Content-Agnostic Backscatter from Thin Air

Yifan Yang
University of Science and Technology of China
ivanyang@mail.ustc.edu.cn

Longzhi Yuan
University of Science and Technology of China
longzhi@mail.ustc.edu.cn

Jia Zhao
Simon Fraser University
zhaojiaz@sfu.ca

Wei Gong*
University of Science and Technology of China
weigong@ustc.edu.cn

ABSTRACT

We present CAB, a content-agnostic backscatter system that can demodulate both tag and ambient data from ambient backscattered WiFi alone. In contrast to prior ambient backscatter systems that use ambient data (content) as tag-data carriers, we focus on zero-subcarriers, which are invariant and independent for any ambient OFDM WiFi. The idea of using zero-subcarriers to convey tag data is simple and elegant. Not only does it for the first time remove the dependency of tag-data demodulation on ambient data, but it also significantly improves the practicality of ambient backscatter.

We prototype CAB using off-the-shelf FPGAs and SDRs. Extensive experiments show CAB is universal as it can work with multi-band, multi-stream, and multi-user ambient traffic, including WiFi 3/4/5/6. CAB is also high-performing since it can deliver 340.9 Mbps aggregate throughput, reaching 97% Shannon capacity. Since CAB is general, we extend it to leverage ambient LTE traffic as excitations, and the achieved tag-data BER is below 0.002%. As the first content-agnostic backscatter that delivers near Shannon-capacity throughput, we believe CAB takes a curial step forward on ubiquitous battery-free IoTs.

CCS CONCEPTS

• **Networks** → **Network design principles**; **Sensor networks**.

KEYWORDS

Backscatter, OFDM, Internet of Things

ACM Reference Format:
Yifan Yang, Longzhi Yuan, Jia Zhao, and Wei Gong. 2022. Content-Agnostic Backscatter from Thin Air. In *The 20th Annual International Conference on Mobile Systems, Applications and Services (MobiSys '22), June 25-July 1, 2022, Portland, OR, USA*. ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3498361.3538930

*Corresponding author: Wei Gong

MobiSys '22, June 25-July 1, 2022, Portland, OR, USA

ACM ISBN 978-1-4503-9185-6/22/06…$15.00
https://doi.org/10.1145/3498361.3538930

1 INTRODUCTION

In recent years, ambient backscatter has attracted ever-growing attention as it is promising to deliver near zero-power communications for billions of tiny computing devices [18, 32, 39, 41, 42, 49, 54, 57, 60, 63]. Different from traditional radio frequency identification (RFID) communications, it has three distinct features. First, it intends to use uncontrolled ambient signals as wireless carriers, expanding excitation sources from dedicated readers to abundant signals from thin air [39, 58]. Second, it can support high-throughput backscatter communication [20, 40], while typical RFID systems support 40-640 kbps [17]. Third, unlike duplex-radios for decoding RFID signals, it only requires standard radios that support general-purpose wireless protocols, e.g., WiFi [34, 55], Bluetooth [59], as receivers.

One of the key features for ambient backscatter is the independence of excitors and receivers, which makes backscatter a general-purpose communication paradigm possible. But it is also the pain point for tag-data demodulation because ambient signals are out of control [39, 58]. In wireless communications, demodulation is all about finding differences between received signals and referenced ones [52]. For example, in active-radio communications, a binary phase-shift keying (BPSK) receiver uses continuous waves (CW) as references to obtain phase differences for demodulation. When it comes to ambient backscatter, a question arises: where to find references in uncontrolled excitation signals? To address this issue, the wisdom in prior ambient backscatter systems is *content-aware*. In particular, they first employ an additional receiver to obtain ambient data (content); then, they use those contents as references to demodulate tag data. For example, in Hitchhike [58], if a received backscatter symbol is '1' and the corresponding ambient symbol is '0', the tag symbol will be demodulated as '1', which means phase rotation π by tag modulation translates the ambient symbol '0' to '1'. FreeRider [59], PLoRa [44], LScatter [20] extend this idea to 802.11n, ZigBee, Bluetooth, LoRa and LTE signals. Those methods, however, share some common drawbacks: 1) their tag-data demodulation is completely dependent on ambient data quality. 2) their channel bandwidths are underexploited as they require two independent bands for demodulation. 3) their practicality is significantly limited due to the extra hardware cost of additional receivers, the high synchronization cost of two receivers, and constrained compatibility with multi-stream and multi-user signals.

After reviewing the problems brought by those content-aware systems, we ask a simple yet difficult question: *is it possible to demodulate tag data from ambient backscattered signals only?* A positive answer to this question would bring us closer to the ambient backscatter vision, where ambient tags can reuse rich ambient

signals as excitations. A single standard radio is adequate to demodulate backscattered signals. However, designing such a system is challenging because demodulation references are missing if no ambient signals are available.

New hope arises when we have an insightful observation: no matter what data OFDM-WiFi signals carry, zero-subcarriers are always the same: single tones. In other words, we have found perfect references because zero-subcarriers are natural-born invariants in all kinds of OFDM-WiFi signals. Based on this key observation, we present CAB, a content-agnostic backscatter system that can demodulate both tag and ambient data from ambient backscattered OFDM signals alone. Its key enabler is to novelly use zero-subcarriers for tag-data backscatter, as shown in Figure 1. Compared to prior systems, it not only removes the dependency of tag-data demodulation on ambient data for the first time, but also improves practicality of ambient backscatter due to reduced cost and expanded excitation sources. Turing the above idea into a practical system, however, faces several key challenges.

Challenge 1: how to estimate zero-subcarriers and demodulate tag data from them?

Theoretically, zero-subcarriers have a nice property that their phases of backscattered signals should be exactly the same as the phase rotations induced by tag modulation. Unfortunately, there is no direct way to obtain those values on typical WiFi receivers. Naive solutions include simple averaging or spline interpolation across pilot subcarriers [53]. Yet, due to magnitude imbalances brought by frequency-selective fading and phase error gradient caused by sampling frequency offsets (SFO), those methods would suffer from low-resolution phase estimates, which cannot support high-order phase-shift keying (PSK) modulation. In addition, the common phase error (CPE) is deeply coupled with tag-data phase rotations. To solve these two issues, we propose a Weighted-Least-Square(WLS)-based phase estimation with iterative CPE separation for demodulation, where WLS is used to obtain fine-grained phase estimates, and CPE separation is further designed to refine tag-data phases. The details are in §3.2.

Challenge 2: how to demodulate ambient data without knowing tag data?

Ambient backscatter inevitably poses new challenges for ambient-data demodulation because all the symbols experience different channel conditions caused by tag-data modulation. Hence, we introduce a customized phase tracking that eliminates tag-data phase rotations without knowing tag data beforehand, making ambient-data demodulation independent. The main difference between traditional phase tracking in standard WiFi and ours is a pre-filter that calibrates CPE and tag-data phase rotations simultaneously on a symbol-by-symbol basis, leading to low bit error rates (BERs) for ambient data. We explain how it works in detail in §3.3.

Challenge 3: how to achieve subsymbol-level synchronization on tags?

As demodulation above requires accurate pilot and data subcarriers, it implicitly demands subsymbol-accuracy synchronization for tag modulation. The main difficulty of doing so is the inability of WiFi demodulation on tags and thus, all the originally designed training fields, e.g., legacy short training field (L-STF), end up unusable. Thanks to high-accuracy independent demodulation in §3.3, we propose a joint synchronization method to break this barrier.

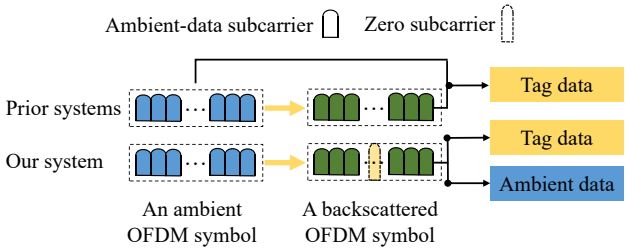

Figure 1: Our system demodulates both tag and ambient data from backscattered signals, while prior systems require both ambient and backscattered signals to demodulate tag data.

Different from prior schemes focusing solely on tags, our method relies on receiver feedback. The key idea is that the ambient-data BER is a quality indicator for synchronization on tags. Hence, we shift most computation burden to receivers while only a simple interval searching is needed for tags. The detailed workflows are presented in §3.4.

We build an FPGA-based communication prototype, a battery-free sensor prototype, and a simulated integrated circuit (IC) prototype to examine CAB's various aspects. Here are the main results.

- CAB can work with all kinds of ambient OFDM-WiFi signals, including WiFi 3/4/5/6. With WiFi 6 excitations of a 40-MHz band and two spatial streams, the maximal aggregate throughput of both ambient and tag data is 340.9 Mbps, which reaches 97% Shannon capacity.
- With WiFi 5 excitations of a 20-MHz band and two spatial streams, the maximal tag-data throughput of CAB is 1 Mbps, which is 269.7× better than that of FS-backscatter [60], the state-of-the-art content-agnostic system.
- The IC simulation shows that CAB can achieve as low as 271 μW power consumption.

Contributions: We make the following contributions:

(1) We provide an insightful observation that zero-subcarriers, which are invariants of all OFDM signals, are perfect tag-data carriers for ambient backscatter.
(2) We propose CAB, a content-agnostic backscatter system that for the first time demodulates both tag and ambient data from backscattered WiFi alone.
(3) We design a novel subsymbol-level synchronization scheme based on receiver cyclic redundancy check (CRC) verification.
(4) We implement three different prototypes to conduct comprehensive evaluations. Results show that CAB is universal, high-performing, and practical, and is ready to support various IoT applications.

