

An Empirical Study of UHF RFID Performance

Michael Buettner*

*Department of Computer Science and
Engineering
University of Washington
Seattle, Washington, USA
buettner@cs.washington.edu

David Wetherall†*

†Intel Research Seattle
Seattle, Washington, USA
david.wetherall@intel.com

ABSTRACT

This paper examines the performance of EPC Class-1 Generation-2 UHF RFID reader systems in a realistic setting. Specifically, we identify factors that degrade overall performance and reliability with a focus on the physical layer, and we explore the degree to which reader configuration options can mitigate these factors. We use a custom software-radio based RFID monitoring system and configurable RFID readers to gather fine-grained data and assess the effects of the different factors. We find that physical layer considerations have a significant impact on reader performance, and that this is exacerbated by a lack of integration between the physical and MAC layers. We show that tuning physical layer operating parameters can increase the read rate for a set of tags by more than a third. Additionally, we show that tighter integration of the physical and MAC layers has the potential for even greater improvements.

Categories and Subject Descriptors

C.2.2 [Network Protocols]

General Terms

Experimentation, Measurement, Performance, Reliability

1. INTRODUCTION

Passive Radio Frequency IDentification (RFID) is an emerging wireless technology that allows small, inexpensive computer chips to be remotely powered and interrogated for identifiers and other information. Many applications that make use of these capabilities have been proposed, ranging from conventional supply chain monitoring where pallets or individual objects are tagged with an electronic product code and scanned as they are moved [4] [10], to future ubiquitous computing scenarios where interactions between instrumented objects trigger processing [12], e.g., to check the compatibility of

a unit of blood plasma with the patient receiving it. Commercial readers and tags that can be used to realize these applications are now on the market, and it is reasonable to assume that their deployment will grow rapidly. In particular, while there are many types of RFID, the EPC Class-1 Generation-2 standard [2] defines readers and passive tags that operate at UHF frequencies and use ‘backscatter’ communication to support read ranges measured in meters. This is in contrast to earlier HF tags based on inductive coupling that only provide read ranges of centimeters, and active tags that require batteries to increase range. UHF tags, combined with growing tag capabilities in terms of data storage and sensing [15], are expected to make inroads in applications where earlier RFID technologies were otherwise limited.

From a networking viewpoint, UHF RFID as defined by the EPC standard is as interesting as more widely studied protocols such as 802.11. At the physical layer, RFID tags communicate by ‘backscattering’ signals that are concurrent with reader transmissions, and use a variety of frequencies and encodings under the control of the reader. At the MAC layer, readers and tags use a variation on slotted Aloha [14] to solve the multi-access problem in a setting where readers can hear tags but tags cannot hear each other. This differs from an earlier standard (EPC Gen-1) which used a protocol based on binary tree-walking [3]. Reliability is achieved via explicit and implicit acknowledgements and repetition. Moreover, the EPC standard defines a set of primitive commands that can be combined in a variety of ways so that readers can implement proprietary algorithms; there is no single standard method for reading tags.

Given the likely importance of RFID in practice, we would like to know how well the RFID protocol works as it is implemented in commercial readers and tags. This will help us understand the levels of reliability, performance and security that applications can reasonably expect and the key factors that affect these levels. In addition, as RFID applications become more widespread and their requirements more intensive, high performance and reliability will become increasingly important.

Surprisingly, there is very little work on this topic seen in the literature. The RFID studies we have found typically correspond to different protocols than the one we focus on, e.g., short-range tags or EPC Gen-1 protocols. Vendor white papers do report performance for their readers, but often provide limited information that cannot be interpreted without details of the reader configuration. Partly, this scarcity is due to the current lack of tools for studying RFID operation. Most readers

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiCom’08, September 14–19, 2008, San Francisco, California, USA.
Copyright 2008 ACM 978-1-60558-096-8/08/09 ...\$5.00.

and tags are black-boxes that report the end result of an overall read operation. While readers typically have many tuning parameters and configuration options, they do not expose the underlying sequence of RFID reader/tag commands and give little indication of what is happening at the MAC and PHY layers of the system.

In this paper, we present what to the best of our knowledge is the first low-level study of EPC Class-1 Generation-2 reliability and performance with commercial tags and readers. Our study is based on measurements of the Alien Technologies and ThingMagic readers interacting with multiple Alien tags in an office setting. To shed light on the operation of RFID at the MAC and PHY protocol layers, we made use of two kinds of experimental tools. First, we used the Universal Software Radio Peripheral (USRP) platform and GNURadio to monitor reader transmissions in order to determine operating parameters and decode individual reader commands. This enables us to reverse-engineer the protocol implementations and deduce information about errors at tags and readers from observations of the subsequent commands. Second, we used RFID readers that provide a high degree of configurability and give detailed results for each read cycle that include the frequency on which each tag was read, and the time it was read with millisecond resolution.

As an initial study, we focus on reliability and performance when reading tags, measuring the sources of error that cause tag misses and attributing the overall read time to its various components. We find that physical layer effects, such as errors and multipath, can significantly degrade performance by increasing the duration of each reader cycle as well as the number of cycles needed to read all tags in a tag set. Our measurements suggest that simply tuning physical layer operating parameters (while still maintaining the same link-rates and coding scheme) can increase the read-rate for a set of tags by more than one third. Moreover, the physical and MAC layers should be considered in conjunction, rather than separately using a statically configured physical layer as is done today. This would allow for greater performance gains while remaining standards compliant, e.g., by switching encodings in a manner analogous to 802.11 rate adaptation.

The rest of this paper is organized as follows. In the next section we provide a primer on RFID, describing the parts of the EPC Class-1 Generation-2 standard that are relevant to our experimentation. In Section 3, we describe the goals of our study and our experimental tools. In Section 4, we describe our experimental setup and present the results of our measurement study. Then, in Section 5, we use these results to estimate the potential improvements of variations on the RFID protocol. We present related work in Section 6, then conclude in Section 7.

2. BACKGROUND

RFID was developed as a replacement for barcode identification systems, as it provides a number of key advantages such as non-line-of-sight operation, higher inventory rates, and rewritable product IDs. While there are a number of different RFID specifications with different operating frequencies and methods of powering tags, this study focuses exclusively on Class-1 Generation-2 (C1G2) RFID systems as defined by EPCglobal [2].

