MUTE: Bringing IoT to Noise Cancellation

Sheng Shen, Nirupam Roy, Junfeng Guan, Haitham Hassanieh, Romit Roy Choudhury

University of Illinois at Urbana-Champaign

{sshenn19, nroy8, jguan8, haitham, croy}@illinois.edu

ABSTRACT

Active Noise Cancellation (ANC) is a classical area where noise in the environment is canceled by producing *anti-noise* signals near the human ears (e.g., in Bose’s noise cancellation headphones). This paper brings IoT to active noise cancellation by combining wireless communication with acoustics. The core idea is to place an IoT device in the environment that listens to ambient sounds and forwards the sound over its wireless radio. Since wireless signals travel much faster than sound, our ear-device receives the sound in advance of its actual arrival. This serves as a glimpse into the future, that we call *lookahead*, and proves crucial for real-time noise cancellation, especially for unpredictable, wide-band sounds like music and speech. Using custom IoT hardware, as well as lookahead-aware cancellation algorithms, we demonstrate *MUTE*, a fully functional noise cancellation prototype that outperforms Bose’s latest ANC headphone. Importantly, our design does not need to block the ear – the ear canal remains open, making it comfortable (and healthier) for continuous use.

CCS CONCEPTS

• **Networks** → *Sensor networks*; • **Human-centered computing** → *Ubiquitous and mobile devices*;

KEYWORDS

Noise Cancellation, Acoustics, Internet of Things, Wearables, Edge Computing, Adaptive Filter, Smart Home, Earphone

ACM Reference Format:

Sheng Shen, Nirupam Roy, Junfeng Guan, Haitham Hassanieh, Romit Roy Choudhury. 2018. MUTE: Bringing IoT to Noise Cancellation. In *SIGCOMM ’18: ACM SIGCOMM 2018 Conference, August*

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 ACM 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.

SIGCOMM ’18, August 20–25, 2018, Budapest, Hungary

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5567-4/18/08...\$15.00

<https://doi.org/10.1145/3230543.3230550>

Figure 1: MUTE leverages the difference between wireless and acoustic propagation delay to provide future lookahead into the incoming sound signals.

20–25, 2018, Budapest, Hungary. ACM, New York, NY, USA, 15 pages.
<https://doi.org/10.1145/3230543.3230550>

1 INTRODUCTION

Ambient sound can be a source of interference. Loud conversations or phone calls in office corridors can be disturbing to others around. Working or napping at airports may be difficult due to continuous overhead announcements. In developing regions, the problem is probably most pronounced. Loud music or chants from public speakers, sound pollution from road traffic, or just general urban cacophony can make simple reading or sleeping difficult. The accepted solution has been to wear ear-plugs or ear-blocking headphones, both of which are uncomfortable for continuous use [22, 31, 41]. This paper considers breaking away from convention and aims to cancel complex sounds without blocking the ear. We introduce our idea next with a simple example.

Consider Alice getting disturbed in her office due to frequent corridor conversations (Figure 1). Imagine a small IoT device – equipped with a microphone and wireless radio – pasted on the door in Alice’s office. The IoT device listens to the ambient sounds (via the microphone) and forwards the exact sound waveform over the wireless radio. Now, *given that wireless signals travel much faster than sound*, Alice’s noise cancellation device receives the wireless signal first, extracts the sound waveform from it, and gains a “*future lookahead*” into the actual sound that will arrive later. When the actual sound arrives, Alice’s ear-device is already aware of the signal and has had the time to compute the appropriate *anti-noise* signal. In fact, this lead time opens various other

Figure 2: MUTE's experimental platform: The center figure (b) shows the full system with the wireless IoT relay taped on the room's inside wall and the (crude) ear-device on the table (composed of a microphone on the human head model, an anti-noise speaker, and a DSP board). The left figure (a) shows our vision of the hollow ear-device, not covering the ear. The right figure (c) zooms into the relay hardware.

algorithmic and architectural opportunities, as will become clear in the subsequent discussions.

In contrast, consider today's noise cancellation headphones from Bose [9, 10], SONY [15], Philips [18], etc. These headphones essentially contain a microphone, a DSP processor, and a speaker. The processor's job is to process the sound received by the microphone, compute the anti-noise signal, and play it through the speaker. This sequence of operations starts when the sound has arrived at the microphone, however, *must complete before the same sound has reached the human's ear-drum*. Given the small distance between the headphone and the ear-drum, this is an extremely tight deadline ($\approx 30 \mu\text{s}$ [13]). The penalty of missing this deadline is a phase error, i.e., the anti-noise signal is not a perfect "opposite" of the actual sound, but lags behind. The lag increases at higher frequencies, since phase changes faster at such frequencies. This is one of the key reasons why current headphones are designed to only cancel low-frequency sounds below 1 kHz [5, 46], such as periodic machine noise. For high-frequency signals (e.g., speech and music), the headphones must use sound-absorbing materials. These materials cover the ear tightly and *attenuate* the sounds as best as possible [10, 33].

Meeting the tight deadline is not the only hurdle to real-time noise cancellation. As discussed later, canceling a sound also requires estimating the inverse of the channel from the sound source to the headphone's microphone. Inverse-channel estimation is a *non-causal* operation, requiring access to future sound samples. Since very few future samples are available to today's headphones, the anti-noise signal is not accurate, affecting cancellation quality.

With this background in mind, let us now return to our proposal of forwarding sound over wireless links. The forwarded

sound is available to our cancellation device *several milliseconds* in advance of its physical arrival (as opposed to tens of microseconds in conventional systems). This presents 3 opportunities:

- (1) **Timing:** The DSP processor in our system can complete the anti-noise computation before the deadline, enabling noise cancellation for even higher frequencies. Hence, sound-absorbing materials are not necessary to block the ear.
- (2) **Profiling:** Lookahead allows the DSP processor to foresee macro changes in sound profiles, such as when Bob and Eve are alternating in a conversation. This allows for quicker multiplexing between filtering modes, leading to faster convergence at transitions.
- (3) **Channel Estimation:** Finally, much longer lookahead improves anti-noise computation due to better inverse-channel estimation, improving the core of noise cancellation.

Of course, translating these intuitive opportunities into concrete gains entails challenges. From an algorithmic perspective, the adaptive filtering techniques for classical noise cancellation need to be delicately redesigned to fully harness the advantages of lookahead. From an engineering perspective, the wireless relay needs to be custom-made so that forwarding can be executed in real-time (to maximize lookahead), and without storing any sound samples (to ensure privacy). This paper addresses all these questions through a *lookahead-aware noise cancellation* (LANC) algorithm, followed by a custom-designed IoT transceiver at the 900MHz ISM band. The wireless devices use frequency modulation (FM) to cope with challenges such as carrier frequency offset, non-linearities, and amplitude distortion.

