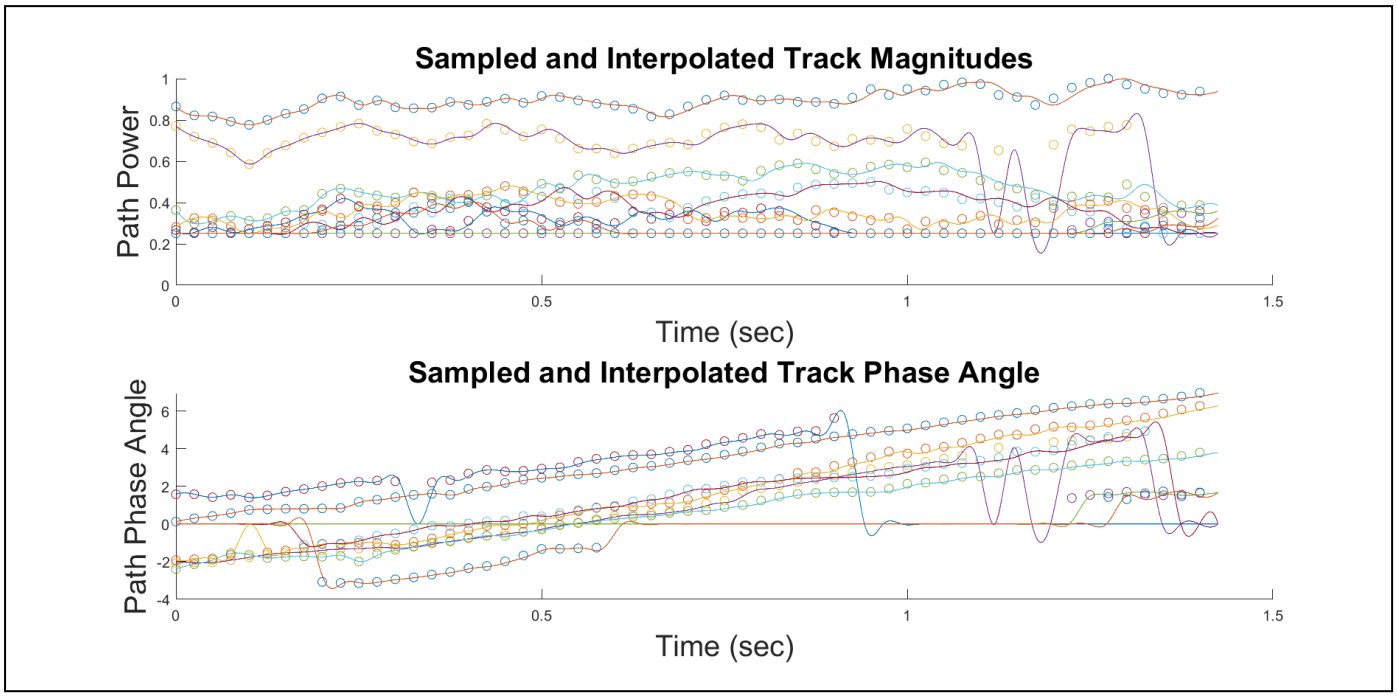

Fig. 6. Amplitude and phase angle using spline interpolations. Circles indicate the measured data (ower two plots of Fig. 5) and the solid lines are the interpolations.