The C1G2 standard defines communication between RFID readers and passive RFID tags in the 900 MHz band, and has a maximum range of approximately 10m. A reader transmits information to a tag by modulating an RF signal, and the tag receives both down-link information and the entirety of its operating energy from this RF signal. For up-link communication, after transmitting a down-link command the reader transmits a continuous RF wave (CW) which assures that the tag remains powered, and the tag then transmits its response by modulating the reflection coefficient of its antenna. By detecting the variation in the reflected CW, a reader is able to decode the tag response. This is referred to as “backscattering”.

2.1 Class-1 Generation-2 Physical Layer

The C1G2 specification defines a number of physical layer options for both down-link and up-link. All down-link communication uses Amplitude Shift Keying (ASK), where bits are indicated by brief periods of low amplitude, and Pulse-Interval Encoding (PIE), where the time between low amplitude periods differentiates a zero or a one, and the reader can choose pulse durations that result in down-link rates ranging from 26.7 kbps to 128 kbps. Through the use of a structured preamble the tag can determine the pulse lengths being used, so that it can decode reader commands, and also what data rate should be used by the tag for up-link communication.

The up-link rate is determined partially by the down-link preamble and partially by a bit field set in the *Query* command which starts each query round, as will be described in the next section. These settings allow for an up-link frequency ranging from 40 kHz to 640 kHz. Along with setting the up-link frequency, the reader also sets one of four up-link encodings, namely FM0, Miller-2, Miller-4, or Miller-8. When using FM0, one bit is transmitted during each cycle and the data rate is equivalent to the link frequency. However, FM0 is highly susceptible to noise and interference which motivated the addition of the Miller encodings. While these are more robust to errors with a higher number being more robust, their link rates are reduced by a factor of 2, 4, or 8, depending on the encoding.

As this study focuses on C1G2 readers configured for use in North America, the readers have a maximum output power of one watt. In accordance with FCC regulations, the readers frequency hop across 50, 500 kHz channels in the 902 - 928 MHz ISM band with a channel occupancy of no greater than 400 ms averaged over a ten second period.

2.2 Class-1 Generation-2 MAC Layer

The MAC protocol for C1G2 systems is based on Framed Slotted Aloha [14], where each frame has a number of slots and each tag will reply in one randomly selected slot per frame. The number of slots in the frame is determined by the reader and can be varied on a per frame basis.

When a reader wishes to read a set of tags, it first powers up and transmits the CW to energize the tags. It then initiates a series of frames, varying the number of slots in each frame to best accommodate the number of tags, as will be described later. After all tags have been read, the reader will power down. We refer to an individual frame as a *Query Round*, and the series of *Query Rounds* between power down periods as a *Query Cycle*.

Figure 1 shows the sequence of commands that make up a *Query Round*. At the beginning of a *Query Round*, the reader

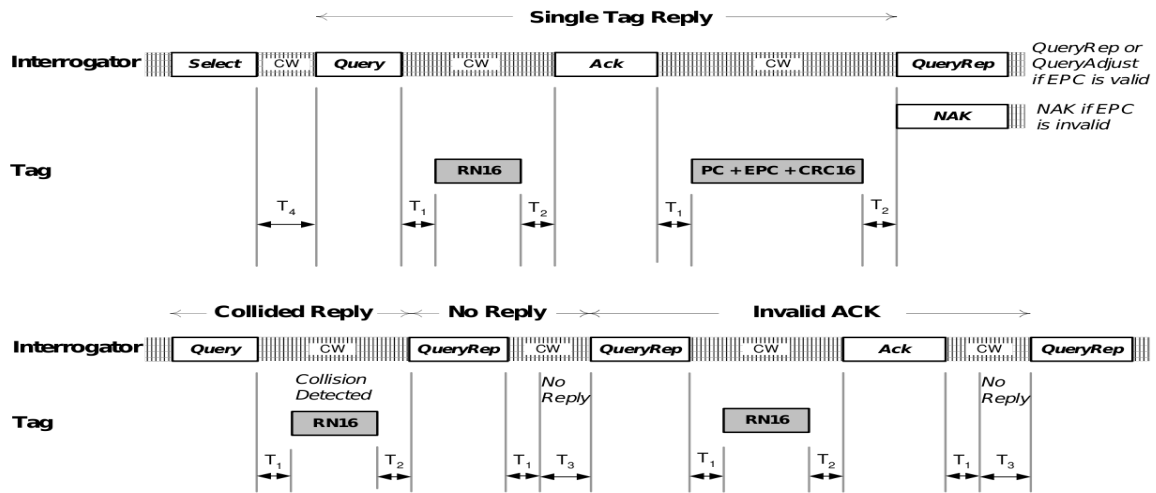


Figure 1: C1G2 Protocol (Courtesy of EPCglobal)

can optionally transmit a *Select* command which limits the number of active tags by providing a bit mask, as only tags with ID's (or memory locations) that match this mask will respond in the subsequent round. A *Query* command is then transmitted which contains the fields that determine the up-link frequency and data encoding, the *Q* parameter which determines the number of slots in the *Query Round*, and a *Target* parameter. When a tag receives a *Query* command, it chooses a random number in the range $(0, 2^Q - 1)$, where *Q* is in the range $(0, 15)$, and the value is stored in the slot counter of the tag.

If a tag stores a 0 in its slot counter, it will immediately transmit a 16 bit random number (*RN16*). Upon receiving the *RN16*, the reader will echo the *RN16* in an *ACK* packet at which point, if the tag successfully receives the *ACK* with the correct random number, the tag will backscatter its ID (referred to as *EPC* in Figure 1). This message sequence is shown in the upper half of Figure 1.

