Progress In Electromagnetics Research B, Vol. 43, 53–72, 2012

PERFORMANCE ENHANCEMENT OF THE RFID EPC GEN2 PROTOCOL BY EXPLOITING COLLISION RECOVERY

D. De Donno[1, *], L. Tarricone[1], L. Catarinucci[1], V. Lakafosis[2], and M. M. Tentzeris[2]

[1]Innovation Engineering Department, University of Salento, Lecce, Italy

[2]School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA

Abstract—Maximizing the Radio Frequency Identification (RFID) performance is one of the main challenges in application domains, such as logistics and supply chain management, where the undesired effect of Tag collisions can significantly degrade the speed of the inventory process. The dominating UHF EPC Class-1 Generation-2 (EPC Gen2) protocol only specifies collision avoidance algorithms but makes no provision for collision resolution. In this paper, performance enhancement of the EPC Gen2 standard exploiting Tag collision recovery is demonstrated, for the first time, in real time with measurements. Three simple and effective approaches to handle successful Tag acknowledgments of recovered collided packets are proposed and implemented on a software-defined Reader and programmable Tags. The attained benefits over the conventional EPC Gen2 MAC scheme are significant: the throughput per time slot is increased by 72% while the overall time required to inventory the Tag population is reduced by 26%. The effectiveness of the proposed approach and the validity of the achieved results are confirmed by the good agreement with simulations reported in the literature.

1. INTRODUCTION

In the effort to automatically identify a vast number of RFID Tags as fast as possible over the inherently broadcast air medium, the undesired effect of communication collisions becomes a more and more

Received 8 June 2012, Accepted 12 August 2012, Scheduled 14 August 2012

* Corresponding author: Danilo De Donno (danilo.dedonno@unisalento.it).

widespread problem. The direct consequences on the overall inventory time and communication throughput, as widely pointed out in the literature [1–3], are potentially so adverse that even partially solving the problem can significantly speedup the operations and significantly improve the performance of many different business sectors, such as airports, warehouses and factories.

The de facto ISO 18000-6C RFID protocol, widely known as the EPC Class-1 Generation-2 ("EPC Gen2" for short hereafter) protocol [4] does not effectively address the problem. This standard specifies collision avoidance algorithms like the dynamic Framed Slotted Aloha (FSA) [5] and the binary tree splitting but makes no provision for collision resolution. In particular, an EPC Gen2 compliant RFID Reader (or interrogator) possesses no capabilities of extracting useful information contained within the collided RFID signals of simultaneous wireless transmissions from more than one Tag; rather, it discards this otherwise exploitable data, renders the communication with the Tags unsuccessful and just requires a retransmission of the Tag packets, incurring inventory delays and a waste of throughput.

As discussed later in the related work section, multi-packet reception, as a result of the implementation of collision recovery techniques, has been a widely investigated topic in the literature. In fact, research groups have demonstrated successful recovery of collided RFID packets under a certain probability with real measurements. Nonetheless, to the best of our knowledge, no extra steps have ever been taken to leverage the coupling of the collision recovery techniques with minor modifications to the EPC Gen2 protocol to realistically improve the RFID communication performance. In this paper, we rely on the multi-packet collision recovery capability and the flexibility provided by a software-defined RFID Reader and programmable Tags to achieve experimentally significant improvements on throughput and inventory time. Essentially, these RFID performance benefits stem from very simple, yet effective changes that we propose in the way successful Tag acknowledgments of recovered collided packets are handled. We demonstrate, for the first time, with a real EPC Gen2 RFID setup the performance boost in the protocol communication that can be drawn by easily and cheaply integrating existing technical contributions to real RFID implementations. Our experimental results agree very well with simulations reported in the literature, thus confirming the validity of the proposed approach.

The paper is organized as follows: Section 2 briefly recalls basics of the EPC Gen2 protocol along with the proposed modifications to exploit Tag collision recovery for multi-Tag acknowledgment; in

Section 3, some details about the adopted equipment are given while in Section 4 the implemented changes to the EPC Gen2 protocol are extensively tested with both real-time measurements and simulations. The related work is discussed in Section 5 and the relevant conclusions are drawn in Section 6.

2. RELATED WORK

RFID multi-Tag signal decoding and collision recovery has been extensively investigated in the literature yielding successful and highly reliable results. As expected, the common aspect of these works is the exploitation of diversity-combining techniques for blind source separation, i.e., the separation of independent sources from a mixed signal without having knowledge of the mixing process. The most classical diversity method has been the use of multiple receiving antennas [7–10]. For instance, Mindikoglu et al. linearly combine the outputs of the elements of an antenna array, model the source signals as Zero Constant Modulus (ZCM) signals in order to remove the mutual interference and test the corresponding ZCM algorithms on synthetic and measured data sets with a simulation setup. Under the single-receive-antenna detection scheme category, Zero Forcing (ZF) and Successive Interference Cancellation (SIC) are used in [6]. Maximum Likelihood (ML) sequence decoders are considered in [11–13].

Although the majority of the aforementioned papers does point out the significant benefits that can stem from implementing the multiple-Tag collision recovery techniques and issuing multiple ACKs to the Tags in a real-world scenario, only few attempt to quantify the total inventory time and throughput improvement in EPC Gen2 via simulation. Specifically, Frey [8] achieves a reduction of the total inventory time on the order of 10% with multiple receive antennas. Kimionis et al. [11] similarly provide simulation results that demonstrate an inventory time reduction on the order of 8–17% with memory-assisted FM0 collided signal detection using a single receive antenna. Angerer et al. [6] identify the theoretical throughput increase of a receiver, which is capable of successfully reading and acknowledging two Tags in the same slot, to be approximately 1.6 times the throughput of a conventional RFID Reader.

To the best of our knowledge, no work has implemented a real setup with an EPC Gen2 Reader and multiple Tags. In this work, not only are we following such an experimental approach, but the reliability of the presented results is corroborated by the agreement with previous simulation works.

