Battery-Free Subsea Internet-of-Things

by

Sayed Saad Afzal

B.Sc., Lahore University of Management Sciences (2018)

Submitted to the Department of Electrical Engineering and Computer
Science in partial fulfillment of the requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2022

Author . . . .
Department of Electrical Engineering and Computer Science
May 13, 2022

Certified by . . . .
Fadel Adib
Associate Professor of Media Arts and Science
and Department of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by . . . .
Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

Battery-Free Subsea Internet-of-Things

by

Sayed Saad Afzal

Submitted to the Department of Electrical Engineering and Computer Science
on May 13, 2022, in partial fulfillment of the
requirements for the degree of
Master of Science

Abstract

Bringing massive connectivity to low-cost, low-power ocean sensors is important for numerous oceanographic applications (across climate/weather modeling, marine biology, aquaculture, and defense). However, standard IoT technologies (e.g, Bluetooth, WiFi, GPS) cannot operate underwater, which has left 70% of our planet (the ocean) beyond their reach. In this thesis, I describe how we can change this reality by introducing IoT technologies that are inherently designed for the ocean. Specifically, I show how by rethinking the entire IoT technology stack in the context of oceans, we introduced low-cost ($<$ \$100), net-zero-power, scalable connectivity technologies that seamlessly operate underwater and pave the way for massive underwater sensing, networking, localization, and machine learning.

The thesis makes four fundamental contributions: First, it introduces ultra-wideband underwater bacskcatter, a technology that enables scalable, battery-free underwater communication. Second, it demonstrates how we can push the network throughput of underwater backscatter through a family of techniques including higher-order modulation techniques, self-interference cancellation, and multi-access protocols. Third, it shows how we can leverage our underwater backscatter nodes to enable a battery-free underwater GPS for localization and navigation. Finally, it demonstrates the feasibility of battery-free inference and machine learning in underwater environments by developing a task-specific deep neural network (DNN) model and deploying it on our battery-free underwater nodes.

I deliver these contributions by designing and building new algorithms, systems and protocols for ultra low-power and scalable underwater sensing, networking, localization and inference. I also implement and evaluate these systems in real underwater environments (including rivers and lakes) and challenging weather conditions (including snow and rain), and discuss how they pave the way for new applications in ocean climate monitoring, underwater navigation, ocean exploration, robotics, aquaculture, and marine discovery.

Thesis Supervisor: Fadel Adib
Title: Associate Professor of Media Arts and Science
and Department of Electrical Engineering and Computer Science

To Allah, who gave me the courage to never give up and helped me in miraculous ways. To my sister and parents, whose encouragement and support of me goes beyond what words can adequately express. To my little angel, Mirha, who lit up my world like never before and whose smile made me trudge on even on the down days.

Acknowledgments

I would first like to thank my thesis advisor, Dr. Fadel Adib, for his unwavering support, his words of encouragement, and most importantly for him believing in me when I could not do so myself. It is people like him that continue to inspire their students each and everyday.

I would also like to express my gratitude towards Dr. Hamed Haddadi for his insightful comments and suggestions. Without his passionate participation and input, part of this thesis wouldn't be possible. Additionally, I would like to extend my sincere thanks to to Prof. Micheal Triantafyllou and Dr. Andrew Bennet for letting me use the MIT sea grant facility, and for helping me during my experiments with creative ideas.

Much credit for the research in this thesis goes to my phenomenal student collaborators. Especially Waleed, Osvy and Reza, who have worked with me round the clock in building and debugging these systems over the past 2 years. I am very fortunate to have learnt from my other coauthors, who were fantastic collaborators, mentors, and mentees: Mario, Unsoo, Yuchen, Vincent and David.

I am very lucky to have an amazing group of fun and supportive colleagues here in Signal Kinetics: Waleed, Nazish, Osvy, Tara, Reza, Laura, Isaac, Unsoo, Aline, Jack, Maisy and Purui. Spending time with them reminds me that family isn't always blood but it's the people in your life who want you in theirs. I owe many thanks to our group admins, Holly and Aimee, for taking care of food, project orders, reimbursements or any last minute emergencies. I really appreciate all that you do for us to make our lives easier.

Finally, I must express my very profound gratitude to my parents, Amira and Jawad, and my sister, Kanza, for standing with me till the very end, giving me every form of support, guidance and love. I can not even begin to thank you for all the sacrifices that you made for me, without which I wouldn't be where I am today; this thesis is dedicated to you.

The research in this thesis is supported by the Office of Naval Research, the National Science Foundation, the Sloan Research Fellowship, the MIT Media Lab, and the MIT Sea Grant.

Previously Published Material

Chapter 2 revises a previous publication [25]: Reza Ghaffarivardavagh*, Sayed Saad Afzal*, Osvy Rodriguez*, and Fadel Adib. Ultra-Wideband Underwater Backscatter via Piezoelectric Metamaterials. ACM SIGCOMM 2020.

Chapter 3 revises a previous publication [7]: Sayed Saad Afzal, Reza Ghaffarivardavagh, Waleed Akbar, Osvy Rodriguez, and Fadel Adib. Enabling higher-order modulation for underwater backscatter communication. IEEE Global Oceans 2020.

Chapter 4 revises a previous publication [26]: Reza Ghaffarivardavagh, Sayed Saad Afzal, Osvy Rodriguez, and Fadel Adib. Underwater backscatter localization. ACM HotNets 2020.

Chapter 5 revises a previous publication [101]: Yuchen Zhao*, Sayed Saad Afzal*, Waleed Akbar, Osvy Rodriguez, Fan Mo, David Boyle, Fadel Adib, and Hamed Haddadi. Towards battery-free machine learning and inference in underwater environments. ACM HotMobile 2022.

Contents

1 Introduction 16

1.1 Underwater Bacskcatter . . . 18

1.2 Systems Developed . . . 20

2 Ultra-Wideband Underwater Backscatter 24

2.1 The Narrowband Problem . . . 28

2.1.1 Underwater Backscatter . . . 28

2.1.2 The Resonance Bottleneck . . . 29

2.2 Ultra-wideband Meta-material Design . . . 33

2.2.1 Coupled Eigenmodes . . . 33

2.3 Fabrication & Implementation . . . 38

2.3.1 Mechanical Fabrication . . . 38

2.3.2 Hardware Design . . . 40

2.3.3 Evaluation Setup . . . 41

2.4 Results . . . 41

2.4.1 Throughput . . . 42

2.4.2 Range . . . 43

2.5 Conclusion . . . 45

3 Pushing the Network Throughput of Underwater Backscatter 46

3.1 System Design . . . 48

3.1.1 Enabling QAM Backscatter . . . 48

3.1.2 Self-Interference Cancellation . . . 51

3.1.3 Scaling to Many Nodes 54
3.1.4 From PHY to Higher Layers of the Networking Stack 55
3.2 Experimental Results 56
3.2.1 Simulation Results 56
3.2.2 Empirical Results 57
3.2.3 Self-Interference Cancellation Performance 59
3.2.4 Scaling Performance 61
3.3 Conclusion 62

4 Underwater Backscatter Localization 63
4.1 The (New) Problem 66
4.2 UBL 68
4.2.1 Backscatter ToA Estimation 69
4.2.2 Adaptive Backscatter Localization 72
4.3 Feasibility Study 76
4.4 Conclusion 77

5 Underwater Battery-Free Machine Learning 78
5.1 Underwater Battery-free Architecture 80
5.1.1 Underwater Energy Harvesting 80
5.1.2 Battery-free Underwater Inference 81
5.1.3 From Inference to Communication 83
5.2 Feasibility Study 83
5.2.1 Prototype Hardware 83
5.2.2 Offline Pipeline 84
5.2.3 Model Accuracy 85
5.2.4 Power Performance 87
5.3 Conclusion 88

6 Conclusion 89
6.1 Looking Forward 89

6.1.1 Underwater Backscatter Localization 90
6.1.2 Underwater Battery-Free Machine Learning 91

List of Figures

1-1 **Underwater Battery-less Sensor.** (a) shows one of our battery-less sensors harvesting energy from sound to power up and communicate. (b) shows two students dipping multiple batteryless nodes in the Charles River at the MIT Sailing Pavilion in Cambridge, MA. 17

1-2 **Underwater Backscatter.** The figure depicts the operation principle of underwater backscatter communication. 18

1-3 **Backscattered BPSK Signal.** The figure plots the amplitude of the BPSK backscattered signal as a function of time. 19

1-4 **Fabricated U^2B transducer.** (a) shows the different components of a U^2B transducer prior to potting. (b) shows one of our potted U^2B transducers which was tested in the Charles River. 20

1-5 **Design of battery-free inference on underwater edge nodes:** The system harvests energy from sound and powers up an on-board sensor. The sensor records animal sounds and feeds them to the processor for inference 22

2-1 U^2B **Synthesizes Resonances to Achieve UWB Performance.** In each of (a) and (b), the left figure shows the transducer architecture, the middle figure shows the SNR as a function of frequency, and the right figure shows one of the vibration modes – or eigenmodes – of the active layers. 26

2-2 **Single-Layer Eigenmode.** The figure shows the amplitude of the piezo vibration across its body when it is excited (a) at resonance, and (b) outside resonance. The vibrations are visualized as a heat map where red and blue represent high and low vibration amplitudes respectively. 30

2-3 **Backscatter and Resonance.** The figure plots the absolute value of the difference in the reflection coefficient (between open and short cases) as a function of normalized frequency and as a function of the normalized imaginary component of piezo-transducer impedance. . . . 31

2-4 **Eigenmodes and Eigenfrequencies of U^2B's Metamaterial Design.** The figure in the top left shows the vibration amplitude across the different eigenfrequencies. (I)-(VII) depict the eigenmodes associated with some of the eigenfrequencies. The vibration amplitude for each of these eigenmodes is visualized as a heatmap across the body of the transducer, where red indicates a high amplitude vibration and navy blue indicates low amplitude vibration. 33

2-5 **Characteristic Curve for Single and Multi-layer Designs.** The figure plots the characteristic curve for single-layer (in red) and multi-layer (in black) designs as a function of frequency. 36

2-6 **Exploded Transducer View.** The figure shows the exploded transducer view and the layered piezoceramic structure for U^2B. 38

2-7 **SNR vs Bitrate.** The figure plots the median SNR as a function of bitrate for both U^2B (in red) and PAB (in blue) at two different downlink center frequencies: 20 kHz (dashed line) and 40 kHz (solid line). The error bars represents the 25^{th} and 75^{th} percentile and the dotted gray horizontal line represents the SNR needed to get a BER of 10^{-2}. 42

2-8 **Backscatter SNR vs Distance.** The figure plots the backscatter SNR as a function of distance for different coding rates. Error bars indicate standard deviation, and solid lines indicate smoothed curves. 44

3-1 ***PAB-QAM*** (a) shows one of our potted *PAB* transducers (b) shows the schematic design for *PAB-QAM*. 47

3-2 **Backscatter Constellation.** This figure plots the constellation points of the received backscatter signal for: (a) BPSK backscatter, and (b) *PAB-QAM*. 49

3-3 **Mitigating Interference and Scaling and the Size of *PAB-QAM* networks.** The figure plots the power of the signal received by a hydrophone as a function of frequency. (a) shows how *PAB-QAM* can eliminate the downlink self-interference (in green) and retain a backscatter response (in red). (b) shows how *PAB-QAM* can scale to multiple nodes by assigning a different frequency channel to each of them. 52

3-4 **Packet Structure.** This figure shows the different fields of the downlink and uplink packet. 55

3-5 **Throughput vs SNR (Simulations).** The figure plots the throughput of QAM (blue) and BPSK (red) as a function of SNR with different code rates. For any given code rate, 4-QAM achieves twice the throughput when compared with BPSK. 57

3-6 **Throughput vs SNR (Empirical).** The figure plots the normalized throughput of QAM (blue) and BPSK (red) as a function of SNR. The error bars represents the 25^{th} and 75^{th} percentile respectively. 58

3-7 **Self-Interference Cancellation**. The figure plots the median SNR with (red) and without (blue) self-interference cancellation as a function of transmitter voltage. The error bars represent the 10^{th} and 90^{th} percentile. The dotted line represents the SNR needed to get a BER value of 10^{-2}. 60

3-8 **BER and SNR vs Concurrent Transmissions.** The figure plots the median BER and SNR as a function of the number of concurrent transmissions. The error bars represent the 25^{th} and 75^{th} percentile. . 62

4-1 **Multipath and Underwater Backscatter.** (a) shows how sound propagates underwater, repeatedly reflecting off the surface and seabed. (b) shows a received backscatter packet in deep (low multipath) and shallow (dense multipath) water. 67

4-2 **Range estimation via frequency hopping.** (a) shows how UBL can isolate the direct path in the time domain using the frequency-hopping localization method and (b) shows the effect of the wake-up lag on conventional ToA based localization schemes. 71

4-3 **ToA Estimation in shallow water.** This figure shows how multipath affects the localization ability for UBL. (a) shows that at a high bit-rate of 2 kbits/s, UBL fails to localize the object. (b) shows that operating at a lower bit-rate of 100 bits/s in multipath rich environments yields better performance. 73

4-4 **ToA Estimation in deep water with mobility.** This figure shows how UBL can adapt to mobility in deep water environments . (a) shows that at a bit-rate of 1 kbits, UBL is unable to localize the object while the object is moving with a speed of 0.3 m/s, so for better accuracy, it is desirable to operate at a higher bit-rates to deal with mobility as shown in (b). 75

4-5 **Preliminary Results for UBL.** The system was tested for three different ranges i.e. 24 cm, 34 cm and 44 cm respectively. 77

5-1 Overview of offline training and on-board inference. Initially, a high-level model is trained, validated, and evaluated offline on pre-collected audio clips from different marine mammals. A model with acceptable offline accuracy is converted to a low-level model that can be supported by standard C libraries on a target battery-free device. The converted model is then deployed in an end-to-end (, from sounds received by a hydrophone to final classifications) marine mammal recognition pipeline on the target device. 80

5-2 **Prototype Hardware** (a) shows the processor that our design uses for on-board inference. (b) shows one of our potted transducers which can be used for energy harvesting and communication. 83

5-3 CNN structure for the classification task. The convolutional layer has 8 filters (size 3 × 3, strides 2 × 2) activated by ReLU. The dense layer is activated by softmax. 85

5-4 Distribution (boxplots and kernel density) of offline test accuracy CNN models with different input layer size (N) and number of mammal classes (C). The converted model of (N-16, C-4) can be deployed on our prototype. 86

List of Tables

5.1 On-board power and energy analysis of the battery-free inference prototype with a data rate of 1 kbps and a sampling rate of 330 samples/s. 87

Chapter 1

Introduction

Energy-efficient underwater networking has recently witnessed mounting interest from academia and industry due to emergent needs for its environmental, defense, and industrial applications [63, 41, 37, 44]. Concerns about the impact of climate change on ocean health have prompted oceanographers and climatologists to seek sensor networks that can be used to monitor ocean vital signs such as carbon balance, coral reef conditions, and biodiversity [52, 58, 72]. On the defense front, DARPA launched the “Ocean of Things” program in 2017 to propel the development of low-cost, low-power distributed sensor networks that can enhance maritime situational awareness [17]. Major industry players – including Google, Microsoft, and Honeywell – have also become increasingly interested in deploying such networks to monitor underwater infrastructures ranging from oil and gas pipelines [28] to submerged datacenters [103].

Despite major advances in underwater technologies, existing proposals remain far from the vision of a low-cost, low-power distributed architecture for a subsea IoT. The majority of today’s systems rely on point-to-point communication between modems that require at least 50-100 Watts for data transmission [95, 60], a power level that quickly drains the batteries of underwater sensors and limits their lifetime. Early workarounds for this high power consumption problem involved heavy duty-cycling which strangled the data rates to within few to tens of bits per second [42, 40]. Recent work on underwater wireless charging [31, 62] and backscatter [37] promises to extend battery life and to reduce the power consumption of underwater communication to

(a) Batteryless underwater sensor

(b) Testing in Charles River

Figure 1-1: **Underwater Battery-less Sensor.** (a) shows one of our batteryless sensors harvesting energy from sound to power up and communicate. (b) shows two students dipping multiple batteryless nodes in the Charles River at the MIT Sailing Pavilion in Cambridge, MA.

sub-milliWatt levels; however, these systems are still limited in their throughput, operational range, and scale. For example, the state-of-the-art system for underwater backscatter is limited to 3 kbps and a maximum distance of 10 m [37].

This thesis presents new systems, algorithms, and protocols to enable ultra-low-power and scalable underwater sensing, networking, localization and inference.

The combination of these algorithms and hardware-software systems enables us to sense the underwater environment at scale at extremely low power using low-cost acoustic sensor nodes. Fig. 1-1(b) shows one of our underwater sensor nodes being tested in the Charles River in Cambridge, MA. Our sensors have three key properties that make their designs particularly powerful: First, they operate entirely based on energy harvested from underwater acoustic signals, which allows them to operate at net-zero power. Second, these sensors employ underwater backscatter for communication, a technology that consumes six orders of magnitude lower power than prior underwater acoustic modems. Third, their communication and sensing capabilities are robust, as evidenced through our implementation and extensive evaluation in real underwater environments.

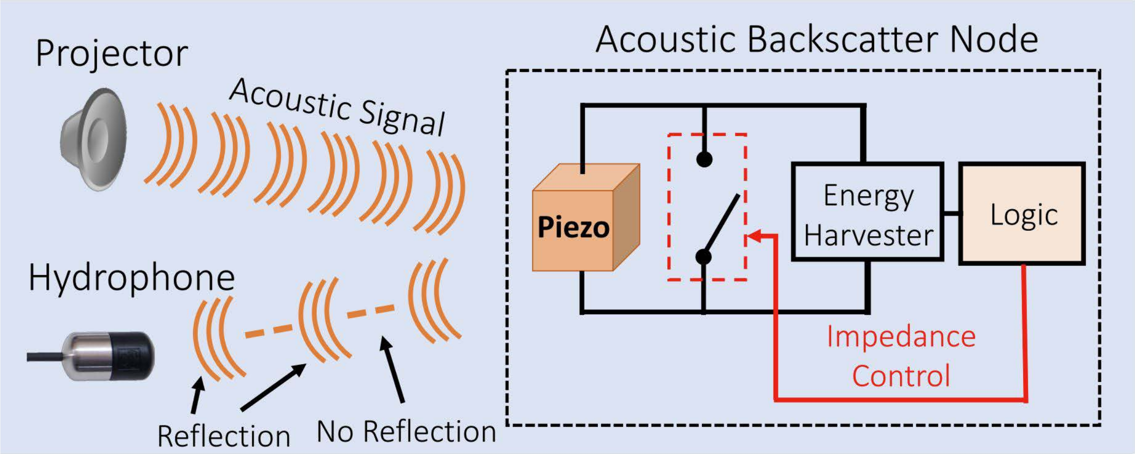

Figure 1-2: **Underwater Backscatter.** The figure depicts the operation principle of underwater backscatter communication.

1.1 Underwater Bacskcatter

The systems developed in this thesis rely on a low-power communication method called underwater backscatter. In this section, we provide a primer on underwater backscatter by briefly describing the principle, and we refer the interested reader to [37] for more details.

Underwater backscatter systems typically consist of a projector, a hydrophone receiver, and a network of batteryless nodes. For simplicity, let us focus on a single backscatter node, as shown in Fig. 1-2. The projector transmits an acoustic signal on the downlink. The backscatter node harvests energy from the downlink signal to power up, and it communicates by modulating the reflections of impinging acoustic signals. A remote hydrophone can sense these reflections and use them to decode the transmitted packets from the backscatter node.

In order to harvest energy and modulate acoustic reflections, underwater backscatter nodes rely on piezoelectric resonators. Piezoelectric materials transform sound into electrical energy, enabling these sensors to harvest energy and power up. Once powered up, they can modulate their acoustic reflections by changing the impedance across the terminals of the piezoelectric transducer. Prior designs change this impedance by simply toggling a switch between the terminals of the piezoelectric resonator. By

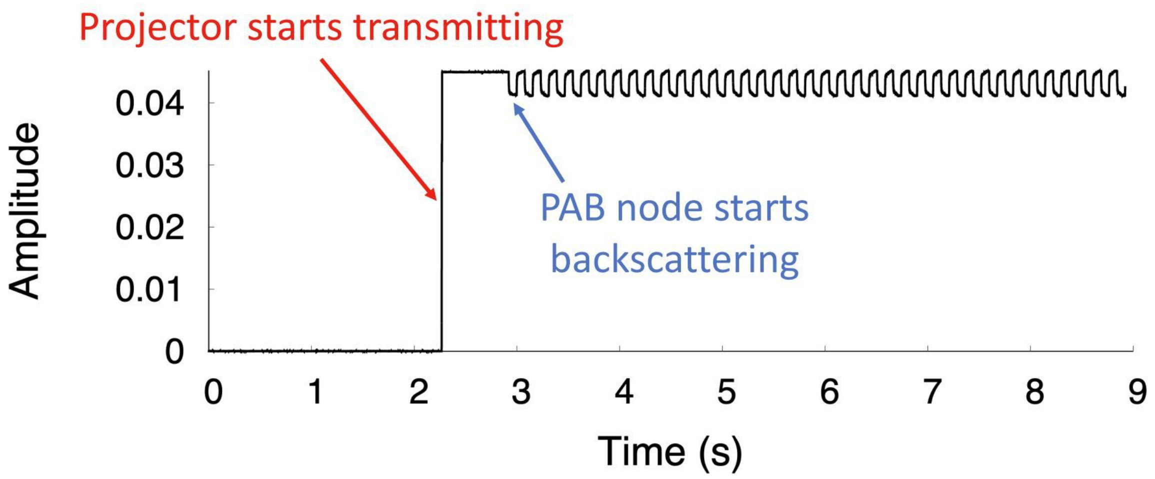

Figure 1-3: **Backscattered BPSK Signal.** The figure plots the amplitude of the BPSK backscattered signal as a function of time.

switching between reflective and non-reflective states, they can communicate bits of zero and one.

Fig. 1-3 shows the received signal by a hydrophone in one of our experimental trials. The figure plots the amplitude as a function of time. At around 2.1s, the amplitude jumps because the projector starts transmitting. Shortly thereafter, we can see that the received signal starts alternating between two amplitudes, corresponding to the two reflective states of the backscatter node. The hydrophone can map these amplitude changes to bits of zero and one, and use them to decode the packets transmitted by a backscatter node. This standard backscatter modulation scheme is inherently limited to only two reflection states.

Prior to this thesis, research on underwater backscatter demonstrated its feasibility as a communication modality, but was significantly limited in its scalability, throughput, and sensing capabilities. For example, state-of-the-art prior systems could achieve throughputs of only a couple of kbps and operate at ranges of few meters. My thesis builds on this foundational past work and is the first to show how underwater backscatter can be scaled toward massive battery-free underwater networking and sensing that achieve orders of magnitudes higher throughput and range, and are capable of new sensing applications that were not feasible before.

