

# AquaHelper: Underwater SOS Transmission and Detection in Swimming Pools

Qiang Yang

qiang.yang@connect.polyu.hk

The Hong Kong Polytechnic University  
Hong Kong SAR, China

Yuanqing Zheng

csyqzheng@comp.polyu.edu.hk

The Hong Kong Polytechnic University  
Hong Kong SAR, China

## Abstract

Drowning incidents can occur in swimming pools even with professional lifeguards present. This is because drowning swimmers often face difficulties in calling for help due to choking, making it challenging for lifeguards to recognize them and provide a timely rescue. To address this problem, this paper presents AquaHelper, an underwater SOS system that can transmit and detect acoustic SOS signals in swimming pools. Specifically, a wearable device (e.g., a smartwatch) serves as an underwater SOS transmitter, with which a swimmer can call for help in emergency situations. Multiple underwater acoustic receivers are deployed to detect SOS signals and promptly alert lifeguards. The main challenge lies in the low transmission power of lightweight wearable devices, which poses difficulties in detecting weak SOS signals, particularly in low-SNR underwater scenarios. To achieve reliable underwater SOS detection, AquaHelper develops novel techniques (e.g., incorporating high-order harmonics, multi-scale window aggregation, and coherent combining of multiple receivers) to fully leverage the spectral, temporal, and spatial diversity of underwater acoustic signals. We also describe lessons learned and our solutions to address practical challenges involved in underwater SOS transmission and detection. Our experiments demonstrate the effectiveness of AquaHelper in detecting SOS signals in typical swimming pool environments.

## CCS Concepts

• **Human-centered computing** → *Ubiquitous and mobile computing design and evaluation methods.*

## Keywords

SOS detection, Underwater communication, Wearable devices

## ACM Reference Format:

Qiang Yang and Yuanqing Zheng. 2023. AquaHelper: Underwater SOS Transmission and Detection in Swimming Pools. In *The 21st ACM Conference on Embedded Networked Sensor Systems (SenSys '23)*, November 12–17, 2023, Istanbul, Turkiye. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3625687.3625816>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

SenSys '23, November 12–17, 2023, Istanbul, Turkiye

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0414-7/23/11...\$15.00

<https://doi.org/10.1145/3625687.3625816>




Figure 1: Underwater SOS transmission and detection.

## 1 Introduction

Each year, there are an average of 390 deaths due to drowning in swimming pools in the US [6]. It is commonly assumed that distressed swimmers could splash and yell for help during drowning, while nearby lifeguards quickly recognize them and provide timely rescue. In reality, however, most swimmers in trouble sink quickly and quietly underwater [3]. Therefore, even when professional lifeguards are present, some individuals still drown due to delayed rescue, often because lifeguards can be distracted or may miss distress signs [19]. The worldwide rescue standard expects lifeguards to discover the drowning incidents within 10 seconds and rescue within 20 seconds, known as the 10/20 principle [12]. The chance of survival increases with earlier rescue efforts—every second counts. Unfortunately, only 16% of the drowning swimmers are discovered within 10 seconds by lifeguards. On average, it takes around 69 seconds to discover a person in distress [28].

To aid lifeguards in detecting drowning individuals, existing solutions [1, 2, 4, 5, 29, 45, 65] mainly utilize specialized cameras installed above or in swimming pools for drowning detection. Most of them generate a warning signal when a swimmer is detected stationary in the water for a period of time, while the drowning person may already have suffocated [44]. Moreover, vision-based solutions are typically susceptible to poor lighting conditions, swimmer occlusions, and highly dynamic backgrounds [29]. As a result, many drowning swimmers may miss the best rescue time if we solely rely on vision-based systems.

As water-proof smartwatches are becoming increasingly popular for swimmers in assessing swimming performance [8, 11], *can we repurpose a smartwatch as an SOS transmitter?* We provide an affirmative answer by developing AquaHelper, an acoustic SOS system that can work in underwater scenarios. As illustrated in Fig. 1, AquaHelper consists of two main components: a wearable device based SOS transmitter and multiple receivers deployed underwater. AquaHelper uses speakers widely available in wearable devices to transmit acoustic SOS signals, which have much better underwater propagation properties than radio signals [22]. A swimmer can easily initiate the SOS transmission by pressing a rescue button on

a smartwatch in case of drowning or sudden discomfort. In addition, AquaHelper provides an API for existing wearable based drowning detection methods [42, 46, 55] to transmit SOS signals underwater. On the other hand, once the SOS signals are detected, AquaHelper will immediately alert lifeguards on duty and visualize the possible location of a drowning swimmer on a screen so that lifeguards can provide a timely rescue.

Turning the concept of AquaHelper into a practical system entails a series of technical challenges. First, lightweight smartwatches have relatively weak transmission power limited by their tiny form factors. In a large swimming pool, reliably detecting weak SOS signals with substantial underwater attenuation is challenging. Furthermore, underwater noise (e.g., water flowing and air bubbles) further reduces the Signal-to-Noise Ratio (SNR) [22]. Additionally, swimming pool environments suffer from interference (e.g., pump operations and intense splashing), which spans across a wide frequency band, severely hampering signal detection. Second, underwater acoustic channels are inherently dynamic, both in time and frequency due to various factors such as water flow and channel selective fading [22]. Consequently, we need a careful design of SOS signals to ensure high signal detectability and strong resilience to frequency selective fading. Third, in contrast to open water, the swimming pool's shallow and confined nature causes complex multipath reflections originating from the pool bottom, side walls, water surface, and swimmers [56]. These multipaths significantly affect SOS localization.

AquaHelper aims to enable lifeguards to recognize and rescue drowning people as early as possible, necessitating reliable SOS detection and visualization in low-SNR swimming pool environments. To achieve this goal, AquaHelper utilizes acoustic chirps as SOS signals. Chirp signals exhibit a linearly increasing frequency across the entire frequency band [43]. This characteristic makes them robust against frequency selective fading and the Doppler effect [39]. On the receiver side, we coherently aggregate the spectral high-order chirps, temporal consecutive observation windows, and multiple spatial receivers in the frequency, time, and spatial domain, respectively, to enhance the SNR of SOS signals that could otherwise be missed. To visualize the location of a drowning swimmer, we first initiate a pre-calibration operation to synchronize the distributed receivers. Next, we leverage the time-frequency linearity of chirp signals to separate the direct path from multipath effects, which allows us to precisely estimate the arrival time of SOS signals. After that, we partition the swimming pool into small grids and generate a visualization that indicates the likelihood of the drowning swimmer's presence within each grid. Thus, lifeguards can quickly identify the drowning user, facilitating swift and effective rescue efforts. To summarize, this paper makes the following contributions:

- To the best of our knowledge, AquaHelper is the first underwater drowning SOS transmission and detection system that transforms lightweight wearable devices (e.g., smartwatches) into SOS signal transmitters for swimmer safety.
- We develop an underwater acoustic SOS detection method that fully leverages the signal diversity in the spectral, temporal, and spatial domains to significantly improve the SNR and achieve reliable SOS detection and visualization.

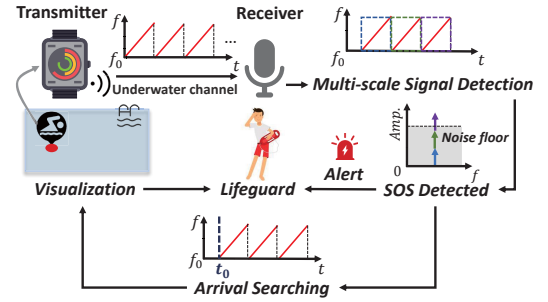


Figure 2: System overview of AquaHelper.

- We conducted a comprehensive evaluation of AquaHelper in swimming pools, and the results show that AquaHelper can effectively detect SOS signals and alert lifeguards to increase the chances of a timely rescue operation.

We emphasize that AquaHelper is not intended to replace professional lifeguards or other swimmer supervision techniques *but to work alongside them*, providing an additional safety net and a help-asking channel for swimmers. The objective is to alert lifeguards in case of an emergency and to shorten the rescue time, which can be of utmost importance for drowning persons in need.

## 2 AquaHelper Overview

Most smartwatches are equipped with one or more physical buttons on the side, offering swimmers a convenient means to trigger a command. AquaHelper allows swimmers to send SOS signals by pressing a rescue button in case of drowning or sudden discomfort, enabling them to actively seek for help. In addition, AquaHelper provides an API for existing wearable-based drowning detection systems [42, 46, 55] to send underwater SOS signals. Such an approach will leverage the ubiquity of wearable devices to democratize swimming safety and make it available to everyone with a water-proof smartwatch. Swimming pools can also provide smartbands to swimmers for temporary uses. Therefore, our system assumes the availability of water-proof smartwatches as transmitters and pre-deployed hydrophones in swimming pools as receivers.

Figure 2 illustrates the system overview of AquaHelper. In an emergency, a swimmer can activate SOS transmission by pressing a rescue button on the smartwatch. The underwater acoustic receivers continuously monitor SOS signals in a multi-scale manner. Upon detection of an SOS signal, AquaHelper will immediately alert lifeguards through audible alarms and visualize the incident area on a screen. This real-time information empowers lifeguards to provide timely rescue to drowning swimmers.

In the following, we describe the key components of AquaHelper including SOS transmission (Sec. 3), SOS detection (Sec. 4), and SOS alerting and visualization (Sec. 5). We also describe lessons learned and our solutions to address some practical issues (Sec. 6).