While powered up, tags maintain an *Inventoried* flag which can be in one of two states, A or B. In the *Query* command the *Target* parameter is set to either A or B, and only tags with a matching *Inventoried* flag will respond during the round. After a tag transmits its ID, a subsequent *QueryRepeat* command will cause the tag to toggle its *Inventoried* flag. If the ID was not successfully received by the reader, a *NAK* command is sent which resets the tag so that a subsequent *QueryRepeat* will not result in the flag being changed. In this case, the tag will be active in the next round.¹

Along with toggling the *Inventoried* flag of a tag, a *QueryRepeat* signals the end of the slot. On receiving the command the remaining tags will decrement their slot counter and respond with a *RN16* if their slot counter is set to 0. The process then repeats, with the number of *QueryRepeats* being equal to the number of slots set using the *Q* parameter.

¹The C1G2 standard specifies a timeout period between tag responses and subsequent reader commands. Consequently, even if a *NAK* is lost the tag will be reset by the timeout and will not toggle the *Inventoried* flag at the next *QueryRepeat*.

Of course, tags may choose the same random number initially with the result that multiple tags reply during a slot. In this case, a collision occurs and the tags will not be *ACKed* during the round. However, these tags will be active in the next round, where they will choose a new random number. In the lower half of Figure 1, a scenario is shown with three tags, two of which choose an initial random value of 0 resulting in a collision in the first slot, and one which chooses an initial value of 2 and so responds after the second *QueryRepeat*.

As well as collisions, if there are no tags with a slot counter of 0, a slot may be empty. In order to reduce empty slots and collisions, the reader changes the *Q* parameter for every *Query Round* in order to best accommodate the number of tags remaining to be read. This is done by determining the number of empty slots, slots with collisions, and slots with only a single tag reply during the course of the *Query Round*. Using this information, the appropriate *Q* parameter can be set for the next round. In [16], it is shown that the optimum is seen when the number of slots is set equal to the number of tags, and approximately 35% of the slots would see only a single tag response.

While there are a number of strategies for reading a tag set, a *Query Cycle* generally consists of a series of *Query Rounds* with the target set to A. As tags are read, and their *Inventoried* flags are set to B, they become inactive during subsequent *Query Rounds*. When a *Query* elicits no tag responses, as all tags have set their *Inventoried* flags to B, the *Query Cycle* ends and the reader powers down. When powered down, the *Inventoried* flag of each tag is reset to A, and all tags will consequently be active during the next *Query Cycle*.²

²Note that a reader can be configured such that the state of a tag's *Inventoried* flag will persist across cycles. As this persistence time is defined to be anywhere between 500 ms to 5 seconds, this configuration option is not considered in this study.

3. METHODOLOGY AND TOOLS

The goal of this study is to assess the performance of current RFID systems to understand how well they work, and how well they can be expected to work, in a realistic setting. In particular, we are interested in how interactions between the physical and MAC layers affect reader performance. To this end, we perform a measurement study in order to answer the following questions:

How well do commercial readers perform? Commercial readers generally state performance in terms of tags read per second. However, this metric is affected by tag set size and reader configurations such as up-link encodings. We wish to understand reader performance in more detail.

What protocol factors degrade reader performance? The EPC specification allows for wide variation in physical layer parameters and reader side algorithms such as Q selection. We wish to understand the key factors, such as how up-link encodings trade performance for robustness, and how errors arise and impact the protocol.

What causes tags to be missed during a read? Ideally, readers would have a well defined range at which they could read a collection of tags with perfect reliability, and beyond that range the read rate would drop to zero. Of course, this is not the case. Instead, a reader generally reads some subset of the tags during each cycle even at relatively short distances. We wish to understand why this is the case.

What can be done to improve performance? The EPC-global specification allows for a high level of flexibility. We wish to know whether there are implementation choices that can substantially improve reader performance and reliability.

3.1 Software-radio Based Monitor

To enable this study, we developed a software-radio based RFID monitor using the Universal Software Radio Peripheral (USRP) hardware platform in conjunction with the GNU Radio software toolkit. The USRP is an open-source, general purpose platform for software radio development. It provides an interface between four high speed analog to digital converters (ADCs), four high speed digital to analog converters (DACs), and a USB 2.0 interface. Our monitor consists of the USRP, a 900MHz USRP daughterboard that converts RF signals to and from baseband, and a Linux based host computer running custom software developed using GNU Radio.

Our software uses the GNU Radio toolkit, and consists of custom C++ signal processing blocks. A Python application composes these blocks into a signal processing flow-graph which detects reader transmissions and decodes the individual commands which are then logged by the host. With this, we can log a complete trace of reader commands. In addition to decoding commands, we are able to determine down-link parameters, such as those seen in Table 1, and inter-command timing with a resolution of 125 nanosecondss.

While the USRP sampling rate is 64 MS/s, the USB interface acts as a bottleneck and the signal must be down sampled at the USRP. This results in a maximum effective sampling rate of 8 MS/s, and approximately a 6 MHz band being received at the host. As such, we can only decode reader transmissions when the current hopping channel of the reader is within +/- 3 MHz of the monitor center frequency. As will be seen, this 6 MHz slice is representative of the full set of channels and is sufficient to answer the questions posed above.

3.2 RFID Readers

| Mode | Reader(kbps) | Tag(kHz) | Tag Coding |
|------|--------------|----------|------------|
| HS | 128 | 250 | FM0 |
| STD | 26.7 | 250 | FM0 |
| DR | 26.7 | 250 | Miller-4 |

Table 1: Alien Reader Modes

The primary reader we use to generate traces is the Alien ALR-9800, an EPC Class 1 Gen 2 UHF reader designed for dock door and conveyor belt deployments. This reader was chosen as it is highly configurable and returns inventory results for individual read cycles, in contrast to over a given time window as is the case with many commercial readers. This allows us to analyze the data on a per cycle basis. The reader can be configured to use one of three modes, namely High-Speed (HS), Standard (STD), and Dense Reader Mode (DR), each of which implies different physical layer parameters as show in Table 1. These three modes allow us to clearly assess the affect of down-link bit-rate and up-link encoding schemes.

Along with the Alien reader, we perform a set of experiments using a ThingMagic Mercury5e Development kit. This development kit returns the specific hopping channel on which a tag was read, which enables us to explore performance variation across different channels. Unfortunately, the development kit provides only one physical layer configuration (which is equivalent to DR mode of the Alien reader) and returns results aggregated over a given time window. This makes cycle level analysis difficult.

