Continuous Measurement of Interactions with the Physical World with a Wrist-Worn Backscatter Reader

ALI KIAGHADI, University of Massachusetts Amherst
PAN HU, Stanford University
JEREMY GUMMESON, SOHA ROSTAMINIA, and DEEPAK GANESAN, University of Massachusetts Amherst

Recent years have seen exciting developments in the use of RFID tags as sensors to enable a range of applications including home automation, health and wellness, and augmented reality. However, widespread use of RFIDs as sensors requires significant instrumentation to deploy tethered readers, which limits usability in mobile settings. Our solution is WearID, a low-power wrist-worn backscatter reader that bridges this gap and allows ubiquitous sensing of interaction with tagged objects. Our end-to-end design includes innovations in hardware architecture to reduce power consumption and deal with wrist attenuation and blockage, as well as signal processing architecture to reliably detect grasping, touching, and other hand-based interactions. We show via exhaustive characterization that WearID is roughly 6× more power efficient than state-of-art commercial readers, provides 3D coverage of 30 to 50 cm around the wrist despite body blockage, and can be used to reliably detect hand-based interactions. We also open source the design of WearID with the hope that this can enable a range of new and unexplored applications of wearables.

CCS Concepts: • **Applied computing** → **Health care information systems**; • **Human-centered computing** → *Ubiquitous and mobile computing systems and tools*;

Additional Key Words and Phrases: Wireless sensor networks, IoT

ACM Reference format:
Ali Kiaghadi, Pan Hu, Jeremy Gummeson, Soha Rostaminia, and Deepak Ganesan. 2020. Continuous Measurement of Interactions with the Physical World with a Wrist-Worn Backscatter Reader. *ACM Trans. Internet Things* 1, 2, Article 7 (April 2020), 22 pages.
https://doi.org/10.1145/3375800

1 INTRODUCTION

The ability to monitor tactile interactions between people and objects is important for a range of applications including home automation, health and wellness, smart spaces, augmented reality, and tele-rehabilitation. Perhaps the simplest way to monitor such interactions is by using passive UHF RFID tags that can be cheaply attached to objects. Recent work on interactive RFID systems has shown that such tags can be used not only for identification but also for sensing the type of interaction by analyzing low-level channel parameters like phase and received signal strength

This research was partially funded by NSF award #1763524 and NIH award #R01MH109319.
Authors' addresses: A. Kiaghadi, P. Hu, J. Gummeson, S. Rostaminia, and D. Ganesan, Computer Science Building, Amherst, MA 01002; emails: akiaghadi@umass.edu, panhu@stanford.edu, {gummeson, srostaminia, dganesan}@cs.umass.edu.

2577-6207/2020/04-ART7 $15.00
https://doi.org/10.1145/3375800

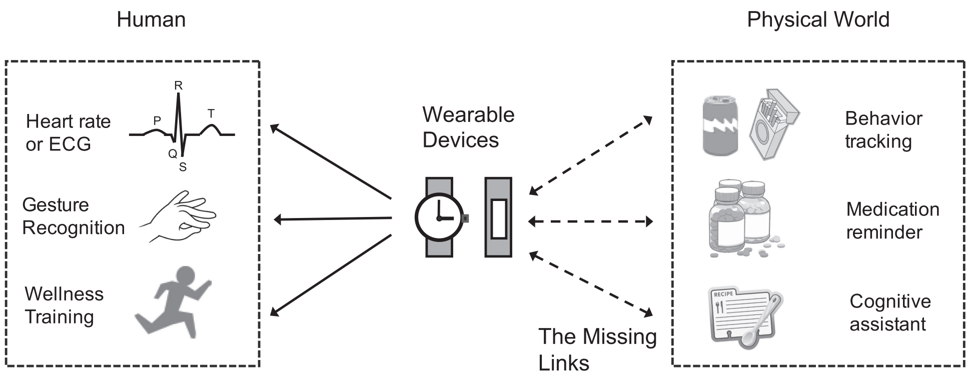

Fig. 1. Wearables provide extensive information about physiological and movement signals at the worn location. But they provide little information about interactions with objects in the physical world.

indication (RSSI). This makes it possible to detect whether the interaction involves touching a tag [26], blocking a tag [27], or moving a tagged object [62], as well as relative orientation with respect to a tagged object [54]. This has applications in many scenarios including interactive smart homes [15], human-robot interaction [29], and battery-less user interfaces [45].

Although an expanding set of applications use RFIDs as a sensor for measuring interaction, the implicit assumption is that it is easy to deploy readers in the infrastructure. This is a restrictive assumption, particularly because it does not extend to ambulatory situations where RFID-based sensing can be very useful. In contrast to infrastructure-based readers, a wearable reader that is integrated into a device like a wristwatch can bridge an important gap in wearable technologies. But wearables do not enable such sensing. As illustrated in Figure 1, today's wearable technologies are primarily focused on measuring physiological and movement signals rather than monitoring interaction with external objects.

The figure shows two classes of applications in which WearID can have significant utility in a wearable scenario. The first class involves monitoring of interaction with tagged objects without explicit user interaction. For example, automated journaling using cheap tags on soda cans, alcoholic beverages, chips, and cigarette packs can help with behavior tracking and modification for alcoholics, smokers, and binge eaters. Such tracking can also be beneficial to monitor medication adherence; studies have shown that persistent and consistent adherence is needed for optimal clinical outcomes [43]. WearID can also be useful as a cognitive assistant that tracks the sequence in which tagged objects are used and can provide instructions for furniture assembly, food preparation using a recipe, and daily routines for the elderly. A cognitive assistant can also be useful at the workplace, for example, to track the sequence of objects a physician might want to interact with during a medical procedure. WearID can also be useful in pill tracking (UHF RFIDs have recently been embedded in smart pills [4]), as well as smart clothing and garments [36]. Thus, everything from smart pills to smart clothing may be equipped with embedded RFIDs, making such a wearable reader a crucial component of the wearable ecosystem.

We face two challenges in achieving this objective. First, we need to design a practical wearable backscatter reader that operates within the form-factor and power constraints of a smartwatch-class device. The downside of a typical UHF RFID reader is that it consumes a lot of power. Second, such a device should have a relatively narrow field of view to capture signals from tags that an individual interacts with rather than a large number of tags that may be in the vicinity of a user. Although the RFID industry has been growing steadily as object tracking and IoT have become more pervasive, commercial readers are intended mostly for tethered operation and need to use high transmission power to read tags over tens of feet in cluttered environments. Our goal is to design a reader that is optimized to measure short-range tactile interactions while being small form factor, low power, and robust to occlusion by the hand.

Our work builds on a large body of research in the area of backscatter circuits and systems [18, 30, 47, 56]; however, the main innovation is in enabling an end-to-end wearable RFID reader that is low power, small form factor, and exposes phase and RSS output for RFID sensing applications. Our work is inspired by early work on the design of glove and bracelet-sized NFC readers for activity recognition [13]. However, an NFC-based approach has two limitations: (1) NFC operates at a long wavelength of 22 m and does not provide useful phase information, and (2) NFC operates at short ranges of a few centimeters (with wearable form-factor antennas) and is therefore more suited for intentional interactions like credit card payment.

Summary of results. In summary, we design WearID, a wearable, low-power UHF backscatter reader that is designed to detect interactions with tagged objects. Our end-to-end implementation of WearID includes a hardware prototype that is optimized for power, form factor, and performance and a signal processing/machine learning pipeline to classify various interactions. Our experiments with WearID show the following:

- WearID can provide IQ (in-phase, quadrature) output to enable a variety of fine-grained RFID applications while consuming 6× less power than best-in-class commercial readers.
- WearID can provide 3D coverage of 30 to 50 cm despite hand blockage, thereby enabling robust monitoring of interactions with objects in the hand.
- WearID can reliably detect various interactions including grasp and release of objects, touch of a tag, and passing near a tag.

2 RELATED WORK

Our work builds on or relates to a substantial body of literature in two broad categories: (1) infrastructure-based interaction detection methods using RFIDs and other modalities, and (2) wearable-based interaction detection using NFC/RFID- and non-RFID-based techniques.

We start by describing methods for detecting interactions by leveraging infrastructure-mounted devices including RFID readers, WiFi APs, and other devices. We also provide a summary of this work in Table 1.

Infrastructure-mounted RFID readers. There has been substantial interest in measuring interaction with RFIDs via infrastructure-deployed RFID readers. The work in this domain aims to detect various gestures and interactions using phase and RSSI information of the read tags. For example, GRFID measures swipe interactions in front of a tag by looking at changes in phase [62], RF-IDRAW tracks the angle of arrival (AoA) of an RFID tag on a user's finger to track text drawn in the air [52], RFCompass mounts an RFID reader on a robot and uses RFID tags on objects to assist with robot navigation and orientation sensing [54], IDSense classifies user interactions with a tagged object by extracting features from phase and RSSI information [28], and Tadar enables through-wall RFID-based tracking by leveraging low-level RF signals and several reference tags [57].