Figure 2(b) shows the overall experimentation platform for our wireless noise cancellation system (MUTE). The custom-designed wireless relay is pasted on the wall, while the

(crude) ear-device is laid out on the table. The ear-device has not been packaged into a wearable form factor, however, is complete in functionality, i.e., it receives the wireless signals from the relay, extracts the audio waveform, and feeds it into a TI TMS320 DSP board running the LANC algorithm. Figure 2(a) visualizes the potential form-factor for such a wearable device (sketched in AutoDesk), while Figure 2(c) zooms into the relay hardware. To compare performance, we insert a “measurement microphone” into the ear position of the human head model – this serves as a virtual human ear. We place Bose’s latest ANC headphone (QC35 [10]) over the head model and compare its cancellation quality against *MUTE*, for different types of sounds, multipath environments, and lookahead times. Finally, we bring in 5 human volunteers to experience and rate the performance difference in noise cancellation. Our results reveal the following:

- *MUTE* achieves cancellation across [0, 4] kHz, while Bose cancels only up to 1 kHz. Within 1 kHz, *MUTE* outperforms Bose by 6.7 dB on average.
- Compared to Bose’s full headphone (i.e., ANC at [0, 1] kHz + sound-absorbing material for [1, 4] kHz), our cancellation is 0.9 dB worse. We view this as a non ear-blocking device with a slight compromise. With ear-blocking, *MUTE* outperforms Bose by 8.9 dB.
- *MUTE* exhibits greater agility for fast changing, intermittent sounds. The average cancellation error is reduced by 3 dB, and human volunteers consistently rate *MUTE* better than Bose for both speech and music.
- Finally, Bose is advantaged with specialized microphones and speakers (with significantly less hardware noise); our systems are built on cheap microphone chips (\$9) and off-the-shelf speakers (\$19). Also, we have designed a mock ear-device to suggest how future earphones need not block the ear (Figure 2(a)). However, we leave the real packaging (and manufacturing) of such a device to future work.

In closing, we make the following contributions:

- Introduce *MUTE*, a wireless noise cancellation system architecture that harnesses the difference in propagation delay between radio frequency (RF) and sound to provide a valuable “lookahead” opportunity for noise cancellation.
- Present a *Lookahead Aware Noise Cancellation* (LANC) algorithm that exploits lookahead for efficient cancellation of unpredictable high frequency signals like human speech. Our prototype compares well with today’s ANC headphones, but does not need to block the user’s ears.

We expand on each of these contributions next, beginning with a brief primer on active noise cancellation (ANC), and followed by our algorithm, architecture, and evaluation.

2 NOISE CANCELLATION PRIMER

An active noise cancellation (ANC) system has at least two microphones and one speaker (see Figure 3). The microphone placed closer to the ear-drum is called the *error microphone* M_e , while the one away from the ear is called the *reference microphone*, M_r . The speaker is positioned close to M_e and is called the *anti-noise speaker*. Ambient noise first arrives at M_r , then at M_e , and finally at the ear-drum. The DSP processor’s goal is to extract the sound from M_r , compute the *anti-noise*, and play it through the speaker such that the *anti-noise* cancels the ambient noise at M_e .

Figure 3: Basic architecture of an ANC headphone, currently designed for a single noise source.

Given that received sound is a combination of current and past sound samples (due to multipath), the DSP processor cannot simply reverse the sound samples from M_r . Instead, the various channels (through which the sound travels) need to be estimated correctly to construct the anti-noise signal. For this, the DSP processor uses the cancellation error from M_e as feedback and updates its channel estimates to converge to a better anti-noise in the next time step. Once converged, cancellation is possible at M_e regardless of the sound sample. So long as the ear-drum is close enough to M_e , the human also experiences similar cancellation as M_e .

■ **The ANC Algorithm:** Figure 4 redraws Figure 3 but from an algorithmic perspective. Observe that the error microphone M_e receives two signals, one directly from the noise source, say $a(t)$, and the other from the headphone’s anti-noise speaker, say $b(t)$. The output of this microphone can be expressed as $e(t) = a(t) + b(t)$. For perfect cancellation, $e(t)$ would be zero.

Now, $a(t)$ can be modeled as $a(t) = h_{ne}(t) * n(t)$, where h_{ne} is the air channel from the noise source to M_e , $n(t)$ is the noise signal, and $*$ denotes convolution. Similarly, $b(t)$ can be modeled as:

$$b(t) = h_{se}(t) * \left(h_{AF}(t) * \left(h_{nr}(t) * n(t) \right) \right) \quad (1)$$

Figure 4: ANC block diagram.

Here, the inner-most parenthesis models the noise signal received by the reference microphone M_r over the channel $h_{nr}(t)$. The ANC algorithm in the DSP processor modifies this signal using an adaptive filter, $h_{AF}(t)$, and plays it through the anti-noise speaker. The speaker's output is distorted by the small gap between the speaker and the error microphone, denoted $h_{se}(t)$. Thus, the error signal $e(t)$ at the output of M_e is

$$e(t) = a(t) + b(t) \\ = h_{ne}(t) * n(t) + h_{se}(t) * \left(h_{AF}(t) * \left(h_{nr}(t) * n(t) \right) \right)$$

For active noise cancellation, the ANC algorithm must design $h_{AF}(t)$ such that $e(t)$ is as close to 0 as possible. This suggests that $h_{AF}(t)$ should be set to:

$$h_{AF}(t) = -h_{se}^{-1}(t) * h_{ne}(t) * h_{nr}^{-1}(t) \quad (2)$$

In other words, ANC must estimate all 3 channels to apply the correct h_{AF} . Fortunately, h_{se}^{-1} can be estimated by sending a known preamble from the anti-noise speaker and measuring the response at the error microphone. However, h_{ne} and h_{nr}^{-1} cannot be easily estimated since: (1) the noise signal $n(t)$ does not exhibit any preamble-like structure, (2) the channels are continuously varying over time, and (3) the inverse channel requires future samples for precise estimation.

To cope with this, ANC uses adaptive filtering to estimate h_{AF} . The high-level idea is gradient descent, i.e., adjusting the values of the vector h_{AF} in the direction in which the residual error $e(t)$ goes down. Thus, ANC takes $e(t)$ as the feedback and feeds the classical Least Mean Squared (LMS) technique [20, 32] – the output is an adaptive filter, $h_{AF}(t)$. With this background, let us now zoom into the lookahead advantage and corresponding design questions.

3 LOOKAHEAD AWARE ANC

MUTE is proposing a simple architectural change to conventional systems, i.e., disaggregate the reference microphone M_r from the headphone, place M_r a few feet away towards

the noise source, and replace the wired connection between M_r and the DSP processor with a wireless (RF) link. This separation significantly increases the lead time (or *lookahead*), translating to advantages in timing and cancellation. We detail the advantages next and then develop the Lookahead Award Noise Cancellation (LANC) algorithm.

3.1 Timing Advantage from Lookahead

Figure 5(a) shows the timeline of operations in today's ANC systems and Figure 5(b) shows the same, but with a large lookahead. Note that time advances in the downward direction with each vertical line corresponding to different components (namely, reference microphone, DSP processor, speaker, etc.) The slanting solid arrow denotes the arrival of the noise signal, while the black dots mark relevant events on the vertical timelines. We begin by tracing the sequence of operations step-by-step in Figure 5(a).