3. EPC GEN2 OVERVIEW AND PROPOSED APPROACH

EPC Gen2 [4] is, nowadays, the most widely adopted RFID standard. The energy required for the passive EPC Gen2 Tags to operate is harvested exclusively from the Continuous Wave (CW) transmitted by the Reader. Passive Tags simply backscatter the CW and modulate it by changing their reflection coefficient. The EPC Gen2 interrogation procedure begins with a *Query* command sent by the Reader. This *Query* packet not only configures the uplink communication parameters, such as the encoding scheme (FM0, Miller-2, Miller-4 or Miller-8) and the frequency offset with respect to the CW (40 to 640 KHz) of the Amplitude Shift Keying (ASK) modulation, but also contains the Q value of the FSA MAC protocol. In particular, based on this Q integer value, which ranges from 0 to 15 and is directly tied to the number of Tags within the interrogation zone of the Reader [4], a Tag randomly selects an integer in the range from 0 to 2^{Q-1} as its own slot number and responds with a random 16-bit number, hereafter referred to as *RN16*, in the corresponding time slot. Upon successful reception, the Reader will echo the *RN16* in the following ACK message. If the Tag successfully receives the ACK with the correct (exactly same) *RN16* number, it will finally backscatter its 96-bits ID in the *EPC* message. After all Tags have been read, the Reader will power down.

We refer to an individual frame as an *Inventory Round*, and the series of *Inventory Rounds* between power-down periods as an *Inventory Cycle*. Fig. 1(a) shows an example of successful Reader-Tag handshake in the first slot of an *Inventory Round*. While powered up, Tags maintain a flag, which can be in one of two states, A or B. A field in the *Query* command is set to either A or B, and only Tags with a matching flag will respond during the round. After a Tag transmits its ID, a subsequent *QRep* command will cause the Tag to toggle its flag. If the ID is not successfully received by the Reader, the Tag will not change its flag, thus remaining active in the next round. When Tags choose the same random number, and, as a result, reply in the same time slot, a collision occurs. In this case, the Reader will not ACK the Tags during the current round (Ref. Fig. 1(b)). However, these Tags will be active in the next round, where they will choose a new random slot. For the rest of the paper, we will refer to *Mode 0* as the conventional vanilla Gen2 MAC protocol just described.

As opposed to the Reader collisions, i.e., commands from different Readers reach the same Tag at the same time, which are always destructive events, this is not the case with Tag collisions. Although a

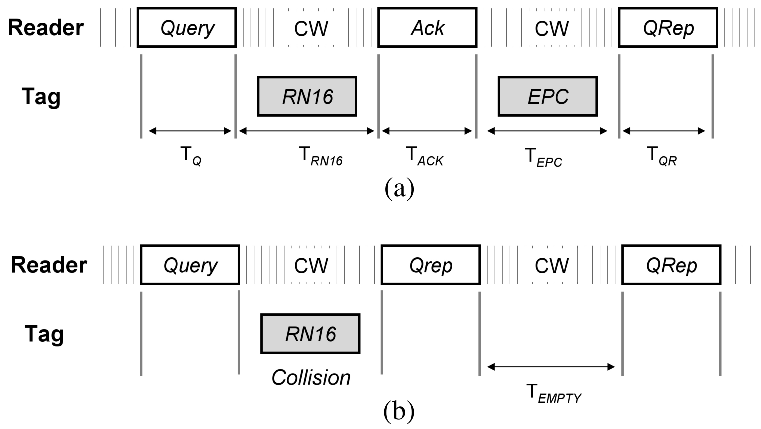

Figure 1. EPC Gen2 protocol in case of (a) single-Tag reply and (b) Tag collision.

typical RFID Reader may only communicate with at most one Tag at a time based on current single-Tag detection techniques, it has been demonstrated that multi-packet reception is practically feasible with various novel ad-hoc algorithms, such as Zero Forcing (ZF) and Successive Interference Cancellation (SIC) [6], Zero Constant Modulus (ZCM) [7], Maximum Likelihood (ML) [11–13], etc.. Hence, it is realistic to assume that on the physical layer the Reader can separate and decode collided Tag signals. The question, however, is how can the Reader use the information extracted from the waveform separation? We refer to the simple case where two Tags transmit during the same slot. Note that in such case the aforementioned techniques for collision recovery achieve excellent performance with, for instance, a probability of 85% to decode both Tag replies reported in [12]. A first possibility for the Reader is to acknowledge only one Tag and discard the other. We refer to this case as *Mode 1*. Alternatively, the Reader can acknowledge both Tags by sending two consecutive *ACK*s (*Mode 2*) or a new command, referred to as *Long ACK* (*L-ACK*), which includes the two decoded *RN16* numbers (*Mode 3*). Fig. 2 illustrates the Reader-Tags communication for each of the three modes. Note that *Mode 3* gives also rise to a collision on the EPC message because both Tags will reply to the *L-ACK*. Such a collision can be recovered with the same technique adopted for the *RN16*. It should be clarified that scenarios involving collisions of more than two Tags are not examined in this work because of the lack of reliable and technically feasible collision-

