Tones for Real: Managing Multipath in Underwater Acoustic Wakeup*

ISI-TR-659

Affan A. Syed
USC/ISI
asyed@isi.edu

John Heidemann
USC/ISI
johnh@isi.edu

Wei Ye
Broadcom Corporation
weiye@broadcom.com

Abstract

The principles of sensor networks—low-power, wireless, in-situ sensing with many inexpensive sensors—are only recently penetrating into underwater research. Acoustic communication is best suited for underwater communication, since other methods (optical and radio) attenuate very quickly. Yet acoustic propagation is five orders-of-magnitude slower than RF, so propagation times stretch to hundreds of milliseconds. A new generation of underwater acoustic modems have added low-power wakeup tones that combat the energy consumption acoustic modems would waste on idle listening. Recently, these tones have been used as an integral part of application layer and MAC protocols. While all wireless data-networks suffer from multipath interference of received data, in this paper, we show that due to large acoustic propagation delay tone echoes cause a unique interference, tone *self-multipath*, for tone-based protocols. To address this interference we introduce *Self-Reflection Tone Learning* (SRTL), a novel approach where nodes use Bayesian techniques to discriminate self-reflections from noise and communication from other nodes. We present detailed experiments using an acoustic modem in two physical environments to show that SRTL's knowledge corresponds to physical-world predictions, that it can cope with reasonable levels of noise, and that it can track a changing multi-path environment. Simulations confirm that these real-world experiments generalize over a wide range of conditions.

1 Introduction

Sensornets are transforming for science and industry by enabling pervasive, in-situ sensing through inexpensive, wireless sensors. This success has sparked interest, bringing sensing with these characteristics underwater to improve our ability to chart the oceans, lakes, and waterways that strongly influence our environment and can provide natural resources [32, 17, 9].

Perhaps the most significant to changes when deploying sensornets underwater is the use of acoustic instead of radio-frequency-based (RF) communication. Radio communication is significantly attenuated underwater, making it impractical for distances over a meters or so. (Low-frequency RF is impractical because it requires a very large antenna and significant energy to operate.) While optical links can provide high-speed communication over short-range, point-to-point links, it requires careful aiming and so cannot fulfill the need for easy deployment as is possible with surface sensornets.

Underwater acoustic communication poses several challenges. The most serious is large propagation delays. With the speed of sound around 1500m/s underwater, and ranges of 100 to 10000m, delays of 100ms are common and multiple seconds are possible. The large delay makes adapting traditional media access protocols difficult, since channel sensing time follows propagation delay, a simple CSMA MAC will consume a great deal of power listening to an idle channel. In addition, the underwater channel poses additional challenges, including temperature-based delays, ducting and significant problems due to multi-path interference. Several papers summarize these challenges [3, 26, 19]

Commercial underwater acoustic modems today primarily target long-range (multi-km) point-to-point communication (for example, the Benthos modem [1]), and recent research efforts have targeted higher throughput using coherent communication [6, 24]. Since energy is a constraint in many underwater networks (for example, stationary networks, or those using battery powered gliders [22])), some modems employ special wake-up tones and low-power wakeup receivers to activate modems [1, 33], and recent work has shown how to integrate tones with a low-power MAC protocol [30] . Tone-based wakeup is also important for wakeup after long-duration sleep [21, 16], or triggering more sophisticated data reception algorithm [1, 6, 33]. Replacing full packet reception with tone detection (detecting energy on the channel) is a very effective way to save energy since it can be done both very quickly, as seen in low-power listening with radios [4, 18]; and very efficiently with dedicated hardware, as seen in radio-based pagers and recent underwater modems [33].

*This research is partially supported by the National Science Foundation (NSF) under the grant NeTS-NOSS-0435517 as the SNUSE project and grant number CNS-0708946, "Open Research Testbed for Underwater Ad Hoc and Sensor Networks", and by Chevron Co. through USC Center for Interactive Smart Oilfield Technologies (CiSoft).

Multipath interference is a problem in all wireless communication. For data, it causes inter-symbol or inter-chip interference. These problems have been extensively studied in radios [20, 14]. Long propagation makes multipath worse underwater, and only in the 1990s was coherent communication demonstrated underwater [27]; managing data multipath is still an area of research [23].

While a problem for data communication, multipath can be crippling for tones. When tones are used for coordination, echos reflecting from stationary objects cause the transmitter to hear *self-reflections* of their tones. If tones indicate contention and their absence indicates a channel clear for transmission (as it does for T-Lohi MAC protocol [30]), self-reflections will *prohibit* all communication, since a sender will always mistake the echo for another contender! We explore this scenario in more detail in Section 2 where we use the T-Lohi MAC protocol to illustrate tone-based self-multipath. This paper is the first to identify the problem of self-multipath for low-power underwater acoustic communication.

The challenge of self-multipath is that a transmitter will always interfere with itself. We show that this challenge provides the means to address the problem: senders can and must *learn their self-multipath patterns*. This goal is difficult because of the unique challenges of underwater sensornets. First, current data-focused multi-path techniques do not directly apply because we are concerned with self-multipath, not sender-receiver multipath. Also wakeup tones are simple signals and are therefore not amenable to sophisticated coding. Finally we must allow for long acoustic propagation times and reflections to hundreds of milliseconds. In Section 3 we describe *Self-Reflection Tone Learning* (SRTL) where we use Bayesian techniques to learn the channel state. In addition to directly applying to underwater MAC protocols, SRTL's ability to estimate the number of reflecting surfaces in the environment also indicates the sparsity of the channel. This information may assist optimizations that exploit sparsity in managing the complexity of multipath in data reception [15].

Since wireless channels are very difficult to model or simulate, we demonstrate the effectiveness of our approach with experiments using an acoustic modem in three different physical environments in Section 4. We establish that what SRTL learns matches to physical-world predictions through tests in a completely controlled anechoic chamber (Section 4.2). We show that SRTL copes with reasonable amounts of noise through tests there and in a less controlled laboratory, and with simulations that let us vary noise precisely (Section 4.3). Our controlled tests require in-air acoustics; underwater tests are currently underway. Finally, we show that SRTL can track changes to the environment (Section 4.4), how environmental parameters and their estimates affect the accuracy of algorithm (Section 5.1 and 5.2), and finally that SRTL remains effective with low-complexity, low-memory implementations (Section 5.3)

The contributions of this paper are to identify self-reflections as a new problem posed in low-power, high-latency underwater communication; to show how Bayesian learning with SRTL can mitigate this problem; and finally to show that it works experimentally. We believe this approach will be essential to energy conserving underwater media access protocols, and also useful for long-duration sleep, and may apply to the broader problem of multipath in high-latency channels.

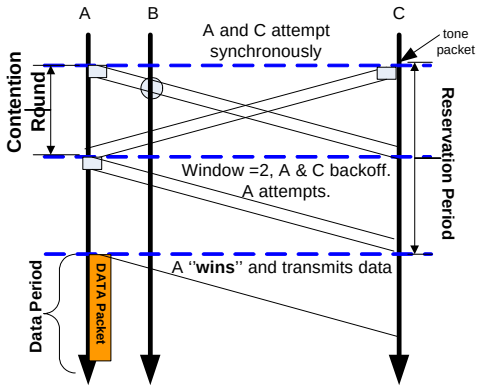

Figure 1: An example T-Lohi protocol contention exchange

2 T-Lohi MAC: Embracing the Propagation Latency of Acoustic networks

Tone-Lohi (T-Lohi) is a contention-based MAC protocol for underwater acoustic networks [30] . Its key features are that all contention is done with short wake-up tones, that can be observed with sub-mW energy consumption [33] and that it converges quickly due to accurate estimates of the number of contenders [30] . As a result, overall energy consumption is minimal even in the face of contention periods that may last up to a second. However, we will show that echos due to self-reflection will compromise these benefits. We next briefly summarize this protocol and then describe the problem self-multipath poses for it.

2.1 Overview of T-Lohi

The primary objective of T-Lohi is to provide a MAC protocol that has efficient channel utilization, stable throughput, and low energy consumption. It is contention-based, so that nodes can adapt to varying traffic patterns. A full description of T-Lohi can be found elsewhere [30] ; we summarize it here to show how tones are key to its energy-efficient operation.