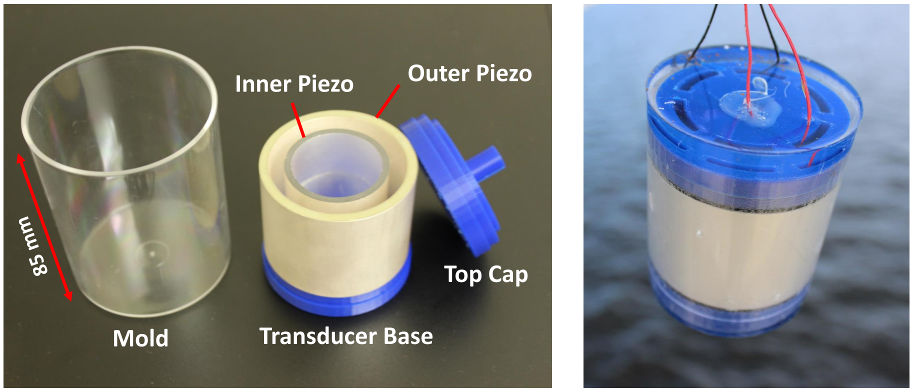

(a) Pre-potting Components (b) Potted Transducer

Figure 1-4: **Fabricated U^2B transducer.** (a) shows the different components of a U^2B transducer prior to potting. (b) shows one of our potted U^2B transducers which was tested in the Charles River.

1.2 Systems Developed

My thesis contributes four main systems for scalable networking and sensing with underwater backscatter.

1. **Ultra-Wideband Underwater Backscatter:** I present U^2B, a technology that enables ultra-wideband backscatter in underwater environments. In contrast to prior approaches for underwater backscatter whose bandwidth (and throughput) was limited by their resonance, U^2B synthesizes a large number of resonances to enable ultra-wideband operation. At the core of its design is a novel metamaterial-inspired transducer for underwater backscatter (see Fig. 1-4 for a sample fabricated node). U^2B's design also introduces new backscatter algorithms that enable long operational range and higher throughput for underwater backscatter. In Chapter 2, I explain how we designed and fabricated U^2B nodes and tested them in a river across different weather conditions, including snow and rain. I also describe our empirical evaluation, which demonstrated that U^2B can achieve throughputs of up to 20 kbps, and an operational range up to 62 m. In comparison to the state-of-the-art prior system for underwater backscatter, this design achieves 5× more throughput and 6× more communication range. Moreover, our evaluation of U^2B represents the first ever experimental validation of underwater backscatter in the wild.

2. **Higher-Order Modulation for Underwater Backscatter:** While U^2B nodes increase the bandwidth of underwater backscatter, their throughput remains limited by their simple binary modulation scheme. To overcome this limitation and enable higher spectral efficiency, I present *PAB-QAM*, the first underwater backscatter design capable of achieving higher-order modulation. *PAB-QAM* exploits the electro-mechanical coupling property of piezoelectric transducers to modulate their reflection coefficients. Specifically, by strategically employing reactive circuit components (inductors), we demonstrate how *PAB-QAM* nodes can modulate the phase and amplitude of acoustic reflections and realize higher-order and spectrally-efficient modulation schemes such as QAM. In Chapter 3, I elaborate on how we designed and built a prototype of *PAB-QAM* and empirically evaluated it underwater. Our empirical evaluation demonstrates that *PAB-QAM* can double the throughput of underwater backscatter without requiring additional power, spectrum, or cost. Additionally, Chapter 3 introduces strategies for concurrent transmissions, full-duplex communication and self-interference cancellation for scalable, reliable and robust underwater networking. Such increased throughput paves way for various subsea IoT applications in ocean exploration, underwater climate monitoring, and marine life sensing.

3. **Underwater Backscatter Localization:** Chapter 4 presents the design and demonstration of the first underwater backscatter localization (*UBL*) system. Our design explores various challenges for bringing localization to underwater backscatter, including extreme multipath, acoustic delay spread, and mobility. We describe how an adaptive and context-aware algorithm may address some of these challenges and adapt to diverse underwater environments (such as deep vs shallow water, and high vs low mobility). We also present a prototype implementation and evaluation of *UBL* in the Charles River in Boston, and highlight open problems and opportunities for underwater backscatter localization in ocean exploration, marine-life sensing, and robotics.

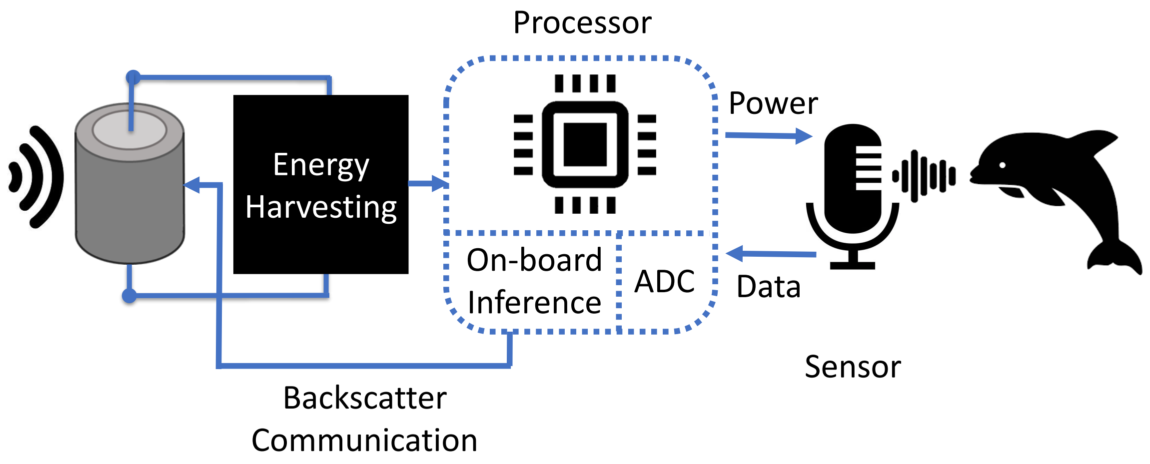

Figure 1-5: **Design of battery-free inference on underwater edge nodes:** The system harvests energy from sound and powers up an on-board sensor. The sensor records animal sounds and feeds them to the processor for inference

4. **Underwater Battery-Free Machine Learning:** Finally, Chapter 5 of this thesis demonstrates how we can enable inference on battery-free edge devices in underwater environments. The key question that this chapter addresses is the following: Can we design and build battery-free devices capable of machine learning and inference in underwater environments? An affirmative answer to this question would have significant implications for a new generation of underwater sensing and monitoring applications for environmental monitoring, scientific exploration, and climate/weather prediction. To answer this question, this chapter explores the feasibility of bridging advances from the past decade in two fields: battery-free networking and low-power machine learning. Our exploration demonstrates that it is indeed possible to enable battery-free inference in underwater environments. Chapter 5 explains how we designed a device that can harvest energy from underwater sound, power up an ultra-low-power microcontroller and on-board sensor, perform local inference on sensed measurements using a lightweight Deep Neural Network, and communicate the inference result via backscatter to a receiver. Fig. 1-5 shows the system architecture for our device, which includes energy harvesting, edge processing and inference, acoustic sensing, and underwater backscatter. We tested our prototype in an emulated marine bioacoustics application, demonstrating the potential to recognize underwater

animal sounds without batteries. Through this exploration, we also uncover challenges and opportunities for making underwater battery-free inference and machine learning ubiquitous.

This remaining chapters discuss the above four systems in chronological order of developments. The techniques presented in these chapter lay the foundation of an ongoing exploration of how we can design and build a battery-free subsea IoT for applications in ocean exploration, marine life sensing, and underwater climate change monitoring.

Chapter 2

Ultra-Wideband Underwater Backscatter

We present Ultra-wideband Underwater Backscatter (U^2B), a system that enables scalable and ultra-low power underwater networking. The system brings the benefits of Ultra-Wideband (UWB) technology to low-power underwater communication. UWB is a mature technology in radio frequency (RF) communication[1] and has been used to extend communication range [24], boost throughput [65], and scale radio networks via concurrent transmissions [35, 36]. However, existing UWB RF systems cannot be used for underwater communication because RF attenuates exponentially in water. This is why today's underwater communication systems primarily rely on acoustic signals, which have good propagation properties in water [33, 90].

Bringing UWB to underwater backscatter faces multiple challenges. First, underwater communication nodes, including backscatter nodes, rely on resonant piezoelectric transducers. These transducers have high efficiency when transmitting and receiving sound at their resonance frequency, but their performance quickly degrades as the frequency moves away from resonance. Hence, they are limited to narrow bandwidths. Second, backscatter communication is inherently full-duplex and suffers from self-interference between the transmitted signal and the backscatter response. This

---

[1] A communication system is defined as UWB if its bandwidth is larger than 20% of its center frequency [92].

problem is exacerbated for UWB backscatter since UWB signals are typically below the noise floor making it more difficult to detect and decode them in the presence of strong self-interference. Finally, scaling underwater UWB backscatter to multiple nodes will introduce new forms of interference and require new mechanisms to disentangle interfering signals and decode them in frequency-selective and time-varying underwater environments.

At the heart of U^2B's approach to overcoming these challenges is a metamaterial-inspired transducer design for underwater backscatter. Metamaterials are artificial materials (or material composites) that exhibit properties which would otherwise not occur in nature. In the context of U^2B, we needed to develop a backscatter transducer that exhibits wideband properties.

Before we describe our new transducers, let us understand why traditional low-power designs have a limited bandwidth. Fig. 2-1(a) shows a typical underwater transducer consisting of a piezoceramic cylinder. Piezoceramic cylinders are used for transmitting and receiving sound underwater because they can transform sound to electric signals, and vice versa. The geometry of these transducers determines their resonance frequency. In the above figure, the resonance frequency is determined by the thickness, radius, and material composition of the piezoceramic cylinder. (By analogy, guitar strings have a resonance frequency that is determined by their thickness, length, and material). Existing underwater communication systems leverage this property to maximize the efficiency of backscatter and energy harvesting by transmitting sound at the resonance frequency of the transducer.[2] However, the signal-to-noise ratio (SNR) and energy harvesting efficiency significantly degrade outside this frequency (as shown in Fig. 2-1(a)). In principle, it is possible to use a non-resonant transducer; however, such a design is undesirable as it would be inefficient and would further limit the range and SNR of underwater communication [74, 19].

To overcome the bandwidth limitation while maintaining high efficiency, U^2B synthesizes different forms of resonance through the multi-layer (metamaterial) design

---

[2] While past systems have tried to shift the resonance electrically [37], they can only do so by a small amount (around 2-3kHz) but remain limited by the mechanical resonance of the piezoceramic cylinder's geometry.

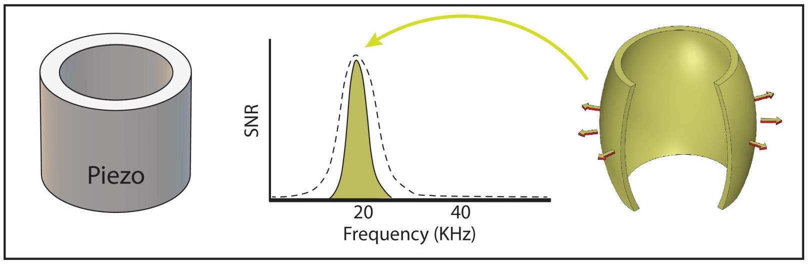

(a) Traditional Single-Layer Resonator

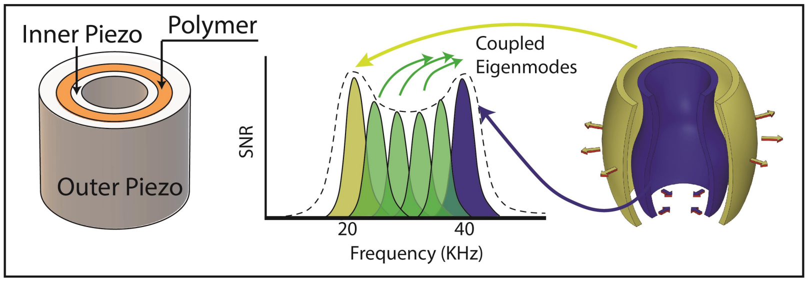

(b) U^2B's Metamaterial Design

Figure 2-1: U^2B **Synthesizes Resonances to Achieve UWB Performance.** In each of (a) and (b), the left figure shows the transducer architecture, the middle figure shows the SNR as a function of frequency, and the right figure shows one of the vibration modes – or eigenmodes – of the active layers.

shown in Fig. 1(b). The design alternates between active (piezoelectric) layers and passive (polymer) layers. It makes use of two kinds of resonance: the first (shown in blue and yellow) is similar to that of standard transducers and arises from the primary resonance of each individual piezoceramic layer. The second and more interesting, kind of resonance arises from the passive coupling between the different active layers. Even though the polymer layers are themselves passive, by sandwiching them between two piezoceramic layers, we impose new constraints on how the active layers may stretch and squeeze concurrently. This coupled interaction yields additional resonances that fill in the gaps between the first two types of resonances, thus resulting in wideband operation (depicted by green curves in Fig. 2-1(b)). Each of these resonances corresponds to a complex vibration mode – called *eigenmode*. The combination of these eigenmodes allows U^2B to achieve high-efficiency backscatter over a wide bandwidth. In 2.2.1, we explain this multi-layered design in detail, and in 2.3.1, we describe how U^2B can independently or jointly activate different layers for communication and energy harvesting.

A major benefit of UWB RF technologies is that they can extend the communication range between two nodes by allowing them to decode packets that are received below the noise floor [77]. Specifically, a UWB transmitter can apply a spread-spectrum code (e.g., CDMA) before transmitting its packet, and a UWB receiver can correlate with the same code to boost the received signal and decode packets that are orders of magnitude below the noise floor.[3] In the context of underwater backscatter, however, the communication range is not only limited by the noise floor, but also by the strong self-interference between the transmitted and received signals (since the nodes communicate by modulating the reflection of a continuous downlink signal sent from a remote transmitter). This self-interference limited past systems' operational range to within a few meters [37].

To overcome this challenge, U^2B exploits the multi-resonant wideband transducer and shifts the backscatter response to an out-of-band channel. Specifically, U^2B's ultra-wide bandwidth allows it to divide its available spectrum to multiple orthogonal frequency channels and to use separate channels for downlink and uplink communication. Let us assume the downlink packet is transmitted at 40 kHz. The backscatter node can leverage the non-linear nature of backscatter to shift the uplink response to a different center frequency, e.g., 20 kHz. A remote receiver that obtains both the strong downlink (at 40 kHz) and weak backscatter uplink (at 20 kHz) can simply filter out the downlink signal in hardware. This approach allows it to cancel self-interference and decode the weak backscatter packet even if it is orders of magnitude weaker than the transmitted signal or the noise floor.

U^2B's design builds on the above primitives to address many of the limitations of prior low-power subsea IoT proposals. Beyond increasing the communication range, it uses the wide bandwidth to boost throughput and increase the number of frequency channels available for communication, thus scaling the subsea IoT to dozens of nodes.

We implemented a prototype of U^2B and tested it in a river and indoor pools. We mechanically fabricated our metamaterial transducers in-house; our fabrication process (detailed in 2.3) involved designing and 3D printing the molds, stacking

---

[3] By analogy, GPS receivers use such coding mechanism to decode signals with SNRs of -20 dBm.

(electrically and mechanically) piezoceramic transducers, and curing the transducers under high-pressure environments. Our prototype evaluation was performed with integrated energy-harvesting micro-controllers as well as with software radio-based nodes (USRP N210 [22] with LFRX daughterboards [2]). Our empirical evaluation across around 400 experimental trials demonstrated the following results:

- In comparison to single-layer backscatter nodes, U^2B's design boosts the throughput from 3kbps (of the state-of-the-art system) to more than 20 kbps. Such throughput is considered on the high end of underwater communication systems, and can enable new applications such as streaming low-resolution images (e.g., for coral reefs or aquafarms).
- U^2B's self-interference cancellation approach allows it to mitigate the direct signal from the projector. Practically, we show that it enables us to communicate at up to 60 m.

Contributions: U^2B is the first system that enables ultra-wideband underwater backscatter. Its design introduces a novel metamaterial-based transducer that synthesizes multiple resonances to achieve wideband operation. The design brings various capabilities to the subsea IoT domain, including out-of-band self-interference cancellation and multi-channel communication. The Chapter also demonstrates a prototype implementation and evaluation in challenging real-world environments.

2.1 The Narrowband Problem

We start by explaining why existing solutions for underwater backscatter have a limited bandwidth. Then, in 2.2, we explain how U^2B overcomes this problem by introducing an ultra-wideband design.

2.1.1 Underwater Backscatter

Recall from section 1.1 that in underwater backscatter the backscatter node communicates by modulating its reflections. Specifically, it can send bits of zero and one by alternating between reflective and non-reflective states.

The signal-to-noise ratio (SNR) of the received signal is determined by the difference between the reflective and non-reflective states. It is given by the following equation:

$$SNR = \frac{|P_{\text{reflective}} - P_{\text{non-reflective}}|^2}{|N|^2}$$

where $P_{\text{reflective}}$ and $P_{\text{non-reflective}}$ refer to the pressure received by the hydrophone[4] in the reflective and non-reflective states and N refers to the amplitude of noise. The square indicates the power.

We can simplify the above equation by rewriting it as a function of the reflection coefficient Γ and the incident pressure P_{incident} as follows:

$$SNR = |\Gamma_{\text{reflective}} - \Gamma_{\text{non-reflective}}|^2 \frac{|P_{\text{incident}}|^2}{|N|^2} \tag{2.1}$$

To modulate the reflection coefficient of sound, underwater backscatter nodes rely on piezoelectric materials. Specifically, by opening and shorting the terminals of that material (via a switch as shown in Fig. **??**), they can toggle between the two reflective states, thus enabling backscatter communication.

2.1.2 The Resonance Bottleneck

While using piezoelectric materials enables underwater backscatter, it limits the bandwidth of underwater communication. Specifically, piezoelectric materials must be operated at a specific resonance frequency – also called the *natural frequency* of the material. At resonance, they vibrate with a large amplitude, which leads to high efficiency in sending and receiving sound. A popular example of resonance is how an opera singer can break glass with her voice [69, 71]; in particular, when the singer's voice matches the resonance frequency of the glass, the glass vibrates with higher and higher amplitudes and eventually breaks [69, 71]. Similarly, when piezoelectric materials are excited by their resonance frequency, they generate stronger signals.

---

[4] Note that the amplitude of the electric field in RF is replaced by the amplitude of the pressure wave in acoustics.

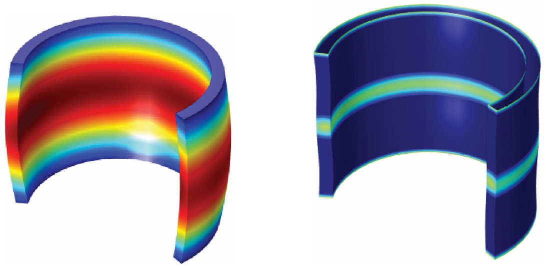

(a) At $f = f_{resonance}$ (b) At $f \neq f_{resonance}$

Figure 2-2: **Single-Layer Eigenmode.** The figure shows the amplitude of the piezo vibration across its body when it is excited (a) at resonance, and (b) outside resonance. The vibrations are visualized as a heat map where red and blue represent high and low vibration amplitudes respectively.

To gain more insight into the bandwidth problem, we simulated a piezoceramic cylinder's vibration at different frequencies and measured the amplitude of vibration. Fig. 2-2 visualizes the amplitude as a heatmap over the cross-section of the piezoelectric cylinder.[5] The heatmap shows the largest amplitude of vibration in red and the lowest amplitude of vibration in blue. Fig. 2-2(a) shows that when the cylinder is excited by its resonance frequency, it expands and shrinks in the radial direction with a large amplitude; this radial vibration is called an *eigenmode* and the corresponding frequency is called an *eigenfrequency*. Fig. 2-2(b) shows that when the same cylinder is excited by a frequency far from its resonance, the vibration amplitude is very small. This simulation verifies that operating at resonance is necessary to achieve high efficiency in transmitting and receiving sound. What is less clear – and has not been well-studied in past work – is why resonance limits the underwater backscatter bandwidth of piezoelectric materials. The rest of this section focuses on understanding this problem.

The impact of resonance on SNR. Recall from Eq. 2.1 that the SNR depends on the difference between the reflection coefficients in the two backscatter states. The reflection coefficient itself is a function of the acoustic impedances of water and the piezoelectric material. Mathematically, the reflection coefficient can be described

---

[5] We used the COMSOL Multiphysics software [1] to simulate the vibrations of the piezoceramic cylinder.

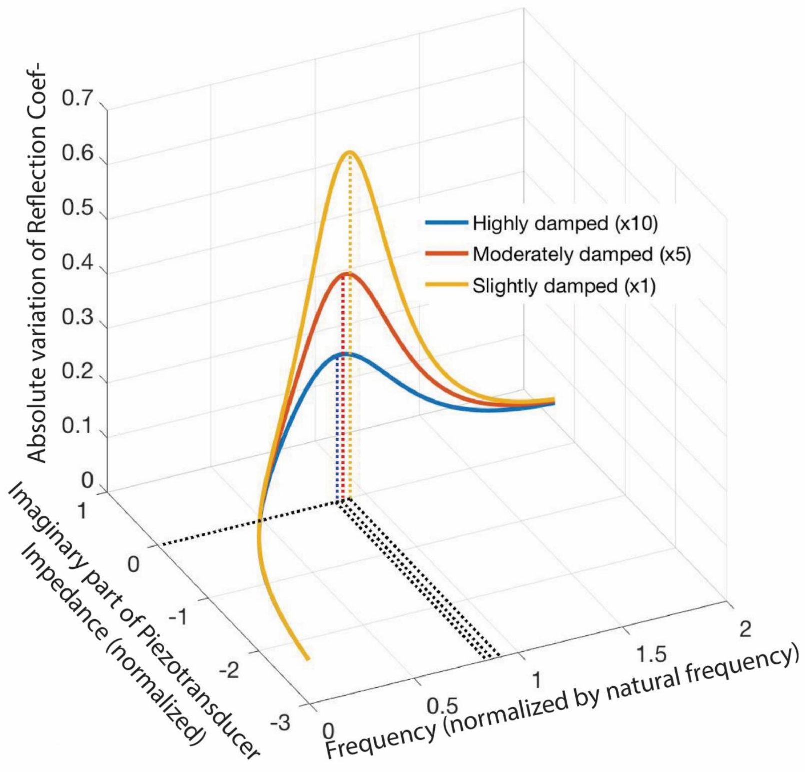

Figure 2-3: **Backscatter and Resonance.** The figure plots the absolute value of the difference in the reflection coefficient (between open and short cases) as a function of normalized frequency and as a function of the normalized imaginary component of piezo-transducer impedance.

using the following equation:

$$\Gamma = \frac{Z_{piezo} - Z_{water}}{Z_{piezo} + Z_{water}} \tag{2.2}$$

where Z_{piezo} and Z_{water} refer to the acoustic impedances of the piezoceramic cylinder and water. During backscatter, the terminals of the piezoelectric are opened and shorted, resulting in different piezoelectric impedances Z_{piezo} across the two states. This, in turn, modulates the reflection coefficient and enables backscatter.

Next, we would like to understand the impact of frequency on the SNR. Fig. 2-3 plots the absolute value of the difference in reflection coefficients (between open and short cases) as a function of frequency and as a function of the impedance of Z_{piezo}. We make the following remarks from the figure:

- At the resonance frequency, the difference in reflection coefficients between the two states is maximized, resulting in the highest SNR. The SNR degrades quickly outside resonance.