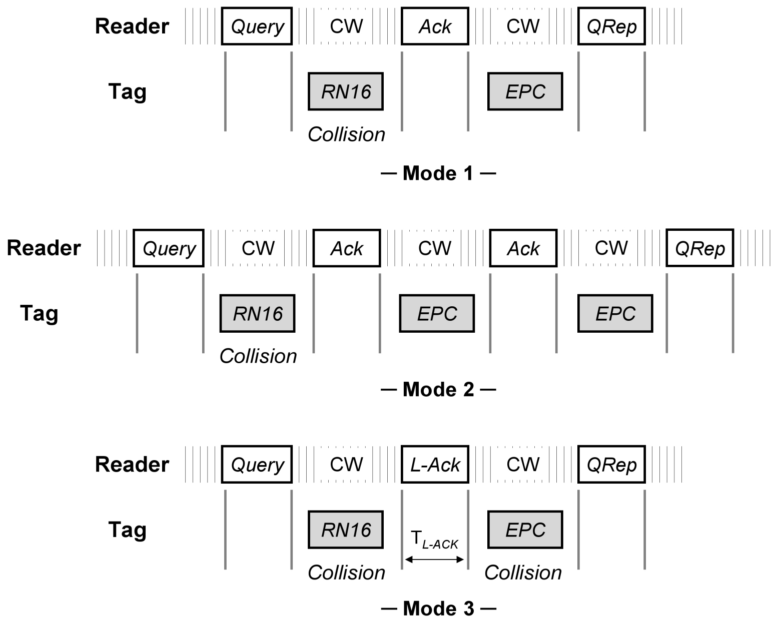

Figure 2. Proposed schemes for the EPC Gen2 MAC protocol in case of two-Tags collision recovery.

recovery techniques that are directly portable to commercial RFID Readers and can operate in real time. Moreover, two-Tag collisions are the most probable events in real scenarios — three times more probable than collisions involving more than two Tags according to the binomial probability for a population of 100 Tags and 100 allocated time slots. In the rest of the paper, we will use the term "collision" to indicate always collision events involving exactly two Tags.

In summary, the two-Tag collision recovery capability allows us to introduce and implement simple (corresponding to only a few lines of code) modifications, namely the aforementioned *Mode 1*, *Mode 2* and *Mode 3*, to the EPC Gen2 standard to acknowledge Tags involved in collisions. The experimental setup used to achieve this goal is introduced in next section.

4. EXPERIMENTAL SETUP

The prototyping of RFID protocols is a very challenging task and can be significantly facilitated by the use of flexible and programmable platforms for both ends involved in the communication: the Tag and the Reader. In this work, we rely on the Intel Wireless Identification and Sensing Platform (WISP) [14] Tags, and on a Software Defined Radio (SDR) implementation of an EPC Gen2-compliant Reader [15, 16].

4.1. WISP RFID Tags

The WISP is a fully passive and programmable RFID Tag developed by Intel Research Seattle. WISP can be powered and read by off-the-shelf UHF RFID Readers and has an on-board microcontroller for sensing and computing functions. The latest firmware version ($hw41\_D41$) that comes with the WISPs is not completely compliant with the EPC Gen2 standard, mainly because of power constraints. For example, the handshake mechanism is not supported and the WISP will always reply to an *ACK* sent by the Reader, whatever the *RN16* contained in it really is. Moreover, as clearly stated in the firmware code, "a pretty aggressive slotting algorithm" is used to preserve power. As a result of this, the slot selection performed by the vanilla version of the WISP code is not genuinely random. In fact, as shown in Fig. 3 in the case of 32 time slots allocated by the Reader, the WISPs tend to pick always the same slots. Since the EPC Gen2 handshake and random slot selection functionality are key features for this work, we implemented them on the WISP. As for the slot selection, a simple and computationally inexpensive pseudorandom

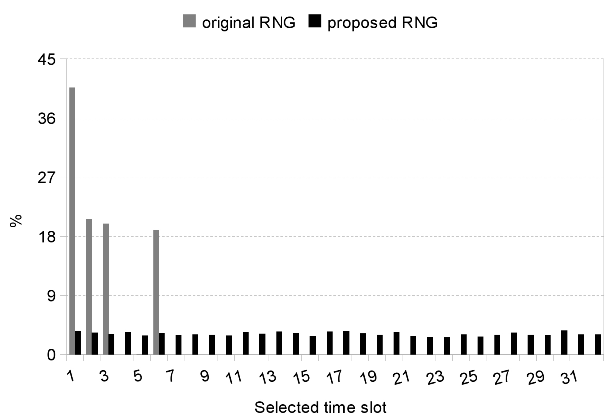

Figure 3. Frequency of time-slot selection by WISP Tags: comparison between the original and the implemented random number generator.

number generator based on a Linear Congruential Generator (LCG) was used. The achieved randomness is fairly good, as shown in Fig. 3. Finally, we programmed the two WISPs used in the experiments with two known *RN16* numbers. In fact, since the purpose of this work is not to propose a new technique for separating and decoding collided Tags but instead to highlight the benefits of such a collision recovery, the transmitted *RN16* numbers will be assumed known at the Reader. This does not cause loss of generality since it is reliable to assume that well-consolidated and effective multi-packet reception techniques will be commonly implemented onto commercially available Readers.

4.2. GNUradio-based RFID Reader

The freely available SDR Reader by Buettner [16] is considered in this work as the basis to implement the collision-aware EPC Gen2 MAC protocol. To our knowledge, this is the first and only cost-effective tool, which allows for easy introduction of changes to the PHY and MAC layer of EPC Gen2. It is based on the low-cost Universal Software Radio Peripheral (USRP) [17] and the open-source GNUradio toolkit [18]. Because the signal processing is completely performed on a standard Linux PC, the SDR Reader enables the modification of MAC and PHY functionalities simply by re-writing user-level software. The effectiveness of GNUradio and the USRP to investigate RFID communication has been demonstrated also in several recent publications [19–22]. The original MAC layer of the SDR Reader has been modified and the three collision-aware *Modes* implemented. For instance, Fig. 4 depicts captured communication between the SDR

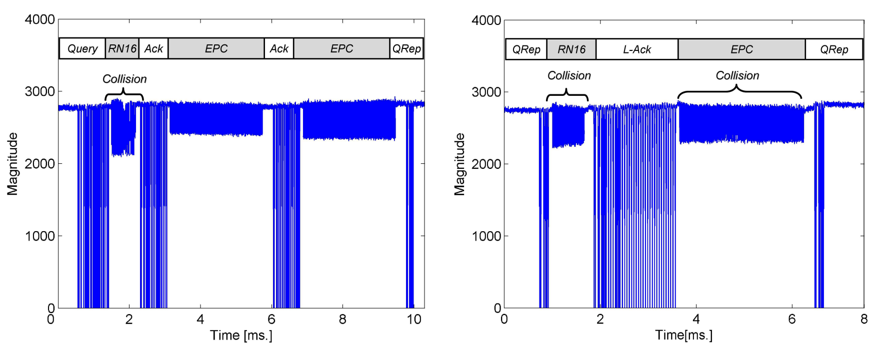

(a) (b)

Figure 4. Example of captured Reader-WISPs communication with (a) *Mode 2* and (b) *Mode 3*.

