Xiaoran Fan *Google, Irvine, CA, USA* **Longfei Shangguan** *University of Pittsburgh, Pittsburgh, PA, USA*
Siddharth Rupavatharam *Rutgers, The State University of New Jersey, NJ, USA*
Yanyong Zhang *University of Science and Technology of China, Hefei, Anhui, China*
Jie Xiong *University of Massachusetts Amherst, MA, USA* **Yunfei Ma** *Alibaba Group, CA, USA*
Richard Howard *Rutgers, The State University of New Jersey, NJ, USA*

Editors: Nicholas D. Lane and Xia Zhou

A NEW DESIGN PARADIGM FOR ENABLING SMART HEADPHONES

Excerpted from "HeadFi: bringing intelligence to all headphones," from *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking* with permission. https://dl.acm.org/doi/10.1145/3447993.3448624©ACM 2021

Headphones continue to grow more intelligent as new functions (e.g., touch-based gesture control) appear. These functions usually rely on auxiliary sensors (e.g., accelerometer and gyroscope) that are available in smart headphones. However, for those headphones that do not have such sensors, supporting these functions becomes a daunting task. This paper presents *HeadFi*, a new design paradigm for bringing intelligence to all headphones. Instead of adding auxiliary sensors into headphones, HeadFi turns the pair of drivers that are readily available inside all headphones into a versatile sensor to enable new applications, spanning across mobile health, user-interface, and context-awareness. HeadFi works as a plug-in peripheral connecting the headphones and the pairing device (e.g., a smartphone). The simplicity (can be as simple as just two resistors) and small form factor of this design lend itself to be embedded into the pairing device as an integrated circuit. We envision that HeadFi can serve as a vital *supplementary* solution to existing smart headphone design by directly transforming large amounts of existing "dumb" headphones into intelligent ones.

Photo, istockphoto.com

Headphones[1] are among the most popular wearable devices worldwide, and are forecast to maintain the leading position in the coming years [4]. Existing smart headphones all build upon advanced hardware components (mostly embedded sensors). However, statistics show that more than 99% of consumer headphones shipped in 2019 are not equipped with embedded sensors, and more than 43% of consumer headphones even lack a microphone [3, 12]. Thus, consumers must purchase a new pair of smart headphones with embedded sensors to enjoy the sensing features.

In this paper, we ask the following question: Can we turn these non-smart headphones in hand into intelligent ones without redesigning the headphone or adding embedded sensors? A positive answer would enable the consumers to enjoy smart features on their "dumb" headphones at a minimal cost. More importantly, it would also pave the way for realizing earable intelligence at an unprecedented scale by transforming the large amount of existing non-smart headphones into intelligent ones. To realize this high-level idea, we need to address both technical and implementation challenges. From the technical point of view, the primary challenge comes from measuring the minute variation in voltage induced by the pressure change. The voltage measurement on the headphones is determined by both the audio input signal (e.g., music) and the excitation signal. In practice, however, the excitation signals are weak and can easily be buried in the audio input signal that is orders of magnitude stronger. From the usability point of view, our design should not break the appearance and the internal structure of the headphones. Besides, as the headphones are usually driven by mobile devices, our design should also be low power, incurring zero or negligible power consumption.

To address these challenges, we are inspired by a null measurement circuit design—*Wheatstone bridge*. Originally Wheatstone bridge was used to measure an unknown resistance by balancing the two arms of the bridge. In HeadFi, we repurpose the Wheatstone bridge to cancel the strong interference of the audio input signals to measure the subtle variations in voltage caused by excitation signals. Specifically, the left and right two drivers of the headphones are connected to the two arms of the bridge using the headphones' stereo jacket. Once the bridge is balanced, its output voltage does not change with the variation of the audio input signal. On the other hand, the output voltage of this bridge still varies with the pressure change around the headphones, which is affected by the excitation signals such as the hand touch and the physiological activities. We prototype HeadFi on PCB, demonstrate four applications (Figure 1), and conduct extensive experiments with 53 volunteers using 54 pairs of non-smart headphones under the institutional review board (IRB) protocols. The results show that HeadFi can achieve 97.2%–99.5% accuracy on user identification, 96.8%–99.2% accuracy on heart rate monitoring, and 97.7%–99.3% accuracy on gesture recognition.

