

Enabling Instantaneous Feedback with Full-duplex Backscatter

Vincent Liu[†], Vamsi Talla[†], Shyamnath Gollakota
University of Washington
{liuv, vamsit, gshyam}@uw.edu
[†]Co-primary Student Authors

Abstract— This paper introduces the first design that enables full-duplex communication on battery-free backscatter devices. Specifically, it gives receivers a way to provide low-rate feedback to the transmitter on the same frequency as that of the backscatter transmissions, using neither multiple antennas nor power-consuming cancellation hardware. Our design achieves this goal using only fully-passive analog components that consume near-zero power. We integrate our design with the backscatter network stack and demonstrate that it can minimize energy wastes that occur due to collisions and also correct for errors and changes in channel conditions at a granularity smaller than that of a packet. To show the feasibility of our design, we build a hardware prototype using off-the-shelf analog components. Our evaluation shows that our design cancels the self-interference down to the noise floor, while consuming only 0.25 μW and 0.54 μW of transmit and receive power, respectively.

CATEGORIES AND SUBJECT DESCRIPTORS

C.2.1 [Network Architecture and Design]: Wireless communication

KEYWORDS

Backscatter; Full-duplex wireless; Wireless

1. INTRODUCTION

Backscatter communication is a promising technique for low-power computing devices. RFID, for instance, has enabled a multitude of useful applications [27, 34] that take advantage of the fact that one side of the link can be completely battery-free. More recently, researchers have shown the feasibility of using backscatter to communicate directly between two low-power devices [18]. The power draw of this technique is such that both devices can operate solely off of harvested energy. By creating a network of battery-free devices that can communicate with each other, this can enable the types of applications envisioned by the Internet-of-Things (e.g., smart homes).

Unfortunately, existing link- and network-layer protocols are ill suited for such communication networks. Extreme power constraints not only magnify the effects of collisions and failed transmissions, they also make it difficult for these devices to use multi-

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

MobiCom'14, September 7-11, 2014, Maui, Hawaii, USA.

Copyright 2014 ACM 978-1-4503-2783-1/14/09 ...\$15.00.

<http://dx.doi.org/10.1145/2639108.2639136>.

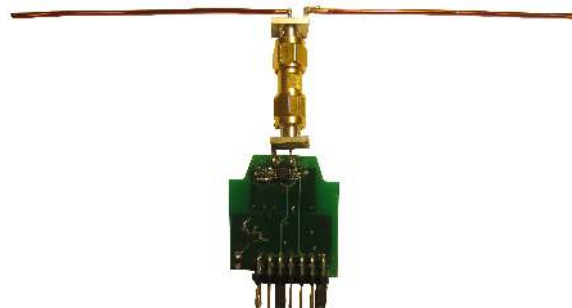


Figure 1: **Full-duplex Backscatter Prototype:** A photo of our Full-duplex Backscatter prototype. This side contains the transmitter switch, while the receiver is implemented on the other side.

round protocols like RTS-CTS—storing increasing amounts of energy becomes exponentially hard. In this paper, we explore the feasibility of an instantaneous feedback channel where a receiver can simultaneously send feedback information while still receiving the original transmission. Such a feedback channel could address many of the problems facing these devices:

- *Packet Collisions:* Instantaneous feedback could be used to terminate transmission as soon as a collision is detected. Because recharge time dominates transmission time [28], enabling early termination increases throughput by orders of magnitude.
- *Rate Adaptation:* Instead of waiting for packet drops to adjust rate, a feedback channel allows us to gather bit error rate statistics and use them to change rate at the level of bits rather than packets. This is particularly important for backscatter devices, as their transmission rates are orders of magnitude lower than traditional radio communication (e.g., Wi-Fi). Thus, they take significantly longer to transmit the same amount of data—long enough that the channel may change within the span of a single packet.
- *Retransmissions:* In the same way, we can use a feedback channel to inform transmitters of bit-level errors. Thus, the transmitter only needs to retransmit the subset of bits that were incorrect, rather than throwing away the packet entirely and trying again.

Traditional radio communication designs have explored the idea of adding a feedback channel, but none of those existing approaches are applicable to backscatter systems. They mainly fall into two categories: frequency-division duplexing and full-duplex communication. Frequency division is not practical because battery-free backscattering tags use simple analog envelope detectors for reception and switches for transmissions. They therefore cannot be designed to be frequency selective while maintaining low power consumption [28].¹ Similarly, recent work on full-duplex communication [9, 5] does not apply because these devices do not have

¹RFID tags can not decode in the presence of concurrent transmissions in non-overlapping frequency bands.

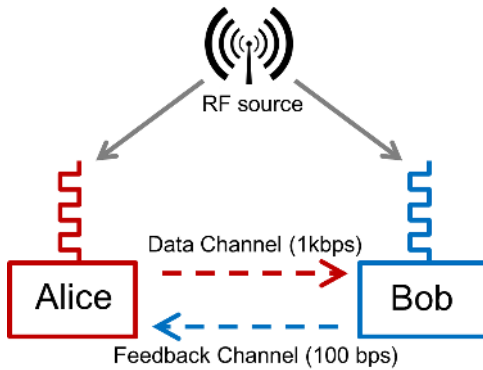


Figure 2: **Full-duplex Backscatter:** Two battery-free devices, Alice and Bob, communicate by backscattering signals from the RF source. Alice transmits to Bob at a high rate on the data channel and receives instantaneous feedback from Bob on the feedback channel.

the computational or energy resources to implement complex interference cancellation techniques or include space-consuming arrays of antennas. In fact, existing full-duplex designs are prototyped on software radios, each of which consumes more than 1-2 Watts of power [11, 2]. More importantly, these designs use components such as ADCs and oscillators that are power-consuming and hence are not applicable for backscattering devices [6].

We introduce a novel technique called Full-duplex Backscatter with which backscatter devices can obtain instantaneous feedback on the same frequency as that of the transmission, without multiple antennas or power-consuming cancellation hardware. Our technique uses fully-passive analog components in order to enable devices to simultaneously maintain a data channel and a low-rate feedback channel in the opposite direction. Because everything is done using simple analog components such as diodes and resistors, we incur near-zero power cost. Finally, we go on to show that the resulting feedback channel can be used to address each of the above problems in the network stack design.

To understand Full-duplex Backscatter, consider the ambient backscatter [18] setup in Fig. 2 with two battery-free devices, Alice and Bob. Alice communicates with Bob by backscattering the transmissions from the RF source; she does so by reflecting or absorbing the incident signals to convey a ‘0’ bit and ‘1’ bit, respectively. Bob receives these bits by tracking the changes in the average amplitude of its received signal.

The challenge in creating a simultaneous feedback channel from Bob to Alice is that the act of backscattering at Bob creates large changes in Bob’s received signal amplitude that greatly degrades his decoding capabilities. Conceptually, since Bob uses the same antenna to receive and transmit, the amplitude of his received signal will change significantly based on whether he is reflecting (sending a ‘0’) or absorbing (sending a ‘1’) on the feedback channel.

Our intuition is as follows: *Since Bob’s receiver decodes by tracking the amplitude of the received signal, if Bob can backscatter information without changing the amplitude, he will not create any self-interference.*² Our approach is therefore to reflect and absorb a *fixed* amount of signal, such that the amplitude of the received signal is constant. To achieve this, Full-duplex Backscatter changes the impedance of the antenna to create phase shifts to the received signal while maintaining the same amplitude. Two reflections with different phases will interfere with the ambient signal at Alice to create two different amplitude levels that Alice can decode using a standard, amplitude-tracking backscatter receiver. We note,

²Decoding both amplitude and phase information requires power-consuming oscillators [13] that are avoided in backscatter devices. Thus, phase is a free-parameter that is available for our purposes.

however, that while in theory one can pick the impedance values to create equal-amplitude signals with different phase shifts, practical circuits have small impedance mismatches that result in residual interference. In §4, we describe how Full-duplex Backscatter leverages the low rate of the feedback channel to eliminate the effect of this residual interference.

To demonstrate the feasibility of our techniques, we built the prototype in Fig. 1 using off-the-shelf components. Our prototypes communicate with each other by backscattering continuous wave transmissions from an RF source in the 920 MHz range. We configure the devices to transmit at 100 bps on the feedback channel and 1 kbps on the data channel; the latter is the state-of-the-art for device-to-device backscatter communication [18]. Our evaluation shows that our hardware reduces the self-interference close to the noise floor across the frequency range and for a range of receiver power levels, while consuming only 0.25 μ W and 0.54 μ W of transmit and receive power respectively.