The noise signal first arrives at the headphone's reference microphone at time t_1 . This sample is conveyed via wire and reaches the DSP processor at time t_2 , where $(t_2 - t_1)$ is the ADC (analog-to-digital converter) delay. The DSP processor now computes the anti-noise sample and sends it to the anti-noise speaker at t_3 , which outputs it after a DAC (digital-to-analog converter) and playback delay. Ideally, the speaker should be ready to play the anti-noise at t_4 since the actual sound wave is also passing by the speaker at this time. However, meeting this deadline is difficult since the distance between the reference microphone and speaker is < 1 cm. With sound traveling at 340 m/s, the available time window is $(t_4 - t_1)$, which is around $30 \mu\text{s}$. Since ADC, DSP processing, DAC and speaker delay can easily be $3\times$ more than this time budget, today's ANC systems miss the deadline. Thus, instead of t_4 , the anti-noise gets played at a later time t_6 , as shown by the red dashed line in Figure 5(a).

For low frequencies, this can still deliver partial noise cancellation, since the phases of the noise and anti-noise would be slightly misaligned. However, for higher frequencies (i.e., smaller wavelengths), the performance would degrade since the excess delay (past t_4) would cause larger phase misalignment. This is the core struggle in today's noise cancellation systems.

Figure 5(b) illustrates how *MUTE* naturally relieves this time pressure. By virtue of being further away, the reference microphone captures the noise signal earlier and forwards it over wireless (as shown by the horizontal dashed arrow at time t_1). The lookahead is far greater now, offering adequate time to subsume the ADC, DSP, DAC, and speaker delays. Hence, *MUTE* can compute the anti-noise sample and be ready to play it exactly when the actual noise arrives at the speaker at t_6 . The anti-noise now coincides with the noise, as

Figure 5: Global timeline with (a) limited lookahead and (b) large lookahead. Time advances in the downward direction, and the slanted arrows denote the sound samples arriving from a noise source to the human ear. With large lookahead in (b), MUTE has adequate time to subsume all delays and play the anti-noise (red arrow) in time.

shown by the black and red arrows in Figure 5(b). It should therefore be possible to cancel higher frequencies too.

To summarize, the following is a necessary condition for overcoming the timing bottleneck in ANC systems.

$$\text{Lookahead} \geq \text{Delay in } \{\text{ADC} + \text{DSP} + \text{DAC} + \text{Speaker}\} \quad (3)$$

This brings the natural question: **how much lookahead does MUTE provide in practice?** Let us assume that noise travels a distance d_r to reach the reference microphone at the IoT relay, and a distance $d_e > d_r$ to reach the error microphone at the ear device. Since wireless signals travel at the speed of light, a million times faster than the speed of sound, forwarding the noise signal from the IoT relay is almost instantaneous. Hence, lookahead can be calculated as:

$$T_{\text{lookahead}} = \frac{d_e}{v} - \frac{d_r}{v} = \frac{(d_e - d_r)}{v} \quad (4)$$

where v is the speed of sound in air (≈ 340 m/s). Translating to actual numbers, when $(d_e - d_r)$ is just 1m, lookahead is ≈ 3 ms, which is $100\times$ larger than today's ANC headphones. This implies that Alice can place the IoT relay on her office table and still benefit from wireless forwarding. Placing it on her office door, or ceiling, only increases this benefit.

3.2 Lookahead Aware ANC Algorithm

The timing benefit discussed above is a natural outcome of lookahead. However, we now (re)design the noise cancellation algorithm to explicitly exploit lookahead. Two key opportunities are of interest:

1. Recall from Equation 2 that the adaptive filter $h_{AF}(t)$ depends on the inverse channel, $h_{nr}^{-1}(t)$. Since this inverse is non-causal, the construction of the anti-noise signal would require sound samples from the future (elaborated soon). Today's systems lack future samples, hence live with sub-optimal cancellation. Large lookahead with MUTE can close this gap.
2. Lookahead will help foresee macro changes in sound profiles, such as when different people are taking turns in speaking. While traditional ANC incurs latency to converge to new sound profiles, MUTE can cache appropriate filters for each profile and "load" them at profile transitions. With lookahead, profile transitions would be recognizable in advance.

We begin with the first opportunity.

(1) Adaptive Filtering with Future Samples

■ **Basic Filtering:** Observe that a filter is essentially a vector, the elements of which are used to multiply the arriving sound samples. Consider an averaging filter that performs the average of the recent 3 sound samples – this filter can be represented as a vector $h_F = [\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]$. At any given time t , the output of the sound passing through this filter would be: $y(t) = \frac{1}{3}x(t) + \frac{1}{3}x(t-1) + \frac{1}{3}x(t-2)$ (which is called the convolution operation “*”). This filter is called *causal* since the output sample only relies on past input samples.

■ **Non-Causality:** Now consider the inverse of this filter h_F^{-1} . This should be another vector which convolved with

$y(t)$ should give back $x(t)$, i.e., $x(t) = h_F^{-1} * y(t)$. Filtering theory says that this inverse needs to be carefully characterized, since they are *non-causal*, *unstable*, or both [38, 42]. With a non-causal inverse, determining $x(t)$ would require $y(t + k)$ for $k > 0$. Thus estimating $x(t)$ in real time would be difficult; future knowledge of $y(t)$ is necessary. The physical intuition is difficult to convey concisely, however, one way to reason about this is that $x(t)$ originally influenced $y(t + 1)$ and $y(t + 2)$, and hence, recovering $x(t)$ would require those future values as well. In typical cases where h_F is the room's impulse response (known to have non-minimum phase property [43]), the future samples needed could be far more [42, 45].

■ **Adaptive Filtering:** Now, let us turn to adaptive filtering (h_{AF}) needed for noise cancellation. The “adaptive” component arises from estimating the filter vector at a given time, convolving this vector with the input signal, and comparing the output signal against a target signal. Depending on the error from this comparison, the filter vector is *adapted* so that successive errors converge to a minimum. Since this adaptive filter is non-causal (due to its dependence on the inverse filter), it would need future samples of the input signal to minimize error. With partial or no future samples (i.e., a truncated filter), the error will be proportionally higher. With this background, let us now design the LANC algorithm to fully exploit future lookahead.

■ **LANC Design:** Recall from Section 2 that the adaptive filter needed for noise cancellation is $h_{AF}(t) = -h_{se}^{-1}(t) * h_{ne}(t) * h_{nr}^{-1}(t)$. This minimizes the error:

$$e(t) = h_{ne}(t) * n(t) + h_{se}(t) * h_{AF}(t) * x(t) \quad (5)$$

where $x(t)$ is the noise captured by the reference microphone, i.e., $x(t) = h_{nr}(t) * n(t)$. Now, to search for the optimal h_{AF} , we use steepest gradient descent on the squared error $e^2(t)$. Specifically, we adapt h_{AF} in a direction opposite to the derivative of the squared error:

$$h_{AF}^{(new)} = h_{AF}^{(old)} - \frac{\mu}{2} \frac{\partial e^2(t)}{\partial h_{AF}} \quad (6)$$

where μ is a parameter that governs the speed of gradient descent. Expanding the above equation for each filter coefficient $h_{AF}(k)$, we have:

$$h_{AF}^{(new)}(k) = h_{AF}^{(old)}(k) - \mu e(t) h_{se}(t) * x(t - k) \quad (7)$$

In the above equation, $h_{se}(t)$ is known and estimated a priori, $e(t)$ is measured from the error microphone, and $x(t)$ is measured from the reference microphone.

