

Towards Higher Throughput Rate Adaptation for Backscatter Networks

Wei Gong, Si Chen, and Jiangchuan Liu
School of Computing Science, Simon Fraser University, Canada
{gongweig, sca228, jcliu}@sfu.ca

Abstract—Recently backscatter networks have received booming interest because, they offer a battery-free communication paradigm using propagation radio waves as opposed to active radios in traditional sensor networks while providing comparable sensing functionalities, ranging from light and temperature sensors to recent microphones and cameras. While sensing data on backscatter nodes has been seen on a clear path to increase in both volume and variety, backscatter communication is not well prepared and optimized for transferring such continuous and high-volume data. To bridge this gap, we propose a high-throughput rate adaptation scheme for backscatter networks by exploring the unique characteristics of backscatter links and the design space of the ISO 18000-6C (C1G2) protocol. Our key insight is that while prior work has left the downlink unattended, we observe that the quality of downlink is affected significantly by multipath fading and thus can degrade the uplink and overall throughput considerably. Therefore, we introduce a novel rate mapping algorithm that chooses the best rate for both the downlink and uplink. Also, we design an efficient channel estimation method fully compatible with the C1G2 protocol and a reliable probing trigger, substantially saving probing overhead. Our scheme is prototyped using a COTS RFID reader and tags. The results show that we achieve up to 2.5x throughput gain over state-of-the-art approaches across various mobility, channel, and network-size conditions.

I. INTRODUCTION

There is a long-standing vision of ultra-low power ubiquitous sensor networks where many tiny sensors are wirelessly connected and can perform continuous sensing tasks without human intervention, e.g., Smart Dust [1]. Backscatter networks are one of most promising candidates to realize this goal as backscatter nodes -like RFID tags- can capture power from propagation radio waves, making battery-free networks possible. Thanks to the advances of energy efficiency scaling for microelectromechanical systems, a wide range of applications that previously are only supported by battery-assisted sensors become available for backscatter networks, such as temperature or light intensity sensing [2], acoustic signal capturing [3], and even video surveillance [4]. While backscatter networks have seen the future of increasing sensing data coming in, backscatter communication that supports continuous and high-throughput transmission is not quite ready yet. Recently there have been several attempts that focus on revamping the traditional backscatter protocols for more efficient transmission [5], [6], [7]. Yet incompatibility with industry standards, e.g., ISO 18000-6C (C1G2) specification, and requirements of customized hardware hinder wide adoption of those proposals.

As such, we aim to design a high-throughput protocol that is fully compatible with C1G2 using Commercial Off-The-Shelf (COTS) devices, which can benefit tons of currently deployed backscatter devices. To achieve this, however, there are several key challenges:

- *Ineffective Rate Selection*: Prior work of rate selection for backscatter networks only focuses on the uplink that is for transmitting sensor data [8], [9], leaving the impact of downlink rates largely uninvestigated. Actually, the downlink is indispensable and implicitly involved in the uplink transmission because any uplink has a downlink as its *predecessor*, which means if the downlink fails due to incorrect rate settings, the uplink would be discontinued. This is the unique characteristic of the backscatter link that a downlink and an uplink are sequentially combined as a backscatter link. Therefore, if the downlink rate is left unattended, even the optimal setting for the uplink may not bring overall throughput gain.
- *Probing Overhead*: In backscatter networks, all transmissions are scheduled by the reader through an ALOHA-like MAC protocol because nodes cannot sense each other. The performance of channel probing would severely degrade due to MAC collisions when the node population increases [8]. Although CARA [9] proposes an estimation algorithm to compensate such collisions, the probing process still needs to follow the above MAC scheduling, prolonging the probing time. In addition, the probing trigger, which is necessary for deciding when to probe, could exacerbate the issue. For example, Blink [8] requires measurements of at least 10 channels for its trigger, and CARA needs to probe at least 5 channels.
- *Limited Visibility for Channel Estimation*: While it is common that PHY hints for channel estimation, e.g., bit error rate (BER), are not available for most of the COTS wireless devices, it becomes even worse when we deal with COTS readers; even the packet level loss rate is very difficult to obtain because COTS readers only report the number of successful reads in a time interval. Previous solutions either use an extra monitoring device, like USRP, to sniff messages transferred in the air, or log commands from the reader into tags' EPC memory using Computational RFIDs (CRFID). Yet these methods not only introduce more cost due to additional hardware but also are inapplicable to situations where only COTS devices are available.

To address the above issues, we propose a high-throughput Rate Adaptation framework for Backscatter networks, RAB. It is fast and efficient while being compatible with the C1G2 protocol and existing commercial RFID readers. To do so, it primarily makes three fundamental optimizations over the current standard. First, our work provides insights that both the uplink and downlink affect the overall throughput significantly, which motivates us to adapt rates for both in contrast to prior work that only focuses on the uplink [5], [8], [9]. Second, we describe a novel channel estimation method that uses filter-based probing to effectively reduce errors brought by MAC-layer collisions and estimates the loss rate by leveraging the link timing features of the C1G2 protocol. Third, we present a correlation-based channel hopping and an accurate mobility detection approach that uses PHY hints to determine when to trigger channel estimation, considerably saving channel-probing overhead.

We build a prototype of RAB using a Thingmagic reader and 20 Alien Higg3 tags. We compare RAB with Blink and CARA and results show that across 80 traces with different mobility, channel, and network-size conditions, RAB achieves overall throughput gains of 2.5x over Blink and 1.9x over CARA on average. This gain comes from two sources: First, RAB reduces probing cost significantly by 8.2x compared to Blink, and by 4.3x compared to CARA; Second, for data transmission, our rate selection scheme achieves throughput gains of 1.8x over Blink and 1.6x over CARA.

The rest of this paper is organized as follows. Section II presents the primer of backscatter communication. Section III gives an overview of our solution. Section IV details how we choose optimal rate that best fits the channel condition. Section V introduces our novel channel estimation that includes filter-based probing and loss rate estimation. Section VI describes the design of our probing triggers. The implementation and evaluations of the proposal are in Section VII and VIII. The related work is discussed and compared in Section IX. The concluding remarks and future work are in Section X.

II. BACKSCATTER PRIMER

Backscatter System. A backscatter system usually is composed of a reader and one or more backscatter nodes¹, e.g., RFID tags. The reader initiates the communication by transmitting carrier waves, which serves two purposes. First, the tag can capture energy from the radio waves and power itself for computation and communication. Second, the tag backscatters information bits by modulating the same carrier waves. While many of the principles are generally applicable to all RFID devices, here we focus on the UHF RFID devices whose behaviors are defined in the C1G2 protocol [10].