Our work on using a wrist-worn RFID reader for interaction detection is inspired by this body of work but adds a dimension that is currently lacking. Although infrastructure-based readers do not require the user to wear any device, they often involve substantial overhead for instrumenting the space. The coverage area of a reader is bottlenecked by the ability to deliver sufficient power to a passive tag, particularly in conditions where there is occlusion due to the human body between the reader and tag. This necessitates careful reader placement with multiple directional antennas to provide adequate coverage. This limitation is evident from a review of how prior work has

Table 1. Overview of Infrastructure-Based Activity Recognition Methods

Name	Infrastructure Needs	Interaction Area	Signal Features
Methods based on UHF RFID technology			
GRFID [62]	Wall-mounted reader; multiple reference tags on wall	2 m reader to tags; 70–90-cm tags to user	Signal phase
RFIDraw [52]	Wall-mounted reader; tagged finger	Small room	AoA
IDSense [28]	Ceiling-mounted reader; tagged objects	Small room	RSSI, phase, Doppler shift, tag ID
RF-Compass [54]	Robot-mounted reader; tagged robot and objects	2–6 m	Multipath profile
Tadar [57]	Reader behind wall; tags on wall	4–6 m in front of wall	Phase and RSS
Methods based on WiFi technology			
CARM [55]	AP and laptop	42 m^2 lab	CSI
WiGest [8]	AP and laptop	≤30 cm above laptop	RSS
WiSee [37]	Multiple APs	Whole home	Doppler shift
WiDraw [46]	AP and laptop	2 ft in front of laptop	AoA
MultiTrack [48]	Multiple WiFi devices	70 m^2 classroom	CSI

evaluated RFID-based sensing methods. The evaluation involves placing directional antennas carefully to point in the direction of the RFID tags being sensed, and ensuring that the human body does not occlude the line of sight (LoS) path between the reader and the tag. For example, GRFID [62] was evaluated when a user performs the actions in field of view of the used directional antenna without occluding the LoS path; the room in Li et al. [28] has a ceiling-mounted RFID reader with antenna facing the tagged toys; the interactions in Li et al. [26] are performed with a reader antenna underneath the table and facing upward; and so on. WearID removes this restriction since the reader is on the hand and not blocked by the body while the user is interacting with an object.

WiFi and vision-based detection. There has been a lot of work on tag-free methods for inferring activities and interactions. One such technique that has been used frequently is the repurposing of WiFi signals to perform activity and gesture classification by leveraging the fact that body movements change the channel. For example, WiSee translates observed Doppler shift into gestures [37]; WiGest relates RSSI of the WiFi signal to gestures [8]; CARM utilizes changes in CSI in WiFi transmissions to identify human activities [55]; and WiDraw [46] uses AoA, RSS, and CSI of each channel to track hand position.

These methods lead to coarse-grained detectors since it is often difficult to separate reflections due to the human and object from a sea of other reflections and multipaths. In particular, the movement of a human hand/fingers is notoriously difficult to resolve using wireless reflections. In contrast, WearID provides very specific information about phase/RSSI of objects in the immediate vicinity of the hand, allowing us to extract relative movement of the hand and the object.

There are many other modalities aside from WiFi that can be used for remotely sensing interaction. One common approach is to leverage image or depth information, and process the stream to extract information regarding objects, people, activities, and gestures (e.g., [25, 44]). Although vision/depth cameras are an excellent tool for such research, two key limitations are that they require LoS and that they are not always easy to deploy in the "wild" due to privacy concerns.

2.1 Wearable-Based Methods

We now turn to work on leveraging wrist-worn wearables to detect interactions. We classify this work into two sub-areas: use of a mobile RFID/NFC reader and passive sensing using a smartwatch.

Activity recognition using mobile RFID reader. An alternative to an infrastructure-based RFID reader is the use of a wearable RFID reader. Much of the work with wearable RFID focuses on the use of near-field NFC technology rather than far-field RFID (unfortunately, both of these are referred to as RFID, leading to confusion in distinguishing between the technologies). The use of a wearable NFC reader for activity monitoring was proposed nearly a decade ago by Smith et al. [42]. For example, RFIDGlove [32] is an NFC reader with a large 10-cm loop antenna around the hand, and with a read range of a few centimeters. This was further extended in iBracelet [13], which shrank the device to the size of a wrist-worn band and used a circular loop antenna on the wrist to get higher range of about 10 cm.

There are two key downsides to using NFC for such applications. The first is that much of the recent work on using RFID as a sensor relies on phase information that can be used to track small changes in relative distance. This is possible because UHF RFID operates at 915 MHz and has a wavelength of about 33 cm. But NFC operates at 13.56 MHz and has a wavelength of about 22 m, so phase changes due to small movements are not observable using this technique. The second is that range is directly dependent on the size of the loop antenna, and it is difficult to place larger antennas on a wristband. This places a hard limit on the distance over which NFC can work from a wearable device (commercial NFC readers on smartwatches operate only at a couple of centimeters). A UHF RFID reader does not have such limitations and can increase range as needed by increasing the transmit power.

There is also some work on small and low-power UHF RFID readers. In the commercial sphere, reader chips have evolved considerably in their ability to power tags at longer distances and scale to large numbers of tags. But power has not been a dominant design consideration. As a result, although handheld readers for inventory management are commonplace, these are not sufficiently low power to be integrated into a wearable device like a smartwatch.

Our work also builds on recent efforts to design COTS-based backscatter readers [18, 33]. WearID improves on these efforts in several ways: WearID provides IQ signals for RFID-based sensing, whereas Braidio and the device of Nikitin et al. [33] use an envelope-based receiver that is non-coherent and insensitive to phase. The devices are also designed without form factor as a constraint and make design choices (e.g., multiple antennas in Hu et al. [18]) that make it difficult to scale down into a wearable form-factor device.

Some chip-level proposals for low-power RFID reader designs have also been presented in the literature. For example, Ye et al. [58] present a chip-level design of a reader that consumes 160 mW and has a output power of 4 dBm, although most are designed with RFID identification rather than sensing and do not provide IQ output. Perhaps the main difference is that our design is meant to be practical to enable prototyping and exploration of new applications.

Wearable-based methods. A class of techniques that has seen significant research in recent years is measurement of interaction using smartwatches and other wearables. Table 2 lists several such approaches—these primarily track gestures but also sense signals emitted or induced from the object via vibration, electromagnetic, acoustic, or capacitive sensors.

Our work complements such passive sensing approaches in that we can provide a direct and unambiguous signal about the interaction, which helps to precisely localize temporal windows when an interaction occurs. This is often a challenge for passive sensing–based methods where

Table 2. Survey of Work on Wearable/Low-Power RFID Readers and Wearable-Based Methods for Detecting Interactions

Modality	Prior Work	Comments
Wearable RFID readers		
Near-field NFC RFID	RFIDGlove [32], iBracelet [13]	NFC reader integrated into a wearable glove and wrist-worn device. Range of 10–15 cm. No phase output and hence not useful for sensing applications. High power output and/or large loop antenna for desired range (≈200 mW for 10-cm range [13]).
Far-field UHF RFID	Simple low-cost RFID reader, Braidio [18]	Only ID output, no phase output for RFID sensing; large form factor and not designed for body wear.
Passive sensing via smartwatches		
Inertial	RisQ [35], ArmTrack [41], TypingRing [34], Osense [9]	Use inertial sensors to infer hand/shoulder trajectory and detect gestures and behaviors. Complementary and can be fused with RF signals from WearID.
Electromagnetics	EMSense [24], Electric Field [11, 61]	Leverage electromagnetic emissions from devices, particularly powered ones like drills.
Acoustics	GestureRing [16], BodyBeat [38]	Sound is used to either detect motion along surfaces or to classify different sounds generated by the human body.
Vibration	Viband [23]	High-speed accelerometers are used to classify bio-acoustic events.
Body	Skintrack [60] Skinput [17]	Classify interactions with human skin and use as a general-purpose input mechanism.
Capacitive	Touche [40], SignetRing [53]	Wearable device capacitively couples with the environment or capacitive touch screen. The resulting channel can be used for communications or to classify the surroundings.

the start and end of an interaction is difficult to precisely isolate from the stream of data. It may, of course, be possible to further improve our ability to characterize interaction with objects by combining these signals with WearID.

3 OVERVIEW OF WEARID OPERATION

The central advantage of WearID is that it provides a wearable form-factor device that can sense interactions with objects that are tagged with backscatter devices such as passive RFID tags or computational RFIDs like the WISP [59] while providing physical layer information regarding the communication link. Figure 2 shows an application example where WearID can detect the particular object being picked up among multiple objects.

Cognitive assistance. A broad class of applications seek to assist a user by remembering a set of tasks that have been previously performed or tracking tasks in real time as they are performed. These tasks could range from following a recipe to prepare a meal, following instructions to assemble a piece of furniture, or a medication adherence reminder service that informs a user as to which medications were taken earlier in the day.