4. RESULTS

Our experiments are designed to answer the questions posed in Section 3. We use different reader configurations to tease apart the effects of physical layer parameters, and fine-grained measurements to assess MAC overheads. Throughout this study, we consider performance only for a single reader environment.

4.1 Experimental Setup

The experiments were conducted in a standard office setting with cubicles of 42 inch height, using Alien 9460-02 “Omni-Squiggle” tags, which are designed for difficult reading environments and are highly orientation-insensitive. Sixteen such tags were adhered to a sheet of poster board in a 4 x 4 grid, with tags spaced approximately six inches apart. The experiments were performed once in a room approximately 30’ x 22’ x 10’, and replicated on a different floor in a room of approximately 40’ x 24’ x 13’. Throughout the remainder of this paper we will refer to these as *Experiment1* and *Experiment2*, respectively.

For *Experiment1*, the grid was suspended vertically from the ceiling with the center being approximately 7’ above the floor, while in *Experiment2* the grid was set on a tripod style poster stand with the center being approximately 4’ above the floor. In both cases, the height of our Alien ALR-9610 circular antenna was aligned with the center of the grid. This setup results in a clear line of sight between the reader antenna and all tags in the grid. Unless otherwise specified, our experiments are performed using the complete set of 16 tags.

To perform experiments at varying distances we move the grid of tags in two foot increments. We conduct all experiments for a given distance before moving the tags. At each distance, 1000 read cycles are performed in each mode with the monitor antenna placed close to the reader antenna, and both reader and monitor data are logged. For all experiments, the readers are configured such that the *Inventoried* flag of each tag is reset after each cycle. This results in each cycle being an independent data point. The Q parameter is set such that the initial number of slots is optimized for the set size, with the number of slots being equal to the number of tags.

4.2 Reader Protocol Implementation

As the C1G2 specification allows for a wide range of implementation choices at both the physical and MAC layers, we first used our monitor system to determine how the specification was realized by the readers used in this study.

By measuring reader PIE durations and decoding the *Query* commands, we determined the physical layer parameters for the Alien reader as shown in Table 1. Additionally, we found that the ThingMagic reader used settings identical to those of DR mode for the Alien reader. We also determined the duration of the power down periods that take place between each *Query Cycle*, and the *CW* duration that is used to power up the tags before commands are sent. These were seen to be approximately 40 ms and 3 ms respectively.

As the C1G2 specification allows for a number of strategies for reading tags, it was necessary to determine exactly what approach was used by our readers. While the specification suggests a scheme where readers first transit all tags into the 'B' state, and then transit all tags back to the 'A' state, we found that both readers opted for a different approach. In both cases, all *Query Rounds* are performed with the target set to 'A', and when no tags respond in a round the reader powers down, ending the *Query Cycle* and resetting all tags back to the 'A' state. The two strategies are functionally equivalent in a single reader environment where the tag state is reset at power down, but it is interesting that both manufacturers opted to deviate from the suggestions of the C1G2 standard.

To determine the channel hopping rate of the readers, we connected a WISP [15] programmable RFID tag to an oscilloscope. As the WISP harvests a different amount of power depending on the reader frequency, we were able to determine how often the reader changed hopping channels. It was seen that in the case of both readers, a series of *Query Cycles* was performed on a given channel with a new channel being used approximately every 400 ms. This is the maximum channel dwell time per FCC regulations.

4.3 Overall Performance

In order to assess reader performance for different set sizes, we use a *Select* command at the start of every cycle to select a subset of the tags. Figure 2 shows the total number of tag reads per second for HS mode as tag set size and distance increase. The results shown are taken from *Experiment1*. With increasing distance, read rates are reduced as would be expected due to errors and missed tags. However, the degree of scaling with set size is worth more consideration. At two feet, missed tags are rare and the increase in absolute read rate can be attributed largely to inter-cycle overhead. This overhead consists of the time between cycles when the reader is powered down, and the carrier time used to power the tags before beginning com-




Figure 2: Read rates for High-Speed mode.




Figure 3: Cycle times and overhead for HS mode at 2 feet

munication. For our reader these times are 40 ms and 3 ms respectively. This overhead, along with the average cycle duration at a distance of two feet, is shown in Figure 3.

For small tag sets the inter-cycle overhead is significant, though readers have a high degree of flexibility as the EPC specification only requires the reader to be powered down for 1 ms. However, inter-cycle duration must be taken into account when considering overall reader performance, as reading a larger number of tags per cycle amortizes the cost of inter-cycle overhead, resulting in lower per tag read time. This highlights the fact that tag misses should be reduced, as a tag miss degrades the efficiency of each cycle.

In Figure 2, read rates are shown in terms of the total number of tag reads per second, as this is generally the metric used by vendors when stating reader performance. However, for most applications the metric of interest is the rate at which the entire tag set can be read, as reading the same tag multiple times is generally of limited value. Figure 4 shows performance with respect to this metric for each of the three reader modes when reading 16 tags. Results are shown for both *Experiment1* and *Experiment2*.

At two feet, where errors are likely rare for all modes, the overhead associated with each mode can be seen clearly. As the down-link rate of HS mode is higher, and the higher rate




Figure 4: Tag set read-rate for varying coding schemes.

up-link encoding is used, HS mode performs significantly better out to approximately four feet. It can also be seen that both STD and HS modes see similar performance out to four feet, with only a constant difference due to their different down-link rates. At greater distances, the two modes converge when error rate likely becomes a dominant factor. Though both HS and STD mode see better performance in *Experiment2*, the trend is the same with the two modes converging at a slightly greater distance. This is likely because the larger room results in less multipath and hence lower error rates.

In the case of DR mode, the low link speeds result in this mode performing worse than both HS and STD mode at lesser distances, but the overhead incurred by DR mode is outweighed by the benefits at greater distances. For both experiments, the performance of DR mode is approximately the same. This suggests that the difference seen for HS and STD modes is due to errors that are largely mitigated by DR mode.