2 MOTIVATION

2.1 Background

A typical ambient backscatter system consists of three parties. The *excitor* provides carriers for tag modulation and is usually unknown, e.g., ambient WiFi or Bluetooth. The *tag* transmits data to a receiver by first synchronizing with the carrier and then backscattering it. The *receiver* demodulates backscattered signals. In traditional RFID

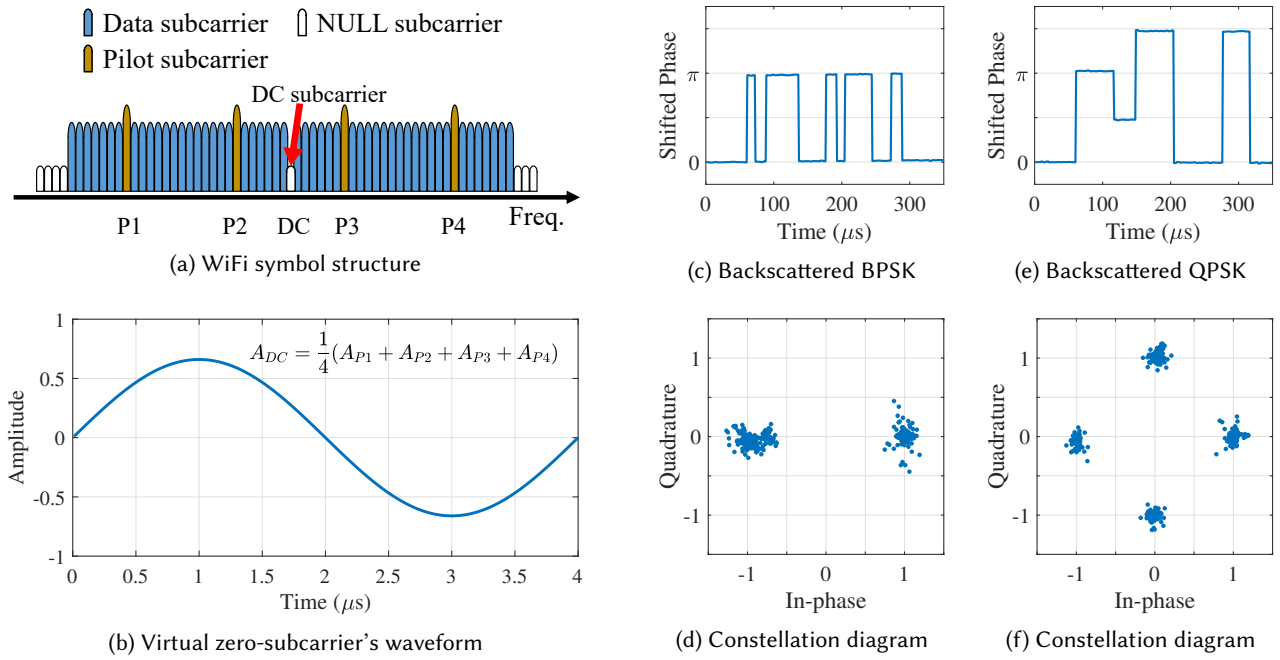

(a) WiFi symbol structure

(b) Virtual zero-subcarrier's waveform

(c) Backscattered BPSK

(d) Constellation diagram

(e) Backscattered QPSK

(f) Constellation diagram

Figure 2: (a) Pilots are centered around the direct current (DC) subcarrier (null). (b) The virtual zero-subcarrier of an uncontrolled WiFi packet is actually a single tone positioned at the DC subcarrier. (c) When the tag backscatters with binary PSK (BPSK), the phases of zero-subcarriers at the receiver have two levels, corresponding to two clusters in (d). We have similar observations in (e) and (f) for quadrature PSK (QPSK) backscatter, showing that the zero-subcarrier is a single tone in uncontrolled packets and can function as continuous waves.

communications, the excitor and receiver are the same, which is usually a reader. While ambient backscatter researchers strive to split the two in the past years, they are still somewhat connected, e.g., requiring a dedicated device to receive excitations [40, 58]. In this paper, our goal is to make them completely detached.

Orthogonal frequency division multiplexing (OFDM) is a digital data transmission technology that uses multiple carrier frequencies to transmit data. Multiple orthogonal subcarriers with overlapping spectra are transmitted in OFDM to carry data parallelly. Conventional techniques are used to modulate each subcarrier, such as PSK or quadeature amplitude modulation (QAM). This work focuses on those OFDM ambient signals.

2.2 Basic Idea

Different from prior systems putting a great deal of efforts to deal with ever-changing contents of excitations [40, 58, 61, 62], we look at this problem from a different perspective: if we can find invariants in variant signals, then we can view those uncontrolled signals as some form of 'continuous waves'. Following this idea, we observe that all the OFDM-WiFi signals surrounding us indeed contain virtual invariants. For instance, Figure 2a shows the structure of a WiFi 4 (802.11n) symbol. It has 64 subcarriers for the 20 MHz bandwidth operation, where 52 subcarriers are used for data transmission, four pilot subcarriers are for fine-grained channel estimation, one null subcarrier is for DC, and seven null subcarriers are used as guard bands. We observe that as four pilots are known and fixed, a virtual subcarrier could be created on their center, which happens to be null (no signals). We call it zero-subcarrier, whose phase and magnitude are estimated from pilots. Figure 2b demonstrates visualization results of zero-subcarriers (the average of 4 pilots) from an uncontrolled excitation packet. Clearly, it is a single tone at 250 kHz, which is perfect for carrying tag data.

Moreover, we let the tag modulate this uncontrolled WiFi packet using BPSK and observe the phases of zero-subcarriers at the receiver side. Figure 2c shows that the phases of backscattered zero-subcarriers are at two distinguishable levels. The corresponding constellation diagram in Figure 2d has two clear clusters, which means those are BPSK signals. Similar observations are made in Figures 2e and 2f when the tag modulates using QPSK. Those two examples reveal that virtual zero-subcarriers can function as continuous waves, i.e., every OFDM-WiFi ambient signal has a hidden single tone.

3 CAB DESIGN

3.1 Overview

CAB is designed to backscatter tag sensing data using ambient OFDM-WiFi signals as carriers. Upon receiving an OFDM-WiFi packet, the tag detects it and tries to synchronize with it, then modulates sensing data onto the payload using PSK. A tag-symbol can be as short as an ambient-symbol long, i.e., 3.6 μs. The tag also shifts the backscatter signals to another WiFi channel to avoid interference on the original channel [60]. The key point is that our

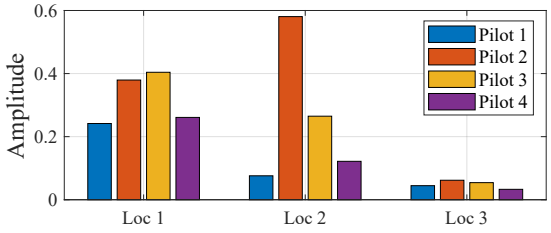

Figure 3: Diverse pilot qualities at different locations.

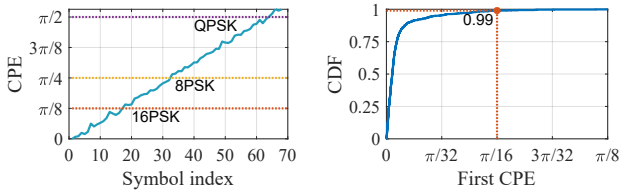

(a) CPE rotation (b) First-CPE distribution

Figure 4: CPE characteristics. (a) CPEs keep rotating across symbols in a packet. (b) Empirical CDF of the first CPEs from 1000 packets.

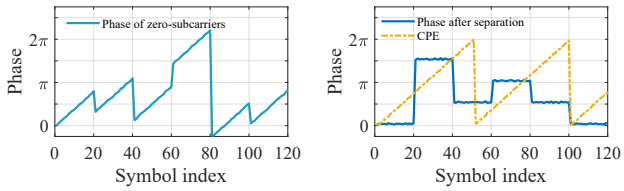

(a) Before separation (b) After separation

Figure 5: Phase before and after CPE separation.

tag modulates sensing data onto this payload-unknown packet as if it modulates a single tone. For example, if an ambient WiFi 4 packet has a guard interval (GI) of 0.4 μs [5] and the tag-symbol duration is 3.6 μs, the backscatter process is like modulating a single tone at 278 kHz. Similarly, when the ambient signal is WiFi 6 with a GI of 3.2 μs [3] and the tag-symbol duration is 16 μs , it is like transmitting PSK signals at a baud rate of 62.5 kHz.

On receiving the backscattered signal, the receiver tries to demodulate tag data without knowing any ambient data. Meanwhile, it demodulates ambient data from the backscattered signal as if it has never been backscattered. Before delving into demodulation details, we give a formal formulation of this problem. Let ϕ^{am}, ϕ^{ch}, ϕ^{cfo}, ϕ^{sto}, ϕ^{sfo}, ϕ^{tag} denote phase rotations induced by ambient data, wireless channels, carrier-frequency offset (CFO), sampling-time offset (STO), sampling-frequency offset (SFO), and tag modulation. For subcarrier j from received symbol i, the received phase is

$$\phi^r_{i,j} = \phi^{am}_{i,j} + \phi^{ch}_{i,j} + \phi^{cfo}_{i,j} + \phi^{sto}_{i,j} + \phi^{sfo}_{i,j} + \phi^{tag}_{i,j}. \quad (1)$$

Next, we present how to extract ϕ^{tag} and ϕ^{am} from ϕ^r.

3.2 Tag-Data Demodulation

As aforementioned, we demodulate tag data by focusing on the phases of zero-subcarriers, $\phi^r_{i,0}$. Doing so has several advantages. First, since no ambient data transmits on this subcarrier [1–5] and it is a virtual single tone, $\phi^{am}_{i,0} = 0$ [52]. Second, STO and SFO affect all the subcarriers except the zero-subcarrier, hence $\phi^{sto}_{i,0} = 0, \phi^{sfo}_{i,0} = 0$ [45, 53]. In addition, as we do not change the preamble of the excitation packet, all the legitimate training fields, e.g., S-LTF and L-LTF, are kept intact. Therefore, after channel estimation and CFO calibrations as in a standard WiFi receiver, the received zero-subcarrier phase before demodulation is

$$\phi^r_{i,0} = \phi^{cpe}_{i,0} + \phi^{tag}_{i,0}, \quad (2)$$

where $\phi^{cpe}_{i,0}$ is the phase rotation caused by the residual CFO [47, 50], aka "common phase error" (CPE).