In T-Lohi, nodes contend to reserve the channel to send data. It requires that nodes first send a *short tone* and then listen for the duration of the *contention round* to decide if reservation is successful. If only one node contends in a contention round, it wins, ending the *reservation period* and then transmitting its data. When nodes detect another concurrent tone, they interpret it as another node's contention and extend the reservation period by randomly backing-off. As tones are short and distributed over a long contention round (due to space-time uncertainty in high-latency communication [28]) each contender can do *contender counting* by counting the tones received. This count gives a good estimate of the number of other contenders, allowing the channel to quickly converge by adapting back-off in proportion to the number of contenders. Figure 1 shows an example of

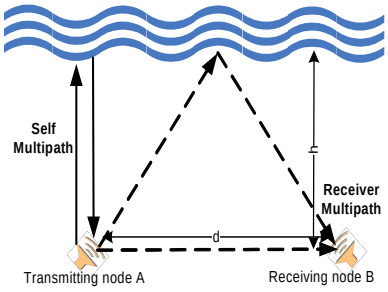

Figure 2: receiver- and self-multipath shown for a single reflecting surface at a distance h from transmitting node A. Similar multipath occurs for each reflecting surface.

this process: nodes A and C have data to transmit but first send tones indicating contention. At the end of the first contention round both A and C have a count of two and backoff to attempt uniformly in one of the next *two* rounds. If no other tone is detected in a given round (like A does not in round two), collision free data transmission occurs in the subsequent round.

Three significant sources of energy consumption in media access are idle listening, collisions, and control overhead [34]. T-Lohi addresses these sources of waste with two different mechanisms. First, in lightly used networks such as many sensornets, the majority of time is spent waiting for traffic, or idle listening [34]. Terrestrial sensornets have reduced this cost with scheduling [34] and low-power listening [4, 18]. In T-Lohi, we exploit very-low power dedicated hardware for tone reception, where tone listening consumes only 500μW [33], Second, to reduce the cost of collisions and control overhead, T-Lohi converges very quickly to a collision free channel reservation when traffic occurs. Quick convergence is possible because of contender counting, allowing the MAC protocol to converge in a small constant number of rounds with high probability, instead of the $O(\log n)$ convergence form binary-exponential backoff (as in Ethernet and 802.11).

2.2 Impact of Multipath

We now show how tone echos result in self-multipath that cripples the T-Lohi protocol.

Figure 2 shows multipath occurring between two contenders, A and B. Tone sent by A reaches B directly and, a little later, via a surface reflection. This regular, or receiver-multipath, will cause multiple tones to be detected at B.

From the perspective of T-Lohi, the additional tones increase the contention count. However, as our simulation and analysis of reservation period show the duration of reservation period does not increase significantly with network density [29] . Therefore, although the throughput decreases due to slightly longer duration before a packet is sent, T-Lohi will successfully contend the medium to send data.

Other than the multi-path reflections received at B there is also the echo that A receives of its own tone; this reflection causes what we call *self-multipath* as it interferes at the transmitter(Figure 2) understanding of the contention status. As opposed to traditional multipath that results in data interference at receivers, self-multipath is essentially echo-interference amplified in the acoustic channel due to large propagation delays. Moreover, while the multipath spread (the maximum delay between direct and multipath signal) of receiver-multipath is proportional to the *difference* in the path taken by a direct and reflected signal (typically 10s of msec for 500m range) self-multipath is proportional to just the delay to a reflecting surface (and thus larger than just the difference, typically 100s of msec for 500m range).

This self-multipath breaks T-Lohi MAC in such a way that contending nodes are *never* able to transmit data. This is because a *self-reflection* [1] for a contention tone sent in any contention round will result in an echo-induced tone detection. Since contenders transmit data in T-Lohi only when no other tone is detected, even a single reflecting surface, results in a contender *always* hearing an echo tone that it interprets as another contender. Thus a contender will always backoff and never be able to transmit data.

While one could work around this problem, perhaps by timing out after non-convergence, or by sending information with the tone, these approaches would increase energy consumption. Instead, we next show how Bayesian techniques allow contenders to *learn* about self-reflections, then choose to *ignore* them.

3 SRTL: Learning to Ignore Echos

We now introduce the Self-Reflection Tone Learning (SRTL) algorithm, our approach to manage self-reflection. We first give an overview of SRTL, then review Bayesian learning, the theory we draw upon (Section 3.2). We then cover SRTL details: what it observes about the channel, sources of error in those observations, and how these factors come together.

3.1 Overview of SRTL

The goal of SRTL is to learn about self-reflection and allow higher-level protocols (such as T-Lohi) to distinguish between echos, random background noise, and tones sent from other contenders.

The key intuition behind SRTL is that echos are deterministic and repeatable relative to a transmission, while noise and tones from other senders are independent and uncorrelated. An example of this observation is shown in Figure 3. When a node transmits a tone that reflects off other surfaces (A and B in Figure 3), the echos return to the sender with the same delay after each transmission (Figure 3). The tone receiver will also trigger in response to ambient noise (as shown by lightening bolts in Figure 3), or other transmitters, but these triggers occur *independently* of transmission times (compare location in Figure 3(a) and 3(b)). SRTL can therefore learn repeated echos using Bayesian learning techniques (reviewed next in Section 3.2).

SRTL makes two assumptions: echos are repeatable, and other triggers are independent. Since echos are dependent on the physical placement of nodes and reflecting surfaces, they will be repeatable in a static network. We confirm both of these assumptions, showing in Section 4.2 that the echos

[1] We refer to each echo as self-reflection, distinct from the broader protocol-interference concept of self-multipath.

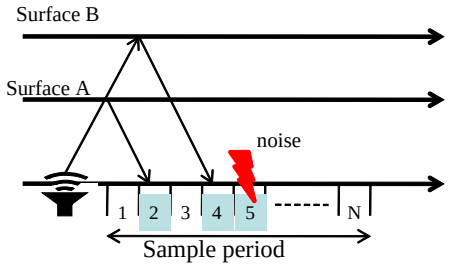

(a) First Sample: Two echo and one noise-triggered tone detection.

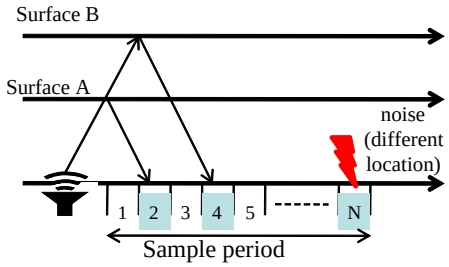

(b) Subsequent Sample: Echo locations repeat, but noise at different location.

Figure 3: Key idea behind the sample collection process for SRTL algorithm: A node sends a tone and waits for a *Sample period*. Bin location of echos repeat but that of non-echo detections do not.

we observe correspond directly with the physical environment. We tolerate minor variation in echo delay by discretizing responses into bins; in Section 5.3 we show our choice of bin size is reasonable and that we tolerate responses on bin edges. Although SRTL is optimized for a static network, we show in Section 4.4 that it can accommodate environmental change due to drift or slow movement.

The second assumption is that other (than echos) tone triggers are independent of transmissions. Most observations confirm that underwater noise is uncorrelated and random, although the noise distribution is different for shallow and deep water channels [3, 26]. We have confirmed this observation in the experiments at the Marina del Rey harbor.

It is possible that higher-level protocols or applications would create synchronization, either intentionally or accidentally [5]. While some variations of T-Lohi intentionally synchronize transmissions similar to slotted ALOHA, MAC protocols in general (and the variant of T-Lohi that we use) can explicitly randomize transmissions to guarantee protocol-level synchronization does not occur, and thus concurrent application-level transmissions are desynchronized at the MAC-level.

Learning about echos is not the only approach to identify self-reflection. We consider related work in detail in Section 6, but two alternate approaches are to use signal processing techniques to correlate echos and transmitted signals for disambiguation, or to code sender identification in tones. For example, Rake receivers handle data multipath [20] using chip-correlations, but because its complexity is proportional to delay, it doesn't easily scale to high-latency acoustic communication. Similarly Girod and Estrin use a similar approach in acoustic localization [7]. Coding sender identification in a data packet can also solve the issue of self-multipath. Using data packets, however, precludes the energy benefit of tone wake-up. Alternatively, we could use time-based codes, pulsing the tone on and off. However, to reach low power, the low-cost tone wake-up circuit uses very simple analog electronics and so is unable to use sophisticated codes or provide bit synchronization guarantees. Also, to maximize sensitivity, the wakeup receiver requires long activation times (up to 5ms in the worst case). Thus, the requirement to minimize energy consumption precludes many alternate approaches. We therefore turn to learning, and next provide background about Bayesian inference as employed by SRTL.

3.2 Bayesian Inference Background

Bayesian Inference is a well established approach that measures a probabilistic belief or knowledge regarding an event or hypothesis [12]. (Note that probability here represents belief in a hypothesis, not probability as frequency of occurrence.) We next briefly introduce the Bayes' theorem and show how it can be extended to perform inference based on empirically collected data.