Backscatter Links. While the reader is usually assumed powerful, the tag is restricted in terms of computation, communication, and hardware capabilities since it can only capture limited power from radio waves. Therefore, the asymmetry exists almost everywhere in backscatter systems including

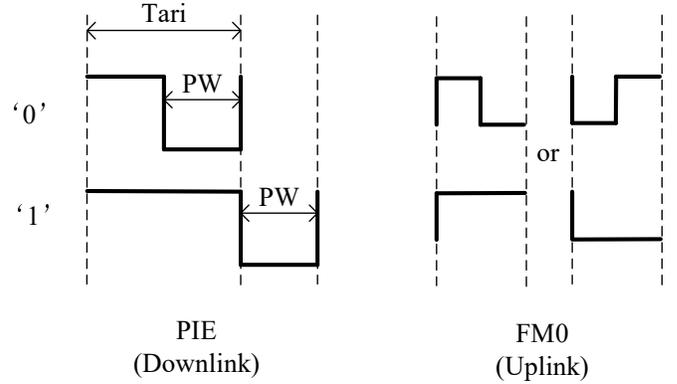


Fig. 1. Examples of downlink and uplink symbols. The downlink rate, ranging from 40 to 160 kbps, is controlled primarily by the length of Tari; The uplink rate, ranging from 5 to 640 kbps, mainly depends on encoding schemes (FM, Miller2/4/8) and backscatter link frequencies.

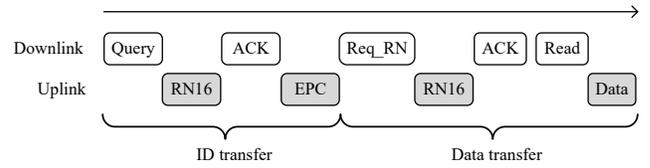


Fig. 2. Reading data from a tag following the C1G2 protocol. The reading process includes an ID transfer phase and a Data transfer phase, each of which has a handshaking through several different commands.

backscatter links. For example, the tag typically has a dipole antenna with a gain of 2.1 dBi and a sensitivity of -13 dBm, while the reader is with a circularly polarized antenna that has a gain of 9 dBi and a sensitivity of -80 dBm. Accordingly, the downlink symbols are amplitude-modulated Pulse Interval Encoding (PIE) symbols, which are easy to decode because an analogy comparator is enough. As shown in Figure 1, downlink symbol ‘0’ is composed of a power-on interval and a power-off interval of equal length. The total length of symbol ‘0’ defines Tari (Type A Reference Interval) and PW (pulse width) is half of Tari. A symbol ‘1’ differs from ‘0’ only in the power-on interval length; The total duration of ‘1’ should be more than 1.5Tari and less than 2Tari. The C1G2 protocol specifies the typical values of Tari: 6.25, 12.5, and 25 μ s, which correspond to downlink rates of 160, 80, and 40 kbps². In contrast, the uplink data rate is configured by setting BLK (Backscatter Link Frequency) and different encoding schemes (FM0, M2/4/8). For example, if the uplink is set at a BLK of 250 kHz using Miller2, its data rate is 250/2 = 125 kbps. Note that both rates of uplink and downlink are controlled by the reader.

C1G2 Protocol. The C1G2 protocol specifies how the reader interrogates tags through several rounds of handshaking. We briefly describe its data reading as follows³. As shown in Figure 2, basically the reading process includes two phases: ID transfer and Data transfer. First, the reader starts by

²These are maximum rates assumed all symbol-0s.

³For more details please refer to [10].

¹We use sensors and tags interchangeably in this paper.

transmitting a *QUERY* command that contains a Q parameter, which specifies how many slots are included in a query round. Then the tag would choose a random number in $[0, 2^Q - 1]$ as its slot counter. If this counter is equal to 0, the tag replies a 16-bit random number (*RN16*); otherwise, the counter decreases 1 after each *QUERY/QUERYREP*. On receiving the *RN16*, the reader sends an *ACK* that contains the decoded *RN16* to the tag. If the tag confirms the reader-decoded *RN16* is correct, it backscatters an identifier, *EPC* (typically 96 bits). This is the end of the ID transfer phase. If the reader needs data from the tag, it starts another round of handshaking through *REQ_RN*, *RN16*, and *ACK* messages. If this round of handshaking goes well, the tag would reply the memory data upon receiving a valid *READ* command.

Our focus in this paper is to choose optimal rates for both the uplink and downlink that can maximize the overall throughput while conforming to the C1G2 protocol. Optimizations from other aspects, such as rateless coding, energy efficiency, or the fairness of MAC, are out of this paper's scope and thus are not considered.

III. OVERVIEW

Figure 3 presents the framework of RAB. The cornerstone of RAB is our observation that we should adapt data rates for both the downlink and uplink to maximize throughput. While common wisdom says that the uplink rate should be properly chosen to improve the throughput of the backscatter link, we argue that the downlink rate should be treated in the same way as there is a tradeoff in setting the downlink rate. Our experiments show that too slow downlink rates could lose the chance to increase throughput when the channel is good, which motivates us to increase the downlink rate. At the same time, we also observe that too aggressive downlink rates can bring down the throughput even to 0 when a bad channel is present because of the well-known sharp transition between low and high loss rates [11] due to multipath fading. By using a rate mapping algorithm, we choose the optimal rates for both the uplink and downlink using overall loss rates and RSSIs that capture multipath fading and path loss, respectively.

While RSSIs are the standard output of most readers, loss rate measurements are not readily available. To measure the loss rate accurately, we introduce a filter-based probing scheme that avoids the potential MAC collisions of multiple tags and thus is able to achieve fast probing regardless of the tag population. To do so, we leverage the built-in *SELECT* command provided by the C1G2 protocol, making our probing lightweight and suitable for point-to-point measuring. In addition, we design a link timing based loss-rate estimation to overcome the invisibility brought by the programming interfaces of COTS readers. Link timing is another unique characteristic of backscatter communication, which ensures the compatibility of devices from different manufacturers. By using such link timing structure, we can accurately approximate how many queries have been sent and thus derive the loss rate.

The final module is to answer a question: when to probe. We design a reliable probing trigger to further reduce the

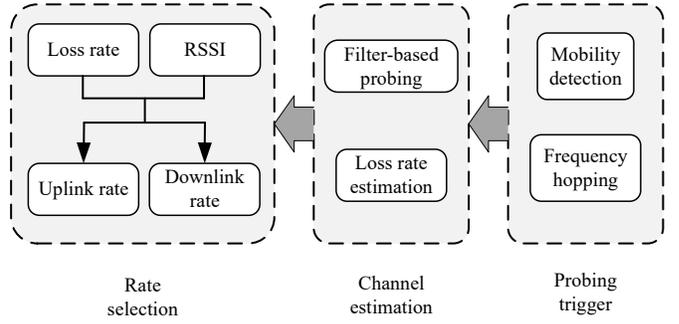


Fig. 3. The framework of our rate adaption scheme including three modules: rate selection, channel estimation, and probing trigger.

probing cost by combing a PHY-assisted mobility detection and a correlation-based channel hopping. In our mobility detection, we mainly make use of a PHY-hint, *phase*, which is widely used in many localization schemes and supported by all COTS readers and the LLRP standard [12]. Differing from [8], [9], it is lightweight and does not need measurements from multiple channels. Channel hopping is another time window for probing. We present a fast channel hopping that is based on the observation that good/bad channels tend to get together instead of being randomly distributed in the spectrum. Therefore, our strategy is that staying away from the probed bad channel and sticking around the good channel.