We also integrate the feedback channel provided on our prototype into the backscatter network stack in order to perform collision detection, implement rate adaptation at sub-packet granularities, and reduce retransmissions by performing in-frame error correction. Our results are as follows:

- By terminating colliding transmissions early, we reduce transmitter recharge time by two orders of magnitude for small 64-byte packets across the operational range of receive power levels. The reduction is higher for longer packet sizes.
- By performing in-frame error correction, we reduce the number of retransmitted bits by an order of magnitude, for a range of bit error rates on the data channel.
- Performing in-frame rate adaptation can increase the throughput of a backscatter communication system by about 33% compared to an idealized, SNR-based, packet-level adaptation algorithm. We note that existing backscatter systems do not have access to SNR and thus the benefits of our system are likely greater.

Contributions. In this paper, we:

- Introduce the first design that enables an instantaneous feedback channel on battery-free backscatter devices. We do so leveraging the properties of backscatter communication to perform self-interference cancellation using only passive analog components.
- Develop and integrate the feedback channel with the network stack to address collisions, fine-grained rate adaptation, and error-correction on battery-free devices.
- Design and build an hardware prototype that demonstrates the feasibility of our techniques and protocols in practice.

Network designers have traditionally believed that a full-duplex feedback channel is difficult to achieve on battery-free devices. In this paper, we not only demonstrate the feasibility of such a channel but also show that it can enable link-layer and network-layer mechanisms that can significantly benefit backscatter communication. We believe that this is an important step towards creating a network of backscatter devices that can communicate with each other.

2. THE CASE FOR FEEDBACK IN BACKSCATTER SYSTEMS

A feedback channel has multiple benefits for both traditional backscatter (between a battery-free device and a powered reader) as well as ambient backscatter communication (between two battery-free devices). The key operational challenge for these devices is that they may not have enough energy to transmit multiple pack-

ets back-to-back due to the fact that they largely rely on harvested power—a resource that is both limited and unpredictable.³

Below, we provide three illustrative examples of core issues in backscatter communication, and discuss how an instantaneous feedback channel can help address these issues.

Wireless Collisions. Collisions in power-constrained scenarios are problematic because the power used to transmit the packet is essentially wasted. This problem is aggravated when backscatter communication is used in general scenarios beyond traditional RFID/NFC systems that read/write small bits of information. As the packet sizes increase in general backscatter communication, the amount of power wasted in transmitting an undecodable packet increases proportionally. The problem is even worse in scenarios like ambient backscatter communication [18] where a network of battery-free devices communicate with each other without a central coordinator. In such networks, the presence of hidden terminals can aggravate the collision problem.

Multi-round protocols such as RTS-CTS are not attractive in this domain because the devices may not have enough power to transmit and receive multiple packets in succession.

An instantaneous feedback channel can help alleviate the problem of collisions and hidden terminals. Specifically, the receiver can send a message back to the transmitter informing it of the presence of a collision. Since the feedback arrives at the transmitter while it still is transmitting, it can terminate transmission immediately and avoid power waste.

Rate Adaptation. Rate adaptation enables devices to adapt to changing channel conditions to achieve good throughput. The problem here is that backscatter devices generally have low transmission rates. So low, in fact, that traditional rate adaptation reacts too slowly to be effective—channel quality can change significantly within the span of a single packet. This is important because, whereas maximum transmission time for an 802.11g packet is 542 μ s [16], backscatter devices such as those proposed in [18] can take twice as much time to transmit a single bit. These devices can take up to 2s to transmit a 256 byte packet. Even an environment that changes at a modest rate (e.g., a person walking) can create changes within that time frame. Note that reducing the packet size is not a desirable solution since it increases the overhead of the preambles and the headers significantly.

An instantaneous feedback channel can address this problem by delivering feedback about channel conditions within the duration of a single packet transmission. This is beneficial in scenarios like that of wearable devices [24], which can experience channel changes as the person moves in the environment.

Retransmissions. Finally, we consider the issue of retransmission. While forward error correction, which may include rate adaptation, can effectively decrease the frequency of failures, there are inevitably cases where certain bits are unrecoverable. This traditionally requires retransmission of the entire packet.

The challenge in backscatter communication is that retransmissions are very undesirable—more than in traditional communication. Not only is the penalty of failed/wasted transmissions much higher in these types of devices, so is the probability of errors. As an example, consider the problem presented in the preceding case. In dynamic channels, packets take so long to transmit that it is likely that the channel will change or interference will occur during the course of a single transmission. Even if only a single bit is un-

³One might think that a solution to this problem would be to simply gather enough energy for multiple packets; in §5.1, we explain why storing more energy is not straightforward due to leakage issues.

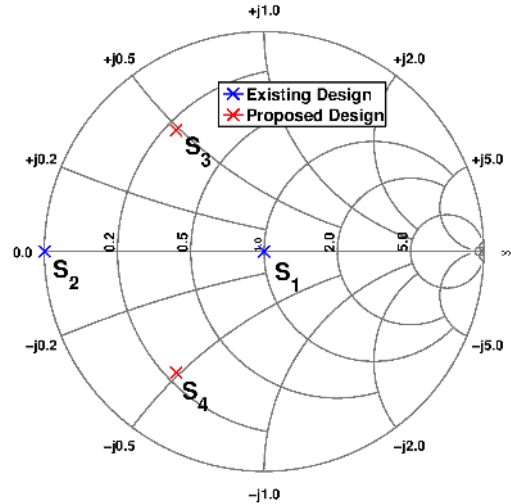


Figure 3: **Smith Chart:** A Smith chart is used to represent impedances and their effect on reflection of incident signals. Each point on the Smith chart corresponds to an input impedance state Z , whose distance to the center of the Smith chart represents the absolute value of the coefficient of reflection, $|\Gamma|$. The center of the Smith chart S_1 corresponds to the matched impedance state, while impedance state S_2 corresponds to the short state. A conventional backscatter transmitter switches between S_1 and S_2 . In our design, transmitters switch between the S_3 and S_4 impedance states instead.

recoverable, retransmitting an entire packet means that the energy spent in the transmission of the original packet was wasted.

An instantaneous feedback channel can be used to address the above issue by allowing the transmitter to only retransmit the set of bits that are incorrectly received. Specifically, the receiver can use the feedback channel to transmit checksums of small subsets of bits [37] during the packet transmission. The transmitter can use these checksums to identify subsets of bits that are incorrectly received. Note that since these checksums are transmitted on the feedback channel at the same time as the packet transmissions, they do not add additional transmission overhead to the system.

Road Map for the Rest of the Paper. We first provide a background of how backscatter communication works. We then describe in detail our design for achieving the instantaneous feedback channel. Finally, we describe how we can integrate the feedback channel with the backscatter network stack to address the above issues.

3. BACKGROUND ON BACKSCATTER

Backscatter communication is markedly different from traditional radio communication. Rather than generating radio signals—a process that requires orders of magnitude more power than is harvestable—many battery-free device designs [23, 18] have adopted the strategy of re-purposing existing signals in order to communicate. These systems leverage existing RF waves, either from a nearby reader or from ambient signal sources [18]. By backscattering RF waves rather than generating their own (as in traditional communication), these devices can decrease the power required to communicate by several orders of magnitude. While the reduction in the power consumption is enough to make charging delays practical, since these systems still work in energy-constrained environments, problems of collisions, rate adaptation, and retransmissions stress the system. In the rest of this section, we provide a general overview of the design of backscatter devices that applies for both RFID as well as ambient backscatter [18] systems.

To illustrate how backscatter works, consider two devices, Alice and Bob, and say Alice wishes to send a packet to Bob. At a

high level, she does so by switching between two states: reflecting and non-reflecting. In the reflecting state, Bob will see a superposition of the original ambient wave as well as Alice’s reflected wave. In the non-reflecting state, Bob will simply see the original wave alone. These two states can be used as two bits, which Bob can distinguish between so as to decode a string of ‘0’s and ‘1’s. We describe in detail how current backscatter systems achieve such transmission and reception capabilities in hardware.

Backscatter Transmitter Designs. The job of a backscattering transmitter is to create two distinct superpositions of waves at the receiver. As mentioned above, it does so by switching between two states, which we call reflecting and non-reflecting. Intuitively, RF signals get reflected when they cross two materials that have different impedance values. Since the impedance of an antenna is different from the air around it, a fraction of the incident RF signals get reflected off the antenna. Backscatter works by creating an *additional* impedance boundary between the antenna and the circuitry of the backscattering device.