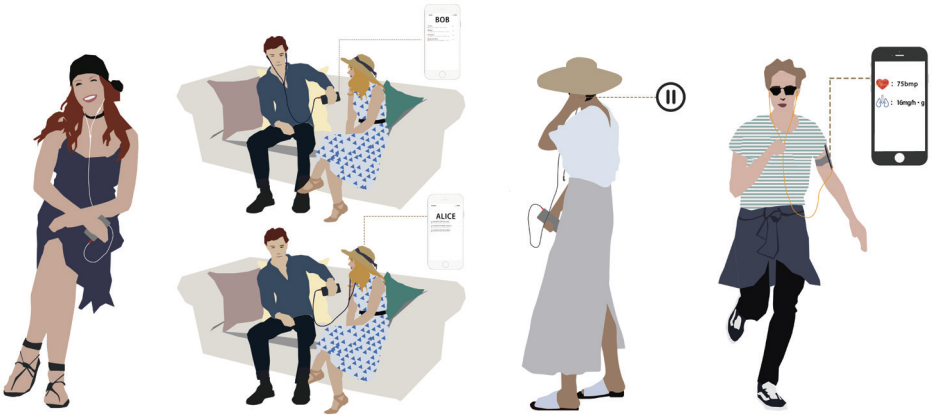

(a) Voice communication (b) User identification (c) Gesture control (d) Heart rate monitoring

FIGURE 1. Mobile applications enabled by HeadFi. (a) voice communication on headphones without a microphone (b) identifying different users and switch the context (c) touch-based gesture control (d) physiological activity monitoring.

TRANSFORMING HEADPHONES INTO SENSORS

HeadFi employs the pair of drivers inside headphones as versatile sensors to realize the functionalities mentioned above. In essence, speakers and microphones are reciprocal in principle [1]. For headphones without a built-in microphone, an intuitive solution would be turning the speaker[2] into a microphone to capture these excitation signals. However, this solution does not work in our case due to the following two reasons. First, the sensitivity of speaker-converted microphone is inferior to purposely built microphones as diaphragms in headphones are well-calibrated for playing sound as opposed to sound recording [14]. Second, the excitation signals are feeble and will be buried in the music signal that is orders of magnitude higher. In HeadFi design, the key observation is that headphones usually come in a pair of matched driver. Based on this observation, we build a differential measurement circuit that enables the headphones to reject the playing music signal but keep the sensing capability at the same time.

Null Measurement

We leverage a passive, null measurement circuit—Wheatstone bridge, to detect the minute variation of $E_{headphones}$.

Wheatstone bridge primer. A Wheatstone bridge consists of two voltage divider arms, each consisting of two simple resistors connecting the power source and ground terminal. Originally this bridge was used to measure an unknown resistance (as small as several milli-Ohms) by tuning the trimpot until the two arms reach to a balanced state (i.e., the output voltage is zero) [8]. As shown in Figure 2, R_1 and R_2 are two identical bridge arm resistors. C_x is the unknown load and C_1 is the trimpot.

[1] We use headphones to represent in-ear, supra-aural (a.k.a., on-ear) and circumaural (a.k.a., over-ear) listening devices throughout the paper.

[2] The driver is the key component of a speaker in headphones and, therefore, driver and speaker are used interchangeably throughout the paper.

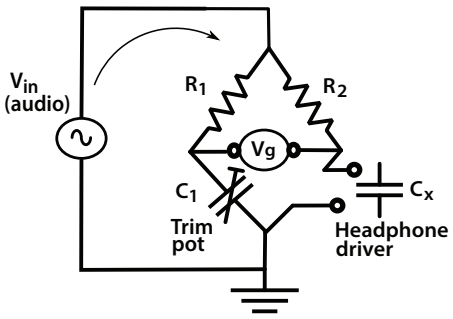

FIGURE 2. Wheatstone bridge.

The trimpot C_1 is tuned until its impedance equals to the impedance of C_x, leading to a "balanced" bridge. In such a balanced state, the voltages on these two loads are the same, resulting in a zero voltage output ($V_g = 0$). Any minute change in the impedance of C_x would alter the voltage on this load and break the balance of the bridge, resulting in a non-zero voltage output (i.e., $V_g \neq 0$).

Detecting minute voltage with the bridge. In HeadFi, we repurpose the Wheatstone bridge to cancel the strong audio input signals and measure the subtle changes in headphone impedance[3] caused by excitation signals. We replace the unknown load C_x in the bridge with the driver of the headphones. The audio input (e.g., music signal) serves as the voltage supply V_{in} to this bridge. Once the bridge is balanced, the voltage output V_g becomes zero. The variation in audio input signal V_{in} does not break the bridge's balance. However, the excitation signals caused by human gestures and physiological activities inherently break the balance of the bridge and alter the voltage measured ($E_{headphones}$). More importantly, the Wheatstone bridge is super sensitive to the voltage variation at the headphones. Thus, we can leverage the variation in the voltage output of this bridge V_g to detect even very subtle excitation signals.

