Bayesian Transmission Strategy for Framed ALOHA based RFID Protocols

Christian Floerkemeier
Auto-ID Labs
Massachusetts Institute of Technology
Cambridge, MA 02139
floerkem@mit.edu

***Abstract*— Transmission control strategies can increase the throughput of the shared wireless channel and thus accelerate the identification of large RFID tag populations. In this paper, we present a Bayesian strategy that minimizes the response time to changes in the number of RFID tags transmitting by updating the tag number estimate after each slot in an ALOHA frame. If the current frame size is no longer considered to be optimal, our control strategy aborts the current frame and triggers the start of a new frame size. The transmission control strategy is evaluated with the help of a scalable RFID simulation engine that implements the ISO 18000-6 C protocol and that supports different pathloss, fading, capture, and tag mobility models. Our evaluation shows that the Bayesian transmission strategy has a higher throughput than other approaches that only update the estimate at the end of the frame. The evaluation also shows that our Bayesian approach outperforms the Q algorithm specified in ISO 18000-6 Part C at the expense of a significant amount of computations.**

I. Introduction

While traditional RFID applications usually feature no more than a single tag in the read range of an RFID reader, this will be different once cases and individual items in supply chain and logistics applications are equipped with RFID tags. A gate reader at a dock door will typically have hundreds of tags in its read range. Since artificially slowing down the loading process or conveyor belt speeds is not desirable from a business perspective, the fast identification of RFID tags is an important issue.

There are a variety of approaches to improving the speed at which RFID tags are identified. Most of them target the physical and medium access control layer of RFID communication protocols. Examples include increased data transfer rates due to more efficient spectrum usage on the physical layer and various so-called anti-collision algorithms that aim to minimize the time it takes to identify all tags in the range. These RFID anti-collision protocols are variants of contention-resolving tree algorithms [4] or ALOHA [1]. In framed ALOHA, which is used in a number of RFID communication protocols [6], [9], [13], the reader begins its interrogation round by announcing the frame size to the tags (cf. Figure 1). Each tag selects one of the available slots at random and transmits a (temporary) identifier. According to [16], the expected throughput U of framed ALOHA with N tags and L slots in a frame is given by:

$$U(N,L) = \frac{N}{L}\left(1 - \frac{1}{L}\right)^{N-1} \quad (1)$$

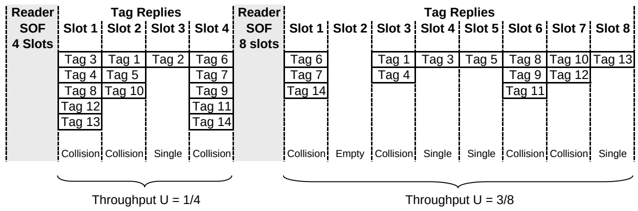

Fig. 1. *Framed Slotted ALOHA.* The reader initiates a frame with a start-of-frame (SOF) signal that broadcasts the frame size. Because of the large proportion of collisions in the first frame, the frame size is increased to 8 slots in the second frame.

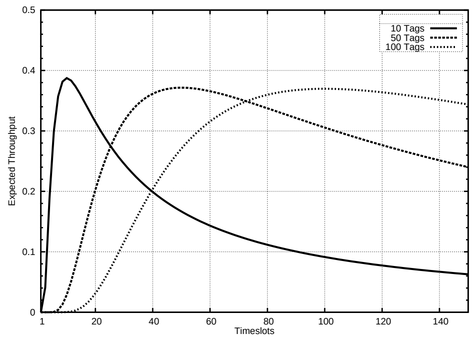

Fig. 2. *Number of Timeslots vs. Expected Throughput in Framed ALOHA.*

It is evident from the above equation that the throughput depends on the appropriate choice of frame length L, given the number of tags N in the read range. Figure 2 shows the well-known upper bound of the throughput of e^{-1} (as N becomes large) that is characteristic for slotted ALOHA and also applies to framed ALOHA. The maximum throughput occurs at $L = N$.

Since the number of tags present is usually not known, the performance of framed ALOHA depends on a transmission scheme that estimates the (unknown) number of tags based on feedback from the reader and chooses a corresponding frame size. This feedback from the reader comprises the number of slots in which no, a single, and more than one tag replied (cf. Figure 1). The latter is referred to as a collision because the data received by the reader are garbled. The number of tags N_t that reply in a frame at time t is given by

$$N_t = \gamma c + s \quad (2)$$

where s and c denote the number of single tag replies and collisions, respectively and $\gamma \geq 2$. The exact value of γ

is usually unknown because the reader cannot detect how many tags replied if there are two or more tag replies. The tags involved in a collision are backlogged and retransmit their identifier in the subsequent frame. The number of tags N_{t+1} that transmit in the subsequent frame is the sum of the backlogged tags that remain powered and newly arriving tags.

$$N_{t+1} = \gamma c - n_D + n_A \tag{3}$$

where n_D and n_A denote the departing tags and newly arriving tags, respectively.