This is where non-causality emerges. Since h_{AF} is actually composed of h_{nr}^{-1} , the values of k in Equation 7 can be negative ($k < 0$). Thus, $x(t - k)$ becomes $x(t + k)$, $k > 0$, implying that the updated $h_{AF}^{(new)}$ requires future samples of $x(t)$. With

lookahead, our LANC algorithm is able to “peek” into the future and utilize those sound samples to update the filter coefficients. This naturally results in a more accurate anti-noise signal $\alpha(t)$, expressed as:

$$\alpha(t) = h_{AF}(t) * x(t) = \sum_{k=-N}^L h_{AF}(k) x(t - k) \quad (8)$$

Observe that larger the lookahead, larger is the value of N in the subscript of the summation, indicating a better filter inversion. Thus, with a lookahead of *several milliseconds* in LANC, N can be large and the anti-noise signal can significantly reduce error (see pseudocode in Alg. 1). In contrast, lookahead is *tens of microseconds* in today's headphones, forcing a strict truncation of the non-causal filter, leaving a residual error after cancellation.

Algorithm 1 LANC: Lookahead Aware Noise Cancellation

```

1: while True do
2:   Play  $\alpha(t)$  at anti-noise speaker
3:    $t = t + 1$ 
4:   Record the error  $e(t)$  at error mic.
5:   Record future sample  $x(t + N)$  at reference mic.
6:   for  $k = -N, k \leq L, k++$  do
7:      $h_{AF}(k) = h_{AF}(k) - \mu e(t) h_{se}(t) * x(t - k)$ 
8:   end for
9:    $\alpha(t) = \sum_{k=-N}^L h_{AF}(k) x(t - k)$ 
10: end while
```

(2) Predictive Sound Profiling

Another opportunity with lookahead pertains to coping with more complex noise sources, such as human conversation. Consider a common case where a human is talking intermittently in the presence of background noise – Figure 6(a) and (b) show an example spectrum for speech and background noise, respectively. Now, to cancel human speech, the adaptive filter estimates the channels from the human to the ear device. However, when the speech pauses, the filter must re-converge to the channels from the background noise source. Re-convergence incurs latency since the h_{AF} vector must again undergo the gradient descent process to stabilize at a new minimum. Our idea is to leverage lookahead to foresee this change in sound profile, and swap the filtering coefficients right after the speech has stopped. Hence, we expect our cancellation to not fluctuate even for alternating sound sources, like speech or music.

■ **Validation:** Figure 7 explains the problem by illustrating the convergence of a toy adaptive filter, h_{AF} , with 7 taps. Initially, the filter is $h_{AF}^{(1)}$, and since this vector is not accurate, the corresponding error in Figure 7(b) is large. The vector

Figure 6: Acoustic spectrum in the (a) presence and (b) absence of speech. LANC recognizes the profile and pre-loads its filter coefficients for faster convergence.

then gets updated to $h_{AF}^{(2)}$ based on Equation 7, in the direction that reduces the error. This makes $h_{AF}^{(2)}$ closer to the ideal filter and $e(t)^2$ closer to zero. The filter continues to get updated until the error becomes nearly zero – at this point, the filter is said to have converged, i.e., $h_{AF}^{(3)}$.

Figure 7: Convergence process of the adaptive filter, h_{AF} . (a) 7-tap h_{AF} filter changes from time (1) to time (3). (b) residual error $e(t)$ converges to a minimum.

For persistent noise (like machine hum), the converged adaptive filter can continue to efficiently cancel the noise, as shown in Figure 8(a). However, for intermittent speech signals with random pauses between sentences, the adaptive filter cannot maintain smooth cancellation as shown in Figure 8(b). Every time the speech starts, the error is large and the adaptive filter needs time to (re)converge again.

■ **Predict and Switch:** With substantial lookahead, LANC gets to foresee the start and stop of speech signals. Thus, instead of adapting the filter coefficients every time, we cache

Figure 8: LANC's convergence timeline showing adaptive filtering with (a) continuous noise, (b) speech, (c) lookahead aware profiling. LANC converges faster due to its ability to anticipate profile transitions in advance.

the coefficient vector for the corresponding sound profiles. A sound profile is essentially a statistical signature for the sound source – a simple example is the average energy distribution across frequencies. For 2 profiles – say speech and background noise – LANC caches 2 adaptive filter vectors, h_{AF}^{speech} and $h_{AF}^{background}$, respectively. Then, by analyzing the lookahead buffer in advance, LANC determines if the sound profile would change imminently. When the profile change is indeed imminent (say the starting of speech), LANC directly updates the adaptive filter with h_{AF}^{speech} , avoiding the overhead of re-convergence.

To generalize, LANC maintains a converged adaptive filter for each sound profile, and switches between them at the right time. So long as there is one dominant sound source at any given time, LANC cancels it quite smoothly as shown in Figure 8(c). Without lookahead, however, the profile-change cannot be detected in advance, resulting in periodic re-convergence and performance fluctuations.

With the LANC algorithm in place, we now turn to bringing together the overall MUTE system.

4 MUTE: SYSTEM AND ARCHITECTURE

Recall that our basic system requires an IoT relay installed near the user; the relay listens to the ambience and streams

the acoustic waveform over its RF interface in real time. The receiver – a hollow earphone – receives the sound signal, applies the LANC algorithm to compute the anti-noise signal, and finally plays it through the speaker. Several components have been engineered to achieve a fully functional system. In the interest of space, we discuss 3 of these components, namely: (1) the wireless relay hardware, (2) automatic relay selection, and (3) privacy protection. Finally, as a conclusion to this section, we envision architectural variants of *MUTE* – such as noise cancellation as a service – to demonstrate a greater potential of our proposal beyond what is presented in this paper. We begin with wireless relay design.

4.1 Wireless Relay Design

Figure 9 shows the hardware block diagram of the wireless relay. *MUTE* embraces an analog design to bypass delays from digitization and processing. Specifically, the relay consists of a (reference) microphone that captures the ambient noise signal, passes it through a low pass filter (LPF), and then amplifies it. An impedance matching circuit connects the audio signal to an RF VCO (voltage controlled oscillator). The VCO outputs a *frequency modulated* (FM) signal, which is then mixed with a carrier frequency generated by a phase lock loop (PLL), and up-converted to the 900 MHz ISM band. The RF signal is then band pass filtered and passed to a power amplifier connected to a 900 MHz antenna. Thus, with audio signal $m(t)$ captured at the microphone, the transmitted signal $x(t)$ is:

$$x(t) = A_p \cos \left(2\pi f_c t + 2\pi A_f \int_0^t m(\tau) d\tau \right) \quad (9)$$

where f_c is the carrier frequency, A_p is the gain of the RF amplifier, and A_f is the combined gain of the audio amplifier and FM modulator¹.

Figure 9: *MUTE*'s RF Relay Design

Why Frequency Modulation (FM)? The significance of FM is three-fold. *First*, it delivers better audio quality because noise mainly affects amplitude, leaving the frequency of the signal relatively less affected. *Second*, since the bandwidth used is narrow, $h_w(t)$ is flat in frequency and hence can be represented with a single tap. As a result, there is no need to

¹The receiver in the ear-device applies a reverse set of operations to the transmitter and outputs digital samples that are then forwarded to the DSP.

estimate the wireless channel since it will not affect the audio signal $m(t)$. *Finally*, any carrier frequency offsets between up-conversion and down-conversion appear as a constant DC offset in the output of the FM demodulator which can easily be averaged out. This precludes the need to explicitly compensate for carrier frequency offset (CFO).