To be more formal, let’s consider the Smith chart shown in Fig. 3. Existing backscatter transmitters switch the impedance presented to the antenna between state S_1 (corresponding to an impedance matched to the antenna) and state S_2 (corresponding to an impedance that is shorted). The difference in reflected power between the two states increases with the square of the conjugate match reflection coefficient:

$$\Gamma^* = \frac{Z_a^* - Z_L}{Z_a + Z_L}$$

where Z_a is the resonant antenna impedance, and Z_L is the complex load impedance. In the matched state S_1 , $Z_L = Z_a^*$, $\Gamma^* = 0$, and there is no reflection due to load impedance. In contrast, in the shorted state S_2 , $Z_L = 0$, $\Gamma^* = \frac{Z_a^*}{Z_a}$, and some fraction of the wave is reflected. By switching between these two impedances, a transmitter can change the amplitude of the reflected signals and convey information to the receiver.⁴

Backscatter Receiver Designs. The reflecting and non-reflecting states at the transmitter create two different resulting signals, which the receiver hardware should differentiate to decode the transmitted bits. The challenge is that the ultra-low-power devices that use backscatter communication typically do not have enough power to perform complex demodulation on the signal. In fact, they often do not have enough power to operate conventional radio components such as Analog-to-Digital Converters (ADCs) and oscillators (that generate the carrier frequency).

Current receiver designs address this problem by relying on the fact that the received signal has two different amplitudes that can be extracted in hardware using a simple envelope detector circuit with passive components (i.e., a diode combined with an RC circuit). A naive approach to decode the bits from the output of such an envelope detector is to use a full-range ADC. However to reduce the power consumption, backscattering devices instead use a low-power comparator circuit. Specifically, the comparator takes as inputs the signal envelope and its long-term average. In the absence of a backscattering transmitter, the envelope does not vary and the output of the comparator is steady. In presence of a transmitter, the signal envelope varies with respect to the long-term average and the comparator uses this difference to decode the bits.

⁴As in any backscatter system (including RFID), it is possible that the two impedance states result in little or no difference in the amplitude of the signals at the receiver. However, the probability of such an event is very low since it requires the channel phase to match exactly to cancel the effect of the two impedance states.

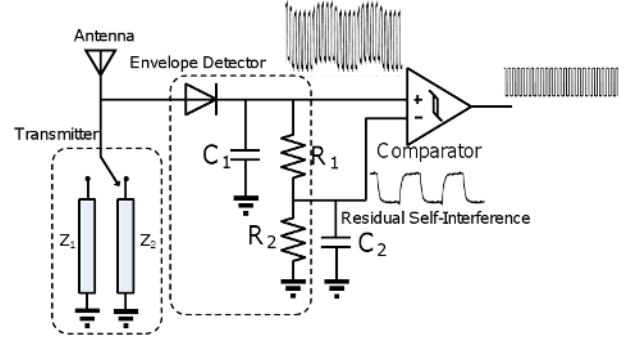


Figure 4: **Bob’s Transmitter and Receiver:** Bob switches between Z_1 and Z_2 at a low rate to transmit data to Alice. Bob’s receiver consists of three main components: an envelope detector to remove the carrier frequency, a low pass filter to isolate the low frequency residual self-interference and a comparator to cancel the residual self-interference from the output of the envelope detector and decode the received bits. In effect by doing so, Bob is implementing a high pass filter by using a low pass filter to track the residual interference and subtracting it from the envelope signal using a low power comparator. The high pass operation cancels the low rate self-interference from the desired high rate signal.

4. FULL-DUPLEX BACKSCATTER DESIGN

We introduce Full-duplex Backscatter, a novel communication technique that creates a low-rate, instantaneous feedback channel for backscatter devices. Our goal is to design a feedback channel that uses a single antenna for transmit and receive and consumes a negligible amount of power.

Creating such a feedback channel is challenging for at least three reasons. First, conventional backscatter transmitters actively change the amplitude of reflected signal to encode data. This is bad for the receiver, which relies on changes in the amplitude of the received signal to decode transmissions. Because the transmitter and receiver share a single antenna, a device’s transmitter will create a great deal of interference for its own receiver. Second, backscatter transmitters and receivers must share the same signal. Conservation of energy means that there is a fundamental tradeoff between the strength of the signal scattered by the transmitter and the strength of the signal received by the receiver. Finally, backscatter devices have severe power, area, and cost constraints that rule out traditional interference-cancellation techniques including multiple antennas and complex adaptive cancellation techniques.

4.1 Overview

Our approach hinges on two ideas. (1) The design of a transmitter that can backscatter information with minimal changes in the amplitude of its reflections. We show that such a design provides significant cancellation in practice. (2) To deal with the residual interference we exploit the rate difference between the data and feedback channels to design custom receiver circuits for the data and feedback channel receivers respectively. At a high level, the receiver of the feedback channel sees high-frequency residual interference, whereas the data channel receiver sees low-frequency residual interference. We use low-pass and high-pass filters designed using passive components to eliminate this interference. We describe these two ideas in more detail below.

4.2 Full-duplex Backscatter Transmitter Design

Rather than modulating the amplitude of its reflections, a Full-duplex Backscatter transmitter reflects/absorbs a fixed amount of signal, such that the amplitude of the received signal is constant. At a high level, the devices transmits information by switching be-

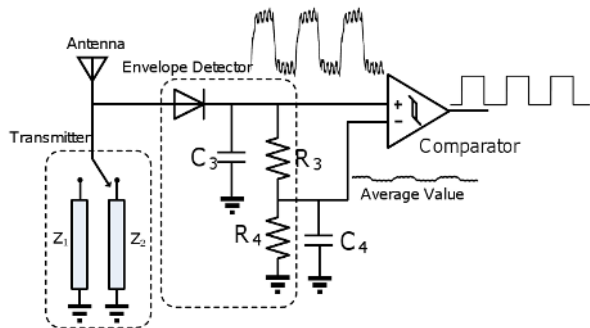


Figure 5: **Alice's Transmitter and Receiver:** Alice switches between Z_1 and Z_2 impedances at a high rate to transmit data to Bob. The receiver on Alice consists of two main components: an envelope detector/low pass filter to remove the carrier frequency and self-interference and another low pass filter to track the average value. These two signals are fed to a comparator to threshold the received signal and decode the digital bits sent by Bob.

tween two impedances that modulate the *phase* of the reflected signal instead of the amplitude.

More formally, we can understand our design using a Smith chart shown in Fig. 3. The transmitter switches between two conjugately-matched impedance states corresponding to S_3 and S_4 . The two impedance states S_3 and S_4 are chosen such that the magnitude of the reflection coefficient $|\Gamma^*|$ and hence the magnitude of the signal reflected in the two states is equal. However, because the Γ in the two states are complex conjugates of each other, the reflected signals are out of phase. The two reflections with different phases from the transmitter will interfere with the ambient signal at the corresponding receiver to create waves with two different amplitude levels. We can thus, in principle, decode the bits on the receiver using a standard, amplitude-tracking backscatter receiver.

Note that using complex impedances to change the phase does not necessarily degrade the quality of transmissions. In particular, the quality of the transmitted signal is a function of the distance between the two impedance states on the Smith chart. Said differently, the greater the distance between the two impedance states, the greater the difference between the two received bit values. Thus, the above approach can achieve equal or greater signal quality than traditional backscatter transmitters.

In theory, the above technique should result in perfect cancellation; however, in practice, impedance mismatches occur due to component variations and other effects. These mismatches result in residual interference. We describe how our receiver addresses such residual interference in the next section.

4.3 Full-duplex Backscatter Receiver Design

A Full-duplex Backscatter receiver has two main components: an envelope detector and a threshold computation stage that either performs a low-pass filter or a high-pass filter operation. The envelope detector is used to remove the RF carrier frequency and extract amplitude information. This operation is similar to that used in traditional backscatter receivers.

The threshold computation stage on the other hand takes the amplitude information from the output of the envelope detector and looks at how the amplitude changes over time. The goal here is to remove the effect of the residual interference. We do so by leveraging the fact that the two devices are transmitting at different rates. As an example, consider Bob from Fig. 2 whose received signal is illustrated in Fig. 4.

We can use a low-pass filter consisting of resistors R_1, R_2 , capacitors C_1 and C_2 and a comparator to track the slowly-varying self-interference and compute a time-varying threshold as shown in the figure. We feed the received signal to the positive terminal and the computed slowly varying threshold to the negative terminal of the comparator. This is essentially performing a high pass filter operation that subtracts the low rate interference from the received signal and enables Bob to efficiently decode Alice's transmissions.

The operations at Alice, as shown in Fig. 5, are analogous but different. Specifically, we implement an envelope detector using different C_3, C_4, R_3 , and R_4 that suppresses both the RF carrier and the high-frequency self-interference. We then use a low pass filter consisting of R_4 and C_4 to compute the long-term average of the envelope and subtract this average from the output of the envelope detector to decode the low-rate feedback transmissions from Bob.