The main contributions of this paper is a Bayesian transmission control strategy for RFID communication protocols, based on framed ALOHA, that minimizes the response time to changes in the number of RFID tags transmitting. It updates the estimate after each slot in an ALOHA frame and, if the current frame size is no longer considered to be optimal, the current frame is aborted and the start of a new frame size with a different number of slots is triggered. The proposed transmission control scheme builds on earlier work by [15] on Bayesian broadcast strategies, but has been adapted to suit the characteristics of RFID. The latter includes the rapid changes in the number of tags transmitting due to power losses as tags move into deep fades and the use of a variant of framed slotted ALOHA, where a running frame can be canceled early. The transmission strategy also makes no assumption about the statistical distribution of the number of tags in the read range. This is important because the number of tags in RFID applications tends to be variable and the traffic is highly correlated rather than caused by many independent point-to-point transmissions – imagine a dock door with the occasional pallet of more than a thousand tagged items moving through. The paper also presents simulation results showing that the proposed transmission control scheme provides a superior throughput when compared to existing approaches. The simulation are carried out with RFIDSim, our RFID simulation engine, which supports different pathloss, fading, capture, and tag mobility models and implements the ISO 18000-6 Part C RFID protocol.

The paper is organized as follows: we first identify the main factors influencing tag estimation and throughput in RFID systems, which include the unusual traffic characteristics of RFID applications and the particularities of the variant of framed ALOHA commonly used in RFID. Based on this analysis, we then present the Bayesian transmission control scheme which specifically take these unique characteristics of RFID systems into account. Before we conclude the paper with an evaluation of our Bayesian transmission scheme, we present related work and our simulation engine RFIDSim.

II. RFID Characteristics Introducing Novel Constraints for RFID MAC Schemes

Before presenting the novel transmission strategy, we describe the characteristics of the RFID domain that introduce novel constraints. We focus in particular on the traffic characteristics, the possibility to interrupt a frame before the last slot of the frame is reached, and the limited available frame sizes.

(a) Portal Reader [12]

(b) Pallet with RFID Tagged Cases [12]

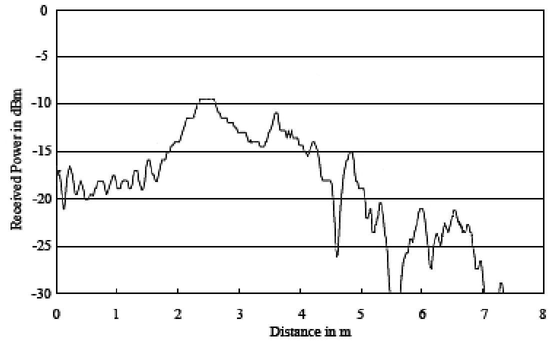

(c) Received Power as Single Tag Passes Reader Antenna [12]

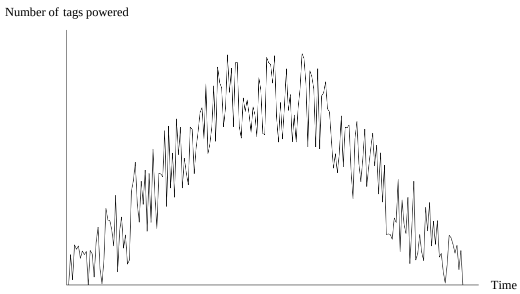

(d) Number of Tags Powered vs. Time

Fig. 3. *Bulk identification in a warehouse application.* The tagged cases are moved through a portal to which an RFID reader antenna is attached. As the tags move past the reader antenna, the received signal strength varies causing the tags to lose power. The exact number of tags powered at any position of the pallet depends on the detailed set-up.

A. Traffic Characteristics

In RFID applications, the tag arrival and departure rates are influenced by application parameters and RFID system design choices. In supply chain operation, UHF readers are often used in a portal configuration as shown in Figure 3(a). Tagged objects are placed on a pallet and moved past the reader antenna. UHF RFID systems are affected by the strong fading component characteristic for indoor wireless channels in this frequency band [14], which leads to frequent field nulls. Since tags do not carry a battery, they will lose power and possibly also their state, as they move past the reader. Figure 3(c) shows the received signal power vs. position for a single tag that is placed inside the pallet [12].

The number of tags that are powered at a certain moment of time thus varies significantly due to the movement of the

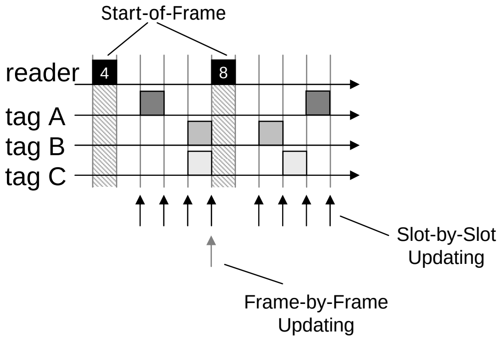

Fig. 4. *Frame-by-Frame and Slot-By-Slot Frame Size Updating.* In our slot-by-slot approach, the estimate of the number of tags transmitting in the frame is updated after each slot.

tags past the reader and the frequent field nulls within the interrogation volume (cf. Figure 3(d)). The actual arrival and departure rates depend on a number of factors, including antenna properties, multipath effects, material properties of the tagged objects, tag orientation, density and speed.