We know from the above equation that if we have accurate phase estimates for the received zero-subcarriers, we can approximate tag-data phases by separating CPE. We design two major steps to do so: phase estimation for zero-subcarriers and CPE separation.

Phase estimation for zero-subcarriers. Naive ways to obtain zero-subcarrier phases include simple averaging or spline interpolation across all pilots [53]. However, both do not work for demodulation purposes. As shown in Figure 6a, the constellation points using simple averaging of pilot phases are too dispersed, leading to poor demodulation accuracy. We observe that the primary reason is the unequal qualities of pilot phases caused by frequency selective fading and in-phase and quadrature (IQ) imbalance [37, 47, 52]. As shown in Figure 3, the signal noise ratios (SNRs) of all the pilots manifest non-negligible differences at multiple locations. This motivates us to take diverse qualities of pilots into consideration for estimating zero-subcarrier phases. Moreover, although the residual STO does not affect zero-subcarriers, it brings additional phase rotations to pilot subcarriers, propagating errors in simple averaging. To solve these issues, we propose a Weighted-Least-Square (WLS)-based phase estimation for zero-subcarriers. In particular, for the k-th symbol, we perform a WLS process as follows,

$$\arg\min_{\beta_0,\beta_1} \Sigma_{i=1}^{n_p} A_i(\Phi_i - \beta_0 - \beta_1 x_i)^2, \quad (3)$$

where n_p is the number of pilots, A_i, Φ_i, and x_i are the magnitude, phase, and index of the i-the pilot. After WLS, we have β_0, the estimated phase for the zero-subcarrier $\widehat{\phi^r_{k,0}}$, and β_1, the slope of the approximated linear equation, which indicates how much STO affects pilots. From Figure 6b, we can see that through WLS, the constellation points are tightly gathered around a circle. The reason that those points are not clustered yet is there are CPE errors for $\widehat{\phi^r_{k,0}}$ and that is our next step.

CPE separation. Removing the CPE from $\widehat{\phi^r_{k,0}}$ is quite challenging due to two main factors. First, it is deeply tangled with tag-data phase rotations, as shown in equation 1, thus prior ambient backscatter systems barely discuss this point [39, 40, 60]. The hands of traditional CPE calibration schemes in WiFi receivers [45, 47, 50] are also tied as they only work when there is no tag-data induced phase involved. Second, it is ever-changing from symbol to symbol. In order to characterize CPE, we perform experiments where the tag only performs frequency shifting in backscatter, which means no

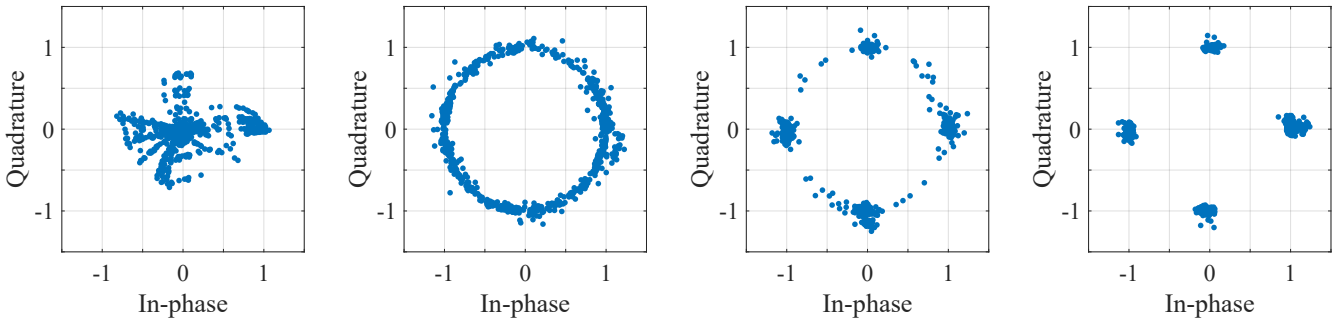

(a) Simple averaging (b) WLS (c) CPE separation (d) Symbol assembly

Figure 6: Our tag-data demodulation process. (a) Simple averaging across pilots suffers from spreading constellation points. (b) After WLS, all the points are gathered around a circle. (c) Further CPE separation almost eliminates incorrect rotations. (d) At last, accurate symbol assembly makes all the points form four tight clusters, leading to perfect QPSK demapping.

tag modulation is involved. This way, the estimated zero-subcarrier phases should be equal to CPEs. Results in Figure 4a show that CPEs keep accumulating due to the residual CFO [47]. CPEs would easily surpass the 16PSK, 8PSK, and QPSK thresholds[1] if not well-handled, bringing about numerous demapping errors [52].

To solve this, we propose an iterative CPE separation scheme based on our key observation: the CPE of the first OFDM symbol is near zero. Figure 4b demonstrates our empirical measurements of first CPEs from 1000 packets. It shows that the 99th percentile of first-symbol CPEs is only 0.2 radian. Consequently, an iterative CPE separation is designed as follows.

(1) Assign the CPE estimate from the previous symbol to that of the current symbol (i), $\widehat{\phi_{i,0}^{cpe}} \leftarrow \widehat{\phi_{i-1,0}^{cpe}}$. If the current symbol is the first one, $\widehat{\phi_{1,0}^{cpe}} = 0$.

(2) Estimate the tag-data phase by equation 2, $\widehat{\phi_{i,0}^{tag}} \leftarrow \widehat{\phi_{i,0}^{r}} - \widehat{\phi_{i,0}^{cpe}}$, and then update $\widehat{\phi_{i,0}^{tag}}$ to the phase of its closest constellation center.

(3) Update CPE using the new tag-data phase estimate, $\widehat{\phi_{i,0}^{cpe}} \leftarrow \widehat{\phi_{i,0}^{r}} - \widehat{\phi_{i,0}^{tag}}$. One iteration completes and moves on to the next symbol.

Through the above iterations done on each symbol, we successfully remove accumulating CPE as shown in Figure 5. We make the tag apply QPSK modulation every 20 symbols, and the phase of zero-subcarriers in every 20-symbol-long segment contains not only a constant tag-data phase but also a growing CPE, as Figure 5a illustrates. Our approach successfully separates the CPE, and the accumulating error disappears in those segments after separation in Figure 5b. Furthermore, as shown in Figure 6c, clusters that are similar to tag-data modulation points are clearly observed. The task of tag-data demodulation is almost complete.

Symbol assembly. One may notice that there are some stragglers in Figure 6c. After analyzing intermediate experimental results, we find those are symbol assembly errors caused by SFO. Recall that SFO happens because the sampling frequency of the transmitter is different from that of the receiver. Such a difference results in additional STO and accumulates over time, which leads to under sampling or over sampling. For example, for a 20 MHz transmission, a WiFi receiver should sample 80 points for a 4-μs OFDM symbol. After the receiver samples a number of points, if the accumulated STO is more than 50 ns, the assembled symbol would be missing or contain extra sample points from neighbor symbols. To address this, we make use of β_1, which is a good indicator of STO to coordinate symbol assembly. When the absolute value of β_1 exceeds our empirical threshold[2], the symbol will choose to drop or duplicate some sample points according to the sign of β_1.

Let us recap the whole process of tag-data demodulation. After receiving the backscattered signal, the receiver performs coarse CFO calibration, timing synchronization and fine CFO calibration, and channel estimation as in a WiFi receiver. It then estimates phases of zero-subcarriers for each symbol and separates tag-data phases from CPEs iteratively. During the iteration, OFDM symbols are accurately assembled by approximating SFO. Finally, tag data is recovered by demapping to the closest constellation centers. As shown in Figure 6d, the recovered data show a clear QPSK map with no spreading or rotation, indicating an ultra-low bit error rate.

3.3 Ambient-Data Demodulation

Now we intend to demodulate ambient data from backscattered signals, which is about data subcarriers. By reviewing equation 1, one may think of a naive solution: since we already have all the tag data, simply compensating those phase rotations brought by tag data would make backscattered signals close to original ambient signals. This way, a standard WiFi demodulation process can handle the rest. However, this solution is feasible but brings about another unpleasant issue: it implicitly creates demodulation dependency of ambient data on tag data, which is similar to the demodulation dependency of tag data on ambient data in prior backscatter systems

[1]The thresholds are half of the resolution of the corresponding PSK order.

[2]The β_1 can be calibrated for $\pm\frac{2\pi}{N}$ for each dropped or duplicated sample, where N is the number of subcarriers [46]. We set the threshold to $0.9 \times \frac{2\pi}{N}$ to cope with drastically changing β_1.

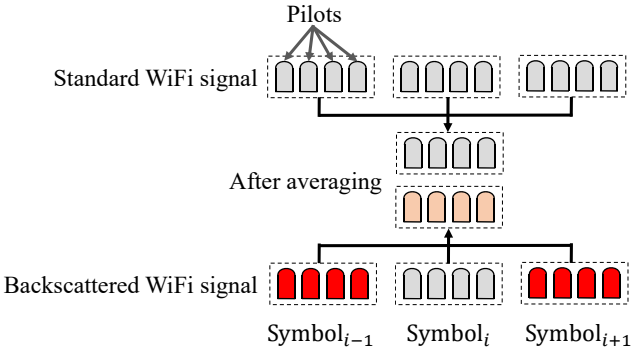

Figure 7: Pilot averaging for two different signals. While averaging works for standard WiFi, it does not work for backscattered WiFi signals as their pilots are already 'polluted' by random tag data.

[40, 58, 59, 61, 62]. Through field experiments, we find that the main difficulty of demodulating ambient data without knowing tag data is that all the symbols experience different channel conditions caused by tag-data modulation, which manifests itself in two main aspects of phase tracking.