## 3 SOS Transmission

### 3.1 Challenges of Underwater SOS Transmission and Detection

The key objective of AquaHelper is to transmit SOS signals using lightweight wearable devices (e.g., smartwatches), reliably detect the signals at underwater receivers, and then alert lifeguards in a




Figure 3: The Chirp signal and dechirp process.

timely manner. Intuitively, we can send a single-tone sine wave as the SOS signal and examine the corresponding frequency power on the receiver side. However, several factors pose significant challenges in accomplishing this goal. The limited transmission power of smartwatches, due to their small size, hinders the strength of SOS signals. Moreover, the underwater environment introduces substantial attenuation, causing the signal to weaken significantly over long distances before arriving at the receiver. During the propagation in underwater channels, SOS signals also suffer from frequency selective fading due to multipath effects that arise from signal reflections and scattering within swimming pools [21]. Additionally, the underwater environment introduces various sources of noise, including bubbles, splashes, and water pumps, which can overpower the already weak SOS signal [22]. Collectively, these factors contribute to an extremely low SNR, making the detection of underwater signals a highly challenging task.

### 3.2 Acoustic Chirp as SOS Signals

To overcome these challenges, we consider both the signal design at the transmitter and the reception strategy at the receiver to jointly improve the SNR. On the transmitter side, we use the acoustic chirp as the SOS signal. As shown in Fig. 3(a), the frequency of a chirp increases linearly with time. The chirp signal spreads within a frequency band and hence is robust to frequency selective fading. More importantly, it features the pulse compression property [18], which means its signal power can be accumulated over time to yield a high SNR after the dechirp-based demodulation [64]. Thanks to this property, the chirp signal can be detected even under the noise floor (Fig. 3(c)), which makes it ideal for long-range communication [63]. Mathematically, a chirp can be represented as follows:

$$S(f', t) = C(t) \cdot e^{j2\pi f' t} = e^{j2\pi(f_0 + \frac{k}{2}t)t} \cdot e^{j2\pi f' t} \quad (1)$$

where  $C(t) = e^{j2\pi(f_0 + \frac{k}{2}t)t}$  denotes a base chirp, starting from  $f_0$ .  $k = \frac{Bw}{T}$  is the increasing rate of the frequency, where  $Bw$  and  $T$  are the bandwidth and the chirp duration, respectively.  $S(f', t)$  is a time-shifted version of the base chirp, starting from  $f_0 + f'$ , and its initial frequency is  $f'$ . The chirp signal adopts Chirp Spread Spectrum (CSS) modulation [64], in which we can demodulate the chirp  $S(f', t)$  through multiplying it by the conjugate of the base chirp, denoted by  $C^{-1}(t)$ . This procedure is called dechirp and is represented below:

$$S(f', t) \cdot C^{-1}(t) = C(t) \cdot e^{j2\pi f' t} \cdot C^{-1}(t) = e^{j2\pi f' t} \quad (2)$$

Then, we can perform a Fast Fourier Transform (FFT) on the dechirped signal, and the power of all samples will be accumulated at the initial frequency  $f'$ .

Fig. 3(b) shows a chirp with the initial frequency  $f' = 0$  (i.e., a base chirp), which is almost overwhelmed by the noise. After




Figure 4: Underwater noise and channel property.

dechirp and FFT (Fig. 3(c)), the power of all chirp samples is concentrated at the initial frequency (0 Hz). In contrast, the noise spreads across the whole frequency band. Consequently, the dechirped signal peak stands out the noise floor. We detect the SOS signal by examining if the highest peak satisfies  $X_{peak} > \bar{X}_{noise} + \beta\sigma(X)$ , where  $X_{peak}$  is the highest FFT peak,  $\bar{X}_{noise}$  is the noise floor, and  $\sigma(X)$  is the standard variance.  $\beta$  is the threshold. This energy accumulation property of chirp signals opens up the possibility of SOS detection in low-SNR underwater scenarios. We will introduce more reception strategies to improve the SNR on the receiver side in Sec. 4.

### 3.3 SOS Chirp Design

The design of the SOS chirp involves careful consideration of the bandwidth and duration parameters. While a wider bandwidth and longer duration could potentially increase the accumulated energy after dechirp in theory, the property of underwater channels must be considered to maximize the SNR improvement.

**Frequency.** To determine the appropriate bandwidth for SOS chirps, we analyze the spectrum of underwater noise, as shown in Figure 4(a), and plot the Cumulative Distribution Function (CDF) of its energy with an increasing frequency in Figure 4(b). We can observe that 96% of underwater noise is distributed below 1 kHz [22]. In this case, a highpass filter with a cutoff frequency of 1 kHz can be used to remove the vast majority of the noise. However, considering the transition width of a non-ideal filter, there is a possibility that residue noise could still overpower the weak SOS signal. Therefore, we leave 500 Hz as the guard band and choose 1.5 kHz as the lower bound of the SOS chirp. We also perform the channel estimation between the transmitter and receiver 5m apart underwater using a chirp. The frequency response is shown in Fig. 4(c). We can see a rapid attenuation above 3.5 kHz since higher frequencies suffer more from the attenuation underwater [22]. As such, we set 3.5 kHz as the upper bound of the SOS chirp. Therefore, we choose 1.5 ~ 3.5 kHz as the bandwidth of SOS chirps, striking a balance between noise interference and signal attenuation.

**Duration.** Considering the low transmission power of smartwatches and the severe attenuation over underwater channels, the signal at the receiver side would have a very low SNR. In this case, a longer chirp duration can theoretically accumulate more energy, increasing the chance of successful SOS detection [64]. However, a longer chirp duration requires a larger detection window and longer waiting time. To strike a balance between effectiveness and efficiency, we set the chirp duration to one second and enable smartwatches to transmit chirps consecutively. Accordingly, we design a dynamic SOS detection scheme utilizing multi-scale detection windows that can adapt to varying SNR scenarios.






Figure 5: Incorporating high-order chirps.

#### 4 SOS Detection

As a drowning rescue assistant, AquaHelper should satisfy two key requirements—high detection reliability and short response time. To strike a balance between the two potentially conflicting goals, AquaHelper makes full use of received chirp samples so that it can aggregate sufficient power to achieve reliable detection in a multi-scale manner. In particular, AquaHelper enhances the chirp SNR by dynamically aggregating the energy from high-order harmonics, multiple detection windows, and different receivers in the frequency, temporal, and spatial domains, respectively.

To process the acoustic signals captured by underwater microphones, we down-convert the real-valued signal  $S_R(f', t)$  to the baseband by multiplying with the cosine and sine components of the carrier frequency and then lowpass filtering:

$$S(f', t) = LPF(S_R \cdot \cos 2\pi f_c t) - j \cdot LPF(S_R \cdot \sin 2\pi f_c t) \quad (3)$$

where  $f_c$  is the center frequency of the base chirp (*i.e.*, 2.5 kHz), and  $LPF(\cdot)$  is the lowpass filter with a cutoff frequency of  $\frac{B_w}{2}$ . Thus, we can obtain the complex chirp signals to facilitate dechirp and subsequent processing.

##### 4.1 Incorporating High-order Harmonics

Acoustic transceiver systems are typically expected to be linear. However, due to hardware imperfection of the diaphragm and amplifier, there is usually some form of non-linearity in practical acoustic systems [15, 53]. As a result, the acoustic chain brings nonlinear distortion, introducing high-order harmonics to the recorded chirps:

$$\begin{aligned} S_R(f', t) &= \sum_{r=1}^{\infty} a_r S_r^*(f', t) \approx \underbrace{a_1 S_*(f', t)}_{\text{Linear}} + \underbrace{a_2 S_2^*(f', t) + a_3 S_3^*(f', t)}_{\text{Nonlinear}} \\ &= (a_1 + \frac{3a_3}{4}) \cos(2\pi(f_0 + kt + f')t) \\ &\quad + \frac{a_2}{2} \cos(2\pi(2f_0 + 2kt + 2f')t) \\ &\quad + \frac{a_3}{4} \cos(2\pi(3f_0 + 3kt + 3f')t) + \frac{a_2}{2} \end{aligned} \quad (4)$$

where  $r$  is the order index, and  $a_r$  is a gain coefficient.  $S_*(f', t) = \cos(2\pi(f_0 + kt + f')t)$  are high-order harmonics produced at microphone recording components. In practice, the higher-order terms (larger than 3) are extremely weak and can be ignored.

As shown in Fig. 5, we can observe that the power of the chirp signal is dispersed to two parts: the fundamental chirp and its high-order chirps which start from  $r(f_0 + f')$  and increase with




Figure 6: Signal Detection with window misalignment.

a rate of  $rk$ . We note that these high-order chirps are generated due to hardware imperfection during *recording at the microphone*, independent of the channel attenuation in high-frequency bands as mentioned in Sec. 3.3. Therefore, we can aggregate these chirps to compensate for the energy loss due to the microphone nonlinearity.

Given that the nonlinearity of most electronics is symmetric (*i.e.*, the same impact on positive and negative signal parts), the odd harmonics will be much higher than the even ones [48]. Hence, we only use the 3rd harmonic chirp to enhance the SNR signal. Specifically, we perform down-conversion on recorded signals but with different center frequencies to separate the fundamental chirp and its 3-order chirp. After that, we can perform dechirp and FFT on them individually. Since the 3-order chirp is always three times higher than the fundamental chirp in the frequency, we can down-sample the dechirp FFT result of the 3-order chirp by one-third and then add it to the dechirp FFT of the fundamental chirp, as shown in Fig. 5. Consequently, two SOS peaks will be aligned and superposed constructively, while the random noise will be added destructively. In this way, we can compensate for the diffused energy because of hardware imperfection to enhance SNR in the frequency domain.