Balancing the Wheatstone bridge. To measure the minute change in $E_{headphones}$, it is important to balance the Wheatstone bridge first. The audio input signal is a wideband AC signal varying over the entire audible band from 20 Hz to 20 kHz.

[3] The voltage change is linearly related to the headphone impedance change.

FIGURE 3. Part of the headphones used in our experiments.

To balance the bridge over this audible band, the trimpot C_1 should be tuned to match C_x —the load of headphones' driver. Accordingly, C_1 should be an RLC type of circuit to match the driver's load. However, in practice, this balancing mechanism is not scalable since different headphones have dramatically different driver impedance values.

We instead leverage the symmetry nature of the drivers to solve this problem. The drivers of headphones come in a pair (*i.e.*, in both left side and right side of the headphones). To ensure a good user experience, each pair of drivers undergo a fine-grained calibration during manufacturing to ensure the impedance of the two drivers are exactly the same. Based on this intuition, HeadFi replaces the trimpot C_1 with the second driver in the headphones, which naturally balances the bridge without introducing any complex tuning circuits. As shown in Figure 3, we experimentally tested the driver symmetry nature by conducting many of our demonstration applications on 54 pairs of different "dumb" headphones with the cost ranging from $2.99 to $15,000.

Physical interpretation of V_g. The pair of drivers in headphones are wired to be in-phase for coherent stereo AC signal. Note that when C_x and C_1 are replaced by the two drivers, the voltage measured at the left driver E_{left} and the right driver E_{right} coming to the bridge are *phase inverted*. That is to say, the voltage output V_g of the bridge characterizes the difference of E_{left} and E_{right}: $V_g = E_{left} - E_{right}$. In some applications (e.g., heartbeat, and breathing monitoring), the excitation signal is picked up by both drivers in the headphones. Hence, a critical question is whether the voltage variation caused by the excitation signals is canceled by the bridge without being detected, *i.e.*, $V_g = 0$. In practice, the excitation signals arrive at these two drivers usually through different paths. Hence, $E_{left} \neq E_{right}$. HeadFi can therefore still leverage this differential voltage measurement to detect the minute excitation signals.

USER IDENTIFICATION

We first demonstrate how HeadFi can be used for user identification. The mainstream identification method, face recognition, does not work well in poor lighting conditions or when the user wears a mark. HeadFi can be leveraged to check the user identity and unlock the phone (pairing device) regardless of the lighting conditions. Face recognition also raises privacy concerns, whereas HeadFi can identify users without taking photos.

Signal Processing

Ideally, an identification service should be non-intrusive, *i.e.*, it should be triggered automatically as long as the user puts

on the headphones. As such, our design should be able to *i)* detect if the user puts on the headphones and *ii)* identify the user automatically.

Headphones ON-OFF detection. Our design is inspired by the *seashell resonance effect* [9]: when a seashell is clasped to the ear, the ambient acoustic noise will resonate within the cavity of the seashell and certain frequency noise will get amplified. One can thus hear ocean-tide-like sounds from the seashell. Similarly, once the user puts on her headphones, the headphones, ear canal, and eardrum establish a resonance chamber, amplifying the ambient acoustic noise. This amplified noise leads to a higher voltage signal output measured at HeadFi. Based on this observation, we use the RSS and its standard deviation (σ) for ON-OFF detection. These two values jump dramatically when the user puts on the headphones.

Identification. Since the headphones can now transmit and receive at the same time, we can proactively probe the ear channel response using the headphones. Specifically, the smartphone sends a chirp signal through the headphones to profile the user's inner ear structure. The two drivers of the headphones receive echo signals that characterize the ear canal's channel response.

As HeadFi measures the voltage difference between the two drivers of headphones, one may wonder whether the channel response from the left ear cancels out that from the right ear. Interestingly, the ear-related physiological uniqueness exists not just between two users, but also between two ears of the same person [10, 11]. Hence the channel response measured at two ears would not be the same. Figure 4 shows the channel response measured by HeadFi on three different persons. We can see the channel responses are dramatically different in frequency bands higher than 3 kHz. This is because the physiological differences between human ears are in the scale of sub-centimeter level, which can be picked up by signals with a wavelength of sub-centimeter (≥ 3 kHz). We further adopt a preamplifier (INA126) to control the output level. As a result, HeadFi can retrieve a clear echo even though the excitation signal is weak.