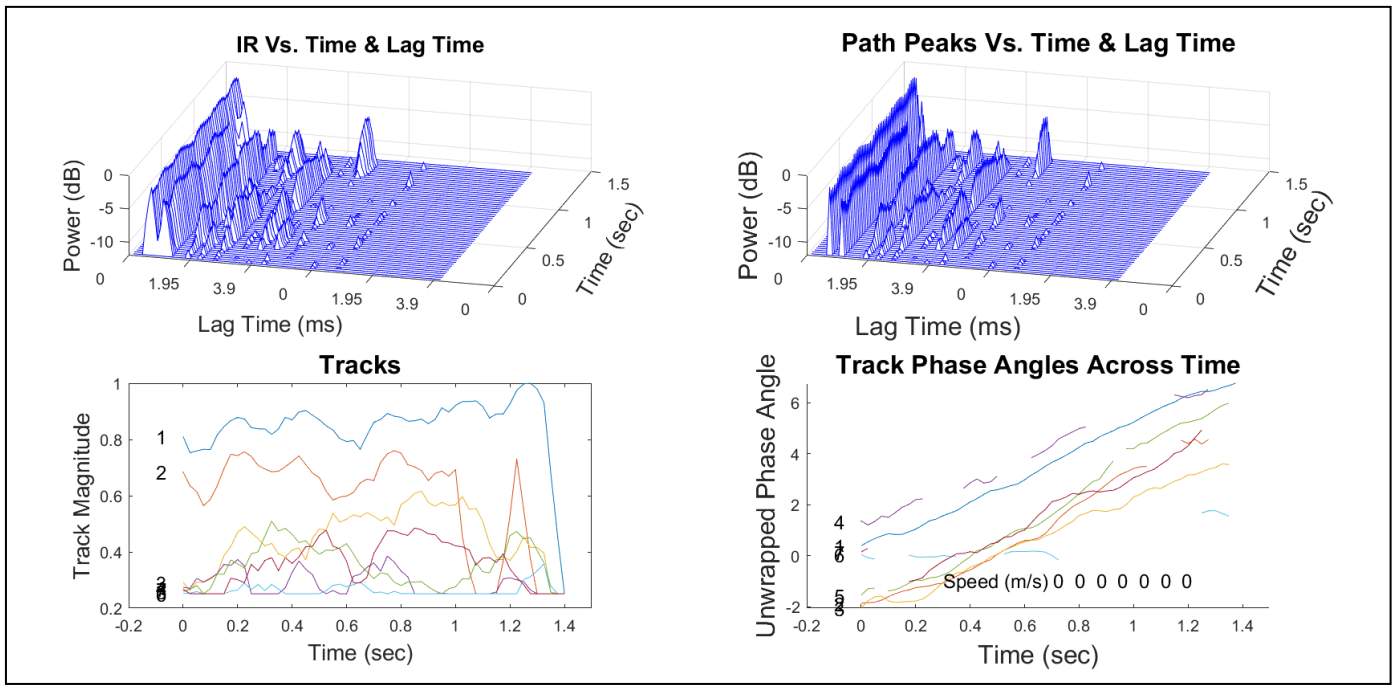

Fig. 7. Analysis of the modeled received signal reflecting ping 10 from the actual at-sea data.

Consider as a second example a much shorter waveform, specifically a 100 ms LFM chirp built using the same bandwidth and center frequency. As the interpolated tracks are much longer than this waveform we randomly extract portions of the tracks, each with a duration equal to that of the LFM waveform, and subject the LFM to the amplitude, phase, and delay characteristics of the channel. To observe the results we employ a replica correlator with each "received" waveform, with results shown in Fig. 8.

One further result of the analyses of the probe data is the significant variation in received signal power over the 20 pings. Fig. 9 is a series of plots showing the peak power variation for each modulation event for a single test period. These variations are observed during all events. This demonstrates that transmitting a much more limited series of pings might have severely mischaracterized performance.

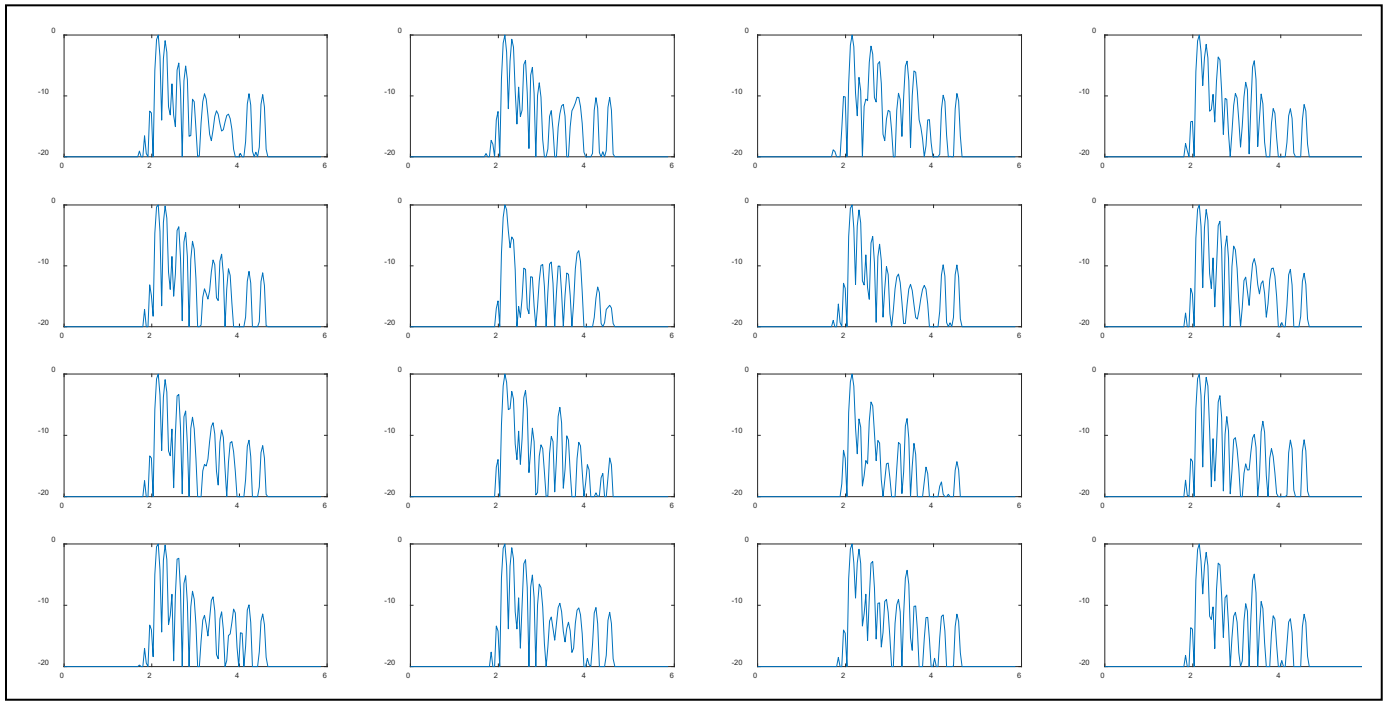

Fig. 8. Twenty realizations of "Channelized" LFM waveforms which have been subjected to the interpolated tracks developed with the broadband channel probe, each "received" LFM is evaluated via replica correlation, resulting in the plots shown here.