We note that in principle, it is possible to use a conventional amplitude modulating transmitter in combination with the presented residual interference cancellation techniques to enable a simultaneous feedback channel. However, in practice, use of such conventional transmitters results in a residual interference of about 100 dB (relative to the noise floor). In order to eliminate such residual interference, we need to implement high-order filters. Given the losses associated with passive components such as capacitors and resistors and loading effects of cascaded passive filters, a practical implementation of such filters using only passive components is not a feasible in practice. Hence, we believe that our proposed phase modulating transmitter architecture followed by the residual cancellation technique is essential to the design of an instantaneous backscatter feedback system.

4.4 Putting Things Together

For simplicity, we have thus far described the Alice and Bob's designs separately. In reality, every device must take on multiple roles depending on whether it has a packet to send, wants to receive a packet, or is simply stay idle. We must therefore answer two questions. First, how do we consolidate the two design illustrated in Fig. 5 and Fig. 4? And second, how do we engineer the system to maximize receive and harvesting efficiency?

A simple way to consolidate the above designs is to include hardware for both the forward and feedback channels between which the device can opportunistically switch based on it's operating state. In particular, devices need to have four impedance values—two for the forward channel and two for the feedback channel. Devices also need two receive chains for each target transmission rate in order to implement the low-pass and high-pass filters. We note that in an ASIC implementation, tunable resistors and capacitors can be used to eliminate hardware redundancy between the different states.

For both reception and harvesting, efficiency is highest when the matched impedance is close to S_1 in Fig. 3 (where it reflects the least amount of energy back on the antenna). Since the transmitter design on the feedback channel picks impedance values S_3 and S_4 that are necessarily distinct from S_1 , the receiver efficiency of the data channel is lower than existing backscatter receivers. To address this issue, we jointly optimize the impedance values at the transmitters of both the data channel (Alice) and feedback channel (Bob) to ensure that the efficiency (receive BER) of the data channel remains the same. This translates into impedances values that are closer to the origin for Bob's feedback transmissions and those that are further from the origin for Alice's data transmissions. In our prototype implementation this results in a feedback rate that is one-tenth of the data rate.

Our design has an additional impedance value for the case when there are no ongoing transmissions. During this state, the device

does not need to transmit any signals and therefore can fully optimize for receive/harvesting efficiency by staying in the S_1 state. This is important because during transmission, the forward channel has a harvesting efficiency of 44.83% compared to 50% in a traditional backscatter system. Because communication periods are typically very short in energy-constrained systems, staying in a matched state during idle periods makes these losses negligible.

5. A LINK-LAYER DESIGN FOR FULL-DUPLEX BACKSCATTER

The hardware design in the preceding section enables an entirely new set of protocols for general-purpose backscatter devices. In this paper, we explore one example instantiation of a link-layer protocol for backscatter communication. In particular, our protocol tackles the three problems described in §2: wireless collisions, error correction, and rate adaptation.

In this section, we first explain how RF power harvesting can affect protocol design. We then describe our link-layer protocol that uses the feedback channel to address the above problems.

5.1 Design Principle

Our protocols strive to minimize the amount of energy required to transmit a single packet successfully. To see why this is desirable, we need to look at the workings of an RF energy harvester.

The purpose of a device’s harvesting circuitry is to charge up a capacitor that can be discharged whenever the device needs to perform computation, sensing, or communication. At a high level, harvesting happens in two stages: rectification and energy storage. When a signal arrives through the antenna, it translates into a voltage value, say, $V(t)$. The rectification stage in addition to converting the time-varying AC signal to a DC voltage, also acts as a voltage multiplier i.e., it raises the voltage of the signal by trading off current. The resulting DC voltage is applied to the storage capacitor in order to charge it up.

In an ideal world, one could convert arbitrarily weak voltage signals into arbitrary amounts of energy in two ways:

- *Increasing the voltage value.* Adding more stages to the rectification stage will increase the voltage multiplication effect [36]. This can potentially be used to increase the voltage applied to the capacitor and therefore the total energy stored by the capacitor.
- *Increasing the capacitor size.* Increasing the capacitance of the capacitor increases the amount of charge it can hold for a given voltage according to $Q_{max} = CV$.

Unfortunately, in practice, both of the above approaches come with significant tradeoffs that limit their applicability. First, adding more stages to the voltage multiplier decreases harvesting efficiency and also significantly increases power losses due to parasitic capacitances and leakage effects [17]. Second, increasing the size of the capacitor increases the charge time disproportionately. For instance, doubling the size of the capacitor would increase the time to fully charge by more than 8x. Furthermore, larger capacitors generally have greater leakage currents, decreasing the sensitivity of the harvester and charge time even further.

The above discussion implies that we need to leverage our feedback channel to minimize the need for multiple packet transmissions that requires significant amounts of energy.

5.2 Link-layer Protocol

An instantaneous feedback channel allows us to design a MAC protocol that minimizes the amount of energy required to send a packet. Our protocol reduces the penalty associated with collisions and aids in detection of hidden terminals, all with minimal energy

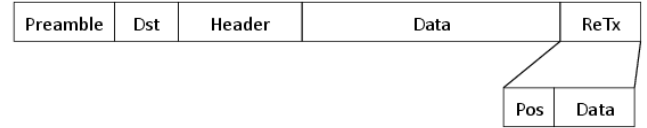


Figure 6: **Packet format:** The packet format used on the data channel. At the end of every packet, the transmitter can optionally append retransmissions of bit chunks and their positions.

requirements. It is inspired by a variety of related work including CSMA/CD [20], busy tones [12], μ ACKs [37], and others [21], but the underlying mechanism is simple: use the feedback channel to acknowledge data transmissions at a fine granularity.

For ease of exposition, we describe the feedback and data channels separately. We begin by describing the feedback channel operation in the common case; we later describe how the data channel takes advantage of feedback.

5.2.1 Feedback Channel

Our system uses the feedback channel to acknowledge ongoing data transmissions. We consider a source sending a packet using the format shown in Fig. 6. The destination uses the feedback channel to perform the following operations:

1. The destination first decodes the transmitted preamble and the first field of the header (i.e., the destination address). As soon as the destination realizes that it is the intended recipient, it begins transmitting a preamble on the feedback channel.
2. The destination will then divide the packet (including header) into chunks of b bits. For each group of b bits, it computes a c -bit checksum. The destination transmits the checksum back to the source on the feedback channel.

The ratio of b to c is determined by the difference in transmission rate between the data and the feedback channels. For instance, a 1 kbps data channel with a 100 bps feedback channel, must satisfy the following condition: $\frac{b}{c} = 10$. The time for both the data and feedback transmissions are therefore approximately equal, with the feedback channel lagging behind the data channel slightly. Our prototype implementation uses 40 and 4 bits for b and c respectively.

5.2.2 Data Channel

The source uses the above feedback channel to adapt its own transmissions to errors and collisions. In particular, it performs the following protocol:

1. The source first listens on the medium to ensure there are no existing transmissions.⁵ If the channel is empty, it begins to transmit the preamble, header, and payload. Otherwise, it exponentially backs off before retrying. See §5.3 for details.
2. While transmitting, the source continues to listen on the feedback channel for errors. It uses the incoming checksums to change rates at the level of bits rather than entire packets. If it detects a collision, it will terminate the transmission and back off accordingly. See §5.4 and §5.3 for details.
3. At the end of the packet, any failed groups of bits are retransmitted along with a short header denoting their position in the original stream. See §5.5 for details.

5.3 Dealing With Collisions

There are two places where collisions can occur: (a) at the beginning of the packet, during the preamble or destination address

⁵In backscatter-style communication, this can be effectively implemented by simply checking the output of the receive chain for bit transitions [18].

and (b) in the middle of the packet. Collisions at the beginning of the packet can occur when two nodes start to transmit simultaneously or when the new transmitter cannot hear the existing connection. These collisions can either interfere with the forward channel or the feedback channel, but in both cases, no feedback will return and the data channel transmitter will assume a collision has happened. If the collision happens in the middle of a packet, we leverage our bit-level rate adaptation to account for the resulting bit errors. Further, devices assume that a collision has occurred after multiple (more than one) consecutive drops happen in the data rate during our bit-level rate adaptation.

Backoff is implemented similarly to existing protocols: the detecting transmitter will wait for a random number of time slots between 0 and 2^r , where r is the number of retries attempted by the device. Our link-layer protocol detects collisions, even from hidden terminals, and minimizes their effects. We note the following about the effect of collisions in our system.

Firstly, by terminating the packet as soon as a collision is detected, our design can minimize the energy penalty associated with collisions. Specifically, recharging the energy spent during the short amount of time before collision detection at the transmitter takes significantly less time than recovering the energy required for an entire packet.

Secondly, the transmitter that detects the collision does not need to jam/inform the any other transmitters because bit-level error correction can correct from any collision-related errors. As long as the preamble and header are decoded correctly, any subsequent bit errors can be corrected as long as the network ends up with a single active transmitter-receiver pair.