Proof-of-concept. As a proof-of-concept, we use support vector machine (SVM), a lightweight classifier for user identification. Specifically, we collect multiple copies of the user's echo chirp as positive samples. We then collect the same amount of negative samples by putting the headphones on the E.A.R.S dummy head. Finally, we train a binary SVM classifier and perform *k*-fold.[5]

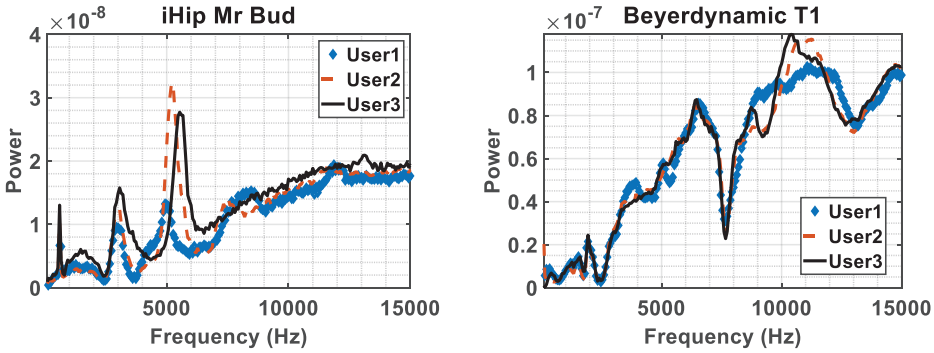

(a) iHip Mr Bud ($3.6)

(b) Beyerdynamic T1 ($999).

FIGURE 4. Channel response of three people characterized by (a) low-end and (b) high-end headphones.

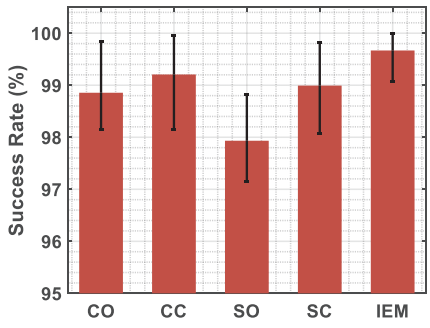

FIGURE 5. ON-OFF detection cross-validation.

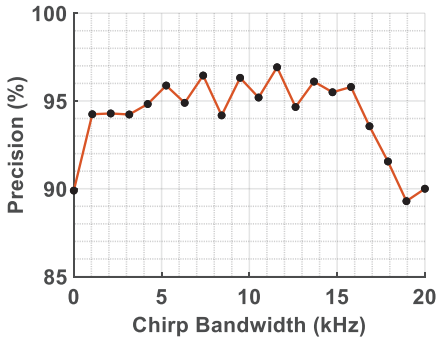

FIGURE 6. Precision test.

Experiment

The experiments involve 27 participants (7 females and 20 males), including one pair of identical twins. By default, we use the Jays U-JAYS supra-aural headphones (MSRP $19.99) as the testing device. The chirp duration is one second throughout the experiments. The participant is asked to put on and then take off the headphones each time we record an echo chirp. We record 50 echo chirps for each of the 25 participants and 100 echo chirps for each of the twins.

ON-OFF detection. We first evaluate the success rate of ON-OFF detection across 54 pairs of headphones. We further categorize the results into five groups based on headphone types and show the results in Figure 5. We observe that the success rate is consistently high (>97.93%) across all five types of headphones. In particular, IEM headphones achieve the highest success rate (99.8% on average) since this type of headphones go deeper into the ear canal and, thus, are less affected by noise.

User Identification. Next, we evaluate the performance of user identification. In each experiment, we adopt *k*-fold (k =5) cross validation to demonstrate the system performance. We adopt precision [6] as our evaluation metric. A high precision value indicates that only the authorized users can successfully pass the verification. Figure 6 shows the precision under different chirp bandwidth settings. When the chirp bandwidth is relatively small (*e.g.*, < 4 kHz),