- The figure also plots the difference in reflection coefficients as a function of the imaginary part of the impedance. It demonstrates that the SNR is maximized when the imaginary component of the impedance is zero. Intuitively, the imaginary part is zero when the vibration of the cylinder's wall and the acoustic wave are in-phase (i.e, the phase ϕ of the complex impedance is zero). In reference to the opera singer analogy, when the singer's voice and the glass vibrate in-phase, the glass's amplitude increases.

  We can also show this mathematically. The impedance is a function of the cylinder's mass M, its stiffness K, and its resistance R. Formally, we show in the appendix that the backscatter acoustic impedance can be given by the following equation:

  $$Z_{piezo} = R + j\omega\left(M - \frac{K}{\omega^2}\right)$$

  The above equation reaffirms that the absolute value of the impedance is minimized when its imaginary part is zero. The equation also shows that the resonance frequency is $\omega = \sqrt{K/M}$.

- The figure repeats the same simulation for different resistance values R, which correspond to different damping coefficients. While more damping leads to lower SNR, it does not impact the resonance frequency.

Before we move on to describe 's approach to overcoming this narrowband problem, it is worth reflecting on the difference in effective bandwidth between underwater piezoelectric backscatter and RF backscatter (e.g., RFIDs). RF backscatter operates at a much higher center frequency (GHz) than acoustic/ultrasonic backscatter ($\sim$ tens of kHz). As a result, a communication bandwidth of 20 kHz for piezoacoustic backscatter would be of the same order as its center frequency, but that same bandwidth would be less than 1% the center frequency of RF communication. As per Fig. 2-3, this makes it significantly more difficult for piezo-acoustic backscatter than for RF backscatter to achieve high SNR across the same bandwidth (since the absolute value of the difference in reflection coefficient is a function of the *normalized* frequency).

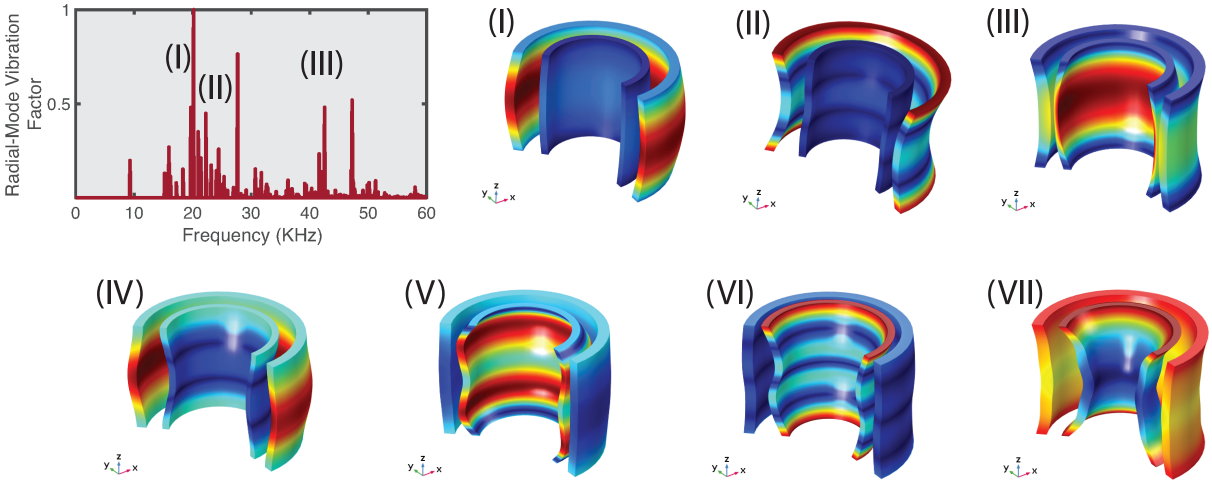

Figure 2-4: **Eigenmodes and Eigenfrequencies of U^2B's Metamaterial Design.** The figure in the top left shows the vibration amplitude across the different eigenfrequencies. (I)-(VII) depict the eigenmodes associated with some of the eigenfrequencies. The vibration amplitude for each of these eigenmodes is visualized as a heatmap across the body of the transducer, where red indicates a high amplitude vibration and navy blue indicates low amplitude vibration.

2.2 Ultra-wideband Meta-material Design

So far, we have explained why existing solutions for underwater backscatter suffer from a narrow communication bandwidth. In this section, we describe how U^2B overcomes this challenge via a metamaterial design, and how its design boosts the throughput, range, and scale of underwater backscatter networking.

2.2.1 Coupled Eigenmodes

A strawman solution to the bandwidth problem is to connect multiple piezoelectric cylinders, each having a different resonance frequency. However, such an approach would result in bulky and costly nodes, and it does not scale well to a wide bandwidth since it requires adding a new cylinder for each frequency. The design would also suffer from a directionality problem since placing multiple cylinders side-by-side would mimic a fixed antenna array (beamforming) behavior.

To provide a scalable, cost-effective, compact, and omnidirectional solution, U^2B adopts a multi-layer architecture that alternates between active and passive layers. This architecture is shown in Fig. 2-1(b) and consists of two active piezoelectric layers and a passive polymer layer sandwiched between them.

Despite its apparent simplicity, U^2B's multi-layer design has powerful properties that enable it to operate over a wide bandwidth. First, because its layers are concentric, the transducer does not suffer from a directionality problem. Second, because the polymer between the two active layers is mechanically compliant (i.e., flexible), it allows the inner and outer layers to vibrate independently to some extent. This enables the overall structure to inherit the Eigenfrequency and Eigenmode of both Piezoelectric layers without suffering from the problems of the strawman approach described earlier. Finally, even though the polymer itself is flexible, it imposes new constraints on how the two layers can vibrate with respect to each other, thus resulting in a coupled behavior. This coupling provides additional degrees of freedom that lead to new Eigenmodes and Eigenfrequencies that would not be present in any of the individual active piezoelectric layers.

We simulated U^2B's design at different frequencies and show some of the resulting Eigenmodes and Eigenfrequencies in Fig. 2-4. Similar to our earlier visualization, we overlay a heatmap of the vibration amplitude on the cross-section of each of the two layers. Consider mode (I) in the center top of the figure. This mode results from exciting the transducer at a frequency of 20 kHz, and it corresponds to the original radial vibration of the outer piezoelectric cylinder. Notice how, in this configuration, the outer layer exhibits large vibrations, while the inner layer is relatively static. This verifies that U^2B can indeed inherent the eigenmode of its active layers.

Next, consider mode (III) in the top right of the figure. This mode is interesting because both the inner and outer piezos exhibit large radial vibrations as demonstrated by the red and yellow central regions in both of them. This eigenmode corresponds to the scenario where both layers vibrate by bending in opposite directions to each other. Physically, this mode arises from the fact that both piezoelectric cylinders have the same height, which results in a bending mode vibration that is enhanced by stacking them together.[6]

It is important to note that not all eigenmodes of this multi-layer design have a simple physical interpretation. In particular, the bottom row of Fig. 2-4 shows a number of Eigenmodes which would not have been possible without the coupling between the two active layers. In mode (V), for example, we can see that the internal layer itself splits into two concurrent radial modes in the top and bottom half of the cylinder. Such a vibration mode would not be natural for a single-layer piezoelectric cylinder in the absence of coupling.

The top left of Fig. 2-4 plots the vibration amplitude as a function of frequency. Surprisingly, the figure demonstrates that U^2B has a large number of eigenmodes between 20 kHz and 60 kHz due to its coupled multi-layer design. Mathematically, it is possible to derive these eigenmodes by solving the following differential equation, along the different degrees of freedom of this meta-material structure [56]:

$$det(\omega^2 M - K) = 0$$

To solve this differential equation, we used the COMSOL Multiphysics software [1], and we defined the coupling between the layers (both active and passive) by applying continuity constraints on both pressure and velocity across the boundaries. The wealth of resonances arising from this multi-layer design signals that U^2B would indeed exhibit the desired ultra-wideband behavior. Specifically, the presence of multiple Eigenmodes and Eigenfrequencies would ensure that the backscatter SNR remains sufficiently high over a wide bandwidth.

Experimental Validation in a River. Next, we would like to empirically verify that U^2B exhibits wideband behavior. We ran an experiment to compare the SNR

---

[6] Note that this bending mode would also exhibit for each cylinder alone, but would be weaker than the radial mode (I).

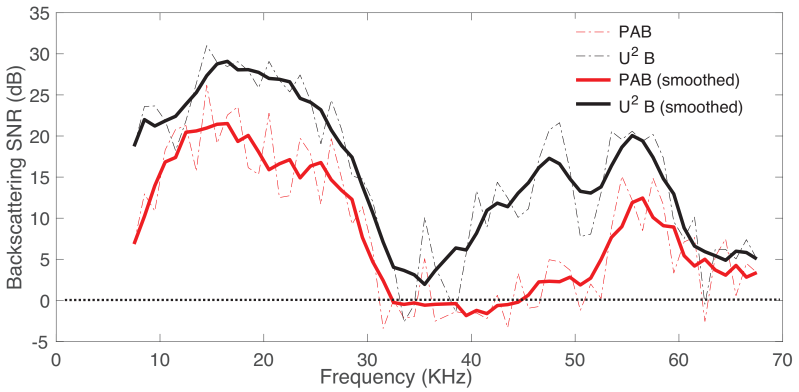

Figure 2-5: **Characteristic Curve for Single and Multi-layer Designs.** The figure plots the characteristic curve for single-layer (in red) and multi-layer (in black) designs as a function of frequency.

of a U^2B node to that of a state-of-the-art underwater backscatter node, PAB, which uses a single-layer piezoelectric resonator. Our experiments were performed in the river which had a depth of around 4 m. In these experiments, the acoustic transmitter sends a downlink signal at a frequency of 37.5 kHz. The backscatter node modulates this signal at different rates, and the hydrophone measures the SNR of the received signal at the corresponding frequency. (The design of the backscatter node is detailed in 2.3).

Fig. 2-5 plots the SNR of the received signal as a function of frequency for both single-layer (PAB, in red) and multi-layer (U^2B, in black) designs. The figure plots the exact values of the measured SNR with dotted-dashed lines, and the smoothed SNR curves with solid lines. In particular, since the experiment is performed in a river (i.e., an uncontrolled testing environment), the communication channel is impacted by various factors – including wind, river current, and depth – and changes rapidly over time. These changes in the channel impact the result by creating high-frequency variations that can be seen in the exact values of the resulted SNR. Therefore, we rely on the smoothed/averaged version of the SNR, shown with the solid line, to characterize the performance of the two designs.

We make the following observations:

- First, the overall SNR of a U^2B node is higher than that of the single-layer node, PAB, across all frequencies. The increased SNR results in higher channel capacity – which enables achieving higher throughputs (as we demonstrate empirically in 2.4.1) – and enables communication over further distances (as we demonstrate empirically in 2.4.2).

- Second, in the single-layer PAB design, the backscatter SNR is relatively limited to around 20 kHz (from 10-30 kHz); outside this band, the SNR significantly decays due to being far from resonance. The limited bandwidth in PAB's design limits both the throughput and the available spectrum for communication. In contrast, U^2B's multi-layered design exhibits stronger SNR performance over a wide bandwidth due to its coupled resonance modes. Specifically, the backscatter signal's SNR in U^2B remains sufficiently high across the entire band spanning 10 kHz to 60 kHz.

- Interestingly, U^2B achieves higher SNR than PAB even between 10 kHz and 30 kHz. There are two reasons for this improvement. First, the acoustic impedance matching for U^2B is higher than that of PAB; this is because the added polymer layer in U^2B reduces its density, making its acoustic impedance closer to that of water. Second, even though the downlink frequency in this region is close to only one of the active piezoceramic layers, the other layer would still vibrate, thus enhancing the overall efficiency of U^2B.

  U^2B It is worth noting that the backscatter performance is not frequency flat and exhibits a dip around 35 kHz. This behavior mimics the theoretical simulation in the top left of Fig. 2-4, which also exhibits a dip around that frequency.

  U^2B Another point worth noting is that the single-layer node exhibits a small peak around 55 kHz. This peak is corresponds to the bending eigenmode due to the length (rather than radius) of the cylinder. This corresponds to mode (III) in Fig. 2-4, and is reinforced in the multi-layer structure due to coupling.

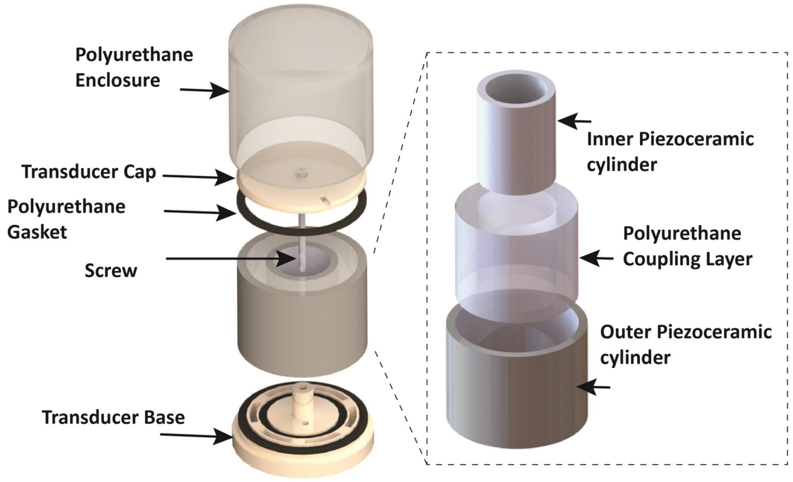

Figure 2-6: **Exploded Transducer View.** The figure shows the exploded transducer view and the layered piezoceramic structure for U^2B.

2.3 Fabrication & Implementation

In this section, we first explain the fabrication process of U^2B, then elaborate on our overall setup.

2.3.1 Mechanical Fabrication

Fig. 2-6 shows an exploded view of our transducer design. The transducer consists of a layered piezoceramic structure. We purchased two different types of piezoceramic cylinders from Steminc [3] to build U^2B. The outer piezoceramic cylinder has a nominal resonance frequency of 17 kHz and has an outer radius of 27mm, inner radius of 23.5mm, and height of 40mm [82]. The inner piezoceramic cylinder has a nominal resonance frequency of 30 kHz and has an outer radius of 18mm, inner radius of 15.5mm, and height of 20mm [81]; we stacked two of the inner piezoceramic cylinders and soldered them together to obtain the same height as the outer cylinder. We then soldered wires to the inner and outer surfaces of each of the cylindrical layers, resulting in a total of four terminals for each U^2B transducer. To couple the vibration of the two piezoceramic cylinders, we filled the space between them with polyurethane. It is very important to ensure that this polyurethane coupling layer is entirely filled such that it is in contact with the walls of both piezoceramic layers. Otherwise, the possible delamination may significantly deteriorate the coupled vibration mode.

The fabrication process of U^2B starts by 3D printing the base and cap of the transducer. Then, we laser cut a polyurethane gasket with 2mm thickness to match the size of each of the two cylinders. We placed the gasket and then the piezoceramic cylinders in their designated locations on the base, then screwed and tightened the cap and base together. This design process ensures that the inside of the internal piezo remains air-backed even after dipping a U^2B node in the water. This design methodology, referred as air-backed design, has been shown to yield to better performance [37].

Next, the capped piezocylinders were placed inside the 3D printed mold, and we poured the polyurethane WC-575A mixture from BJB enterprise into the cylinder mold [4]. We used a transparent cylinder mold to ensure that the outer surface of the entire structure is covered by the polyurethane mixture, which is needed to insulate it from the surrounding environment. Notice that the transducer base, as shown in Fig. 2-6, has several openings in the region sandwiched between the two piezoceramic cylinders. These openings guide the poured polyurethane polymer to completely fill the gap and spacing between two piezoceramic layers. Afterwards, the structure is placed for 12 hours inside a pressure chamber at a pressure of 60psi to remove any bubbles from the polyurethane solution. We also experimented with using a vacuum chamber instead of a pressure chamber to remove the bubbles; however, this approach was less effective in removing the bubbles because of the high viscosity of the polyurethane mixture.

Fig. 1-4 shows the transducer before and after potting. In Fig. 1-4(a), the inner and outer piezoceramic cylinders are placed on the transducer's base, and are separated by empty space (this separation is eventually filled with the polyurethane mixture for coupling); the figure also shows the outer mold and top cap of the transducer. Fig. 1-4(b) shows the potted transducer with four electrodes; the electrodes are connected to the inner and outer surfaces of each of the two piezoceramic layers. This enables experimenting with different configurations for backscatter and energy harvesting. The total cost of each fabricated node is around \$120, which includes the cost of the outer piezo (\$41) and the cost of the two inner piezos (\$35 each).

2.3.2 Hardware Design

To communicate via backscatter, underwater nodes short and open two terminals of a piezoelectric transducer as shown in Fig. **??**. While a single-layer transducer has only two terminals (positive and negative terminals of the piezoceramic cylinder), U^2B's multi-layer design has four terminals – two for each of its active layers. This provides different options for backscattering by shorting and opening any of the two terminals.

To enable backscatter while at the same time allowing U^2B to continuously harvest energy, we determined that the best approach is to connect the modulating transistor across the terminals of the outer layer while connecting energy harvesting circuit across the terminals of the inner layers. This approach allows U^2B to continuously operate without the need to alternate periods of energy harvesting and backscatter. Aside from the front-end, U^2B adopts a similar approach to prior designs in terms of rectification for energy harvesting.

In terms of communication protocols, U^2B adopts similar techniques to prior backscatter designs, where the downlink projector encodes bits via PWM and uplink (backscatter) communication is performed via FM0 modulation. In our evaluation, we experimented with both energy-harvesting backscatter and battery-assisted backscatter. Since energy-harvesting is not the focus of this Chapter, most of our results are reported from battery-assisted backscatter designs. The power consumption for U^2B is similar to PAB [37], which ranges from 120μW to 500μW. It is important to note that that battery-assisted nature does not discount the importance of backscatter since, being the lowest-power communication technology, it would significantly extend the battery life of underwater communication.

2.3.3 Evaluation Setup

Our evaluation setup is composed of three main components: an acoustic projector, a hydrophone receiver, and a number of U^2B nodes. Since we explained the design and fabrication of U^2B in previous section, below we will explain the other two main components along with the environment where testing is performed.

(a) Transmitter. To generate an acoustic signal, we used one of our in-house fabricated U^2B transducers (discussed in the previous section) as a projector. To drive the node as a projector, a sine wave with the desired amplitude and frequency is first generated using the Siglent SDG1032X arbitrary waveform generator [75]. Next, the output of the signal generator is fed to an XLi2500 Two-channel 750W power amplifier [16]. The output of the amplifier is connected to the terminals of the outer piezo in differential mode.

(b) Receiver. To record the signal, we used omnidirectional Reson TC 4014 hydrophone [86] (flat response from 15 Hz to 480 kHz) with the sensitivity of -180dB re: $1V/\mu Pa$ in the differential mode. The output of the hydrophone is connected to USRP N210 [22] and the data is collected with a sampling rate of 1MSps. To capture the differential output of the hydrophone, we used two LFRX daughterboards [2] whose grounds are connected together. Our decoding is performed offline in MATLAB.

(c) Testing Environments. Unless otherwise noted, all of our experiments were performed in a river. In the river, we chose the test location based on a number of factors including accessibility, depth, and water current. Some of our experiments were performed on rainy and snowy days. While changes in weather impacted the SNR and channel coherence time, U^2B was capable of operating correctly across different weather conditions.

2.4 Results

To evaluate the performance of U^2B, we performed controlled and uncontrolled experiments in the river as described in 2.3.3. Our experiments tested U^2B's through-

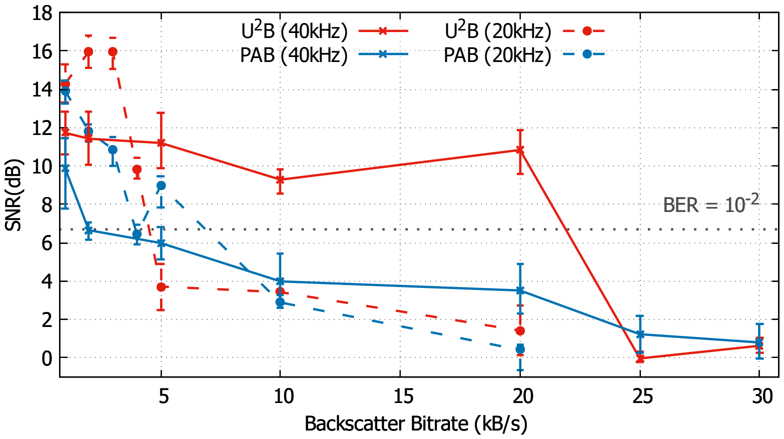

Figure 2-7: **SNR vs Bitrate.** The figure plots the median SNR as a function of bitrate for both U^2B (in red) and PAB (in blue) at two different downlink center frequencies: 20 kHz (dashed line) and 40 kHz (solid line). The error bars represents the 25^{th} and 75^{th} percentile and the dotted gray horizontal line represents the SNR needed to get a BER of 10^{-2}.

put, communication range, and its ability to enable concurrent transmissions. We performed around 400 experimental trials in total. We varied the backscatter rate, location, and depth of U^2B's transducers throughout these experiments.

2.4.1 Throughput

First, we are interested in evaluating U^2B's throughput, and comparing its throughput to the state-of-the-art underwater backscatter system, PAB. To evaluate U^2B's ability to communicate across different bitrates, we fixed the locations of the projector, hydrophone, and backscatter node, and we varied the backscatter bitrate. We repeated the same evaluation for both U^2B and PAB. We also tested each of them at two different downlink center frequencies: 20 kHz and 40 kHz. In each of our experimental trials, the hydrophone receives the backscatter node's packet and computes the SNR of the received signal. The SNR is computed as the ratio of the signal power (i.e., the square of the channel estimate) to the noise power.

Fig. 2-7 plots the median SNR as a function of the bitrate for PAB (in blue) and U^2B (in red) at the two center frequencies mentioned earlier. The dotted horizontal line in the figure represents the SNR needed to achieve a throughput of 10^{-2}, which is considered a standard threshold for reliable underwater communication [96]. We make the following remarks:

- U^2B maintains high SNR (i.e., above the dashed lines) at throughputs up to 20 kbps. In contrast, the PAB baseline cannot achieve good SNR above 5 kbps. These results are in line with those reported for PAB [37] and demonstrate that U^2B increases the throughput by about $4-5\times$ in comparison to the state-of-the-art baseline. This result is expected because U^2B's transducer enjoys a much wider bandwidth as we demonstrated experimentally in Fig. 2-5. U^2B The baseline, PAB, performs better when the downlink frequency is at 20 kHz vs at 40 kHz. This behavior is also expected according to the characteristic curve in Fig. 2-5, which showed that the system has a higher SNR around the 20 kHz region, which matches its primary resonance.[7] In contrast, U^2B performs better when the center frequency is 40 kHz, a result that is also expected since it has high SNR between 10 kHz and 60 kHz according to its characteristic curve.
- Both U^2B and PAB follow similar SNR curves when the downlink frequency is at 20 kHz. This is also expected since their characteristic curves are much closer in the bandwidth surrounding 20 kHz. Here, U^2B achieves slightly higher SNR because it is more efficient even at lower frequencies as described in 2.2.1.
- Beyond 20 kbps, the SNR drops close to 0 dB for both U^2B and PAB. This is expected because the backscatter node becomes inefficient beyond 60 kHz (i.e., 40 kHz center frequency + 20 kHz backscatter rate) which can also be seen from the characteristic curve in Fig. 2-5.