WearID is suited for cognitive assistance applications since it can be used to track tasks that involve the manipulation of physical objects with attached RFID tags. The unique identifier stored in each tag is used to determine *which* tag is being interacted with, whereas signal-level information is used to determine *when* the tag is interacted with and *what* type of interaction is being

Fig. 2. WearID being used to distinguish various objects.

performed. Although inertial sensors can be used to capture and label specific gestures such as hand to mouth [20, 35], it can be quite difficult to correctly label an interaction without additional information; WearID can provide both of these functions simultaneously.

Just-in-time interventions. Another use case of WearID is in the context of health monitoring. There has been substantial interest in just-in-time adaptive interventions (JITAI) for addictive behavior [39, 51]. For example, the Sense2Stop intervention for quitting smoking involves monitoring stress and hand-to-mouth gestures, and triggering support at vulnerable times when a user may lapse [49]. The challenge is the lack of a reliable signal immediately prior to a smoking lapse—hand-to-mouth gestures are detected *after the fact*, for instance, after the individual has lapsed, and stress can occur at many times other than just before a smoking lapse, making it difficult to isolate the specific event. By tagging the cigarette pack, WearID can be used to provide such a signal prior to a smoking lapse. We evaluate this use case in Section 5.2.

Smart pill bottle. A major challenge in the healthcare system today is tracking prescription medication use, particularly for opioids. Compulsive opioid abuse is a significant problem in the United States and has led to several large initiatives that seek to curb and manage their use. Since pill bottles can be easily and cheaply instrumented with RFID tags, WearID can be useful to track medication use among patients who are newly prescribed opioids to manage pain. The combination of timing information about opioid medication intake together with information about anxiety and craving from physiological parameters like heart rate and breathing rate that wearables provide can potentially lead to enhanced understanding of the propensity for addiction. In turn, this can lead to better addiction management strategies.

4 WearID: A LOW-POWER WRIST-WORN READER

The design of WearID involves several optimizations in terms of power consumption, form factor, and range. In this section, we describe the salient design and performance results that pertain to the interaction performance of WearID and defer a detailed look at the low-level performance optimizations and benchmarks to Section 6.

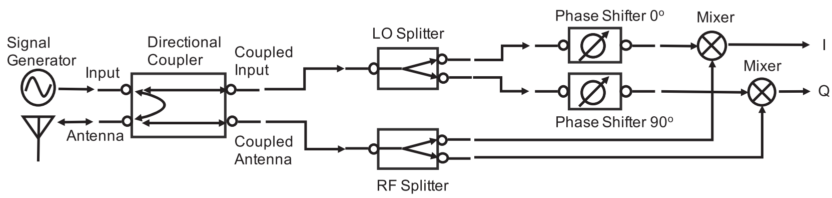

Fig. 3. Depiction of the WearID backscatter transceiver. All components (except signal generator) are passive. A microwave component called a *directional coupler* removes most of the leakage between the transmitter and receiver while still sharing a single antenna. The coupled signal is split for the I and Q mixer input. The Q channel local oscillator (LO) is phase shifted by 90 degrees before being fed into the mixer. A delay line is used as a passive phase shifter to avoid the use of active components. The coupled antenna signal is then fed into the mixers to obtain I and Q in a passive manner.

4.1 Optimizing WearID Power Consumption

The first key consideration in the design of WearID is optimizing power consumption. A typical smartwatch has a 200- to 300-mAh battery, whereas a typical commercial RFID reader chip with IQ output consumes between 640 mW and 2 W [2]. This implies a lifetime of at most 1 hour.

WearID leverages the fact that interaction detection requires limited working range—unlike tethered readers that need to power and read tags tens of feet away, the desired working range of WearID is only a few tens of centimeters to power and read tags on objects near the hand. Our design of WearID leverages this observation to scale down both transmit power and receive sensitivity to optimize power consumption while not compromising its ability to obtain signals of interest for interaction detection. Specifically, WearID should be able to extract signal information that is needed for RF-based interaction classification. In particular, phase information is essential and has been shown to be valuable for many RF-based classification problems [26, 27, 29].

At a high level, WearID relies on a low-power design that is almost entirely constructed from passive components. Figure 3 shows WearID's low-power receiver pipeline. It consists of a directional coupler, splitters, delay components, and mixers, all of which are passive components and consume zero power. The I and Q signals are fed into a baseband amplifier that does consume power, although only a few hundreds of micro-watts. The carrier emitter is shared with the transmitter circuitry and therefore does not add additional cost to the receiver.

Although the building blocks of our passive receiver are also commonly used in RF circuits [18, 19, 21], we synthesize these modules into a practical, open -source platform [7] that exposes useful low-level IQ information, and can enable new research in wearable RFID sensing and interaction.

Our implementation of WearID is shown in Figure 4. The size of this PCB is 42 × 42 mm, which is perhaps on the larger side for a watch form-factor device, but we believe that can be further optimized through integration. In the figure, we can see that the PCB mainly consists of three modules: the carrier emitter module, our optimized passive receiver module, and the main controller module. At the maximum output power of WearID (9 dBm), the system power consumption is 72 mW, of which more than 98% is the carrier generator and less than 2% is the IQ receiver. This is about 6× less power than low-power off-the-shelf readers like AS3993. More detailed performance benchmarks of WearID are provided in Section 6.

4.2 Detection Range of WearID

We now turn to the primary question from an interaction perspective: Does WearID have good coverage along directions that we care about for detecting interaction with tagged objects?

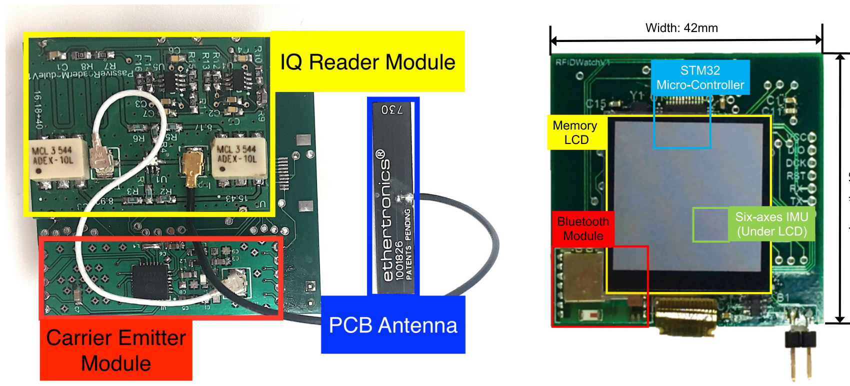

(a) Back View (b) Front View

Fig. 4. Back (a) and front (b) view of the WearID PCB.

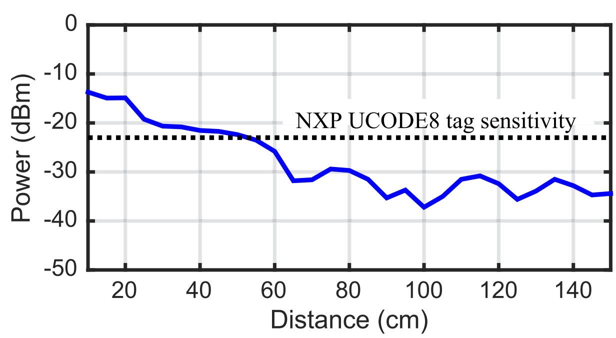

Fig. 5. Power in front of WearID.

Specifically, we care most about the direction in front of the hand, as this is how a user approaches an object to interact with it.

Signal strength across distance. To measure the amount of overall signal attenuation, we perform an experiment where we measure signal strength at different distances in front of WearID. We use the rigid PCB antenna for WearID since it performs better under dynamic conditions, and we measure distance along the front of the hand, which is the most useful direction for interaction sensing. At each distance, we take an average of 10 measurements.

The results are shown in Figure 5. State-of-the-art passive RFID tags require −23 dBm of received power (e.g., the NXP UCODE 8 tag [3]), which results in around 50 to 55 cm of read distance as shown in the figure. Other popular RFID tags such as the Impinj Monza R6 and the Alien Higgs4 are based on slightly older technology (2014 release date, and hence they have slightly lower sensitivity of −22.1 dBm and −20.5 dBm, respectively. For these tags, WearID achieves a range of roughly 30 to 45 cm.