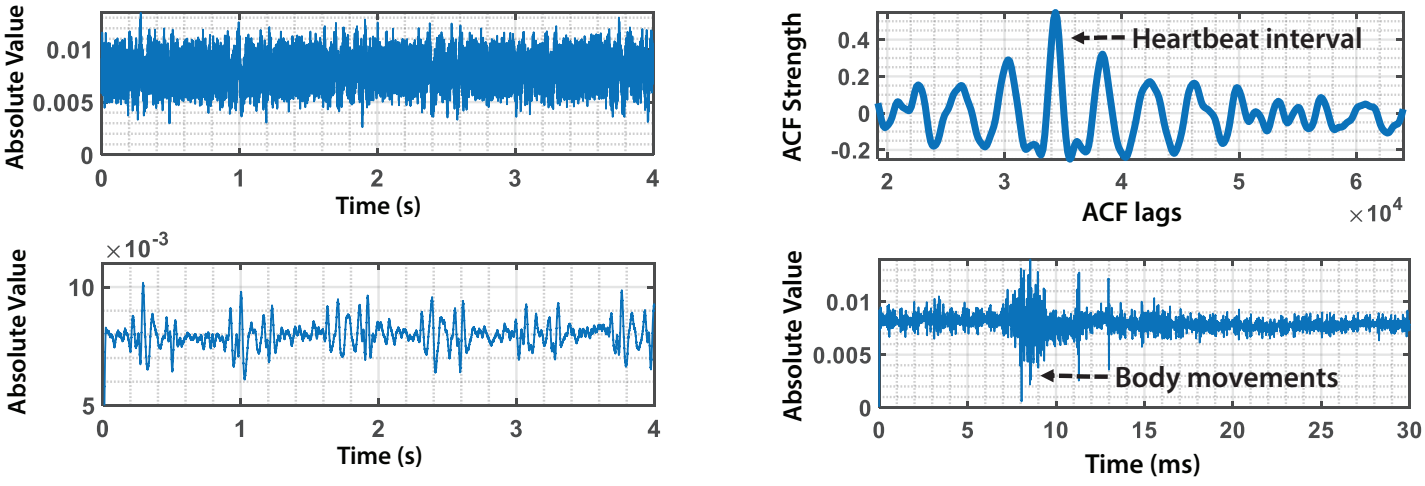

(a) The voltage output before and after filtering

(b) The ACF plot and time domain interference

FIGURE 7. (a) Heartbeat signal becomes clear after filtering (b) ACF is adopted to calculate the heartbeat rate (top) and an example of time domain interference caused by body movements (bottom).

we observe that the precision grows with increasing chirp bandwidth. The precision then fluctuates around 95% as we further increase the bandwidth to 15 kHz. It then drops to around 90% as the frequency bandwidth goes beyond 15 kHz. Such precision variation is due to the subtle changes during multiple rounds of putting on headphones: the sub-*mm* level changes can be captured by the high-frequency (higher than 15 kHz) signal, which disturbs the user identification. Suggested by this study, we employ a frequency band from 100 Hz to 10 kHz as the default chirp bandwidth. We exclude the frequency band below 100 Hz because most mechanical movement-induced noise is in this frequency range.

PHYSIOLOGICAL SENSING

Next, we demonstrate the feasibility of applying HeadFi to detect subtle physiological signals. Vital physiological sign sensing plays a key role in human health monitoring. HeadFi can empower users to monitor a variety of key physiological activities continuously and accurately (*e.g.*, heartbeat rate) using their non-smart headphones. In Figure 7, we take heartbeat monitoring as an illustrative example.

Signal Processing

Monitoring heartbeat is challenging due to the extremely weak excitation signal induced by the subtle blood vessel deformation in the ear canal. As shown in Figure 7(a) (top), such a minute excitation signal can be buried in the noise and interfered by user motions. To solve this challenge, we first pass the signal output from HeadFi through a low-pass filter with a very low cut-off frequency (F_c =24Hz) to remove the high-frequency noise introduced by the echoes of audio input signals and environment excitations. The result is shown in Figure 7(a) (bottom). We then leverage the auto-correlation function (ACF) to identify the periodicity, which corresponds to the heartbeat rate:

$$r_{xx}(k) = \frac{1}{N-k}\sum_{n=0}^{N-1-k} x(n)x(n+k). \qquad (1)$$

where $x(n)$ is a copy of the signal output from HeadFi and k is the lag. N is the length of the received signals. Figure 7(b) (top) shows an example of the auto-correlation output. The location of peak values reflects the time period of one heartbeat cycle. Blindly enumerating all choice of k in hopes of finding the peak is computationally intractable. It may also introduce false positives. We thus set the upper (U) and lower (L) bounds of k based on the possible heartbeat rate of human beings (35 – 200 bpm [17]). Our goal can be represented by the following function:

$$k^{\star} = \underset{k\subseteq(L,U)}{\arg\max}\, r_{xx}(k). \qquad (2)$$

We then calculate the heartbeat rate using the equation $R_{BPM} = 60 \cdot \frac{F_s}{k^{\star}}$, where F_s is the sampling rate. In reality, however, body movements may also introduce strong excitation signals that can overwhelm the minute heartbeat signals, as shown in Figure 7(b) (bottom). We thus truncate the voltage output from HeadFi into windows and calculate R_{BPM} within each window. We then apply an outlier detection algorithm [15] to filter out those outlier estimations and average the remaining to obtain the heartbeat rate.