Pilot averaging over symbols. Standard data-subcarrier demodulation in WiFi usually employs pilot averaging over a sliding window to combat varying STO caused by SFO [47]. It works well for standard WiFi but not for backscattered signals. As shown in Figure 7, if Symbol$_{i-1}$ and Symbol$_{i+1}$ are modulated by tag phase π, their pilots have different channel conditions as others. Thus, their averaged pilots are not the expected mean and are uncertain due to random tag data. Such 'polluted' averaged pilots cause incorrect demodulation for data subcarriers.

Inter-symbol phase unwrapping. Another challenging issue is inter-symbol phase unwrapping, a standard operation in WiFi for computing how fast SFO-induced phase errors rotate across different subcarriers [47]. Such a phase unwrapping scheme fails as the phases of data subcarriers from neighbor symbols are not continuous anymore. Tag modulation introduces unexpected and random phase jumping, causing errors in tracking phases.

It seems that the above difficulties make demodulating ambient data without knowing tag data nearly impossible to achieve. The situation turns around when we take a different perspective: is it possible to remove tag-data phase rotations without knowing them? The answer is yes. From §3.2, we know that the tag-data induced phase is deeply coupled with CPE in the zero-subcarrier phase. So, if we subtract the zero-subcarrier phase from all the subcarriers in a symbol, the tag-data phase rotation is implicitly removed regardless of different tag data modulation. Meanwhile, as we observe that cross-symbol operations are the key for incorrect pilot averaging and phase unwrapping of backscattered WiFi signals, the subtraction should be performed on a symbol-by-symbol basis. As a result, we propose a customized phase tracking scheme containing a pre-filter that calibrates CPE and tag-data phase rotations together. After our pre-filter, a standard WiFi phase tracking can successfully recover all the ambient data. Note that the groundtruth

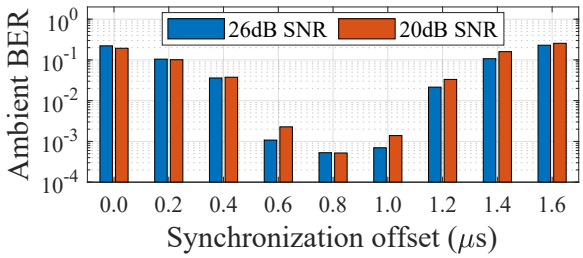

Figure 8: Impact of synchronization offsets on BERs. Results of demodulating ambient data show that the ambient-data BER can be a good indicator of synchronization qualities.

CPE will not be deducted twice because the estimated CPE in a standard phase tracking would be near-zero after the pre-filter.

3.4 Subsymbol-Level Synchronization

CAB requires the tag to accurately synchronize with ambient signals on the subsymbol level, ensuring zero-subcarriers and data-subcarriers can be successfully recovered on the receiver. Traditionally, synchronization with WiFi signals is not a problem for a WiFi receiver because there are a number of purpose-built training fields [1–3, 5]. For example, L-STF, which consists of ten short repetitions (0.8 μs), can be autocorrelated for initial timing synchronization [47]. Legacy long training field (L-LTF) contains two long and same repetitions, of which each has 64 samples (3.2 μs). Cross-correlation using this field can make fine timing synchronization [47]. At present, however, tags are unable to decode those fields and thus render them unusable. As a result, prior systems approach this in different ways. Hitchhike [58], FreeRider [59], and their variants [61, 62] adopt energy detection, which is a coarse synchronization and does not work on the OFDM-subcarrier level. TScatter [40] introduces pseudo-noise (PN) sequences for synchronization on the sample level. Yet, it incurs non-negligible decoding overhead. More importantly, it does not work with our requirement: demodulating tag data using only backscattered signals. SyncScatter [22] designs a hybrid of low-bandwidth and high-bandwidth detectors, which requires one of the most power-hungry components, a low-noise amplifier (LNA). Unlike those solutions that focus only on tags, we propose a joint timing synchronization scheme that uses the receiver feedback. Thanks to accurate ambient-data demodulation in §3.3, we observe the demodulated ambient-data BER is an excellent indicator for synchronization errors, as shown in Figure 8. Based on this observation, we design joint timing synchronization as follows.

Tag side. The tag uses an energy detector to discover a coarse start for the ambient WiFi packet. For fine-grained synchronization, our synchronization starts an interval search. The search starting point is τ seconds after the coarse start, where τ is the time length of everything before the physical layer conformance procedure (PLCP) service data unit (PSDU) and the medium access control (MAC) header. The step size is $\frac{1}{2}$GI, as the length of GI can only achieve symbol-level synchronization. The search window size is 1 symbol long. This size is enough because the error of our coarse start is about 2 μs, which is the same as prior systems [58, 61]. Based on

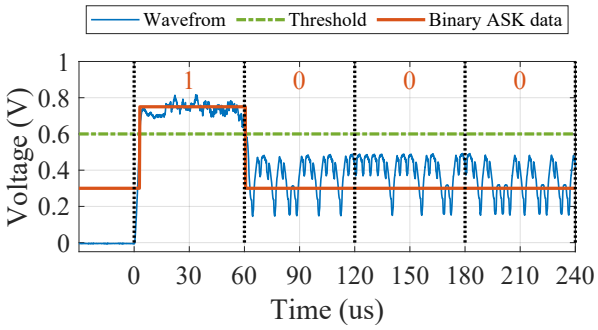

Figure 9: RX-tag communication. Tag can demodulate ASK-like signals from the receiver.

this search pattern, the tag starts modulation and waits for the receiver response for each searching point. If the message from the receiver is 'S-ACK', the tag would know the current search point is accurate enough. Otherwise, the tag moves on to the next search point. Our search process can be multi-resolution, where the step sizes of two layers are $\frac{1}{2}$GI and $\frac{1}{20}$GI.

Receiver side. Upon receiving a backscattered packet, the receiver demodulates ambient data as described in §3.3. If the receiver can measure the BER of this packet, we can set a threshold based on the channel quality and the requirement for synchronization accuracy to determine whether to send an 'S-ACK' to the tag. Unfortunately, in our design, the receiver is entirely detached from excitors, which means there is no way to obtain the groundtruth bits for ambient data and then compute the BER. To get around this, we make use of the CRC field of the backscattered packet. Although the CRC field is originally designed to detect accidental changes to raw data, it is also a good indicator for the quality of wireless demodulation [52]. Therefore, if the backscattered packet can pass CRC verification, the receiver sends an 'S-ACK' to the tag. Otherwise, it sends an 'F-ACK', which means a CRC failure.

RX-tag communication. For RX-tag communication, we employ the observation from Interscatter [31] that OFDM symbols can be transformed into ASK-modulated signals. Specifically, we employ 15 random OFDM symbols to denote a bit '1' and 15 constant OFDM symbols to denote a bit '0'. Only a passive envelope detector and a low-pass filter are adequate for the tag to demodulate these WiFi-emulated ASK signals [14]. Figure 9 shows how a low-power tag does this job. A comparator can separate the low-energy part and the high-energy part of the waveform from the filter, yielding binary signals.

From the above, we can see that the distinct feature of our synchronization scheme is the shifted computation burden from the tag to the receiver. Such an arrangement simplifies tag designs and lowers power consumption. Meanwhile, it adds negligible overhead to the receiver because CRC verification is a built-in function.

One may think that a single-pilot solution that uses a pilot to carry tag data could also do the job. It is simpler but no better than ours for at least three reasons. First, a single pilot is not robust as it experiences different channel fading at different locations as shown in §3.2 while our zero-subcarriers combine the strengths from all pilots. Second, it is hard to fix pilot positions for different OFDM-WiFi signals. For example, pilot positions for a 20 MHz transmission are -21,-7,7,21, but they become -53,-25,-11,11,25,53 for a 40 MHz transmission. In contrast, zero-subcarriers are always the center of pilots for OFDM-WiFi signals.

4 IMPLEMENTATION

Receiver prototype. The WiFi receiver is prototyped using ZedBoard ZYNQ-7000 [16] and AD-FMCOMMS3 [8], which can support 2x2 multi-input multi-output (MIMO) for a 40-MHz band. We implement 802.11g as WiFi 3, 802.11n as WiFi 4, 802.11ac as WiFi 5, and orthogonal frequency-division multiple access (OFDMA) 802.11ax as WiFi 6 using GNU radio [9]. All the baseband signal processing algorithms, e.g., tag and ambient data demodulation, are realized on-board in real-time.

Tag prototype for verification. We first build a functional verification prototype. As shown in Figure 10a, it mainly consists of a Xilinx ZYNQ-7010 FPGA, a 25 MHz crystal oscillator, a modulator, and an energy detector. For high-order PSK modulation, the modulator connects five switches with lumped terminations as a 16-to-1 multiplexer. We modulate the circuit impedance between 16 impedance states according to mapped raw bits. Each impedance state is chosen by adjusting its reflection coefficient $\Gamma = \frac{Z_a - Z_b}{Z_a + Z_b}$, where Z_a is the antenna impedance and Z_b is the complex impedance of backscatter circuits. This method can achieve significantly lower power consumption than bias-currents-based solutions [51] with almost power consumption of 80 mW.

In addition, single-sideband modulation can be achieved in this design as well. For example, assume that there are four reflection states obtained from properly matched networks: $\Gamma_1 = -\frac{1}{2} - j\frac{1}{2}$, $\Gamma_2 = \frac{1}{2} - j\frac{1}{2}$, $\Gamma_3 = \frac{1}{2} + j\frac{1}{2}$, $\Gamma_4 = -\frac{1}{2} + j\frac{1}{2}$. If we need to encode a bit '0', which means frequency shift without phase change. We toggle the RF-switch at frequency Δf_s in each period as (Γ_1,Γ_2,Γ_3,Γ_4). When we need to do frequency shift and 180-degree phase change for a bit '1', the RF-switch changes the states in each period as (Γ_3, Γ_4,Γ_1,Γ_2).