The classical Bayes' theorem as defined for a bimodal hypothesis H, incorporating current evidence X is:

$$P(H|X) = \frac{P(X|H)\,P(H)}{P(X)} = \frac{P(X|H)P(H)}{P(X|H)P(H)+P(X|\overline{H})P(\overline{H})} \quad (1)$$

We assume, in all cases, either the hypothesis holds (H) or is not true ($\overline{H}$). The *prior probability*, $P(H)$ is the confidence before considering current evidence, while $P(H|X)$ is the confidence after observing X. $P(X|H)$ is the *conditional probability*, the likelihood that event X implies that hypothesis H is true, while $P(X|\overline{H})$ is the likelihood X occurs even though the hypothesis is *not* true. $P(X)$ is the probability of witnessing the new evidence X under the two mutually-exclusive hypotheses.

The factor $\frac{P(X|H)}{P(X)}$ represents the *impact that evidence has on belief in the hypothesis*. When evidence strongly indicates the hypothesis, this factor is large, but if evidence is inconclusive, perhaps because the environment is noisy, it will be small. This ability to incorporate evidence lends naturally to using Bayesian inference for interactive sensing.

The Bayes theorem describes how a single observation modifies a belief. We build on work in landmark-based localization in robotics [31] to use successive observations to learn about the environment. This work uses robots with sensors that identify landmarks and move or change their environment. Observations of landmarks represent evidence increasing belief in the current location; manipulation and movement change the environment and decrease belief. Changes in belief are scaled by models of accuracy of sensing and actuation. Next, we describe how SRTL "senses" echos, and how we model its accuracy to scale the belief update, thus learning echo locations.

3.3 Sampling in SRTL

To apply Bayesian reasoning, and inspired by work in robotics (Section 3.2), we must decide how to sense the environment, and how our samples correspond to our hypothesis. Each time we send a tone, we follow that transmission with a *sampling period* (as in Figure 3). The duration of this period is defined by the maximum range of our tone hardware. Since transmissions attenuate, we listen until any additional echos would be too faint to detect.

To manage observations, we divide the sampling period into fixed-duration *bins*. For each bin i, we track the hypothesis H representing belief that the i^{th} bin corresponds to a self-reflection; we call such bins self-reflection (SR) bins. For each bin the hypothesis $\overline{H}$, that a bin does *not* have echos, is simply complementary to H due to bimodality of the hypothesis space.

After a transmission and the entire sampling period, we have an array of evidences, one sample per bin. Each sample can take on two values, either E_i, indicating a tone detection (an apparent echo), or $\overline{E_i}$, indicating absence of any detection in that bin.

Bayesian learning has a rich mathematical background for estimating hypothesis by incorporating empirical data. While we currently model static nodes, the incremental Bayesian learning can also incorporate motion if an appropriate model is provided [31]. Thus we believe that a Bayesian learning approach is appropriate for learning tone-echoes.

3.4 Modeling Truth and Observations

A transmitter's observation corresponds to four possible real-world events: True echo detection, true silence detection, Non-echo Detection (ND), and Tone Cancellation (TC). The first two events are the accurate observations about the world from the point of view of the transmitter. However, we model the next two events because they represent error introduced into our observations.

Non-echo detection corresponds to an incorrect observation where our tone detector triggers, but it is not due to an echo of our transmission. Channel noise is a common source of a ND. A ND can also be caused by a valid tone transmission from another node. We cannot distinguish between these sources of incorrect observations, but our approach can filter both out, since both are effectively random sources of noise (as justified in Section 3.1).

A second source of error is *Tone cancellation*. While unlikely, it's possible that channel noise or another node's transmission can interfere with and cancel reception of a tone. In these cases, we are unable to observe an echo that should be there. Missing tones also occur when a tone lies on the boarder of a bin as we discuss in Section 5.3.

We model the event ND and EC with parameters p_{nd} and p_{tc}. These parameters are engineering estimates of the probabilities that these events will occur in for a given hardware and physical environment. We show in Section 5.1 that our algorithm tolerates a wide range of parameters, so they need not be set exactly, but its estimate tracks reality best when the parameters closely approximate the actual probability of these events.

3.5 The SRTL Algorithm

We now apply Bayesian learning to our system. We apply our algorithm in parallel to the hypothesis H and $\overline{H}$ for each bin, using the samples array observed in each sample period. For simplicity we stop using the subscript and describe a single bin. We refer to E and $\overline{E}$ as positive and negative evidence of a tone activation in that bin. We next describe how we initialize our algorithm, then how we substitute E and $\overline{E}$ into X from Equation 1 to update, as we learn, $Bel(H)$ which is our confidence in the hypothesis, then finally how we make decisions from our observation.

3.5.1 Initializing the Algorithm

We initialize the algorithm with $Bel(H)^{init}$, representing our initial belief of the bin being a self-reflection (SR) bin. This value seeds the Bayesian algorithm and is the same for all bins. Subsequent samples update the current belief using the update equations we describe next.

From the perspective of Bayesian learning, the initial value should reflect our assumptions about the environment, presumably from prior experiments for a given transmitter. Since in general we do not have such knowledge, we instead start with an arbitrary value of $Bel(H)^{init} = 0.3$. The algorithm is largely unaffected by choice of initial value because it rapidly aligns with reality using experience; Section 4.4 confirms that we can quickly track environmental changes, and we have confirmed in experiments (omitted here due to space) that we are largely unaffected by the initial value.

We next explore how negative and positive evidence change our estimate $Bel(H)^{posterior}$ and mark bins as self-reflecting.

3.5.2 Update for Negative Evidence

We consider absence of a tone as negative evidence ($\overline{E}$) that indicates a bin does not receive echos. We therefore update our estimate from Equation 1, replacing X with $\overline{E}$. Also, we replace the term $P(H|X)$ with $Bel(H)^{posterior}$ (belief after incrementally incorporating current evidence) and $P(H)$ with $Bel(H)^{prior}$ (current belief incorporating all prior evidence) to reflect the standard Bayesian inference terminology [31].

$$Bel(H)^{posterior} = \frac{P(\overline{E}|H)Bel(H)^{prior}}{P(\overline{E}|H)Bel(H)^{prior}+P(\overline{E}|\overline{H})Bel(\overline{H})^{prior}}$$

We next explore how this update equation for negative evidence is modeled using our parameters p_{nd} and p_{tc}.

$P(\overline{E}|H)$ is the conditional probability for the event when no tone is detected, given that we *know* an echo tone should be detected. This event is essentially the failure of our tone detection hardware to be triggered in the presence of tone energy. Assuming that bin duration is small enough to allow only a single detection, such an event can happen only if the tone echo is canceled by noise or interference (the event $TC|H$, which reduces to TC since TC is independent of the hypothesis H). This assumption simplifies our modeling as we no longer consider the more complicated case when an additional tone detection occurs, in the same bin, after the first. Thus, we can now define:

$$P(\overline{E}|H) \quad \equiv \quad p_{tc} \tag{2}$$

$P(\overline{E}|\overline{H})$ is the probability for the event when no tone is detected given that we know that no echo can occur in that bin. However, this knowledge does not rule out non-echo detections caused by noise or other transmitters. Thus the event can be described by the *union* of two disjoint events: no non-echo noise is detected (the $\overline{ND}$ event) union with the event that although wake-up-triggering noise could have been detected it was canceled (the $ND \bigcap TC$ event). Noting that the above events are disjoint, we formulate the following probability:

$$P(\overline{E}|\overline{H}) \quad \equiv \quad (1-p_{nd}) + p_{nd} \times p_{tc} \tag{3}$$

Finally, using the definitions from Equation 2 and 3 the update equation for negative evidence becomes:

$$Bel(H)^{posterior} = \frac{p_{tc} Bel(H)^{prior}}{p_{tc} Bel(H)^{prior} + (p_{nd} p_{tc} + 1 - p_{nd}) Bel(\overline{H})^{prior}} \tag{4}$$

This update is applied to our current belief $Bel(H)^{prior}$ to realize a new belief $Bel(H)^{posterior}$ when no tone is detected within that bin (negative evidence).

3.5.3 Update for Positive Evidence

Detection in any bin during a sample period is considered positive evidence (E) of the bin being a SR bin. Similar to the negative update, we update our belief according to Equation 1, replacing E for the update event X. We next define the components of the update equation for positive evidence based on the input parameters p_{nd} and p_{tc}.

$P(E|H)$ is the conditional probability for the event that tone detection occurs in a bin with known self-reflection. This probability is simply the complement of $P(\overline{E}|H)$; Thus:

$$P(E|H) \quad \equiv \quad 1 - p_{tc} \tag{5}$$

On the other hand, $P(E|\overline{H})$ is the conditional probability for a detection occurring in a bin we know has no echos. This detection can, therefore, occur only because of wake-up-triggering noise (ambient noise and other contention tones) that is *not* canceled; the event $ND \bigcap \overline{TC}$. Thus this probability becomes:

$$P(E|\overline{H}) \quad \equiv \quad p_{nd}(1-p_{tc}) \tag{6}$$

Finally, using the definitions from Equation 5 and 6 the update equation for positive evidence becomes:

$$Bel(H)^{posterior} = \frac{(1-p_{tc}) Bel(H)^{prior}}{(1-p_{tc}) Bel(H)^{prior} + (p_{nd}(1-p_{tc})) Bel(\overline{H})^{prior}} \tag{7}$$

This update is applied to our current belief $Bel(H)^{prior}$ to realize a new belief $Bel(H)^{posterior}$ when a tone is detected within a bin (positive evidence).