##### 4.2 Handling Window Misalignment

The aforementioned detection approach assumes that the detection window aligns well with the chirp edges. However, such an assumption rarely holds since the receiver does not know the chirp's accurate arrival time due to the dynamic underwater channel. As shown in Fig. 6, if the detection window and the chirp signal are not aligned, there would be two fragments from adjacent chirps in the window. These two chirp fragments have different initial frequencies. Consequently, the SOS peak in the FFT result after dechirp will be split into two lower peaks:  $f_A$ , and  $f_B = f_A - Bw$ .

To address this issue, we adopt a downsampling technique [64]. Specifically, we downsample the dechirped signal with a sampling rate equal to the bandwidth  $Bw$ . Thus, the second frequency peak  $f_B$  will be merged with the first peak  $f_A$  due to frequency aliasing:

$$f_B^{\text{alias}} = |f_B - \ell \cdot Bw|, \ell \in \mathbb{Z} \quad (5)$$

We can see that  $f_B^{\text{alias}} = f_A$  when  $\ell = -1$ . However, since the fundamental chirp and its 3-order chirp have different bandwidths, we cannot downsample the recorded signal directly with a single sampling rate of  $Bw$ . Instead, we downsample the fundamental chirp and its 3-order chirp using different downsampling factors: 24 for the fundamental chirp ( $48\text{kHz} \rightarrow 2\text{kHz}, Bw$ ), and 8 for the 3-order chirp ( $48\text{kHz} \rightarrow 6\text{kHz}, 3Bw$ ). After that, we can perform




Figure 7: Multi-scale signal detection.

FFT and add their respective results to incorporate the 3-order chirp. By doing so, AquaHelper can address the misalignment issue and detect the SOS chirp regardless of the window's location.

### 4.3 Aggregating Multi-scale Windows

The dynamic underwater channel is affected by many factors, such as distance, noise, occlusion, and interference. If the SNR is high, it is possible to detect the SOS signal within a single window. However, if the SNR decreases, the signal becomes very weak and cannot be detected even if we accumulate the energy of an entire chirp. Thus, we adopt a multi-timescale signal detection approach, which keeps the latest consecutive windows (e.g., three windows) in the buffer and successively aggregates multiple windows to detect the SOS signal. As shown in Fig. 7, if the SOS signal is detected in the first window, AquaHelper will alert lifeguards immediately. Otherwise, we buffer this window and concatenate it with the following windows for iterative detection until the SOS signal power goes beyond the noise floor. This multi-scale approach can adaptively detect the SOS signal in the underwater scenario with dynamic SNRs, which strikes a balance between detection effectiveness and efficiency.

### 4.4 Combining Multiple Receivers

In addition to aggregating multiple windows, we also utilize multiple receivers to further enhance SNR. The basic idea is to coherently combine the signals received on multiple receivers. Specifically, we align and superpose the signals from multiple receivers:

$$S = \sum_{n=1}^N S_n(f', t - \frac{h_n}{F_s}) \quad (6)$$

where  $N$  is the number of receivers,  $F_s$  is the sampling frequency, and  $S_n$  is the received signal on the  $n^{th}$  receiver.  $H = \{h_1, h_2, \dots, h_N\}$  is the sample offset vector of receivers, where  $h_1 = 0$ , and  $h_2$  is the sample delay between receiver  $Rx_1$  and receiver  $Rx_2$ . However, the sample offset  $H$  cannot be obtained before detecting the SOS signal, creating an egg-and-chicken problem. A simple but inefficient solution can be a brute-force search for all possible sample delays to maximize the FFT peak after dechirp. However, if different receivers are not synchronized, the sample delays can have a large possible range which would impose significant computational overhead.

To deal with this problem, we employ a pre-calibration process to synchronize different receivers during system launch, which will be described in Sec. 5.1. By doing so, we can substantially reduce the search space by exploiting the prior knowledge of a maximum




Figure 8: Sample-level synchronization.

time delay as a constraint based on the swimming pool geometry. For example, if two receivers  $Rx_1$  and  $Rx_2$  are synchronized with their respective sampling indices  $m_1$  and  $m_2$ , then the search range of  $h_2$  can be reduced as:

$$h_2 \in \left[ m_2 - \frac{\max(d_{1,2})}{v_s}, m_2 + \frac{\max(d_{1,2})}{v_s} \right] \quad (7)$$

where  $\max(d_{1,2})$  is the maximum distance difference between the path from the signal source to the two receivers, which is known beforehand based on the deployed receiver locations.  $v_s$  is the speed of acoustics in water. In this way, we can significantly reduce the computational overhead and expedite the SOS detection process.

## 5 SOS Alerting and Visualization

Upon detecting an SOS signal, AquaHelper immediately notifies lifeguards through sound alerts. To further reduce every second of the rescue time, AquaHelper also visualizes the possible area where the drowning person is located in the swimming pool on a nearby screen. To this end, AquaHelper should achieve accurate location estimation and effective visual coverage. However, there are practical challenges that need to be addressed.

**Challenges of SOS Localization and Visualization.** Intuitively, we may compare the timestamps of transmission and reception to calculate the time of flight (ToF) of acoustic signals to estimate the swimmer's location. This solution does not work well because of the following reasons. (1) Lack of synchronization. The transmitter and receivers are not synchronized due to their different local clocks and OS latency. Previous works [21, 47] address this challenge using two-way communication and timestamp exchange. However, building additional communication channels from the hydrophones back to a smartwatch is challenging due to the low reception sensitivity of a lightweight smartwatch. Moreover, the limited local computational capability of a smartwatch makes it hard to efficiently execute complex algorithms to detect weak reply messages from hydrophones over low-SNR underwater channels. (2) Arrival time estimation. SOS localization requires precise estimation of the arrival time of SOS signals at the receiver. However, due to the dynamic underwater channels, the receivers may not be able to detect the SOS chirp at the very beginning and miss a few samples, leading to a localization error. (3) Multipath reflections. The swimming pool environment introduces significant multipath reflections due to its shallow and confined nature. To avoid obstructing swimming activities, the receivers are deployed on the pool wall, exacerbating the multipath effect. We present our solutions to address these practical challenges in the following.




Figure 9: Searching the arrival time of SOS signals.

### 5.1 Receiver Synchronization

To deal with the first challenge, we propose a sample-level synchronization protocol tailored for underwater distributed receivers. The high-level idea is to calibrate the receiving buffer of different receivers before localization with the known distance between them. As shown in Fig. 8, for the receiver  $Rx_A$ , in addition to an underwater microphone, we also attach a high-power underwater loudspeaker to broadcast a synchronization signal at  $t_0$ . Subsequently, the signal arrives at the receivers  $Rx_A$  and  $Rx_B$  at the sample index  $m_{A_1}$  and  $m_{B_1}$  in their recording buffer, respectively. Given the known distance  $d_{AB}$  between  $Rx_A$  and  $Rx_B$ , we can synchronize the two receivers by aligning the buffer of  $Rx_B$  with  $Rx_A$  by the calibrated index  $m'_{B_1} = m_{B_1} - \frac{d_{AB} \cdot F_s}{v_s}$ . AquaHelper uses a down chirp with a decreasing frequency as the synchronization signal, orthogonal to SOS up-chirps.

This calibration-based synchronization approach offers two advantages. First, it effectively reduces the search space for coherently combining signals from multiple receivers (Sec. 4.4) and opens up opportunities for SOS localization and visualization. Second, unlike the existing ToF-based methods, this method is transparent to users (*i.e.*, transmitter) and does not require their active involvement.

### 5.2 Arrival Time Estimation

After synchronization, we can calculate the Time Difference of Arrival (TDoA) of SOS chirps. Suppose the SOS signal transmitted by the smartwatch  $Tx$  reaches  $Rx_A$  and  $Rx_B$  at the sampling indices  $m_{A_2}$  and  $m_{B_2}$ , we can calculate TDoA  $t_{AB}$  between the two receivers by sample counting [47]:

$$t_{AB} = \frac{m_{B_2} - m'_{B_2}}{F_s} = \frac{(m_{B_2} - m_{B_1}) - (m_{A_2} - m_{A_1})}{F_s} + \frac{d_{AB}}{v_s} \quad (8)$$

Considering the high transmission power of underwater loudspeakers, we can use a matched filter to detect the synchronization signal (*i.e.*,  $m_{A_1}$  and  $m_{B_1}$  in Fig. 8). However, the exact arrival time of weak SOS signals ( $m_{A_2}$  and  $m_{B_2}$ ) is hard to measure due to the dynamic nature of underwater channels. Figure 9 illustrates the SOS chirp signals after down-conversion in two typical scenarios. Theoretically, the FFT peak should be  $f_0 = -\frac{Bw}{2}$  Hz if the detection window is exactly aligned with the chirp arrival time  $m_0$ . In practice, due to the low SNR, the SOS signal may not be detected at the beginning (*i.e.*,  $m_0$ ) but at a later time (*e.g.*,  $m_1$ ). Consequently, the frequency peak after dechirp becomes  $f_1$  instead of the expected frequency peak  $f_0$ . In high-SNR situations, the SOS signal can be detected with a fractional chirp at an earlier time (*e.g.*,  $m_2$ ) as illustrated in the figure.

To address the misalignment problem in arrival time detection, we leverage the frequency-time linearity of chirp signals (Eq. 1). By




Figure 10: Multipath effect mitigation with (a) Cross-correlation and (b) Dechirp.

analyzing the frequency peak after dechirp, we can infer the actual arrival time using the following equation:

$$m_0 = m_* \pm \frac{f_* + \frac{Bw}{2}}{k} F_s \quad (9)$$

where  $m_*$  is the sample index of the current window, and  $f_*$  is the corresponding frequency peak after dechirp. The operation  $\pm$  indicates the track direction, with the positive direction corresponding to the high-SNR case and the negative direction corresponding to the low-SNR case.