IV. RATE SELECTION

A. Backscatter Link Characteristics

As discussed before, a backscatter link consists of a downlink that is Reader-to-Tag and an uplink that is Tag-to-Reader. Prior work mainly focuses on adapting appropriate rates for the uplink for two reasons. First, the path loss fading of an uplink is more severe than its corresponding downlink because, while power decays with the square of distance for the downlink, it decays with the fourth power of distance for the uplink. Second, the uplink is supposed to transfer more important data, like sensing information, while the downlink is more viewed as a way to disseminate parameters/commands. However, a key point that is largely ignored is that if there is anything wrong with the downlink, e.g., decoding errors, the corresponding uplink would be discontinued, leading to handshaking failures.

From previous sections, we know that the downlink rate can be set by adjusting the value of T_{ari} . To examine the impact of different T_{ari} values on the throughput, we keep a tag at a fixed place and $BLK=250$ kHz. Then we vary different encoding schemes for the uplink link. The results are shown in Figure 4a. This is a link with good channel quality where faster rates have better throughput. The optimal rates in this case are $T_{ari}=6.25$ for the downlink and FM0 for the uplink. Therefore in the case of good channels, we would miss the chance to increase throughput if a conservative T_{ari} is chosen. For example, with M2 for the uplink, the throughput of $T_{ari}=6.25$ is 171 reads/s, but it drops to 120 reads/s with $T_{ari}=25$. This observation motivates us to use the

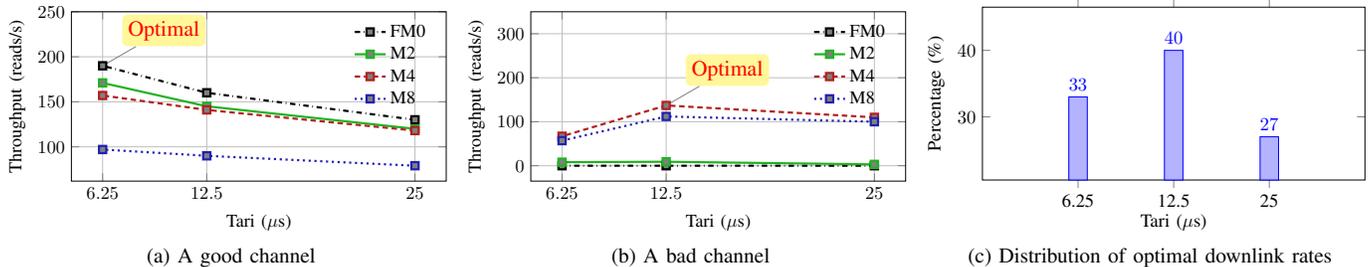


Fig. 4. To examine the impact of data rates of both the uplink and downlink, we measure throughput with various settings. (a) is an example of a good channel, which favors the fastest uplink rate (FM0) and downlink rate (Tari=6.25); (b) is an example of a bad channel. Specifically, both FM0 and M2 encoding settings do not work, and the performance of Tari 6.25 is even worse than that of Tari 12.5, which suggests Tari 6.25 is an aggressive choice. (c) is the distribution of optimal Tari values across 100 random locations, showing that there is no single Tari value that is dominating.

fastest rate for maximizing throughput. However, this is not always the case. As we move the tag to an 1-meter away location, we observe different behaviors. As shown in Figure 4b, this time the link is experiencing some difficulties because the throughput of both FM0 and M2 encoding schemes is almost 0. In this case, the optimal rates become that Tari=12.5 for the downlink and M4 for the uplink. This case tells us that too aggressive rates would not benefit but hurt overall throughput in the case of not good channels. In addition, we measure links at 100 random locations and plot the distribution of optimal Tari values in Figure 4c, which shows that there is no single Tari value that is dominating. To summarize, the above observations suggest that the optimal Tari should be carefully chosen to maximize the throughput based on the quality of channels.

B. Rate Mapping

To find the optimal rates for the uplink and downlink, we adopt a classification-based approach that takes loss rates and RSSIs as input. Although RSSIs are inaccurate in measuring backscatter signal strength due to self-interference [8], they are still useful in indicating path loss. At the same time, the overall loss rate entails multipath fading for both the uplink and downlink. This feature is very important because our hypothesis is that multipath fading is the main reason that the aggressive rate, Tari=6.25, would not always be the optimal rate for the downlink where path loss is less of a problem.

Our rate selection map is built as in Figure 5. The intuition behind this mapping is that when the loss rate increases, more complex encoding schemes should be introduced for resisting channel errors; when the RSSI decreases, the lower-throughput uplink is used to combat path loss. In addition, the impact of both the uplink and downlink under multipath fading is accounted into the loss rate. Therefore, this mapping essentially is able to deliver accurate and fast rate selection. While classes in Figure 5 are only for illustration, the real sizes and types of classes are empirically learned through a training set collected in indoor environments. After all the classes are established (class center and distance), we map a new pair of measured loss rate and RSSI to the closest class.

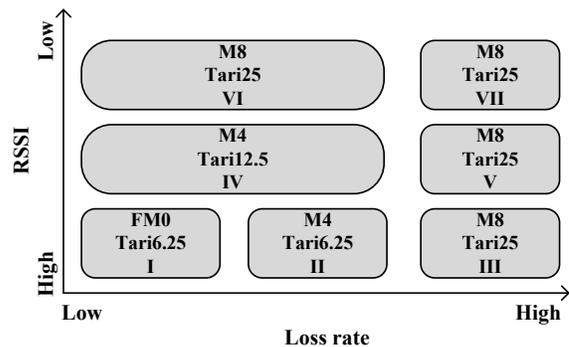


Fig. 5. Optimal rate map of the uplink and downlink. When RSSIs decrease, we choose the downlink with lower throughput. When loss rates increase, we use slower encoding schemes of the uplink to combat the interference. Note that BLK is not considered here for simplicity.

V. CHANNEL ESTIMATION

For rate selection, we assume that the loss rate is known. However, it is not readily available in practice. In this section, we show how to efficiently probe and estimate the loss rate.

A. Filter-based Probing

Previous work of backscatter channel probing is neither accurate nor efficient. The inefficiency of Blink and CARA comes from the C1G2 MAC that is designed for tags that cannot sense each other because probing packets still need to follow the same MAC. There have been many solutions on how to overcome such inefficiency [5], [13]. While those efforts achieve significant efficiency by overhauling the C1G2 MAC, they are overkill for just channel probing. Furthermore, those solutions bring inevitable incompatibility with the C1G2 protocol and thus lose interoperability with many COTS tags.