Table 1: Payoff table used in determining decision threshold that maximizes payoff.

	Reality	
Decision	*SR bin*	*not a SR bin*
SR bin (Ignore tones)	10	2
not a SR bin (Count tones)	1	9

While Bayes works as described in theory, in our experiments we observed that long runs of consistent evidence would saturate the bins with perfect positive or negative belief ($Bel(H) = 1$ or $Bel(H) = 0$). We expect saturation occurs because of floating point rounding error. These equations are unable to shift from certainty, even in the face of later contrary evidence. We therefore cap the belief for each bin at a maximum value of 0.999 and a minimum value of 0.0001 to avoid saturation.

3.5.4 Identification of SR Bins

We intend for the SRTL algorithm to work continuously in the background to our MAC protocol (or any other networking protocol using tones). In T-Lohi, we automatically get one sensing period for each contention round, so observations about the environment occur automatically as a side effect of MAC operation, incurring no additional overhead. We next evaluate how to translate these continuous observations into decisions about which bins represent self-reflections and therefore should be ignored.

To reach a decision, we bias our estimates by the payoffs of correct or incorrect decisions, then select the most profitable. We believe this decision threshold is reasonable, and we show later (Section 5.3) that our assumptions about the environment (p_{nd} and p_{tc}) have much stronger influence on correctness. For the values given in Table 1, we can derive a fixed threshold of 0.45. Thus if the $Bel_i(H)$ is greater than 0.45 we will consider that to be a SR bin. We use this threshold in our experiments. In principle, one could adapt these values to the environment.

3.5.5 Selecting Bin Duration

Our algorithm uses fixed size bins; bin granularity is one factor to the sensitivity of our algorithm. In practice, bin size is limited by hardware. As a lower bound, our microcontroller has a millisecond level clock granularity, and interrupt debouncing causes a 2ms delay between successive tone detections. We therefore set bin size conservatively at 3ms in both simulations and experiments.

4 Experimental Evaluation of SRTL

We now evaluate SRTL through experiments, both in the controlled setting of an anechoic chamber and a less controlled open lab. Table 2 summarizes our research questions, but our overall goal here is to show that SRTL can successfully manage echos. To do that, we first confirm SRTL's conclusions are justified by the physical environment (Section 4.2), and evaluate how to copes with different levels of noise (Section 4.3). Finally we verify that the algorithm adapts to changes in environment, either due to movement of the node or other objects (Section 4.4). We begin by summarizing our experimental methodology.

Table 2: Research questions asked about the merits of our Bayesian learning algorithm.

Questions asked about SRTL	Section	Environment		
		Controlled	Uncontrolled	Underwater
Correctly ID known surface?	4.2	Yes	n/a	n/a
Is robust to Noise?	4.3	Yes	Yes	Planned
Can handle dynamic environment?	4.4	Yes	n/a	n/a
How sensitive to parameters?	5.1	Yes	Yes	Planned
What's the impact of discrete bins?	5.3	Yes	n/a	n/a

(a) Transmitter Setup

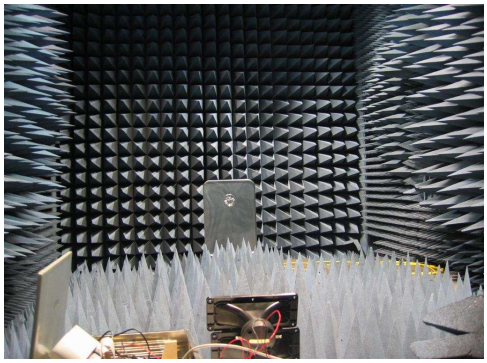

(b) Reflecting Surface

Figure 4: **Controlled Experiments:** The Setup in the anechoic chamber. On the left was the transmitter/receiver setup and the right figure shows a reflecting surface whose location was varied in our experiments.

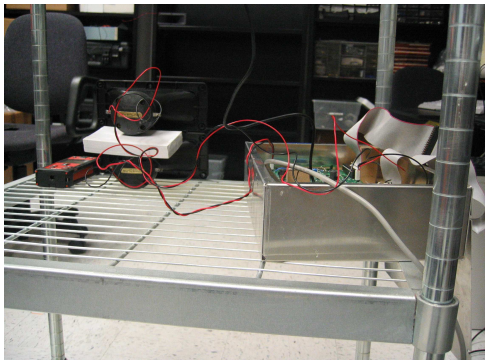

(a) Lab/office test setup

(b) Underwater test setup

Figure 5: **Uncontrolled Experiments:** Uncontrolled experiments were performed at two locations. The lab/office location provided for in-air experiments, while the test setup off the docks in Marina del Rey harbor provided for underwater experiments.

4.1 Experimental Methodology

We evaluate SRTL using our acoustic modem in different environments: a controlled, anechoic chamber and a less controlled laboratory, both using in-air acoustics. In addition, we are currently carrying out underwater tests. We next describe details in common to the three environments we report here, how they differ, and bounds on our ground truth.

4.1.1 Details common to experiments

We run experiments using the SNUSE acoustic modem [33], hardware revision 2. We use tweeters for in-air tests, and hydrophones when underwater. The modem is driven by a custom data collection program running on a Mica-2 mote. The microcontroller directly controls modem operation via I/O through a custom digital interface.

Each experiment consists of 200 sample periods of tone transmission followed by echo observation. For each sample, the mica2 configures the modem to transmit a wake-up tone, then switches to *tone-sleep* where it is quiescent until woken up by a tone. We timestamp each tone reception on the Mica-2 with 1ms resolution, then compute delay between initial transmission and echo. We later map detection delay into a corresponding 3ms bin (Section 3.5.5).

The SRTL algorithm currently runs in a host PC connected to Mica-2, although in principle it could run on the mote itself. After each transmission we record all tone triggers as positive evidence (E), and assume negative evidence

($\overline{E}$) for all other bins. We then update SRTL belief estimates based on Equations 4 and 7.

We set the sensing duration based on the maximum observed in-air range of our modem. We measure in-air range at 20m, so we anticipate reflections from objects up to 10m from the transmitter. We therefore anticipate echos arriving with up to 60ms delay (20m, with speed of sound as 343m/s at 24° Celsius). We conservatively extend sensing duration to 100ms after each transmission.

We next describe details specific to our three experimental locations.

4.1.2 *Location-specific experimental details*

We carried out experiments at three locations, each providing us with different level of complexity in the reflective nature of the environment.

The first experiment location is an *anechoic chamber*at USC's UltraLab Laboratory. The chamber is designed to absorb all RF radiation for controlled radio experiments, but it also provides a good acoustically neutral environment. When necessary we place a large metallic pan in the chamber to act as a reflecting surface (Figure 4). We measure the distance from our transmitter as described in Section 4.1.3. Since the anechoic chamber is designed to be reflection free, this configuration lets us confirm against a strong ground truth: the physics of the measured location of our reflecting surface.

Our second environment is an *office laboratory* (Figure 5(a)). This location is much more complex, with multiple possible reflecting surfaces (walls, file cabinets, machine rack doors, etc.). We therefore observe more complicated channel response and so cannot provide firm ground truth. However, this more complex environment provides a richer level of noise and signal response.

For both in-air tests (anechoic chamber and laboratory), the modem uses high efficiency, Motorola piezoelectric tweeters that were impedance matched for both transmission and reception.

Our final experimental environment explores underwater performance. We test at the docks in Marina del Rey harbor (Figure 5(b)). Our current experimental modems are not yet packaged for underwater use, so we operate them on dock, connected to a Benthos AT-18AT hydrophone lowered about 1m underwater off the docks. Our initial underwater experiments were not successful because the wake-up receiver was too sensitive, receiving near-continuous activations from background noise. We have recalibrated it and are in the process of collecting additional data underwater.

4.1.3 *Estimating ground truth*

We estimate ground truth based on the physical distance between transmitter and reflector and compare this distance to measured echo distance (converted from measured echo delay). Both these measurements, however, have potential sources of error due to our modem hardware and the measurement process.

The largest source of error in measurement of echo delay is the detection circuit of our modem. We time-stamp the transmit time of tone and detection time of echos to calculate the distance to reflector. Due to transmit side warm-up, the actual transmission time of the tone can vary by about 1ms.