B. Early Cancellation of the Current Frame

Framed ALOHA usually means that acknowledgements are only sent after the end of each frame [16]. This is however not true for the RFID domain. There is usually a reader command after each slot [6], [9], [13]. Transmission control schemes consequently do not have to wait until the end of a frame to change broadcast probabilities by setting the appropriate frame size. The schemes can simply cancel a running frame and initiate a new one.

C. Limited Number of Frame Sizes Available

The RFID domain not only introduces unusual traffic characteristics; the variant of framed ALOHA used in RFID protocols such as [6], [9], [13] also differs from the framed ALOHA commonly described in the networking literature [16] because not all frame sizes are available. To reduce the complexity of the tags, the available frame sizes are limited to powers of two. This results in a reduction in the maximum throughput to 35% from the maximum of $e^{-1} \approx 37\%$ for some tag estimates.

III. Bayesian Transmission Control Strategy

In this section, we present our Bayesian transmission strategy that addresses the characteristics of the RFID domain mentioned in the previous section. The scheme explicitly models medium access in framed ALOHA and computes the probability that a certain number of tags are present based on the feedback from the reader. It makes no restrictive assumption about the probability distribution of the random variable that represents the number of tags powered and ready to transmit. The strategy presented here differs from previous work because the estimated probability distribution of tags present is updated after each slot, while previous work updated the probability distribution after each frame only (cf. Figure 4).

The individual steps of the broadcast scheme are adapted from [15] to suit the nature of framed Aloha and RFID:

1) Compute the frame length L based on the probability distribution of the number of RFID tags transmitting $Pr(N)$.
2) Start frame with L slots and wait for tag replies.
3) Update $Pr(N)$ based on evidence from the reader at the end of each slot.
4) Adjust $Pr(N)$ for tags that are departing during the current frame because they lost power.
5) If frame length L is optimal, given $Pr(N)$, continue with the next slot and go back to step 3. Otherwise, cancel current frame.
6) At the end of a frame (aborted or not), adjust $Pr(N)$ by considering the arrival of "new" tags and the departure of tags that were successfully identified.

A. Computing the Optimum Frame Size

In step 1 of our procedure the optimum frame length is computed, given the probability distribution of N. We choose the frame length L which maximizes the expected throughput U (cf. Eqn. 1).

$$E(U(L)) = \sum_{i=0}^{n_{max}} U(N=i, L) Pr(N=i) \tag{4}$$

This approach is computationally feasible because the available frame sizes are limited to powers of 2 in RFID protocols using framed ALOHA (cf. Section II-C). To reduce the number of computations, one can otherwise also approximate the above by choosing the frame size as:

$$L_{t+1} = 2^{round(\log_2 E(N))} \tag{5}$$

B. Bayesian Updating of the Probability Distribution

In step 3, Bayes' rule is used to update the probability that n tags are replying in the current frame, given all evidence $z_{1:t}$ from previous frames and the evidence $y_{1:j}$ from the first j slots in the current frame:

$$\begin{aligned} Pr(N|y_{1:j}, z_{1:t}) &= \alpha Pr(N|y_{1:j-1}, z_{1:t}) \\ &\quad \cdot Pr(y_j|N, y_{1:j-1}, z_{1:t}) \end{aligned} \tag{6}$$

where α denotes a normalizing constant. Since consecutive frames are considered to be independent given the number of tags transmitting, the following holds

$$Pr(y_j|N, y_{1:j-1}, z_{1:t}) = Pr(y_j|N, y_{1:j-1}) \tag{7}$$

To compute the conditional probability distribution $Pr(y_j|N, y_{1:j-1})$, let us first consider the problem of determining the number of ways T_C to distribute n distinguishable tags into L distinguishable slots $1, 2, 3, ..., L$ with the first c slots containing at least 2 tags, the next s slots containing exactly a single tag, the next h slots with no tag reply, the $j = (c+h+s+1)$ th slot containing at least 2 tags, and the remaining $L-c-h-s-1$ slots containing an unconstrained number of tags. The exponential generating function for $T_C(n,c,h,s,L)$ is given by

$$\begin{aligned} F_C(x) &= \left(\frac{x^2}{2!}+\frac{x^3}{3!}+\frac{x^4}{4!}+...\right)^{c+1} x^s \\ &\cdot \left(1+x+\frac{x^2}{2!}+\frac{x^3}{3!}+\frac{x^4}{4!}+...\right)^{L-c-s-h-1} \\ &= \left(e^x-(1+x)\right)^{c+1} x^s e^{(L-c-s-h-1)x} \end{aligned} \quad (8)$$