Our solution for this is that we make use of the built-in *SELECT* command of the C1G2 protocol to create a filter for probing. The *SELECT* command is designed for choosing a tag population for inventory and access. One or more tags are selected by the reader according to user-specified criteria, which is analogous to selecting records from a database. In a *SELECT* command, the reader can specify which *Memory Bank* to match, the associated starting address and length, and

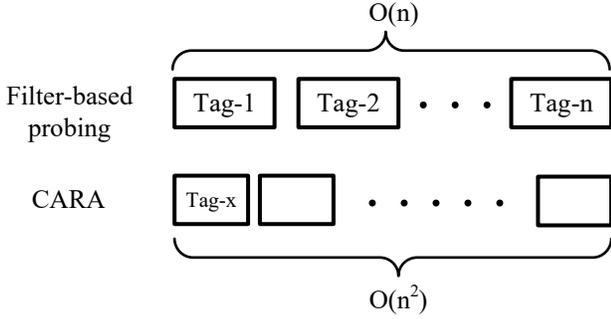


Fig. 6. Probing cost comparison between ours and CARA [9]. While our probing time grows linearly with the number of tags, n , the probing time of CARA grows quadratically with n .

a *MASK*. There are four types of memory banks: Reserved, EPC, TID, and User memory. For example, if we know a tag's ID in advance, then we can easily make it selected by simply sending a *SELECT* command specifying the memory bank as EPC, starting address as 0, length as 96, and *MASK* as the wanted tag's ID. This way, only the tag that matches the mask would reply. Note that this method requires the ID information before probing. As our goal is to maximize the throughput for reading sensor data, we should know which sensor we would like to collect data from in advance. Even sometimes we may not know the sensor's ID beforehand, as shown in Figure 2, the data transfer phase is always preceded by an ID transfer phase. Therefore, knowing the ID of a sensor before transferring the data is not a problem for us. For the rest of the paper we assume the IDs of tags are known before reading sensor data.

Now by using the *SELECT* command, we enable a point-to-point probing style that avoids MAC collisions completely. Usually, a *SELECT* command is about 45-bit long (excluding the *MASK*), which incurs some extra cost. However, such cost is considerably less than the waste due to the inefficient MAC, as shown in Figure 6. As seen from the figure, although each of our filter-based probing slots is larger than that of CARA's scheme. But CARA's probing time increases quadratically with the number of tags while ours grows linearly, which means with more and more tags coming in, the probing overhead we save would be even greater. Note that Blink suffers from the problem as CARA does.

B. Loss Rate Estimation

After probing, the next step is to estimate the loss rate of the link. Unlike USRP-based readers, COTS readers do not offer the way to directly measure the loss rate and are more like a black box. The only result from probing is the number of successful reads in a given time interval. Therefore, we need to estimate how many probes/queries sent in a given period of time. While many prior efforts try to solve this, they all need extra hardware. For example, Flit [13] logs all the message counts into EPC using CRFIDs; [14] uses an extra USRP-based monitor. To solve this without additional hardware, we observe an opportunity of making use of precise timing

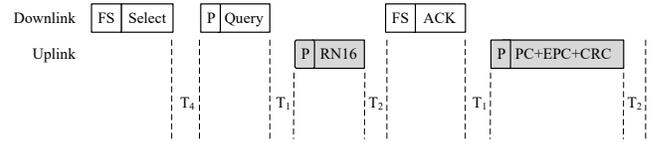


Fig. 7. Link timings of a probe. The C1G2 protocol has strict timing requirements for each message, giving us opportunities to estimate loss rates. P denotes either an uplink or downlink Preamble. FS denotes the Frame-Sync symbol.

structures that are specified in the C1G2 protocol. The original intent of such timing structures is to ensure the compliance and interoperability of devices from different vendors. While it is mainly used for conforming tests for backscatter devices, we here use its preciseness of the structures as a new way to do estimation because all the timings of downlink and uplink messages are strictly bounded.

The timing of probe includes two parts: data transmission delays for the uplink and downlink, and built-in protocol delays. Hence, our first step is to take into account of the data rate and the amount of data to be sent over both the uplink and downlink. Then we need to find certain delays built in the protocol, as shown in Figure 7. The first specified timing limitation is T_4 , which is the time that the reader has to wait before issuing another command. The length of T_4 is $2RT_{cal}$, where $RT_{cal} = 0_{length} + 1_{length}$. After the *QUERY* command, the tag needs to wait for T_1 , of which the nominal value is $\text{MAX}(RT_{cal}, 10T_{pri})$, where $T_{pri} = 1/BLK$. If there is a reply from the tag, the reader must acknowledge it within T_2 , ranging from $[3T_{pri}, 20T_{pri}]$. T_1 and T_2 also apply to the *ACK* and *EPC* messages.

Now let us take a case study to examine the probing process. Table I gives an example showing the timings of a probe by walking through all the messages in Figure 7. From the table we know that a probe using $T_{pri}=6.25$ and FM0 would take about 2.5 ms, corresponding to 400 probes/second. However, in the field study, our measured result is around 250. This is because there is a hardware-dependant command delay between two probes. Besides this uncertain hardware-dependent delay, we model all uncertain parameters in the protocol into a linear system, including T_1 , T_2 , T_4 , and 1_{length} . To build the linear system, we make multiple measurements across different settings and use the *constrained least square* method to estimate unknowns. After we have loss-rate estimates, the final question is when to probe, which is detailed in the next section.

VI. PROBING TRIGGER

The probing trigger decides when to probe the channel, which is very important because too often probing poses unnecessary overhead and too rare probing would lose the chance to adapt rates. Our probing trigger includes two indicators: mobility detection and channel hopping.

⁴For the details of the message format, please refer to [10].

⁵This includes preambles and FrameSync(FS) symbols.

⁶This includes a 16-bit PC, a 96-bit EPC, and a 16-bit CRC.

TABLE I

EXAMPLE OF PROBE TIMING ESTIMATION. THE SETTINGS ARE TARI=6.25 μ s, BLK=250 KHZ, RTCAL=2.75TARI, TREXT=0, ENCODING=FM0, $l_{\text{length}}=1.75\text{TARI}$, FS=12.5 μ s + 3.75TARI, $P_{\text{down}}=FS + 2.05\text{RTCAL}$, $P_{\text{up}}=6$ BITS.

Messages	Length (bits) ⁴	time (μ s) ⁵	Cumulative time (μ s)
Select	141	1247.7	1247.7
T ₄	-	31.4	1279.1
Query	22	260.2	1539.3
T ₁	-	40	1579.3
RN16	16	88	1667.3
T ₂	-	46	1713.3
ACK	18	190.6	1903.9
T ₁	-	40	1943.9
EPC	128 ⁶	536	2479.9
T ₂	-	46	2525.9

A. Mobility Detection

When a sensor moves to another location, its channel inevitably changes. At this time, a reader may need to choose the optimal rate for this new position to maximize the throughput. While many localization schemes have been proposed for RFID devices, they either require a number of antennas [15], or are not fast and lightweight enough for channel estimation purposes [16]. Blink uses link signatures to detect mobility, yet it requires measurements from at least 10 channels. Because the channel switching on COTS readers takes at least 30 ms, such multiple-channel detection introduces too much overhead.