Reader and two collided WISP Tags when *Mode 2* and *Mode 3* are adopted.

Preliminary experiments revealed that one of the most critical issues for the Reader is to detect a collision event. In other words, for a given time slot the Reader should be able to determine if no Tag has replied (*empty slot*), one Tag has replied (*single-reply slot*), or more than one Tag (exactly two Tags in our case) have replied (*collision slot*). In the last case (*collision slot*), two possible events may occur at the signal-decoding stage of the Reader:

(i) **Preamble not found.** Tag sequences start with a defined preamble. Hence, all tags modulate the same bits at the beginning. However, each Tag exhibits its own backscatter delay, due, for example, to imperfections in the hardware design, thus making preamble recognition difficult at the Reader.

(ii) **Invalid bit sequences.** A collision between Tags in the air interface produces invalid bit sequences at the Reader. The SDR Reader implements Tag decoding using a correlator, i.e., to make a hard decision (bit '0' or '1') according to a correlation score. If the symbols' score is extremely low, the *RN16* packet is marked as *invalid*.

The tweaked SDR Reader denotes a given slot as *collision slot* when either one of the following two situations arise:

(i) the Signal-to-Noise Ratio (SNR) in the slot is above a fixed threshold (this ensures that at least one tag has selected the slot) and a valid preamble is not found, or

(ii) the SNR is above the aforementioned threshold, the preamble is found and the bit sequence is *invalid*.

In both cases, the chosen SNR threshold is fundamental to minimize the number of situations where a *single-reply slot* is erroneously marked as *collision slot*. Based on the analysis of the decoder's performance, we have selected the level of 6 dB as the SNR threshold. The validity of this choice is also confirmed by Buettner's measurements in [15] where more than 95% of *single-reply slots* with errors exhibited an SNR of less than 4 dB, i.e., Tag responses with SNR greater than 4 dB were generally decoded successfully. Therefore, a threshold of 6 dB ensures that when the preamble is not found or the bit sequence is *invalid*, a *collision slot* occurred with high probability.

Once a collision event is detected, we envisage that the Reader will implement a multi-packet reception algorithm to separate and decode the Tags' *RN16* strings in a real situation, as discussed earlier. Without loss of generality, however, the *RN16* values are at this point fixed numbers programmed on the WISPs, so we can consider them known at the Reader. In order to emulate a real scenario, we added the possibility to choose a collision-recovery probability p_{cr} for the Reader. In other words, the Reader will not always acknowledge Tags, based on the running *Mode*, when a collision is detected, but only with a certain probability. This mechanism takes into account that in real scenarios successful waveform separation can fail because of noise and multipath fading.

5. PERFORMANCE EVALUATION

In order to evaluate the gains achieved by the collision-aware EPC Gen2, we run a series of experiments in a real Tag-Reader communication scenario. We fix two WISP Tags on a polystirene sheet at a distance of 1 meter from the SDR Reader antennas — note that the USRP-based Reader uses a bistatic configuration, i.e., one antenna is for transmission and one for reception. The Tags and antennas are mounted on easels 1.5 meters above the ground. It is worth emphasizing that, as demonstrated below, there is essentially no difference in performance between considering a population of two or more Tags.

We instruct the Reader to perform 100 *inventory cycles* with 5 *inventory rounds* each. Recall that if a Tag is not successfully singulated by the Reader during a round (for instance, because of a collision event), the Tag will not change its flag, thus resulting active in the next round. Conversely, a successfully singulated Tag in a round will be inactive in the next rounds. In order to minimize the inventory time, we force the Reader to stop an *inventory cycle* and pass to the next once all Tags has been read. In other words, we assume that the Reader has acquired an accurate estimate of the total number of Tags (two Tags in our case). Such information can be inferred by well-known algorithms proposed in literature, based on deterministic [23], probabilistic [24, 25], or recursive [26] approaches. We repeat the experiments and average out the results provided by the Reader logs for 5 different frequencies in the United States UHF RFID band ($905 + 5i$ MHz with $i \in [0, 1, 2, 3, 4]$) and for 5 different seeds of the pseudo-random number generator implemented on the WISP Tags for the slot selection.

5.1. Throughput Analysis

We consider a population of N Tags and *inventory rounds* consisting of K slots. The Tags randomly select one slot $k \in [1, .., K]$ for transmission. As previously discussed, it may happen that certain slots are either not used (*empty slots*), or used by one Tag (*single-reply slot*) or used by more than one Tag (*collision slots*). The probability of q Tags transmitting in a given time slot is described by the binomial coefficient:

$$P(q)_{N,K} = p_q = \binom{N}{q} \left(\frac{1}{K}\right)^q \left(1 - \frac{1}{K}\right)^{N-q} \tag{1}$$

For *Mode 0*, a successful Tag transmission occurs if exactly one Tag transmits in a slot (*single-reply slot*). In that case, the number of Tag reads per slot, i.e., the throughput, is given by:

$$T_P^{Mode\,0} = P(q=1)_{N,K} = N\left(\frac{1}{K}\right)\left(1 - \frac{1}{K}\right)^{N-1} \tag{2}$$