Thirdly, the feedback channel is compatible with our collision detection technique because transmitters should see a checksum of their own transmissions. Any competing transmission will cause the bits on the feedback channel to differ from the expected bits significantly, allowing the transmitter to detect a collision.

Finally, we note that false positives in detecting collisions can occur due to bit errors; however a conservative collision detection estimate is preferable to retransmission of the entire packet.

5.4 Adapting Rate at a Fine Granularity

Rate adaptation proceeds in a fashion similar to existing techniques, except at the level of b -bit chunks rather than entire packets. The exact algorithm used is orthogonal to our work, but we take [30] as a baseline. At a high level, the goal of the rate control algorithm in [30] is to maximize throughput while occasionally testing alternative rates ($\sim 10\%$ of the time). Throughput for a particular rate is defined by (*Probability of success* * *Transmission rate*), where the probability of success is based upon an exponentially-weighted, moving average of chunk-level success rates. If through the occasional probes, we find an alternative rate with a higher throughput, we will switch to that rate. If the line code is self-clocking, the receiver will adjust to rate changes automatically.

The above protocol allows us to adjust rates at the level of chunks of b bits, rather than entire packets. Adjusting the rate at a fine granularity gives devices the ability to react very quickly to changes in channel state. More importantly, we can react within the span of a single packet—a useful benefit since bit rates are often very low in backscatter systems. It also allows us to react after fewer packet drops, as looking at small chunks of bits gives us a much greater sample size for rate adaptation rather than just coarse packet drops.

5.5 In-Packet Error Correction

Finally, the source can use the feedback channel to decrease the energy penalty of bit errors. Specifically, error correction in our sys-

tem proceeds in a manner similar to that of [37]: the transmitter re-sends failed bit sequences at the end of each packet along with their position in the original stream.

This type of error correction is necessary to make the above bit-level rate adaption useful—without it, errors in the packet would make the entire packet useless. The additional benefit to being able to tolerate bit errors in this fashion is (as mentioned above) that collisions no longer necessarily cause all transmissions to be wasted.

6. PROTOTYPE IMPLEMENTATION

We implement a Full-duplex Backscatter prototype shown in Fig. 1 on four-layer printed circuit boards (PCB) using off the shelf components. The PCBs were designed in Altium software and were fabricated by Sunstone circuits. Our prototype uses a dipole antenna consisting of two sections of 3-inch long 16 AWG magnetic copper wire. We implemented prototypes capable of 1 kbps transmission on the data channel and 100 bps on the feedback channel. The hardware is tuned to operate at frequencies in the 3 MHz band centered at 920 MHz. Since the bit rates used by backscatter data communication are significantly lower than traditional radio communication (e.g., Wi-Fi), a 3 MHz data bandwidth does not limit the underlying RF source bandwidth. Specifically, as long as the data bandwidth is less than 3 MHz, the system would work even with an underlying RF source with higher bandwidth (e.g., TV signals). This is because the circuit would low-pass filter out the higher frequency components. Evaluating feedback hardware for the TV band is however not in the scope of this paper.

The transmitter is implemented using the ADG919 RF switch [1] connected directly to the antenna terminals. On the data channel, the transmitter’s S_3 impedance states are implemented using a series combination of a $10\ \Omega$ resistor and a $1\ nH$ inductor, while the S_4 state is implemented using a $1.8\ pF$ capacitor. For the feedback channel, the transmitter’s S_3 impedance states are implemented using a parallel combination of a $6.8\ nH$ inductor and a $50\ \Omega$ resistor, while the S_4 state is implemented using a series combination of a $15\ \Omega$ resistor and a $0.8\ pF$ capacitor. Note that these impedance values are specific to the choice of the switch and impedance network on the PCB. The power consumption of our analog transmitter design (including the switch) is about $0.25\ \mu W$.

On the receiver we use a TS881 [3] as the ultra-low power comparator. On the data channel’s receiver, the capacitor and resistor values R_1 , R_2 , C_1 , and C_2 shown in Fig. 4 are set to $100\ k\Omega$, $10\ M\Omega$, $4.7\ nF$ and $1.47\ nF$, respectively. On the feedback channel’s receiver, the capacitor and resistor values R_3 , R_4 , C_3 , and C_4 are set to $100\ k\Omega$, $10\ M\Omega$, $27\ nF$, and $220\ nF$ respectively. The analog prototype receiver (including the comparator) consumes $0.54\ \mu W$.

7. EVALUATION

Next, we evaluate various practical aspects of our Full-duplex Backscatter prototype. In particular, we first measure the effect of our low-power backscatter cancellation technique described in §4.2 as a function of the input power level. Next, we evaluate the effect of the feedback channel on the bit error rate (BER) of the data channel at the receiver. We also evaluate the bit error rate (BER) achieved on the feedback channel as a function of the distance between the backscatter devices. Finally, we evaluate the feedback channel with collisions, retransmissions, and rate adaptation.

7.1 Full-duplex Backscatter Cancellation Effectiveness

In this section, we evaluate how well our design in §4.2 reduces the self-interference from the feedback channel.

Experiments: To do this, we examine the degree to which we decrease self-interference at the destination (i.e., the device sending

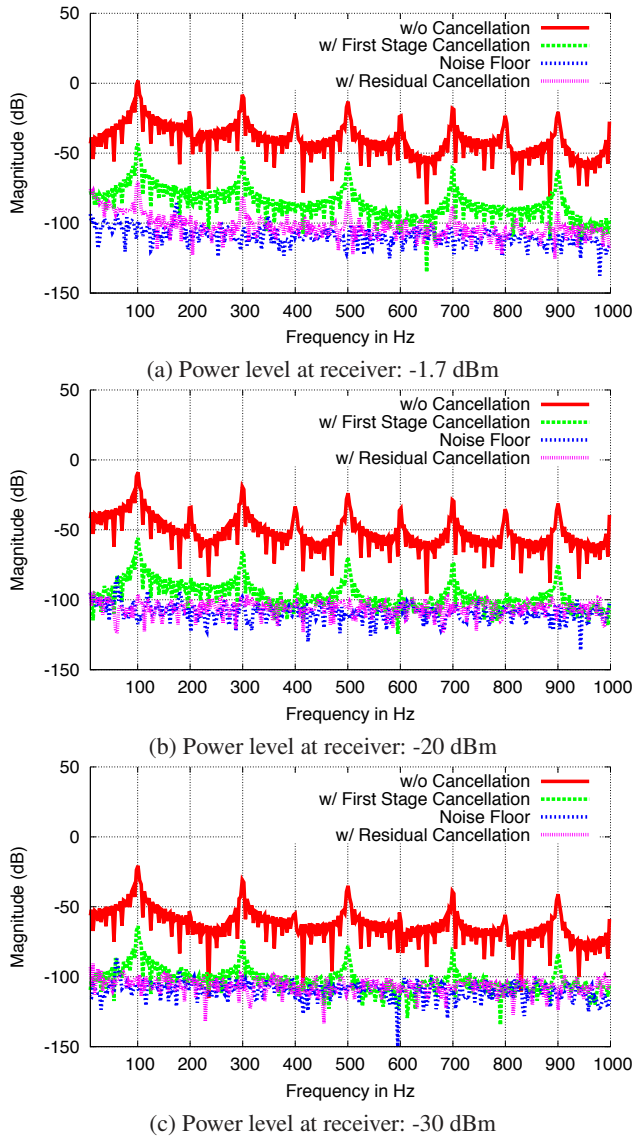


Figure 7: **Self-interference Cancellation in Full-duplex Backscatter:** The graphs show the strength of the voltage signal received over a frequency spectrum due to the device’s own 100 bps transmissions. Our technique reduces the self-interference from the feedback channel down to the noise floor of the device across the frequency range. We note that the typical power level at the receiver is less than -15 dBm.

the feedback) using the transmitter cancellation (described in §4.2) and the residual cancellation techniques (described in §4.3). We place the Full-duplex Backscatter prototype device in the presence of a continuous wave transmission from an RFID reader at 920 MHz and configure it to continually transmit an alternating sequence of bits at 100 bps. Note that this is one feedback bit for every 10 data bits on an ambient backscatter communication system, which as we show later is sufficient to address the networking issues described in §2. We then tap into the outputs of the envelope detector and the low pass filter on our prototype board and connect it to an ADC to get direct access to the voltage values.