To address this issue, we propose a zero-overhead mobility detection on a single channel. The solution is to use phase, a PHY-hint, which is supported in COTS readers as specified in the LLRP standard. For every successful read, the reader outputs a phase reading and an RSSI value, making it virtually zero-overhead. The reported phase is an effective way to measure the distance between the reader and tag, R . The relationship between such distance and measured phase, θ , is as follows [16],

$$\theta = 2\pi \frac{2R}{\lambda} + \theta_D + \theta_R + \theta_M + N\pi,$$

where λ = is the wavelength, θ_D , θ_R , θ_M , are phase errors brought by tag and antenna diversity, reflection characteristics, and multipath, respectively, N is the integer ambiguity as the measured phase is with period π . Therefore the distance between two locations is approximated as

$$\Delta R \approx \frac{\lambda}{4\pi} \Delta\theta.$$

To set up a threshold that detects mobility, we conduct an empirical study. Figure 8 shows 500 phase measurements when a tag is static. We observe that when the tag is stationary, the phase measurement is highly concentrated. Specifically, the variance is only 2.2°, and the gap between the min value and

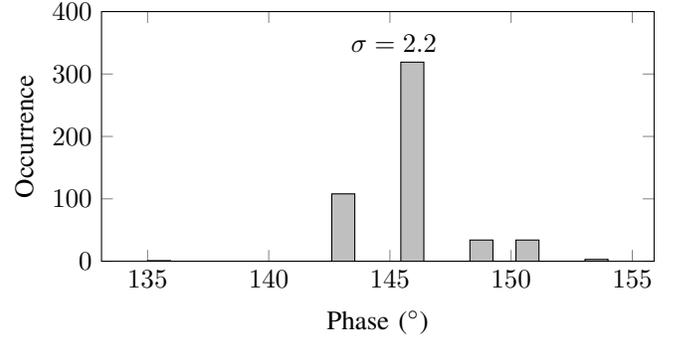


Fig. 8. 500 phase measurements across different uplink and downlink rates when the sensor is static. The high concentration of these measurements shows that phase difference is a good indicator for mobility detection.

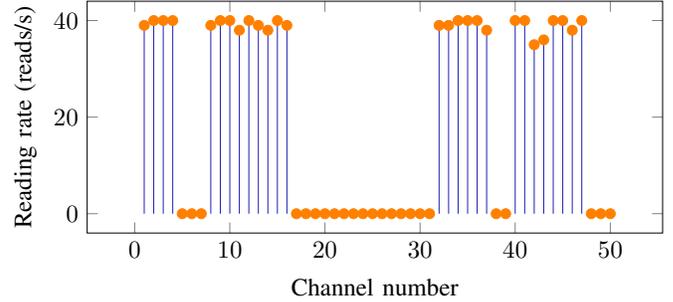


Fig. 9. Throughput measurements across 50 channels. We observe that a strong channel correlation exists. For example, channels 17-31 have 0 reading rate while channels 7-16 have high reading rates. Another observation is the sharp transition between high and low loss rates.

max value is only 19°(0.33 radians), which only corresponds to 0.8 cm. Therefore, we set up a threshold $\theta_{th} = 0.33$.

Note that to ensure that N is the same for two consecutive phases, the phase rotation between the two should be less than π . This requirement is equal to that when the reading rate is 50 reads/s, it can handle moving objects at velocity up to 4 m/s, which is fairly enough for indoor applications. When the reading rate is below this threshold, it could make false negative alarms. To reduce this alarm, we use RSSIs as a second metric and set its threshold at $\text{RSSI}_{th} = 1$, which is the granularity of RSSIs from COTS readers. Therefore, our mobility detection works as follows. First, we check whether the phase difference is greater than θ_{th} , if so, we label it as a positive location change; otherwise, we check whether the RSSI difference is greater than RSSI_{th} , if so, it is positive, otherwise negative.

Note that environmental mobility, e.g., human/metal objects moving nearby, could be misidentified as location changes because link characteristics, e.g., RSSIs and phases, are easily affected by multipath. In fact, such misidentification is beneficial to our system because it is the channel change that causes misidentification and thus makes probing necessary.

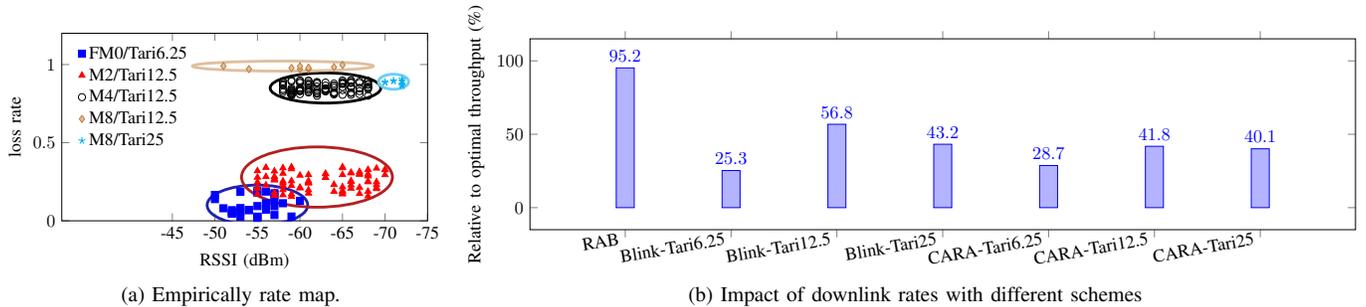


Fig. 10. We learn an empirical rate map from over 200 samples as in (a), which can be used to guide the rate selection for measured RSSI and loss-rate pairs; then we compare RAB’s rate selection against BLINK and CARA, showing that RAB has significant improvement thanks to the optimal rate selection of downlink rates.

B. Channel Hopping

Our second trigger is based on channel hopping, which is mandatory as defined in the C1G2 protocol that the reader can only stay on a channel in a time window. The quality of channel may change due to hopping so that it is the chance the reader needs to adapt rates. Prior work, such as *selection* in [8], needs to probe all the channels to choose top ones, incurring substantial unnecessary overhead.

Our hopping scheme is based on the observation that neighbor channels tend to get together, exhibiting channel correlation. We conduct an empirical study of channel correlation and plot results in Figure 9. We observe a strong channel correlation, i.e., good or bad channels could be clustered by channel indexes. This motivates us to design a correlation-based hopping scheme. Specifically, when the current channel is good, we choose to probe the next channel that is within h_g -hop of the current one; if the probed channel one is good, we stay, otherwise, we will switch to another one that is far away from the probed one, say h_b -hop distance. The channel gap is empirically set at $h_g = 3$ and $h_b = 5$. To decide a channel is good or bad, we use a very conservative threshold 5 reads/s. The rationale of this setting is the observation that the transition between high and low loss rates is sharp, as shown in Figure 9, which is also confirmed in [11].