Experiment

In this section, we evaluate the performance of heartbeat rate monitoring. Each measurement lasts for 40 seconds. We truncate a recording session using a window size of

WE ASK THE QUESTION: CAN WE TURN NON-SMART HEADPHONES IN HAND INTO INTELLIGENT ONES WITHOUT REDESIGNING THE HEADPHONE OR ADDING EMBEDDED SENSORS?

four seconds, with an overlapping of two seconds. We measure the heartbeat rate of the participant in two conditions: *i*) with audio input signal on (*i.e.*, listening to the music during the testing); and *ii*) with audio input signal off. The ground-truth is obtained by a CONTEC CMS50D1A pulse oximeter [2]. We use error rate (ER) to measure the performance of our heartbeat monitoring: $ER = \frac{|R_{HF} - R_{PO}|}{R_{PO}}$, where R_{HF} and R_{PO} are the heartbeat rate reported by HeadFi and the oximeter, respectively.

Heartbeat rate monitoring accuracy. We evaluate the accuracy of the heartbeat rate estimation using all 54 pairs of headphones. In particular, we categorize these headphones into three groups, namely, circumaural headphones (C), supra-aural headphones (S), and in-ear model (IEM). The result is shown in Figure 8(a). We observe that HeadFi achieves consistently low error rate across all three groups of headphones. Circumaural headphones (C) achieve the lowest error rate both in the absence (1.37%, C) and presence (1.42%, C-M) of audio input signals, followed by supra-aural headphones (1.40% and 1.68% in these two cases, respectively). HeadFi achieves the highest error rate for the IEM headphones: around 1.64% and 2.42% in the absence (IEM) and presence (IEM-M) of audio input signals, respectively. While the intrinsic reason behind this performance drop is unknown, one possible reason could be that IEM headphones have less contact area with skins and thus receive the weakest vibration signals compared to the other two types of headphones. The maximum error rate achieved by HeadFi is around 3%, which still satisfies the requirement (less than 5%) for commercial heartbeat monitoring systems [13]. These results demonstrate the feasibility of using HeadFi to measure user's heartbeat rate even in the presence of music.

Impact of body movement. In this experiment, 27 participants (including 7 females and 20 males between 27 to 55 years old) are asked to put on/off the headphones occasionally during the testing, which brings in a strong interference signal. Figure 8(b) shows the error rate. We also show the error rate in the absence of body movement for comparison. We observe a slight increase (0.59% on average) in the error rate in the presence of body movements, while the overall error rate is still less than 3%, well below the requirement for commercial heartbeat monitoring systems (< 5%).

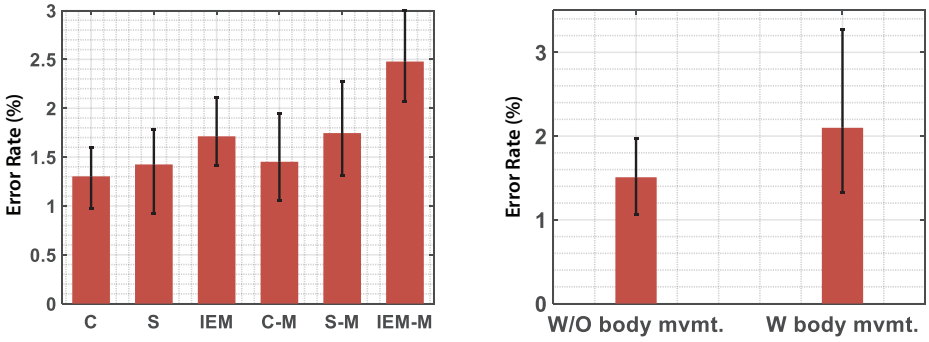

(a) Impact of ambient noise.

(b) Impact of body movement.

FIGURE 8. Error rate of the heartbeat rate estimation. (a) We measure the error rate both in the absence (the first three columns) and in the presence (the last three columns) of the audio input signal. (b) We measure the error rate both in the absence and presence of strong body movements.