Tag prototype for applications. We also make a prototype using low-power devices, as shown in Figure 10b. It is a battery-free backscatter camera for live streaming using uncontrolled ambient excitations. This prototype consists of a tag for transmitting data, an OV7670 camera for capturing images, a 2-in-1 (RF&light) harvester using a TI BQ25570, and a solar cell for harvesting indoor ambient light. The harvested energy is stored in a 0.1 F supercapacitor. In particular, the battery-free tag is mainly composed of low-power components, namely an ADI LTC6930 oscillator, a Microchip Igloo Nano AGLN250 FPGA, RF switches, a passive rectifier, and an NCS2200 comparator. This prototype aims to investigate how CAB can benefit battery-free live-streaming in real-world applications.

Tag prototype with IC. As shown in Table 1, the overall power consumption of the battery-free prototype is 8.45 mW, which is 34× lower than that of the verification prototype. Although such a prototype can support battery-free camera sensing, it still does not release CAB's potential for ultra-low power consumption. As a result, we simulate a prototype using Cadence IC6.17 Virtuoso software and TSMC 0.18μm CMOS process design kits.

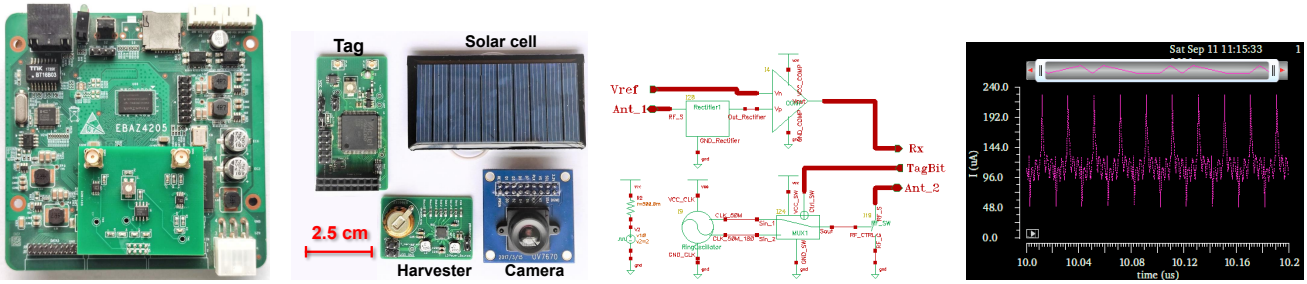

(a) Verification prototype (b) Battery-free prototype (c) IC simulated prototype (d) Currents of IC prototype.

Figure 10: We first build a functional verification prototype using FPGAs. Then we make a battery-free prototype using low-power devices for live-streaming applications. We also simulate an IC prototype to study its power consumption limits.

Table 1: Power consumptions of three prototypes.

	power consumption breakdown				Total
	Digital core	Oscillator	RF-switch	Detector	
Verification	209.00 mW (100%)	18.15 mW (100%)	1.83 mW (100%)	78.00 mW (100%)	307.03 mW (100%)
Battery-free	5.85 mW (2.80%)	0.59 mW (3.24%)	1.83 mW (100%)	0.18 mW (2.3%)	8.45 mW (2.75%)
IC	39.00 μW (0.02%)	196.00 μW (1.08%)	2.41 μW (0.13%)	33.79 μW (0.04%)	271.20 μW (0.09%)

- **Front-end**. Our RF front-end is composed of RF switches and a detector that contains a rectifier and a comparator. We use diodes generated using pMOS transistors to build a two-stage passive rectifier circuit[21]. An open-loop comparator is used for its high reaction speed to reduce synchronization error. Together, the detector consumes 33.79 μW. RF switches are composed of nMOS transistors and are connected as a multiplexer so that the digital core controls them for modulation. The measured power consumption for RF switches is 2.41 μW.
- **Oscillator**. We connect an odd number of inverters in series to compose a ring oscillator, which generates a 50 MHz clock for frequency shifting. Different from phase-shifted clocks that generate different time delays using inverters [60], we only need to toggle between 16 impedance states, which has fewer requirements for clocks. Hence, the measured power consumption for the oscillator is 196 μW.
- **Digital core.** As mentioned above, the digital core manipulates the multiplexer to modulating data at 250 kHz instead of running at 50 MHz. Such a low rate brings about significant power savings. The simulated baseband consumes 39 μW, which is 150× lower than that of the battery-free prototype.

Overall, the power consumption is 271.20 μW, which is 1133× and 31× lower than the verification and battery-free prototypes, respectively. The average current is about 116.10 μA, as shown in Figure 10d. Most of the power is dissipated by the oscillator, so further optimizations can be made through advanced IC techniques [33, 36].

5 EVALUATION

We first use the functional verification prototype to evaluate CAB's end-to-end performance and then investigate the contributions of various individual algorithms. At last, we show how the battery-free prototype can realize the ambient backscatter vision with real-world ambient traffic.

5.1 End-to-End Performance

As CAB is designed to work with all kinds of OFDM-WiFi signals, we would like to examine its overall performance with a variety of WiFi traffic, including multi-band, multi-stream, and multi-user WiFi.

WiFi 3. First, we examine how CAB performs with a typical single-stream OFDM-WiFi. The excitations are transmitted at -10 dBm in a 20 MHz band where the GI is 0.8 μs. From the results in Figure 11, we have two critical observations.

- The tag-data demodulation is near perfect where the maximal tag-data throughputs are 0.25 Mbps, 0.5 Mbps, and 0.99 Mbps for BPSK, QPSK, and 16PSK tag modulations. Those stable performances also show that the tag-data and ambient-data demodulations are independent.
- Rich ambient data are recovered with extremely high accuracy. Specifically, the maximal ambient-data throughputs are 6 Mbps, 18 Mbps, 34.96 Mbps, and 53.94 Mbps for modulation coding scheme (MCS) 13 7 11 3, which reach 100% ($\frac{6}{6}$), 100% ($\frac{18}{18}$), 97.1% ($\frac{34.96}{36}$), and 99.8% ($\frac{53.94}{54}$) of their MCS capacities.

WiFi 4. Next, we test how CAB works with multi-band signals of WiFi 4. We keep tag-data modulation at 16PSK and shorten the GI to 0.4 μs. As shown in Figure 11, the ambient-data throughputs of a 40 MHz band are nearly twice those with a 20 MHz band. This indicates that CAB is fully compatible with WiFi's high-throughput designs, e.g., multi-band, and works almost the same as a standard WiFi receiver. In addition, due to the short GI introduced, the maximal throughput for tag-data boosts up to 1.1 Mbps.

	MCS index				Excitation modulation	WiFi 3			WiFi 4		WiFi 5		WiFi 6	
	Non-HT[1]	HT[2]	VHT[3]	HE[4]		T/BPSK[5]	T/QPSK	T/16-PSK	20MHz	40MHz	Single stream	Double streams	1 user with 484 tones	2 users, 242 tones each
Ambient-data	13	0	0	0	BPSK	6.00[6]	6.00	6.00	7.20	14.96	14.99	30.00	34.40	34.35
	7	2	2	2	QPSK	18.00	18.00	17.99	21.70	44.42	44.91	89.58	102.20	102.11
	11	4	4	4	16QAM	34.93	34.96	34.94	43.30	89.00	89.70	179.83	204.04	205.96
	3	7	7	7	64QAM	53.94	53.90	53.75	72.06	145.78	148.63	294.70	340.61	327.66
Tag-data	13	0	0	0	BPSK	0.25	0.50	0.98	1.08	1.09	1.08	1.07	0.29	0.21
	7	2	2	2	QPSK	0.25	0.50	0.98	1.08	1.09	1.07	1.06	0.29	0.22
	11	4	4	4	16QAM	0.25	0.50	0.99	1.10	1.09	1.08	1.10	0.29	0.22
	3	7	7	7	64QAM	0.25	0.50	0.98	1.10	1.10	1.09	1.09	0.29	0.22

[1] Non-high-throughput transmission (WiFi 3). [2] High-throughput transmission (WiFi 4). [3] Very-high-throughput transmission (WiFi 5).
[4] High-efficiency transmission (WiFi 6). [5] BPSK for tag modulation. [6] Throughput (Mbps)

Figure 11: Comprehensive evaluation of CAB. With all kinds of OFDM-WiFi excitations, including WiFi 3/4/5/6, CAB exhibits full compatibility. Meanwhile, both tag and ambient data throughputs are consistent and stable across different modulation schemes. More importantly, the highest aggregate throughput is 340.9 Mbps with WiFi 6, reaching 97% Shannon capacity.

WiFi 5. By keeping all the high-throughput parameters from WiFi 4 experiments (40-MHz band and short GI), we continue to push CAB to work with multi-stream signals of WiFi 5. While the results of single-stream traffic are pretty much the same as WiFi 4, the maximal ambient-data throughput increases to 294.70 Mbps. Note that such ambient-data throughputs are achieved by demodulating only on backscattered signals, whereas prior systems require a dedicated receiver to obtain those ambient data [40, 59, 61].

WiFi 6. To push the envelope of ambient WiFi backscatter, we want to see whether CAB can take over the last castle, WiFi 6. Not surprisingly, CAB works smoothly with WiFi 6 excitations as shown in Figure 11. Specifically, the maximal throughput of ambient data grows to 340.61 Mbps. Such a throughput gain is owing to that the shortest equivalent GI of WiFi 6 is only 0.2 μs ($\frac{0.8}{4}$) [3], which allows more ambient data transmissions. In addition, CAB is compatible with multi-user OFDMA (MU-OFDMA) excitations of 2 users, each taking 242-tone resource units (RUs), where the maximal tag-data throughput is 0.22 Mbps. The main reason for less tag-data throughputs is that the length of our WiFi 6 symbol is 13.6 μs, much larger than the traditional one. Besides, the highest achieved aggregate throughput is 340.9 Mbps, where the excitation is 64QAM WiFi 6, and the tag modulates using 16PSK. It reaches 97% ($\frac{340.9}{351.38}$) Shannon capacity [4].