Radiation pattern around the wrist. A more complete view of the radiation pattern across all directions is shown in Figure 6. The 3D surface represents the region within which the signal strength is higher than the target threshold of −23 dBm and can be scaled appropriately for other thresholds. These results were collected for one user in a Qualisys motion capture facility with 0.5-mm ranging accuracy. We place WearID at the origin and move a tag to different points in the 3D space around the tag. The coverage volume of WearID varies in 3D space due to hand blockage,

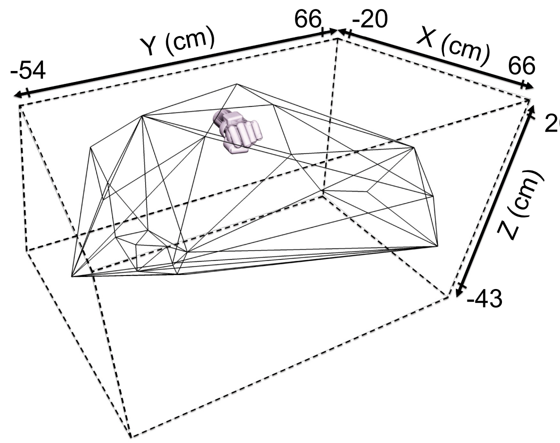

Fig. 6. We measured the radiation pattern of WearID while worn on a user's wrist. The origin is at the top of the smartwatch. Although the hand blocks the signal in some directions, WearID provides excellent coverage of roughly 1.2 m lateral, 86 cm along the hand, and 63 cm vertical. This provides focused coverage of short-range interactions without being cluttered by tags in the vicinity but that are not part of the interaction.

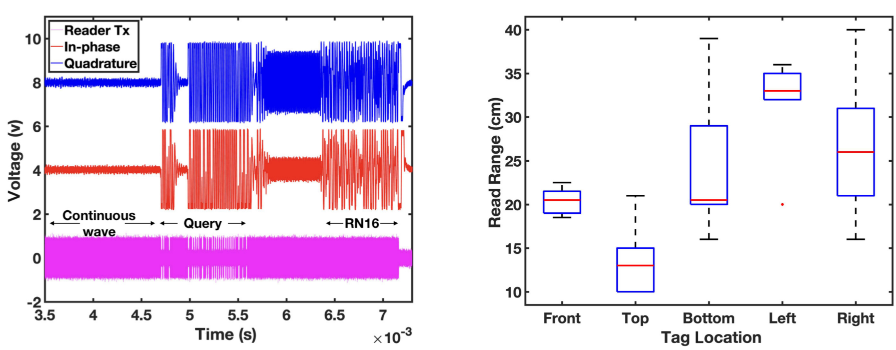

(a) Query msg/response from commercial tag to WearID (b) Read range for commercial tags

Fig. 7. Back and front view of the WearID PCB.

antenna placement, and curvature, as well as skin adjacency, so this result is the accumulation of all of these factors.

The results show that we can obtain the desired coverage in regions where it is needed. For example, the coverage is very good not only in front of the hand but also along the lateral and vertical axes. This shows that the area where the hand is approaching a potential object has good coverage, and WearID can be used for a variety of interaction detection applications.

Detection range for commercial tags. So far, we have looked at signal strength rather than the ability to read commercial tags. To evaluate the tag reading coverage of WearID, we implemented the query message of the EPC Gen-2 protocol to activate commercial tags and obtain a response containing RN16. Figure 7(a) shows a query message from the WearID reader to a commercial tag and the RN16 response from the tag. To determine detection range, we place an Avery Dennison AD-237r6 tag [1] around the wrist in different directions and record the distances in which the tag successfully responds to the query message.

Table 3. Examples of Interaction Scenarios Involving WearID and the Detection Primitives Necessary to Support These Applications

Applications	**Primitives**	**Description**
Home Automation	Touch	RFID tag as a temporary-use home automation switch—for example, "touch" interaction with a tag can be mapped to closing the garage door or turning on lights
Food Journaling Medication Reminder Behavior Tracking	Grab, Release	Detect when an individual picked up a food item or beverage ("grab") and when they set it down ("release").
Cognitive Assistant Lifelogging	Grab, Release, Pass	Monitor interaction and proximity to tagged objects to provide reminders and guidance—for example, as a component of an elder care facility or for lifelogging.

Figure 7(b) shows the result. We see that the read range is highest when approaching it from the left where the range is about 35 cm; front, bottom, and right approaches also have good performance with range between 20 and 30 cm. The worst case is when the tag is directly on top of the reader where the range drops to 10 to 15 cm. In general, the performance is quite good along most directions.

This result shows that we should be able to detect close proximity to tagged objects, as well as the interaction, allowing us to explore a variety of interaction-based applications. We expect these numbers to improve further with technology trends and further hardware improvements. For example, technology trends in the past 10 years suggest that tag sensitivity improves by roughly 1 dBm per year and should increase up to about −30 dBm within the next decade [5]. This will approximately double the detection range of our wrist-worn reader at the current output power level or alternately allow us to reduce transmit power to achieve the same range.

5 INTERACTION DETECTION

We now look at how we can detect a variety of interactions that are uniquely enabled with WearID. Although interaction classification using RFID-based sensing has been explored in prior work, this has been in the context of infrastructure-based readers. The signal from a wrist-worn reader is very different in that it is primarily affected by mobility of the wrist and occlusions due to the hand rather than multipath due to surrounding objects and occlusions due to the body blocking the path from reader to tag in the case of infrastructure sensing. Our goal is to demonstrate that various interactions can be robustly classified with the signal from a wearable reader to a tag.

5.1 Interaction Primitives

Table 3 gives an overview of several interactions and breaks them down into specific detection problems that we need to solve to be able to determine that a particular type of interaction occurred. For example, in the context of food journaling, we need to be able to detect when an individual picked up a food item or beverage ("grab" primitive) and when they set it down ("release" primitive). We find that many applications of WearID can be enabled if we can classify five specific types of interactions: grabbing a tagged object (grab), releasing a tagged object (release), briefly touching a tag (touch), passing near a tag (pass), and none of the above (idle). These interactions can occur in a sequence—for example, grabbing an object may be followed after a period of time by releasing it. We now look at how to classify these five types of interactions by using RF-based RSSI and phase features.

Which RF-based features are useful? The IQ output of WearID can provide instantaneous and relative phase information, as well as signal strength, from which many features can be derived:

- *Phase-based features*: Although our initial expectation was that the instantaneous phase would be a useful proxy for distance (given that the working range is only one to two wavelengths), this turned out not to be the case. The main issue is that the antenna placement in WearID introduces signal disturbances that can vary over time due to factors like tightness of the watch band, sharpness of the antenna curvature, its distance from the skin, and antenna rotations.

  Many RF-based classification systems use relative phase—that is, they assume that signal disturbances do not vary over a short timescale and can be subtracted away. We find that relative phase is better than absolute phase, but we also find that multipath effects vary considerably as a user approaches a tag. Despite these issues with the phase signal, the different types of interactions do seem to produce different types of changes in relative phase, so this is still a useful signal for our classifier.
- *Signal strength-based features*: The RSS is often too noisy to use when the reader and tag are far apart but quite useful in our case. We find that RSS contains reliable information regarding presence of a nearby tag. The SNR (difference between RSS and noise floor) is also particularly useful as an indicator of presence or absence of tags, particularly when we want to quickly detect tag presence without packet exchanges at the protocol level.
- *Temporal features*: The time series of phase and RSS can be used to extract many useful time-series features. When approaching, changes in phase and RSS can be used to differentiate between activities like catching a ball, grabbing a pen, or picking up a tea cup; when moving away, they can help differentiate between placing, dropping, or throwing an object. We can also extract temporal features regarding how long a tag remains in the vicinity of the hand, when an object was grabbed versus dropped, and what the phase and RSS variation were while the object was held in the hand. These can help differentiate between activities like pressing a key, drinking, or passing a tag.

Feature extraction. To extract useful features from noisy signals, several signal processing steps are required. Initially, raw signals need to be segmented so that feature extraction is performed over a meaningful period. To detect the presence of a tag, we use a sliding window over the RSS of the received signal until a tag is detected. Then the window is adjusted so that local maximum value is placed in the middle of the window, and RSS and phase features are extracted over this window.

To extract RSS features, a low-pass filter is applied to the signal to eliminate high-frequency noise. Next, an envelope detector is used as an indication of relative RSS, as shown in Figure 8. The signal is then normalized and separated into segments of equal duration. Local peaks and valleys in each sub-segment are considered as features to be fed into the classifier pipeline.

Phase features are extracted in a similar manner. In periods where RSS is weak, phase is ignored, and phase is tracked when a tag is detected and RSS passes a dynamic threshold relative to the noise floor. We used five samples from the overall pattern of the phase signal as features for our classifier.

Classifying interactions. We now turn to classifying interactions using the preceding features. We look at four interactions—grab, hold, touch, and pass—and show how well we can distinguish between these classes. To accomplish this, we place a backscatter tag in five different settings:

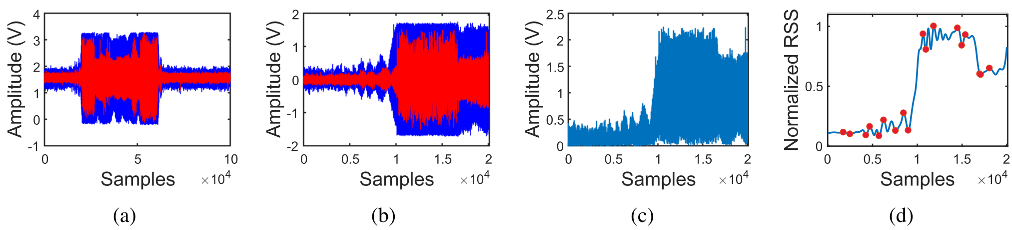

Fig. 8. Signal processing pipeline to capture RSS features. (a) The I/Q signal of a 10-second recorded holding interaction with 10 KSa/s. (b) The 2-second window chosen for feature extraction. (c) Amplitude of the signal derived from I/Q. (d) Obtained after a low-pass filter, an envelope detector, and a normalizer. The dotted points are the chosen RSS features for classification.