4.2 Automatic Relay Selection

MUTE is effective only when the wireless relay is located closer to the sound source than the earphone. This holds in scenarios such as Figure 1 – the relay on Alice's door is indeed closer to the noisy corridor. However, if the sound arrives from an opposite direction (say from a window), the relay will sense the sound *after* the earphone. Even though the relay forwards this sound, the earphone *should not* use it since the lookahead is negative now (i.e., the wirelessly-forwarded sound is lagging behind). Clearly, *MUTE* must discriminate between positive and negative lookahead, and in case of the latter, perhaps nudge the user to reposition the relay in the rough direction of the sound source.

■ **How to determine positive lookahead?** *MUTE* uses the GCC-PHAT cross-correlation technique [21]. The DSP processor periodically correlates the wirelessly-forwarded sound against the signal from its error microphone. The time of correlation-spike tells whether the lookahead is positive or negative. When positive, the LANC algorithm is invoked. Correlation is performed periodically to handle the possibility that the sound source has moved to another location.

■ **Multiple Relays:** Observe that a user could place multiple relays around her to avoid manually repositioning the relay in the direction of the noise source. The correlation technique would still apply seamlessly in such a scenario. The relay whose correlation spike is most shifted in time is the one *MUTE* would pick. This relay would offer the maximum lookahead, hence the best cancellation advantage.

4.3 Architectural Variants

The basic architecture thus far is a wireless IoT relay (closer to the sound source) communicating to an ear-device around the human ear. We briefly sketch a few variants of this architecture aimed at different trade-offs and applications.

1. **Personal TableTop:** *MUTE* removes the reference microphone from the headphone, which in turn eliminates the noise-absorbing material. As mentioned earlier, this makes the ear-device light and hollow. Following this line of reasoning, one could ask what else could be stripped off from the ear-device. We observe that even the DSP can be extracted and inserted into the IoT relay. In other words, the IoT relay could compute the anti-noise and wirelessly transmit to the ear-device; the ear-device could play it through

Figure 10: Architectural variants: (a) Personal tabletop device includes DSP and reference microphone; sends anti-noise signal to ear-device, which responds with error signal. (b) Noise cancellation as a edge service: the DSP server is connected to IoT relays on the ceiling and computes the anti-noise for all users. (c) Smart noise, where noise sources attach a IoT relay while users with *MUTE* ear-devices benefit.

the anti-noise speaker, and transmit back the error signal from its error microphone. Observe that the IoT relay can even become a portable table-top device, with the ear-device as a simple “client”. The user can now carry her personal *MUTE* tabletop relay (Figure 10(a)), eliminating dependencies on door or wall mounted infrastructure.

2. **Public Edge Service:** Another organization is to move the DSP to a backend server, and connect multiple IoT relays to it, enabling a *MUTE* public service (Figure 10(b)). The DSP processor can compute the anti-noise for *each* user and send it over RF. If computation becomes the bottleneck with multiple users, perhaps the server could be upgraded with multiple-DSP cores. The broader vision is an edge cloud [47] that offers acoustic services to places like call centers.
3. **Smart Noise:** A third architecture could be to attach IoT relays to noise sources themselves (and eliminate the relays on doors or ceilings). Thus, heavy machines in construction sites, festive public speakers, or lawn mowers, could broadcast their own sound over RF. Those disturbed by these noises can wear the *MUTE* ear-device, including the DSP. Given the maximal lookahead, high quality cancellation should be feasible.

We conclude by observing that the above ideas may be viewed as a “disaggregation” of conventional headphones, enabling new future-facing possibilities. This paper is an early step in that direction.

4.4 Privacy Awareness

Two relevant questions emerge around privacy:

- **Will the IoT relay record ambient sounds and conversations?** We emphasize that the relays are analog and

not designed to even hold the acoustic samples. The microphone’s output is directly applied to modulate the 900 MHz carrier signal with no recording whatsoever. In this sense, *MUTE* is different from Amazon Echo, Google Home, and wireless cameras that must record digital samples for processing.

- **Will the wirelessly-forwarded sound reach certain areas where it wouldn’t have been audible otherwise?** This may be a valid concern for some scenarios, e.g., a person outside a coffee shop may be able to “hear” inside conversations. However, with power control, beamforming, and sound scrambling, the problem can be alleviated. We leave a deeper treatment of this problem to future work. On the other hand, this may not be a problem in other scenarios. For instance, with personal table-top devices, the wireless range can be around the user’s table, resulting in almost no leakage. For smart noise, the noise need not be protected at all, while for call center-like settings, acoustic privacy is relatively less serious.

5 EVALUATION

We begin with some details on experimental setup and comparison schemes, followed by performance results.

5.1 Experimental Setup

MUTE’s core algorithms are implemented on the Texas Instrument’s TMS320C6713 DSP board [6], equipped with the TLV320AIC23 codec. The microphones are SparkFun’s MEMS Microphone ADMP401 and the anti-noise speaker is the AmazonBasics computer speaker. Ambient noise is played from an Xtrememac IPU-TRX-11 speaker. All microphones and speakers are cheap off-the-shelf equipment. For

Figure 11: *MUTE+Passive*: (a) Bose headphone on the 3D head model, with DSP output connected to the headset. (b) The measurement microphone inside the ear, and the reference microphone nearby.

performance comparison, we purchased Bose’s latest ANC headphone, the QC35 [10] (pictured in Figure 11).

For experimentation, we insert a separate “measurement microphone” at the ear-drum location of a 3D head model (Figure 2(b)) – this serves as the approximation of what the human would hear. We play various sounds from the ambient speaker and measure the power level at this microphone. We then compare the following schemes:

- ***MUTE_Hollow***: Our error microphone is pasted outside the ear while the anti-noise speaker and DSP board are placed next to it, as shown in Figure 2(b).
- ***Bose_Active***: We place the Bose headphone on the 3D head model and measure cancellation, first with ANC turned OFF, and then with ANC turned ON. Subtracting the former from the latter, we get Bose’s active noise cancellation performance.
- ***Bose_Overall***: We turn on ANC for Bose and measure the net cancellation, i.e., the combination of its ANC and passive noise-absorbing material.

Finally, we bring human volunteers to compare Bose and *MUTE*. In the absence of a compact form factor for *MUTE*, we utilize Bose’s headphone. Specifically, we feed the output of our DSP board into the AUX input of the Bose headphone (with its ANC turned OFF), meaning that our LANC algorithm is executed through Bose’s headphone (instead of its native ANC module). Of course, the passive sound absorbing material now benefits both Bose and *MUTE*, hence we call our system ***MUTE+Passive*** (see Figure 11). We report cancellation results for various sounds, including machines, human speech, and music.

5.2 Performance Results

Our results are aimed at answering the following questions:

- (1) Comparison of overall noise cancellation for *MUTE_Hollow*, *Bose_Active*, *Bose_Overall*, and *MUTE+Passive*.
- (2) Performance comparison for various sound types.

- (3) Human experience for *Bose_Overall* and *MUTE+Passive*.
- (4) Impact of lookahead length on *MUTE_Hollow*.
- (5) Accuracy of relay selection for *MUTE_Hollow*.