The above evaluation presents a comprehensive study of CAB with all kinds of OFDM-WiFi excitations, demonstrating its universality, high throughput, and modulation independence.

Competitions. We compare CAB against FS-backscatter [60], FreeRider [59], and MOXcatter [61]. While CAB and FS-backscatter are content-agnostic systems, MOXcatter and FreeRider can also demodulate only from backscattered signals if the ambient data is prearranged. We deliver all-zero packets using a single spatial stream on 20 MHz band of MCS 0 for tag-data measurements, and send random data

[3] A WiFi symbol (excluding GI) can be made as long as 12.8 μs, which is four times the traditional length.
[4] Given a channel of 40 MHz with SNR of 20 dB, the Shannon capacity for a single stream is $40 * \log_2(1 + 20) = 175.69$ Mbps.

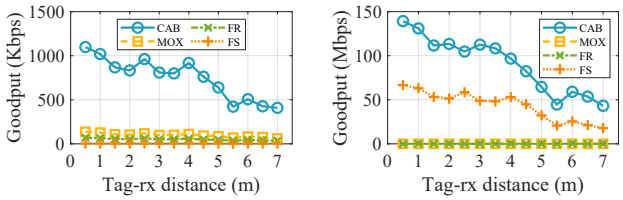

(a) Tag goodput (b) Ambient goodput

Figure 12: Goodput comparison of CAB and other advanced systems.

with two spatial streams on 20 MHz of MCS 7 to measure the ambient goodput. We depict the results in Figures 12a and 12b. From this experiment, we observe that for tag-data goodput, CAB consistently outmatches other systems. In particular, when the tag-rx distance is 7 m, the goodput of CAB is 408.56 kbps, which is 6.7×, 12.5× and 207.4× better than those of MOXcatter, FreeRider, and FS-backscatter respectively. Such goodput gains climb to 8.1×, 16.0× and 269.7× at a tag-rx distance of 1 m. This is primarily because our baselines modulate on a multi-symbol or packet level while CAB backscatters on a single-symbol level. The other factor is that CAB supports fine-grained PSK modulation, whereas our baselines only employ two-level AM modulation.

For ambient-data goodput, CAB outperforms again. For instance, when the tag-rx distance is 6.5 m, CAB's goodput is 53.59 Mbps, which is 2.5× better than FS-backscatter's. The main reason is that FS-backscatter has to lower the packet RSSI when it needs to modulate a bit '0', which inevitably reduces SNRs and demodulation quality, and FreeRider and MOXcatter have zero ambient goodput in our random-data setup.

5.2 Micro Benchmarks

Next, we examine contributions from individual modules.

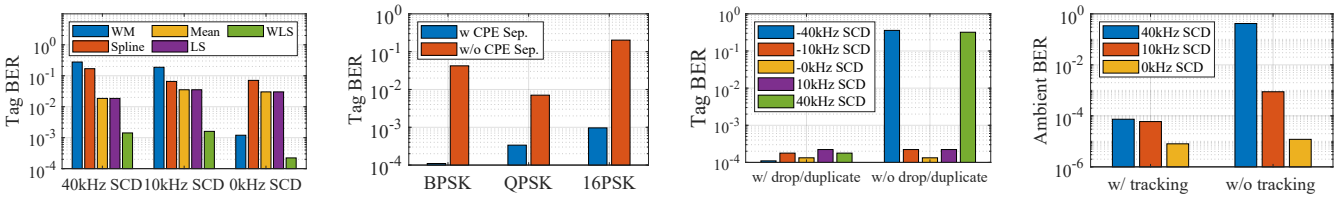

(a) Phase estimation (b) CPE separation (c) Symbol assembly (d) Phase tracking

Figure 13: Our zero-subcarrier phase estimation, CPE separation, symbol assembly, and customized phase tracking contribute significant BER improvements for demodulating tag and ambient data.

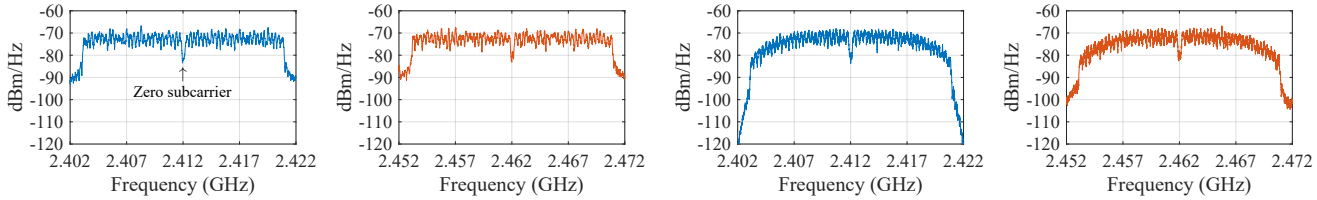

(a) Ambient signal (b) Backscatter signal (c) Ambient signal (d) Backscatter signal

Figure 14: Power spectral density (PSD) of ambient and backscattered 802.11n signals. (a) and (b) are simulations, while (c) and (d) are real-world measurements.

Estimating zero-subcarrier phases. First, we study how different phase-estimation schemes work with tag-data demodulation. Besides WLS, we include mean (simple averaging), weighted mean, LS (least square), cubic spline interpolation for estimating zero-subcarrier phases, and choose QPSK for tag modulation. We program sampling-clock differences (SCD)[5] between the excitor and the receiver at 0, 10 kHz, and 40 kHz. Results in Figure 13a show that WLS is the best among all and exhibits robustness at different SCD conditions. Specifically, the tag-data BERs of WLS are always below 0.16% for different SCDs. In addition, when there is no SCD, the tag-data BER of WLS is 0.022%, which is 135.7×, 5.4×, 135.7×, 319.7× lower than mean, weight mean, LS, and spline, respectively.

Impact of CPE separation. Next, we plan to examine the impact of CPE separation on tag-data demodulation. The tag modulation adopts BPSK, QPSK, and 16PSK, and results are depicted in 13b. We have two observations. First, our CPE separation significantly improves tag-data demodulation accuracy. In particular, with BPSK backscattering, the tag-data BER with CPE separation is 0.011%, and it increases to 4.2% when there is no CPE separation, which achieves 383× performance gain. Second, our CPE separation degrades gradually with the increasing order of PSK, which is similar to high-order modulation in active radios [52]. The tag-data BER for 16PSK is still less than 0.97%, demonstrating the effectiveness of our iterative CPE separation.

Impact of symbol assembly. Furthermore, we want to check how our symbol assembly works with various clock differences. The programmed SCDs include -40 kHz, -10 kHz, 0, 10 kHz, and 40 kHz. Figure 13c demonstrates that our accurate symbol assembly considerably reduces tag-data BERs. For example, when SCD is -40 kHz, the tag-data BER without drop/duplicate is 36%, then it plummets to 0.011% after introducing our symbol assembly. Similar trends are observed for other SCDs. Meanwhile, the tag-data BERs are always under 0.024%, showing our symbol assembly is consistent and reliable.

Customized phase tracking. Also, we conduct experiments to validate our proposed phase tracking for ambient-data demodulation. Results in Figure 13d manifest the ambient-data BERs decrease substantially after customized phase tracking. For instance, when the SCD is 40 kHz, the ambient-data BER is 42% before phase tracking. Later it sharply lowers to 0.01% after phase tracking. Such an improvement is because our phase tracking is processed on a symbol-by-symbol basis, avoiding incorrect pilot averaging and phase unwrapping across symbols. Moreover, the ambient-data BERs are all less than 0.01% for various SCDs. Such high-quality demodulation facilitates our synchronization design substantially.

Interference at DC. One may be concerned that using zero-subcarriers is improper since typical WiFi uses simple yet cheap direct-conversion RF receivers that have strong interference at DC and thus significantly spoil the tag-data accuracy in the backscatter system. However, the zero-subcarriers used in our modulation are purely virtual, which means the DC subcarrier after backscattering is still NULL. Figure 14 manifests the power spectral density (PSD) of 802.11n signals before and after backscattering in simulations and real-world measurements. The power of zero-subcarriers can be demonstrated by the PSD value at the center frequency, which has a negligible difference before and after backscattering. In detail, they are -83.5

[5]SCD is the difference of programmed values for clocks, which differs from the groundtruth clock difference, SFO.

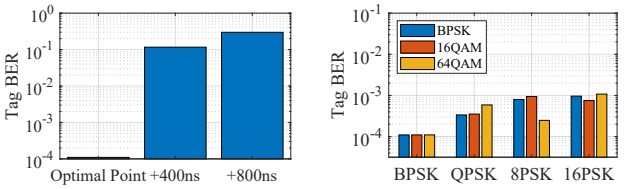

(a) Joint sync. (b) High-order PSK

Figure 15: Our tag-data BERs keep low with accurate synchronization and high-order PSK modulation.

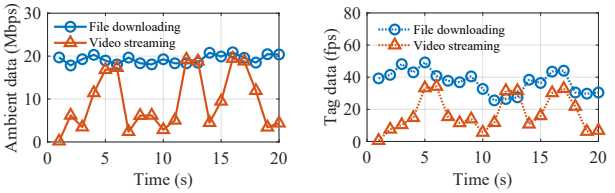

(a) Ambient traffic (b) Tag streaming data

Figure 16: Tag streaming with uncontrolled traffic.

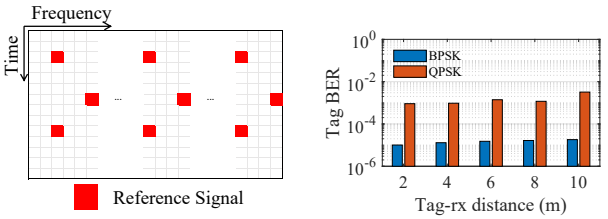

(a) LTE slot structure (b) CAB with ambient LTE

Figure 17: CAB extension with ambient LTE.