2.4.2 Range

We would like to evaluate the range at which we can reliably communicate with U^2B. In this experiment, we varied the distance between the backscatter node and the hydrophone/projector pair. The hydrophone and projector were co-located (1m

---

[7] Note that the SNR from Fig. 2-5 is that of a single-carrier as described in 2.2, which is different from the communication SNR described in this section.

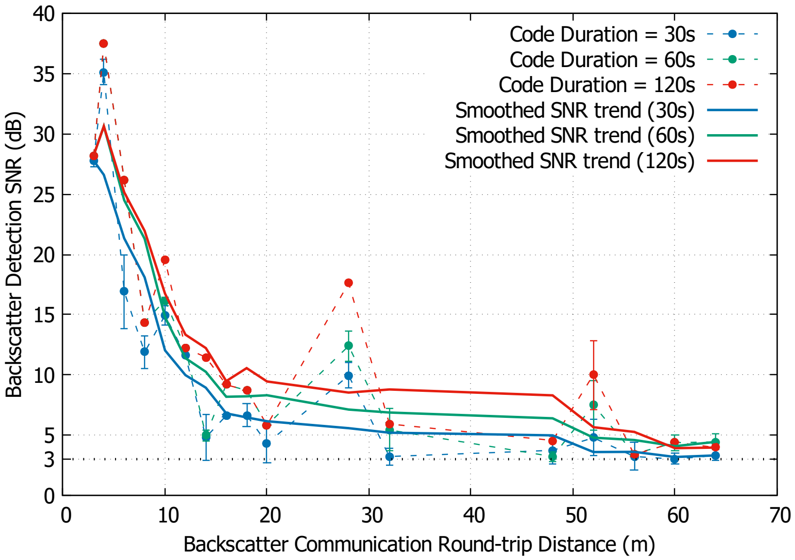

Figure 2-8: **Backscatter SNR vs Distance.** The figure plots the backscatter SNR as a function of distance for different coding rates. Error bars indicate standard deviation, and solid lines indicate smoothed curves.

apart) and the backscatter device was moved further away (up to 62m round trip). At each distance, we applied a CDMA-style code and measured the received signal's SNR after coding. To deal with the low power of the backscatter response, we applied codes of different lengths: 30s, 60s, and 120s. We repeated each experiment at least 3 times at each rate and distance.

Fig. 2-8 plots the SNR as a function of distance for each of the coding rates. We make the following remarks:

- The figure shows that U^2B nodes can be used to communicate up to 60 m (SNR $>$3dB).[8] Such range represents more than 6× improvement over the baseline (PAB) [37].

- The SNR decreases with distance across different coding rates. This is expected because the signal power decays with distance. Moreover, the trend follows the theoretically expected decay resulting from the round-trip path loss.

- Applying longer codes enables us to communicate over further distances. Specifically, 30s, 60s, and 120s codes can communicate up to 32m, 48m and 60m

---

[8] Note that we used the 3 dB threshold here instead of the 7 dB one since we use coding.

respectively. Moreover, since our USRP receiver is limited to a 14-bit ADC, it would be possible to communicate over longer ranges with higher-bit ADCs due to the higher dynamic range.

2.5 Conclusion

Motivated by recent advances in underwater backscatter and ultra-low power networking, this Chapter introduces a new design, U^2B, that marks a significant step forward in this domain. The design bridges recent advances in metamaterials to the underwater backscatter problem, and demonstrates significant improvements over state-of-the-art proposals. As the research evolves, we hope that these techniques would enable truly ubiquitous subsea IoT systems that can be used for climate change monitoring, marine life sensing, and ocean exploration.

Chapter 3

Pushing the Network Throughput of Underwater Backscatter

In this chapter, we ask the following question: *Can we bring higher-order modulation to underwater backscatter sensors while maintaining their battery-free nature?* A positive answer to this question would allow us to achieve higher throughput while maintaining the ultra-low-power and low-cost nature of backscatter communication.

To understand the difficulty in achieving higher-order modulation with underwater backscatter, let us first briefly recall how existing backscatter sensors operate. These sensors differ from traditional communication modems in that they do not expend their own energy for transmission. Instead, they communicate by modulating the reflections of existing underwater acoustic signals. Specifically, they transmit bits of '0' and '1' by switching between reflective and non-reflective states. A remote receiver can sense these acoustic reflections to decode the transmitted packets from the backscatter node.[1]

While the backscatter approach described above enables ultra-low-power communication, it suffers from limited modulation capabilities. In particular, existing underwater backscatter nodes, like that shown in Fig. 3-1(a), can alternate between only *two* reflective states (as discussed in section 1.1). This can be done by toggling a

---

[1] Note that such reflections can be omnidirectional. Moreover, backscatter nodes apply a code on their transmissions, thus enabling a receiver to use this code in order to isolate their packets from other unmodulated reflections in the environment.

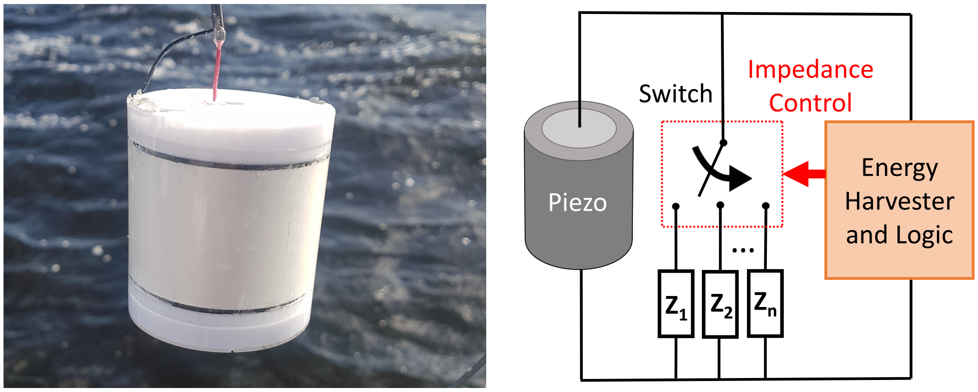

(a) Underwater Backscatter Node (b) Higher-Order Modulation

Figure 3-1: ***PAB-QAM*** (a) shows one of our potted *PAB* transducers (b) shows the schematic design for *PAB-QAM*.

switch (transistor) on and off between the two terminals of their piezoelectric transducer. As a result, they cannot realize higher-order modulation schemes like QAM (Quadrature Amplitude Modulation) which require modulating both the phase and the amplitude of their transmissions. Unfortunately, this limitation has prevented prior designs from achieving higher throughputs via higher-order modulations.

We present *PAB-QAM*, the first underwater backscatter sensor design capable of higher-order modulation. *PAB-QAM* exploits the electro-mechanical coupling property of piezoelectric transducers in order to achieve complex reflection coefficients and realize a larger number of reflection states. Fig. 3-1(b) shows the schematic of our design for *PAB-QAM*, which consists of a piezo-electric transducer connected to a multi-load switching network. At a high level, the system is able to achieve multiple reflection states by toggling between different complex load impedances Z_i connected between the terminals of the backscatter sensor node. Specifically, the electrical impedance of the load determines the vibration amplitude and phase of the piezo-electric resonator (i.e. reflection). A receiver can sense these different reflection states and use them to decode the transmitted messages. Hence, by enabling more than two reflection states, this design allows us to transmit multiple bits per symbol, thereby achieving higher throughput than prior underwater backscatter designs.

In the rest of this chapter, we describe the theoretical and practical aspects of the design and implementation of *PAB-QAM*. We also present an empirical evaluation that demonstrates how our approach can double the throughput of underwater backscatter communication, thus paving way for richer sensing applications in the context of subsea IoT. Finally, we demonstrate how we can enable robust full-duplex backscatter communication by employing self-interference cancellation and we also describe how our design can scale to a large number of nodes using a MAC protocol with concurrent transmissions.

3.1 System Design

3.1.1 Enabling QAM Backscatter

In this section, we describe how *PAB-QAM* builds on prior underwater backscatter designs (as discussed in 1.1) to achieve higher-order modulation. Recall that the challenge in achieving higher-order modulation is that existing underwater backscatter systems can only switch between two states, reflective and non-reflective. However, in order to achieve higher-order modulation, one needs to be able to switch between more states.

To overcome this challenge and realize a larger number of reflection states, *PAB-QAM* exploits the electro-mechanical coupling properties of piezoelectric transducers. Specifically, the reflection properties of a piezoelectric material are determined both by the material itself (and geometry), and by the electrical load connected to it. In particular, due to the coupling between the electrical and mechanical components of piezoelectric materials, changing their electric loads would result in changing the amplitude and phase of their vibration. Since each vibration mode results in a different reflection state, this allows us to encode more than two types of symbols (i.e., more than 1 bit/symbol). Interestingly, the reflective (closed switch) and non-reflective (open switch) states of prior underwater bacskcatter designs correspond to two special cases of this interpretation, where the load impedance is 0 and ∞ respectively.

Mathematically, we can describe the backscatter signal as a function of the reflection coefficient Γ. The reflection coefficient itself can be expressed as a function of

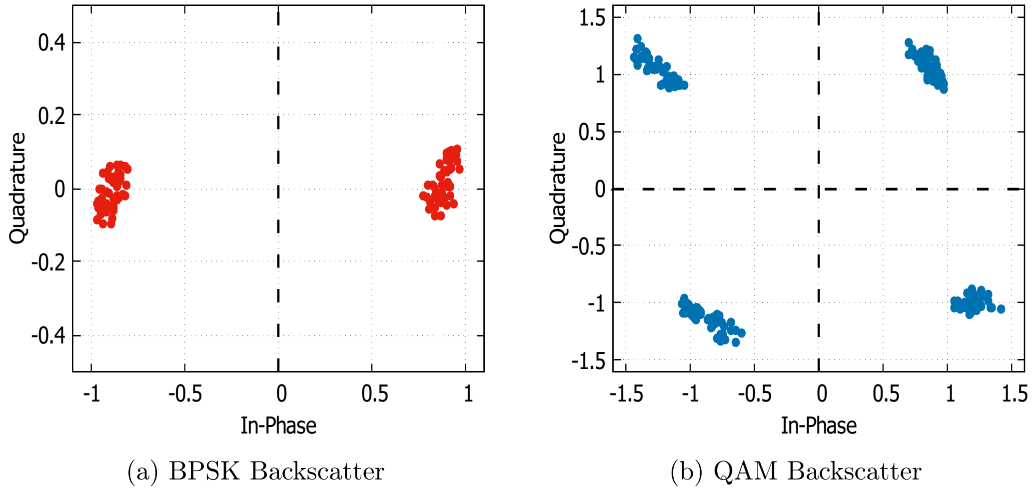

(a) BPSK Backscatter

(b) QAM Backscatter

Figure 3-2: **Backscatter Constellation.** This figure plots the constellation points of the received backscatter signal for: (a) BPSK backscatter, and (b) *PAB-QAM*.

the source impedance Z_s (i.e., that of the piezo) and load impedance Z_l (i.e., that of the electrical load) through the following equation:

$$\Gamma = \frac{Z_l - Z_s^*}{Z_l + Z_s^*} \tag{3.1}$$

Since the source impedance Z_s is constant for a given carrier frequency, we need to modify the load impedance Z_l to obtain different reflection coefficient values. This can be done by choosing reactive components (e.g., inductors) as impedance loads. Fig. 3-1(b) shows how such a design can be realized in practice by having a switch that can toggle between unique reactive loads.

The above description demonstrates how one can achieve random backscatter reflection states. In practice, we would like to strategically choose these states to achieve high spectral efficiency and high modulation depth (i.e., to separate the reflection coefficients in order to increase their resilience to noise). To make this more concrete, we ran an experiment using a prior (open-close) backscatter design [37]. Fig. 3-2(a) plots the constellation of the received symbols in the I/Q plane. The plot shows that the received clustered are diametrically opposite at +/-1. This matches the expected constellations from standard BPSK modulation.

When adopting higher-order modulations, we would like to similarly separate the symbols from each other by the maximum possible distance in the I/Q plane. Such separation can typically be achieved by implementing higher-order modulation schemes like Quadrature Amplitude Modulation (QAM) [91]. Thus, we strategically chose the inductor values of the impedance loads in our implementation such that they resulted in QAM constellations. Then, we ran an experiment with a *PAB-QAM* node, and we plot the constellation of the received symbols in Fig. 3-2(b). The plot shows *PAB-QAM* can indeed encode four different symbols spanning the four quadrants of the I/Q plane – i.e QAM. Thus, it can assign two bits to each symbol (reflection state), which in turn allows it to achieve twice the throughput compared to standard underwater backscatter systems, as we demonstrate empirically in 3.2.

Few additional points are worth noting:

- So far, our discussion has focused on the implementation of QAM by using multiple complex loads. In principle, this same idea can be extended to achieve other high-order modulation schemes such as PAM, PSK, 16-QAM, etc.[2]

- An alternative way to achieve higher throughput is to increase the switching/backscatter rate. However, such an approach may be less desirable for multiple reasons: first, backscattering at a higher rate would consume more energy because the oscillator would need to run at a higher frequency. Second, higher backscatter rate would require larger bandwidth; however, since piezo-transducers are typically narrow-band, this would result in lower SNR (signal-to-noise ratio). Finally, since the underwater channel is frequency-selective, spreading the signal over a wider bandwidth would complicate the decoding process [84, 10]. Despite these drawbacks, one might still want to adopt higher rates for certain data-intensive applications. In such scenarios, one can combine *PAB-QAM* with recent ultra-wideband underwater backscatter designs (discussed in Chapter 2) [25] as two complementary mechanisms to boost throughput.

---

[2] It is worth noting that prior work has demonstratted QAM for RFIDs [88]. *PAB-QAM* is inspired by this line of work and brings higher-order modulation to underwater backscatter.

- Finally, it is worth noting that our above discussion has ignored the impact of the channel on the constellations. In practice, one can estimate and invert the channel using standard channel estimation techniques from the communication literature [91]. This is typically done by transmitting a training sequence in the preamble p_n, and estimating the channel h from the received signal y_n as follows:

  $$h = \sum_n y_n . p_n^* \tag{3.2}$$

  Once the channel is estimated, it can be used to decode the payload of the packet.

3.1.2 Self-Interference Cancellation

Backscatter communication is inherently full duplex, since the backscatter node communicates by modulating the reflections of a downlink signal. As a result, the hydrophone receives a strong downlink signal from the projector as well as a weak backscatter reflection from a low-power node. Unfortunately, the projector's direct signal can overwhelm the receiving hydrophone and prevent it from detecting the weak backscatter response. Note that increasing the transmit power from the projector would not solve the problem because it boosts both the downlink signal as well as the backscatter response. In fact, boosting the transmit power would saturate the analog-to-digital converter (ADC) of the receiver and prevent it from being able to sense the backscatter response even if the backscattering node is nearby.

To better understand this challenge, let us assume that the projector transmits the downlink signal at $f_c = 40\, kHz$. Additionally, let us assume that the backscatter node wishes to transmit a packet $p(t)$ on the uplink. To do so, the node can simply apply $p(t)$ across the backscatter switch, resulting in the following signal being received by the hydrophone:

$$y(t) = I\cos(2\pi f_c t) + hp(t)cos(2\pi f_c t)$$

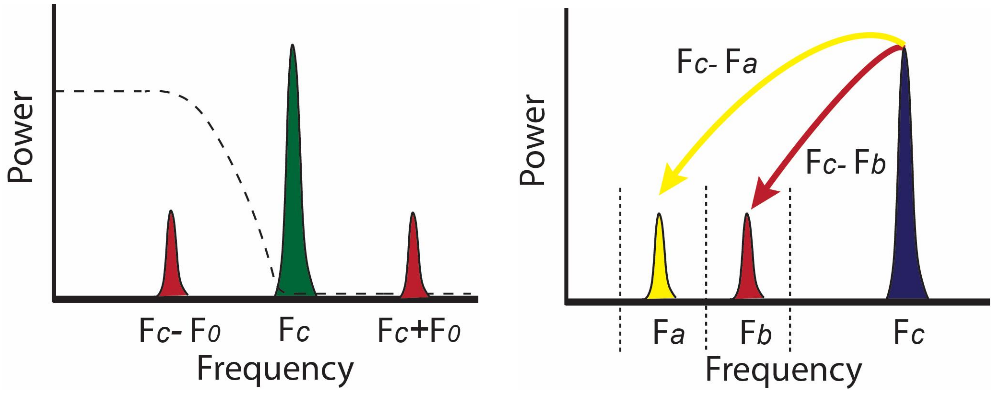

(a) Self-interference Cancellation

(b) FDMA MAC

Figure 3-3: **Mitigating Interference and Scaling and the Size of *PAB-QAM* networks.** The figure plots the power of the signal received by a hydrophone as a function of frequency. (a) shows how *PAB-QAM* can eliminate the downlink self-interference (in green) and retain a backscatter response (in red). (b) shows how *PAB-QAM* can scale to multiple nodes by assigning a different frequency channel to each of them.

where I is the amplitude of the direct path from the projector and h is the channel from the backscatter node to the receiver. Because $I >> h$, we will have self-interference which limits the ability to sense the backscatter response of a faraway node.

To overcome this challenge, $PAB-QAM$ leverages its wideband front-end, shifts the backscatter response out-of-band, and filters out the strong in-band downlink signal as shown in Fig. 3-3(a). Specifically, rather than just applying $p(t)$ across the backscattering switch, it can apply $p(t)\cos(2\pi f_o t)$ to shift the response away from the downlink signal in the frequency domain. In such a scenario, the received signal by the hydrophone is given by the following equation:

$$\begin{aligned} y(t) &= I\cos(2\pi f_c t) + hp(t)cos(2\pi f_c t)cos(2\pi f_o t) \\ &= I\cos(2\pi f_c t) + 0.5hp(t)cos(2\pi(f_c - f_o)t) \\ &\qquad + 0.5hp(t)cos(2\pi(f_c + f_o)t) \end{aligned}$$

So, the received signal contains the downlink frequency at f_c and the backscatter response at $f_c - f_o$ and $f_c + f_o$. Assuming $f_o = 25\ kHz$, then the backscatter response is at 15 kHz and 65 kHz, while the downlink signal remains at 40 kHz as shown in

Fig. 3-3(a). Thus, the receiver can simply apply a low-pass filter whose cutoff is 30 kHz to eliminate the downlink signal. This leaves it with only the backscatter response at 15 kHz allowing it to easily decode it.

Few points are worth mentioning about *PAB-QAM*'s approach for self-interference cancellation:

- First, by mitigating the strong downlink signal, the projector can transmit higher power and receive a stronger response from further away backscatter node. We can also place a low-noise amplifier (LNA) after the filter but before the ADC to further improve the sensitivity of our receiver to a further away node.

- In our above discussion, we employed a low-pass filter to eliminate the downlink signal. Alternatively, it is possible to employ a high-pass filter (e.g., around 50 kHz in the above example) to mitigate self-interference. In such an implementation, the receiver retains the backscatter response at 65 kHz and can still decode it. Naturally, it can also employ a notch filter around 40 kHz and retain both 15 kHz and 65 kHz to obtain an even better SNR.

- One might wonder whether it is possible to simply apply a notch filter without shifting the backscattered response out of band. While this is possible in theory, in practice, it would require a very high-Q filter which may be infeasible and it would result in mitigating not only the self-interference, but also the backscatter response [49].

- Another interesting question is whether applying $f_o = 25\ kHz$ would defeat the purpose of backscatter since it is of the same order of the downlink frequency $f_c = 40\ kHz$. Recall that the primary benefit of backscatter is that the receiver does not need to amplify its own generated signal since it communicates by reflecting a powerful downlink signal. In practice, we found that if the hydrophone is placed very close to a *PAB-QAM* node, it would indeed sense a much weaker signal at f_o in comparison to a much stronger response at $f_c - f_o$ and $f_c + f_o$. This verifies the benefits of backscatter despite the apparent concern.

3.1.3 Scaling to Many Nodes

Next, we describe how *PAB-QAM* can scale underwater backscatter networks to a larger number of nodes. At a high level, the system employs an FDMA style MAC protocol to enable concurrent transmissions. It divides the overall available bandwidth into different sub-channels and allocates each sub-channel to a different backscatter node.

To understand how *PAB-QAM* supports concurrent transmissions, let us consider a simple network with two backscatter nodes A and B, and let us assume that the projector wants to allocate channels centered at f_A and f_B to these two nodes, as shown in Fig. 3-3(b). To do so, the projector needs to instruct each of the two nodes to shift by the appropriate frequency in order to occupy their respective channels. Hence, it sends a downlink command signal, instructing node A to apply a shifting frequency of $f_c - f_A$ and node B to apply a shifting frequency of $f_c - f_B$, which would result in the desired shifts as shown in the figure.

Mathematically, the hydrophone receives the following signal:

$$y(t) = (I + h_A S_A(t)\cos(2\pi(f_c - f_A)t) + h_B S_B(t)\cos(2\pi(f_c - f_B)t))\cos(2\pi f_c t)$$

This results in a shifting behavior similar to that described in 3.1.2, as demonstrated in the figure. Subsequently, the receiver can simply apply a bandpass filter around f_A and f_B to decode the two concurrent transmissions. The same idea can be extended to a large number of nodes, each occupying a different subchannel f_i.

Due to the on-off switching nature of backscatter communication, the resulting backscatter signal is a square wave (rather than a simple cosine). One challenge, here, is that the harmonics of the square wave may interfere with other nodes by overlapping with their channels. To this end, the projector can pre-determine orthogonal channels (i.e., ones whose harmonics don't interfere) in its channel allocation process. Moreover, *PAB-QAM* adapts its MAC protocol to maximize the spacing between the allocated channels, thus minimizing interference.

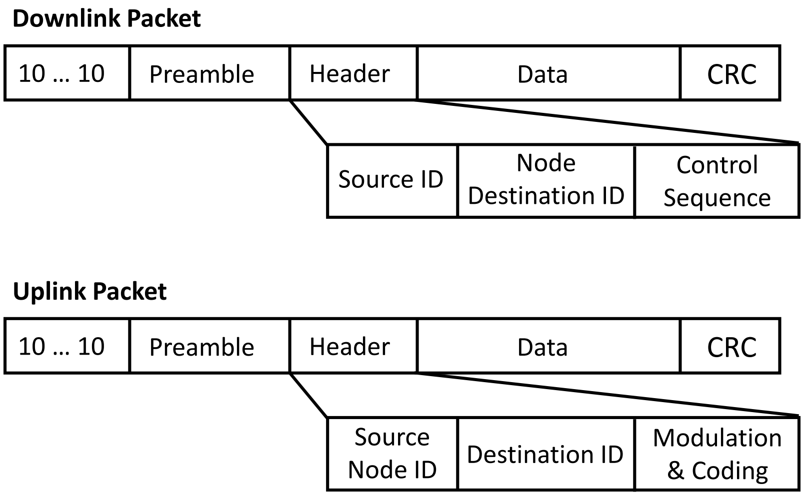

Figure 3-4: **Packet Structure.** This figure shows the different fields of the downlink and uplink packet.