Overall Noise Cancellation

Figure 12: *MUTE* and Bose’s overall performance.

Figure 12 reports comparative results when wide-band white noise (which is most unpredictable of all noises) is played from the ambient speaker. The noise level is maintained at 67 dB at the measurement microphone. Four main points are evident from the graph. (1) *Bose_Active* is effective only at lower frequency bands, implying that Bose must rely on passive materials to cancel sounds from 1 kHz to 4 kHz. (2) The ear-blocking passive material is effective at higher frequencies, giving *Bose_Overall* a -15 dB average cancellation. (3) *MUTE_Hollow* is almost comparable to *Bose_Overall* even without passive materials, indicating that our LANC algorithm performs well (*Bose_Overall* is just 0.9 dB better on average). (4) When *MUTE+Passive* gains the advantage of passive materials, the cancellation is 8.9 dB better than *Bose_Overall*, on average.

In summary, *MUTE* offers two options in the cancellation versus comfort tradeoff. A user who values comfort (perhaps for long continuous use) can prefer lightweight, open-ear *MUTE* devices at a 0.9 dB compromise from Bose, while one who cares more about noise suppression can experience 8.9 dB improvement over Bose.

We briefly discuss two technical details: (1) *MUTE*’s cancellation is capped at 4 kHz due to limited processing speed of the TMS320C6713 DSP. It can sample at most 8 kHz to finish the computation within one sampling interval. A faster DSP will ease the problem. (2) The diminishing cancellation at very low frequencies (<100 Hz) is due to the weak response of our cheap microphone and anti-noise speaker – Figure 13 plots the combined frequency response.

Figure 13: The combined frequency response of our anti-noise speaker and the microphone.

■ Varying Ambient Sounds (Speech, Music)

Figure 14 shows *MUTE*'s cancellation performance across 4 different types of real-world noises with different spectral characteristics: male voice, female voice, construction sound, and music. The results are a comparison between *MUTE_Hollow* and *Bose_Overall*. Our lookahead-aware ANC algorithm achieves mean cancellation within 0.9dB to Bose's native ANC combined with its carefully perfected passive sound-absorbing materials [10].

Figure 14: Comparison between *MUTE_Hollow* and *Bose_Overall*, measured for 4 types of ambient sounds.

■ Human Experience

We invited 5 volunteers to rate *MUTE+Passive*'s performance relative to *Bose_Overall*. Recall that for *MUTE+Passive*, we use the Bose headset with ANC turned OFF. Now, since we have only one DSP board, we were able to run *MUTE+Passive* only on the right ear – for the left ear, we use both an earplug and the headset (with ANC turned OFF). For *Bose_Overall*,

we turned ON native ANC on both ears. In this setup, we played various human voices and music through the ambient speaker. Since fine grained (per-frequency) comparison is difficult for humans, we requested an overall rating between 1 to 5 stars. We did not tell the volunteers when *MUTE* or Bose was being used for cancellation.

Figure 15: User feedback of music and voice noise.

Figure 15 shows the comparison for music and human voice. Every volunteer consistently rated *MUTE* above Bose. Their subjective opinions were also strongly positive. However, almost all of them also said that “Bose was superb at canceling hums in the environment”, and *MUTE* did not perform as well. One reason is the weak response of the speaker and microphone at low frequencies, as mentioned before. Upon analyzing, we also realized that the background hums are from various sources. With Bose's microphone array, they are equipped to handle such scenarios, while our current system is aimed at a single noise source (the ambient speaker). We have left multi-source noise cancellation to future work, as discussed later in Section 6.

■ Impact of Shorter Lookahead

Lookahead reduces when the wireless relay gets closer to the user, or when the location of the noise source changes such that the time-difference between direct path and wireless-relay path grows smaller. For accurate comparison across different lookaheads, we need to ensure that the physical environment (i.e., multipath channel) remains identical. Therefore, instead of physically moving the noise source or the wireless relay (to vary lookahead time), we fix their positions, but deliberately inject delays into the reference signal within the DSP processor (using a delayed line buffer).

Figure 16 plots the results for *MUTE_Hollow*. The lookahead times are expressed relative to the “Lower Bound” from Equation 3 (recall that lookahead must be greater than ADC + DSP processing + DAC + speaker delay, as explained in Section 3.1). Evidently, as the lookahead increases, the performance improves due to better inverse filtering.

Figure 16: As lookahead becomes smaller, the system performance degrades.

■ Profiling and Cancellation

To highlight the efficacy of sound profiling and filter switching, we run a separate experiment where wide-band background noise is constantly being played from one ambient speaker, while mixed human voice (with pauses) is being played from another speaker. We compare the residual error of *MUTE*'s filter selection mechanism with that of using only one adaptive filter. Figure 17 shows the cancellation gain in *MUTE_Hollow* with profiling and switching turned ON. Evidently, the cancellation improves by 3 dB on average. We could not compare with Bose in this case since Bose uses at least 6 microphones to cope with scattered noise sources. Upgrading *MUTE* with that many microphones is bound to offer substantial advantage.

Figure 17: Lookahead enabled filter switching provides additional gain for intermittent noise cancellation.

■ Wireless Relay Selection

Does the correlation technique to identify (maximum) positive lookahead work in real environments? Figure 18 shows two typical examples of GCC-PHAT based cross-correlation between the forwarded sound waveform and the directly-received sound. Observe that one case is positive lookahead while the other is negative. *MUTE* was able to correctly determine these cases in every instance.

Now consider multiple relays and different locations of the noise source. Figure 19 shows *MUTE*'s ability to correctly pick the wireless relay depending on the ambient speaker location in the room. We place the *MUTE* client at the center

Figure 18: *MUTE* client chooses the relay with largest positive lookahead (i.e., earliest correlation).

of the room, and three wireless relays around the edges and corners. We observe that when the ambient speaker is near the i -th relay, *MUTE* selects that relay consistently. We also observe that when the noise source is closer to the *MUTE* client location, no relay is selected because all of them offer negative lookahead.

Figure 19: *MUTE* client associates with appropriate RF relays, depending on the location of the noise source.

6 CURRENT LIMITATIONS

Needless to say, there is room for further work and improvement. We discuss a few points here.

- **Multiple Noise Sources:** Our experiments were performed in natural indoor environments, with a dominant noise source (such as a human talking on the phone, or music from an audio speaker). With multiple noise sources, the problem is involved, requiring either multiple microphones (one for each noise channel), or source separation algorithms that depend on statistical independence among sources. Today's ANC headphones utilize at least 6 microphones and source separation algorithms to mitigate such issues. We believe the benefits of looking ahead into future samples will be valuable for multiple sources as well – a topic we leave to future work.
- **Cancellation at the Human Ear:** We have aimed at achieving noise cancellation at the measurement microphone, under the assumption that the ear-drum is also located close to the error microphone. Bose, Sony, and other companies take a step further, i.e., they utilize anatomical ear models (e.g., KEMAR head [4]) and design for cancellation at the human ear-drum. Thus, Bose's performance may have

been sub-optimal in our experiments. However, even without ear-model optimizations, our human experiments have returned positive feedback. Of course, a more accurate comparison with Bose would require *MUTE* to also adopt human ear-models, and then test with large number of human subjects. We have left this to future work. Finally, companies like *Nura* [14] are leveraging in-ear acoustic signals to build personalized ear models. Embracing such models are likely to benefit both *MUTE* and Bose.