VII. IMPLEMENTATION

In this section, we present details of our evaluation.

Reader: We mainly use a Thingmagic M6e reader for implementation, which is fully compatible with the C1G2 protocol. Same as [8], the COTS reader has three limitations due to API constraints: First, the data rate can only be set up at the beginning of a query round; Second, the channel switching is not lightweight and takes about 30 ms; Third, the minimum probing time is 30 ms. We hope these factors will be addressed in the future readers. Currently, we only use trace-driven studies to examine the aspects that are bounded by the above limitations, such as channel switching.

Tag: Although we have tested many tags from different vendors, such as Impinj, NXP, we do not observe significant performance differences. Thus we choose a representative, the Alien Higgs 3 tag, AZ-9640. One of the main reasons that we

extensively use this tag is that it has the largest user memory, which is 512 bits, among tags in the same price range. As the content of sensor data does not affect our protocol at all, we write 512 random bits into the user memory of each test tag in advance.

Parameter: The Thingmagic M6e provides two BLK options, 640 kHz and 250 kHz, but only FM0 and Tari 6.25 are allowed with 640 kHz. Thus we mainly use 250 kHz for BLK on this reader, which allows Tari 6.25, 12.5, 25 and FM0/M2/4/8 on this frequency. For probing, we set up $Q=1$ to avoid MAC collisions and a filter of which the memory bank is EPC, the starting address is 32, the length is 96, and the mask is the target tag’s ID. The rates of probing packet are fixed at the slowest: M8 and Tari 25. The reader power is fixed at 30 dBm.

Competition: We compare RAB with two state-of-the-art schemes, Blink [8] and CARA [9]. To ensure a fair competition, rate adaptation schemes from other wireless networks, e.g., SampleRate [17], are not included as no clear standards or publications have specified how to adapt them to backscatter networks, because a backscatter link is two-way not one-way for other wireless networks.

VIII. EVALUATION

Rate selection: To begin with, we investigate how our rate selection scheme works. As Figure 5 only shows the intuition how rates would adapt to different locations, the actual boundaries of different classes could be irregular. Figure 10a is the empirical rate map we learn from 230 randomly sampled locations in our testbed of size $4m \times 5m$. At each location, we measure all possible combinations of downlink and uplink rates. As expected, we observe that not every class is on the map and the boundaries are not regular. In addition, the trend of different classes does go with our prediction that when the RSSI decreases, the lower throughput of the downlink is favored; when the loss rate increases, a slower encoding scheme should be used. Note that our classifier has some errors. For example, some points of FM0/Tari6.25 and M2/Tari12.5 are mixed, because the throughput of both is similar.

To further check the impact of downlink rates, we compare it with Blink and CARA. Since both Blink and CARA do not

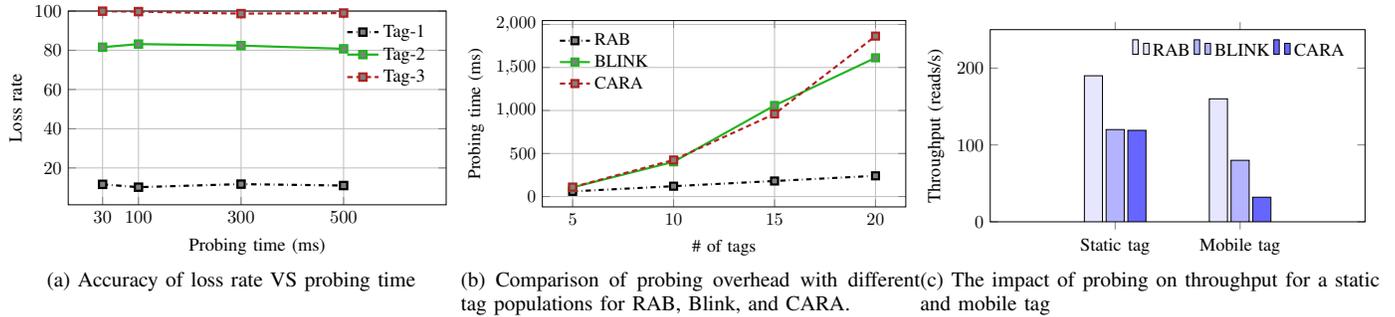


Fig. 11. We examine our probing scheme in detail. (a) shows that a time interval of 30 ms is enough to accurately estimate loss rates; (b) shows that the probing costs of Blink and CARA are way larger than that of RAB; (c) shows our lightweight probing benefits the throughput in both static and mobile scenarios.

TABLE II
APPLYING THE LEARNED MAP TO DIFFERENT SCENARIOS ACROSS TIME AND PLACES.

	Accuracy (%)	relative to optimal throughput (%)
Testbed - 1st day	93.4	96.4
Testbed - 2nd day	94.5	98.1
Testbed - 3rd day	92.5	93.1
Classroom	83.2	90.2
Library	76.5	86.3
Lounge	77.9	85.7

consider the downlink rate, we make three variants for them, each of which has a distinct Tari. The results are plotted in Figure 10b. Not surprisingly RAB outperforms all the variants of Blink and CARA because a single fixed Tari cannot bring too much gain across different location and channel conditions. One interesting thing to note is that the fastest downlink rate, Tari 6.25, performs even worse than other Tari values. It is mainly because that the too aggressive rate hurts the downlink and makes uplink and overall throughput suffered.

To verify the effectiveness of our rate map, we apply it to various scenarios that are with different dates and places. The results are shown in Table II. First, we test this rate map for three consecutive days in our testbed and obtain testing data of 200 samples for each day. We achieve more than 90% rate selection accuracy and more than 90% of the optimal throughput for three days, which shows the robustness of our scheme against time. Then, we apply the map at three different places including classroom, library, and lounge. The rate selection accuracy decreases a bit due to the different background of the place, yet the achieved throughput is still more than 85% of the optimal one. This is because the boundary errors in the empirical rate map make the rate selection accuracy degraded, but the similar performance of boundary points keeps the overall throughput not affected too much.

Probing cost: Next, we examine the impact of our probing scheme. First, we need to determine how long should we

probe. Figure 11a shows the probing results across different time intervals for 3 different tags. We observe that the accuracy of probing is not sensitive to the time interval for low and high loss rates. Therefore, we set the probing interval at 30 ms. Note that 30 ms is the minimal time window that is allowed on COTS readers.

Furthermore, we compare our probing cost against Blink and CARA with different tag populations. To avoid the negative effect of 30 ms minimal window that severely degrades the probing performance of Blink and CARA, this comparison is done with traces. Figure 11b demonstrates that the probing cost of Blink and CARA grows quadratically with the number of tags while that of RAB increases linearly. Specifically, the probing costs of Blink and CARA are 1612 ms and 1864 ms, corresponding to 6.7x and 7.8x more than that of RAB when there are 20 tags. This is primarily due to the filter-based probing paradigm that probes tags sequentially while Blink and CARA need more time to deal with MAC collisions.