### 5.3 Multipath Interference Mitigation

Swimming pool environments are characterized by significant multipath reflections, which pose challenges for synchronization and localization. Figure 10 illustrates the cross-correlation envelope and the dechirp result for an SOS chirp signal. The presence of reflections causes multiple peaks. We observe that some reflections can be stronger than the LoS path due to constructive superposition, but they always arrive later than the LoS path. As such, we adopt the earliest peak instead of the highest as the LoS path [56]. Specifically, we find all local maxima above a certain threshold empirically set as  $0.5 \times \max(\text{peak})$ . Theoretically, AquaHelper can distinguish the two signals whose arrival time difference is greater than  $\frac{1}{Bw} = 0.5$  ms. Therefore, among these maxima, we choose the peak with a minimum distance  $\frac{1}{Bw}$  to other maxima to remove noisy peaks. Finally, the first maximum is selected as the LoS peak. It is important to note that the temporal direction in the dechirp result is opposite to the correlation, as the delayed reflections will be translated into peaks from back to front frequency bins after the dechirp operation.

### 5.4 Likelihood-based SOS Visualization

With the TDoAs of multiple receiver pairs obtained by Eq. 8, a naive way to locate the swimmer is to find the intersection of hyperbolas with different TDoA measurements. However, this method is sensitive to measurement error, and a single location result may not indicate the exact position of the swimmer, which could mislead lifeguards. To address this problem, we employ a likelihood-based hologram to visualize the possible SOS area. The hologram visualizes the likelihood that a drowning swimmer is in a particular position. Specifically, we partition the 2D swimming pool into grids and calculate the likelihood of each grid by measuring the difference between the estimated and theoretical TDoAs:

$$\mathcal{L}(\theta \mid \mathbf{x}_{\in G}) = 1 - \text{Normalize}(\sum_{p,q \in \text{Pairs}} \|t_{p,q} - t_{p,q}^x\|^2) \quad (10)$$




Figure 11: SOS visualization.

where  $\mathbf{x}$  is a grid in the grid set  $G$ , and  $p, q$  are receivers from all hydrophone pairs.  $t_{p,q}$  denotes the measured TDoA between receivers  $p$  and  $q$ , while  $t_{p,q}^x$  represents the theoretical TDoA for a virtual transmitter at this grid  $\mathbf{x}$ . *Normalize* is a min-max normalization function. Intuitively, the smaller the residue between the theoretical and the estimated TDoA is, the higher the likelihood  $\mathcal{L}(\theta|\mathbf{x}_{\in G})$  should be.

Figure 11 illustrates a hologram with different likelihood contours. Here light colors denote higher likelihood values. We can see that the swimmer is within the area of the highest likelihood, which decreases as the increasing distance. The likelihood threshold for visualization (e.g., 90%) can be adjusted as needed, which ensures effective coverage as the swimmer may be located at relatively low-likelihood areas due to estimation errors. In this way, AquaHelper can visualize the possible area of the drowning swimmer with the likelihood hologram to assist lifeguards in prompt rescue.




Figure 12: (a) A chirp with bursting noise. (b) Dechirp with one window. (c) The SOS peak outstands the noise by aggregating energy from more windows.

## 6 Practical Considerations

**Bursting Noise.** While most underwater acoustic noise, such as water flow and air bubbles, is typically under 1 kHz, in our experiments, we observe some types of noise (e.g., intense splashes, water pump launching) span across the entire spectrum and cannot be removed by a lowpass filter. As shown in Fig. 12(a), the bursting pump noise can generate significant interference across the frequency spectrum, resulting in a notable increase in the noise floor and overpowering the SOS peak (Fig. 12(b)). Fortunately, the multi-scale detection of AquaHelper can handle this issue. In practice, bursting noise is sporadic and lasts for a short period. Thus, by combining multiple windows for the dechirp, the SOS peak can be proportionally amplified, as illustrated in Fig. 12(c). Meanwhile, multi-window aggregation can also effectively spread the noise and reduce the likelihood of false alarms.

**Sampling Frequency Offset.** In practice, there may be slight variations between different receivers in the sampling frequencies due to hardware imperfections, resulting in a sampling frequency offset ( $SFO = F_{sA} - F_{sB}$ ). This small offset can accumulate over time,




Figure 13: Doppler Effect on SOS signals. (a) The theoretical chirp (red) and the actual chirp (blue). (b) The FFT peak before and after Doppler compensation.




Figure 14: (a) Two SOS chirps are overlapped in multi-user scenarios. (b) Multiple users can be separated in the frequency domain after dechirp.

causing an estimation error in SOS localization. Performing calibration frequently can address this issue, but the synchronization signal can potentially disturb swimmers. To strike a balance between accurate synchronization and minimizing disturbance, AquaHelper first measures the offset between several calibrations within a fixed duration after the system launch. Consequently, we can fit the SFO and compensate for the measured TDoA by  $\frac{(m_{B_2} - m_{B_1})SFO}{F_{sB}}$ . Moreover, AquaHelper also performs periodic synchronization every hour and during pool-cleaning times to calibrate the compensation residue.

**Doppler Effect.** When a drowning swimmer struggles in the water, the arm and hand movement inevitably affects the frequency of SOS signals due to the Doppler effect. As shown in Fig. 13(a), compared to the original chirp, the SOS signal has a Doppler frequency shift  $\Delta f = \frac{v}{v_s}(f_0 + kt)$ , where  $v$  is the velocity of the swimmer-worn smartwatch. Assuming the maximum  $v$  is 1 m/s, the corresponding maximal  $\Delta f$  will be  $\pm 3$  Hz. Based on our experiment, the SOS peak can still exceed the noise floor by a wide margin in most cases, as shown in Fig. 13(b). The frequency shift range of low-frequency signals is quite narrow, which is also why the chirp signal is resilient to the Doppler effect [18]. To maximize the SNR and increase the possibility of SOS detection, AquaHelper compensates for the Doppler effect by multiplying the received signal with a reversed frequency shift  $e^{-j2\pi \frac{v}{v_s}(f_0 + \frac{k}{2}t)t}$ . If the FFT resolution is 1 Hz,  $v$  can be accordingly discretized to  $\frac{\Delta f v_s}{f_0 + kt}$ , where  $\Delta f \in \mathbb{Z} = [-3, 3]$  Hz. Thus, we can search for a speed  $v$  to maximize the height of the FFT peak after dechirp. As shown in Fig. 13(b), we can observe that the FFT peak becomes sharper and higher after Doppler compensation. Although the arm speed may vary continuously within a chirp duration, such a compensation method can partially rectify the frequency shift and enhance the SNR. This Doppler shift can also be compensated at the sub-window level to achieve better performance.



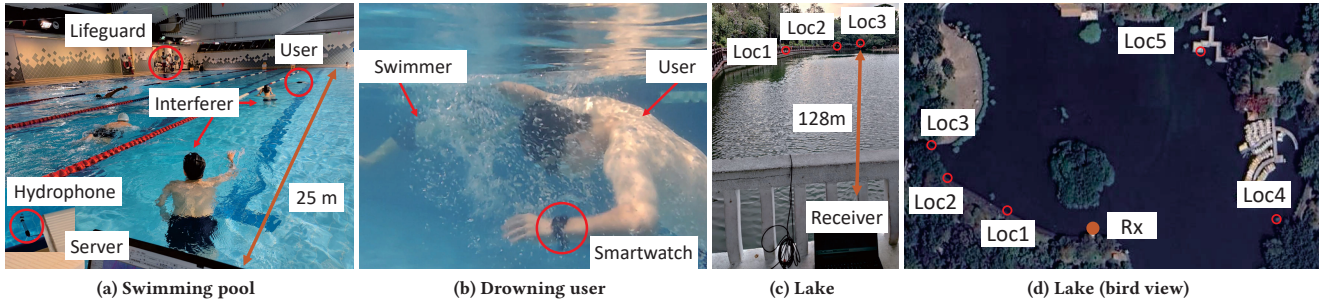


Figure 15: Experiment scenarios.

**Scaling to Multiple Users.** Although multiple swimmers rarely drown simultaneously in a swimming pool, we consider such a multi-swimmer detection scenario in AquaHelper. As shown in Fig. 14(a), two chirps may collide if two swimmers send SOS signals simultaneously. As a matter of fact, the chirp-based signal design naturally supports multiple users. In most cases, there is a tiny time offset between the SOS chirps from two swimmers. This time offset can be translated into different FFT peaks after dechirp, as shown in Fig. 14(b). Thus, we can resolve the collision by separating the two different frequency peaks. In addition, we can also assign chirps with different slopes to different swimmers at the entrance counter and register these slopes in the server. These chirps with different slopes are orthogonal and will not affect each other in SOS detection and localization. Thus, AquaHelper can detect different chirps in parallel to handle multiple users.

## 7 Evaluation

**System implementation.** We implemented AquaHelper and tested it with two commercial smartwatches: Huawei Watch 3 and OPPO Watch 3 pro<sup>1</sup>. Four Shuimi SN005 hydrophones (28 USD per unit) were deployed in the water around a swimming pool to receive SOS signals sent by the smartwatches. We use an underwater loudspeaker to transmit synchronization signals. The underwater receivers are connected to Raspberry Pi and forward the data to a laptop for post-processing. As the smartwatch broadcasts SOS signals consecutively, we use a sliding window with a step of 0.5s to detect SOS signals. With a processing time of approximately 65ms per detection window on the laptop, AquaHelper can achieve a maximum detection rate of 15 Hz, which is sufficient for detecting SOS signals.