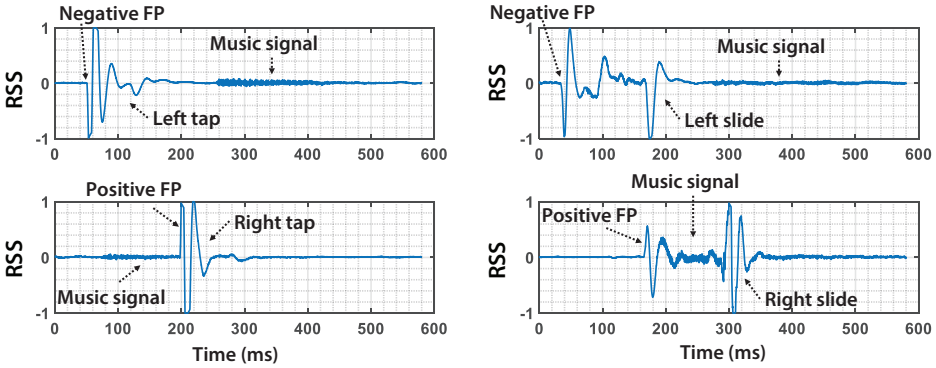

(a) Tapping

(b) Sliding

FIGURE 9. The voltage output signals V_g caused by different touch-based gestures. (a) tapping the left (top) and right (bottom) enclosure (b) sliding on the left (top) and right (bottom) enclosure.

TOUCH-BASED GESTURE RECOGNITION

We next demonstrate the feasibility of transforming the enclosures of the non-smart headphones into virtual touchpads using HeadFi. The rationale behind this is that the variation in the output voltage V_g caused by different gestures manifests unique features in both spatial and temporal domains.

Design intuitions. We invite a volunteer to tap the left and right enclosure of one pair of headphones and record the RSS out of HeadFi. As shown in Figure 9(a), when there is a tap on the headphones, we can always observe multiple RSS peaks. In particular, when the user taps the left enclosure, there is a negative peak followed by a positive peak, as shown in Figure 9a (top). In contrast, the positive peak shows up ahead of the negative peak when the user taps the right enclosure as shown in Figure 9a (bottom). This is because the Wheatstone bridge measures the differential voltage between the two drivers of headphones. Consequently, the excitation signals measured at the bridge are *phase inverted* for right tap and left tap. Note that the echoes of input music signal been recorded by HeadFi are orders of magnitude weaker and would not overwhelm peaks introduced by tapping gestures. Similar to the tapping gestures, left and right sliding can also be easily distinguished based on the

same principle. On the other hand, sliding gestures usually last longer than tapping in time domain, as shown in Figure 9b. We can thus leverage the peak interval to distinguish them. Please refer to our detailed gesture recognition design and evaluation results in the HeadFi full paper.[7]

VOICE COMMUNICATION

Finally, we demonstrate the feasibility of using HeadFi to enable full-duplex voice communication on those headphones without a built-in microphone. As discussed in "Null Measurement," the human voice signals will not be canceled out by the bridge since the voice signals propagate to left and right headphone drivers through two complicated but independent channels determined by air, bones, tissue, etc.

The impact of echoes. One interesting issue that may exist with our design is the echo (here we assume the user on the other side is not using HeadFi for easier explanation). This is because during a voice call, HeadFi captures the voice from not just the HeadFi user side but also the other side at the headphone's diaphragm. Both captured voices will be transmitted to the other side. Thus, the other side may hear an echo of her own voice. Fortunately, this issue is already addressed by the service providers. To provide high-quality voice communication, service providers usually run sophisticated signal cancellation algorithms at the base station to remove echoes before transmitting the voice signals to the receiver [16]. Therefore, echoes would not be a problem and the evaluation results also confirm this. Detailed evaluation results in voice communication are presented in the HeadFi full paper [7].

CONCLUSION

We have presented the design, implementation, and evaluation of HeadFi, a low-power peripheral to bring intelligence to headphones. HeadFi employs the pair of drivers inside headphones as a versatile sensor to enable new functionalities as opposed to adding embedded sensors. This design can potentially upgrade existing non-smart headphones into intelligent ones. We prototype HeadFi on PCB board and demonstrate the potential of HeadFi by showcasing four representative applications using 54 pairs of headphones. ■

Xiaoran Fan is an experimental scientist at Google. He is interested in hearables, including hardware architectures, signal processing algorithms, and software-hardware co-design for novel applications. His research works won three best paper awards, turned into multiple patents, and lead to a start-up company (Ohmic). He received his Ph.D. in ECE from WINLAB, Rutgers University in 2020.

Longfei Shangguan is an assistant professor at the University of Pittsburgh. His research interests are in all aspects of IoT systems: from building novel IoT applications, solving security issues all the way down to optimizing the network stack, and designing low-power IoT hardware. He received his Ph.D. from the Hong Kong University of Science and Technology in 2015.