Similarly, for the jth slot featuring a single reply, the exponential generating function for $T_S(n,c,h,s,L)$ is given by

$$F_S(x) = \left(e^x-(1+x)\right)^c x^{s+1} e^{(L-c-s-h-1)x} \quad (9)$$

In the event that the jth slot is empty, the exponential generating function for $T_H(n,c,h,s,L)$ is given by

$$F_H(x) = \left(e^x-(1+x)\right)^c x^s e^{(L-c-s-h-1)x} \quad (10)$$

The number of ways T_T to distribute n distinguishable tags into L distinguishable slots $1,2,3,...,L$ with the first c slots containing at least 2 tags, the next s slots containing exactly a single tag, the next h slots with no tag reply, and the remaining $L-c-h-s$ slots containing an unconstrained number of tags, can be computed with the following exponential generating function

$$F_T(x) = \left(e^x-(1+x)\right)^c x^s e^{(L-c-s-h)x} \quad (11)$$

The number of ways of distributing the n tags into the L slots with the above constraints is given by the coefficient of $\frac{x^n}{n!}$ in the expansion of the corresponding generating function $F(x)$. The conditional probability distribution $Pr(y_j = collision|y_{1..j-1}, N)$ that the jth slot is a collision slot given that there were c collisions, h empty slots, and s single occupied slots in the $j-1$ previous slots in the frame is then given by

$$Pr(y_j = collision|y_{1..j-1}, N) = \frac{T_C(n,c,h,s,L)}{T_T(n,c,h,s,L)} \quad (12)$$

Similarly for the conditional probability $Pr(z_j = empty|C,H,S,N)$ and $Pr(z_j = single|C,H,S,N)$,

$$Pr(y_j = empty|y_{1..j-1}, N) = \frac{T_H(n,c,h,s,L)}{T_T(n,c,h,s,L)} \quad (13)$$

$$Pr(y_j = single|y_{1..j-1}, N) = \frac{T_S(n,c,h,s,L)}{T_T(n,c,h,s,L)} \quad (14)$$

Note that we only consider a single arrangement of the c collision, h empty slots, and s single slots when we compute $T(n,c,h,s,L)$. This is feasible because the factor that captures the number of ways the collision, single, and empty slots can be arranged is identical for T_C, T_S, T_H, and T_T.

C. Adjusting $Pr(N)$ for Tags that are Departing During the Current Frame

Because of the frequent field nulls, some tags which received the initial start of frame signal will lose power before they can reply in the slot randomly selected. The result is that the average number of tags that reply decreases as the frame progresses. To take into account this early departure of d transmitting nodes, we simply need to drop the first d entries of the posterior tag distribution in order to compute $Pr(N_{t+1} = n|y_{1:t})$:

$$Pr(N_{t+1} = n|y_{1:t}) = Pr(N_t = (n+d)|y_{1:t}) \quad (15)$$

No new tags can arrive during a frame because tags which power up during the frame will have missed the initial start of frame signal. In practise, it is difficult to find a good estimate for the number of tags that lost power, however.

D. Evaluation of the Current Frame Size

After each slot, we compute the expected throughput in the next slot based on the updated probability distribution of N for the different frame size according to Eqn. 4. If the expected throughput with a different frame size is larger by a certain margin, we cancel the current frame and initiate a new one with the modified size. We use hysteresis to prevent oscillations between two frame sizes that exhibit marginally different performance.

E. Modelling Newly Arriving and Successfully Identified Tags at the End of a Frame

At the end of frame (aborted or not), we still need to incorporate the successful transmissions of the last frame. This only applies to RFID protocols where tags transition to a quiet state after successful identification. Under these circumstances, successful transmissions result in a reduction in the number of tags which reply in the next frame just like tags that departed because they lost power. This means that we simply need to drop the first s entries of the posterior tag distribution as before. The number of tags that transmit their ID in the next frame also changes because new tags arrived. The exact probability distribution of newly arriving, $P_A(n)$, depends on the application characteristics and technology parameters as mentioned earlier. Newly arriving tags can be incorporated by computing the probabilities for N'_{t+1} as:

$$Pr(N'_{t+1} = n) = \sum_{k=0}^{n} Pr(N_{t+1} = k) P_A(n-k) \quad (16)$$

where the conditioning evidence $z_{1:t}$ is omitted.

IV. Related Work

In the previous section, we presented a Bayesian transmission strategy that controls the frame size to optimize the throughput. The scheme addresses in particular the challenges of the RFID domain mentioned earlier. The idea of controlling the ALOHA channel with a transmission strategy is not new, however [11]. Work focussing in particular on controlling an ALOHA channel with an additional frame structure has been carried out by Schoute [16] and Wieselthier [18]. Schoute developed a backlog estimation technique for framed ALOHA which is exact under the assumption that the frame size is chosen in such a way that the number of stations which transmit in each time slot is Poisson distributed with mean 1. The backlog after the current frame B_t is then simply given by:

$$B_t = 2.39c \tag{17}$$

where c is the number of collisions in the current frame. Due to the unknown distribution of the number of tags and limited number of available frame sizes, the assumption made by Schoute leads to deviations between the estimate and the true number of tags present, whenever the above assumption is not valid.

In [18], Wieselthier et al. present a performance evaluation of framed ALOHA with capture. It is based on a combinatorial technique that computes the probability that there are i single occupied slots, j slots with two replies, k slots with three replies, etc., in a frame. Their work is based on the assumption that the number of slots per frame L is fixed and access to the channel can be controlled by adapting the probability with which backlogged stations respond in subsequent frames. In our analysis, the frame size is variable – though limited to powers of two – but all backlogged tags respond in subsequent frames.