- **Head Mobility:** We have side-stepped human head mobility since our error microphone is static around the head model. Of course, head mobility will cause faster channel fluctuations, slowing down convergence. While this affects all ANC realizations (including Bose and Sony headphones), the issue has been alleviated by bringing enhanced filtering methods known to converge faster. We plan to also apply such mobility-aware LMS techniques in our future versions of *MUTE*.
- **Portability:** While Bose and Sony headphones are easily portable, *MUTE* requires the user to be around the IoT relay. While this may serve most static use cases (e.g., working at office, snoozing at the airport, sleeping at home, working out in the gym, etc.), headphones may be advantageous in completely mobile scenarios, like running on the road.
- **RF Interference and Channel Contention:** Our system will occupy the RF channel once the IoT relay starts streaming. However, it only occupies 8 kHz bandwidth, far smaller than the 26 MHz channel in the 900 MHz ISM band. Further, covering an area requires few relays (3 for any horizontal noise source direction, 4 for any 3D direction), hence, the total bandwidth occupied remains a small fraction. Even with multiple co-located users, channel contention can be addressed by carrier-sensing and channel allocation.

7 RELATED WORK

The literature in acoustics and active noise control is extremely rich, with connections to various sub-fields of engineering [20, 24, 25, 30, 35, 36, 39, 49, 51]. In the interest of space, we directly zoom into two directions closest to *MUTE*: wireless ANC, and ANC with lookahead.

Wireless ANC: An RF control plane has been proposed in the context of multi-processor ANC, mainly to cope with various sound sources in large spaces [23, 26–29, 34]. In this body of work, distributed DSP processors communicate between themselves over wired/wireless links to achieve real-time, distributed, noise cancellation. The notion of “piggybacking” sound over RF, to exploit the propagation delay difference, is not a focus in these systems. Moreover, most of the mentioned systems are via simulations [23, 26–28].

ANC with Lookahead: Certain car models [1–3] and airplanes [7, 8] implement ANC inside their cabins – reference microphones are placed near the engine and connected via wires to the DSP devices. While this offers promising lookahead, observe that the problems of inverse-channel estimation are almost absent, since the noise source positions are known, the noise signal is well structured, and the acoustic channel is stable. Moreover, these systems have no notion of at-ear feedback (from headphone microphones), since they are canceling broadly around the passenger’s head locations. This is the reason why cancellation is feasible only at very low frequencies (<100 Hz in Honda vehicles [2]). In contrast, *MUTE* introduces *wireless forwarding*, embeds lookahead-awareness in the ANC pipeline, and integrates a personal architecture for 4 kHz cancellation. Said differently, the intersection of “personal” ANC and “wireless” lookahead is both technically and architecturally new, to the best of our knowledge.

The idea of sound forwarding over RF has been applied to very different contexts, such as acoustic communication across sound-proof boundaries [37], wireless acoustic MIMO and beamforming [19], and even walkie-talkies and wireless microphones [16, 17]. However, they are not aimed towards noise cancellation. Finally, we should mention that some systems have leveraged the propagation delay difference between RF and sound, albeit for other applications. *Cricket* [44], *AHLoS* [48], and *Dolphin* [40] have all used the Time-of-Arrival (ToA) difference between RF and sound for ranging and localization. [50] uses RF signals as a tool to avoid acoustic collision in wireless sensor networks. Overall, this is similar to how earthquake and tsunami sensors [11, 12] work, by utilizing the fact that wireless signals travel much faster than ocean waves and tectonic vibrations.

8 CONCLUSION

This paper exploits the velocity gap between RF and sound to improve active noise cancellation. By anticipating the sound milliseconds in advance, our proposed system is able to compute the anti-noise signal in time, better estimate sound channels, and ultimately attain wider-band cancellation. In addition, the core idea opens a number of architectural possibilities at the intersection of wireless networks and acoustic sensing. This paper is a first step towards these directions.

ACKNOWLEDGMENTS

We sincerely thank our shepherd, Dr. Ranveer Chandra, and the anonymous reviewers for their insightful comments and suggestions. We are also grateful to Intel, Google, Qualcomm, HP and NSF (grant NSF 1619313) for partially funding this research.