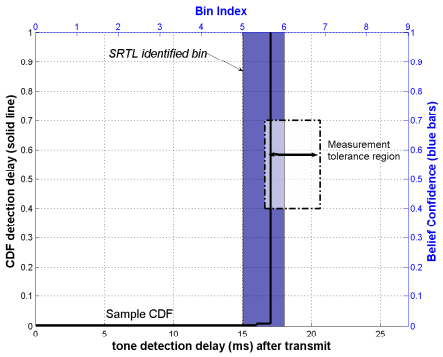

(a) Reflecting Surface at 3.36m

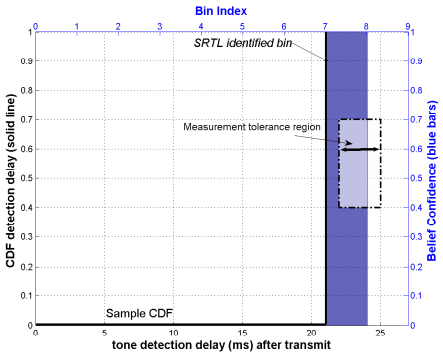

(b) Reflecting Surface at 4.11m

Figure 6: Experimental results showing the CDF of 200 samples and SRTL response with objects at known location, adjusted for measurement error (shown as the shaded area with dashed boundary).

Similarly the actual detection time can vary by 2ms based on the strength of echo.

We measure the physical distance with a HILTI PD-40 high-precision laser range-finder [11]. Accuracy is $\pm$1mm, so we believe error in the distance measurement is minimal. However, the most significant source of error is in our measurement process when we approximate the angle for the line-of-sight measurement between the piezoelectric crystal located inside the transducer and the reflecting surface. This measurement error is approximately $\pm$2cm and results in a corresponding delay error of about $\pm$0.5ms (with speed of sound as 343m/s at 24° Celsius).

When comparing the ground truth to identified echos, we have to reconcile the above independent errors, in both measured distance and echo location. In the figures, we show the tolerance region that accounts for the worst case error in each measurement. The identified location can safely be considered to match the ground truth if these regions overlap. Thus, due to the resulting overlap of the error bounds, the tolerance region varies on a case-to-case basis for each measurement.

4.2 SRTL Correctness

We first seek to confirm that SRTL can *correctly* identify the location of a known reflecting surface: does our algorithm and experimental setup match the physical configura-

tion of the world?

Since this experiment needs knowledge of the ground-truth, we use the anechoic chamber to perform controlled experiments. For this experiment we place a reflecting surface perpendicular to the piezoelectric tweeters, measure its distance and compute the expected delay. We take several measurements (as described in Section 4.1) with the surface at a particular location. Since there is only one reflecting surface in the room, our algorithm should only identify the bin corresponding to the measured distance. We then compare SRTL's estimate with our prediction from the physical distance.

Figure 6 shows the result of our experiments for reflectors at two different distances. Each figure combines three different values: prediction from physics, all observations, and the SRTL belief distribution. The dotted box indicates the prediction from our distance measurements, including estimated error. The solid black line represents the cumulative distribution function of delay values for all the samples considered by SRTL, measured against the left axis. Finally, solid blue bars represents the belief distribution ($Bel_i(H)$) for each bin in the 100ms sensing period (we show only the first 25ms and omit the remainder since there is no belief there). Bin indices are given on the top axis.

Figure 6(a) shows the result of our experiment with the surface measured to be at 3.36m from the transmitter. The sample CDF shows that nearly all samples are received at a delay of 17ms, which corresponds to the fifth bin. SRTL is able to identify this bin with complete confidence and we can see that the identified bin lies within the error bounds of the physically measured surface location.

Figures 6(b) shows results from the same experiment with the surface at 4.11m (the seventh bin). We observe that the bin identified by SRTL matches the location indicated by the CDF and predicted by our range measurements.

From these experiments we conclude that we can completely explain SRTL performance under known conditions; SRTL places known reflections in their correct bin as predicted by the physical setup.

4.3 Robustness to Noise

Although we verified that SRTL works as expected in perfect conditions, we also care about performance in the face of environmental noise.

We investigate this question by adding an artificial noise source to our experiments. This noise source is a second modem that transmits tones that trigger non-echo tone detections at our original sender (on which SRTL algorithm is running). Our artificial noise source transmits tones repeatedly, with inter-transmission times chosen uniformly randomly within a fixed interval. We then vary this interval to adjust the degree of noise, with smaller intervals causing greater (more frequent) noise. Since timing of noise is random, we expect SRTL to ignore such noise and still be able learn detections from known surfaces. (We select this noise model to provide simple, controlled tests. Exploration of richer noise sources is an area of future work.)

We performed our experiments with different levels of noise in both controlled and uncontrolled environments (anechoic chamber and laboratory). Figure 7 shows the result for

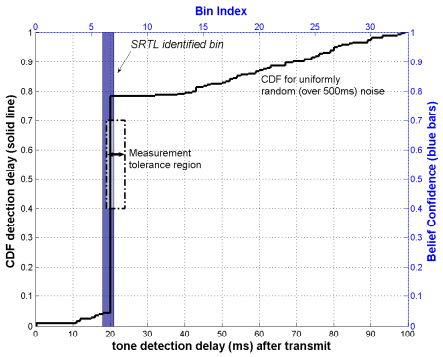

(a) Uniform noise source over 500ms

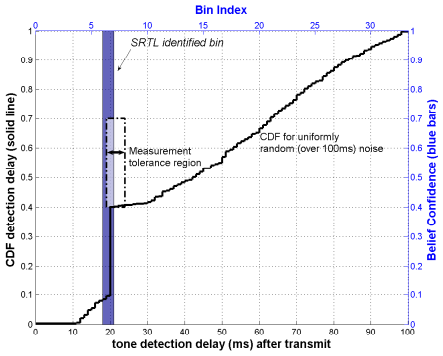

(b) Uniform noise source over 100ms

Figure 7: Results of a controlled experiment (object at 3.87m) showing the CDF of 200 samples and SRTL response, with varying levels of noise.

the controlled environment of the anechoic chamber.

The presence of a gradual slope in the sample CDF (the solid black line) indicates the presence of noise. The slope of the CDF indicates the level of noise; a steeper slope indicate greater noise, as can be seen comparing Figures 7(a) and 7(b). However, we observe that SRTL identifies the correct bin of the reflector even with substantial interference. We looked at several noise levels, Figure 7(b) shows the case where there are, on average, two noise triggers in each sample period for the single true echo. However, because noise is randomly distributed, SRTL can suppress it and learn the true echos. We conclude from this experiment that while SRTL will learn its environment, it will not be fooled by competing contenders (in the case of T-Lohi) or some levels of environmental noise.

We anticipate that real-world noise and multipath will be more complex, so we next reproduce this experiment in the uncontrolled environment of our lab where bin identification is more challenging for SRTL. Figure 8 shows the result of our experiments. Over many experiments with and without noise (not shown) SRTL detects two likely echos in bins 4 and 8. Although we do not know the ground truth surface, these indicate empirically the presence of two reflecting surfaces.

Again, we see that the algorithm consistently identifies

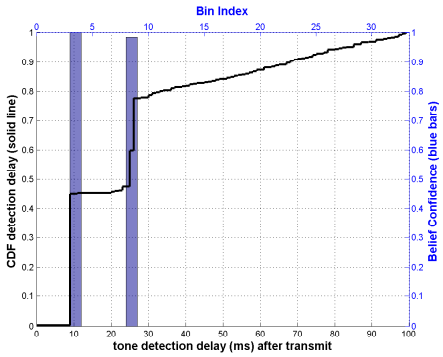

(a) Uniform noise source over 500ms

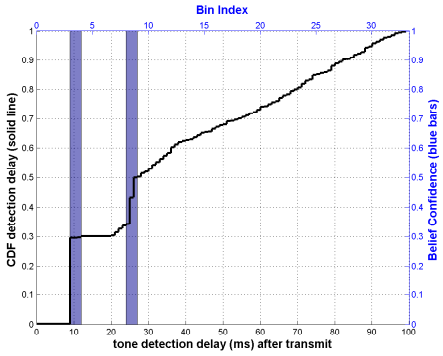

(b) Uniform noise source over 100ms

Figure 8: Laboratory experiment empirically find two stationary reflections with two levels of artificial noise.

these as self-reflection bins and is able to see through and suppress random noise.

Our underwater experiments provide a final evaluation of noise. Our preliminary tests in the Marina show tone triggers every few milliseconds. Assuming a single reflecting surface, this corresponds to 50 false triggers per sensing period, far more noise than regular reflections. SRTL is unable to track known surfaces with this level of noise. We are currently working to better characterize SRTL tolerance, and to improve our wakeup receiver's ability to filter brief noise in hardware.