More recently, Vogt [17] and Zhen et al. [19] also studied framed ALOHA in the context of RFID. Zhen et al. use the approach proposed by Schoute to estimate the number of tags. Vogt presents a backlog estimation procedure that selects the tag number estimate that minimizes the error between the observed number of empty h, singly-occupied s, and collision slots c and the expected values $E(H)$, $E(S)$, $E(C)$. While this approach does not assume a fixed multiplicity of conflict as the scheme proposed by Schoute does, it only updates the estimate N at the end of a frame. Vogt compares the above estimation algorithm to an estimation strategy that represents a lower bound:

$$B_t = 2c \tag{18}$$

In [7], a Bayesian transmission scheme was presented that does not rely on the Poisson assumption used by Schoute. It estimates the number of tags N present at time t given the feedback from the current frame z_t:

$$Pr(N_t = n|z_{1..t}) = \alpha Pr(N_t, z_{1..t-1}) \cdot Pr(S, C, H|N_t) \tag{19}$$

The algorithm simply chooses the frame size L that maximizes the expected throughput U for the consecutive frame. This scheme differs from the Bayesian scheme presented in this paper because the probability distribution is only updated at the end of the frame.

The Q Algorithm defined in ISO 18000-6 C [6] represents another transmission control strategy. It keeps a representation of the current frame size which is multiplied by a constant β whenever a collision occurs and which is divided by β whenever an empty slot is detected. A successful slot leaves the estimate unchanged. The estimated backlog B_t after the current frame is given by

$$B_t = 2^{Q_{fp}} = 2^{Q_t + a(c-h)} = 2^{Q_t}\beta^{c-h} \tag{20}$$

The new frame size L_{t+1} is then chosen as

$$L_{t+1} = 2^{round(\log_2 B_t)} \tag{21}$$

While the Q algorithm requires only modest computational resources, it does not specify a method to compute the crucial control parameter $\beta = 2^a$. It only provides a range of suitable values ($1.07 \leq \beta \leq 1.41$). This backlog estimation procedure can also be carried out after each slot. This allows the early cancellation of frames with suboptimal frame sizes, whenever the optimum frame length computation of Eqn. 21 indicates this. After Q_{fp} has been updated accordingly, the algorithm compares the updated Q_{fp} against the Q of the current frame size (2^Q). If the current frame size is deemed suboptimal, it is cancelled and a new frame with an optimized length is initiated.

The work by Rivest [15] is also closely related to our work on transmission control strategies for slotted ALOHA. Rivest introduces an elegant pseudo-Bayesian transmission strategy by approximating the probability distribution of the number of stations N with a Poisson distribution with mean v. Each station keeps a copy of v and during each slot, transmits a packet with a probability of $1/v$. It decrements v by 1 if the current slot is empty or a success, increments v by 1.39 if the current slot is a collision, and sets v to max($v+\lambda$), where λ is an estimate of the arrival rate. While our transmission scheme builds on the Bayesian approach outlined in [15], our mathematical model has been adapted to suit framed ALOHA and does not assume that the random variable N denoting the number of tags is Poisson distributed. Due to the large number of arriving and departing tags in RFID applications, the Poisson assumption leads to a slow response to changes in the number of tags present. In [8], Frigon et al. present a pseudo-Bayesian algorithm which extends to framed ALOHA and mixed priorities. It also assumes that the number of tags present follow a Poisson distribution.

Krohn et al. [10] recently presented an approach which presents a fast technique to estimate approximately the number of tags present in the read range. Their work is based on the assumption that there are empty slots and occupied slots only.

V. Evaluation Method

In this section, we present the simulation engine RFIDSim which we developed to compare the Bayesian transmission strategy presented earlier against the schemes outlined in the related work section. The simulation engine implements the ISO 18000-6 Part C (EPCglobal UHF Class 1 Generation 2) protocol and supports different path loss, fading, capture, and tag mobility models.

RFIDSim runs as a collection of simulation entities in the simulation runtime JiST [3]. JiST is a discrete event simulation engine that runs over a standard Java virtual machine. The JiST simulation engine is efficient and performs

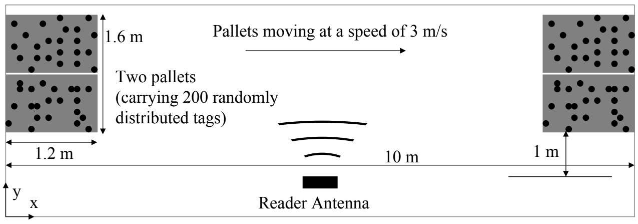

Fig. 5. *Loading dock application.* Pallets containing 200 randomly distributed RFID tags are moved past a single reader at a constant speed of 3 m/s.

Parameter Name	Value
Path Loss Exponent	2
Rician factor	6 dB
Reader power (EIRP)	3300 mW
Reader antenna 3dB beamwidth	60^o
Reader capture ratio	32 dB
Reader sensitivity	-80 dBm
Minimum tag input power	-15.22 dBm
Tag backscatter factor	0.25
BER	10^{-3}

TABLE I

Simulation parameters.

well when compared to existing highly optimized simulation runtimes both in time and memory consumption [3]. The pathloss, fading, and mobility models we implemented build on the Scalable Wireless Ad-hoc Network Simulator [2] that is also implemented on JiST. We also implemented support for different capture models, such as a stochastic model [18] and different power models [14]. The simulation engine supports the command set of the EPCglobal UHF Generation 2 Class 1 Protocol. RFIDSim currently only supports a single reader. While the directivity of the reader antenna can be specified, it is assumed that all tag antennas are isotropic. The simulation is also limited to two dimensions in space. Multipath effects cannot be modelled explicitly, but need to be modelled by statistical fading models, such as Rician or Rayleigh fading.