Results: Fig. 7 shows the magnitude of the voltage values (in dB) across the frequency spectrum. Each plot shows the frequency-domain representation of the received voltage values with just the

transmitter cancellation technique, with both transmitter and receiver cancellation techniques (to cancel the residual interference as described in §4.3) and without our cancellation technique. They also depict noise values when there is no transmission on the feedback channel. We plot the results for three different received power levels of the RFID continuous wave signals at the receiver. The graphs show the following:

- In the absence of our cancellation circuit, the plots show spikes at frequencies corresponding to the odd multiples of 100 Hz. This is expected because the bit rate on the feedback channel is 100 bps which results in a spike at 100 Hz. Further since the transmitted signal on the feedback channel is approximately a square wave, we also see spikes at frequencies that are odd multiples of 100 Hz.
- Our technique reduces the interference from the feedback channel to the noise floor of our device across the frequency range. This is very significant and is impressive since this approach requires near-zero power compared to a conventional device.
- The reduction in self-interference is similar for different power levels of the received signal. This results in higher power levels having higher residual interference (in comparison to noise). In the next section, we will demonstrate how this cancellation technique results in comparable receive bit error rates for systems with and without the feedback channel.
- The self-interference reduction varies by about 5-8 dB across time and also the 3 MHz operational bandwidth. This variation is expected because the conjugate impedance network used for the cancellation is implemented using off the shelf resistors, capacitors and inductors that have tolerances and variations that change with environmental conditions.

7.2 Effect on Data Channel BER

Experiments: We place two prototype devices at different distances from each other. The devices are both configured to transmit an alternating sequence of bits at the same time—the sender transmits data bits at a rate of 1 kbps, and the receiver of the data packet transmits bits at a rate of 100 bps on the feedback channel. We place the two devices in the presence of an RF signal source that broadcasts continuous wave RFID signals centered at 920 MHz. The source has a dipole antenna and is placed equidistant from both the tags. To analyze the BER, we connect an NI myDAQ to the output of the receiver’s receive chain. We capture 100 seconds of data at each distance value, which corresponds to a total of 10^5 bits; when no bit errors occur we set the BER to 10^{-5} . We consider distance values of up to 1.5 meters, which spans the communication range of the ambient backscatter devices in [18]. We also vary the RF power to span the receive power levels from -18 dBm to -4 dBm; the lower value is the minimum power level at which the harvester [25] works.

Results: Fig. 8 shows BER versus the average received power at our prototype devices for three different tag-to-tag distances. We also plot as a baseline the bit error rate results for the ambient backscatter prototype that we replicate from [18] that does not have a feedback channel. The results show the following:

- For both Full-duplex Backscatter and existing backscatter systems, as the distance between the devices increases the BER increases. Similarly, BER reduces as the RF power level at the prototype devices increases.
- We do not observe any bit errors for both the systems when the distance between the tags is less than or equal to 1 m. As the distance increases to 1.5 m, we start observing bit errors for lower power values. We note that in these cases, the observed bit error

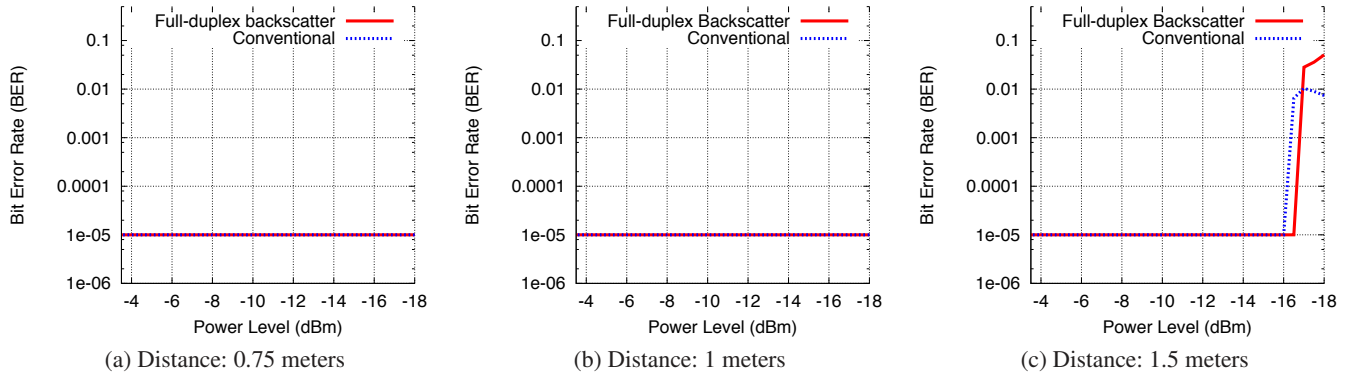


Figure 8: **Data BER versus power:** BER as observed by the receiver of the forward data channel. We show the variation in BER versus power at the tag for three different tag-to-tag distances. The plots show that the feedback channel does not significantly affect the data BER.

rate for our system across all three distances is comparable to conventional backscatter.

- In some cases, the bit error rate for our system is worse than existing backscatter systems and is better in other scenarios. This is because existing backscatter systems modulate information by changing the amplitude of the transmission (i.e., using ASK modulation). In contrast, we encode information in phase values (i.e., effectively use PSK modulation). Thus, the signals in our design interfere differently than those of a traditional backscatter system and hence exhibit the above noted behavior. The key point however is that, the feedback channel in our results does not significantly affect the performance of the forward data channel.

7.3 BER of the Feedback Channel versus Distance

Next, we evaluate the performance of the feedback channel. Specifically, we measure the bit error rate on the feedback channel by measuring its BER.

Experiments: As before, we place two prototype devices in the presence of an RF source broadcasting an RFID continuous wave transmission at 920 MHz. The data rate and feedback rate are set to 1 kbps and 100 bps respectively. We vary the distance between the two devices and measure the bit error rate on the feedback channel. We again measure the BER with the assistance of an NI myDAQ and compare the received feedback bits with those transmitted. We transmit a total of 10^{-4} bits in each experiment. Note that this value is less than that of the previous set of experiments since the feedback channel has a lower bit rate than the data channel. We set the bit error rate to 10^{-4} in experiments which do not see any bit errors. We run the experiments at a fixed transmit power of 7.5 dBm at the RF source. The BER trends are similar at other power levels.

Results: Fig. 9 shows the BER of our feedback channel at different distances between the two prototype devices. The figure shows that the feedback channel sees bit errors as we approach 2 meters between the transmitter and receiver. However, we do not see any bit errors even at distances greater than 1.7 meters. Further, the observed feedback BERs match very well with the observed BERs on the data channel. This implies that the feedback channel is reliable enough to not be a limiting factor.

8. EVALUATING FULL-DUPLEX BACKSCATTER'S NETWORK STACK

Finally, we evaluate how our feedback-channel-enabled network protocol addresses collisions, retransmissions, and rate adaptation.

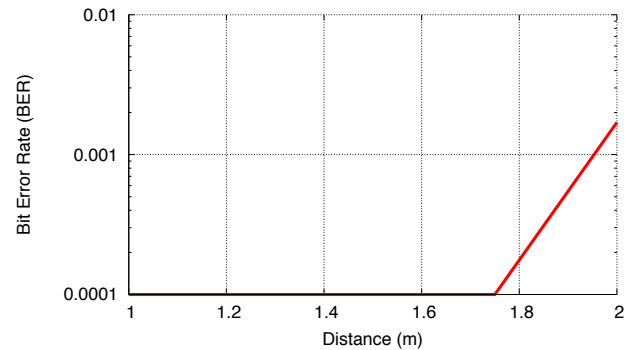


Figure 9: **Feedback BER versus distance:** BER observed on the feedback channel. We show the variation in BER versus distance between the two devices.

8.1 Recharge-time Reduction during Collisions

As described in §5.3, a feedback channel allows our prototype device to terminate its transmission when it detects a collision. We use the metric of recharge time to evaluate the benefit of this feature. Terminating wasteful transmissions can result in a reduction in the recharge time at the transmitter since it can conserve power that is otherwise wasted in transmitting a undecodable packet. This metric is very useful because recharging and duty cycling tend to be the bottleneck in virtually all energy-constrained systems.

Experiments: To evaluate the benefits of Full-duplex Backscatter in this context, we measure the time it takes to recover from a collision. Our prototype devices use a preamble length of eight bits and transmit at a rate of 1 kbps on the data channel (similar to the design in [18]). The feedback channel also uses a preamble length of eight bits and has a bit rate of 100 bps. We measure the time it takes for the transmitter to detect an existing transmission and then terminate its own. We then calculate the total delay from the beginning of the canceled transmission to the time at which the device has enough power to retry transmitting the packet, given a packet size of 64 bytes. We repeat the measurements for both scenarios, with and without the feedback channel. In the absence of the feedback channel the transmitter continues transmitting the whole packet even in the presence of the collision.