In the outlined scenario, i.e., when two-Tag collisions can be recovered, the throughput for the three *Modes* is given by:

$$T_P^{Mode\,1} = P(q=1)_{N,K} + p_{CDR} P(q=2)_{N,K} \tag{3}$$

$$T_P^{Mode\,2} = T_P^{Mode\,3} = P(q=1)_{N,K} + 2p_{CDR} P(q=2)_{N,K} \tag{4}$$

where p_{CDR} is the probability that a collision event can be detected (with probability p_{CD}) and recovered (with probability p_{CR}) by the Reader. The achieved throughput per slot measured in our experiments (recall that the Tag population is $N = 2$) is reported in Fig. 5 when 0.9 is chosen as probability of collision recovery at the Reader. Achieved throughput for *Mode 3* is indistinguishable from *Mode 2* and, therefore, it has been omitted for clarity. Fig. 5 contains

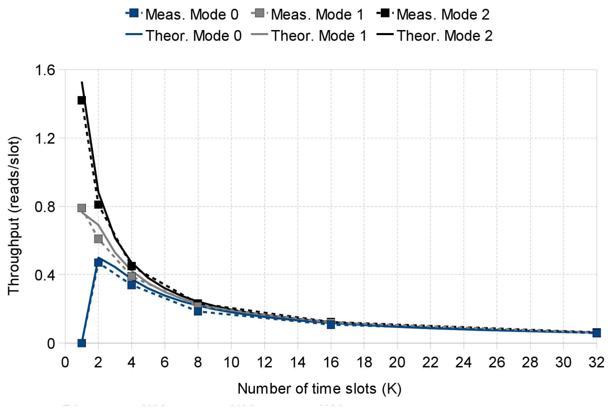

Figure 5. Comparison between theoretical and measured throughput ($N = 2$ Tags).

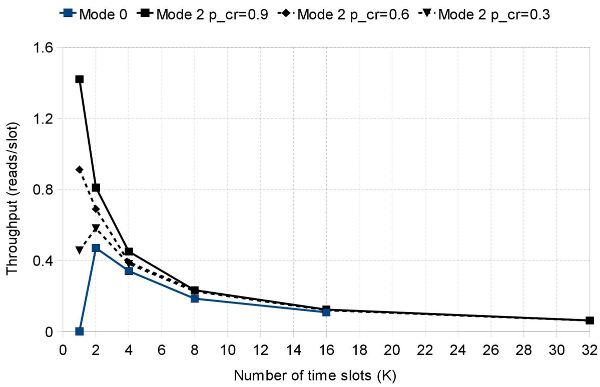

Figure 6. Measured throughput per slot ($N = 2$ tags) with three different values of collision-recovery probability set for the Reader.

also the theoretical throughput calculated by Equations (2), (3) and (4) when p_{CD} and p_{CR} are set to 0.83 and 0.9 respectively. The rationale behind choosing these values is that, regarding the former, a collision event is correctly recognized in 83% of all cases in a set of separate extensive experiments where we forced both Tags to pick the same slot. Regarding the latter, as previously stated, we instructed the reader to implement the collision-aware EPC Gen2 functionalities in 90% of detected collision events. Imposing in (3) and (4) the empirical values of p_{CD} and p_{CR} is the most suitable way to compare the mathematical formulation to the real experiments. Differently from *Mode 0* where the maximum throughput is achieved when $N = K = 2$, the collision-aware *Modes* achieve best performance when $K = 1$ — no read is reported in this case for *Mode 0* because Tags are always forced to pick the same slot. When $K = 2$ slots are allocated by the Reader, *Mode 1* and *Mode 2* attain a throughput gain of 30% and 72% respectively over *Mode 0*. This is in quite good agreement with the simulation results in [6], where a 60% increase of the expected throughput for the equivalent of our *Mode 1* is reported — recall that the probability of collision detection and recovery are not taken into account in [6]. A parametric analysis for three different values of p_{CR} is also conducted. Fig. 6 shows the measured throughput per slot for *Mode 0* and *Mode 2* when 0.3, 0.6, and 0.9 are set as collision-recovery probabilities for the Reader. It is worth highlighting that even when the Reader can separate a waveform with merely 30% of probability, a throughput gain of 23% is achieved. We evaluate the performance improvement, which can be prospectively achieved by the proposed collision-aware

EPC Gen2 when a generic Tag population $N > 2$ is considered. Fig. 7 shows the theoretical throughput per slot when varying the ratio of the number of slots per round K over the Tag population N. Also in this case empirical values of p_{CD} and p_{CR} (0.83 and 0.9 respectively) are imposed for (3) and (4). The theoretical throughput gains achieved by *Mode 1* and *Mode 2* over *Mode 0* when $K = N$ are 38% and 76%, respectively.

5.2. Inventory-time Analysis

Besides increasing the throughput per slot, the proposed modifications to the EPC Gen2 protocol significantly reduce the inventory time, i.e., the time needed for the Reader to read the whole Tag population N. In the following experiments, we assume a fixed number of time slots K allocated in each of the 5 *inventory rounds* that make up an *inventory cycle*. We recall that the Reader knows exactly how many Tags are in the area and stops the inventory cycle once all Tags have been read. Fig. 8 shows the average inventory time measured for each *Mode* of operation when 0.9 is chosen as probability of collision recovery at the Reader. It can be seen that the inventory time gain of *Mode 1*, *Mode 2* and *Mode 3* over the conventional *Mode 0* reduces with an increase in the number of allocated time slots. The reason for this is that the probability of a collision event drops as the number of time slots grows. The achieved inventory-time reduction when $K = N = 2$ is 13%, 19% and 26% respectively for *Mode 1*, *Mode 2* and *Mode 3* over *Mode 0*. These results are in agreement with those provided via

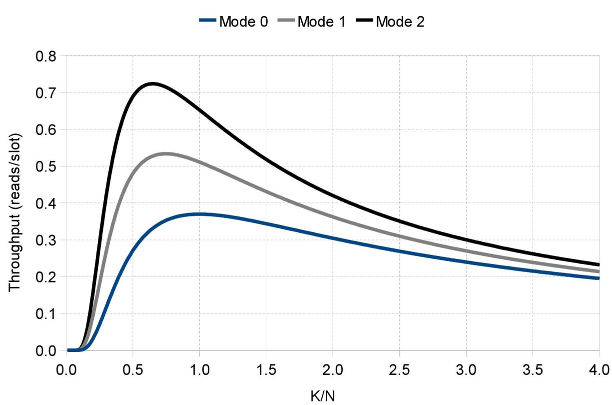

Figure 7. Expected throughput per slot with probabilities of collision detection and recovery set to 0.83 and 0.9, respectively.