To investigate the impact of our lightweight probing scheme on the throughput, we compare it under static and mobile scenarios. To eliminate the impact of MAC collisions and channel hopping, we only use 1 tag and 1 channel. Figure 11c shows that the throughput of RAB is considerably better than those of Blink and CARA. Also, while there is no much difference between Blink and CARA in the static setting, CARA suffers more degradation than Blink does in the mobile scenario because CARA is not mobility-aware.

Loss rate estimation: Now we look to check link timing based loss rate estimation. As the number of successful reads is known from the reader output, we only need to examine the accuracy of query estimation. For the ground truth, we use a USRP-based monitor at a very close distance, 10 cm, to capture messages between the reader and the tag. The results in Table III show that our estimation achieves less than 5% errors all the time and thus are quite robust across a range of different rate settings. Such errors do not affect the rate selection as shown in Figure 10a. Note that while prior methods can also obtain loss-rate estimates, they require either a USRP monitor or CRFID tags [13]. In contrast, our method is accurate and does not need any extra hardware because we make use of the link timing feature of backscatter communication.

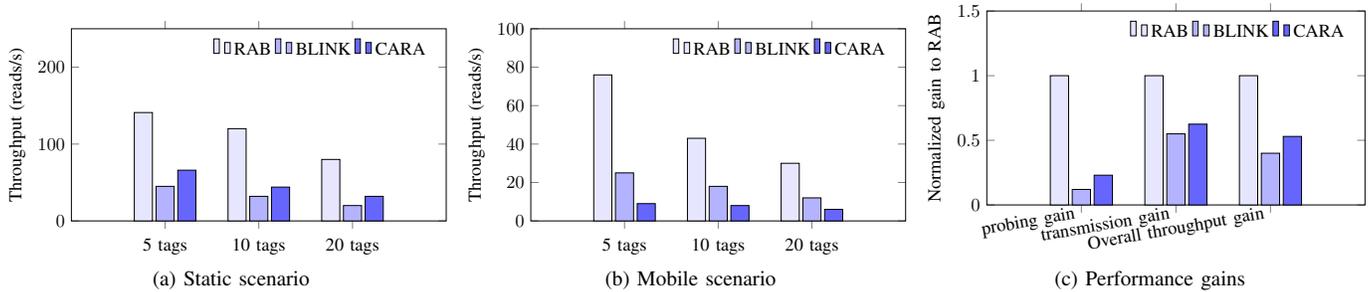


Fig. 12. Overall performance comparison under static and mobile scenarios with different tag populations.

TABLE III
QUERY ESTIMATION ACROSS DIFFERENT RATES.

	query measured	query predicted	relative error (%)
FM0/Tari6.25	248.8	258.3	3.8
Miller2/Tari6.25	244.6	256.4	4.8
Miller4/Tari6.25	235.7	224.4	4.7
Miller8/Tari6.25	127.1	130.9	3
FM0/Tari12.5	246.2	255.8	3.9
Miller2/Tari12.5	245.1	241.5	1.5
Miller4/Tari12.5	209.8	214.4	2.2
Miller8/Tari12.5	122.0	123.8	1.4
FM0/Tari25	244.6	233.8	4.4
Miller2/Tari25	243.8	241.1	1.1
Miller4/Tari25	175.6	182.9	4.1
Miller8/Tari25	106.3	105.6	0.6

TABLE IV
ACCURACY OF MOBILITY DETECTION ACROSS DIFFERENT RATES.

	False positive (%)	False negative (%)
FM0/Tari6.25	6.6	1.3
Miller2/Tari6.25	3.6	0.9
Miller4/Tari6.25	0	0.5
Miller8/Tari6.25	0	0
FM0/Tari12.5	5.8	1.0
Miller2/Tari12.5	4.9	0.8
Miller4/Tari12.5	0	1.0
Miller8/Tari12.5	0	0
FM0/Tari25	3.9	0.3
Miller2/Tari25	4.2	0.2
Miller4/Tari25	0	0
Miller8/Tari25	0	0

Mobility Detection: The accuracy of mobility detection is very important since it decides when to probe. In this evaluation, we compare its accuracy across all available rates. Table IV shows that by using RSSIs and phases together, our mobility detection achieves less than 7% false positive rates and less than 1.5% false negative rates with various data rates. The false positive rate is a bit higher because sometimes phases could be affected by even minor interferences and internal hardware imperfections, such as carrier frequency offset. Yet overall, our mobility detection is very robust and accurate enough for triggering probe, because such low false positive rates marginally bring down overall throughput.

Overall performance: We now look at the overall performance of the whole framework and compare it with state-of-the-art systems. First, we study the static case where all tags are placed randomly. Figure 12a shows that when there are 5 tags, the throughput of RAB is 3.1x and 2.1x better than Blink and CARA, respectively. The same trend can be observed when the number of tags increases. As expected, all schemes degrade with the increasing number of tags because of more coordination time needed.

When it turns to the mobile case in Figure 12b, all of the three systems are affected by mobility differently, but RAB

is still the best across different tag populations. Particularly, when the number of tags is 20, RAB achieves 2.5x and 5x throughput gains over Blink and CARA. CARA is the worst due to its lack of mobility detection module.

Then we conduct over 80 tests across different mobility, channel, and network-size conditions. For mobility, we vary the velocity of tags from 0 to 1 m/s. For channels, we collect the data across 1-week at two different places. The tag population varies from 1 to 20. The overall gains and its breakdown on average are reported in Figure 12c. RAB achieves overall throughput gains of 2.5x over Blink and 1.9x over CARA. We break down this gain and find that RAB reduces probing cost by 8.2x and 4.3x over Blink and CARA. The majority of this probing gain comes from the filter-based probing design as it successfully avoids MAC collisions while being compatible with the CIG2 protocol. Meanwhile, regarding data transmission, RAB is 1.8x and 1.6x better than Blink and CARA. This transmission gain is mainly brought by the downlink-aware rate selection scheme while all prior systems, like Blink, leave the downlink unattended.

IX. RELATED WORK

Backscatter Communication Efficiency: Backscatter communication optimizations can be roughly classified into two

categories: C1G2-compatible and C1G2-incompatible. Buzz [5] introduces a rateless coding for backscatter nodes, which achieves lossless transmission. Flit [13] designs a new MAC that enables burst transferring bulk data, significantly reducing wasted time by the C1G2 MAC. Laissez-Faire [18] and BiGroup [19] propose to decode parallel transmissions by analyzing signals in the both time and IQ domains, which can work at moderate and high SNR scenarios. Those C1G2-incompatible optimizations achieve substantial performance gain but fall short of accommodating billions of deployed RFID readers and nodes. Some C1G2-compatible improvements have been proposed recently. Blink [8] makes use of unique backscatter link signatures to detect mobility and adapt rates. CARA [9] observes the opportunity that throughput can be improved by channel-aware rate selection. Unlike both that focus on the uplink rate selection, we observe that the downlink rate could greatly affect the overall throughput as well. In addition, our filter-based probing tries to efficiently estimate channels and avoid collision problems that are not well considered before.