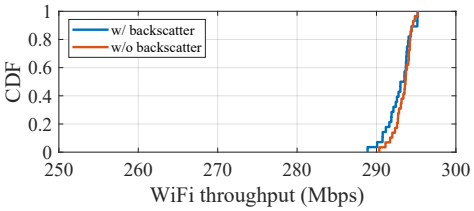

Figure 18: Co-existence with ambient WiFi.

and -83.1 dBm/Hz in simulation, and -83.7 and -82.6 dBm/Hz in real-world measurements.

Joint synchronization. Additionally, we check whether accurate joint synchronization helps reduce tag-data BERs. In Figure 15a, we show tag-data BERs at different search points in an interval. We can see that the achieved lowest BER happens at the optimal synchronization point, which confirms that the receiver feedback correctly guides the search process. When the point is offset by 800 *ns*, its BER rises rapidly to 30% from 0.01% at the optimal point.

High-order PSK. Finally, we plan to look into the compatibility of our high-order PSK with varying modulation schemes on excitations. In this experiment, the modulation schemes of excitations include BPSK, 16QAM, and 64QAM, and those of tag modulation are BPSK, QPSK, and 16PSK. Results in Figure 15b reveal that our tag modulation works well with all different excitations. Such pervasiveness is as expected because our key insight of zero-subcarriers is to view every OFDM-WiFi packet a virtual single tone. It reconfirms that our tag modulation is independent of excitation modulation, so are tag and ambient demodulation.

Ambient traffic for tag streaming. Next, we show how CAB enables battery-free live-streaming with uncontrolled ambient traffic. We let a Dell laptop connect to an AP through WiFi 4 and generate normal WiFi traffic in various ways. We first download a 10 GB file, then generate video streaming traffic by playing a 1080p, 30 fps blue-ray movie on an online video site. The traffic lasts while the file is downloading or the video is playing, but the rate may vary significantly from time to time. Our battery-free prototype takes those traffic as carriers, and backscatters image data from the OV7670 camera at 240p (426*240). Upon receiving backscattered signals, the receiver demodulates these backscattered live streams in real-time. We adopt the intra-frame algorithm in [43] and achieve 69× compression ratio on average. Another average compression ratio of 38 × can be realized by inter-frame compression [43] while maintaining a peak signal-to-noise ratio (PSNR) of more than 30 dB.

The streaming results are plotted in Figure 16. We can see that the ambient WiFi traffic of file downloading is quite stable, around 20 Mbps, whereas that of video streaming has ups (~ 19 Mbps) and downs (~ 2 Mbps). This is because cache servers for improving quality of experience (QoE) are standard nowadays for most http-streaming platforms [13, 15]. Also, we observe that the tag-streaming traffic is similar to excitation traffic. This manifests the characteristic of ambient backscatter: CAB can exploit it for backscatter communication whenever there is usable ambient traffic. For instance, with file-downloading excitations, the average achieved tag-streaming rate is 37.1 fps (approximately 33.4 kbps), showing CAB's decent capability for battery-free live streaming. As this demo is general and ambient WiFi is pervasive, it opens the door for high-throughput battery-free sensing by leveraging uncontrolled ambient traffic. Future tag-streaming rates could be further improved by several means, e.g., denser tag modulation [52], advanced compression algorithms [56], and more ambient traffic types [30].

CAB with ambient LTE. Finally, we demonstrate that CAB's principle can be generalized to other OFDM ambient traffic by extending it to ambient LTE [10]. For ambient LTE, we use downlinks from cell towers to user equipments (UEs) as excitations. Like pilots in WiFi, LTE signals include predefined reference signals, from which we create virtual zero-subcarriers to convey tag data. The LTE eNodeB and UEs are implemented using LTE SDR OpenAirInterface [11] and srsLTE [12]. From the results in Figure 17b, we observe that tag data can be transmitted reliably via ambient LTE. For example, the tag-data BERs are below 0.002% with BPSK modulation across different tag-rx distances. As ambient LTE signals are ubiquitous

worldwide, it shows CAB's readiness for a range of novel sensing paradigms for LTE devices.

Co-existence with ambient WiFi. Next, we plan to look into whether the CAB can co-exist with WiFi networks. We transmitted 300 Mbps WiFi 4 signals on channel 1, then put the CAB tag 1 m away from the WiFi receiver. With and without backscatter on the tag, we measure the throughput of ambient WiFi. Results in Figure 18 reveal that the CAB tag has a negligible impact on ambient WiFi traffic. In more detail, the median throughput with and without backscatter is 293.53 and 293.65 Mbps, respectively.

6 RELATED WORK

Backscatter communication has the cutting-edge of low-power transmission at backscatter nodes [23–29, 38, 64]. Ambient backscatter systems further employ standard radios in backscatter transmission. Ambient backscatter systems can be broadly classified into three categories based on the types of excitations.

Content-agnostic. The seminal work, ambient backscatter [39], proposes to leverage TV signals in the air to realize battery-free backscatter communication. WiFi backscatter [34] builds on top of this idea and enables the first general-purpose WiFi backscatter for Internet connectivity. FS-backscatter [60] observes that frequency shifting is the key to improve demodulation accuracy. Although those systems can demodulate tag data without knowing ambient data, they suffer low throughput due to packet-level modulation, up to thousands of bps. As opposed to those prior works, CAB supports symbol-level tag modulation and agnostic demodulation, which achieves over 300 Mbps aggregate throughput and 97% Shannon capacity.

Content-aware. For uncontrolled ambient traffic, the pioneering work, Hitchhike [58] proposes to make uncontrolled ambient 802.11b signals as carriers and uses codeword translation for tag modulation and demodulation. It not only achieves decent throughput (~ 300 kbps) due to symbol-level modulation but also can work with off-the-shelf devices. A number of works receive inspiration from this, e.g., MOXcatter [61], X-Tandem [62], PLoRa [44], LScatter [20], TScatter [40]. BackFi [19] builds a 'WiFi reader' using full-duplex radios and achieves a maximal throughput of 6.67 Mbps. Yet, those systems are content-based because they require ambient data to decode tag data. In contrast, CAB is agnostic and can decode both ambient and tag data from backscattered signals alone, greatly boosting band efficiency and practicality of ambient backscatter.

CW-based. The systems of this category mainly use a helper device that can generate CW as carriers. Passive Wi-Fi provides a low-power tag design that backscatters CW into standard 802.11b signals [35]. LoRa backscatter extends the uplink range to 2.8 km for long-range backscatter communication using a dedicated single-tone generator [48]. Interscatter novelly uses reverse whitening techniques to turn a Bluetooth signal into a partial single tone and further backscatter it into ZigBee and 802.11b signals [31]. While those systems contribute to general-purpose backscatter communication in different ways, the requirement of single tones as excitations poses challenges on the pathways to widespread deployment.

CAB builds on all these works and proposes the first content-agnostic ambient backscatter that delivers near Shannon-capacity throughput.

7 DISCUSSION AND FUTURE WORK

COTS receiver. The factor limiting commodity off-the-shelf (COTS) devices as receivers is that current WiFi network interface controllers (NICs) do not provide access to pilot subcarriers on the PHY layer. Although CAB relies on PHY data, we believe it is promising to implement CAB in commodity NICs in the future since more and more APIs for PHY data provided by NICs enable various functions and applications. For example, Linux CSI Tool [7] supplies a convenient way to obtain CSI from the modified firmware, the output has selective 30 data subcarriers. Besides, CTE fields in Bluetooth 5 direction-finding packets, for example, convey IQ data to the receiver host, allowing it to analyze the channel and locate the transmitter [6].

Ambient traffic patterns. Although we design CAB to work with various ambient traffic patterns, its ambient-data demodulation would degrade for sporadic WiFi traffic because the synchronization accuracy would be affected when the intervals between packets are not stable. Yet, it barely affects tag-data demodulation. Future work includes designing proper low-power WiFi demodulation schemes to obtain S-LTF and L-LTF fields.

For low-power IoTs to reach widespread deployment, ambient backscatter is one of the most promising solutions. CAB's main contribution is to provide a novel way to exploit all the uncontrolled OFDM-WiFi signals as virtual single tones. Based on this insight, we have built several prototypes to show that CAB is universal with WiFi 3/4/5/6, and its throughput reaches near Shannon capacity. We believe that CAB takes a curial step forward on ubiquitous battery-free IoTs and paves the way for fast and wide adoption of general-purpose backscatter communication.

ACKNOWLEDGMENTS

We thank the shepherd Kate Ching-Ju Lin and the anonymous reviewers for their helpful comments. This work was supported by NSFC Grant No. 61932017 and 61971390.

REFERENCES

[1] [n. d.]. 802.11a. https://standards.ieee.org/standard/802_11a-1999.html. ([n. d.]).
[2] [n. d.]. 802.11ac. https://standards.ieee.org/standard/802_11ac-2013.html. ([n. d.]).
[3] [n. d.]. 802.11ax. https://standards.ieee.org/standard/802_11ax-2021.html. ([n. d.]).
[4] [n. d.]. 802.11g. https://standards.ieee.org/standard/802_11g-2003.html. ([n. d.]).
[5] [n. d.]. 802.11n. https://standards.ieee.org/standard/802_11n-2009.html. ([n. d.]).
[6] [n. d.]. Bluetooth. https://www.bluetooth.com/specifications/bluetooth-core-specification/. ([n. d.]).
[7] [n. d.]. CSI Tool. https://dhalperi.github.io/linux-80211n-csitool/. ([n. d.]).
[8] [n. d.]. FMCOMMS3. https://wiki.analog.com/resources/eval/user-guides/ad-fmcomms3-ebz. ([n. d.]).