To compare the different transmission strategies, RFIDSim is used to model the dock door scenario shown in Figure 5, where the movement and identification of two pallets carrying 200 UHF tags is simulated. The multipath fading is modelled statistically by a Rician distribution, which is commonly used to describe the small-scale fading envelope, when there is a strong line-of-sight component [14]. The other simulation parameters are listed in Table I.

VI. Evaluation of the Bayesian Transmission Strategy

To analyze the impact of reducing the response time of the transmission scheme from an entire frame to a few slots only, we compare the Bayesian algorithm presented in this thesis with the Bayesian frame-by-frame scheme presented in [7]. The characteristics of the RFID domain outlined in Section II suggest that a transmission scheme is desirable that can interrupt a running frame, if the current frame size is considered non-optimal. Our analysis is split into two parts. We begin by showing the performance improvement that results from our Bayesian slot-by-slot updating scheme, and continue by comparing our Bayesian algorithm against the Q Algorithm of ISO 18000-6 C (EPCglobal UHF Class 1 Generation 2) protocol [6].

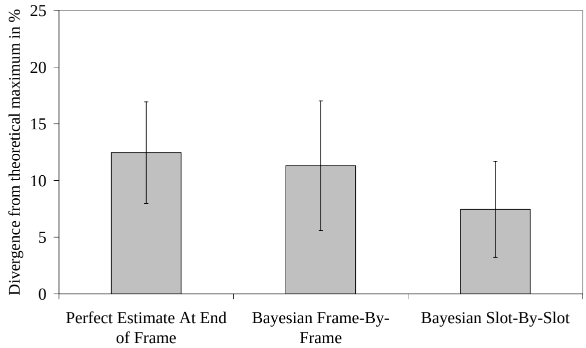

Fig. 6. *Divergence of the simulated throughput from the theoretical maximum of 37% for three different transmission schemes.* The figure illustrates the improvement that can be expected from transmission schemes which operate on a slot-by-slot basis.

A. Comparison with Frame-By-Frame Schemes

We use the simulation scenario outlined in Section V to compare the performance of the Bayesian frame-by-frame of [7] and our slot-by-slot transmission strategy quantitatively in the presence of unknown arrival and departure rates. Two pallets carrying 200 UHF tags are moved past a single antenna and the transmission control strategy changes the frame size accordingly, until all 200 tags are identified.

The two different transmission schemes are compared to what we termed the “perfect estimate at the end of frame” (cf. Figure 6). This transmission scheme knows at the end of each frame exactly how many tags responded in the last frame and chooses the next frame size accordingly. The divergence from the theoretical maximum throughput evident in Figure 6 results from the unknown number of newly arriving and departing tags in the consecutive frame. The simulation results also show that the performance of the Bayesian frame update algorithm presented in [7] is essentially the same as the perfect estimator. Its throughput is 11% below the maximum theoretical throughput of framed ALOHA. The Bayesian scheme that updates the probability distribution of tags transmitting in the frame after every slot achieves the highest throughput (34% on average), which is close to the theoretical maximum of 37%. Due to the limited number of different frame sizes, the actual theoretical maximum will even be slightly less.

B. Comparison with the Q Algorithm

In the previous subsection, we demonstrated the throughput improvement resulting from the use of our Bayesian slot-by-slot technique. In this subsection, we compare the Bayesian slot-by-slot scheme with the Q Algorithm, which is part of the ISO 18000-6 C (EPCglobal UHF Class 1

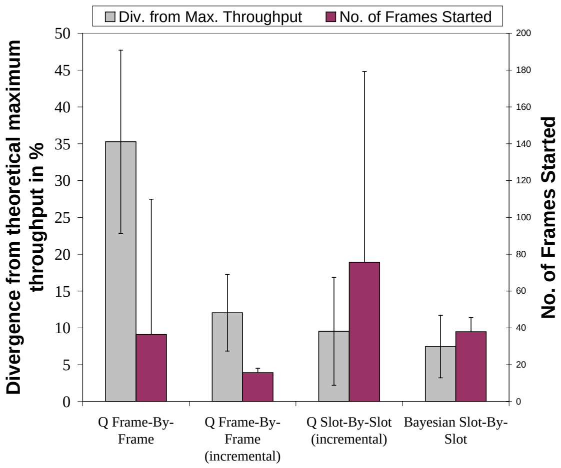

Fig. 7. *Performance of the Q Algorithm, which is part of the ISO 18000-6 C [6] specification, vs. the performance of our Bayesian approach.*

Generation 2) specification [6] and which can also operate on a slot-by-slot basis.