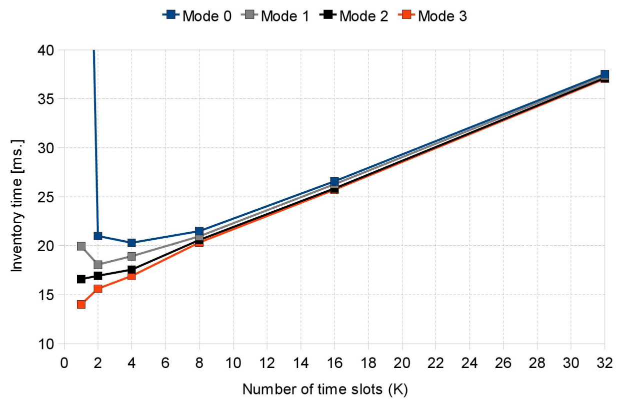

Figure 8. Average inventory time measured at the Reader with Tag population $N = 2$.

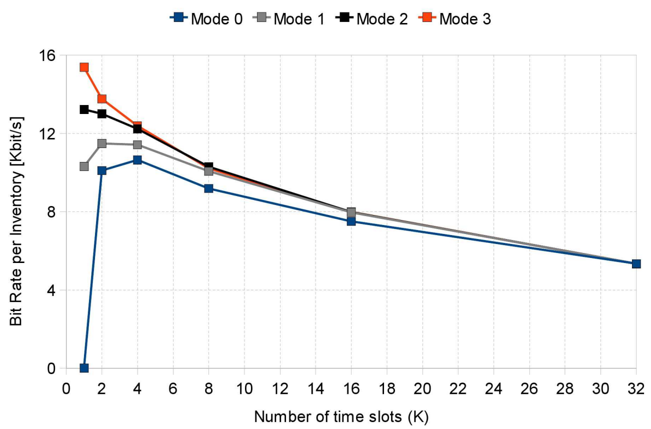

Figure 9. Average bit rate per inventory cycle measured at the Reader with Tag population $N = 2$.

simulation in [8, 11], where the inventory time for the equivalent of our *Mode 1* is reduced by 15% and 10% respectively. It can be seen also that for $K = 1$, *Mode 0* requires an infinite time to complete the inventory because the Tags will always collide. Conversely, the collision-aware schemes carry out the inventory in a finite time. The average bit rate per inventory cycle in the uplink channel (from Tag to Reader) is shown in Fig. 9. It has been calculated using the following formula:

$$\text{Bitrate}^{Mode\,i} = \frac{n_{bit,EPC} N_{read}}{t_{inv}^{Mode\,i}} \tag{5}$$

where $n_{bit,EPC} = 96$ is the number of bits of information contained in the Tag's EPC message, N_{read} the average number of Tag reads per inventory, and $t_{inv}^{Mode\,i}$ the measured inventory time for *Mode i*. In order to evaluate the inventory-time reduction that can be prospectively achieved with a generic Tag population $N > 2$, the following mathematical formulation is considered. For *Mode 0*, which does not explore collision recovery, the average duration of an inventory round can be approximated by the following equation:

$$\begin{aligned} t_{round}^{Mode\,0} &= t_0 + K\left\{p_0 t_{EMPTY} + p_1(t_{RN16} + t_{ACK})\right. \\ &\quad \left. + (1 - p_0 - p_1) t_{RN16})\right\} + K p_1 t_{EPC} \end{aligned} \tag{6}$$

where:

$$t_0 = 2t_{CW} + t_Q + (K - 1)t_{QR} \tag{7}$$

is computed by taking into account the time periods reported in Table 1. t_0 is fixed for each round and comprises the duration of two *CW* periods (one at the beginning and one at the end of the round), a *Query* command and $K - 1$ *QRep* commands. For the proposed collision-aware schemes, the duration of an *inventory round* can be approximated by the following formulas:

$$\begin{aligned} t_{round}^{Mode\,1} &= t_0 + K\left\{p_0 t_{EMPTY} + p_1(t_{RN16} + t_{ACK}) + p_2[p_{CD}(1 - p_{CR})t_{RN16}\right. \\ &\quad + p_{CD}p_{CR}(t_{RN16} + t_{ACK}) + (1 - p_{CD})t_{RN16}] \\ &\quad \left. + (1 - p_0 - p_1 - p_2)t_{RN16}\right\} + K(p_1 + p_{CDR}p_2)t_{EPC} \end{aligned} \tag{8}$$

$$\begin{aligned} t_{round}^{Mode\,2} &= t_0 + K\left\{p_0 t_{EMPTY} + p_1(t_{RN16} + t_{ACK}) + p_2[p_{CD}(1 - p_{CR})t_{RN16}\right. \\ &\quad + p_{CD}p_{CR}(t_{RN16} + 2t_{ACK}) + (1 - p_{CD})t_{RN16}] \\ &\quad \left. + (1 - p_0 - p_1 - p_2)t_{RN16}\right\} + K(p_1 + p_{CDR}p_2)2t_{EPC} \end{aligned} \tag{9}$$

$$\begin{aligned} t_{round}^{Mode\,3} &= t_0 + K\left\{p_0 t_{EMPTY} + p_1(t_{RN16} + t_{ACK}) + p_2[p_{CD}(1 - p_{CR})t_{RN16}\right. \\ &\quad + p_{CD}p_{CR}(t_{RN16} + 2t_{ACK}) + (1 - p_{CD})t_{RN16}] \\ &\quad \left. + (1 - p_0 - p_1 - p_2)t_{RN16}\right\} + K(p_1 + p_{CDR}p_2)t_{EPC} \end{aligned} \tag{10}$$

Since in each round and for each *Mode i* $K \cdot T_P^{Mode\,i}$ Tags are successfully read, we can calculate the total inventory time by:

$$t_{inv}^{Mode\,i} = \frac{N}{K T_P^{Mode\,i}} t_{round}^{Mode\,i} \tag{11}$$

Figure 10 shows the theoretical inventory time for a population of $N = 100$ Tags while Table 2 shows the gain achieved when the collision-aware mechanisms are exploited. Differently from what one can expect,

Table 1. Link timing parameters for the considered scenario.