3.1.4 From PHY to Higher Layers of the Networking Stack

Now that we have discussed how *PAB-QAM* works at the physical layer, we describe how we can extend it to the higher layers of the networking stack.

Packet Structure: First, we describe the packet structure of underwater backscatter networks. Fig. 3-4 shows the different components of the uplink (backscatter node to hydrophone) and downlink (projector to backscatter node) packets in *PAB-QAM*. Each packet consists of a preamble, header, payload, and CRC. The preamble contains a training sequence, which is used for synchronization and channel estimation. The training sequence is followed by a header which contains the source ID, node destination ID, and control information. On the downlink, the control sequence may contain a specific time slot or frequency channel for communication. On the uplink, the control sequence is typically the modulation and coding rate of the payload. Note that the preamble and header for the uplink packet is typically in BPSK (or FM0 modulation), while the payload can be in a higher-order modulation scheme. Each packet ends with a CRC sequence which is used to verify if the entire packet was received correctly.

Multiple Access Control (MAC): Our system incorporates a MAC protocol to scale *PAB-QAM* to multiple nodes while minimizing interference. The MAC protocol is arbitrated by the downlink projector. In its simplest form, the projector implements time division multiple access (TDMA) by leveraging the node ID field in the packet header. Specifically, at the beginning of each communication session, the projector can select the node it wishes to communicate with. If multiple backscatter nodes receive the downlink packet, only the node whose ID matches the header field would respond with its sensor data on the uplink, thus avoiding interference. Alternatively, the projector may implement other MAC protocols such as frequency division multiplexing (FDMA) (as discussed in 3.1.3). In such scenarios, the projector would use the control sequence field in its downlink packet to assign different frequency channels to different backscatter nodes. Upon decoding the downlink packets, each backscatter node can tune its oscillator to the corresponding frequency channel before starting transmissions.

3.2 Experimental Results

To assess the benefits of higher-order modulation, we evaluated the overall throughput of our proposed design and compared it to prior work in underwater backscatter in both simulations and empirical evaluation.

3.2.1 Simulation Results

First, we evaluated the performance of our system in simulation. We simulated its throughput at 521 different SNR values ranging from -15dB to 40dB. At each SNR value, we performed 50 simulation trials, each with random additive white gaussian noise (AWGN). For each SNR value, we averaged the throughput across all simulation trials. We repeated the same evaluation for *PAB-QAM* and PAB (our baseline) and at two coding rates: 1/2-rate and 3/4-rate. Across all these trials, we fixed the backscatter link frequency (oscillator) to 3 kHz.

Fig. 3-5 plots the throughput as a function of SNR for 4-QAM (in blue) and BPSK (in red). The dashed lines correspond to a 1/2-rate code while the solid lines correspond to 3/4-rate code. We make the following remarks:

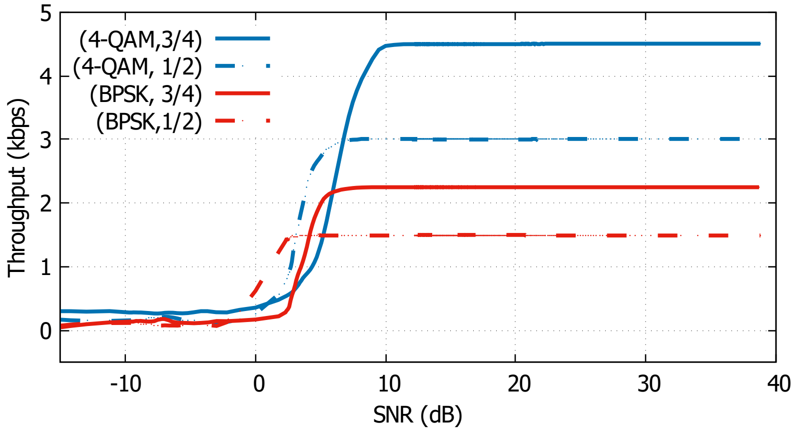

Figure 3-5: **Throughput vs SNR (Simulations).** The figure plots the throughput of QAM (blue) and BPSK (red) as a function of SNR with different code rates. For any given code rate, 4-QAM achieves twice the throughput when compared with BPSK.

- All modulation schemes start from 0 throughput at low SNRs, then increase with SNRs until they reach a plateau. At high SNRs, *PAB-QAM* outperforms the baseline (i.e., BPSK). Specifically, for all SNR values beyond 10dB, *PAB-QAM* achieves twice the throughput compared to prior underwater backscatter modulation schemes at the same rate.

- For certain low SNR values (between 0dB and 5dB), it may be more desirable to use BPSK than QAM. This is expected since, at the same power level (SNR), BPSK is more resilient to noise than QAM because it has larger minimum distance between its constellation symbols.

3.2.2 Empirical Results

Next, we evaluated our system by performing real world experiments in a controlled underwater environment (a medium-sized water tank). In these experimental trials, the projector was configured to transmit a downlink signal at a carrier frequency of 20kHz, and we measured the throughput of the backscatter node on the uplink. We performed experimental trials at three different SNR regimes: low, medium, and high (representing different SNR regions of Fig. 3-5). To achieve different SNRs, we

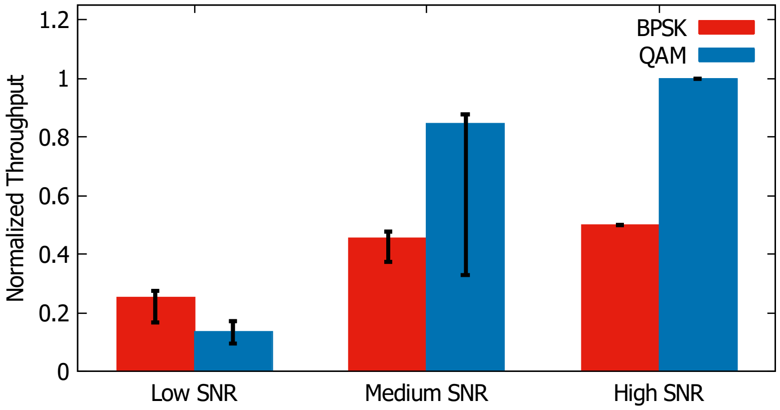

Figure 3-6: **Throughput vs SNR (Empirical).** The figure plots the normalized throughput of QAM (blue) and BPSK (red) as a function of SNR. The error bars represents the 25^{th} and 75^{th} percentile respectively.

changed the location of the backscatter node and/or the power transmitted by the downlink projector. For each configuration (of power and location), we repeated the same experimental trial for both *PAB-QAM* and our baseline (BPSK). We performed 120 trials in total.

Fig. 3-6 plots the median normalized throughput for both QAM and BPSK as a function of SNR. Specifically, we normalized the throughput by the highest achieved median throughput of *PAB-QAM* at high SNR. The error bars represent the 25^{th} and 75^{th} percentile respectively. We make the following remarks:

- At high SNR, *PAB-QAM* achieves twice the throughput of BPSK. This is expected because higher-order constellations such as QAM allow us to transmit more bits per symbol (i.e. 2 bits instead of 1 in case of BPSK) which increases our throughput by a factor of 2. Moreover, this result matches the expected performance based on our simulations in section 3.2.1.

- For low SNR values, BPSK achieves slightly higher throughput than 4-QAM. This is also expected (and matches simulations) because BPSK is more resilient to noise since it has larger minimum distance between two adjacent symbols in its normalized constellation.

These results demonstrate the importance of *PAB-QAM*'s ability to achieve higher order modulation. They also show that it may be desirable to incorporate bit rate adaptation schemes in underwater backscatter in order to enable them to gracefully scale their bit rate to channel conditions.

3.2.3 Self-Interference Cancellation Performance

Recall that *PAB-QAM* employs self-interference cancellation in order to mitigate the direct path from projector to hydrophone and enable communication with further away nodes. Specifically, by canceling self-interference, the projector can transmit signals at a higher power (without saturating the hydrophone) and receive stronger SNR from the backscatter node.

To assess the benefit of self-interference cancellation, we ran two kinds of experiments: one with and one without the filtering mechanism described in 3.1.2. In each kind of experiment, the projector transmits a downlink signal at 40 kHz, and the backscatter node is set to shift its response out-of-band by 20 kHz. The projector, hydrophone, and backscatter node were all placed within a short distance of about 1 m in order to better understand the effect of the direct bath and the benefits of self-interference cancellation. In each of the two experiments, we increased the transmitted power (voltage) fed into the power amplifier, and we measured the SNR of the received backscatter signal.

Fig. 3-7 plots the SNR as a function of the input voltage to the transducer. We make the following observations:

- Without self-interference cancellation, the SNR starts around 15 dB but experiences a sharp decline around 50 V and plateaus around -7dB beyond 80V. This is because when the input voltage to the transducer exceeds 50V, the receiver's ADC clips, preventing us from being able to decode the received packet. In such scenarios, the strong self-interference pushes the backscatter signal outside of the dynamic range of the receiver, and the backscatter signal becomes undetectable. As a result, the system can achieve reliable communication (i.e., BER

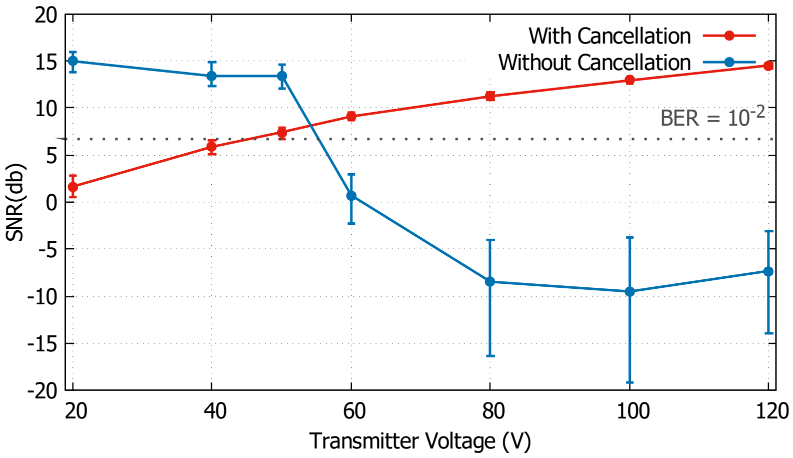

Figure 3-7: **Self-Interference Cancellation**. The figure plots the median SNR with (red) and without (blue) self-interference cancellation as a function of transmitter voltage. The error bars represent the 10^{th} and 90^{th} percentile. The dotted line represents the SNR needed to get a BER value of 10^{-2}.

below 10^{-2}) only for low values of transmitter voltages.

- With self-interference cancellation, the SNR starts around 2 dB and continuously increases with increased transmit voltage, following the theoretically predicted improvement in backscatter SNR. This is because a higher transmit voltage results in a stronger reflection, which in turn leads to high backscatter SNR. This result demonstrates that self-interference cancellation can indeed mitigate the direct path and keep the backscatter response within the dynamic range of the receiver's ADC. This result also demonstrates that *PAB-QAM*'s self-interference cancellation mechanism allows increasing the backscatter SNR and indicates the potential to communicate with further away backscatter nodes.

- Interestingly, at low transmit voltage, the configuration without self-interference cancellation outperforms that with self-interference cancellation. Specifically, the SNR around 20V is around 15 dB without cancellation but only 2 dB with cancellation. This is because the low-pass filter adds an insertion loss which degrades the receive SNR. Note that it is straightforward to overcome this loss by adding an LNA at the receiver (or by amplifying the transmit signal).

- One might wonder whether self-interference cancellation is indeed useful if the highest SNR with amplification (i.e., 120V) is similar to that with no amplification (i.e., 20V) in the absence of cancellation. Note, however, that this is due to an reduction in SNR arising from the filter's insertion loss. By accounting for this loss and adding an LNA, the received signal with cancellation would be around 13dB higher than the highest backscatter SNR without cancellation and results in increasing the communication range.

3.2.4 Scaling Performance

Finally, we would like to evaluate *PAB-QAM*'s ability to scale to a large number of nodes and enable concurrent decoding. Recall from 3.1.3 that *PAB-QAM* employs an FDMA-based mechanism that allows it to allocate different frequency channels to different backscatter nodes and decode each independently.

To evaluate this capability, we ran a controlled experiment in an indoor pool whose dimensions are 3m×4m×1.5m. In each experimental trial, we varied the number of backscattering nodes from 1 to 10. Each node was assigned a different frequency channel as per the protocol described in 3.1.3. The hydrophone receives the concurrent responses, applies appropriate software-based filters to isolate them, and decodes each of them separately. We repeated each experimental trial at least 10 times for each number of nodes. For any given number of nodes, we varied the location of the nodes across different trials to average out the effect of the frequency selective channel.

Fig. 3-8 plots the median BER and SNR as a function of the number of nodes. We make the following observations:

- The SNR starts at around 18 dB for one node and gradually drops as we add more *PAB-QAM* nodes to the environment. This gradual degradation is due to the residual interference caused by neighboring nodes. Nonetheless, even with 10 nodes, *PAB-QAM* can achieve a median BER of 10^{-1}. This result shows that *PAB-QAM* can indeed scale up to 10 concurrent nodes despite their close proximity (and interference), thus demonstrating a 5× improvement over the state-of-the-art baseline (PAB) [37].

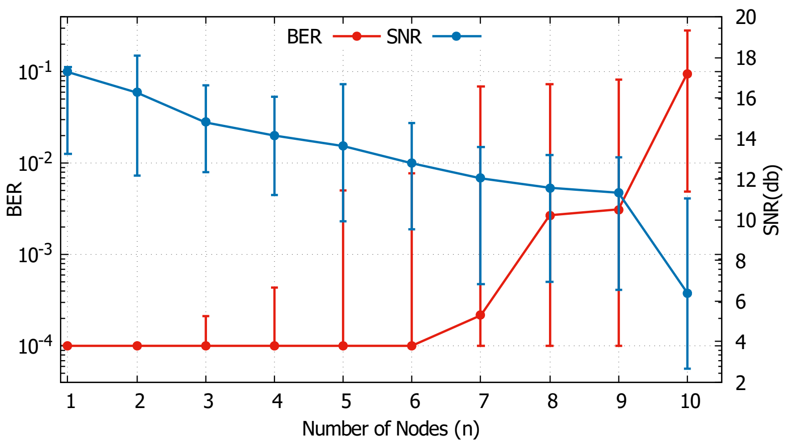

Figure 3-8: **BER and SNR vs Concurrent Transmissions.** The figure plots the median BER and SNR as a function of the number of concurrent transmissions. The error bars represent the 25^{th} and 75^{th} percentile.

- It is also interesting to note here that the BER doesn't change much as we add more nodes. This is because we operate these nodes at a relatively low bitrate of 100 kbps.

- Finally, it is worth noting that the above performance is achieved even though all the nodes are in close proximity. It is possible to scale the system to more concurrent transmissions by exploiting spatial reuse.

3.3 Conclusion

This chapter presents *PAB-QAM*, the first underwater backscatter design capable of achieving higher-order modulation. This design marks an important step forward in ultra-low-power underwater networking by demonstrating the potential of doubling the throughput of backscatter communication nodes through QAM modulation. As underwater backscatter continues to evolve, we hope that it will pave way for ultra-low-power and low-cost subsea IoT for ocean monitoring, scientific exploration, and marine life sensing.

Chapter 4

Underwater Backscatter Localization

There is significant interest in low-power and distributed underwater localization systems for environmental, industrial, and defense applications [53, 87, 42, 18]. Climatologists and oceanographers are interested in deploying such systems to obtain location-tagged ocean measurements for constructing subsea heatmaps [53], understanding ocean processes [70], and developing accurate weather and climate prediction models [64]. Marine biologists are interested in such systems for tracking schools of fish and studying their behavior and migration patterns [42, 14]. Accurate and low-power localization is also a key enabler for various underwater robotic tasks including navigation, tagging, and object manipulation [51, 18].

Unfortunately, prior designs for underwater localization remain far from the vision of a low-power, low-cost, and scalable architecture. Since standard GPS signals do not work in water,[1] most existing underwater positioning systems rely on acoustic signals [12, 13, 46]. These systems typically require their nodes to repeatedly transmit acoustic beacons (which are used by a remote receiver for triangulation). Such repetitive transmissions can quickly drain a sensor's battery, thus requiring a frequent and expensive process of battery replacement [21, 20]. To avoid this problem, existing localization systems either heavily duty-cycle their transmissions [98, 23] or tether the localization beacons to a large power source on a ship or submarine [100, 43]. Unfortunately, such workarounds prevent these systems from accurately tracking fast-moving

---

[1] GPS relies on RF signals which decay exponentially underwater [61, 85].

objects (like fish or drones) and/or scaling to large areas of the ocean.

We introduce Underwater Backscatter Localization (UBL), an ultra-low power and scalable system for underwater positioning. UBL builds on our recent work in underwater networking (discussed in section 1.1), which has demonstrated the potential to communicate at near-zero power via acoustic backscatter [37, 25]. By bringing localization to underwater backscatter, UBL would enable us to build a long-lasting, scalable, battery-free underwater GPS.

Before explaining how UBL works, let us understand why it cannot easily adopt traditional underwater localization techniques. State-of-the-art underwater localization systems rely on computing the time-of-arrival (ToA) between two nodes [99, 11].[2] In these systems, a transceiver sends out an acoustic pulse, and waits for a response from the transponder beacon. The time difference between the initial pulse and the reply is used to determine the separation between the two nodes (by multiplying it with the sound speed in water). Unfortunately, this ToA estimation technique does not work for battery-free nodes. These nodes require an additional wake-up time to harvest energy from acoustic signals before they can start backscattering. This wake-up time cannot be determined a priori and varies with location and environment. As a result, it adds an unknown offset to the time difference between the transmitted and received pulse, preventing us from accurately estimating the ToA and using it for localization.[3]

To overcome this challenge, UBL adopts a time-frequency approach to estimate the ToA. Specifically, instead of estimating the ToA entirely in the time domain, it also collects frequency domain features by performing frequency hopping. Since time and frequency are inversely proportional, hopping over a wide bandwidth would enable UBL to estimate the ToA with high-resolution [48, 93, 6]. Transforming this idea into a practical underwater localization system still requires dealing with multiple confounding factors:

---

[2] In contrast, using angle-of-arrival typically requires expensive and bulky antenna arrays and results in poorer accuracy than ToA [34].

[3] Note that an approach that introduces pauses between a reader's transmissions is also undesirable since the backscatter node requires continuous signals to stay awake [37].

- **Multi-path:** When acoustic signals are transmitted underwater, they repeatedly bounce back and forth between the seabed and the water surface before arriving at a receiver. Such dense multipath reflections make it difficult to isolate the direct path to a backscatter node for ToA estimation.

- **Delay Spread:** The slow speed of sound propagation spreads out the above multipath reflections over time, resulting in a large delay spread. This delay spread causes different backscatter bits – *even from the same node* – to interfere with each other. As we show in 4.1, such *inter-backscatter interference* is unique to acoustic backscatter and exacerbates the ToA estimation problem.[4]

- **Mobility:** Performing accurate localization becomes more challenging for mobile nodes (fish, drones). This is because mobility distorts the estimated frequency features (due to Doppler shift [91]) and because frequency hopping increases the latency of localization, during which a mobile node may have moved to a new location.

Addressing the above challenges simultaneously requires satisfying competing design requirements. For example, reducing the backscatter bitrate would increase the separation between symbols in a packet (thus mitigating inter-symbol interference), but it also slows down the channel estimation process, making it difficult to localize fast-moving objects. In a similar vein, dealing with multipath and mobility results in conflicting design constraints (for the bitrate and hopping sequence). We argue that designing a robust underwater backscatter localization system requires context-aware algorithms that can adapt their bitrates and hopping sequence to their operating domains. In 4.2.2, we describe the fundamental constraints arising from these different challenges and how our design of UBL aims to strategically adapt to its surrounding environment.

We implemented a proof-of-concept prototype for UBL and tested in a river. Our prototype consists of a mechanically fabricated backscatter node and a custom-made

---

[4] In contrast, in RF backscatter localization, due to the high propagation speed, all multipath reflections arrive in the same backscatter state [48, 47].

PCB with a micro-controller and backscatter logic. Our experimental evaluation across three different locations demonstrates the feasibility of achieving centimeter-level accuracy using UBL. Our empirical evaluation is complemented with simulations that demonstrate how the system can adapt to different speeds and multipath environments.

This chapter presents the first design and demonstration of underwater backscatter localization. Our design can deal with unique challenges that arise from the interaction between underwater multipath and acoustic backscatter, and it can can adapt to various underwater conditions (depth, mobility). The chapter also contributes a proof-of-concept implementation and evaluation of underwater backscatter localization.

4.1 The (New) Problem

Before describing UBL's design, it is helpful to understand why underwater backscatter localization poses new challenges that are different from prior work in RF backscatter localization (e.g., RFID localization [32, 94, 48, 47]). To answer this question, in this section, we provide background on underwater acoustic channels, then explain how these channels pose interesting new challenges for backscatter localization.

Underwater Acoustic Channel. The underwater channel is a confined environment bounded with air on one side and sediment on the other side as shown in Fig. 4-1(a). When acoustic signals are transmitted underwater, they can travel over very long distances (tens to hundreds of kilometers [83]) due to two factors: (1) the small attenuation of sound in water; and (2) the fact that sound entirely reflects off the air/water and water/sediment boundaries because of the large impedance mismatch between these media. Thus, an acoustic signal travels on various paths from a transmitter to a receiver, most of which involve multiple reflections off the air and water boundaries. As a result, the receiver obtains multiple copies of the signal, which we refer to as underwater multi-path.

Impact of Multipath on Underwater Backscatter. To understand the impact

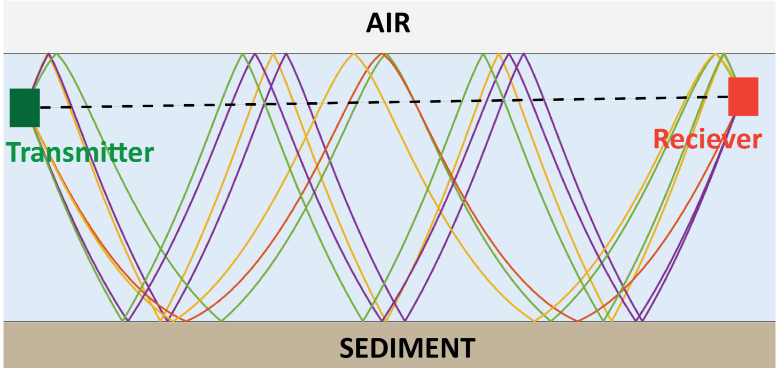

(a) Multipath in underwater channel

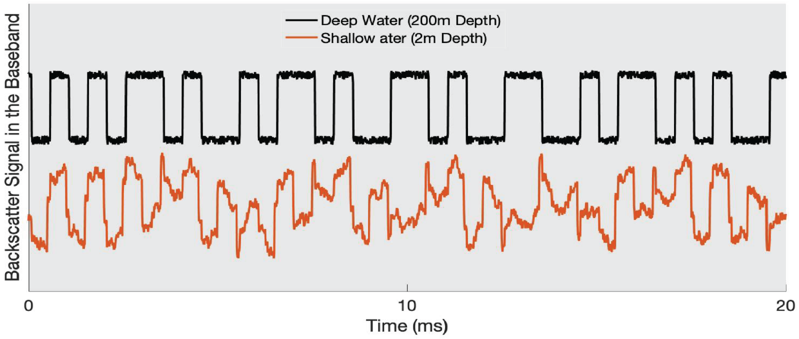

(b) Backscatter signal in deep and shallow water

Figure 4-1: **Multipath and Underwater Backscatter.** (a) shows how sound propagates underwater, repeatedly reflecting off the surface and seabed. (b) shows a received backscatter packet in deep (low multipath) and shallow (dense multipath) water.

of the underwater channel on acoustic backscatter, we simulated backscatter communication in two different environments corresponding to deep water (depth$>$200m) and shallow water (depth$<$10m). In both of these environments, the backscatter node and the receiver are separated by the same distance (4 m).