While the Q algorithm requires only modest computational resources, it does not specify the value of the crucial control parameter β, which is used to update the estimated number of RFID tags:

$$\begin{array}{lrl} \text{for a collision slot:} & N_t = & \beta N_{t-1} \\ \text{for an empty slot:} & N_t = & \frac{1}{\beta} N_{t-1} \\ \text{for a single-occupied slot:} & N_t = & N_{t-1} \\ \text{with} & 1.07 \leq & \beta \leq 1.41 \end{array}$$

Since this parameter depends on a past estimate that incorporates past evidence, the evidence from the current slot, and the current frame size, its choice is not trivial. A value of β that is chosen too large will lead to significant overshoots, while a β that is too small will reduce the swiftness of a response to a change. The Q Algorithm also assumes that the frame size is optimal, once the ratio of collision to empty slots is equal to one. However, computing the expected number of collisions and empty slots at $L = N$, which is the criterion for optimum throughput, implies a ratio of collision to empty slots that is smaller than 1.

Our Bayesian slot-by-slot updating algorithm, on the other hand, explicitly models framed ALOHA. It indirectly uses all available information, i.e. all past evidence, including the evidence from the last slot and the current frame size, to compute the multiplicity of conflict. However, the Bayesian algorithm does require significant computing resources, although some of the computations can be precomputed and stored in the memory of the reader device.

In Figure 7, both slot-by-slot transmission schemes are compared. Our implementation of the Q algorithm uses a value for β which is set to $2^{\frac{0.8}{\log_2 L}}$. We found that this approach provides a reasonable performance across different frame sizes. The throughput achieved with the Q algorithm in our simulation scenario is on average 33%, which represents a 9.5% divergence from the maximum throughput of 37%. The Bayesian slot-by-slot algorithm provides only a slightly better performance (Figure 7), when the throughput is considered. However, the overall performance is not only dependent on the throughput, i.e. the number of successful single replies, but also the number of frames started. Each frame that is started results in an additional overhead because of the reader 'Query'-Command that needs to be transmitted at the beginning of each frame. Figure 7 shows that the Q algorithm changes the frame size very frequently in our dock door simulation scenario. This leads to a large overhead which reduces the overall performance significantly.

Figure 7 also shows the performance of the Q Algorithm, if applied after each frame. The figure demonstrates that the performance of the Q algorithm is poor in our scenario, given our choice of the constant β. The performance can, however, be significantly improved when changes to Q are restricted to incremental changes (denoted (incremental) in Figure 7). Under these conditions the oscillations of the Q algorithm are damped and the simulated throughput is similar to the other frame-based transmission schemes.

VII. Discussion

The increased throughput that can be achieved with our transmission scheme comes at the expense of a significant amount of computations. Alternative transmission schemes make certain assumptions about the distribution of the number of tags present or simply assume a fixed multiplicity of conflict. This reduces the resources required to estimate the number of tags transmitting and to choose the frame size accordingly. Our Bayesian approach permits some time-consuming computations to be made a-priori though, e.g. the conditional probability distribution defined in Section III-B. As the number of tags in the range is increased, the memory storage capacity increases significantly, however.

To evaluate the different schemes, we exclusively relied on simulations. Our approaches have thus not been validated experimentally with a UHF RFID system. Future work should thus aim to implement the transmission schemes presented in this paper in a UHF RFID reader. The simulation engine could also be upgraded to include tag antenna directivity and multiple reader antennas.

In our evaluations, we assumed that a reader can operate independently in a given channel. In practice, there will be other readers operating in close vicinity, which will possibly interfere [5]. Future work might thus also consider the effect of such reader collisions on the performance of the transmission schemes. Furthermore, a transmission scheme that chooses an appropriate frame size for a number of readers which are synchronized in order to deal with the limitations of the listen-before-talk schemes introduced in some countries could be part of future investigations.

Our Bayesian model assumes a feedback model where the reader can successfully distinguish between no, a single, a single, but corrupted tag reply, and more than a single tag reply (the corrupted tag replies were not mentioned explicitly before for simplicity's sake). In practice, it might be difficult to always successfully distinguish a corrupted single tag reply from a collision where more than a single tag replied. Future versions of our transmission schemes should thus include the possibility and ideally the likelihood of such a wrong classification. Due to the large capture ratio set in Table I, we observed no capture in our simulations. RFID communication protocols can benefit from this near-far effect because the throughput can exceed the theoretical maximum throughput associated with slotted and framed ALOHA e^{-1}. In our transmission scheme, the occurrence of captured replies would result in an underestimate of the number of tags because the captured reply would be interpreted as a single reply and not a collision. The influence of this misinterpretation is small, however, since the underestimate is compensated by not including the capture effect in the computation of the expected throughput, which would mandate choosing a frame size that is not equal, but smaller than the number of tags present.

In our model, we also assume that all slots have the same fixed length. In practice, a reader might close an empty or even collision slot early. The cost of an empty or collision slot might thus vary. Future versions of the above transmission strategies might anticipate this possibility and explicitly model the early closure of empty and collision slots.

The throughput measured in our simulations refers to the contention slots only. In RFID protocols, such as ISO 18000-6 C, which use a reservation system the overall throughput is much higher, since framed ALOHA is only used as the contention mode for short packets that reserve longer noncontending slots for the transmission of the unique tag ID. The overall throughput achieved with our Bayesian transmission scheme is then as high as 73% under the assumption that the length of the ALOHA slots is about 1/8 of the reserved slots.