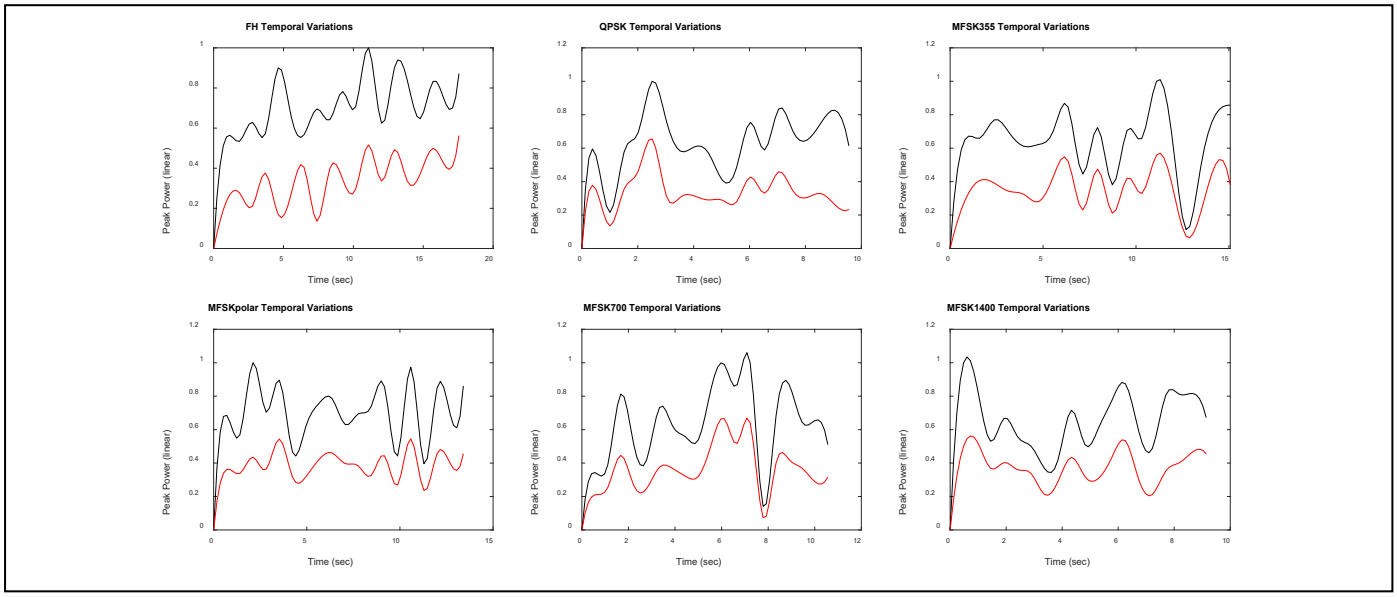

Fig. 9. Peak power variations across 20 received pings. Red curve is power from the single dominant correlation peak, black curve includes added power from all significant peaks

IV. Conclusion

We have presented a method for characterizing a time-varying underwater acoustic channel using a long-duration QPSK probe. The received probe is analyzed using overlapping, quasi-independent short-duration replicas extracted from the transmitted probe. The transmit probe has bandwidth W and duration T, while the short replicas have duration $T_r << T$ and overlap adjacent replicas by $\tau_r < T_r$. For this experiment we arbitrarily chose $T_r = 100$ ms, and $\tau_r = 25$ ms, but we note the following: a) a longer T_r improves sidelobe suppression and gain against additive noise, and b) longer T_r increases the integration interval which might mask rapidly time-varying channel events. Further, the overlap among replicas is only constrained by the requirement that $\tau_r > 1/W$ for statistical independence, although small τ_r will result in more replicas and require more computation. However, nothing precludes applying multiple sets of replicas with varying characteristics to a received time series. We did not do this with the present data set. We also observe than any pseudo-random sequence may be used in place of our QPSK, which was selected for this experiment because it also forms the training vector for our QPSK modulation.

The analysis of the probe data provides a description of the time-varying amplitude, phase, and RR of individual CIR paths, which then permits construction of modeled "received" versions of arbitrary waveforms of comparable (or smaller) bandwidth and duration < T.

V. References

[1] J. Potter, J. Alves, D. Green, G. Zappa, I. Nissen, K. McCoy, "The JANUS underwater communications Standard," UCOMMS 2014 Sestri Levante, Italy, September 2014.