Rate Adaptation: Rate adaptation has been widely researched in active-radio based wireless networks, like 802.11. BER[20], SNR [21], [22], and loss rate [23] are the most commonly used metrics. While our work shares the same idea that chooses the optimal rate that maximizes the network throughput by estimating the channel quality. Those methods have limited applicability to backscatter systems, especially for the C1G2 protocol. For example, the limited visibility of current COTS readers makes even loss rates hard to observe. To solve this, we use the link timing features specified by the C1G2 protocol to approximate the loss rate. In addition, we accurately deduce mobility hints using RSSI and phase measurements together.

New Backscatter Paradigms: Recently several novel backscatter systems where nodes are powered by various sources have been proposed, e.g., WiFi-backscatter [24], [25], [26], Bluetooth-backscatter [6], FM-backscatter [7]. Those systems largely extend the operating range of traditional readers and see a bright future of interconnecting more and more wireless devices. Yet, their interpretability with C1G2 is worth further investigation.

X. CONCLUSION AND FUTURE WORK

We have presented RAB, a protocol that is to optimize throughput within the C1G2 standard from many aspects, including downlink-aware rate selection, filter-based probing, and lightweight probing triggers. Our prototype has shown that considerably throughput gains have been achieved over state-of-the-art schemes. With more and more backscatter sensors have been invented, we believe RAB can benefit a range of Internet-of-Things applications. Our future work includes 1) introducing multiple-antenna to further improve throughput as currently only a single antenna is used; 2) investigation of reading performance with large user memory, e.g., 2K-bit of ImpinJ Monza tags; 3) extension and interoperability with other backscatter paradigms that connect more wireless devices, like WiFi/FM-backscatter.

ACKNOWLEDGMENT

We would like to thank the anonymous reviewers for valuable and insightful comments. This work was supported by the Canada Technology Demonstration Program, a Canada NSERC Discovery Grant, and an NSERC E.W.R. Steacie Memorial Fellowship.

REFERENCES

- [1] V. Hsu, J. M. Kahn, and K. S. Pister. *Wireless communications for smart dust*. Electronics Research Laboratory, College of Engineering, University of California, 1998.
- [2] J. R. Smith, A. P. Sample, P. S. Powlledge, S. Roy, and A. Mamishev. A wirelessly-powered platform for sensing and computation. In *Proc. of UbiComp*, 2006.
- [3] Y. Zhao and J. R. Smith. A battery-free RFID-based indoor acoustic localization platform. In *Proc. of IEEE INFOCOM*, 2013.
- [4] S. Naderiparizi, A. N. Parks, Z. Kapetanovic, B. Ransford, and J. R. Smith. WISPCam: A Battery-Free RFID Camera. In *Proc. of IEEE RFID*, 2015.
- [5] J. Wang, H. Hassanieh, D. Katabi, and P. Indyk. Efficient and Reliable Low-Power Backscatter Networks. In *Proc. of ACM SIGCOMM*, 2012.
- [6] V. Iyer, V. Talla, B. Kellogg, S. Gollakota, and J. R. Smith. Inter-technology backscatter: Towards internet connectivity for implanted devices. In *Proc. of ACM SIGCOMM*, 2016.
- [7] A. Wang, V. Iyer, V. Talla, J. R. Smith, and S. Gollakota. FM-Backscatter: Enabling Connected Cities and Smart Fabrics. In *Proc. of USENIX NSDI*, 2017.
- [8] P. Zhang, J. Gummesson, and D. Ganesan. Blink: A high throughput link layer for backscatter communication. In *Proc. of ACM MobiSys*, 2012.
- [9] W. Gong, H. Liu, K. Liu, Q. Ma, and Y. Liu. Exploiting Channel Diversity for Rate Adaptation in Backscatter Communication Networks. In *Proc. of IEEE INFOCOM*, 2016.
- [10] EPC C1G2 Standard. <http://www.gs1.org/epcrfid/epc-rfid-uhf-air-interface-protocol/2-0-1>, 2017.
- [11] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *Proc. of ACM SenSys*, 2003.
- [12] Low Level Reader Protocol. <http://www.gs1.org/epcrfid/epc-rfid-llrp/1-1-0>, 2017.
- [13] J. Gummesson, P. Zhang, and D. Ganesan. Flit: a bulk transmission protocol for RFID-scale sensors. In *Proc. of ACM MobiSys*, 2012.
- [14] M. Buettner and D. Wetherall. An empirical study of UHF RFID performance. In *Proc. of ACM MobiCom*, 2008.
- [15] J. Wang, D. Vasishth, and D. Katabi. RF-IDraw: Virtual Touch Screen in the Air Using RF Signals. In *Proc. of ACM SIGCOMM*, 2014.
- [16] L. Yang, Y. Chen, X. Li, C. Xiao, M. Li, and Y. Liu. Tagoram: Real-time tracking of mobile RFID tags to high precision using COTS devices. In *Proc. of ACM MobiCom*, 2014.
- [17] J. C. Bicket. *Bit-rate selection in wireless networks*. PhD thesis, Massachusetts Institute of Technology, 2005.
- [18] P. Hu, P. Zhang, and D. Ganesan. Laissez-faire: Fully Asymmetric Backscatter Communication. In *Proc. of ACM SIGCOMM*, 2015.
- [19] J. Ou, M. Li, and Y. Zheng. Come and Be Served: Parallel Decoding for COTS RFID Tags. In *Proc. of ACM MobiCom*, 2015.
- [20] M. Vutukuru, H. Balakrishnan, and K. Jamieson. Cross-layer wireless bit rate adaptation. *ACM SIGCOMM CCR*, 39(4):3–14, 2009.
- [21] H. Rahul, F. Edalat, D. Katabi, and C. G. Sodini. Frequency-aware rate adaptation and MAC protocols. In *Proc. of ACM MobiCom*, 2009.
- [22] S. Sen, N. Santhapuri, R. R. Choudhury, and S. Nelakuditi. AccuRate: Constellation Based Rate Estimation in Wireless Networks. In *Proc. of USENIX NSDI*, 2010.
- [23] S. H. Wong, H. Yang, S. Lu, and V. Bharghavan. Robust rate adaptation for 802.11 wireless networks. In *Proc. of ACM MobiCom*, 2006.
- [24] B. Kellogg, A. Parks, S. Gollakota, J. R. Smith, and D. Wetherall. Wi-fi backscatter: Internet connectivity for rf-powered devices. In *Proc. of ACM SIGCOMM*, 2014.
- [25] B. Kellogg, V. Talla, S. Gollakota, and J. R. Smith. Passive wi-fi: Bringing low power to wi-fi transmissions. In *Proc. of USENIX NSDI*, 2016.
- [26] D. Bharadia, K. R. Joshi, M. Kotaru, and S. Katti. Backfi: High throughput wifi backscatter. In *Proc. of ACM SIGCOMM*, 2015.