VIII. Conclusion

As the number of objects which are equipped with RFID tags increases, it is becoming increasingly important to identify large tag populations quickly. This mandates among other things a high throughput over the shared radio channel. The throughput performance of RFID medium access protocols, such as framed ALOHA, depends, however, on a transmission scheme that estimates the (unknown) number of stations transmitting. The number of RFID tags transmitting remains uncertain, since RFID readers cannot detect the multiplicity of conflict if more than two RFID tags transmit simultaneously.

In this paper, we showed that a Bayesian transmission scheme that updates the tag number estimate after each slot of an ALOHA frame outperforms those transmission schemes that only update the estimate at the end of the ALOHA frame. The comparison with ‘perfect estimator’ shows that this is due to the unknown number of tags arriving and departing during an ALOHA frame. Our Bayesian transmission scheme minimizes the response time to any changes in the number of tags transmitting, since the estimate is updated after every slot.

Our simulation results also indicate that the Q algorithm, which is part of the ISO 18000-6 C, leads to serious overshoots in the tag number estimate and thus a low throughput, unless frame size changes are restricted to incremental modifications. When used on a slot-by-slot basis, the Q algorithm is outperformed by our Bayesian scheme in terms of overall throughput. The improved performance of our Bayesian comes at the expense of an extensive amount of computations though, which makes it unrealistic to implement the scheme in its current form on a resource-constrained RFID reader platform.

References

[1] Norman Abramson. The ALOHA System–Another alternative for computer communcations. In *Proc. 1970 Fall Joint Computer Conference*, volume 37, pages 281–285, 1970.

[2] Rimon Barr. *SWANS - Scalable Wireless Ad hoc Network Simulator, User Guide*, 2004.

[3] Rimon Barr, Zygmunt J. Haas, and Robbert van Renesse. JiST: An efficient approach to simulation using virtual machines. *Software - Practise and Experience*, 35(6):539–576, May 2005.

[4] John I. Capetanakis. Tree Algorithms for Packet Broadcast Channels. *IEEE Transactions on Information Theory*, IT-25(5):505–515, September 1979.

[5] Dan Engels. The Reader Collision Problem. Technical Report MIT-AUTOID-WH-007, Auto-ID Center, 2001.

[6] EPCglobal. Class 1 Generation 2 UHF Air Interface Protocol Standard Version 1.0.9, 2005.

[7] Christian Floerkemeier. Transmission control scheme for fast RFID object identification. In *Proceedings of the Pervasive Wireless Networking Workshop at IEEE PERCOM 2006*, Pisa, Italy, 2006.

[8] Jean-Francois Frigon and Victor C. M. Leung. A Pseudo-Bayesian ALOHA algorithm with mixed priorities. *Wireless Networks*, 7(1):55–63, 2001.

[9] International Organization for Standardization. Information technology – Radio frequency identification for item management – Part 6: Parameters for air interface communications at 860 MHz to 960 MHz, 2004.

[10] Albert Krohn, Tobias Zimmer, Michael Beigl, and Christian Decker. Collaborative Sensing in a Retail Store Using Synchronous Distributed Jam Signalling. In Hans W. Gellersen, Roy Want, and Albrecht Schmidt, editors, *Third International Conference, PERVASIVE 2005*, number 3468 in Lecture Notes in Computer Science (LNCS), pages 237–254, Munich, May 2005. Springer-Verlag.

[11] S.S. Lam and L. Kleinrock. Packet switching in a multi-access broadcast channel: Dynamic control procedures. *IEEE Trans. Commun.*, COM-23, 1975.

[12] Jin Mitsugi and Hisakazu Hada. Experimental Study on UHF passive RFID Readability Degradation. In *Proceedings of the RFID and the Extended Network Workshop at SAINT 2006*, pages 52–55, Phoenix, AZ, USA, Jan 2006.

[13] Philips Semiconductor. I-CODE1 Label ICs Protocol Air Interface Datasheet, January 2005.

[14] Theodore S. Rappaport. *Wireless Communications*. Prentice Hall PTR, second edition, 2002.

[15] Ronald L. Rivest. Network Control by Bayesian Broadcast. *IEEE Transactions on Information Theory*, IT-33(3):323–328, May 1987.

[16] Frits C. Schoute. Dynamic Frame Length ALOHA. *IEEE Transactions on Communications*, COM-31(4):565–568, Apr 1983.

[17] H. Vogt. Efficient Object Identification with Passive RFID Tags. In F. Mattern and M. Naghshineh, editors, *First International Conference, PERVASIVE 2002*, volume 2414 of *Lecture Notes in Computer Science (LNCS)*, pages 98–113, Zurich, Switzerland, August 2002. Springer-Verlag.

[18] Jeffrey E. Wieselthier, Anthony Ephremides, and Larry A. Michaels. An Exact Analysis and Performance Evaluation of Framed ALOHA with Capture. *IEEE Transactions on Communications*, COM-37(2):125–137, 1989.

[19] Bin Zhen, Mamoru Kobayashi, and Masashi Shimizu. Framed ALOHA for Multiple RFID Objects Identification. *IEICE Transaction on Communications*, E88-B(3):991–999, March 2005.