Fig. 4-1(b) shows the received backscatter signal in each of these two scenarios. In deep water (black plot of Fig. 4-1(b)), the received signal shows clear transitions between reflective and non-reflective states. Recall that these states encode bits of 0's and 1's that are used to communicate data. In contrast, in shallow water (orange plot of Fig. 4-1(b)), the backscatter response is highly distorted and the transitions are significantly obscured.[5] It is worth noting that the difference between these two scenarios is not due to difference in the signal-to-noise ratio, since the

---

[5] We observed similar behavior when empirically testing our system in a real river.

distance separation between the backscatter node and the receiver is the same in both cases.

Instead, the difference between the two different scenarios arises from the multipath reflections mentioned earlier. Specifically, in deep water, the direct path is much stronger than the reflected paths because it travels a smaller distance and experiences less attenuation (4m vs 200m). In contrast, in shallow water, the direct path and reflected paths have similar lengths and thus have similar amplitudes; this leads to interference between subsequent symbols (i.e., between different backscatter states). Unless this distortion is accounted for, it will be difficult to estimate the wireless channel in the frequency domain (which UBL needs for localization).

This inter-symbol interference (ISI) is unique to underwater backscatter and does not exhibit in RF backscatter.[6] The difference between RF and acoustic backscatter arises from significant disparity between the speed of RF in air (3×10^8m/s) and that of sound in water ($1,500$m/s). In RF backscatter, the nearest reflector that may cause ISI is more than 3 km away (i.e., significantly attenuated), while in acoustic backscatter even a reflector that is 1.5 m away can cause ISI. This difference motivates a new principled approach for underwater localization that differs from standard RF backscatter localization techniques.

4.2 UBL

UBL is an accurate underwater localization system for ultra-low-power and battery-free nodes. The system can achieve centimeter-scale positioning even in multipath-rich underwater environments. To locate a backscatter sensor, UBL performs the following steps:

- A UBL reader sends a query searching for backscatter nodes in the environment.
- When a node replies, UBL sends a downlink commands specifying the backscatter bitrate.

---

[6] Note that ISI is known in wireless communication, and standard protocols like OFDM can be used to address it [91]. However, OFDM is too complex for battery-free underwater nodes.

- As the node replies, the reader performs frequency hopping to estimate the node's channel over a wide frequency.

- Finally, UBL uses the acquired bandwidth to estimate the time-of-arrival (ToA) to the backscatter node and uses the ToA for localization.

Since prior work has demonstrated the ability to query and command an underwater backscatter node [37], in this section, we focus on how UBL uses frequency hopping to estimate the ToA (4.2.1) and how it selects the bitrate and hopping sequence (4.2.2).

4.2.1 Backscatter ToA Estimation

UBL performs localization by estimating the time-of-arrival (ToA) of a backscatter node's signal. ToA estimation is particularly useful in multipath-rich environments. Specifically, in the presence of multiple reflections, a receiver can determine the direct path as the one having smallest ToA (since it travels along the shortest path).

The main challenge in backscatter ToA estimation arises from the random wake-up time of battery-free nodes. Specifically, recall from **??** that battery-free nodes need to harvest energy in order to power up before they can start backscattering. Moreover, this wake-up delay varies with distance and environment; thus, it cannot be determined a priori.

To overcome this challenge, instead of estimating ToA directly in the time domain, UBL does so in the frequency domain. Since time and frequency are inversely related, a wide bandwidth can be used to separate different paths and identify the direct path as the one that arrives earliest. Specifically, the resolution to determine the direct path is given by the following equation:

$$\text{resolution} = \frac{speed}{Bandwidth}$$

Thus, for a bandwidth of 10kHz, UBL can localize the node to within 10 cm. In the rest of this section, we describe the different steps of UBL's ToA estimation approach.[7]

Stage 1: Wideband Channel Estimation. *UBL* estimates the backscatter channel over a wide bandwidth by performing frequency hopping. Specifically, it transmits a downlink signal at a frequency k and obtains the backscatter response. Once the receiver obtains the response y_t, it performs the following two steps:

1. First, it cross-correlates the received signal with the known backscatter packet preamble to determine the beginning of a packet, denoted τ^*, using the following equation:
$$\tau^* =_\tau \sum_{t \in T} y_{t+\tau} p_t \tag{4.1}$$
where p_t is the known preamble and T is the length of the preamble. By identifying the beginning of the packet, this correlation can be used to eliminate the wake-up lag.

2. Subsequently, *UBL* estimates the backscatter channel H_k using the packet's preamble. This can be done using standard channel estimation as per the following equation:
$$H_k = \frac{1}{T} \sum_{t \in T} y_{t+\tau^*} p_t \tag{4.2}$$
where $y_{t+\tau^*}$ corresponds to the received signal shifted to the beginning of a packet.

UBL repeats the above procedure for different frequencies (each time hopping to a different frequency and computing the corresponding channel) until is has obtained the channels for across a wide bandwidth $[H_1, H_2, \ldots H_N]$.

Stage 2: Obtaining the Time-Domain Channel. After concatenating the different frequencies, *UBL* performs an inverse Fourier transform (IFFT) on the channels. This allows it to obtain an expression of the channel in the time domain. Importantly, this time-domain representation is independent of the random wake-up time since it is obtained entirely from the channel estimates.

---

[7] We note that this technique is similar to that employed in [48] and will be adapted in 4.2.2 to underwater backscatter.

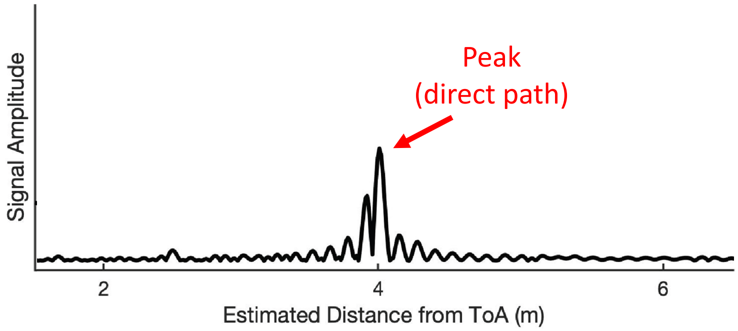

(a) Time-domain channel from frequency hopping

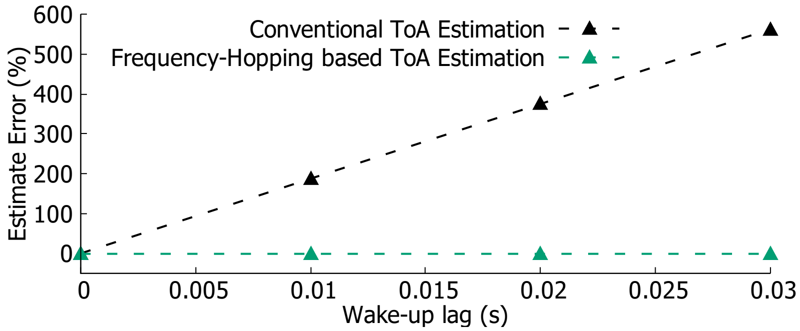

(b) Wake-up lag effect of localization

Figure 4-2: **Range estimation via frequency hopping.** (a) shows how UBL can isolate the direct path in the time domain using the frequency-hopping localization method and (b) shows the effect of the wake-up lag on conventional ToA based localization schemes.

One might wonder whether eliminating the wake-up lag would also eliminate the impact of the round-trip delay on the channel estimates. In practice, this does not happen because UBL estimates the channel in the frequency domain. To see why this is true, consider a simple setup with a single line-of-sight path from the backscatter node to the receiver. Here, the baseband received signal y_t can be expressed as:

$$y_t = e^{j2\pi f_c \tau_r} b(t - \tau_r - \tau_w) \tag{4.3}$$

where τ_r and τ_w correspond to the round-trip delay and wakeup lag respectively. By shifting the received signal in the time domain (by $\tau_r + \tau_w$), UBL eliminates the delays in the time domain but not the impact of the round-trip delay on the frequency-domain channel. Hence, it is able to recover this delay upon performing an IFFT.

To demonstrate this idea in practice, we simulated the localization problem where a UBL reader and a backscatter node were separated by 4 m in a deep underwater environment. Fig.4-2(a) plots the channel amplitude as a function of distance after performing the above procedures. The plot demonstrates a clear peak amplitude in the channel around 4 m, which is aligned with the actual distance of backscatter node. Note that because the simulated environment corresponds to a deep sea where multipath is distance with respect to the line of sight, the plot does not show other peaks from nearby reflections in the environment.

Next, to investigate the effect of wake-up lag on UBL's ToA estimation approach, we simulated localization after introducing different time delays (betweem 0-30 ms), and compared the outcome of UBL with that of conventional time-domain methods for localization. Fig. 4-2(b) plots the percentage error in distance estimation as a function of the wake-up lag for both schemes. The figure shows that while UBL's error remains small irrespective of the wake-up lag, conventional (time-based) ToA estimation systems are significantly affected by this delay and suffer from a large margin of error. This demonstrates that UBL's ToA estimation approach is robust to the random wake-up lags of battery-free backscatter sensors. In 5.2, we empirically verify this result as well.

4.2.2 Adaptive Backscatter Localization

So far, we have described how UBL can estimate the ToA robustly despite a random wake-up lag. However, the above description focused on deep sea environments with little multipath. In this section, we describe how UBL's design can be extended to deal with extreme multipath and mobility in underwater environments.

Dealing with Extreme Multipath

To understand the impact of extreme multipath, we repeated the same simulation of as our earlier experiment but this time in shallow water (depth of 4 m) rather than in deep water. Fig. 4-3(a) plots the signal amplitude as a function distance. Unlike the previous experiment, we are unable to see a sharp peak around 4 m, making it difficult

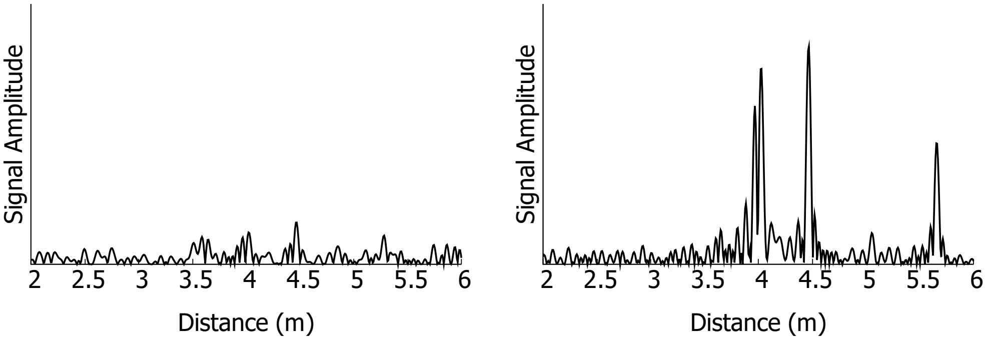

(a) Shallow water, Bit-rate: 2 kbits/s (b) Shallow water, Bit-rate: 100 bits/s

Figure 4-3: **ToA Estimation in shallow water.** This figure shows how multipath affects the localization ability for UBL. (a) shows that at a high bit-rate of 2 kbits/s, UBL fails to localize the object. (b) shows that operating at a lower bit-rate of 100 bits/s in multipath rich environments yields better performance.

to robustly estimate the time-of-arrival in extreme multipath environments. This is because inter-symbol interference (ISI) makes it difficult to obtain accurate channel estimates. This challenge can be seen visually in the orange plot of Fig. 4-1(b).

To mitigate the impact ISI on ToA estimation, UBL can command the backscatter node to lower its bitrate. Intuitively, doing so increases the separation between any two backscatter symbols, thus reducing the interference between the reflection of the former with the direct path of the latter. From a communication perspective, reducing the bitrate results in a more narrowband channel, which increases robustness to frequency selectivity [91].

To test this idea, we repeated the same simulation, but this time at a bitrate of 100 bps instead of 2 kbps. Fig. 4-3(b) plots the resulting output for this experiment. The figure shows a much sharper peak around 4 m than that obtained when the same experiment was performed at a higher bitrate. The figure also shows a second peak around 4.5 m, which corresponds to the first (primary) multipath reflection off the surface and sediment. Note that both experiments used the same bandwidth and are simulated at the same distance (i.e., the latter did not benefit from more resolution or higher SNR). Rather the difference is localization robustness arises from the lower bitrate. Formally, we can prove the following lemma. To ensure the inter-symbol interference from any single path is no larger than k dB, the backscatter bitrate must be less than $\frac{c}{(10^{0.05k}-1)\ 4r}$ To prove this lemma, let us denote the largest delay caused by the reflected path as T_r and the delay caused by the direct path as T_l. We further assume that the power of the reflected path is attenuated by k dB compared to the power of the direct path. This gives us the following relation:[8]

$$T_r = 10^{0.05k}\, T_l$$

Since T_l corresponds to the round trip distance, it can be written as a function of the separation r and the speed of sound c as $T_l = 2r/c$. To ensure that any strong symbol reflection arrives before the next symbol is received, the symbol period T_s (or bit period) should be greater than twice the largest delay T_d, which gives an upper bound for bitrate R:

$$R \leq \frac{c}{(10^{0.05k} - 1)\; 4r}$$

Dealing with Mobility

Next, we are interested in extending UBL to deal with mobility of underwater backscatter (e.g., in tracking fish, AUVs). To understand the impact of mobility on localization, we simulated the localization process in deep water for a node moving at a speed of 0.3m/s. Fig. 4-4 plots the signal amplitude as a function of distance. Unlike the earlier experiment in deep water (i.e., Fig. 4-2(b)), we are unable to see a sharp peak around 4 m, making it difficult to robustly estimate the time-of-arrival in the presence of mobility. This is because mobility causes a change in the channel estimates over time. As a result, the resulting channel estimates $[H_1(t_1), H_2(t_2), \ldots H_N(t_N)]$ cannot be coherently combined to obtain an accurate location estimate.

To mitigate the impact of mobility on ToA estimation, UBL needs to reduce the overall time required for localization. This can be done by commanding the backscatter node to increase its bitrate and the reducing the number of frequencies in the frequency hopping sequence. We can formalize the mobility constraint through the following lemma.

To localize a mobile node moving with the speed of v, backscatter and frequency

---

[8] This comes from standard spherical loss $P \propto 20 \log_{10}(1/T)$.

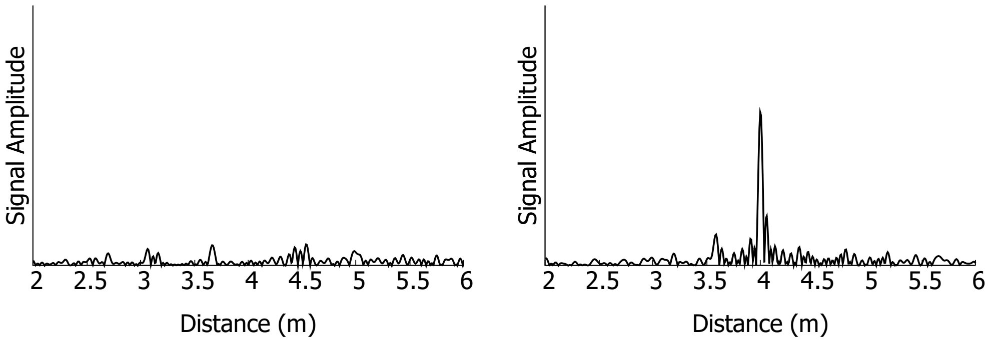

(a) Deep water, Bit-rate: 1 kbits/s (b) Deep water, Bit-rate: 10 kbits/s

Figure 4-4: **ToA Estimation in deep water with mobility.** This figure shows how UBL can adapt to mobility in deep water environments . (a) shows that at a bit-rate of 1 kbits, UBL is unable to localize the object while the object is moving with a speed of 0.3 m/s, so for better accuracy, it is desirable to operate at a higher bit-rates to deal with mobility as shown in (b).

hopping properties should satisfy the condition of : $\frac{R}{N_f L_p} \geq \frac{2vB}{c}$ where R is the backscatter bitrate, N_f is the number of frequency in frequency hopping, L_p is the bit length of the preamble, v is the relative speed of the mobile node , B is the bandwidth and c denotes the speed of sound.

Lemma 4.2.2 is derived considering the fact that to localize a mobile node with the resolution of x, frequency hopping process must be accomplished before the node get displaced more than x. Assuming the backscatter bitrate of R and preamble bit length of L_p, the minimum required time to estimate the channel for each frequency is $\frac{L_p}{R}$ and the minimum required time for frequency hopping duration is $\frac{N_f L_p}{R}$. To localize the node, The duration of frequency hopping should be less than the time it takes for the node move more than x. This gives us the following relation:

$$\frac{N_f L_p}{R} \leq \frac{x}{v} \tag{4.4}$$

Additionally, the resolution x is function of bandwidth and may be written as $x = \frac{c}{2B}$, completing the lemma.

To test the relationship, we repeat the same experiment as above, but this time with a backscatter bitrate of 10kbps. (Here, $B = 10KHz$, $L_p = 20$, $N_f = 100$, this requiring a minimum bitrate of 8kbps according to the lemma). Fig. 4-3(b) plots the resulting output for this experiment. The figure shows a much sharper peak around 4 m than that obtained when the same experiment was performed at a lower bitrate, demonstrating that UBL's adaptation enables it to accurately localize despite mobility.

We make few additional remarks about how UBL chooses its backscatter bitrate and hopping sequence:

- Lowering the bandwidth (B), decreases the resolution of localization. Therefore, UBL always tries to exploit the full bandwidth allowed by the backscatter node's mechanical characteristics.
- Decreasing the bit length of the preamble (L_p), leads to the lower SNR. UBL utilize the preamble length of at least 20 bit to ensure the channel is estimated reliably.
- The longest distance that can be localized is determined by the length of the IFFT. Therefore, decreasing the N_f, limits the range of the localization

4.3 Feasibility Study

In this section, we explain how we implemented and validated the feasibility of UBL for underwater localization. Similar to the system design discussed in section 1.1 for underwater backscatter [37, 25, 7], UBL's implementation leverages a projector to transmit an acoustic signal on the downlink, a backscatter node that decodes the downlink signal and transmits a backscatter packet on the uplink, and a hydrophone (Omnidirectional Reson TC 4014 hydrophone [86]) that receives and decodes the backscatter packets. The projector and backscatter node were fabricated in house from piezoceramic cylinders, following the procedure elaborated in our prior work [37].

In our experiment, the projector was programmed to hop its carrier frequency from 7.5 kHz to 15 kHz (at 75 Hz intervals, each for 6 seconds). This range of frequency is selected based on bandwidth of the backscatter node [37] and, subsequently, the

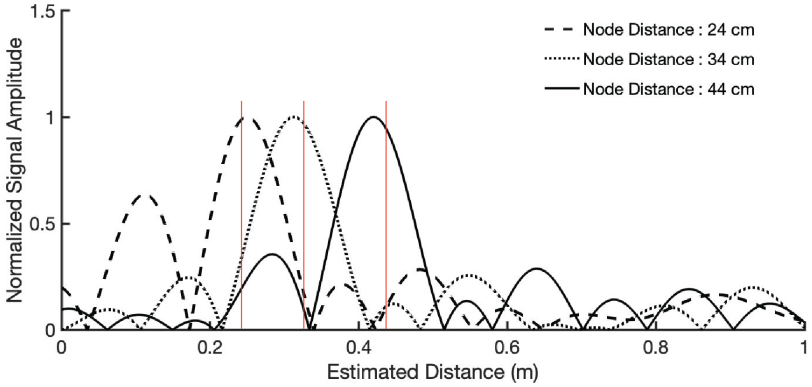

Figure 4-5: **Preliminary Results for UBL.** The system was tested for three different ranges i.e. 24 cm, 34 cm and 44 cm respectively.

expected resolution is 10cm. To account for the effect of multipath in such shallow environment, the backscatter bitrate of 100 bit/s is adopted. Notably, since the node was relatively stationary in the water, the bitrate of 100 bit/s was sufficient to estimate the channel. The received signal recorded by the hydrophone was then processed by first estimating the channel at each of the frequencies, and subsequently, the time-domain channel is computed to estimate ToA per our discussion in 4.2.1.

The outputs of UBL for three different node locations are shown in Fig.4-5. The x-axis corresponds to distance and the y-axis represents the normalized time-domain channel amplitude. In this result, the ground truth is marked using red vertical lines and the peak amplitude for each distance is within 10 cm from the ground truth. Note that, due to the limited bandwidth , our resolution was 10 cm and to achieve finer precision, UBL can emulate a wider bandwidth.

4.4 Conclusion

To conclude, this chapter takes a first step toward ultra-low-power and batteryless underwater localization. As the UBL evolves, we envision it will enable various applications in environmental monitoring, marine life understanding, and underwater exploration.

Chapter 5

Underwater Battery-Free Machine Learning

In recent years, advances in battery-free sensing have enabled a range of novel applications including localization and environmental sensing. These advances have enabled long-term use of wireless sensors without the need for power supplies or batteries, hence enabling a transition towards a pollution-free sensing infrastructure for a variety of environmental sensing and monitoring applications such as detecting pollution and monitoring biodiversity. These applications are not limited to smart cities and urban environments, but also extend to climate change monitoring and environmental sensing in uninhabited areas such as forests, mines, remote areas, and space [5, 97].

The vast majority of existing battery-free technologies have been designed for land applications and very few have been engineered for the ocean. Only recently have researchers started looking into battery-free sensors for underwater environments. However, existing underwater battery-free technologies stop at energy harvesting for powering up sensors and backscatter for communicating raw sensory data [38, 27, 29]. Meanwhile, research into energy-efficient edge-based machine-learning (ML) models have enabled lightweight models to be used for on-device ML inference for a variety of applications. There have been a number of advances in model reduction, compression, and layer-wise quantization for various objectives such as privacy, energy, and efficiency [50, 8, 59]. A fundamental challenge that remains is to perform time-series data analysis and ML on underwater battery-free sensors. One benefit of such capabilities is that they would enable us to limit data transmission, hence increasing the energy budget on underwater devices and their operation longevity. This is particularly important for usages of underwater battery-free sensors in ambient and remote monitoring applications for the ocean, where communication bandwidth is narrow and the sources for energy harvesting are limited [79, 66].

In this chapter, we propose and present a new challenge, investigating the feasibility of battery-free ML in underwater environments, where extremely lightweight and task-specific deep neural network (DNN) models are executed on dedicated and highly-efficient underwater sensor nodes. We further investigate the tradeoffs between lightweight on-device analytics requirements for an exemplar case, which is marine mammal recognition in the ocean. This use case has many applications including monitoring biodiversity, understanding marine animal migration patterns, and even supporting the discovery of new species [55].