From the above results conclude that SRTL can tolerate random noise up to at least two false triggers per true echo, and shows the need to further characterize the limits of noise tolerance. We do not characterize further due to hardware limitation, but explore greater noise tolerance using simulations in Section 5.2.

4.4 SRTL in a Changing Environment

Most underwater environment change, either due to tides or currents, precession on an anchor, or movement of human artifacts or fish. We therefore next wish to evaluate how well SRTL adapts to a changing environment. Properly configured, Bayesian learning can track changes in belief so we expect SRTL to track changes in the environment successfully.

To investigate SRTL response to environmental changes, we return to the anechoic chamber. We place a reflecting surface at a known location (at 3.32m, corresponding to the fifth bin). We then take twenty consecutive samples at that location to train the SRTL algorithm to identify that location with maximum confidence. We then move the reflecting surface to a different location (4.11m, corresponding to bin 7) relative to the transmitter. We then observe how SRTL's belief distribution (about which bins are SR) evolves as it collect additional observations. We expect SRTL to track the changed environment and identify the new bin after incorporating a few samples. The decision threshold is set using the mechanism described in Section 3.5.4.

Figure 9(a) shows how SRTL's belief changes as it takes more observations. Initially SRTL is certain that bin five is the surface, but after 5 samples (the second figure from the left), this believe begins to waver. At 10 samples it has begun recognizing bin 7, the new location, as a self-reflection, although it still remembers the old location. Finally, after about 15 samples, SRTL has nearly completely shifted its understanding of the environment. This experiment demonstrates that SRTL will adapt to changes in environment. In the next section we consider how SRTL parameter settings affect this convergence time.

5 SRTL Parameter Sensitivity

The previous section describes experiments that show SRTL works well, even with noise and environmental change. SRTL has several parameters that affect its operation, including estimates of *ND* and *TC* events and choice of bin size. We next evaluate how sensitive SRTL is to choice of these parameters.

5.1 Estimates of Observation Errors

SRTL's learning algorithms takes two parameters, p_{nd} and p_{tc}, that are used to adapt its belief to new observations. In Section 3.4 we describe how these parameters model our estimates of observation errors. To understand how the parameters affect SRTL, we continue our experiment with changing locations of reflecting surfaces first for a controlled environment. We then reexamine parameter setting for the more complex and uncontrolled laboratory environment.

5.1.1 Parameter Choices: Controlled Dynamic Environments

We next repeat the experiments from Section 4.4 where we move the reflective surface. Figure 9 shows the result for three different estimates of error probability.

We make four observations from these results. First, in every case SRTL is able to adapt to the changed environment and identifies the new bin location (after 15 samples). However, parameter choice affects the rate of adaptation and the transient behavior.

For moderate and equal parameters (Figure 9(a), $p_{nd} = 0.4$, $p_{tc} = 0.4$), beliefs change relatively slowly. However, when tone cancellation is lower than non-echo detection (Figure 9(b), $p_{nd} = 0.4$, $p_{tc} = 0.1$), confidence in the old bin decays quickly. So after 5 samples, neither bin is identified as self-reflective. We explain this result because a low p_{tc} means that cancellation is unlikely, so negative evidence (absence of tone triggers in the old bin) is quickly accepted.

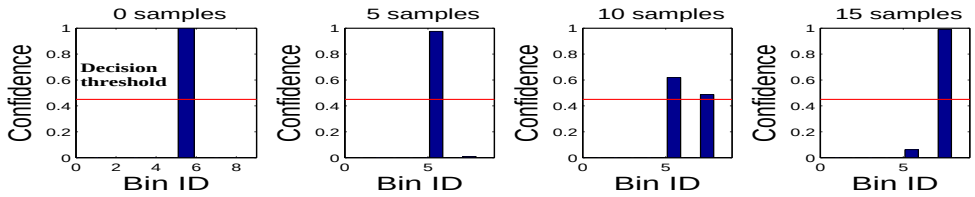

(a) p_{nd}=0.4, p_{tc}=0.4

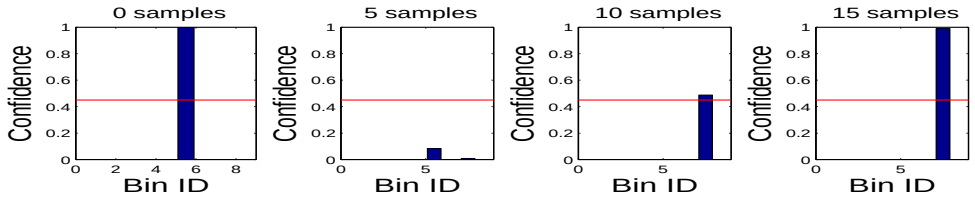

(b) p_{nd}=0.4, p_{tc}=0.1

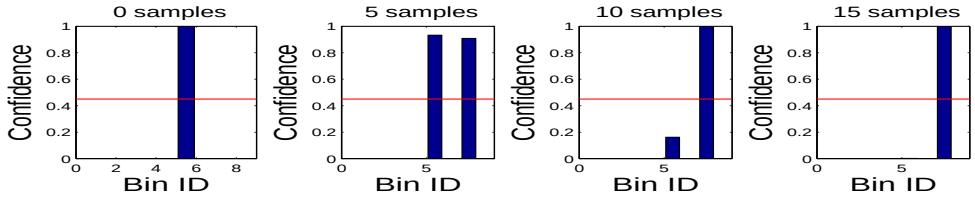

(c) p_{nd}=0.1, p_{tc}=0.4

Figure 9: SRTL response to change in location of reflecting surface from 3.36m (bin 5) to 4.11m (bin 7) with different input parameters.

In Figure 9(c) we see the opposite case, with low non-echo detections but likely cancellation ($p_{nd} = 0.1$, $p_{tc} = 0.4$). Now new evidence is rapidly accepted (bin 7 is detected after five samples), but old assumptions decay slowly. Here, a low value of p_{nd} indicates a low-noise environment, so new triggers are quickly taken as echos.

Overall, we conclude that high p_{tc} values act to damp the response of our algorithm to missed reflections, while large p_{nd} dampens response to new bins. We believe that keeping a moderate value of p_{tc} is essential, because our lab experiments that show that sometimes there are several consecutive absent echos in even regular locations. Also, a moderate value of p_{nd} helps suppress short, transient noise. We therefore suggest moderate and equal values of both parameters (like, $p_{nd} = 0.4$, $p_{tc} = 0.4$). We expect to review these settings as we gather additional data from underwater experiments.

Finally, detection theory allows one to trade certainty of detection for time; we observe this tradeoff in the value of parameters. Thus, a higher value of p_{tc} confirms removal of a reflecting surface slowly, but also increases accuracy as we now do not react to transient changes.

5.1.2 *Parameter Choices: Uncontrolled Environments*

We have shown that balanced values are appropriate in a controlled setting. We next evaluate their impact in the uncontrolled laboratory setting. To do so, we return to our laboratory experiments with relatively complex multipath. We performed experiments for a wide range of parameters sets, including the ones shown in Figure 9. Figure 10, however shows the results for two additional sets of parameters as they depict an interesting case where the choice of parameters becomes important.

Figure 10(a) shows response with moderate parameters ($p_{nd} = 0.4$, $p_{tc} = 0.5$). Here the algorithm identifies just a single bin (bin 4) as self-reflective. However, the same observations but with ($p_{nd} = 0.7$, $p_{tc} = 0.7$) causes SRTL to also identify an additional bin (bin 11, Figure 10(b)). Which parameter set is best?

The CDF of the observations suggest that both potential reflections are consistent reflection and above the noise threshold. Thus it appears that the second parameter set with high values is a better choice. However, a more detailed look at the data shows that detections in bin 11 were consistent at the beginning of the experiment, but after 160 of the 200 samples, detections in that bin became infrequent (just five more times in the next 40 sample periods). These different observations suggest that multipath changes even in an apparently static room over a few minutes. (This change does not show up in the CDF because it sums the entire experiment.) We conclude therefore that the moderate parameters are a better choice since they correctly dismiss the second bin by the end of the experiment to reflect the change in environment.

5.2 Parameter Alignment with Environmental Conditions

Hardware limitations prevent us from introducing large amounts of non-echo or tone cancellation noise in these ex-

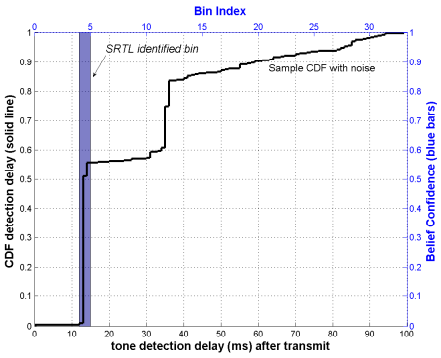

(a) p_{nd}=0.4, p_{tc}=0.5

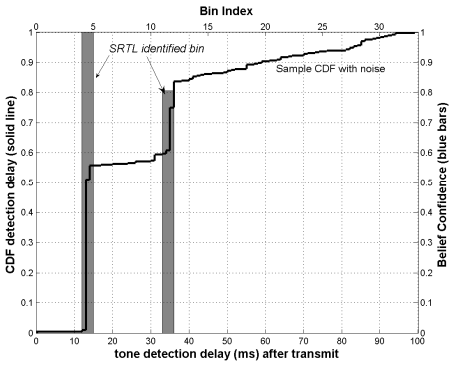

(b) p_{nd}=0.7, p_{tc}=0.7

Figure 10: Impact of parameter choice for an uncontrolled environment.