We have shown that tag set size does affect read rates due to the cost of inter-cycle overhead, and that different reader configurations perform differently depending on the distance from the reader and the RF environment. While Figure 4 gives an overview of reader performance and shows the effect different physical layer parameters can have on read rates, to gain a deeper understanding of *why* these parameters affect performance we must consider their impact on two factors: the duration of each *Query Cycle* and the number of *Query Cycles* needed to read the complete tag set.

4.4 Cycle Duration

As the different modes have different down-link rates and up-link encodings, the duration of reader commands and tag responses and hence cycle duration will differ for different modes. Consequently, there is a trade-off between cycle duration and robustness to error. As inter-cycle overhead is fixed and varies for different reader implementations, it is excluded throughout the remainder of this study. For our purposes, we consider a cycle as beginning with the first reader command after the reader powers up, and ending after the last command before powering down.

Figure 5 shows the distribution of cycle durations seen in *Experiment1* for the different reader modes, normalized for the number of tags read in the cycle. For readability, results

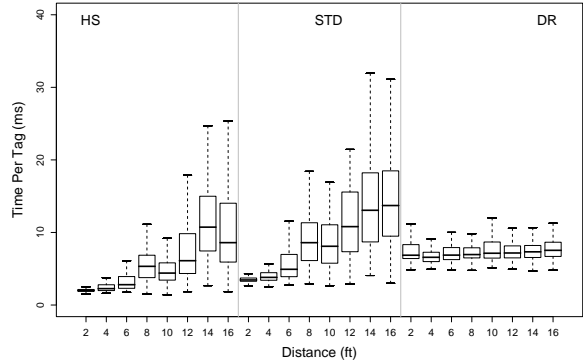


Figure 5: Distribution of cycle time for 16 tags

are shown only up to 16 feet, and outliers have been elided. As can be seen, HS mode has lower average cycle durations than STD mode, but beyond six feet the variance and overall duration increase dramatically in both cases. In contrast, DR mode maintains approximately the same cycle duration at all distances. This suggests that errors become a dominant factor as distance increases, and that the up-link encoding of DR mode effectively mitigates these errors. Additionally, errors appear to be the cause of not only an increase in overall cycle duration, but also in the variance of cycle duration. In the following sections, we will consider error rates directly, and show how errors interact with the protocol to degrade performance.

4.4.1 Error Rates

As mentioned in Section 2.2, after a tag receives a valid *ACK* a subsequent *QueryRepeat* will cause the tag to toggle its inventory flag. If the reader does not successfully receive the ID, as in the case of errors, it must transmit a *NAK* which resets the tag. This results in the inventory flag of the tag not being toggled by the next *QueryRepeat*, and the tag will be active during the next round. Consequently, we can calculate error rates by determining how often *ACKs* are followed by a *NAK* instead of a *QueryRepeat*. In addition, when the Alien reader fails to decode the tag ID it will resend the *ACK*, ostensibly in hopes that the subsequent tag transmission will be error free and the slot will not be wasted. The reader will send the *ACK* up to three times at which point, if it has still failed to decode the ID, the *NAK* will be sent.

Errors can arise either when the tag transmits its *RN16* during the singulation process, or when the tag transmits its ID in response to the reader *ACK*. While a completely garbled *RN16* may be detected by the reader, a bit flip will not as there is no CRC used during singulation. This will result in the reader *ACK* containing an invalid *RN16*, to which no tag will reply. In the case of errors in the ID, these will be detected by the CRC check. While we can accurately detect errors in the traces, we are unable to disambiguate a *RN16* with a bit flipped and an ID with errors.

To calculate error rates from the trace data, we consider each *ACK* immediately followed by a *QueryRepeat* as a success, and every other *ACK* as an error. The error rates for both experiments are shown in Figure 6.




Figure 6: Error rates

The results confirm that error rates are generally lower in *Experiment2*, likely due to the difference in room size. In the case of HS and STD mode, error rates increase rapidly, nearly doubling every two feet up to six feet. Beyond six feet they generally continue to increase, albeit more rapidly in the case of *Experiment1*. For a given experiment, the error rates for HS and STD mode are nearly identical, suggesting that down-link rate does not significantly affect error rate. Interestingly, the error rate does not strictly increase with distance, as position specific multipath can have an effect on error rate.

For DR mode, the more robust up-link encoding reduces the number of failed *ACKs* considerably. For both experiments error rate generally increases slowly with distance. While DR mode significantly reduces errors, it does have a baseline error rate of 10 - 20% in both experiments, and in the case of *Experiment1* the effectiveness of Miller-4 encoding is reduced significantly beyond 16 feet. It is possible that more robust encodings would further reduce the error rate. However, the reader used in this study does not support this.

4.4.2 Effects of Errors

We now consider how errors impact the protocol to lengthen cycle duration. As cycle duration increases due to an increase in the number of commands, the relative duration of each command must be considered along with the rate at which each command becomes more numerous in the presence of errors.

Figure 7 shows the duration of each command, and the amount of time after each command that the reader waits for a tag response. The times shown are for HS mode, but the command times will decrease as down-link rate increases, and wait time will increase proportionally with more robust up-link encodings. *ACKs* are by far the most heavyweight command, largely because the reader waits for the duration of the entire ID, even though *Query* and *QueryRepeats* are more numerous (as there are far more slots than *ACKs*).

Figure 8 shows the average time per cycle accounted for by *ACKs* and slots in *Experiment1*, normalized for the number of tags read in the cycle. As error rate increases, the amount of time spent on both *ACKs* and slots increases considerably, as seen for HS mode. For this mode, at only six feet the time to read each tag has increased three-fold, which results in the majority of the cycle being wasted due to errors.




Figure 7: Duration of each command




Figure 8: *ACK* and slot times per cycle

While slots generally account for less cycle time than *ACKs*, it is unclear why this trend is reversed for HS mode beyond 16 feet. It is possible that this is a result of the *RN16* of the tag having errors, in which case the reader will not send an *ACK* but may be able to detect that a tag replied. Consequently, the reader would not reduce the number of slots in the next round as it would if an empty slot were detected.