Results: Fig. 10 shows results for the above experiments. We plot the recharge time it takes to collect enough power to retransmit the collided packet as a function of the power level received at the prototype device. The figure shows the graphs for both conventional and our systems. The results show the following:

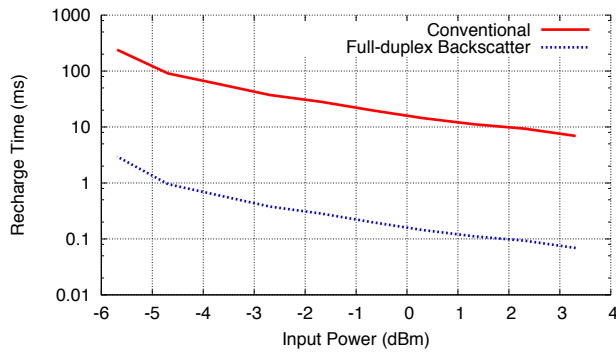


Figure 10: **Recharge time reduction:** The time needed to recover from a packet collision in Full-duplex Backscatter versus conventional transmitters. The x-axis plots the available power at the device. Our system reduces the recharge time by two orders of magnitude.

- Our system decreases the time between retries by two orders of magnitude. The savings are expected to be higher for packet sizes longer than 64 bytes. This is because recharge times are nonlinear—as the number of wasted bits grows, the recharge time non-linearly increases. Thus terminating the transmission earlier minimizes the recharge time it takes to collect enough energy for packet retransmissions.
- The recharge time is higher at lower power levels. This is expected because in power-constrained scenarios, the capacitor takes longer to collect the required energy. This is further complicated by leakage issues as explained in §5.1. We note that Full-duplex Backscatter can provide orders of magnitude reduction across all power levels, specifically in power-constrained scenarios.

8.2 In-frame Error Correction

We next look at the effectiveness of our feedback channel in achieving in-frame error correction. The error correction mechanism described in §5.5 decreases latency and increases throughput. This is because it allows us to recover from bit errors by sending a few extra bits at the end of a packet, rather than re-sending the entire packet. We evaluate these benefits in practice.

Experiments: The backscatter transmitter sends bits at a bit rate of 1 kbps using a packet size of 64 bytes. The receiver computes a 4-bit checksum for every chunk of 40 received bits. It then transmits these checksum bits on the feedback channel to the transmitter. The transmitter uses the checksum to detect when chunks of bits are incorrectly decoded at the receiver. The transmitter then retransmits these chunks at the end of the packet as described in §5.5. We change the distance between the transmitter and the receiver to span a range of bit error rates. We calculate the additional bytes of data that are necessary to deliver a fully correct packet using our prototype. We also repeat the experiments for existing backscatter systems that have to retransmit complete packets until at least one is successfully decoded.

Results: Fig. 11 shows the results of our experiments. The figure plots the overhead, i.e., the number of extra bytes necessary to deliver a single, correct packet as a function of the bit error rate observed on the data channel. The average overhead is plotted in log-scale. The figure shows the following:

- The average overhead incurred by our system is at least an order of magnitude less than in conventional systems, even for low error rates. This is because, for a single error, our system would need to transmit an additional 47 bits. On the other hand, a conventional system would need to retransmit the entire packet (i.e., 64 bytes) assuming the second packet arrives error-free.

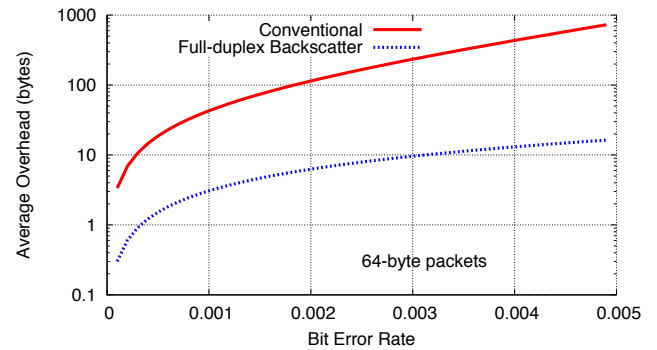


Figure 11: **Overhead of transmitting a single packet:** This graph shows the average number of extra bits transmitted for a given error rate. The extra bits are either in the form of a retransmitted packet in the case of conventional receivers or retransmitted bit sequences in the case of Full-duplex Backscatter. Our system reduces the overhead by at least an order of magnitude.

- Further, as the average bit error rate on the data channel increases, the overhead for existing backscatter systems increases much faster than the overhead incurred on our system. In Fig. 11, the Full-duplex Backscatter system has an order of magnitude less overhead for low error rates, but as the BER increases to 0.005, the overhead increases to two orders of magnitude. This is due to the fact that the probability that a 64-byte packet is correctly decoded becomes increasingly small at higher BERs. On the other hand, the probability that a chunk of 40 bits is decoded without error grows at a much slower rate.
- We note that Full-duplex Backscatter’s advantage over conventional backscatter systems in both of the above points increases with larger packet size. We also note that the actual cost of this overhead in terms of time, energy, and throughput is much higher than shown here. The underlying issue is recharge time. As we saw in the previous section, charge time magnifies the penalty associated with retransmission of extra bits. Full-duplex Backscatter can reduce this charge time and hence alleviate these issues.

8.3 Rate Adaptation

Finally, we evaluate our intra-packet rate adaptation.

Experiments: We place a transmitter and a receiver in the presence of our 920 MHz signal source. We adapt the rate of the transmitter between three rates of 100 bps, 1 kbps, and 10 kbps. We compare three rate adaptation strategies: “slowest”, which always chooses 100 bps; “packet-level”, which models an idealized algorithm that chooses the best rate for the SNR at the beginning of the packet; and “bit-level”, which implements the algorithm described in §5.4. Since we would like to compare the three algorithms in the same scenarios, we connect our receiver to a myDAQ that continually takes measurements of the received voltage values. The transmitter is set to continuously transmit bits as we move the receiver in the area around the transmitter at an average speed of about 3 m/s. From the captured traces, we compute the SNR at every time instance. Since the SNR determines the achievable bit rate, we use it to compute the achievable throughput of the three algorithms given the captured voltage traces.

Results: In Fig. 12, we plot the throughput of each rate adaptation strategy. Our results show the following:

- Bit-level rate adaptation increases performance by $\sim 33\%$ compared to an idealized SNR-based packet-level algorithm. Further, when compared to the slowest-rate strategy, the speedup is almost 3x. This is because devices can take advantage of and protect against channel changes. We also note that since backscatter

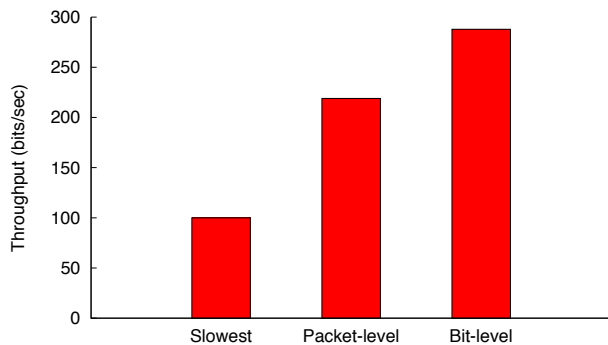


Figure 12: **Effect of Rate Adaptation:** A graph of the throughput for different rate adaptation algorithms. The throughput is calculated using real channel traces with an average mobility of 3 m/s.

devices do not have ADCs, they do not have SNR information. Thus, we can expect higher throughput gains in practice.

- We note that about 30% of packets in the packet-level strategy failed due to changes in the channel within the span of a single packet. This is due to the fact that transmission rates are extremely low in these types of systems.

9. RELATED WORK

There has been recent interest in improving the performance of backscatter systems [7, 26, 15]. Prior work has also proposed coding techniques to address the problem of collisions and increase the efficiency of RFID networks [6, 22, 35]. More recently, work on ambient backscatter [18] enabled two battery-free devices to communicate by backscattering signals from an RF source. Our work builds on these foundational works and achieves the first instantaneous feedback channel for backscatter communication systems.

Our work is also related to recent research on full-duplex communication that uses passive cancellation techniques [12], active cancellation (both in the analog and digital domain) [5, 29], or a combination of the above [9, 8]. While these approaches are effective for traditional radio communication, their reliance on space-consuming antenna arrays and/or power-hungry cancellation techniques are not applicable to ultra-low power backscatter communication, that has a power budget which is orders of magnitude lower. We also note that our goal is not to design a full-duplex radio. Instead we set out to create a low-rate feedback channel for backscatter devices that can address many of the higher-layer issues facing these devices.

Similarly, there has also been previous work on implementing full-duplex communication through clever combinations of signal processing, modulation and coding techniques [19, 10, 4]. Like the above approaches, these assume that at least one of the participants are powered and can therefore perform relatively complex synchronization and decoding that are not applicable to the types of devices that we investigate in this paper.