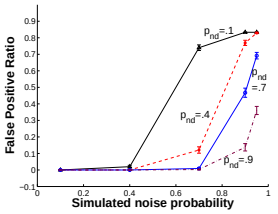

(a) False Positives

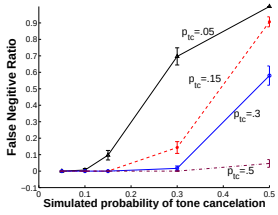

(b) False Negatives

Figure 11: Fraction of false positive and false negative as *simulated* noise and chosen noise estimates vary. Error bars show 95% confidence intervals. (The y-axis starts below zero to show values along the origin.)

periments. To study both types of noise, we simulate the algorithm with artificial noise, allowing us to explore arbitrarily high noise levels under controlled conditions. Our goal is to understand what levels of noise cause SRTL to fail, and how SRTL behaves when our noise estimates (p_{nd} and p_{tc}) differ from the actual amount of noise. Because of space limits, we summarize simulation results here, while details can be found in a technical report [29].

To observe false positives, we vary simulated (wake-up tone triggering) noise in the environment. In Figure 11(a) we fix both the simulated and algorithm parameters of tone

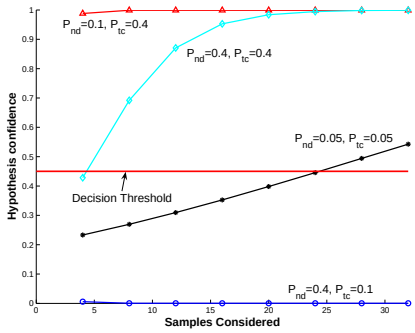

Figure 12: SRTL belief for a single bin being SR for a worst case scenario of reflection time alternating between adjacent bins.

cancellation probability at 0.05, we then vary noise (the x-axis) and observe the fraction of false identifications for different values of p_{nd}. Figure 11(a) shows that the fraction of incorrectly identified bins increases as the simulated environmental noise increases. At greater than 80% noise, nearly all bins are incorrectly identified as self-reflections; SRTL cannot tolerate this level of noise. However, this exact threshold is a function of our noise estimate p_{nd}, because larger estimates make SRTL more skeptical that triggers indicate self-reflections. In general, SRTL performs reasonably as long as the estimate is at least as large as true noise, with some leeway when noise is low. Thus, with $p_{nd} = 0.1$, SRTL handles up to 40% noise (the leftmost line), with $p_{nd} = 0.4$, it works reasonably to 70% noise (the second line from the left). In fact, with $p_{nd} = 0.9$, there are only 10% false positives at 90% noise.

To observe false negatives we vary the simulated probability of a tone cancellation. In Figure 11(b) we fix both the simulated and algorithm input values of noise detection probability at 0.4. Figure 11(b) shows that the fraction of SR bins that are SRTL fail to identify (false negative) increases with the simulated tone cancellation probability. Just as with false positives, performance is best when the estimate (p_{tc}) is close to actual tone cancellation probability, but at low levels SRTL will tolerate noise two or three times that estimate. The main difference is that SRTL is less forgiving of tone cancellation than noise. Because there are few echos, moderate cancellation makes them difficult to identify. Fortunately, as it is rarer for interference to completely suppress channel energy, so cancellation is rarer than non-echo noise.

Thus we conclude that SRTL can tolerate a wide range of noise in the environment, provided that corresponding estimates (both p_{tc} and p_{nd}) are reasonable approximations. Moreover, over-estimating the actual observation errors has a lesser penalty than under-estimating these parameters.

5.3 Impact of Bin Discretization

SRTL uses discrete bins to track belief and provide efficient and low-complexity operation even on mote-class devices. Our final question is to explore how bin size affects our algorithm. We have three concerns: bins that are too

small may disperse observations, bins too large may cause echos to hide real tones from other transmitters, and even with correctly sized bins, tones might fall on the border between two bins. Our observations about system sensitivity (Section 3.5.5) limit our bins to at least 3ms, and we have not seen drift by more than a few ms, so we believe our bins are neither too large nor too small. We next look at echos that lie on the border to investigate the worst case impact of using such a low-complexity, low-memory implementation.

For a reflecting surface exactly on the edge of two bins, minor observation jitter (hardware delays, software locks, clock granularity, or simply very slight motion of either the node or reflecting object) can easily move observations in either direction. To simplify evaluation, we consider the worst case scenario: each sample produces detection at an alternate of two adjacent bins. We emulate this scenario by artificially providing such a sample input to our SRTL algorithm.

Since the bins are symmetric, belief in each bin is identical. Figure 12 plots the belief in one of the adjacent bins as more samples are incorporated by the algorithm using SRTL for four different parameter sets. We see that combining a low value of tone cancellation probability with a higher value of noise detection prevents both bins from *ever* being declared as SR bins ($p_{nd} = 0.4, p_{tc} = 0.1$). This result is because low value of p_{tc} implies any confidence is lost quickly, while a higher value of p_{nd} requires long time for new identification to be made.

Flipping these parameters results in an almost immediate identification of both bins as SR bins ($p_{nd} = 0.1, p_{tc} = 0.4$). This choice is very aggressive and thus inappropriate as a small value of p_{nd} makes for a higher possibility of false positives as discussed in Section 5.1. A more appropriate choice of parameters is possible with both values at a moderate and equal value such as 0.4 (as suggested in Section 5.1). In this case the algorithm requires just six samples (three positives for each bin) before successfully identifying both as being SR bins.

The SRTL input parameters provide us with a fine-tuning-knob to adapt our algorithm to the need of our environment. The results in this section and Section 5.1 suggest setting these parameters is more critical to the performance than the decisions threshold, since belief increases exponentially with positive reinforcements (so a threshold of 0.45 and 0.8 might be reached in just one additional sample). Our results also suggest that using moderate and equal values (for example, 0.4 and 0.4) for both noise-detection and tone-cancellation events provides good performance. This combination damps both the increase in confidence for a new bin as well as the decrease in confidence for bin for which reflections come and go.

6 Related Work

Our work on tone self-multipath is related to three areas: data multipath, echo-detection techniques, and Bayesian learning. We next describe background for each related area and highlight the novelty of our approach.

Combating the large multipath spread to achieve robust data communication is considered the most challenging task of an underwater acoustic (UWA) communication system [25, 3]. Coherent systems, bandwidth efficient for the band-limited UWA channel, are much more sensitive (than non-coherent systems) to this large multipath spread which can result in inter-symbol or inter-chip interference. Stojanovic et al. were the first to propose a suboptimal, and therefore less complex, Decision Feedback Equalizer (DFE) jointly optimized with a Phased Locked Loop (PLL) that enabled coherent underwater communication [27]. Recently Time-Reversal-Mirror (TRM) has also been considered as a mechanism to handle multipath in an underwater environment [23]. Such existing physical-layer techniques, including Rake receiver [20] or TRM [23], distinguish between several copies of the same signal at a *receiver*. However, we are faced with the challenge of *self-multipath*, where a transmitter must identify echos of its own signal, while using low-power acoustic tones, not the more energy consuming data. We therefore require different approaches at the transmitter.

Perhaps the closest to our work are acoustic signal processing techniques based on interpreting echos from the environment, including echolocation, echo-sounding, and sonar [13, 8]. While sonar is used to detect the presence of obstacles for navigation [13], echo-sounder are widely deployed for bathymetric data collection [2]. In both active sonar and echo-sounding, sound pulses, often called "pings", are sent for echos detection that measures the distance to a reflection surface. In these techniques signal correlation or matched filters are used to disambiguate echos. Furthermore, sonar identifies targets by comparing echos with pre-recorded signatures. However, such techniques are not designed for energy-efficient applications. Underwater sensornets use low-power, wake-up circuit that focus on energy and cost-efficiency. Thus the wake-up circuit is a simple analog part that precludes the use of any complex signal-processing techniques. Our work tackles this unique problem of echo identification for low-cost, low-power acoustic tone signals by using the simple tone activation signal provided by such devices to probe the medium and learn to accurately identify echos.