**Experiment setting.** As shown in Fig. 15(a), the experiments were conducted in a 6-lane 25m × 15m swimming pool. Participants are required to wear smartwatches, which send out SOS signals underwater at different locations. All experiments were carried out with the close monitoring of professional lifeguards and approved by our university authority. In our evaluation, we first positioned a receiver at one of the short edges of the pool to qualitatively assess the impact of various factors on SOS detection performance across various settings, including different distances, depths, orientations, and interference scenarios. Then, we strategically positioned four




Figure 16: Overall performance of AquaHelper.

receivers at the midpoints of all four swimming pool edges at a consistent depth of 0.5m. The signals from all four microphones were employed for the collective performance and localization assessment. In addition, we conducted an experiment in a lake (Fig. 15(c)) to evaluate the maximum supporting distance of AquaHelper.

**Evaluation Method.** We evaluate two key metrics for the SOS detection performance: recall and specificity.  $Recall = \frac{TP}{TP+FN}$ , which quantifies the system performance of detecting SOS signals.  $Specificity = \frac{TN}{FP+TN}$  quantifies the ability to tolerate false alarms. TP, FN, TN, and FP are true positives, false negatives, true negatives, and false positives, respectively. In our target application scenario, higher recall and specificity indicate better performance. We empirically set the detection threshold  $\beta$  as 5 to strike a balance between recall and specificity. We maintain a consistent detection criterion throughout the evaluation. Thus, we present recall in most experiment results, as they consistently exhibit similar levels of specificity, as shown in Sec. 7.3.

### 7.1 Overall Performance

As shown in Fig. 15(a), we first evaluate AquaHelper in a real-world swimming pool environment during regular operating hours. To mimic potential disruptions, we engaged several volunteers to swim and make splashes in close proximity to the receiver. Meanwhile, one volunteer wearing a smartwatch simulated frantic drowning movements at varying distances (5m to 25m) from the receiver to evaluate AquaHelper's performance in a realistic drowning scenario, as depicted in Fig. 15(b). During the experiment, we had non-participant swimmers (mainly university students and faculty members) swim freely in adjacent lanes in the swimming pool. All experiments were closely monitored by professional lifeguards nearby. The evaluation consisted of two sessions: one with the

<sup>1</sup>We attempted to implement AquaHelper on an Apple Watch 6, but we found that the Apple Watch automatically interrupts the audio playback when submerged in water.






Figure 17: Signal illustration. Figure 18: Recall in the lake.

smartwatch emitting SOS signals to assess recall and another with the smartwatch turned off to evaluate specificity. We collected 50-min recordings in total.

The result of recall is shown in Fig. 16(a). AquaHelper achieves the highest detection performance at a distance of 5m, with a recall of 96.9%. This detection gradually decreases with increasing distances. At the maximum distance of 25m, the recall remains 87.4%. When employing a single window for SOS detection, the overall recall rate is 93.1%. This performance degradation can be attributed to two main factors. Firstly, SOS signals significantly attenuate as the distance increases. Secondly, interference from other swimmers and the user's frantic drowning movements influence the SOS detection. Fig. 17 illustrates a received signal clip. We can observe that the drowning movements cause strong splashes. Although most energy is distributed in the low-frequency range, it can still increase the overall noise floor. Notably, the user's arm movements intermittently submerge and raise the smartwatch in and out of the water. As shown in the dashed box of Fig. 17, when the smartwatch is out of the water surface, it results in fragmented chirps. This, in turn, reduces the energy peak of SOS signals after dechirping. An interesting observation is that human speech and ambient noise have a minimal impact on SOS signals, particularly when the receiver is positioned at a greater depth below the water surface. By expanding the size of detection windows to 3, the overall detection recall increases to 96.7%. This improvement is mainly because of the aggregation of more chirps. Similarly, utilizing a detection window size of 5 further enhances the detection recall to 98.3%.

We also evaluate the specificity with the smartwatch turned off, and the results are presented in Figure 16(b). We can observe that AquaHelper achieves a specificity rate of 97.8% at a close distance of 5m, as the user's frantic splashes at such close proximity could potentially trigger false positive cases. As the distance between the user and the receiver increased, the specificity gradually improved, reaching 99.5% at a distance of 25m. At greater distances, the influence of the user's movements became negligible, and false positives primarily originated from splashes generated by other swimmers when they approached the receiver, particularly near the pool wall. By utilizing a larger window size (e.g., 3), the energy of interference was effectively dispersed, leading to an increased system specificity of 99.8%. In Sec. 7.3, we present more evaluation results with distances less than 5m.

In summary, these evaluation results demonstrate the effectiveness of AquaHelper in achieving a high level of detection performance in practical swimming pool environments. To further enhance its robustness, we can strategically deploy more receivers around the pool to alleviate the impact of signal attenuation.




Figure 19: Performance under interference.

## 7.2 Effect of Distances

In the previous section, we evaluated AquaHelper's performance at different distances, revealing that the system's effectiveness diminishes with increasing distance due to energy attenuation. To evaluate AquaHelper's maximum operational range, we conducted experiments in a lake, as illustrated in Fig. 15(c). Fig. 15(d) showcases the five distinct test locations within the lake. During these tests, both the smartwatch and the receiver were submerged in water using ropes. We calculated the detection performance based on a single detection window, and the results are presented in Fig. 18. At a distance of 58m (Loc 1), AquaHelper demonstrated a detection recall of 97.1%, which only experienced a slight decrease to 95.8% at 107m (Loc 2). Even at a distance of 128m (Loc 3), AquaHelper maintained a commendable detection performance of 94.1%. Notably, these results surpassed those obtained in the swimming pool scenario, as we observed significantly less noise and interference in the lake environment. To further challenge AquaHelper and investigate its limits, we conducted additional measurements at two more distant locations, situated at distances of 137m (Loc 4) and 144m (Loc 5). At 137m, the detection recall experienced a significant drop to 21.84%. This decline primarily resulted from the weakened SOS signals at this considerable distance, which struggled to surpass the system's detection threshold. At a distance of 144m, the receiver could no longer detect the SOS signal. In light of these results, our findings suggest that AquaHelper effectively detects SOS signals in open water environments at distances of up to 128m. This extensive range adequately covers most application scenarios of AquaHelper (i.e., swimming pools).

## 7.3 Effect of Interference

We also conducted experiments to evaluate the system's robustness under various forms of interference.

**Splash of swimmers on SOS detection.** We first evaluated the impact of splashes on SOS signal detection. To this end, we placed a smartwatch transmitting SOS signals underwater at 10m from a receiver, while a swimmer (not wearing the smartwatch) swims and splashes water at a distance varying from 1m to 5m from the receiver. As illustrated in Fig. 19(a), AquaHelper achieves 100% detection recall when the swimmer is 5m away, which indicates that the impact of splash interference was negligible. When this swimmer moves to 1m, the detection recall reduces to 99.5% owing to strong splashing interference to the receiver. By extending the detection window size to three, the performance improves to 100%. This result indicates that AquaHelper can detect SOS signals under typical splash interference. We also observe some weak human




Figure 20: Performance with different SNR levels.

voice in the received signal, which can be mitigated by the multi-scale detection of AquaHelper as well.

**Blockage of Line-of-sight path.** Human bodies in the swimming pool may block the LoS path between the transmitter and receiver. Therefore, we asked a swimmer to stay at 1m away from a receiver to intentionally block the direct propagation path and evaluate the performance in non-LoS scenarios. The result in Fig. 19(a) (Occlusion-1m) shows that AquaHelper can successfully detect all SOS chirps. This is because low-frequency SOS signals with diffraction can bypass human bodies. Thus, the blockage of LoS paths has a minimal influence on SOS detection.

**Splash of swimmers on system specificity.** We also evaluated the impact of interference on specificity, *i.e.*, to what extent the splash can cause false alarms. Similar to the experiment before, the swimmer swims and splashes at 1m and 5m in front of the receiver, while the smartwatch does not transmit signals in this experiment. The result is shown in Fig. 19(b). AquaHelper has a 100% specificity when no swimmer is present since normal underwater noise can be removed by the lowpass filter. However, performance decreases when there is a swimmer nearby. Specifically, when the distance between the swimmer and the receiver is 5m, the specificity decreases to 97.8%. This value further drops to 96.6% as the distance reduces to 1m. This is because the intense splash can produce high-energy noise resulting in false alarms if we only monitor the signal within one window. When increasing the window size to three, the energy of splash noise is dispersed, and the specificity improves to 98.3% and 100% at 1m and 5m, respectively. By combining five detection windows, the specificity at 5m also increases to 100%. Except for strong interference, another reason for false alarms is that we use a relatively low threshold and trade the specificity for a higher detection recall due to the utmost importance of saving a life. We believe occasional false alarms are acceptable in practice as long as we do not miss any potential SOS signals. In real-world scenarios, we can adjust the detection window size and the detection threshold to balance the detection recall and specificity. It is also possible to automatically calibrate them periodically by transmitting a known number of pseudo-SOS signals based on a constant false alarm rate for different swimming pools and environmental conditions.