REFERENCES

- [1] 2013. Cadillac ELR Takes Active Noise Cancelling to the Limit. Retrieved January 29, 2018 from <http://media.gm.com/media/us/en/cadillac/vehicles/elr/2014.detail.html/content/Pages/news/us/en/2013/Oct/1017-elr-anc.html>
- [2] 2014. Honda's Active Noise Cancellation. Retrieved January 29, 2018 from <https://www.honda.co.nz/technology/driving/anc/>
- [3] 2016. 2017 Buick LaCrosse Quiet Tuning. Retrieved January 29, 2018 from <http://media.buick.com/media/us/en/buick/bcportal.html/currentVideoId/5161313633001/currentChannelId/Running%20Footage/Cars.gsaOff.html>
- [4] 2017. KEMAR, For ear- and headphone test, 2-ch. Retrieved January 29, 2018 from <https://www.gras.dk/products/head-torso-simulators-kemar/kemar-for-ear-headphone-test-2-ch>
- [5] 2017. Review and Measurements: Bose QC25 Noise-Cancelling Headphone. Retrieved January 30, 2018 from <https://www.lifewire.com/bose-qc25-review-specs-3134560>
- [6] 2017. TMS320C6713 DSP Starter Kit. Retrieved January 29, 2018 from <http://www.ti.com/tool/TMDSDSK6713>
- [7] 2018. Bombardier Q300 DHC-8 Dash 8. Retrieved January 29, 2018 from https://www.aerospace-technology.com/projects/bombardier_q300/
- [8] 2018. Bombardier Q400 Active Noise and Vibration Suppression. Retrieved January 29, 2018 from <http://commercialaircraft.bombardier.com/en/q400/Technology.html#1397740891734>
- [9] 2018. Bose QuietControl 30 Wireless Noise Cancelling Earphone. Retrieved January 29, 2018 from https://www.bose.com/en_us/products/headphones/earphones/quietcontrol-30.html
- [10] 2018. Bose QuiteComfort 35 Wireless Noise Cancelling Headphone. Retrieved January 29, 2018 from https://www.bose.com/en_us/products/headphones/over_ear_headphones/quietcomfort-35-wireless-ii.html
- [11] 2018. Deep Ocean Tsunami Detection Buoys. Retrieved January 30, 2018 from http://www.bom.gov.au/tsunami/about/detection_buoys.shtml
- [12] 2018. Earthquake Early Warning. Retrieved January 30, 2018 from <https://earthquake.usgs.gov/research/earlywarning/nextsteps.php>
- [13] 2018. Here Active Listening. Retrieved January 30, 2018 from <https://www.kickstarter.com/projects/dopplerlabs/here-active-listening-change-the-way-you-hear-the/description>
- [14] 2018. Nura: Headphones that learn and adapt to your unique hearing by Nura. Retrieved June 18, 2018 from <https://www.kickstarter.com/projects/nura/nura-headphones-that-learn-and-adapt-to-your-unique>
- [15] 2018. Sony 1000X Wireless Noise Cancelling Headphone. Retrieved January 29, 2018 from <https://www.sony.com/electronics/headband-headphones/mdr-1000x>
- [16] 2018. Walkie-talkie. Retrieved January 29, 2018 from <https://en.wikipedia.org/wiki/Walkie-talkie>
- [17] 2018. Wireless microphone. Retrieved January 29, 2018 from https://en.wikipedia.org/wiki/Wireless_microphone
- [18] 2018. Wireless noise cancelling headphones SHB8850NC/27 | Philips. Retrieved January 29, 2018 from https://www.usa.philips.com/c-p/SHB8850NC_27/wireless-noise-cancelling-headphones
- [19] 2018. WS800 Wireless Microphone System. Retrieved January 29, 2018 from http://www.clearone.com/products_wireless_microphone_system
- [20] Elias Bjarnason. 1995. Analysis of the filtered-X LMS algorithm. *IEEE Transactions on Speech and Audio Processing* 3, 6 (1995), 504–514.
- [21] Michael S Brandstein and Harvey F Silverman. 1997. A robust method for speech signal time-delay estimation in reverberant rooms. In *Acoustics, Speech, and Signal Processing, 1997. ICASSP-97, 1997 IEEE International Conference on*, Vol. 1. IEEE, 375–378.
- [22] John G Casali and James F Grenell. 1990. Noise-attenuating earmuff comfort: A brief review and investigation of band-force, cushion, and wearing-time effects. *Applied Acoustics* 29, 2 (1990), 117–138.
- [23] Stephen J Elliott. 2005. Distributed control of sound and vibration. *Noise control engineering journal* 53, 5 (2005), 165–180.
- [24] Stephen J Elliott and Philip A Nelson. 1993. Active noise control. *IEEE signal processing magazine* 10, 4 (1993), 12–35.
- [25] LJ Eriksson. 1991. Development of the filtered-U algorithm for active noise control. *The Journal of the Acoustical Society of America* 89, 1 (1991), 257–265.
- [26] Miguel Ferrer, Maria de Diego, Gema Piñero, and Alberto Gonzalez. 2015. Active noise control over adaptive distributed networks. *Signal Processing* 107 (2015), 82–95.
- [27] Kenneth D Frampton. 2002. Decentralized vibration control in a launch vehicle payload fairing. In *ASME 2002 International Mechanical Engineering Congress and Exposition*. American Society of Mechanical Engineers, 155–160.
- [28] Kenneth D Frampton. 2003. The control of rocket fairing interior noise with a networked embedded system. *The Journal of the Acoustical Society of America* 113, 4 (2003), 2251–2251.
- [29] Kenneth D Frampton. 2005. Advantages and challenges of distributed active vibro-acoustic control. *The Journal of the Acoustical Society of America* 118, 3 (2005), 1950–1950.
- [30] Woon S Gan and Sen M Kuo. 2002. An integrated audio and active noise control headset. *IEEE Transactions on Consumer Electronics* 48, 2 (2002), 242–247.
- [31] Samir NY Gerges. 2012. Earmuff comfort. *Applied acoustics* 73, 10 (2012), 1003–1012.
- [32] Simon Haykin and Bernard Widrow. 2003. *Least-mean-square adaptive filters*. Vol. 31. John Wiley & Sons.
- [33] Yeh-Liang Hsu, Chung-Cheng Huang, Chin-Yu Yo, Chiou-Jong Chen, and Chun-Ming Lien. 2004. Comfort evaluation of hearing protection. *International Journal of Industrial Ergonomics* 33, 6 (2004), 543–551.
- [34] Vinod Kulathumani and Bryan Lemon. 2013. Sufficiency of Local Feedback for Sensor-Actuator Network-Based Control Systems with Distance Sensitivity Properties. *Journal of Sensor and Actuator Networks* 2, 3 (2013), 409–423.
- [35] Sen M Kuo, Sohini Mitra, and Woon-Seng Gan. 2006. Active noise control system for headphone applications. *IEEE Transactions on Control Systems Technology* 14, 2 (2006), 331–335.
- [36] Sen M Kuo and Dennis Morgan. 1995. *Active noise control systems: algorithms and DSP implementations*. John Wiley & Sons, Inc.
- [37] Lichuan Liu and Sen M Kuo. 2013. Wireless communication integrated active noise control system for infant incubators. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*. IEEE, 375–378.
- [38] Dimitris G Manolakis and Vinay K Ingle. 2011. *Applied digital signal processing: theory and practice*. Cambridge University Press.
- [39] Wenguang Mao, Jian He, and Lili Qiu. 2016. CAT: high-precision acoustic motion tracking. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*. ACM, 69–81.
- [40] Masateru Minami, Yasuhiro Fukuju, Kazuki Hirasawa, Shigeaki Yokoyama, Moriyuki Mizumachi, Hiroyuki Morikawa, and Tomonori Aoyama. 2004. DOLPHIN: A practical approach for implementing a fully distributed indoor ultrasonic positioning system. In *UbiComp*. Springer, 347–365.
- [41] J.P. Mobley, C. Zhang, S.D. Soli, C. Johnson, and D. O'Connell. 1998. Pressure-regulating ear plug. <https://www.google.com/patents/>

- US5819745 US Patent 5,819,745.
- [42] John N Mourjopoulos. 1994. Digital equalization of room acoustics. *Journal of the Audio Engineering Society* 42, 11 (1994), 884–900.
 - [43] Stephen T Neely and Jont B Allen. 1979. Invertibility of a room impulse response. *The Journal of the Acoustical Society of America* 66, 1 (1979), 165–169.
 - [44] Nissanka B Priyantha, Anit Chakraborty, and Hari Balakrishnan. 2000. The cricket location-support system. In *Proceedings of the 6th annual international conference on Mobile computing and networking*. ACM, 32–43.
 - [45] Biljana D Radlovic and Rodney A Kennedy. 2000. Nonminimum-phase equalization and its subjective importance in room acoustics. *IEEE Transactions on Speech and Audio Processing* 8, 6 (2000), 728–737.
 - [46] Ben Rudzyn and Michael Fisher. 2012. Performance of personal active noise reduction devices. *Applied Acoustics* 73, 11 (2012), 1159–1167.
 - [47] Mahadev Satyanarayanan, Paramvir Bahl, Ramón Caceres, and Nigel Davies. 2009. The case for vm-based cloudlets in mobile computing. *IEEE pervasive Computing* 8, 4 (2009).
 - [48] Andreas Savvides, Chih-Chieh Han, and Mani B Strivastava. 2001. Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proceedings of the 7th annual international conference on Mobile computing and networking*. ACM, 166–179.
 - [49] Ying Song, Yu Gong, and Sen M Kuo. 2005. A robust hybrid feedback active noise cancellation headset. *IEEE transactions on speech and audio processing* 13, 4 (2005), 607–617.
 - [50] Qixin Wang, Rong Zheng, Ajay Tirumala, Xue Liu, and Lui Sha. 2008. Lightning: A hard real-time, fast, and lightweight low-end wireless sensor election protocol for acoustic event localization. *IEEE Transactions on Mobile Computing* 7, 5 (2008), 570–584.
 - [51] Sangki Yun, Yi-Chao Chen, and Lili Qiu. 2015. Turning a mobile device into a mouse in the air. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 15–29.