Our investigation focuses on two critical aspects of underwater battery-free ML, which are (1) the capability (i.e, model accuracy) of lightweight DNN models on the tasks to be conducted in underwater environments, and (2) the feasibility of hosting the DNN models on low-power underwater computational devices such as microcontrollers. Specifically on the task of marine mammal recognition, we design an end-to-end pipeline (i.e, from recording sounds to classification results on an underwater battery-free device) and evaluate our prototype's accuracy and power consumption. Our preliminary results indicate that, for a task of classifying four marine mammals, lightweight DNN models can achieve decent accuracy. With the help of existing energy harvesting technologies and customized circuit design, it is possible to run lightweight ML on underwater battery-free devices. These results have important implications for both the mobile and ML communities, and pose exciting new opportunities for ubiquitous underwater battery-free ML.

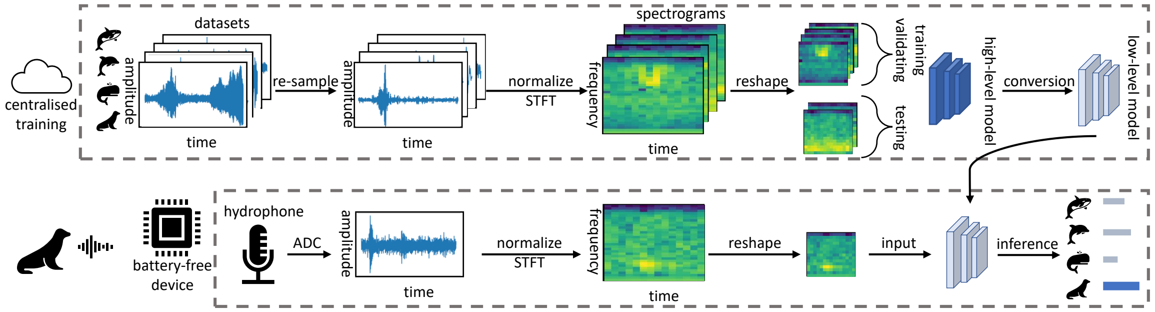

Figure 5-1: Overview of offline training and on-board inference. Initially, a high-level model is trained, validated, and evaluated offline on pre-collected audio clips from different marine mammals. A model with acceptable offline accuracy is converted to a low-level model that can be supported by standard C libraries on a target battery-free device. The converted model is then deployed in an end-to-end (, from sounds received by a hydrophone to final classifications) marine mammal recognition pipeline on the target device.

5.1 Underwater Battery-free Architecture

We design a wireless, battery-free underwater system to perform sensing and inference on edge nodes and communicate with a receiver. To this end, we prototype a system that can perform marine mammal recognition in the ocean. Fig. 1-5 shows the system design for our prototype. The system harvests energy from underwater sound to power up its processor and on-board sensor that captures animal sounds (5.1.1), performs on-board inference (5.1.2), and transmits the result on the uplink via backscatter (5.1.3).

5.1.1 Underwater Energy Harvesting

To enable battery-free operation, our sensor needs to harvest energy from ambient underwater sources such as sound, waves, or temperature gradients. As discussed in 1.1, sound provides sufficient energy to power up low-power microcontrollers [38]; hence, we design our node to harvest energy from underwater sound. It is worth noting that our node architecture is general and can be adopted to other underwater ambient energy sources.

To harvest energy from sound, our node employs piezoelectric materials, which convert mechanical energy (i.e. sound) to electrical energy. Since the harvested signal is typically an alternating current (sound is a wave), our sensor node rectifies it to a DC voltage using a multi-stage rectifier. Once the node harvests sufficient energy, it can power up an on-board microcontroller. The microcontroller has an integrated ADC (analog to digital converter) which can be interfaced to a hydrophone (as shown in Fig. 1-5). The ADC samples received sound and stores it in memory.

5.1.2 Battery-free Underwater Inference

The next stage in our design is to recognize animals from the recordings. To do this, one option is to program the node to transmit its recordings to a remote receiver that has a dedicated power source (e.g., an underwater drone or a coastal base station) and perform inference there. However, such an approach is undesirable for multiple reasons. First, since the throughput of underwater acoustic channels is limited, transmitting the full recording would require our node to remain powered up for an extended period of time, which will drain its harvested energy. Second, the limited throughput would incur a large delay in data transmission, which is detrimental in time-sensitive scenarios (e.g., detection of endangered species). Furthermore, if a desired receiver is not within the communication range of our node, the node would need to store the recording in its constrained memory, which limits its ability to process new recordings.

To overcome these challenges, instead of transmitting the recordings for inference at a powered base station, our node performs on-board inference, as we describe in this section. This results in a lower-power, time-and-resource efficient, and more scalable system design (as we show in 5.2).

Offline Training:

In order to have an accurate DNN model that can recognize marine mammals from their sounds, we need to use pre-collected audio clips (obtained from publicly available databases) to train the model until it achieves acceptable accuracy. Given the limited memory and computational resources of battery-free devices, it is difficult to perform the training process on them. Thus, we first train the high-level DNN models in a centralized manner (, in the cloud) with offline training and validating data.

To ensure that our model would perform well on the target underwater battery-free device, we introduce a number of pre-processing steps (, re-sampling, normalization, and reshaping) as shown in Fig. 5-1. These steps are necessary because the audio files (from databases) were collected using different equipment than our battery-free nodes (, different ADC sampling rates and resolution). To mitigate the impact of these differences, the raw data from pre-collected clips is re-sampled with the sample rate of the nodes. The re-sampled data is then normalized and converted into a spectrogram through Short-Time Fourier transform (STFT).

To make the models work on a battery-free device, the high-level models need to be converted into a low-level model that satisfies the constraints of the architecture and memory of the target device. Our model conversion process uses techniques such as model compression and quantization to make the resulting low-level model work on the target device independently without support from run-time libraries.

On-board Edge Inference:

After the model is trained and converted, we can load the model parameters to the memory of the target device and then use it for on-board inference. We use the same pre-processing pipeline (normalization, STFT etc.) as for offline training which was discussed in 5.1.2. One challenge with ultra-low-power microcontrollers is that, unlike standard audio cards, their ADCs cannot sample audio with both positive and negative amplitude. The lack of negative amplitude introduces another discrepancy between the pre-collected audio files and sound recordings collected using the microcontroller. To overcome this challenge, our design adds a clamping circuit to hardware design in order to passively add a DC offset to our input signal before we feed it to the on-board ADC. This DC offset is subsequently filtered in software by applying a high-pass filter; this ensures that the DC offset does not bias the STFT output or negatively impact the model performance.

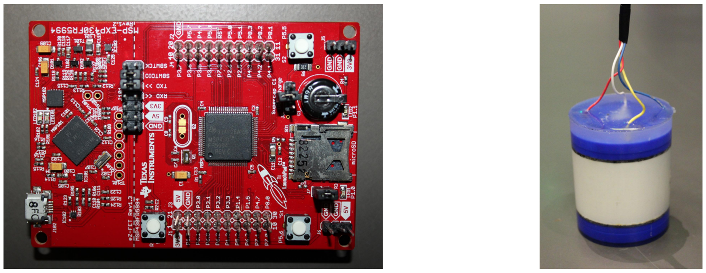
(a) MSP430FR5994 MCU development board (b) Potted Transducer

Figure 5-2: **Prototype Hardware** (a) shows the processor that our design uses for on-board inference. (b) shows one of our potted transducers which can be used for energy harvesting and communication.

5.1.3 From Inference to Communication

The last step in our design is to communicate the inference result to a remote receiver. To do this, we leverage underwater backscatter (discussed in section 1.1) to communicate on the uplink, which allows us to reuse the same piezoelectric material for both harvesting (5.1.1) and communication.

5.2 Feasibility Study

We evaluated our prototype in terms of both its model accuracy and its power requirements.

5.2.1 Prototype Hardware

Backscatter node: Each backscatter node consists of a piezoelectric transducer and a hardware controller as shown in Fig. 5-2. To build our transducers, we used piezo-ceramic cylinders with a nominal resonance frequency of 17 KHz [80]. Similar to fabrication techniques discussed in Chapter 2, the nodes were potted, encapsulated for insulation and matching to water, and housed in 3D-printed mold and end-caps [67]. **Hardware controller:** The node hardware is used for energy harvesting, processing/inference, and interfacing with a sensor (e.g., a hydrophone). Our harvesting architecture consists of a standard multi-stage rectifier followed by a super-capacitor, and low dropout voltage regulator which drives the digital logic unit. We implemented the logic on a MSP430-FR5994 microcontroller [89]. The microcontroller has a 12bit-ADC, which samples the (sound) signal at the rate of 330 samples/sec and stores a window of 512 samples in its SRAM. The microcontroller logic, including the code for downlink and uplink decoding as well as the inference model, is all stored in the FRAM. The classifier output is transmitted on the uplink by controlling the backscatter logic which modulates the load impedance of the backscatter node to change between reflective and non-reflective state.

Receiver: In addition to a backscatter node for energy harvesting and communication, our design uses a hydrophone (Omnidirectional Reson TC4014) that receives and decodes the FM0 encoded backscatter packets [9] .

5.2.2 Offline Pipeline

For the evaluation, we used the Watkins Marine Mammal Sounds Database [68] and chose the sounds from eight mammals including Atlantic spotted dolphins, bearded seals, Beluga white whales, bottlenose dolphins, bowhead whales, harp seals, narwhals, and walruses. Overall we used 364 sound files and in each trial of our offline evaluation, we shuffled and split them as 80%, 10%, and 10% for training, validating, and testing. We re-sampled each file with the sample rate (330) of our prototype's input device. The decoded sound signals were first normalized to $[-1.0, 1.0]$ and then the first 512 signals were transformed into a spectrogram through STFT (window size = 64, window step = 32).

We used TensorFlow to train high-level Keras models and then used Keras2C [15] to convert them into low-level models that are supported by standard C libraries. This library generates smaller DNN models than TensorFlow Lite does and requires fewer third-party libraries when generating models. We used a lightweight convolutional neural network (CNN) model as shown in Fig. 5-3. The spectrogram of normalized input signals is first reshaped as a 2-D input layer with size of $N \times N$. The input layer is connected to a convolutional layer that has 8 filters, each of which has the kernel size of 3×3 and strides of 2×2. The activated output of the convolutional layer is flattened and connected to a dense layer, which outputs the probability distribution

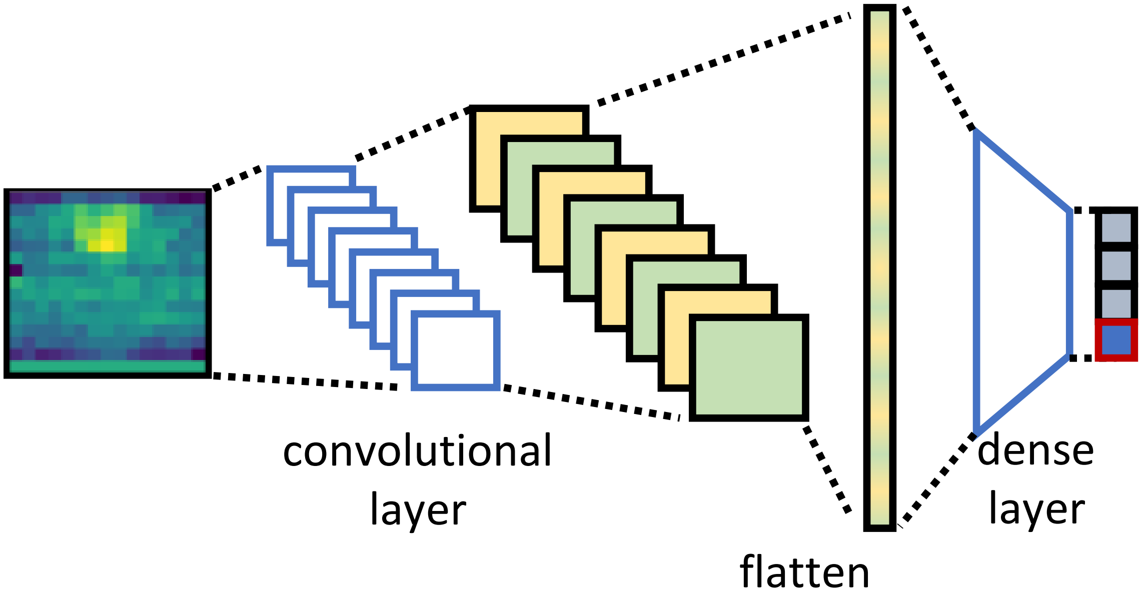

Figure 5-3: CNN structure for the classification task. The convolutional layer has 8 filters (size 3 × 3, strides 2 × 2) activated by ReLU. The dense layer is activated by softmax.

of a classification.

5.2.3 Model Accuracy

We evaluated the offline accuracy of the model with different size N of the input layer and number of mammal classes, as these two hyperparameters affected the size of the model. For 4-mammal classification, we used the data from Atlantic spotted dolphins, bearded seals, beluga white whales, and narwhals. Fig. 5-4 shows the offline accuracy distribution of each model hyperparameter configuration from 64 trials. For 4-mammal classification, the converted models when $N = 8$ and $N = 16$ can fit in our prototype, while $N = 32$ demands a larger model which didn't fit in the memory of our microcontroller. The average accuracy for $N = 16$ and $N = 32$ is 84% and 88% respectively. For 8-mammal classification, with $N = 32$, the average accuracy is 76%, which requires larger memory than that on our prototype.

We further evaluated the model on the online data generated from the ADC on our prototype. To this end, we deployed the converted 4-mammal classification model into the pipeline on our prototype. We then randomly selected 16 raw sound files (4

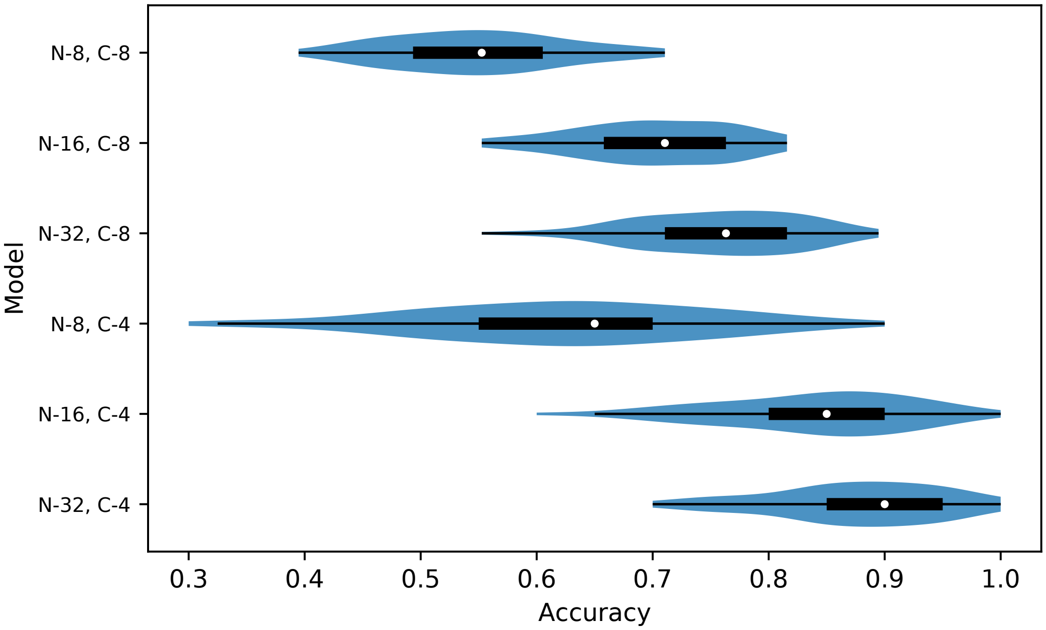

Figure 5-4: Distribution (boxplots and kernel density) of offline test accuracy CNN models with different input layer size (N) and number of mammal classes (C). The converted model of (N-16, C-4) can be deployed on our prototype.

samples for each mammal class) from the dataset. We used SIGLENT SDG 1032X function generator [76] to convert raw sound files (animal sounds [68] of whales, seals, dolphins, and narwhals) to analog signals that were fed to the ADC. For each input signal, the microcontroller sampled the input analogue signal at 330 samples/s, normalized it, and computed the spectrogram of the processed data before running inference on it. The average accuracy of the online model was 63% (while random guess is 25%) which is lower than the offline accuracy. The lower online accuracy can be attributed to low number of testing samples and loss of data resolution during sampling. Nonetheless, this result demonstrates that our prototype was successful in on-board classification of four marine animals with a reasonable accuracy at extremely low power, and this accuracy may be improved as the research evolves and with more comprehensive online evaluation.

5.2.4 Power Performance

Next, we used the MSP-EXP430FR5994 development kit board to evaluate the power consumption of the microcontroller while it was running the end-to-end pipeline. We computed the power as the product of the current and the voltage when the board was connected to a power supply. We used SIGLENT SDG 1032X function generator [76] to power on the board at 1.9 V. To measure the current, we connected Nordic Semiconductor's Power Profiler kit [73] as an ammeter in series with the board. Table 5.1 shows the power consumption of each stage of the pipeline. At 1.9 V input voltage, the microcontroller consumed 932 μW during ADC sampling, 1.3 mW while running inference on the data, and 902 μW when it was backscattering the inference result. These power measurements are in line with those expected from the microcontroller's datasheet. Since prior work has demonstrated the potential to harvest up to few milliwatts from underwater acoustic and ultrasonic signals [30], these results demonstrate that it would be possible for our node to operate entirely based on harvested energy.

Stage	**Power (μW)**	**Time (s)**	**Energy (mJ)**
ADC Sampling	932	1.6	1.49
Inference	1300	3.0	3.91
Backscatter	902	0.012	0.0108
Total Energy Consumption: 5.40 mJ			

Table 5.1: On-board power and energy analysis of the battery-free inference prototype with a data rate of 1 kbps and a sampling rate of 330 samples/s.

Since the power budget for ADC sampling and backscatter communication is similar, one could argue that it may be more efficient to send the entire data to the base station rather than performing the on-board inference. However, recall from 5.1.2 that this approach will be highly inefficient because of the limited throughput of the underwater channel. To demonstrate this inefficiency, consider the following example: For a backscatter node transmitting at 1 kbps (which is standard throughput for underwater acoustic modems), it will take around 6.14 seconds to send 512 samples of raw data (total # of bits = 512×12). The ADC consumes 1.49 mJ of energy to sample 512 samples at a sampling rate of 330 samples/s. This translates to 7.03 mJ of required energy. On the other hand, performing on-board inference, which typically takes 3 seconds (consuming 3.91 mJ) then transmitting the 12-bit-long inference result (consuming 0.0108 mJ) would require a total of 5.40 mJ, which makes sending raw data 30.19% more power hungry than sending inference data. This result demonstrates that battery-free edge inference is feasible, and that it makes on-board inference both more efficient and faster, even with this preliminary inference model. We envision that the gains can be significantly improved as the research evolves.

5.3 Conclusion

In this chapter, we proposed and prototyped battery-free ML inference in underwater environments as an exciting challenge. We have shown that a lightweight CNN model can achieve decent accuracy while its on-board energy requirements can be supported by these devices. Our initial results indicate feasibility of performing sensing and inference.

Chapter 6

Conclusion

This thesis discussed the different algorithms, protocols and fabrication techniques to enable salable underwater sensing, networking, localization and inference. Specifically, this thesis delivers four fundamental contributions: First, it presents the meta-material transducer design capable of performing ultra-wideband underwater backscatter. Second, it demonstrates how the spectral efficiency of this design can be improved by employing higher order modulation. Third, it introduces a technique for battery-free underwater positioning. Finally, it investigates the feasibility of battery-free ML in underwater environments.

6.1 Looking Forward

We have only started scratching the surface of possibilities for using backscatter as a low-power communication mechanism for underwater environments. While this thesis has taken major steps in unlocking some of these possibilities, the presented designs exhibit limitations which would be interesting to explore in the future. In this section, we revisit these limitations and highlight some of the exciting open problems and challenges for future research.

6.1.1 Underwater Backscatter Localization

Chapter 4 introduces underwater backscatter localization, a technology for ultra-low-power and scalable underwater positioning. Below, we highlight a number of open challenges and opportunities that would enable underwater backscatter localization to realize its full potential:

- *From 1D to 3D:* We demonstrated the feasibility of 1D localization by estimating the distance between two nodes. There are various avenues to extend UBL's design to 3D localization. Potential solutions include adding two or more nodes (e.g., hydrophones) to perform trilateration or incorporating phased-array transmitters or receivers for obtaining the angle-of-arrival (and combining it with distance information).

- *Mobility in Shallow Water:* In 4.2.2, we described how UBL can adapt its bitrate and hopping sequence to deal with challenges arising from extreme multipath (in shallow water) or from mobility. Unfortunately, addressing both types of challenges simultaneously leads to competing design requirements (as the former requires designs with a lower bitrate, while the latter requires higher bitrate). Developing underwater backscatter localization systems for such environments is an important open problem to explore in the future.

- *Long-range Backscatter Localization:* UBL inherits the range limitation of prior underwater backscatter systems. Specifically, in contrast to traditional acoustic communication which incurs one-way path-loss, backscatter suffers from a combination of path-loss on the downlink and the uplink. The signal degradation from round-trip path-loss has limited state-of-the-art designs to an operational range of around 60 meters (as demonstrated in section 2.4.2) [25]. As underwater backscatter evolves to operate over longer ranges, we expect new challenges to arise in the context of localization, which would need to be addressed in future designs.

- *Towards Tracking, Navigation, and Robotic Manipulation:* UBL's ability to localize batteryless nodes is a fundamental primitive for a variety of other tasks such as tracking, navigation, and robotic manipulation. For example, backscatter nodes can be used to tag underwater objects or marine animals and track them in real-time to understand mobility and migration patterns. Alternatively, UBL's localization primitive can enable novel navigation systems for underwater drones (AUVs and ROVs) using batteryless GPS anchors. A third application involves tagging underwater assets with backscatter nodes, and using their location to enable complex robotic manipulation tasks (e.g., grasping) in underwater environments. Realizing this capabilities will require addressing an exciting array of challenges with tools from a variety of disciplines ranging from networking to robotic perception, learning, and control.

6.1.2 Underwater Battery-Free Machine Learning

In Chapter 5, we proposed and prototyped battery-free ML inference in underwater environments as an exciting challenge for the community. Our initial results indicate feasibility of performing sensing and inference. Future potential directions include:

- *Battery-free on-device personalization:*Our design assumed that the same low-level model would be deployed on all battery-free devices deployed underwater. However, when implementing battery-free ML at a large scale, we expect that the sensed data by different underwater devices may differ at different locations (e.g, animal sounds near or far from a coast) and with environmental factors (e.g, temperature, pressure, multipath all impact sounds). Therefore, models on the devices would need to be locally tuned to fit different testing data distributions. Previous research [54] proposes to use early exit during offline training so that an end-device can personalize a part of the model to improve its testing accuracy. To enable such on-device personalization, light-weight model training algorithms that can satisfy the computational constrains on battery-free devices are needed.