#### 7.4 Performance v.s. SNR

Figure 20(a) shows the SNR distribution of collected SOS signal samples. The SNR is defined as the signal peak to the sum of the remaining noise after dechirp. We can see that 90% of chirps have an SNR higher than -15.4 dB, which can be handled by our chirp-based SOS detection. To test the performance of AquaHelper in even

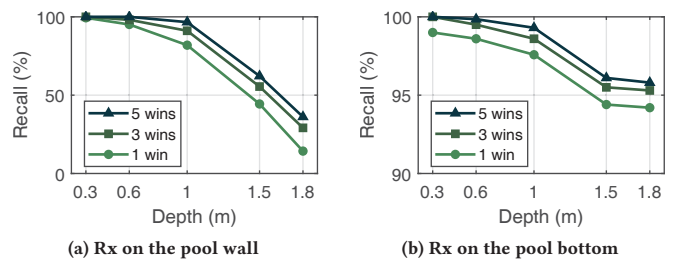



Figure 21: Effect of depths in the water.

more challenging scenarios, we conducted a trace-based experiment with data from different SNR levels. In particular, we collected SOS chirps at a distance of 1m from the transmitter in the swimming pool as clear traces. Subsequently, Gaussian noise was introduced at various levels to generate the required SNR. The result is shown in Fig. 20(b). We can observe that the performance of AquaHelper increases along with SNRs, and can achieve 100% detection recall at -16 dB. Using three windows for dechirp, AquaHelper can reliably detect the SOS chirps at -20 dB. Aggregating five detection windows further decreases the minimum SNR requirement to -22 dB. This result demonstrates that AquaHelper can successfully detect SOS signals with a low SNR, and multi-window aggregation can further improve its performance in ultra-low SNR scenarios.


#### 7.5 Effect of Depths

We conducted an evaluation of AquaHelper at varying depths with a distance of 15m from the receiver, and the results are illustrated in Fig. 21. At a shallow depth of 0.3m, the detection recall reached 99.38%. However, as the depth increased to 0.6m and 1m, the detection performance gradually declined to 95.12% and 81.9%, respectively. A substantial drop in recall rate occurred at a depth of 1.5m, decreasing to 44.37%. When the smartwatch was submerged at the bottom of the pool (1.8m), the detection recall further decreased to 14.3%. This decline can be attributed to the fact that commercial smartwatch speakers are designed for use in the air, and as the water depth increases, so does the water pressure, hindering speaker vibration and SOS transmission. In this case, we enlarged the detection window size to three, which improves detection recalls to 98.1% at 0.6m and 91.1% at 1m, respectively. By aggregating five windows, AquaHelper achieved a recall rate of 97.9% at a depth of 1m. This result suggests that AquaHelper can effectively detect SOS signals using a single window during the early stages of drowning. As the depth increases, AquaHelper needs more detection windows to accumulate and detect weaker SOS signals. Nevertheless, the performance after multi-window aggregation remained unsatisfactory for deeper locations (>1m). This observation underscores the hardware limitations inherent in commercial smartwatches. Yet, it is worth mentioning that some companies are introducing smartwatches specifically designed for underwater activities, such as Garmin Descent Watch [11] and Apple Watch Ultra 2 [8]. These advancements hold promise for achieving improved performance in deeper waters.


A cost-effective strategy to address this depth impact is deploying multiple receivers at the pool bottom. To this end, we repeat the experiment with the receiver positioned at the pool bottom,




**Figure 22: Baseline comparison.**



**Figure 23: Effect of orientations.**



**Figure 24: Contribution of different components.**



**Figure 25: Location estimation performance.**

and the result is displayed in Fig. 21(b). In this case, the detection recall at depths of 1m and 1.5m exhibited substantial improvements, reaching 97.5% and 94.4%, respectively—significantly higher than the performance observed previously. This improvement is mainly because the distance between the transmitter and receiver is dramatically reduced. When we applied a window size of three and five, the detection recalls were further enhanced to 95.5% and 96.1% at 1.5m, respectively. At a depth of 1.8m (*i.e.*, the bottom), AquaHelper still maintains comparable performance due to the closer distance. As hydrophones become affordable (28 USD per unit), we believe it will be economical to deploy a few more at the bottom of the swimming pool to alleviate the depth impact.

## 7.6 Baseline Comparison

We compare our work with AquaApp [22], which modulates SOS signals using Frequency Shift Keying (FSK). AquaApp builds a generic underwater communication channel for mobile devices and is not optimized for SOS transmission in low-SNR scenarios. In addition, it requires a two-way channel estimation to determine the optimal channel frequency. The two-way channel estimation, however, is not feasible over a long range as the low SNR may cause the excitation signal to be lost during transmission. Therefore, we use three uniformly distributed frequencies in the frequency band (1.5, 2.5, and 3.5 kHz) in our experiments, where the transmitter is 25 m away from the receiver. Fig. 22 shows the detection performance comparison between AquaHelper and AquaApp. The chirp-based design of AquaHelper achieves the highest detection recall (99.5%) due to its energy accumulation property and resilience to channel-selective fading. While low-frequency signals experience less attenuation, the part of underwater noise at 1.5 kHz results in a low detection recall (91.8%). The performance improves to 93.8% at a frequency of 2.5 kHz due to less noise. However, the detection recall decreases to 71.6% at a frequency of 3.5 kHz. Despite less noise at this frequency, the signal suffers from severe attenuation underwater. As a result, the detection of weak acoustic signals becomes more challenging. In comparison, benefiting from a series of optimization (*e.g.*, unique chirp signals, high-order chirp aggregation), our work AquaHelper can achieve almost 100% in the same experiment settings.

## 7.7 Effect of Orientations

We conducted orientation evaluations with a distance of 15m between the transmitter and receiver. Five orientations were tested, as illustrated in Fig. 23. The detection recall is 99.8% and 99.5% when the smartwatch speaker is facing toward the receiver (0°)

or horizontally perpendicular to the receiver (horizontal 90°), respectively. As the speaker turns to horizontal 180°, which indicates that the speaker faces backward to the receiver, the performance decreases slightly to 98.5%. This is because the sound of the smartwatch speaker attenuates significantly in the opposite direction. When the smartwatch is facing toward the bottom of the swimming pool (vertical 90°), the detection recall is 98.9%. This is mainly because the signal power is radiated to the deep bottom of the pool and attenuated more. When the smartwatch faces toward surfaces (vertical -90°), the performance increases to 99.4%. Our results show that the orientation of the smartwatch speaker can affect the detection performance because of its directional nature. To achieve full coverage, more receivers can be deployed around the pool wall and bottom.

## 7.8 Contribution of Components

We use the Accumulated SNR (ASNR) [62] to quantify the detectable strength of SOS signals, which is defined as the power of the signal peak over the noise floor after the dechirp process. To evaluate the contribution of different components of AquaHelper, we measure the ASNR gain of each component in Fig. 24. The baseline for comparison is the ASNR obtained from standard dechirp with a single window. Multiple window aggregation demonstrates the highest ASNR gain, with a boost of 2.9 dB and 4.6 dB for two and three windows, respectively. This gain can be attributed to the proportionally increasing signal power after combining multiple windows. The coherent combination of two receivers results in a lower gain of 1.4 dB due to different channel losses for each receiver. The spectral high-order chirp contributes the lowest gain of 0.2 dB because the high-order chirp with weak power is susceptible to noise interference. Note that the results of receiver combining and high-order chirp incorporation are obtained using only one window and can be further increased by aggregating more windows. Overall, AquaHelper effectively improves the ASNR, maximizing the likelihood of SOS detection and facilitating early rescue efforts. By leveraging the diversity of chirps in the frequency, temporal, and spatial domains, the weak SOS signal can be detected as long as we accumulate chirps with enough windows and receivers.

## 7.9 Location Estimation

Our SOS visualization can be extended for swimmer localization by finding the position with the highest likelihood. We pre-marked 20 positions uniformly distributed in the pool as the ground truth, and users were asked to wear the smartwatch and transmit five



SOS chirps at each location. We compare AquaHelper with a TDoA-based localization baseline [17]. Figure 25 shows the CDF of the localization error of two methods, which indicates the Euclidean distance between the ground truth and estimated 2D locations. The median localization error is 1.36m, outperforming the baseline by 2.01m. It is because the baseline suffers from multipath interference. Considering the large size of a swimming pool, this result indicates that AquaHelper is capable of achieving a satisfactory level of localization accuracy, sufficient to assist lifeguards in locating drowning swimmers more efficiently. Given the low cost of hydrophones, we can deploy more receivers to achieve better performance. Thus, AquaHelper has the potential to serve as a building component for various applications in swimming pools. For instance, it can be used to track kids or seniors in water parks and evaluate the performance of amateur swimmers. We also believe our method can be extended to support 3D localization by deploying hydrophones at different depths, which will be explored in the future.

## 8 Related Work

**Underwater Communication and Localization.** While the field of underwater acoustic communication and localization has witnessed a substantial body of work [16, 20, 24, 32, 41, 51, 52, 54, 58], most of them are designed for specialized high-power underwater communication equipment. Recent developments have seen the utilization of specialized piezoelectric transducers to achieve low-power underwater communication and localization [14, 27, 31, 37], enabling numerous applications, including Underwater GPS [32], underwater imaging [13], and low-power machine learning underwater [66]. Unlike the existing works, our research takes on the distinct challenges brought about by the integration of lightweight smart mobile devices, including limited bandwidth, low transmission power, and noise interference in practical swimming pools. To this end, we develop AquaHelper, a system capable of transmitting acoustic SOS signals using compact wearable devices and detecting them even in extremely low-SNR scenarios. We accomplish this through a series of novel techniques, including incorporating high-order harmonics, multi-scale window aggregation, and multi-receiver synchronization and coherent combining.