Siddharth Rupavatharam is a final-year Ph.D. student at WINLAB, Rutgers University. His dissertation focuses on developing novel sensing systems for human-robot and robot-environment interactions. He is interested in developing new sensing modalities for robotics, wireless sensing and mobile computing and has previously worked with modalities, such as visible light, capacitance, acoustics and ultrasonics. He received his M.S. at Rutgers University in 2018.

Yanyong Zhang obtained her B.S. from USTC in 1997, and her Ph.D. from Penn State in 2002, both in Computer Science. In July 2002, she joined the Electrical and Computer Engineering Department and WINLAB, Rutgers University as an assistant professor. She was promoted to an Associate Professor with tenure in 2008, and a professor in 2015. In July 2018, she moved back to her alma mater, School of Computer Science at USTC. Her research interests lie in the intersection of sensing, scheduling, edge computing, and control. Example application areas include autonomous driving and AIoT.

Jie Xiong is currently an assistant professor in the College of Information and Computer Sciences at University of Massachusetts Amherst. His recent research interests include wireless sensing, mobile health, smart IoT and cyber-physical systems. He received his Ph.D., M.S. and B.Eng. degrees from University College London, Duke University and Nanyang Technological University, respectively.

Yunfei Ma is a staff engineer and R&D manager at Alibaba, where his current research focuses on video conferencing, video transport, and QUIC-based Cloud gateway. His research has been deployed in Alibaba's core services, such as Taobao and Dingtalk, and has transformed into several products at AliCloud. Before joining Alibaba, He was a postdoctoral researcher at MIT Media Lab. He received Ph.D. in ECE from Cornell University and B.S. from USTC.

Rich Howard received his B.S. from Cal Tech and his Ph.D. from Stanford University 1970 and 1977, respectively. He is currently the chief technical officer of the Inpoint Systems, Inc., and a research professor at WINLAB, Rutgers University. Before this, he spent 22 years in the research area at Bell Labs, retiring in 2001 as the vice-president of wireless research. His current joint projects with the Rutgers University include developing a new class of wireless nodes for inventory tracking and ubiquitous sensor networks.

REFERENCES

[1] Gig fix: Turn your headphones into a mic. Webpage, 2014.
[2] Cms50d1a gehp040ahus pulse oximeter. Webpage, 2019.
[3] Earphones and headphones market size, industry report. Webpage, 2019.
[4] Global unit sales of headphones and headsets from 2013 to 2017. Website, 2019.
[5] Yoshua Bengio and Yves Grandvalet. 2004. No unbiased estimator of the variance of k-fold cross-validation. *Journal of Machine Learning Research*.
[6] Jesse Davis and Mark Goadrich. The relationship between precision-recall and roc curves. 2006. *Proceedings of the 23rd international Conference on Machine Learning*, 233–240.
[7] Xiaoran Fan, Longfei Shangguan, Siddharth Rupavatharam, Yanyong Zhang, Jie Xiong, Yunfei Ma, and Richard Howard. 2021. Headfi: bringing intelligence to all headphones. *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*.
[8] Karl Hoffmann. 1974. Applying the Wheatstone bridge circuit. HBM Germany.
[9] Liang-Ting Jiang and Joshua R. Smith. 2012. Seashell effect pretouch sensing for robotic grasping. IEEE Robotics and Automation Society, *Proceedings of the International Conference on Robotics and Automation (ICRA)*.
[10] Agnès Job, Paul Grateau, and Jacques Picard. 1998. Intrinsic differences in hearing performances between ears revealed by the asymmetrical shooting posture in the army. *Hearing Research*.
[11] F. Laurain King and Doreen Kimura. 1972. Left-ear superiority in dichotic perception of vocal nonverbal sounds. *Canadian Journal of Psychology/ Revue canadienne de psychologie*.
[12] Sean Olive, Omid Khonsaripour, and Todd Welti. A survey and analysis of consumer and professional headphones based on their objective and subjective performances. 2018. *Audio Engineering Society Convention 145*. Audio Engineering Society.
[13] Alexandros Pantelopoulos and Nikolaos G. Bourbakis. A survey on wearable sensor-based systems for health monitoring and prognosis. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*.
[14] Melih Papila, Raphael T. Haftka, Toshikazu Nishida, and Mark Sheplak. 2006. Piezoresistive microphone design pareto optimization: Tradeoff between sensitivity and noise floor. *Journal of Microelectromechanical Systems*.
[15] Peter J. Rousseeuw and Annick M. Leroy. *Robust Regression and Outlier Detection*. John Wiley & Sons, 2005.
[16] Voice Quality Enhancement and Echo Cancellation.
[17] Joseph C. Volpe Jr. Apr. 8, 2008. Heart rate monitor for controlling entertainment devices. US Patent 7,354,380.