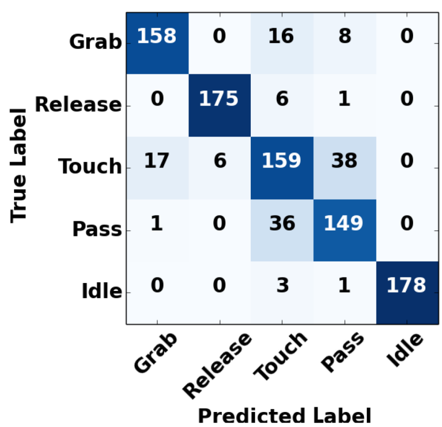

Fig. 9. Confusion matrix for an SVM classifier.

(1) tag on a soda can and grab the can, (2) tag on a soda can and release the can, (3) tag on the wall and touch an object (similar to touching a light switch), (4) tag on an object and pass by the object, and (5) other random non-interactions.

We collected approximately 1,000 data traces from 14 users (8 male, 6 female) across all interactions. We use a leave one subject out (LOSO) method for classification. We train an SVM classifier, tune the hyper-parameters, and report the results of five-fold cross validation.

The confusion matrix in Figure 9 shows that we can achieve very good classification performance (more than 85% accuracy). The primary sources of confusion are between touching a tag and passing near a tag. This is intuitive since both are fleeting events, and hence it is possible to confuse one for the other.

5.2 Case Study: Smoking Intervention

We now look at the use of WearID in the context of JITAI for smoking cessation. We ask a user to wear WearID and perform a routine setup for smoking, which includes reaching for a cigarette pack, picking out a cigarette, putting the pack back, and putting the cigarette to the mouth. A passive tag is attached to the top of the cigarette pack. During the experiment, the user takes out the pack with his or her device-free hand and takes out a cigarette. Ground truth is labeled as three distinct events by pressing a button that is time synchronized with WearID's output.

Figure 10 shows the signal from WearID. We can see the dramatic change in the signal during interaction with the cigarette pack. The classifier output reflects the sequence of events

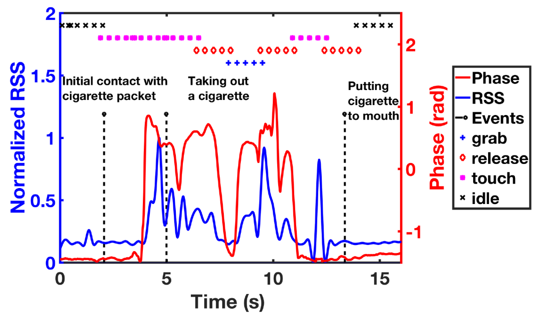

Fig. 10. Time series of interaction with the cigarette pack. The grab event is detected at around 5 seconds, whereas the first puff happens at around 13 seconds, so we have about 8 seconds to trigger a just-in-time intervention.

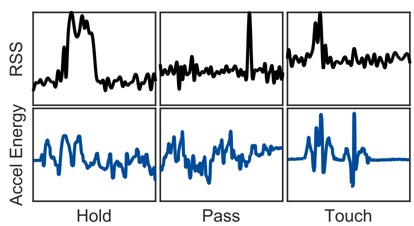

Fig. 11. Figure shows RSS data along with IMU data. (Left) Signals for two consecutive hold gestures, one on a tagged coffee mug and another on a tag-free water bottle. (Middle) Passing a tag; the IMU sees no relevant information. (Right) A tag is touched and then a tag-free surface is touched; whereas the IMU sensor may confuse those cases, WearID clearly distinguishes them.

accurately—the classifier outputs a sequence of touch, grab, and release events. Given the complex nature of interaction, the classifier output sequence is not exactly in that order but switches between the three states. The pattern, however, is clearly visible, and a JITAI can be initiated when a sustained interaction is detected (around 5 to 6 seconds). As the plot shows, WearID can provide up to 7 to 8 seconds of opportunity for initiating the JITAI before the lapse.

The medication pill bottle interaction use case that we described in Section 3 performs similarly since it also involves a sequence of touch, grab, and release events of the pill bottle. Thus, we do not separately report the results for this case study.

We note that although other readers such as the AMS AS3993 can also be used instead of WearID, ours offers the lowest power operation while still providing a sufficient read range to detect interaction with the cigarette pack.

WearID versus IMU for detecting interactions. One question that might arise is whether the interactions that we described in Section 5 are also detectable using inertial signals from a typical smartwatch. Figure 11 shows the RSS from WearID on the top and accelerometer magnitude on the bottom during each of the three gestures. The main observation is that the signal from WearID is a much more localized measure of the interaction, whereas the signal from the accelerometer is more smeared over time. (Of course, the accelerometer signal carries no information relevant to detecting passing near a tag.) This illustrates one of the key benefits of WearID—that it provides a more direct measure of interaction compared to indirect measurements using an inertial sensor.

Combining WearID and IMU signals. Since WearID and the IMU provide different types of signals, they can be used in a complementary manner to improve detection accuracy. For example, prior

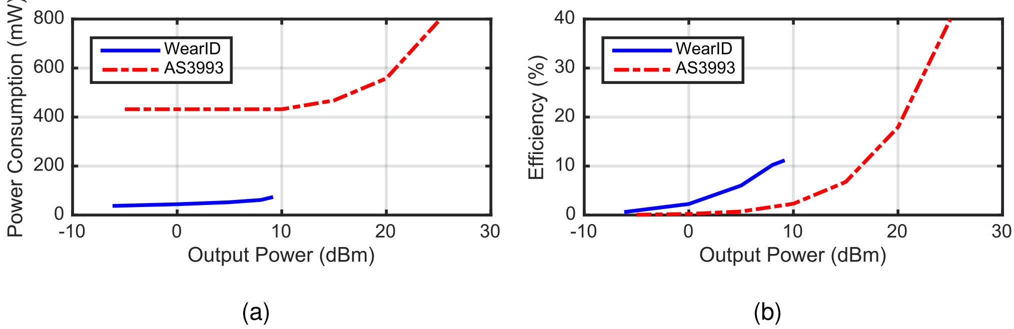

Fig. 12. In these results, we evaluate the power consumed by WearID. (a) The overall radio power consumption as a function of transmit power for WearID and AS3993. (b) The power consumption efficiency as a function of transmit power.

work has also looked at detecting hand-to-mouth gestures corresponding to smoking by using inertial sensors [35]. Fusing RF signals from WearID with motion signals from the IMU can be particularly useful to deal with confounders. For example, hand-to-mouth gestures corresponding to smoking are often confounded by eating and drinking gestures, but the RF signal can provide additional information about interaction with a cigarette pack. In addition, heart rate information from a wrist-worn PPG can be useful to measure signals of craving [10] that can be fused with information about interaction with the cigarette pack or pill bottle from WearID.

6 POWER AND WEARABILITY BENCHMARKING OF WearID

We now turn to low-level benchmarks of various design choices that we made in optimizing WearID. Specifically, we compare WearID against other RFID readers, and provide power, sensitivity, and phase monitoring accuracy benchmarks for the device.

RF power consumption. We now show that our receiver is considerably more power efficient than state-of-the-art reader receivers and considerably reduces power consumption while providing adequate receive sensitivity. Figure 12(a) compares the power consumption of WearID against the state-of-the-art AS3993 commercial RFID reader. We note that WearID is a special-purpose low-power reader, whereas AS3993 is a general-purpose commercial reader, and hence this comparison may be unfair to the commercial reader. However, our goal is to show that there is a substantial performance gap and make the case for a specialized device like WearID to bridge this gap. WearID is currently capable of achieving a maximum of 9 dBm output power, whereas AS3993 can reach 25 dBm, and hence the two plots have different spans.

Our results show that WearID is more than 6× more efficient than AS3993 at equivalent power levels. At the maximum output power of WearID (9 dBm), the system power consumption is 72 mW, of which more than 98% is the carrier generator and less than 2% is the IQ receiver. At the same output power, the power consumption of AS3993 is more than 430 mW; we believe that much of this power is consumed by the RF receiver.

Figure 12(b) compares the efficiency of the two devices, where efficiency is defined as RF output power divided by system power. Our result shows that the commercial reader is optimized to be efficient at higher output power but has very low efficiency at lower output power. This optimization is not surprising since commercial readers are designed to read many distant tags as quickly as possible, as opposed to a small number of nearby tags. In contrast, WearID has higher efficiency at low output power levels that are more appropriate for a battery-powered wearable.