We have shown that errors result in greater cycle durations, as they increase the number of *ACKs* and slots, the cost of which quickly outweighs the overhead of more robust encoding schemes. While we cannot show directly the cause of increased variance as errors increase, it is likely due to the fact that a tag which fails to be read during a given *Query Round* will remain active during the next round. Consequently, if a cycle sees a particular tag with a high error rate, the cycle will be extended significantly as each round will incur the cost of additional *ACKs* and a wasted slot.

While the overall time per tag read for HS mode only surpasses that of DR mode at 12 feet, the overall performance of HS mode falls below that of DR mode at only six feet, as shown in Figure 4. In addition, the error rate for DR mode is approximately constant out to 16 feet, but the performance degrades considerably over this range. This suggests that cycle duration is not the only cause of performance degradation.




Figure 9: Number of cycles for different coding schemes

4.5 Number of Cycles

The fact that DR mode error rates are approximately constant, while performance degrades considerably, suggests that tags are being missed and the number of *Query Cycles* is increasing.

Figure 9 shows the average number of cycles needed to read all tags in the set, for each of the three modes in both experiments. For HS and STD modes, the number of cycles increases sharply at six feet in the case of *Experiment1* and 10 feet in the case of *Experiment2*. By referring to Figure 6, it can be seen that this corresponds with the distance at which the error rate approaches 50%. This tells us that error rates above this threshold result in the reader effectively “giving up” during a cycle.

In both experiments, STD mode generally requires fewer cycles to read the tag set. As the error rates for HS and STD mode were seen to be approximately equal, the fact that STD mode requires fewer cycles shows that down-link rate has an effect on the number of tags read per cycle. This is likely due to a lower down-link rate having longer high amplitude pulses and longer cycles, which results in more tags harvesting sufficient power to respond. Additionally, as seen in the case of error rates, the number of cycles does not strictly increase with distance due to position specific multipath effects.

As can be seen, past 12 feet the number of cycles not only increases sharply for HS and STD mode, but also begins increasing rapidly for DR mode. Of particular interest is the fact that in *Experiment1*, between 12 and 16 feet the error rate of DR mode remains constant while the number of cycles increases.

To isolate the cause of increased cycles for DR mode in *Experiment1*, we look at the hit rate of each tag in terms of the fraction of cycles in which the tag was read. This is shown in Figure 10, for all distances beyond 12 feet. At 14 feet, the distance at which cycles increase sharply for DR mode, tag 16 is read during less than 60% of the cycles. However, at 16 feet the same tag is read over 90% of the time while tag two is read during less than 40% of the cycles. This, along with the general variation in hit rate, suggests that multipath effects not only cause missed tags due to an increase in error rate, but also due to some additional effect.




Figure 10: Hit rate for each tag using DR mode




Figure 11: Hit rates for tags at 18 Feet

4.5.1 Effects of Frequency Selective Fading

Multipath in wireless systems results in frequency selective fading, the effects of which are highly location specific. In the case of RFID, frequency selective fading results in tags not harvesting enough energy to power up when positioned such that destructive interference lowers their received energy. This would result in tags being missed even when error rates are low. In addition, the frequency hopping physical layer used by the C1G2 protocol means that the set of tags which are powered up will vary depending on which channel is being used, as multipath effects are frequency specific.

Using the ThingMagic reader, which returns the frequency on which each tag was read, we can determine how well each tag responds to different frequencies. Figure 11 shows the hit rates for each tag at 18 feet when using the ThingMagic reader in the same location and setup as *Experiment1*. As we cannot gather per cycle data for this reader, the results are in terms of the percentage of all tag hits that each tag accounted for over the course of a 15 minute experiment. As both readers use the same link rates and up-link encoding, the robustness to error should be approximately the same, and the data show a similar degree of variation between tags.

To determine the effect of different frequencies, we calculated the percentage of the 50 channels on which each tag re-




Figure 12: Frequency response of tags at 18 feet

sponded at least once over the course of the 15 minute experiment. The results are shown in Figure 12. There is a strong mirroring between Figure 11 and Figure 12, of which the strongest example is tag two which accounts for very few of the total reads and only responded on 6% of the channels. As the ThingMagic reader uses a deterministic hopping pattern and changes channels every 400 ms, all frequencies were visited approximately an equal number of times. While there is a strong correlation between the number of channels a tag responds on and the number of times that tag was read, it does not appear to be the case that a tag either always responds, or never responds, on a given channel.

Figure 13 shows, for each tag, the percentage of total reads accounted for by that tag on each of the 50 channels. Each bar represents one channel, with the channel frequency increasing along the x-axis in each sub-figure. It should be noted that the position of each tag in the figure is identical to the position of that tag in the grid used for our experiments.

While some tags respond well to most frequencies, there are no frequencies on which all tags respond with a high probability. In addition, it is not the case that a tag either always responds, or never responds, for a given channel. For all tags, even tag two, there are a number of frequencies on which the tag will reply, though the probability of response is highly variable. It is possible that this is a function of harvested energy, and variation in cycle time results in a variation in harvested energy at each tag during different cycles on the same frequency. Regardless of the exact effect on the tag, it is apparent that frequency selective fading contributes to tags being missed during a cycle, and consequently the number of cycles needed to read the complete tag set.

4.6 Summary

Overall, we have found that:

- The size of the tag set affects performance, largely because larger tag sets are more efficient with respect to inter-cycle overhead.
- Slower, but more robust, up-link encodings are more effective at greater distances, as the overhead is quickly outweighed by reduced error rates.

- Different multipath environments result in different error rates as distance increases, and these effects are location specific.
- Errors increase both the variance and overall duration of cycles by increasing the number of *ACKs* and the number of slots.
- We observe that *ACKs* as well as *Query* and *QueryRepeat* commands account for a significant amount of overall time; the former because they are long and the latter because they are numerous.
- Errors not only affect cycle duration, they also result in missed tags when a reader “gives up” during a cycle.
- Lower down-link rates result in fewer cycles needed to read the complete tag set, likely because more tags are able to power up.
- Frequency selective fading is a dominant factor in missed reads, particularly at greater distances.

5. ENHANCEMENTS

With a better understanding of the factors that attribute to performance degradation in RFID systems, we now suggest specific ways in which performance could be increased. These enhancements are achieved either by tuning the physical layer operating characteristics of the reader, or through a closer coupling of the physical and MAC layers. The analyses in this section are based on the data from *Experiment 1*.