- *Model-optimized hardware/software design.* Currently there are limited compilers that can build DNN models on ultra-low power MCUs (e.g, the MSP430 class) such as [78]. Co-design of task-specific hardware components alongside the ability to compile advanced models for lower-power MCUs can enable a wealth of new environmental monitoring applications. Specifically, one direct approach is to adopt existing optimization techniques for online inference on a target platform, such as post-training quantization [45, 57]. It can reduce the sizes of on-board models and increase computational efficiency by avoiding floating operations. Another "engineering" solution is to optimize the implementations of online inference for different target platforms by using hardware-specific instructions and libraries that offer faster or more energy-efficient operations. An auto-design pipeline for such individualization would be worthwhile considering the diversity of IoT devices and used low-power MCUs.

- *Battery-free distributed ML training.* Given the potentially large scale deployment of low-cost and low-pollution battery-free nodes in the ocean, another future direction is to continuously sense data and train ML models in a decentralized manner among participating nodes. This will enable life-long ML with up-to-date data that reflects the changing nature of underwater environments, which is critical for many application domains including climate change and biodiversity monitoring. Similar to on-device personalization, the key challenge is to have light-weight ML training algorithms on the devices. Another challenge is the difficulty of obtaining labelled data from underwater environments. One option is to use unsupervised ML to only train feature extractors (e.g, autoencoders) that can help train classifiers on a cloud server [102]; another option is to generate pseudo labels on local data through data augmentation [39]. In addition, distributed/federated ML solutions that can coordinate devices, stations, and cloud servers are also needed to manage the local training tasks, to make sure that they can satisfy the energy constraints of battery-free underwater devices.

Beyond these, various open problems remain in pushing the technological abilities of underwater backscatter in throughput, range, and sensing capabilities, and in bringing the technology to real-world applications by developing end-to-end systems for scalable underwater sensing, networking, localization, and inference. As our understanding of these low-power underwater sensors evolve, we envision that these sensors will become ubiquitous to enable a battery-free subsea IoT.

Bibliography

[1] COMSOL Multiphysics. https://www.comsol.com/.

[2] LFRX daughterboard. http://www.ettus.com. ettus inc.

[3] Steminc. https://www.steminc.com/, 2020.

[4] BJB Enterprise. 2015. Water clear shore 70 a polyurethane elastomer. https://bjbenterprises.com/index.php/wc-575-a-b/, 2015.

[5] Mohamed R Abdelhamid, Ruicong Chen, Joonhyuk Cho, Anantha P Chandrakasan, and Fadel Adib. Self-reconfigurable Micro-implants for Cross-tissue Wireless and Batteryless Connectivity. In *MobiCom '20*, pages 1–14, 2020.

[6] Fadel Adib, Zachary Kabelac, Dina Katabi, and Robert C. Miller. 3d tracking via body radio reflections. In *Usenix NSDI*, 2014.

[7] Sayed S. Afzal, Reza Ghaffarivardavagh, Waleed Akbar, Osvy Rodriguez, and Fadel Adib. Enabling higher-order modulation for underwater backscatter communication. In *Global Oceans 2020: Singapore - U.S. Gulf Coast*, 2020.

[8] Ranya Aloufi, Hamed Haddadi, and David Boyle. Privacy-preserving voice analysis via disentangled representations. In *CCSW '20*, 2020.

[9] Aquarian Audio. H2a hydrophone. http://www.aquarianaudio.com/h2a-hydrophone.html, 2019.

[10] Minja Bae, Jihyun Park, Jongju Kim, Dandan Xue, Kyu-Chil Park, and Jong Rak Yoon. Frequency-selective fading statistics of shallow-water acoustic communication channel with a few multipaths. *Japanese Journal of Applied Physics*, 55(7S1):07KG03, 2016.

[11] T. Casey, B. Guimond, and J. Hu. Underwater vehicle positioning based on time of arrival measurements from a single beacon. In *OCEANS 2007*, pages 1–8, 2007.

[12] V. Chandrasekhar and W. Seah. An area localization scheme for underwater sensor networks. In *OCEANS 2006 - Asia Pacific*, pages 1–8, 2006.

[13] Xiuzhen Cheng, Haining Shu, Qilian Liang, and David Hung-Chang Du. Silent positioning in underwater acoustic sensor networks. *IEEE Transactions on vehicular technology*, 57(3):1756–1766, 2008.

[14] J. H. Christensen, L. V. Mogensen, R. Galeazzi, and J. C. Andersen. Detection, localization and classification of fish and fish species in poor conditions using convolutional neural networks. In *2018 IEEE/OES Autonomous Underwater Vehicle Workshop (AUV)*, pages 1–6, 2018.

[15] Rory Conlin, Keith Erickson, Joseph Abbate, and Egemen Kolemen. Keras2c: A library for converting keras neural networks to real-time compatible c. *Engineering Applications of Artificial Intelligence*, 100:104182, 2021.

[16] CROWN. Crown Audio Power Amplifier. Part No. XLi-2500 750W https://www.crownaudio.com/en/products/xli-2500.

[17] DARPA. Ocean of Things Aims to Expand Maritime Awareness across Open Seas. https://www.darpa.mil/news-events/2017-12-06, 2017.

[18] DARPA. Angler. https://www.darpa.mil/program/angler, 2020.

[19] H. Dong, H. Wang, X. Shen, Z. Huang, and S. Ma. Detection of underwater weak signal via matched stochastic resonance. In *OCEANS 2017 - Aberdeen*, pages 1–7, June 2017.

[20] Dukane. Ulb-362b,ulb362/pl series underwater locator beacon. https://www.rjeint.com/wp-content/uploads/2017/01/ULB-362-ULB-362PL-Users-Manual-JAN-2018.pdf, 2020.

[21] Dukane. Underwater acoustic beacons. https://shop.boeing.com/medias/2598.pdf, 2020. Part No. 8102007K.

[22] Ettus Research, Usrp n210. https://www.ettus.com/.

[23] Xiaoning Feng, Zhuo Wang, Xiulong Liu, and Jiajie Liu. Adcnc-mac: asynchronous duty cycle with network-coding mac protocol for underwater acoustic sensor networks. *EURASIP Journal on Wireless Communications and Networking*, 2015(1):1–10, 2015.

[24] R. Fontana, A. Ameti, E. Richley, L. Beard, and D. Guy. Recent advances in ultra wideband communications systems. In *2002 IEEE Conference on Ultra Wideband Systems and Technologies (IEEE Cat. No.02EX580)*, pages 129–133, May 2002.

[25] Reza Ghaffarivardavagh, Sayed S. Afzal, Osvy Rodriguez, and Fadel Adib. Ultra-wideband underwater backscatter via piezoelectric metamaterials. In *Proceedings of the ACM Special Interest Group on Data Communication*. ACM, 2020.

[26] Reza Ghaffarivardavagh, Sayed Saad Afzal, Osvy Rodriguez, and Fadel Adib. Underwater backscatter localization: Toward a battery-free underwater gps. In *HotNets '20 125–131*, 2020.

[27] Reza Ghaffarivardavagh, Sayed Saad Afzal, Osvy Rodriguez, and Fadel Adib. Underwater backscatter localization: Toward a battery-free underwater gps. In *HotNets '20*, pages 125–131, 2020.

[28] Google. Oil and Gas Equipment Monitoring and Analytics. https://cloud.google.com/solutions/oil-and-gas-equipment-monitoring-and-analytics, 2019.

[29] Raffaele Guida, Emrecan Demirors, Neil Dave, and Tommaso Melodia. Underwater ultrasonic wireless power transfer: A battery-less platform for the internet of underwater things. *IEEE TMC*, 2020.

[30] Raffaele Guida, Emrecan Demirors, Neil Dave, and Tommaso Melodia. Underwater ultrasonic wireless power transfer: A battery-less platform for the internet of underwater things. *IEEE Transactions on Mobile Computing*, 2020.

[31] Raffaele Guida, Emrecan Demirors, Neil Dave, Jennie Rodowicz, and Tommaso Melodia. An acoustically powered battery-less internet of underwater things platform. In *2018 Fourth Underwater Communications and Networking Conference (UComms)*, pages 1–5. IEEE, 2018.

[32] Dirk Hahnel, Wolfram Burgard, Dieter Fox, Ken Fishkin, and Matthai Philipose. Mapping and localization with rfid technology. In *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, volume 1, pages 1015–1020. IEEE, 2004.

[33] J. Heidemann, Wei Ye, J. Wills, A. Syed, and Yuan Li. Research challenges and applications for underwater sensor networking. In *IEEE Wireless Communications and Networking Conference, 2006. WCNC 2006.*, volume 1, pages 228–235, April 2006.

[34] Huai Huang and Yahong Rosa Zheng. Node localization with aoa assistance in multi-hop underwater sensor networks. *Ad Hoc Networks*, 78:32–41, 2018.

[35] X. Huang, D. Lowe, R. Gandia, and E. Dutkiewicz. An impulse ultra-wideband system capable of concurrent transmission and reception, part i: Requirements and innovations. In *2006 International Conference on Communications, Circuits and Systems*, volume 2, pages 1251–1255, June 2006.

[36] X. Huang, D. Lowe, R. Gandia, and E. Dutkiewicz. An impulse ultra-wideband system capable of concurrent transmission and reception, part ii: Design and performance. In *2006 International Conference on Communications, Circuits and Systems*, volume 2, pages 1256–1260, June 2006.

[37] Junsu Jang and Fadel Adib. Underwater backscatter networking. In *Proceedings of the ACM Special Interest Group on Data Communication*, pages 187–199. ACM, 2019.

[38] Junsu Jang and Fadel Adib. Underwater backscatter networking. In *SIGCOMM '19*, pages 187–199. 2019.

[39] Wonyong Jeong, Jaehong Yoon, Eunho Yang, and Sung Ju Hwang. Federated Semi-Supervised Learning with Inter-Client Consistency & Disjoint Learning. In *ICLR '21*, 2021.

[40] Niels Jepsen, Eva B Thorstad, Torgeir Havn, and Martyn C Lucas. The use of external electronic tags on fish: an evaluation of tag retention and tagging effects. *Animal Biotelemetry*, 3(1):49, 2015.

[41] Chien-Chi Kao, Yi-Shan Lin, Geng-De Wu, and Chun-Ju Huang. A comprehensive study on the internet of underwater things: applications, challenges, and channel models. *Sensors*, 17(7):1477, 2017.

[42] H. Li, C. Tian, J. Lu, M.J. Myjak, J.J. Martinez, R.S. Brown, and Deng Z.D. An energy harvesting underwater acoustic transmitter for aquatic animals. *Scientific Reports*, 6(33804), 2016.

[43] Cho-Chung Liang. A study of a short-baseline acoustic positioning system for offshore vessels. *Marine Geodesy*, 22(1):19–30, 1999.

[44] En-Cheng Liou, Chien-Chi Kao, Ching-Hao Chang, Yi-Shan Lin, and Chun-Ju Huang. Internet of underwater things: Challenges and routing protocols. In *2018 IEEE International Conference on Applied System Invention (ICASI)*, pages 1171–1174. IEEE, 2018.

[45] Zhenhua Liu, Yunhe Wang, Kai Han, Wei Zhang, Siwei Ma, and Wen Gao. Post-training quantization for vision transformer. *Advances in Neural Information Processing Systems*, 34, 2021.

[46] H. Luo, Y. Zhao, Z. Guo, S. Liu, P. Chen, and L. M. Ni. Udb: Using directional beacons for localization in underwater sensor networks. In *2008 14th IEEE International Conference on Parallel and Distributed Systems*, pages 551–558, 2008.

[47] Zhihong Luo, Qiping Zhang, Yunfei Ma, Manish Singh, and Fadel Adib. 3d backscatter localization for fine-grained robotics. In *16th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 19)*, pages 765–782, 2019.

[48] Yunfei Ma, Nicholas Selby, and Fadel Adib. Minding the billions: Ultrawideband localization for deployed rfid tags. *ACM MobiCom*, 2017.

[49] Yunfei Ma, Nicolas Selby, and Fadel Adib. Drone relays for battery-free networks. In *ACM SIGCOMM*, 2017.

[50] Mohammad Malekzadeh, Richard Clegg, Andrea Cavallaro, and Hamed Haddadi. Dana: Dimension-adaptive neural architecture for multivariate sensor data. *ACM IMWUT*, 5(3), September 2021.

[51] Markets and Markets. Autonomous Underwater Vehicle (AUV) Market worth $1,638 million by 2025. https://www.marketsandmarkets.com/PressReleases/autonomous-underwater-vehicles.asp, 2020.

[52] MBARI. Wave Glider Hot Spot Enabling autonomous devices to talk back. https://www.mbari.org/technology/emerging-current-tools/communications/wave-glider-based-communications-hotspot/, 2019.

[53] Clive R McMahon, Emmanuelle Autret, Jonathan DR Houghton, Phil Lovell, Andrew E Myers, and Graeme C Hays. Animal-borne sensors successfully capture the real-time thermal properties of ocean basins. *Limnology and Oceanography: Methods*, 3(9):392–398, 2005.

[54] Fan Mo, Hamed Haddadi, Kleomenis Katevas, Eduard Marin, Diego Perino, and Nicolas Kourtellis. Ppfl: Privacy-preserving federated learning with trusted execution environments. In *MobiSys '21*, page 94–108, 2021.

[55] Camilo Mora, Derek P Tittensor, Sina Adl, Alastair GB Simpson, and Boris Worm. How many species are there on earth and in the ocean? *PLoS biology*, 9(8):e1001127, 2011.

[56] Peter C Müller and Metin Gürgöze. Natural frequencies of a multi-degree-of-freedom vibration system. In *PAMM: Proceedings in Applied Mathematics and Mechanics*, volume 6, pages 319–320. Wiley Online Library, 2006.

[57] Markus Nagel, Rana Ali Amjad, Mart Van Baalen, Christos Louizos, and Tijmen Blankevoort. Up or down? adaptive rounding for post-training quantization. In *International Conference on Machine Learning*, pages 7197–7206. PMLR, 2020.

[58] OHI. Ocean Health Index. http://www.oceanhealthindex.org/, 2020.

[59] Seyed Ali Osia, Ali Shahin Shamsabadi, Sina Sajadmanesh, Ali Taheri, Kleomenis Katevas, Hamid R. Rabiee, Nicholas D. Lane, and Hamed Haddadi. A hybrid deep learning architecture for privacy-preserving mobile analytics. *IEEE IoT Journal*, 7(5):4505–4518, 2020.

[60] Hassan M Oubei, Jose R Duran, Bilal Janjua, Huai-Yung Wang, Cheng-Ting Tsai, Yu-Cheih Chi, Tien Khee Ng, Hao-Chung Kuo, Jr-Hau He, Mohamed-Slim Alouini, et al. 4.8 gbit/s 16-qam-ofdm transmission based on compact

450-nm laser for underwater wireless optical communication. *Optics express*, 23(18):23302–23309, 2015.

[61] A. Palmeiro, M. Martín, I. Crowther, and M. Rhodes. Underwater radio frequency communications. In *OCEANS 2011 IEEE - Spain*, pages 1–8, 2011.

[62] L. M. Pessoa, M. R. Pereira, H. M. Santos, and H. M. Salgado. Simulation and experimental evaluation of a resonant magnetic wireless power transfer system for seawater operation. In *OCEANS 2016 - Shanghai*, pages 1–5, April 2016.

[63] Rajeswari Jayaraman. Beyond IoT: Internet of Underwater Things to Network the Oceans. https://www.prescouter.com/2017/06/internet-of-underwater-things/, 2017.

[64] Dean Roemmich, Gregory C Johnson, Stephen Riser, Russ Davis, John Gilson, W Brechner Owens, Silvia L Garzoli, Claudia Schmid, and Mark Ignaszewski. The argo program: Observing the global ocean with profiling floats. *Oceanography*, 22(2):34–43, 2009.

[65] J. Romme and B. Kull. On the relation between bandwidth and robustness of indoor uwb communication. In *IEEE Conference on Ultra Wideband Systems and Technologies, 2003*, pages 255–259, Nov 2003.

[66] Nasir Saeed, Abdulkadir Celik, Tareq Y Al-Naffouri, and Mohamed-Slim Alouini. Energy harvesting hybrid acoustic-optical underwater wireless sensor networks localization. *Sensors*, 18(1):51, 2018.

[67] Sayed Saad Afzal, Osvy Rodriguez, Reza Gaffarivardavagh, JunSu Jang, Jose Muguira Iturralde, Fadel Adib. Underwater backscatter open source, 2020.

[68] Laela Sayigh, Mary Ann Daher, Julie Allen, Helen Gordon, Katherine Joyce, Claire Stuhlmann, and Peter Tyack. The watkins marine mammal sound database: An online, freely accessible resource. *POMA*, 27(1):040013, 2016.

[69] Schrock. Fact or fiction?: An opera singer's piercing voice can shatter glass. https://www.scientificamerican.com/article/fact-or-fiction-opera-singer-can-shatter-glass/.

[70] Michael S Schwendeman and Jim Thomson. Sharp-crested breaking surface waves observed from a ship-based stereo video system. *Journal of Physical Oceanography*, 47(4):775–792, 2017.

[71] Hard Science. How to break glass with your voice. https://www.youtube.com/watch?v=rRZT7xO5KN4.

[72] Jenny Seifert. Ocean vital signs are stable, but bill of health isn't clean. NCEAS. https://www.nceas.ucsb.edu/news/ocean-vital-signs-are-stable-no-clean-bill-health, 2017.

[73] NORDIC Semiconductor. Nordic semiconductor power profiler kit ii. https://www.nordicsemi.com/Products/Development-hardware/Power-Profiler-Kit-2, 2019. PPK2.

[74] A. A. Shkel, L. Baumgartel, and E. S. Kim. A resonant piezoelectric microphone array for detection of acoustic signatures in noisy environments. In *2015 28th IEEE International Conference on Micro Electro Mechanical Systems (MEMS)*, pages 917–920, Jan 2015.

[75] SIGLENT. Siglent Arbitrary Waveform Generators. Part No. SDG1032X https://siglentna.com/product/sdg1032x/.

[76] SIGLENT. Siglent sdg1032x function generator. https://siglentna.com/product/sdg1032x/, 2019.

[77] K. Siwiak. Ultra-wide band radio: introducing a new technology. In *IEEE VTS 53rd Vehicular Technology Conference, Spring 2001. Proceedings (Cat. No.01CH37202)*, volume 2, pages 1088–1093 vol.2, May 2001.

[78] Benjamin Sliwa, Nico Piatkowski, and Christian Wietfeld. Limits: Lightweight machine learning for iot systems with resource limitations. In *ICC 2020-2020 IEEE International Conference on Communications (ICC)*, pages 1–7. IEEE, 2020.

[79] B Sai Srujana, Princy Mathews, VP Harigovindan, et al. Multi-source energy harvesting system for underwater wireless sensor networks. *Procedia Computer Science*, 46:1041–1048, 2015.

[80] Steminc. Piezo ceramic cylinder 54.1x47x40mm 17 khz, 2019. Part No. SMC5447T40111.

[81] Steminc. Piezo ceramic cylinder 36x31x20mm 30 khz. (2020). Part No. SMC3631T20111. https://www.steminc.com/PZT/en/piezo-ceramic-cylinder-36x31x20mm-30-khz, 2020.

[82] Steminc. Piezo ceramic cylinder 54.1x47x40mm 17 khz. (2020). Part No. SMC5447T40111. https://www.steminc.com/PZT/en/piezo-ceramic-cylinder-541x47x40mm-17-khz, 2020.

[83] Milica Stojanovic. Underwater acoustic communication. *Wiley Encyclopedia of Electrical and Electronics Engineering*, pages 1–12, 1999.

[84] Milica Stojanovic and James Preisig. Underwater acoustic communication channels: Propagation models and statistical characterization. *IEEE communications magazine*, 47(1):84–89, 2009.

[85] Hwee-Pink Tan, Roee Diamant, Winston KG Seah, and Marc Waldmeyer. A survey of techniques and challenges in underwater localization. *Ocean Engineering*, 38(14-15):1663–1676, 2011.

[86] TELEDYNEMARINE. Reson TC-4014 Hydrophone. Part No. TC-4014 http://www.teledynemarine.com/reson-tc-4014.

[87] The Verge. Alphabet's Tidal moonshot tracks individual fish to help sustainably feed humanity. https://www.theverge.com/2020/3/2/21161029/tidal-moonshot-sustainable-fish-farms, 2020.

[88] Stewart Thomas and Matthew S Reynolds. Qam backscatter for passive uhf rfid tags. In *2010 IEEE International Conference on RFID (IEEE RFID 2010)*, pages 210–214. IEEE, 2010.

[89] TI. Msp430fr5994. https://www.ti.com/product/MSP430FR5994, 2021.

[90] Francesco Tonolini and Fadel Adib. Networking across boundaries: enabling wireless communication through the water-air interface. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, pages 117–131. ACM, 2018.

[91] David Tse and Pramod Viswanath. *Fundamentals of wireless communication*. Cambridge university press, 2005.

[92] Paul A Van Walree and Roald Otnes. Ultrawideband underwater acoustic communication channels. *IEEE Journal of Oceanic Engineering*, 38(4):678–688, 2013.

[93] Deepak Vasisht, Swarun Kumar, and Dina Katabi. Decimeter-level localization with a single wifi access point. In *Usenix NSDI*, 2016.

[94] Jue Wang and Dina Katabi. Dude, where's my card? rfid positioning that works with multipath and non-line of sight. In *ACM SIGCOMM*, 2013.

[95] WHOI. WHOI Micro-modem. https://acomms.whoi.edu/micro-modem/.

[96] B. Woodward and H. Sari. Digital underwater acoustic voice communications. *IEEE Journal of Oceanic Engineering*, 21(2):181–192, 1996.

[97] Xu Yang, Xiaojie Yu, Chi Zhang, Sen Li, and Qiang Niu. Minegps: Battery-free localization base station for coal mine environment. *IEEE Communications Letters*, 25(8):2579–2583, 2021.

[98] Muhammad Azfar Yaqub, Syed Hassan Ahmed, Safdar Hussain Bouk, and Dongkyun Kim. Towards energy efficient duty cycling in underwater wireless sensor networks. *Multimedia Tools and Applications*, 78(21):30057–30079, 2019.

[99] C. H. Yu, S. H. Min, and J. W. Choi. Performance evaluation of toa-based sensor localization system in underwater sensor networks. In *2015 15th International Conference on Control, Automation and Systems (ICCAS)*, pages 493–498, 2015.

[100] Jucheng Zhang, Yunfeng Han, Cuie Zheng, and Dajun Sun. Underwater target localization using long baseline positioning system. *Applied Acoustics*, 111:129–134, 2016.

[101] Yuchen Zhao, Sayed Saad Afzal, Waleed Akbar, Osvy Rodriguez, Fan Mo, David Boyle, Fadel Adib, and Hamed Haddadi. Towards battery-free machine learning and inference in underwater environments. In *Proceedings of the 23rd Annual International Workshop on Mobile Computing Systems and Applications*, pages 29–34, 2022.

[102] Yuchen Zhao, Hanyang Liu, Honglin Li, Payam Barnaghi, and Hamed Haddadi. Semi-supervised Federated Learning for Activity Recognition, 2021. arXiv: 2011.00851.

[103] Xiaoyong Zhu. Monitoring environmental conditions near underwater datacenters using Deep Learning. Microsoft, 2018.