Table 4. Survey of Work on Wearable/Low-Power RFID Readers and Wearable-Based Methods for Detecting Interactions

Function	CPU Time per Tag Read (μs)	Required Energy (μJ)
Pre-processing (down-sampling the received signal)	≈14	≈3
Feature extraction (calculating RSS, phase, envelope)	≈560	≈125
Classification (applying a linear SVM)	≈5 [14]	≈1

Computation power consumption. In this section, we analyze the processing overhead in WearID and associated power consumption on the STM32F103 MCU [6]. There are several processing steps that we take to obtain the user's interaction with tagged objects: down-sampling the baseband signal, calculating RSS, calculating the envelope of RSS, normalization, phase calculation, averaging, and applying a linear SVM. Such computations can be executed on many low-power MCUs today given significant advances in machine learning on resource-constrained devices (e.g., [22]). In addition, specialized low-power accelerators are available on the market for such purposes [50].

Table 4 provides a breakdown of power consumption. Feature extraction is the most time-consuming step in the computational process and dominates the time and energy consumption. Since we have only 24 features, a linear SVM is small in size and consumes around 1 KB of memory [14]. We note that training is performed *a priori* and the model is pre-loaded onto WearID, so we only need to perform prediction on the device.

Both power consumed and latency for computation are relatively small. The power consumed for processing is dictated by the frequency of performing the classification operation. We only trigger classification if a tag is detected, and therefore this only needs to execute intermittently. The latency is also small and the classification executes in about a millisecond, so it is not the bottleneck for real-time interaction detection.

Overall power consumption. Overall, we see that WearID consumes about 100 mW in fully active mode (i.e., when RF and compute are operating continuously). In addition, we note that WearID can perform 80 tag reads per second, which means that we can easily duty cycle WearID since typical interactions occur over a window of a few seconds. With 1% to 10% duty cycling, WearID can be optimized to consume 1 to 10 mW for continuous operation.

These numbers are comparable to power consumption of typical smartwatches and fitness bands. Smartwatches tend to be more power hungry and consume more than 100 mW in almost all active modes and consume roughly 14 mW in sleep mode [31]. Basic fitness trackers and pedometers expend less energy on the display and therefore operate in the 1 to 10 mW range, which is similar to duty-cycled operation on WearID.

Comparing receive sensitivity. WearID consumes less power than a commercial reader, but how much receive sensitivity does it sacrifice to obtain this advantage? To evaluate this, we performed a controlled experiment where we placed a tag at different distances in front of WearID and the commercial reader, and measured the SNR of the IQ output in each case.

Figure 13 shows the results. We observe that WearID is comparable or even superior to the commercial reader at distances less than 15 cm and has about 10 dB lower SNR toward the end of the desired range (i.e., at around 30 cm). At even longer distances, the difference becomes more apparent and is between 8 and 20 dB. Thus, although we lose sensitivity overall when choosing a passive receiver, the difference is not significant in the desired working range.

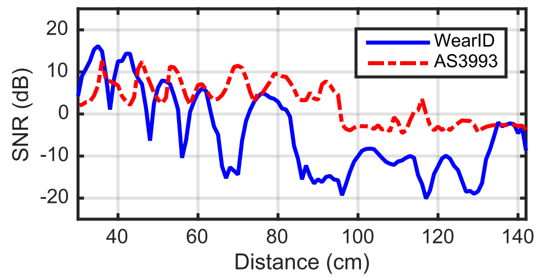

Fig. 13. SNR of the received baseband signal for different distances between the tag and reader.

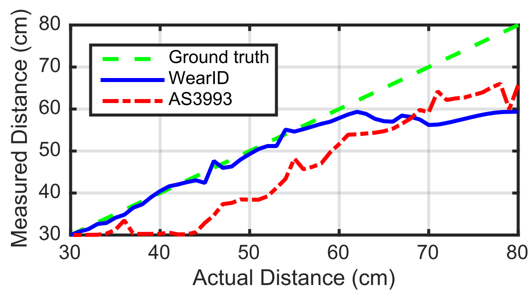

Fig. 14. Phase accuracy for WearID and AS3993. WearID performs very well at reasonably short distances where interactions occur.

Table 5. WearID versus Other Readers

	WearID	Braidio	AS3993
Power consumption	~67 mW	~129 mW	>640 mW
Single antenna operation	Yes	No	Yes
Support for IQ	Yes	No	Yes
Sensitivity	Medium	Low	High
Hardware cost	Low	Low	High

WearID is considerably lower power than a state-of-the-art low-power commercial reader while supporting much of the functionality of the reader. WearID has equivalent power consumption as other research designs while having a much smaller form factor and greater functionality.

Comparing accuracy of phase estimate. Since we use a passive IQ detector implemented using a delay line, our phase output may not be as accurate as an active IQ detector. To determine if this is the case, we perform an experiment where we fix the location of WearID (and AS3993) and move a tag away from them. At each distance, we record the phase measured by WearID and the AS3993 reader, as well as ground truth.

Figure 14 shows the results. At distances less than 60 cm, we find that the phase computed by WearID closely matches the measured ground truth, whereas AS3993 has a fixed offset error until around 40 cm. At distances greater than 80 cm, we find that the phase detected by WearID saturates when the detected energy drops below receive sensitivity. In conclusion, we show that WearID achieves good phase performance over desired operating distances.

Differences from other readers. WearID's receiver design differs from existing reader designs in several ways (Table 5). We focus on two readers in particular: a state-of-the-art commercial reader IC, AS3993, and a low-power reader that was proposed as part of an active-passive radio called

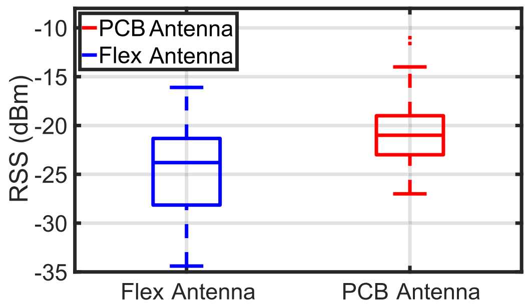

Fig. 15. We plot RSS for the flexible antenna versus the rigid PCB antenna on the PCB board on top of the wrist. The rigid PCB antenna has slightly better and more stable antenna performance due to the reduced impact of variable body-induced detuning.

Braidio [18]. AS3993 is the lowest-power commercially available reader IC and consumes about 0.6 W, whereas other readers consume more than a watt [18]. However, even these numbers are far too high for a portable wrist-worn reader. As discussed in Section 2.1, Braidio operates using two antennas and does not provide IQ output that is crucial for RFID-based sensing. The Braidio receiver is also more power hungry since it uses an instrumentation amplifier to boost the output signal from an envelope detector. In all, WearID presents a different design point compared to existing readers.

Wrist-mounted antenna challenges. We now turn to the practical considerations that emerge when integrating our reader into a smartwatch form-factor wearable. One practical issue that we needed to tackle is that a wristwatch is a device that needs to be comfortable for daily wear. Users prefer to wear wristwatches with varying degrees of looseness. Depending on how tightly the watch is worn, the human body will attenuate the transmitted or received signal strength by varying amounts because of its RF absorption and capacitive properties. Prior studies on signal attenuation for mobile phones have shown that antenna impedance can change significantly when the device is held in the hand [12]. The antenna detunes in the presence of the human body, which changes its impedance and degrades performance.

To facilitate an understanding of the wireless channel near the wrist while conforming to the form factor of a typical smartwatch, we explored two design options for WearID: (1) integrating a flexible dipole antenna directly into the wristband and (2) placing a relatively smaller rigid PCB antenna inside WearID's enclosure (depicted in Figure 7). Rigid PCB antennas can be carefully tuned based on their rigid placement and separation from the human body in the device enclosure; larger, flexible antennas have higher theoretical gain but can be more difficult to tune because detuning effects of the human body vary dynamically. Since the antenna used in WearID needs to provide power in addition to communication, we use the largest antenna possible within our form-factor constraints. We now look at the design tradeoffs between these antenna types more carefully.

To determine the impact of antenna type and positioning, we first determined the best placement for the flexible antenna by placing it at different locations on the wrist and then compared this against the PCB antenna placed on the rigid PCB board. We plot the RSS of several different interactions where a user grabs a tag, touches a tag, and passes near a tag. We repeat these interactions more than 25 times per antenna position per interaction, which leads to a total of 300 measurements. We show all RSS in a single box plot for each configuration in Figure 15.

We see that the PCB antenna has both higher signal strength and less signal variation compared to the flexible antenna–based configuration. This is because the tuning with the human body constantly varies for the flexible antenna due to small movements of the wristwatch. Depending on how tightly the watch is worn, the human body will attenuate the transmitted or received signal strength by varying amounts because of its RF absorption and capacitive properties [12]. The antenna detunes in the presence of the human body, which changes its impedance and degrades performance. Thus, we found that the PCB antenna seems to be a better choice than the flexible antenna for WearID.