5.1 Changes at the Physical Layer

We have found that reader performance is largely determined by the duration of each *Query Cycle*, and the number of cycles necessary to read the complete tag set. By tuning the physical layer characteristics of the reader, the impact of both factors could be significantly reduced.

5.1.1 Reducing Slot Times

As the Q algorithm results in many empty slots, even in the best case, having the reader truncate the listen time for empty slots would reduce overall cycle times. While there are practical limitations to how quickly a reader can determine if a slot is in fact empty, the EPC specification bounds the jitter of the tag response to approximately four μs . Hence, a reader should be able to at least partially truncate the listen time for empty slots, particularly for low up-link data-rates.

Figure 14 shows the potential performance gains, normalized for the number of tags read in the cycle, if the reader were to truncate the wait time for empty slots. While the EPC specification defines a minimum wait time after each command, the figure shows the performance gain if the wait time after empty slots was twice this minimum. As can be seen, the overall performance gain increases as the proportion of read time spent on slots increases, as in the presence of errors seen in the STD and HS cases. Even in the DR case, where errors are low, per tag read time could be reduced by approximately 10% using this scheme. In the case of more robust encodings, the benefits of reducing listen times would be greater, and would help outweigh their larger overhead.




Figure 13: Frequency response of each tag at 18 feet




Figure 14: Read time with reduced listen period




Figure 15: Effects of per cycle frequency hopping

5.1.2 Reducing Missed Tags Due to Fading

While multipath effects cannot be avoided and the frequencies at which tags will respond cannot be predicted, the variation in frequency response can be smoothed by channel hopping at a more rapid rate. The FCC requires that frequency hopping systems in the 902-928 MHz ISM band have a dwell time of no more than 400 ms on each of 50 channels, averaged over a 10 second window. Both of the readers used in this study remain on a given channel for approximately 400 ms before hopping. As the set of tags that respond to a given frequency is relatively static, the reader is effectively reading the same subset of tags repeatedly. However, the EPC specification allows hopping on a per cycle basis, which would result in the full set of channels being covered more rapidly.

To show the benefit that per cycle frequency hopping would have, we took the data for per cycle hit rates and randomized the ordering of the cycles. This results in a trace that approximates a system that performs per cycle frequency hopping with a random hopping sequence. The results are shown in Figure 15, for only the DR mode case in order to exclude the effects of errors. The results for the original trace are shown,

along with the results for the average of ten randomized cycle sets. The figure shows that per cycle hopping significantly decreases the number of cycles needed to read the complete tag set, particularly when more tags fail to be read as distance increases. In the best case, per cycle hopping reduced the number of cycles by 30%, with an average gain of 23% after 14 feet which is where effects of selective fading become prominent. As faster hopping could be combined with reduced listen times for empty slots, at 16 feet the strategies we have presented would result in a 37% increase in reader performance if inter-cycle overhead is ignored.

5.2 Physical/MAC Layer Coordination

The flexibility of the C1G2 MAC protocol enables a high degree of coordination between the physical and MAC layers. Specifically, as link-rates and up-link encodings are determined entirely by the *Query* commands which begin each *Query Round*, fine-grained adaptation can be achieved. By changing the physical layer parameters for each *Query Round*, the reader can choose the mix of parameters that best matches the operating environment.




Figure 16: Chance of reading a tag on an ACK retry

5.2.1 Reducing ACKs

As the Alien reader retries *ACKs* three times, we can determine the success rate for each subsequent *ACK* to assess the value of retries. Figure 16 shows the rate at which retry *ACKs* fail, showing the average success rate for the first retry and the second retry for all modes. At nearer distance, or in the case of DR mode, there appears to be some value in retrying *ACKs* a single time, though the chance of success is generally below 15% even for DR mode. However, when using HS or STD mode, the chance of a successful retry is very low beyond six feet, dropping to below 5%. This suggests that retrying *ACKs* even once is likely to have very little benefit when using these modes at larger distances.

Additionally, in almost all cases the chance of success for a second retry is exceedingly low, less than 3% for STD and HS modes and generally lower than 4% for DR mode. Of course, simply not retrying *ACKs* would not reduce the overall number of *ACKs*, as the tag will remain active and the retry *ACK* will effectively end up in the next round. However, as retries generally do not provide a significant benefit, a more appropriate response would be to not waste time on retries, but instead change the physical layer parameters used in the next round.

5.2.2 Hybrid Reader Modes

By changing physical layer parameters on a per *Query* basis, there is potential for hybrid reader modes which balance the benefits of high data rate and robustness to errors. For example, using HS mode to read the subset that it can very quickly, and then changing to DR mode and reading the remainder.

Figure 17 shows the expected performance gains using this hybrid scheme. This analysis assumes that DR mode is able to read the entire tag set in one cycle. The graph is shown out to eight feet, as up to eight feet the number of DR cycles that successfully read the complete set is above 90%, while beyond eight feet the number of cycles increases considerably. The analysis was generated by determining the duration of each cycle and the number of tags that were read in each cycle using HS mode. To this, we add the time that DR mode would take to read the remainder of the tags, based on the average time per tag. We show the original results for HS mode, along with the results for DR mode and our hybrid scheme, assuming both read the complete tag set in a single cycle.




Figure 17: Effects of intra-cycle mode change

As can be seen, combining the positive attributes of HS and DR mode has the potential to increase performance significantly. At six feet, hybrid mode outperforms DR mode by more than 30%. However, as the error rate increases, the cycle time of HS mode will increase beyond that of DR mode. At this point, DR mode will outperform hybrid mode.

Far more subtle schemes than that shown above are possible, where per *Query* adaptation is triggered based on error rates. For instance, a single *Query Round* could be performed using HS mode, with retry *ACKs* being turned off. Based on the rate of error, the reader could make an informed decision as to what up-link encoding to use for the next round to read the remainder of the tags. However, as the result of a given round determines the tags that are active in subsequent rounds, an analysis of such schemes cannot be done based on our data.