Bayesian learning as a field builds upon the rich mathematical history of Bayesian statistics. Bayesian learning uses empirical data (collected by sensors) to progressively improve estimates of system parameters [31, 10]. Researcher have used Bayesian inference to learn complex CGI (computer graphic imagery) models from human actors [10] as well as to localize robots using landmarks [31]. The problem of self-multipath is unique for low-power acoustic tone-based communications. We are the first to use Bayesian learning mechanism to overcome the impact of such tone echos within the constraints of low-cost, energy-efficient hardware necessitated in underwater sensornets.

7 Future Work and Conclusions

Although this paper reports on extensive in-air experiments, our most pressing area of future work is completing our underwater experiments. We expect hardware revisions in our version 3 modem will enable underwater experiments, and we are very interested in confirming our choices underwater.

There are several other possible future directions. Al-

though we focus on Bayesian approaches, other possibly simpler approaches like exponential weighted moving averages, may also work. We also plan to integrate our approach into T-Lohi. Finally we can use SRTL to suppress receiver multipath of contention tones and explore using our algorithm in conjunction with other application layer protocols for underwater sensornets.

In this paper we first identified the issue of tone self-multipath unique to the large propagation delays of acoustic networks. We then used T-Lohi as an example to show that this delay can significantly reduce the throughput of our MAC protocol. We then introduced a Bayesian learning algorithm, Self-Reflecting Tone Learning (SRTL), that can be used to learn-and-ignore these self-multipath or echo tones. We performed experiments to verify that our algorithm is correct, robust to noise, and can adapt to a dynamic environment.

8 References

[1] Benthos, Inc. *Fast And Reliable Access To Undersea Data*. http://www.benthos.com/pdf/Modems/ModemBrochure.pdf .

[2] B.S. Bourgeois, A.B. Martinez, P.J. Alleman, J.J. Cheramie, and J.M. Gravley. Autonomous bathymetry survey system. *IEEE Journal of Oceanic Engineering*, 24(4):414–423, Oct 1999.

[3] J.A. Catipovic. Performance limitations in underwater acoustic telemetry. *IEEE Journal of Oceanic Engineering*, 15(3):205–216, Jul 1990.

[4] A. El-Hoiydi and J.-D. Decotignie. Wisemac: an ultra low power mac protocol for the downlink of infrastructure wireless sensor networks. In *Proceedings of ISCC"04*, pages 244–251, Washington, DC, USA, 2004. IEEE Computer Society.

[5] S. Floyd and V. Jacobson. The synchronization of periodic routing messages. *ACM/IEEE Transactions on Networking*, 2(2):122–136, April 1994.

[6] L. Freitag, M. Grund, S. Singh, J. Partan, P. Koski, and K. Ball. The WHOI micro-modem: an acoustic communications and navigation system for multiple platforms. *OCEANS, 2005. Proceedings of MTS/IEEE*, pages 1086–1092 Vol. 2, 2005.

[7] Lewis Girod and Deborah Estrin. Robust range estimation using acoustic and multimodal sensing. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2001)*, Maui, Hawaii, October 2001.

[8] E. Hammerstad, S. Asheim, K. Nilsen, and H. Bodholt. Advances in multibeam echo sounder technology. *OCEANS '93. 'Engineering in Harmony with Ocean'. Proceedings*, pages I482–I487 vol.1, Oct 1993.

[9] John Heidemann, Yuan Li, Affan Syed, Jack Wills, and Wei Ye. Underwater sensor networking: Research challenges and potential applications. In *Proceedings of IEEE WCNC'06*, pages 228–235, Las Vegas, Nevad, USA, April 2006.

[10] Aaron Hertzmann. Introduction to bayesian learning. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Course Notes*, page 22, New York, NY, USA, 2004. ACM.

[11] HILTI, Inc., http://www.us.hilti.com . *PD-40 Laser Range Meter*.

[12] E. T. Jaynes. *Probability Theory : The Logic of Science*. Cambridge University Press, April 2003.

[13] W.C. Knight, R.G. Pridham, and S.M. Kay. Digital signal processing for sonar. *Proceedings of the IEEE*, 69(11):1451–1506, Nov. 1981.

[14] W. A. Kuperman, W. S. Hodgkiss, H. C. Song, T. Akal, C. Ferla, and D. R. Jackson. Phase conjugation in the ocean: Experimental demonstration of an acoustic time-reversal mirror. *Acoustical Society of America Journal*, 103:25–40, January 1998.

[15] Weichang Li and J.C. Preisig. Estimation of rapidly time-varying sparse channels. *IEEE Journal of Oceanic Engineering*, 32(4):927–939, Oct. 2007.

[16] Yuan Li, Wei Ye, and John Heidemann. Energy efficient network reconfiguration for mostly-off sensor networks. In *Proceedings of the IEEE SECON*, pages 527–535, Reston, Virginia, USA, September 2006. IEEE.

[17] Jim Partan, Jim Kurose, and Brian Neil Levine. A survey of practical issues in underwater networks. In *Proceedings of WUWNet '06*, pages 17–24, New York, NY, USA, 2006. ACM Press.

[18] Joseph Polastre, Jason Hill, and David Culler. Versatile low power media access for wireless sensor networks. In *Proceedings of the 2nd ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 95–107, Baltimore, MD, USA, November 2004.

[19] James Preisig. Acoustic propagation considerations for underwater acoustic communications network development. In *Proceedings of WUWNet'06*, pages 1–5, New York, NY, USA, 2006. ACM Press.

[20] R. Price and P.E. Green. A communication technique for multipath channels. *Proceedings of the Institute of Radio Engineers*, 46(3):555–570, March 1958.

[21] Nithya Ramanathan, Mark Yarvis, Jasmeet Chhabra, Nandakishore Kushalnagar, Lakshman Krishnamurthy, and Deborah Estrin. A stream-oriented power management protocol for low duty cycle sensor network applications. In *Proceedings of the IEEE Workshop on Embedded Networked Sensors*, pages 53–62, Sydney, Australia, May 2005. IEEE.

[22] Sumit Roy, Payman Arabshahi, Dan Rouseff, and Warren Fox. Wide area ocean networks: architecture and system design considerations. In *Proceedings of WUWNet '06*, pages 25–32, New York, NY, USA, 2006. ACM Press.

[23] Hee-Chun Song, William S. Hodgkiss, and William A. Kuperman. MIMO time reversal communications. In *Proceedings of WuWNet '07*, pages 5–10, New York, NY, USA, 2007. ACM.

[24] Ethem Mutlu Sözer and Milica Stojanovic. Reconfigurable acoustic modem for underwater sensor networks. In *Proceedings of WUWNet '06*, pages 101–104, New York, NY, USA, 2006. ACM.

[25] M. Stojanovic. Recent advances in high-speed underwater acoustic communications. *IEEE Journal of Oceanic Engineering*, 21(2):125–136, Apr 1996.

[26] Milica Stojanovic. *Wiley Encyclopedia of Telecommunications*, chapter Acoustic (underwater) Communications. John Wiley & Sons, Inc., 2003.

[27] Milica Stojanovic, Josko Catipovic, and John Proakis. Phase coherent digital communications for underwater acoustic channels. *IEEE Journal of Oceanic Engineering*, 19(1):100–111, January 1994.

[28] Affan Syed, Wei Ye, Bhaskar Krishnamachari, and John Heidemann. Understanding spatio-temporal uncertainty in medium access with aloha protocols. In *Proceedings of the Second ACM International Workshop on UnderWater Networks (WUWNet)*, Montreal, Quebec, Canada, September 2007. ACM.

[29] Affan A. Syed. *Understanding and Exploiting the Acoustic Propagation Delay in Underwater Sensor Networks*. PhD thesis, University Of Southern California, Los Angeles, CA, April 2009.

[30] Affan A. Syed, Wei Ye, and John Heidemann. T-Lohi: A new class of MAC protocols for underwater acoustic sensor networks. In *Proceedings of the IEEE Infocom*, pages 231–235, Pheonix, AZ, April 2008.

[31] S. Thrun. Bayesian landmark learning for mobile robot localization. *Machine Learning*, 33(1):41–76, 1998.

[32] I. Vasilescu, K. Kotay, D. Rus, M. Dunbabin, and P. Corke. Data collection, storage, and retrieval with an underwater sensor network.

In *Proceedings of ACM SenSys'05*, pages 154–165, New York, NY, USA, 2005. ACM Press.

[33] Jack Wills, Wei Ye, and John Heidemann. Low-power acoustic modem for dense underwater sensor networks. In *Proceedings of WUWNet05*, pages 79–85, Los Angeles, California, USA, September 2006. ACM.

[34] Wei Ye, John Heidemann, and Deborah Estrin. Medium access control with coordinated adaptive sleeping for wireless sensor networks. *ACM/IEEE Transactions on Networking*, 12(3):493–506, June 2004.