7 CONCLUSION

In this article, we present the design of a wearable backscatter reader that we refer to as WearID. Our work tackles a gap in available wearable technologies—although many wearables are designed to measure body signals, we lack a wearable RFID reader to sense RFID tags on objects. We present the design and implementation of a full hardware prototype of such a wearable RF reader that tackles several challenges including reader power consumption and body blockage. Our experimental results show that our design is almost an order of magnitude more power efficient than best-in-class RFID readers on the market while offering ranges of roughly 40 cm around the wrist for power delivery and communication. By leveraging our prototype, we show that we can reliably classify a variety of interactions, including touching a tag, grabbing a tagged object, releasing a tagged object, and passing near a tagged object. These detectors can be leveraged for enabling interactions with tagged objects, for IoT applications, and for mobile health applications. There are many directions that we are continuing to explore, including fusion of RF and IMU signals for more holistic tracking of interactions, improving WearID range to enable a broader range of applications such as augmented reality, and using WearID in new applications including stroke rehabilitation.

REFERENCES

[1] Avery Dennison Corp. 2019. UHF RFID Inlay: AD-237r6 and AD-237r6-P. Retrieved February 26, 2020 from https://rfid.averydennison.com/en/home/innovation/rfid-inlay-designs/AD-237r6-AD-237r6-P.html.

[2] Impinj. 2020. Impinj Indy R2000 UHF Gen 2 RFID Reader Chip. Retrieved February 26, 2020 from https://support.impinj.com/hc/en-us/articles/202755828-Indy-R2000-Datasheet.

[3] NXP. 2019. NXP UCODE 8 Tag Chip Datasheet. Retrieved February 26, 2020 from https://www.nxp.com/docs/en/data-sheet/SL3S1205-15-DS.pdf.

[4] Proteus Digital Health. 2020. Smart Pills for Medication Adherence. Retrieved February 26, 2020 http://www.proteus.com/.

[5] Chris Diorio. (n.d.) Engineering RAIN RFID Solutions. Retrieved February 26, 2020 from http://www.rainrfid.org/wp-content/uploads/2015/07/Diorio-RAIN-solutions-presentation-2015-06-24.pdf.

[6] ST. 2016. STM32L151CBT6 Ultra-Low-Power ARM MCU Datasheet. Retrieved February 26, 2020 from http://www.st.com/st-web-ui/static/active/en/resource/technical/document/datasheet/CD00277537.pdf.

[7] (n.d.) WearID. *Anonymous for review.*

[8] Heba Abdelnasser, Moustafa Youssef, and Khaled A. Harras. 2015. WiGest: A ubiquitous WiFi-based gesture recognition system. arXiv:1501.04301.

[9] Thisum Buddhika, Haimo Zhang, Chamod Weerasinghe, Suranga Nanayakkara, and Roger Zimmermann. 2019. OSense: Object-activity identification based on gasping posture and motion. In *Proceedings of the 10th Augmented Human International Conference 2019 (AH'19)*. ACM, New York, NY, Article 13, 5 pages. DOI: https://doi.org/10.1145/3311823.3311841

[10] Soujanya Chatterjee, Karen Hovsepian, Hillol Sarker, Nazir Saleheen, Mustafa al'Absi, Gowtham Atluri, Emre Ertin, et al. 2016. mCrave: Continuous estimation of craving during smoking cessation. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, New York, NY, 863–874.

[11] Gabe Cohn, Sidhant Gupta, Tien-Jui Lee, Dan Morris, Joshua R. Smith, Matthew S. Reynolds, Desney S. Tan, and Shwetak N. Patel. 2012. An ultra-low-power human body motion sensor using static electric field sensing. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, New York, NY, 99–102.

[12] Ettore Lorenzo Firrao, Anne-Johan Annema, and Bram Nauta. 2004. Antenna behaviour in the presence of human body. In *Proceedings of the 15th ProRisc Workshop on Circuits, Systems, and Signal Processing (ProRisc'04).*

[13] Kenneth P. Fishkin, Matthai Philipose, and Adam Rea. 2005. Hands-on RFID: Wireless wearables for detecting use of objects. In *Proceedings of the 9th IEEE International Symposium on Wearable Computers (ISWC'05)*. IEEE, Los Alamitos, CA, 38–43. DOI: https://doi.org/10.1109/ISWC.2005.25

[14] Tomas Fredriksson and Rickard Svensson. 2018. *Analysis of Machine Learning for Human Motion Pattern Recognition on Embedded Devices*. Master's Thesis. KTH.

[15] Jeremy Gummeson, James Mccann, Chouchang (Jack) Yang, Damith Ranasinghe, Scott Hudson, and Alanson Sample. 2017. RFID light bulb: Enabling ubiquitous deployment of interactive RFID systems. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 2 (2017), Article 12.

[16] Jeremy Gummeson, Bodhi Priyantha, and Jie Liu. 2014. An energy harvesting wearable ring platform for gesture input on surfaces. In *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, New York, NY, 162–175.

[17] Chris Harrison, Desney Tan, and Dan Morris. 2010. Skinput: Appropriating the body as an input surface. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 453–462.

[18] Pan Hu, Pengyu Zhang, Mohammad Rostami, and Deepak Ganesan. 2016. Braidio: An integrated active-passive radio for mobile devices with asymmetric energy budgets. In *Proceedings of the 2016 ACM SIGCOMM Conference*. ACM, New York, NY, 384–397.

[19] Mayank Jain, Jung Il Choi, Taemin Kim, Dinesh Bharadia, Siddharth Seth, Kannan Srinivasan, Philip Levis, Sachin Katti, and Prasun Sinha. 2011. Practical, real-time, full duplex wireless. In *Proceedings of the 17th Annual International Conference on Mobile Computing and Networking (MobiCom'11)*. ACM, New York, NY, 301–312.

[20] Haik Kalantarian, Nabil Alshurafa, and Majid Sarrafzadeh. 2016. Detection of gestures associated with medication adherence using smartwatch-based inertial sensors. *IEEE Sensors Journal* 16, 4 (2016), 1054–1061.

[21] Wan-Kyu Kim, Moon-Que Lee, Jin-Hyun Kim, Hyung-Sun Lim, Jong-Won Yu, Byung-Jun Jang, and Jun-Seok Park. 2006. A passive circulator with high isolation using a directional coupler for RFID. In *2006 IEEE MTT-S International Microwave Symposium Digest*. IEEE, Los Alamitos, CA, 1177–1180.

[22] Liangzhen Lai, Naveen Suda, and Vikas Chandra. 2018. CMSIS-NN: Efficient neural network kernels for arm Cortex-M CPUs. arXiv:1801.06601.

[23] Gierad Laput, Robert Xiao, and Chris Harrison. 2016. ViBand: High-fidelity bio-acoustic sensing using commodity smartwatch accelerometers. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, New York, NY, 321–333.

[24] Gierad Laput, Chouchang Yang, Robert Xiao, Alanson Sample, and Chris Harrison. 2015. EM-Sense: Touch recognition of uninstrumented, electrical and electromechanical objects. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, 157–166.

[25] Bruno Lepri, Nadia Mana, Alessandro Cappelletti, Fabio Pianesi, and Massimo Zancanaro. 2010. What is happening now? Detection of activities of daily living from simple visual features. *Personal and Ubiquitous Computing* 14, 8 (Dec. 2010), 749–766. DOI: https://doi.org/10.1007/s00779-010-0290-z

[26] Hanchuan Li, Eric Brockmeyer, Elizabeth J. Carter, Josh Fromm, Scott E. Hudson, Shwetak N. Patel, and Alanson Sample. 2016. PaperID: A technique for drawing functional battery-free wireless interfaces on paper. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 5885–5896.

[27] Hanchuan Li, Can Ye, and Alanson P. Sample. 2015. IDSense: A human object interaction detection system based on passive UHF RFID. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, 2555–2564.

[28] Hanchuan Li, Can Ye, and Alanson P. Sample. 2015. IDSense: A human object interaction detection system based on passive UHF RFID. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI'15)*. ACM, New York, NY, 2555–2564. DOI: https://doi.org/10.1145/2702123.2702178

[29] Hanchuan Li, Peijin Zhang, Samer Al Moubayed, Shwetak N. Patel, and Alanson P. Sample. 2016. ID-Match: A hybrid computer vision and RFID system for recognizing individuals in groups. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 4933–4944.

[30] Vincent Liu, Aaron Parks, Vamsi Talla, Shyamnath Gollakota, David Wetherall, and Joshua R. Smith. 2013. Ambient backscatter: Wireless communication out of thin air. In *Proceedings of the 2013 ACM SIGCOMM Conference (SIGCOMM'13)*, Vol. 43. ACM, New York, NY, 39–50.

[31] Xing Liu, Tianyu Chen, Feng Qian, Zhixiu Guo, Felix Xiaozhu Lin, Xiaofeng Wang, and Kai Chen. 2017. Characterizing smartwatch usage in the wild. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, New York, NY, 385–398.