Reader command	**Formula**	**Duration [ms.]**
CW period	t_{CW}	1.16
Query	t_Q	0.88
ACK	t_{ACK}	0.76
Long *ACK* (*L-Ack*)	t_{L-ACK}	1.52
QRep	t_{QR}	0.21
Power down period	t_{PD}	1.53
Tag reply	**Formula**	**Duration [ms.]**
RN16	t_{RN16}	0.85
EPC	t_{EPC}	3.45
Empty slot	t_{EMPTY}	0.50

Table 2. Inventory-time reduction for $N = 100$ Tags.

	Min. inventory time [ms.]	**Max. gain over *Mode 0***	**Opt. # of timeslots (*K*)**
Mode 0	678.98	–	119
Mode 1	609.25	10.0%	88
Mode 2	570.30	15.5%	76
Mode 3	507.32	26.1%	63

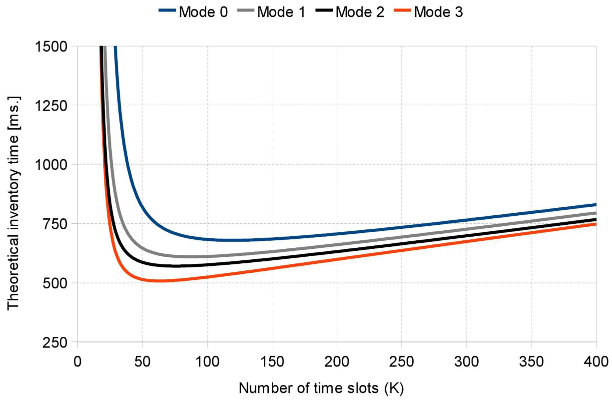

Figure 10. Theoretical inventory time with Tag population $N = 100$.

the minimum inventory time in the conventional case (*Mode 0*) is not achieved when $K = N$. Even if the maximum throughput is obtained for $K = N$ (Ref. Fig. 7), we recall that throughput does not take into account the time required to carry out an inventory but only the number of time slots. Consequently, when $K = N$ a considerable number of collisions still occurs, thus, delaying the inventory process. The theoretical bit rate calculated by (5) is finally shown in Fig. 11.

To validate our experiments, we compare the maximum throughput increase and inventory time reduction in the Tag population cases of $N = 2$ and $N > 2$. To this end, we consider the Equations (2)–(4) and (8)–(10) and calculate for $2 \leq N \leq 500$ the optimum number of time slots K that yields the maximum performance gain for each *Mode i* ($i > 0$) over *Mode 0*. Surprisingly enough, we find

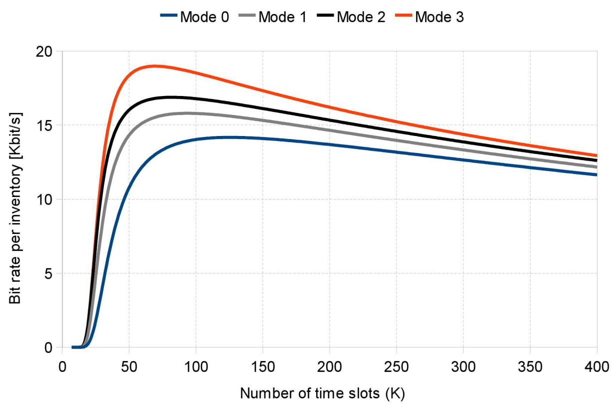

Figure 11. Theoretical bit rate with Tag population $N = 100$.

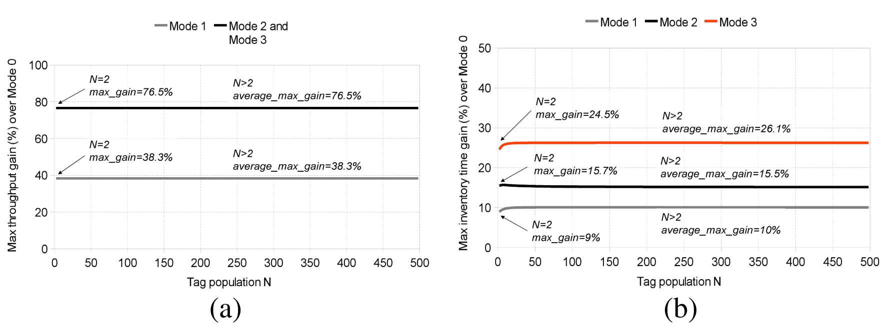

Figure 12. Simulated maximum gain in terms of (a) throughput increase and (b) inventory time reduction for *Modes 1, 2* and *3* over *Mode 0*.

that throughput increase (see Fig. 12(a)) and inventory time reduction (see Fig. 12(b)) are unvarying over the Tag population N. This is a clear indication that our experiment setup, consisting of just $N = 2$ Tags, approximates very closely more complex scenarios with a generic Tag population $N > 2$.