[9] [n. d.]. gnuradio. https://www.gnuradio.org/. ([n. d.]).
[10] [n. d.]. LTE specification. https://www.3gpp.org/IMG/pdf/2009_10_3gpp_IMT.pdf. ([n. d.]).
[11] [n. d.]. OpenAirInterface. https://gitlab.eurecom.fr/oai/openairinterface5g/. ([n. d.]).
[12] [n. d.]. srsLTE. https://github.com/srsran/srsran. ([n. d.]).
[13] [n. d.]. Twitch. https://www.twitch.tv. ([n. d.]).
[14] [n. d.]. WISP 5.0. https://github.com/wisp/wisp5. ([n. d.]).
[15] [n. d.]. YouTube. https://www.youtube.com. ([n. d.]).
[16] [n. d.]. zedboard. https://digilent.com/reference/_media/zedboard:zedboard_ug.pdf. ([n. d.]).
[17] 2017. EPC C1G2 Standard. http://www.gs1.org/epcrfid/epc-rfid-uhf-air-interface-protocol/2-0-1. (2017).
[18] M. R. Abdelhamid, R. Chen, J. Cho, A. P. Chandrakasan, and F. Adib. 2020. Self-reconfigurable micro-implants for cross-tissue wireless and batteryless connectivity. In *Proc. of ACM MobiCom*.
[19] D. Bharadia, K. Joshi, M. Kotaru, and S. Katti. 2015. Backfi: High throughput wifi backscatter. In *Proc. of ACM SIGCOMM*.
[20] Z. Chi, X. Liu, W. Wang, Y. Yao, and T. Zhu. 2020. Leveraging ambient lte traffic for ubiquitous passive communication. In *Proc. of ACM SIGCOMM*.
[21] Daniel Dobkin. 2012. *The rf in RFID: uhf RFID in practice*. Newnes.
[22] M. Dunna, M. Meng, P.H. Wang, C. Zhang, P.P. Mercier, and D. Bharadia. 2021. SyncScatter: Enabling WiFi like synchronization and range for WiFi backscatter Communication.. In *Proc. of USENIX NSDI*.
[23] W. Gong, S. Chen, and J. Liu. 2017. Towards higher throughput rate adaptation for backscatter networks. In *Proc. of IEEE ICNP*.
[24] W. Gong, S. Chen, J. Liu, and Z. Wang. 2018. Mobirate: Mobility-aware rate adaptation using phy information for backscatter networks. In *Proc. of IEEE INFOCOM*.
[25] W. Gong, H. Liu, J. Liu, X. Fan, K. Liu, Q. Ma, and X. Ji. 2018. Channel-aware rate adaptation for backscatter networks. *IEEE/ACM Transactions on Networking* 26, 2 (2018), 751–764.
[26] W. Gong, H. Liu, K. Liu, Q. Ma, and Y. Liu. 2016. Exploiting channel diversity for rate adaptation in backscatter communication networks. In *Proc. of IEEE INFOCOM*.
[27] W. Gong, J. Liu, and Z. Yang. 2016. Fast and reliable unknown tag detection in large-scale RFID systems. In *Proc. of ACM MobiHoc*.
[28] W. Gong, J. Liu, and Z. Yang. 2017. Efficient unknown tag detection in large-scale RFID systems with unreliable channels. *IEEE/ACM Transactions on Networking* 25, 4 (2017), 2528–2539.
[29] W. Gong, I. Stojmenovic, A. Nayak, K. Liu, and H Liu. 2015. Fast and scalable counterfeits estimation for large-scale RFID systems. *IEEE/ACM Transactions on Networking* 24, 2 (2015), 1052–1064.
[30] W. Gong, L. Yuan, Q. Wang, and J. Zhao. 2020. Multiprotocol backscatter for personal IoT sensors. In *Proc. of ACM CONEXT*.
[31] V. Iyer, V. Talla, B. Kellogg, S. Gollakota, and J. Smith. 2016. Inter-technology backscatter: Towards internet connectivity for implanted devices. In *Proc. of ACM SIGCOMM*.
[32] J. Jang and F. Adib. 2019. Underwater Backscatter Networking. In *Proc. of ACM SIGCOMM*.
[33] P. Kamalinejad, K. Keikhosravy, R. Molavi, S. Mirabbasi, and Vcm Leung. 2014. An ultra-low-power CMOS voltage-controlled ring oscillator for passive RFID tags. In *Proc. of IEEE NEWCAS*.
[34] B. Kellogg, A. Parks, S. Gollakota, J. R. Smith, and D. Wetherall. 2014. Wi-Fi backscatter: Internet connectivity for RF-powered devices. In *Proc. of ACM SIGCOMM*.
[35] B. Kellogg, V. Talla, S. Gollakota, and J. R. Smith. 2016. Passive wi-fi: Bringing low power to wi-fi transmissions. In *Proc. of USENIX NSDI*.
[36] K. K. Lee, K. Granhaug, and N. Andersen. 2014. A study of low-power crystal oscillator design. In *Proc. of IEEE NORCHIP*.
[37] K. C.-J. Lin, S. Gollakota, and D. Katabi. 2011. Random access heterogeneous MIMO networks. (2011).
[38] H. Liu, W. Gong, L. Chen, W. He, K. Liu, and Y. Liu. 2014. Generic composite counting in RFID systems. In *Proc. of IEEE ICDCS*.
[39] V. Liu, A. Parks, V. Talla, S. Gollakota, D. Wetherall, and J. R. Smith. 2013. Ambient backscatter: wireless communication out of thin air. In *Proc. of ACM SIGCOMM*.
[40] X. Liu, Z. Chi, W. Wang, Y. Yao, P. Hao, and T. Zhu. 2021. Verification and Redesign of OFDM Backscatter. In *Proc. of USENIX NSDI*.
[41] Y. Ma, Z. Luo, C. Steiger, G. Traverso, and F. Adib. 2018. Enabling deep-tissue networking for miniature medical devices. In *Proc. of ACM SIGCOMM*.
[42] M. H. Mazaheri, A. Chen, and O. Abari. 2021. mmTag: a millimeter wave backscatter network. In *Proc. of ACM SIGCOMM*.
[43] S. Naderiparizi, M. Hessar, V. Talla, S. Gollakota, and J. R. Smith. 2018. Towards battery-free HD video streaming. In *Proc. of USENIX NSDI*.
[44] Y. Peng, L. Shangguan, Y. Hu, Y. Qian, X. Lin, X. Chen, D. Fang, and K. Jamieson. 2018. PLoRa: A passive long-range data network from ambient LoRa transmissions. In *Proc. of ACM SIGCOMM*.
[45] H. Rahul, H. Hassanieh, and D. Katabi. 2010. SourceSync: A distributed wireless architecture for exploiting sender diversity. In *Proc. of ACM SIGCOMM*.
[46] Hamid Shafiee, Behzad Nourani, and M Khoshgard. 2004. Estimation and compensation of frequency offset in DAC/ADC clocks in OFDM systems. In *Proc. of IEEE ICC*.
[47] Michael Speth, Stefan A Fechtel, Gunnar Fock, and Heinrich Meyr. 1999. Optimum receiver design for wireless broad-band systems using OFDM. I. *IEEE Transactions on Communications* 47, 11 (1999), 1668–1677.
[48] V. Talla, M. Hessar, B. Kellogg, A. Najafi, J. R. Smith, and S. Gollakota. 2017. Lora backscatter: Enabling the vision of ubiquitous connectivity. In *Proc. of ACM IMWUT*.
[49] V. Talla, B. Kellogg, B. Ransford, S. Naderiparizi, S. Gollakota, and J. R. Smith. 2015. Powering the next billion devices with Wi-Fi. In *Proc. of ACM CONEXT*.
[50] J. K. Tan. 2006. *An adaptive orthogonal frequency division multiplexing baseband modem for wideband wireless channels*. Ph.D. Dissertation.
[51] Stewart J Thomas, Eric Wheeler, Jochen Teizer, and Matthew S Reynolds. 2012. Quadrature amplitude modulated backscatter

in passive and semipassive UHF RFID systems. *IEEE Transactions on Microwave Theory and Techniques* 60, 4 (2012), 1175–1182.

[52] D. Tse and P. Viswanath. 2005. *Fundamentals of wireless communication*. Cambridge university press.

[53] D. Vasisht, S. Kumar, and D. Katabi. 2016. Decimeter-level localization with a single WiFi access point. In *Proc. of USENIX NSDI*.

[54] J. Wang, H. Hassanieh, D. Katabi, and P. Indyk. 2012. Efficient and reliable low-power backscatter networks. In *Proc. of ACM SIGCOMM*.

[55] Q. Wang, S. Chen, J. Zhao, and W. Gong. 2021. RapidRider: Efficient WiFi Backscatter with Uncontrolled Ambient Signals. In *Proc. of IEEE INFOCOM*.

[56] C.-Y. Wu, N. Singhal, and P. Krahenbuhl. 2018. Video compression through image interpolation. In *Proc. of ECCV RFID*.

[57] M. Zhang, S. Chen, J. Zhao, and W. Gong. 2021. Commodity-level BLE backscatter. In *Proc. of ACM MobiSys*.

[58] P. Zhang, D. Bharadia, K. Joshi, and S. Katti. 2016. Hitchhike: Practical backscatter using commodity wifi. In *Proc. of ACM SenSys*.

[59] P. Zhang, C. Josephson, D. Bharadia, and S. Katti. 2017. Freerider: Backscatter communication using commodity radios. In *Proc. of ACM CONEXT*.

[60] P. Zhang, M. Rostami, P. Hu, and D. Ganesan. 2016. Enabling practical backscatter communication for on-body sensors. In *Proc. of ACM SIGCOMM*.

[61] J. Zhao, W. Gong, and J. Liu. 2018. Spatial Stream Backscatter Using Commodity WiFi. In *Proc. of ACM MobiSys*.

[62] J. Zhao, W. Gong, and J. Liu. 2018. X-tandem: Towards multi-hop backscatter communication with commodity wifi. In *Proc. of ACM MobiCom*.

[63] J. Zhao, W. Gong, and J. Liu. 2020. Towards scalable backscatter sensor mesh with decodable relay and distributed excitation. In *Proc. of ACM MobiSys*.

[64] J. Zhao, W. Gong, and J. Liu. 2021. Microphone array backscatter: an application-driven design for lightweight spatial sound recording over the air. In *Proc. of ACM MobiCom*.