[32] Leire Muguira, Juan Ignacio Vázquez, Asier Arruti, Jonathan Ruiz de Garibay, Izaskun Mendia, and Silvia Rentería. 2009. RFIDGlove: A wearable RFID reader. In *Proceedings of the 2009 IEEE International Conference on e-Business Engineering (ICEBE'09)*. IEEE, Los Alamitos, CA, 475–480. http://dblp.uni-trier.de/db/conf/icebe/icebe2009.html#MuguiraVAGMR09.

[33] Pavel V. Nikitin, Shashi Ramamurthy, and Rene Martinez. 2013. Simple low cost UHF RFID reader. In *Proceedings of the 2013 IEEE International Conference on RFID*. 126–127.
[34] Shahriar Nirjon, Jeremy Gummeson, Dan Gelb, and Kyu-Han Kim. 2015. TypingRing: A wearable ring platform for text input. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, New York, NY, 227–239.
[35] Abhinav Parate, Meng-Chieh Chiu, Chaniel Chadowitz, Deepak Ganesan, and Evangelos Kalogerakis. 2014. Risq: Recognizing smoking gestures with inertial sensors on a wristband. In *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, New York, NY, 149–161.
[36] Ivan Poupyrev, Nan-Wei Gong, Shiho Fukuhara, Mustafa Emre Karagozler, Carsten Schwesig, and Karen E. Robinson. 2016. Project Jacquard: Interactive digital textiles at scale. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 4216–4227.
[37] Qifan Pu, Sidhant Gupta, Shyamnath Gollakota, and Shwetak Patel. 2013. Whole-home gesture recognition using wireless signals. In *Proceedings of the 19th Annual International Conference on Mobile Computing and Networking (MobiCom'13)*. ACM, New York, NY, 27–38. DOI: https://doi.org/10.1145/2500423.2500436
[38] Tauhidur Rahman, Alexander Travis Adams, Mi Zhang, Erin Cherry, Bobby Zhou, Huaishu Peng, and Tanzeem Choudhury. 2014. BodyBeat: A mobile system for sensing non-speech body sounds. In *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys'14)*, Vol. 14. 2–13.
[39] Nazir Saleheen, Amin Ahsan Ali, Syed Monowar Hossain, Hillol Sarker, Soujanya Chatterjee, Benjamin Marlin, Emre Ertin, Mustafa al'Absi, and Santosh Kumar. 2015. puffMarker: A multi-sensor approach for pinpointing the timing of first lapse in smoking cessation. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp'15)*. ACM, New York, NY, 999–1010. DOI: https://doi.org/10.1145/2750858.2806897
[40] Munehiko Sato, Ivan Poupyrev, and Chris Harrison. 2012. Touché: Enhancing touch interaction on humans, screens, liquids, and everyday objects. In *Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems*. ACM, New York, NY, 483–492.
[41] Sheng Shen, He Wang, and Romit Roy Choudhury. 2016. I am a smartwatch and I can track my user's arm. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, New York, NY, 85–96.
[42] Joshua R. Smith, Kenneth P. Fishkin, Bing Jiang, Alexander Mamishev, Matthai Philipose, Adam D. Rea, Sumit Roy, and Kishore Sundara-Rajan. 2005. RFID-based techniques for human-activity detection. *Communications of the ACM* 48, 9 (Sept. 2005), 39–44. DOI: https://doi.org/10.1145/1081992.1082018
[43] Michael C. Sokol, Kimberly A. McGuigan, Robert R. Verbrugge, and Robert S. Epstein. 2005. Impact of medication adherence on hospitalization risk and healthcare cost. *Medical Care* 43, 6 (2005), 521–530.
[44] Yan Song and Yang Lin. 2015. Combining RGB and depth features for action recognition based on sparse representation. In *Proceedings of the 7th International Conference on Internet Multimedia Computing and Service (ICIMCS'15)*. ACM, New York, NY, Article 49, 5 pages. DOI: https://doi.org/10.1145/2808492.2808541
[45] Andrew Spielberg, Alanson Sample, Scott E. Hudson, Jennifer Mankoff, and James McCann. 2016. RapID: A framework for fabricating low-latency interactive objects with RFID tags. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 5897–5908.
[46] Li Sun, Souvik Sen, Dimitrios Koutsonikolas, and Kyu-Han Kim. 2015. WiDraw: Enabling hands-free drawing in the air on commodity WiFi devices. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking (MobiCom'15)*. ACM, New York, NY, 77–89. DOI: https://doi.org/10.1145/2789168.2790129
[47] Vamsi Talla, Mehrdad Hessar, Bryce Kellogg, Ali Najafi, Joshua R. Smith, and Shyamnath Gollakota. 2017. LoRa backscatter: Enabling the vision of ubiquitous connectivity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 105.
[48] Sheng Tan, Linghan Zhang, Zi Wang, and Jie Yang. 2019. MultiTrack: Multi-user tracking and activity recognition using commodity WiFi. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI'19)*. ACM, New York, NY, Article 536, 12 pages. DOI: https://doi.org/10.1145/3290605.3300766
[49] ClinicalTrials.gov. (n.d.). Sense2Stop: Mobile Sensor Data to Knowledge. Retrieved February 26, 2020 from https://clinicaltrials.gov/ct2/show/NCT03184389.
[50] GreenWaves Technologies. (n.d.). GAP8. Ultra-Low Power, Always-On Processor for Embedded Artificial Intelligence. Retrieved February 26, 2020 from https://greenwaves-technologies.com/ai_processor_GAP8/.
[51] MD2K. (n.d.). MD2K: NIH Center of Excellence on Mobile Sensor Data-to-Knowledge. Retrieved February 26, 2020 from https://md2k.org/.
[52] Deepak Vasisht, Jue Wang, and Dina Katabi. 2014. RF-IDraw: Virtual touch screen in the air using RF signals. In *Proceedings of the 6th Annual Workshop on Wireless of the Students, by the Students, for the Students (S3'14)*. ACM, New York, NY, 1–4. DOI: https://doi.org/10.1145/2645884.2645889

[53] Tam Vu, Akash Baid, Simon Gao, Marco Gruteser, Richard Howard, Janne Lindqvist, Predrag Spasojevic, and Jeffrey Walling. 2012. Distinguishing users with capacitive touch communication. In *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking*. ACM, New York, NY, 197–208.

[54] Jue Wang, Fadel Adib, Ross Knepper, Dina Katabi, and Daniela Rus. 2013. RF-compass: Robot object manipulation using RFIDs. In *Proceedings of the 19th Annual International Conference on Mobile Computing and Networking (MobiCom'13)*. ACM, New York, NY, 3–14. DOI: https://doi.org/10.1145/2500423.2500451

[55] Wei Wang, Alex X. Liu, Muhammad Shahzad, Kang Ling, and Sanglu Lu. 2015. Understanding and modeling of WiFi signal based human activity recognition. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking (MobiCom'15)*. ACM, New York, NY, 65–76. DOI: https://doi.org/10.1145/2789168.2790093

[56] Chouchang Yang, Jeremy Gummeson, and Alanson Sample. 2017. Riding the airways: Ultra-wideband ambient backscatter via commercial broadcast systems. In *Proceedings of the IEEE Conference on Computer Communications (INFOCOM'17)*. IEEE, Los Alamitos, CA, 1–9.

[57] Lei Yang, Qiongzheng Lin, Xiangyang Li, Tianci Liu, and Yunhao Liu. 2015. See through walls with COTS RFID system! In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking (MobiCom'15)*. ACM, New York, NY, 487–499. DOI: https://doi.org/10.1145/2789168.2790100

[58] L. Ye, H. Liao, F. Song, J. Chen, C. Li, J. Zhao, R. Liu, et al. 2010. A single-chip CMOS UHF RFID reader transceiver for Chinese mobile applications. *IEEE Journal of Solid-State Circuits* 45, 7 (July 2010), 1316–1329. DOI: https://doi.org/10.1109/JSSC.2010.2049459

[59] Daniel J. Yeager, Alanson P. Sample, Joshua R. Smith, and Joshua R. Smith. 2008. Wisp: A passively powered UHF RFID tag with sensing and computation. In *RFID Handbook: Applications, Technology, Security, and Privacy*, S. A. Ahson and M. Ilyas (Eds.). CRC Press, Boca Raton, FL, 261–278.

[60] Yang Zhang, Junhan Zhou, Gierad Laput, and Chris Harrison. 2016. SkinTrack: Using the body as an electrical waveguide for continuous finger tracking on the skin. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 1491–1503.

[61] Junhan Zhou, Yang Zhang, Gierad Laput, and Chris Harrison. 2016. AuraSense: Enabling expressive around-smartwatch interactions with electric field sensing. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, New York, NY, 81–86.

[62] Yongpan Zou, Jiang Xiao, Jinsong Han, Kaishun Wu, Yun Li, and Lionel M. Ni. 2017. GRfid: A device-free RFID-based gesture recognition system. *IEEE Transactions on Mobile Computing* 16, 2 (2017), 381–393. DOI: https://doi.org/doi.ieeecomputersociety.org/10.1109/TMC.2016.2549518

Received March 2019; revised September 2019; accepted November 2019