While wireless signals have seen widespread use for communication [57, 62] and ubiquitous sensing [26, 33, 38] in the air, they encounter significant attenuation when propagating through water. In contrast, acoustic signals demonstrate superior propagation characteristics underwater [23]. In the research community, acoustic sensors in mobile devices have garnered attention, particularly in the context of human-computer interaction [25, 59, 61, 67]. In recent years, efforts have also been made to integrate underwater communication and localization with mobile devices. AquaApp [22] designs the first underwater messaging system compatible with commercial mobile devices, which supports a communication range of up to 100m with a smartphone. AquaRanger [21] and [23] can measure the range between two smartphones and locate their relative positions underwater. However, these works require smartphones with considerably higher transmission power compared to commercial smartwatches. Furthermore, they rely on a send-and-reply process for channel estimation and distance measurement. This operation is impractical for wearable devices

used in swimming pools due to their low reception sensitivity and urgent response requirements. In contrast, AquaHelper synchronizes different receivers through one-way transmission, making the transmitter transparent and minimizing the SOS response time.

**Drowning Detection.** Vision-based drowning detection has been extensively studied by industry [1, 4, 5, 10] and academia [2, 36, 44, 65]. Most of them detect the drowning incidents by examining if a person is stationary for long time in the water. Some machine learning-based works [29, 30, 34, 35, 45] extract motion features such as speed, posture, and limb movement to build a supervised classifier and recognize drowning events. However, it is challenging to achieve accurate and reliable detection due to the inherently complex drowning behaviour and lack of training samples [19]. Moreover, vision-based methods suffer from poor lighting conditions, highly dynamic backgrounds, and dim visibility of targets in the water [29]. As a complementary approach, AquaHelper allows swimmers to actively call for help in case of an emergency such as severe discomfort, enabling the rescue in the early stage.

Wearable sensor-based solutions [7, 9, 46, 49, 50] are more resilient to lighting conditions. Some works [40, 42, 46, 60] monitor the sensor readings of human motion or vital signs (e.g., oxygen level, respiration) to detect drowning events and then inflate an airbag. Although these systems can detect the drowning risk, they require swimmers to wear heavy life kits and cannot alert lifeguards for additional help, which may miss the best rescue time. AquaHelper retrofits these systems with capabilities to transmit SOS signals to lifeguards, and also provides necessary means to actively call for help in an emergency. In this paper, we focus on SOS transmission and detection, and leave the integration with drowning event detection for future work, which is important yet orthogonal to this research.

## 9 Conclusion

Drowning incidents in swimming pools represent a pressing public safety concern around the world, with devastating consequences for individuals and families each year. In response to this issue, we propose AquaHelper, the first underwater SOS transmission and detection system that can work in extremely low-SNR scenarios. By leveraging lightweight wearable devices (e.g., water-proof smartwatches) as transmitters and multiple low-cost underwater acoustic receivers, AquaHelper can alert lifeguards promptly to significantly increase the chances of survival for drowning swimmers. Our experiments demonstrate the effectiveness of AquaHelper in detecting SOS signals with high accuracy in swimming pools. As smartwatches are becoming increasingly popular for swimmers, we expect such a life-saving feature will become indispensable and can potentially save hundreds of lives worldwide every year in the near future.

## Acknowledgments

We thank the anonymous shepherd and reviewers for their valuable comments and feedback. We thank Leming Shen and Shiming Yu for their assistance in the experiments. We also thank the Student Affairs Office of PolyU for providing the experimental facilities. This work is supported by Hong Kong GRF under grant 15206123. Yuanqing Zheng is the corresponding author.

## References

- [1] 2022. About Coral Manta - Coral Drowning Detection Systems. <https://coraldrowningdetection.com/about-coral-manta/>. (Accessed on 08/21/2022).
- [2] 2022. DROWNING DETECTION | hkusail. <https://www.sail.hku.hk/drowning-detection>. (Accessed on 08/21/2022).
- [3] 2022. Facts & Stats About Drowning - Stop Drowning Now. <https://www.stopdrowningnow.org/drowning-statistics/>. (Accessed on 08/20/2022).
- [4] 2022. Poseidon - Drowning detection system for swimming pools. <https://poseidon-tech.com/en-GB/technology-2/>. (Accessed on 08/21/2022).
- [5] 2022. SwimEye - a drowning detection system for swimming pools. <https://swimeye.com/>. (Accessed on 08/21/2022).
- [6] 2022. Swimming Injury Statistics - Swimming Pool Accidents. <https://www.edgarsnyder.com/statistics/swimming-pool-statistics.html>. (Accessed on 08/20/2022).
- [7] 2022. w10 | WAVE Drowning Detection Systems | United States. <https://www.wavedds.com/w10>. (Accessed on 08/21/2022).
- [8] 2023. Apple Watch Ultra 2 - Apple (HK). <https://www.apple.com/hk/en/apple-watch-ultra-2/>. (Accessed on 09/28/2023).
- [9] 2023. Drowning Detection System from Sentag - Sentag. <https://www.sentag.com/>. (Accessed on 01/11/2023).
- [10] 2023. Drowning Detection Systems Technology - AngelEye. <https://www.angeleye.tech/en/en-technology/>. (Accessed on 01/11/2023).
- [11] 2023. Garmin Descent™ MK2i | Dive Computer | Dive Smartwatch. <https://www.garmin.com/en-US/p/632320>. (Accessed on 09/28/2023).
- [12] 2023. A short history of the 10:20 protection standard. <https://www.thewirh.com/blog/10-20>. (Accessed on 03/17/2023).
- [13] Sayed Saad Afzal, Waleed Akbar, Osmy Rodriguez, Mario Doumet, Unsoo Ha, Reza Ghaffarivardavagh, and Fadel Adib. 2022. Battery-free wireless imaging of underwater environments. *Nature communications* 13, 1 (2022), 5546.
- [14] Sayed Saad Afzal, Reza Ghaffarivardavagh, Waleed Akbar, Osmy Rodriguez, and Fadel Adib. 2020. Enabling higher-order modulation for underwater backscatter communication. In *Global Oceans 2020: Singapore-US Gulf Coast*. IEEE, 1–6.
- [15] Zhenlin An, Qiongzhen Lin, and Lei Yang. 2018. Cross-frequency communication: Near-field identification of uhf rfids with wifi. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*. 623–638.
- [16] Prasad Anjanti, Amy Gibson, Manu Ignatius, Chinmay Pendharkar, Anne Kurian, Alan Low, and Mandar Chitre. 2020. Diver Communication and Localization System using Underwater Acoustics. In *Global Oceans 2020: Singapore-US Gulf Coast*. IEEE, 1–8.
- [17] Tao Bian, Ramachandran Venkatesan, and Cheng Li. 2009. Design and evaluation of a new localization scheme for underwater acoustic sensor networks. In *GLOBECOM 2009-2009 IEEE Global Telecommunications Conference*. IEEE, 1–5.
- [18] Chao Cai, Rong Zheng, and Jun Luo. 2022. Ubiquitous acoustic sensing on commodity iot devices: A survey. *IEEE Communications Surveys & Tutorials* 24, 1 (2022), 432–454.
- [19] Aida Carballo-Fazanes, Joost JLM Bierens, and International Expert Group to Study Drowning Behaviour. 2020. The visible behaviour of drowning persons: A pilot observational study using analytic software and a nominal group technique. *International journal of environmental research and public health* 17, 18 (2020), 6930.
- [20] Thomas Casey, Brian Guimond, and James Hu. 2007. Underwater vehicle positioning based on time of arrival measurements from a single beacon. In *OCEANS 2007*. IEEE, 1–8.
- [21] Tuochao Chen, Justin Chan, and Shyamnath Gollakota. 2022. Underwater Acoustic Ranging Between Smartphones. *arXiv preprint arXiv:2209.01780* (2022).
- [22] Tuochao Chen, Justin Chan, and Shyamnath Gollakota. 2022. Underwater messaging using mobile devices. In *Proceedings of the ACM SIGCOMM 2022 Conference*. 545–559.
- [23] Tuochao Chen, Justin Chan, and Shyamnath Gollakota. 2023. Underwater 3D positioning on smart devices. In *Proceedings of the ACM SIGCOMM 2023 Conference*. 33–48.
- [24] En Cheng, Xizhou Lin, Shengli Chen, and Fei Yuan. 2016. A TDofA localization scheme for underwater sensor networks with use of multilinear chirp signals. *Mobile Information Systems* 2016 (2016).
- [25] Haiming Cheng and Wei Lou. 2021. Push the limit of device-free acoustic sensing on commercial mobile devices. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*. IEEE, 1–10.
- [26] Kaiyan Cui, Qiang Yang, Yuanqing Zheng, and Jinsong Han. 2023. mmRipple: Communicating with mmWave Radars through Smartphone Vibration. In *Proceedings of the 22nd International Conference on Information Processing in Sensor Networks*. 149–162.
- [27] Aline Eid, Jack Rademacher, Waleed Akbar, Purui Wang, Ahmed Allam, and Fadel Adib. 2023. Enabling Long-Range Underwater Backscatter via Van Atta Acoustic Networks. In *Proceedings of the ACM SIGCOMM 2023 Conference*. 1–19.
- [28] Jeff Ellis. 2002. Lifeguards May Look But They Don't Always See: A study points to causes contributing to more than 400 deaths annually in U.S. public swimming facilities. <https://www.lib.niu.edu/2002/ip020538.html>. (Accessed on 08/21/2022).
- [29] How-Lung Eng, Kar-Ann Toh, Alvin H Kam, Junxian Wang, and Wei-Yun Yau. 2003. An automatic drowning detection surveillance system for challenging outdoor pool environments. In *Computer Vision, IEEE International Conference on*, Vol. 2. IEEE Computer Society, 532–532.
- [30] How-Lung Eng, Kar-Ann Toh, Wei-Yun Yau, and Junxian Wang. 2008. DEWS: A live visual surveillance system for early drowning detection at pool. *IEEE transactions on circuits and systems for video technology* 18, 2 (2008), 196–210.
- [31] Reza Ghaffarivardavagh, Sayed Saad Afzal, Osmy Rodriguez, and Fadel Adib. 2020. Ultra-wideband underwater backscatter via piezoelectric metamaterials. In *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication*. 722–734.
- [32] Reza Ghaffarivardavagh, Sayed Saad Afzal, Osmy Rodriguez, and Fadel Adib. 2020. Underwater backscatter localization: Toward a battery-free underwater GPS. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks*. 125–131.
- [33] Mingda Han, Huanqi Yang, Tao Ni, Di Duan, Mengzhe Ruan, Yongliang Chen, Jia Zhang, and Weitao Xu. 2023. mmSign: mmWave-based Few-Shot Online Handwritten Signature Verification. *ACM Transactions on Sensor Networks* (2023).
- [34] Saifeldin Hasan, John Joy, Fardin Ahsan, Huzaifa Khambaty, Manan Agarwal, and Jinane Mounsef. 2021. A Water Behavior Dataset for an Image-Based Drowning Solution. In *2021 IEEE Green Energy and Smart Systems Conference (IGESSC)*. IEEE, 1–5.
- [35] Lixing He, Haozheng Hou, Zhenyu Yan, and Guoliang Xing. 2022. Demo Abstract: An Underwater Sonar-Based Drowning Detection System. In *2022 21st ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, 493–494.
- [36] Xinyu He, Fei Yuan, and Yi Zhu. 2021. Drowning Detection Based on Video Anomaly Detection. In *Image and Graphics: 11th International Conference, ICIG 2021, Haikou, China, August 6–8, 2021, Proceedings, Part III 11*. Springer, 700–711.
- [37] Junsu Jang and Fadel Adib. 2019. Underwater backscatter networking. In *Proceedings of the ACM Special Interest Group on Data Communication*. 187–199.
- [38] Sijie Ji, Yaxiong Xie, and Mo Li. 2022. SiFall: Practical Online Fall Detection with RF Sensing. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems*. 563–577.
- [39] Soonwon Ka, Tae Hyun Kim, Jae Yeol Ha, Sun Hong Lim, Su Cheol Shin, Jun Won Choi, Chulyoung Kwak, and Sunghyun Choi. 2016. Near-ultrasound communication for tv's 2nd screen services. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*. 42–54.
- [40] Mohamed Kharrat, Yuki Wakuda, Noboru Koshizuka, and Ken Sakamura. 2013. Automatic waist airbag drowning prevention system based on underwater time-lapse and motion information measured by smartphone's pressure sensor and accelerometer. In *2013 IEEE International Conference on Consumer Electronics (ICCE)*. IEEE, 270–273.
- [41] Benjamin Kuch, Giorgio Buttazzo, Elaine Azzopardi, Martin Sayer, and Arne Sieber. 2012. GPS diving computer for underwater tracking and mapping. *Underwater Technology* 30, 4 (2012), 189–194.
- [42] Aboli Kulkarni, Kshitij Lakhani, and Shubham Lokhande. 2016. A sensor based low cost drowning detection system for human life safety. In *2016 5th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions)(ICRITO)*. IEEE, 301–306.
- [43] Patrick Lazik, Niranjini Rajagopal, Oliver Shih, Bruno Sinopoli, and Anthony Rowe. 2015. ALPS: A bluetooth and ultrasound platform for mapping and localization. In *Proceedings of the 13th ACM conference on embedded networked sensor systems*. 73–84.
- [44] Fei Lei, Hengyu Zhu, Feifei Tang, and Xinyuan Wang. 2022. Drowning behavior detection in swimming pool based on deep learning. *Signal, Image and Video Processing* (2022), 1–8.
- [45] Wenmiao Lu and Yap-Peng Tan. 2004. A vision-based approach to early detection of drowning incidents in swimming pools. *IEEE transactions on circuits and systems for video technology* 14, 2 (2004), 159–178.
- [46] Bello Kontagora Nuhu, Buhari Ugbede Umar, Adishetu Khadijat Abu, Guda Blessed, and Kim Jinsul. 2022. Development of a Smart Wearable Antidrowning System for Swimmers. (2022).
- [47] Chunyi Peng, Guobin Shen, Yongguang Zhang, Yanlin Li, and Kun Tan. 2007. Beepbeep: a high accuracy acoustic ranging system using cots mobile devices. In *Proceedings of the 5th international conference on Embedded networked sensor systems*. 1–14.
- [48] Ramon Pinyol. 2015. Harmonics: Causes, effects and minimization. In *Salicru white papers*.
- [49] J Geetha Ramani, J Gayathri, R Aswanth, and M Gunasekaran. 2019. Automatic prevention of drowning by inflatable wrist band system. In *2019 5th International Conference on Advanced Computing & Communication Systems (ICACCS)*. IEEE, 346–349.
- [50] Muhammad Ramdhan, Muhammad Ali, Samura Ali, MY Kamaludin, et al. 2018. An early drowning detection system for Internet of Things (IoT) applications. *TELKOMNIKA (Telecommunication Computing Electronics and Control)* 16, 4 (2018), 1870–1876.