While by no means conclusive, the above analysis shows that there is potential for reader strategies that mix link speeds and up-link encodings within a single cycle. In general, per *Query* adaptation could be used in a manner akin to rate adaptation seen in 802.11 networks.

6. RELATED WORK

There have been a number of studies that attempt to improve reader performance by focusing on the C1G2 MAC protocol. In [17], the author describes a technique for choosing optimal frame sizes. In [5], further work is done to refine frame size selection, with a comparison of different *Q* selection algorithms, while in [6] the authors present a Bayesian strategy for dynamically determining frame size. While these studies go far towards increasing performance at the MAC layer, they do not take into account the impact of the physical layer.

In [11], RFID read range is considered by analysis. It is shown that RFID reader systems are limited by the reverse-link and not the forward link as had been suspected. Our study validates this result experimentally, and shows that errors on the reverse-link are a major limiting factor in practice.

To show the effect errors have on the C1G2 protocol, Mitsugi et al. [9] present simulation results which show that bit-errors significantly degrade C1G2 performance. Our results confirm their findings, and additionally show that error rates are highly location dependent and the level of degradation is implementation specific.

In [8] [7], read performance and reliability is evaluated for a Class-1 Generation-1 RFID reader and the effect on performance of different materials and tag arrangements are considered. While the studies show that reliability is relatively poor and highly variable, it does not address the underlying protocol or the effect of physical layer parameters.

Most relevant to our work are the real-world measurement studies of C1G2 tag and reader systems seen in [13] and [1]. The studies show that performance varies widely for different readers, though they lacked the instrumentation to determine what caused these effects. As shown in our study, this is due to physical layer operating parameters having a large effect on overall read performance.

7. CONCLUSION

We have presented what we believe is the first detailed, low-level measurement study of EPC Class-1 Generation-2 UHF reader technology in a real world setting. Through the use of commercial RFID readers and a custom software-radio based monitoring system, we have shown how low level parameter choices affect reader performance, and how physical layer effects interact with the C1G2 MAC protocol.

Our results show that the RFID physical and MAC layers should be considered in conjunction rather than separately as is done at present. Independent of the MAC layer, we find physical layer effects, such as errors and multipath, to be significant factors that degrade the overall performance of commercial readers. These effects increase both the duration of each reader cycle and the number of cycles needed to read all tags in a tag set.

Our analysis suggests that better physical layer implementation choices can improve performance while remaining standards compliant. Specifically, we expect that reducing the listen time for empty slots and increasing the rate of frequency hopping would increase the read-rate for a set of tags by at least one third. Greater gains can be achieved by using the flexibility of the C1G2 standard to coordinate the physical and MAC layers. As one example, our hybrid mode switches encodings to provide both high performance and reliability.

8. REFERENCES

- [1] S. Aror and D. Deavours. Evaluation of the state of passive uhf rfid: An experimental approach. In *IEEE Systems Journal*, volume 1, pages 168–176, 2007.
- [2] A.-I. Center. 860mhz-930mhz class 1 radio frequency identification tag radio frequency and logical communication interface specification candidate recommendation, version 1.0.1. 2004.
- [3] EPCglobal. Epc radio-frequency identity protocols class-1 generation-2 uhf rfid protocol for communications at 860 mhz-960 mhz version 1.0.9. 2005.
- [4] K. Finkenzeller. *RFID Handbook: Fundamentals and Applications in Contactless Smart Cards and Identification*. John Wiley & Sons, Inc., New York, NY, USA, 2003.
- [5] C. Floerkemeier. Transmission control scheme for fast rfid object identification. In *Fourth IEEE International Conference on Pervasive Computing and Communications Workshops*, pages 457–462, 2006.
- [6] C. Floerkemeier. Bayesian transmission strategy for framed aloha based rfid protocols. In *Proceedings of the IEEE International Conference on RFID*, pages 228–235, 2007.
- [7] C. Floerkemeier and M. Lampe. Issues with rfid usage in ubiquitous computing applications. In *Lecture Notes in Computer Science - Pervasive Computing*, volume 3001, pages 188–193, 2004.
- [8] S. Hodges, A. Thorne, H. Mallinson, and C. Floerkemeier. Assessing and optimizing the range of uhf rfid to enable real-world pervasive computing applications. In *Lecture Notes in Computer Science - Pervasive Computing*, volume 4480, pages 280–297, 2007.
- [9] Y. Kawakita and J. Mitsugi. Anti-collision performance of gen2 air protocol in random error communication link. In *International Symposium on Applications and Internet Workshops (SAINTW'06)*, pages 68–71, 2006.
- [10] B. Nath, F. Reynolds, and R. Want. Rfid technology and applications. *Pervasive Computing, IEEE*, 5(1):22–24, Jan.-March 2006.
- [11] P. Nikitin and V. Rao. Performance limitations of passive uhf rfid systems. In *Proceedings of the IEEE Antennas and Propagation Symposium*, pages 1011–1014, 2006.
- [12] M. Philipose, K. Fishkin, M. Perkowitz, D. Patterson, D. Fox, H. Kautz, and D. Hahnel. Inferring activities from interactions with objects. *Pervasive Computing, IEEE*, 3(4):50–57, Oct.-Dec. 2004.
- [13] K. Ramakrishnan and D. Deavours. Performance benchmarks for passive uhf rfid tags. In *Proceedings of the 13th GI/ITG Conference on Measurement, Modeling, and Evaluation of Computer and Communications Systems*, pages 137–154, 2006.
- [14] L. G. Roberts. Aloha packet system with and without slots and capture. *SIGCOMM Comput. Commun. Rev.*, 5(2):28–42, 1975.
- [15] A. P. Sample, D. J. Yeager, P. S. Powledge, and J. R. Smith. Design of a passively-powered, programmable sensing platform for uhf rfid systems. In *Proceedings of IEEE International Conference on RFID*, pages 149–156, 2007.
- [16] F. Schoute. Dynamic frame length aloha. *IEEE Transaction on Communications*, 31(4):565–568, 1983.
- [17] H. Vogt. Efficient object identification with passive rfid tags. *Lecture Notes in Computer Science - Pervasive Computing*, 2414:98–113, 2002.