6. CONCLUSION

We have considered a flexible testbed, made up of an open-source Software-Defined Radio (SDR) Reader and the programmable Intel WISP Tags, to implement techniques for multi-Tag acknowledgment under the assumption that collided Tag packets can be successfully recovered. The simple modifications to the EPC Gen2 protocol proposed in this paper experimentally demonstrate that performance of current RFID systems can be considerably enhanced when Tag collision recovery is performed at the Reader. Specifically, the average throughput per time slot is increased by 72% while the overall inventory time is reduced by 26% over the conventional EPC Gen2 MAC scheme, which discards rather than exploits the information contained within collided Tag packets.

Among the advantages that such improved RFID system can bring in real-world applications, we envision the speed increase in conveyor belt and significant gains in power consumption. Our attained results are in very good agreement with those provided via simulation in the literature, thus demonstrating the validity and effectiveness of the proposed approach. To the best of our knowledge, this is the first time that RFID performance with collision recovery is analyzed in real time with actual measurements.

ACKNOWLEDGMENT

The authors would like to thank the Intel Labs Seattle for providing the WISP samples used in this work.

REFERENCES

1. Simon, L., P. Saengudomlert, and U. Ketprom, "Speed adjustment algorithm for an RFID reader and conveyor belt system performing dynamic framed slotted aloha," *Proceedings of 2008 IEEE International Conference on RFID*, April 2008.
2. Singh, J., E. Olsen, K. Vorst, and K. Tripp, "RFID Tag readability issues with palletized loads of consumer goods," *Packaging Technology and Science*, 2009.

3. Lei, X., et al., "Efficient tag identification in mobile RFID systems," *Proceedings of IEEE INFOCOM*, March 2010.
4. EPCglobal, "EPC radio-frequency identify protocols — Class-1 generation-2 UHF RFID protocol for communications at 860 MHz–960 MHz," Version 1.2.0, 2008.
5. Roberts, L. G., "Aloha packet system with and without slots and capture," *SIGCOMM Computer Communications Review*, 1975.
6. Angerer, C., R. Langwieser, and M. Rupp, "RFID reader receivers for physical layer collision recovery," *IEEE Transactions on Communications*, Vol. 58, No. 12, 3526–3537, December 2010.
7. Mindikoglu, A. F. and A.-J. Van Der Veen, "Separation of overlapping RFID signals by antenna arrays," *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2737–2740, 2008.
8. Frey, B., "Source separation for UHF RFID," Master Thesis, Delft and ETH, Mentor Geert Leus, Supervisor Helmut Bolskei, 2008.
9. Zhai, J. and G.-N. Wang, "An anti-collision algorithm using two-functioned estimation for RFID Tags," *International Conference on Computational Science and Its Applications (ICCSA)*, 702–711, Springer, Berlin, Heidelberg, 2005.
10. Nikitin, P. V. and K. V. S. Rao, "Theory and measurement of backscattering from RFID Tags," *IEEE Antennas and Propagation Magazine*, Vol. 48, No. 6, 212–218, 2006.
11. Kimionis, J., A. Bletsas, A. G. Dimitriou, and G. N. Karystinos, "Inventory time reduction in Gen2 with single-antenna separation of FM0 RFID signals," *Proceedings of 2011 IEEE International Conference on RFID*, April 2011.
12. Fyhn, K., R. M. Jacobsen, P. Popovski, A. Scaglione, and T. Larsen, "Multipacket reception of passive UHF RFID Tags: A communication theoretic approach," *IEEE Transactions on Signal Processing*, Vol. 59, No. 9, 4225–4237, 2011.
13. Khasgiwale, R. S., R. U. Adyanthaya, and D. W. Engels, "Extracting information from tag collisions," *Proceedings of 2009 IEEE International Conference on RFID*, April 2009.
14. Sample, A., et al., "Design of an RFID-based battery-free programmable sensing platform," *IEEE Transactions on Instrumentation and Measurement*, 2608–2615, 2008.
15. Buettner, M. and D. Wetherall, "A software radio-based UHF RFID Reader for PHY/MAC experimentation," *Proceedings of 2011 IEEE International Conference on RFID*, April 2011.
16. Michael Buettner, Gen2 RFID Reader project page, https://ww-

w.cgran.org/wiki/Gen2/.
17. Ettus Research LLC, http://www.ettus.com/products/.
18. GNU Radio, http://gnuradio.org/.
19. Catarinucci, L., D. De Donno, M. Guadalaupi, F. Ricciato, and L. Tarricone, "Performance analysis of passive UHF RFID Tags with GNU radio," *IEEE International Symposium on Antennas and Propagation (APSURSI)*, 541–544, Spokane, WA, USA, July 2011.
20. Catarinucci, L., D. De Donno, R. Colella, F. Ricciato, and L. Tarricone, "A cost-effective SDR platform for performance characterization of RFID Tags," *IEEE Transactions on Instrumentation and Measurement*, Vol. 61, No. 4, 903–911, April 2012.
21. De Donno, D., F. Ricciato, L. Catarinucci, A. Coluccia, and L. Tarricone, "Challenge: Towards distributed RFID sensing with software-defined radio," *ACM MobiCom*, September 2010.
22. De Donno, D., F. Ricciato, L. Catarinucci, and L. Tarricone, "Design and applications of a software-defined listener for UHF RFID systems," *IEEE MTT-S International Microwave Symposium Digest*, June 2011.
23. Schoute, F. C., "Dynamic frame length aloha," *IEEE Transactions on Communications*, Vol. 31, No. 4, 565–568, 1983.
24. Vogt, H., "Efficient object identification with passive RFID Tags," *PERVASIVE 2002 — Lecture Notes in Computer Science (LNCS)*, Vol. 2414, 2002.
25. Floerkemeier, C., "Bayesian transmission strategy for framed ALOHA based RFID protocols," *Proceedings of 2007 IEEE International Conference on RFID*, Grapevine, TX, USA, 2007.
26. Eom, J. B. and T. J. Lee, "Accurate tag estimation for dynamic framed slotted ALOHA in RFID systems," *IEEE Communication Letters*, Vol. 14, No. 1, 60–62, 2010.