- [51] Lynn T Rauchenstein, Abhinav Vishnu, Xinya Li, and Zhiqun Daniel Deng. 2018. Improving underwater localization accuracy with machine learning. *Review of Scientific Instruments* 89, 7 (2018), 074902.
- [52] Marwane Rezzouki, Mohamed Amine Ben Temim, and Guillaume Ferré. 2021. Differential chirp spread spectrum to perform acoustic long range underwater localization and communication. In *OCEANS 2021: San Diego-Porto*. IEEE, 1–9.
- [53] Nirupam Roy, Haitham Hassanieh, and Romit Roy Choudhury. 2017. Backdoor: Making microphones hear inaudible sounds. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*. 2–14.
- [54] DJ Schott, M Faisal, F Hoeflinger, LM Reindl, J Bordoy Andreu, and C Schindelbauer. 2017. Underwater localization utilizing a modified acoustic indoor tracking system. In *2017 IEEE 7th International Conference on Underwater System Technology: Theory and Applications (USYS)*. IEEE, 1–5.
- [55] Abdelaziz M Shehata, Eslam M Mohamed, Khaled L Salem, Ahmed M Mohamed, Mustafa Abdul Salam, and Mennatullah M Gamil. 2021. A Survey of Drowning Detection Techniques. In *2021 International Mobile, Intelligent, and Ubiquitous Computing Conference (MIUCC)*. IEEE, 286–290.
- [56] Fabian Steinmetz, Jan Heitmann, and Christian Renner. 2018. A practical guide to chirp spread spectrum for acoustic underwater communication in shallow waters. In *Proceedings of the 13th International Conference on Underwater Networks & Systems*. 1–8.
- [57] Zehua Sun, Huanqi Yang, Kai Liu, Zhimeng Yin, Zhenjiang Li, and Weitao Xu. 2022. Recent advances in LoRa: A comprehensive survey. *ACM Transactions on Sensor Networks* 18, 4 (2022), 1–44.
- [58] Francesco Tonolini and Fadel Adib. 2018. 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*. 117–131.
- [59] Yanwen Wang, Jiaxing Shen, and Yuanqing Zheng. 2020. Push the limit of acoustic gesture recognition. *IEEE Transactions on Mobile Computing* (2020).
- [60] Jie Wen, Dan Zhou, Haoran Feng, Yongcai Wang, Xiongfei Geng, Hengzhe Ma, and Zongwei Yang. 2019. LifeTag: Vital Sign Detection for Drowning People in Sea Accidents by Wearable Device. In *Proceedings of the 2019 8th International Conference on Networks, Communication and Computing*. 57–64.
- [61] Kaishun Wu, Qiang Yang, Baojie Yuan, Yongpan Zou, Rukhsana Ruby, and Mo Li. 2020. Echowrite: An acoustic-based finger input system without training. *IEEE Transactions on Mobile Computing* 20, 5 (2020), 1789–1803.
- [62] Xianjin Xia, Qianwu Chen, Ningning Hou, and Yuanqing Zheng. 2022. HyLink: Towards High Throughput LPWANs with LoRa Compatible Communication. (2022).
- [63] Xianjin Xia, Ningning Hou, Yuanqing Zheng, and Tao Gu. 2021. PCube: scaling LoRa concurrent transmissions with reception diversities. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*. 670–683.
- [64] Zhenqiang Xu, Shuai Tong, Pengjin Xie, and Jiliang Wang. 2023. From Demodulation to Decoding: Toward Complete LoRa PHY Understanding and Implementation. *ACM Transactions on Sensor Networks* 18, 4 (2023), 1–27.
- [65] Chi Zhang, Xiaoguang Li, and Fei Lei. 2015. A novel camera-based drowning detection algorithm. In *Chinese Conference on Image and Graphics Technologies*. Springer, 224–233.
- [66] Yuchen Zhao, Sayed Saad Afzal, Waleed Akbar, Osmy Rodriguez, Fan Mo, David Boyle, Fadel Adib, and Hamed Haddadi. 2022. Towards battery-free machine learning and inference in underwater environments. In *Proceedings of the 23rd Annual International Workshop on Mobile Computing Systems and Applications*. 29–34.
- [67] Yongpan Zou, Qiang Yang, Yetong Han, Dan Wang, Jiannong Cao, and Kaishun Wu. 2019. AcouDigits: Enabling users to input digits in the air. In *2019 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 1–9.