Full-duplex Backscatter is also related to prior work that implements QAM transmissions on RFID tags by using complex impedances [33]. Our work is similar in that we also use complex impedances to adjust the phase of the reflected signal; however, the purpose of our system is complementary and is to enable a feedback channel between the backscatter devices and demonstrate the benefits of such a channel for the link- and the network-layers.

Finally, there is a large body of work related to network-layer protocols for wireless sensor network or other types of mesh networks. Some of these works deal with half-duplex mesh networks and how to implement MAC protocols [32], rate adaptation [14], and hidden terminal detection [31]. Others are designed for du-

plexed communication [12]. In contrast, the focus of this paper is to enable a instantaneous feedback channel for battery-free backscatter devices, which is a goal that is complementary to prior efforts.

10. CONCLUSION

Energy, and specifically the energy required for communication, is a key bottleneck in the design of computing devices. Recent developments in backscatter communication promise to remove this bottleneck. Specifically, they promise to provide a way for devices to send bits to one another using orders of magnitude less power than is required today. However, an essential question remains: how do we design link- and network-layer protocols for these networks. Existing protocols are ill-suited to these devices where even the transmission of a single packet can exhaust all available energy.

In this paper, we introduce a novel technique called Full-duplex Backscatter that enables almost-zero-power, instantaneous feedback and use it to implement a network stack tailored for battery-free devices. Our technique uses passive, analog circuitry in order to allow for a low-rate feedback channel that operates alongside any data transfer. In addition, we presented a network stack that uses Full-duplex Backscatter to minimize the energy wastes associated with MAC, rate adaptation, and error correction. We believe that our technique improves the practicality of battery-free devices and brings us closer to having a practical, generalized backscatter communication system.

Acknowledgements: We would like to thank Matt Reynolds for helpful discussions during the design of Full-duplex Backscatter's hardware and the anonymous MOBICOM reviewers for their helpful comments. This research is funded in part by Qualcomm Innovation Fellowship, Washington Research Foundation gift, NSF, and University of Washington.

11. REFERENCES

- [1] Adg919 rf switch datasheet, analog devices.
- [2] Software defined radio hardware survey, scott johnston.
- [3] Ts 881 datasheet, smtmicroelectronics, july 2012.
- [4] D. Bharadia, K. R. Joshi, and S. Katti. Full duplex backscatter. In Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks, 2013.
- [5] D. Bharadia, E. McMillin, and S. Katti. Full duplex radios. In Proceedings of the ACM SIGCOMM, pages 375–386. ACM, 2013.
- [6] M. Buettner. Backscatter Protocols and Energy-Efficient Computing for RF-Powered Devices. PhD thesis, University of Washington, Seattle, WA, 2012.
- [7] M. Buettner, B. Greenstein, and D. Wetherall. Dewdrop: An energy-aware runtime for computational rfid. In Proceedings of NSDI'11, 2011.
- [8] M. Duarte. Full-duplex Wireless: Design, Implementation and Characterization. PhD thesis, Rice University, May 2012.
- [9] M. Duarte and A. Sabharwal. Full-duplex wireless communications using off-the-shelf radios: Feasibility and first results. In 2010 ASILOMAR, pages 1558–1562, Nov 2010.
- [10] D. Guo and L. Zhang. Virtual full-duplex wireless communication via rapid on-off-division duplex. CoRR, abs/1010.2667, 2010.
- [11] E. Inc. Universal software radio peripheral. <http://ettus.com>.
- [12] M. Jain, J. I. Choi, T. Kim, D. Bharadia, S. Seth, K. Srinivasan, P. Levis, S. Katti, and P. Sinha. Practical, real-time, full duplex wireless. In Proceedings of the 17th

- annual international conference on Mobile computing and networking, pages 301–312. ACM, 2011.
- [13] B. Kellogg, V. Talla, and S. Gollakota. Bringing gesture recognition to all devices. In NSDI, 2014.
- [14] S. Kim, R. Fonseca, P. Dutta, A. Tavakoli, D. Culler, P. Levis, S. Shenker, and I. Stoica. Flush: a reliable bulk transport protocol for multihop wireless networks. In Proceedings of the 5th international conference on Embedded networked sensor systems, pages 351–365. ACM, 2007.
- [15] M. Kodialam and T. Nandagopal. Fast and reliable estimation schemes in rfid systems. In Proceedings of Mobicom 2006.
- [16] C.-J. M. Liang, N. B. Priyantha, J. Liu, and A. Terzis. Surviving wi-fi interference in low power zigbee networks. In Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems, pages 309–322. ACM, 2010.
- [17] M. M. Liu. Demystifying switched capacitor circuits. Newnes, 2006.
- [18] V. Liu, A. Parks, V. Talla, S. Gollakota, D. Wetherall, and J. R. Smith. Ambient backscatter: Wireless communication out of thin air. In Proceedings of the ACM SIGCOMM, 2013.
- [19] J. MacLellan, R. Shober, G. Vannucci, and G. Wright. Tag for use in a radio communication system, July 15 1997. US Patent 5,649,296.
- [20] R. M. Metcalfe and D. R. Boggs. Ethernet: Distributed packet switching for local computer networks. Commun. ACM, 19(7):395–404, July 1976.
- [21] A. Mishra, S. Rayanchu, D. Agrawal, and S. Banerjee. Supporting continuous mobility through multi-rate wireless packetization. In HotMobile, 2008.
- [22] C. Mutti and C. Floerkemeier. Cdma-based rfid systems in dense scenarios: Concepts and challenges. In RFID, 2008.
- [23] P. V. Nikitin, S. Ramamurthy, R. Martinez, and K. Rao. Passive tag-to-tag communication. In RFID, 2012.
- [24] C. Occhiuzzi, S. Cippitelli, and G. Marrocco. Modeling, design and experimentation of wearable rfid sensor tag. Antennas and Propagation, IEEE Transactions on, 58(8):2490–2498, 2010.
- [25] A. N. Parks, A. P. Sample, Y. Zhao, and J. R. Smith. A wireless sensing platform utilizing ambient rf energy. In BioWireless, 2013 IEEE Topical Conference on, pages 154–156. IEEE, 2013.
- [26] B. Ransford, J. Sorber, and K. Fu. Mementos: system support for long-running computation on RFID-scale devices. SIGPLAN Not., 46(3):159–170, Mar. 2011.
- [27] S. Roy, V. Jandhyala, J. Smith, D. Wetherall, B. Otis, R. Chakraborty, M. Buettner, D. Yeager, Y.-C. Ko, and A. Sample. Rfid: From supply chains to sensor nets. Proceedings of the IEEE, 98(9):1583–1592, sept. 2010.
- [28] A. Sample, D. Yeager, P. Powlledge, A. Mamishev, and J. Smith. Design of an rfid-based battery-free programmable sensing platform. IEEE Transactions on Instrumentation and Measurement, 57.
- [29] S. Sen, R. Roy Choudhury, and S. Nelakuditi. CSMA/CN: Carrier sense multiple access with collision notification. In MobiCom, 2010.
- [30] D. Smithies and F. Fietkau. minstrel: MadWiFi and Linux kernel rate selection algorithm. 2005.
- [31] J. So and N. H. Vaidya. Multi-channel mac for ad hoc networks: handling multi-channel hidden terminals using a single transceiver. In Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing, pages 222–233. ACM, 2004.
- [32] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis. An empirical study of low-power wireless. ACM Transactions on Sensor Networks (TOSN), 6(2):16, 2010.
- [33] S. Thomas and M. Reynolds. A 96 mbit/sec, 15.5 pj/bit 16-qam modulator for uhf backscatter communication. In RFID., 2012 IEEE International Conference on, pages 185–190, april 2012.
- [34] E. Welbourne, L. Battle, G. Cole, K. Gould, K. Rector, S. Raymer, M. Balazinska, and G. Borriello. Building the internet of things using rfid: The rfid ecosystem experience. Internet Computing, IEEE, 13(3):48–55, may-june 2009.
- [35] L.-C. Wu, Y.-J. Chen, C.-H. Hung, and W.-C. Kuo. Zero-collision rfid tags identification based on cdma. In International Conference on Information Assurance and Security, 2009.
- [36] J. Yi, W.-H. Ki, and C.-Y. Tsui. Analysis and design strategy of uhf micro-power cmos rectifiers for micro-sensor and rfid applications. Circuits and Systems I: Regular Papers, IEEE Transactions on, 54(1):153–166, 2007.
- [37] J. Zhang, H. Shen, K. Tan, R. Chandra, Y. Zhang, and Q. Zhang. Frame retransmissions considered harmful: improving spectrum efficiency using micro-acks. In Proceedings of MOBICOM 2012.