1

Energy-Conserving Access Protocols for Identification Networks

Imrich Chlamtac[a], Chiara Petrioli[b], Jason Redi[c]

[a] The University of Texas at Dallas, Richardson, TX, U.S.A., {*chlamtac,radha*}@*utdallas.edu*

[b] Università di Roma "La Sapienza", Italy, *petrioli@dsi.uniroma1.it*

[c] Boston University, Boston, MA, U.S.A., *redi@bu.edu*

Abstract

A myriad of applications are emerging in which energy conservation is a critical system parameter for communications. Radio Frequency Identification Devices (RFID) networks, smart cards, and even mobile computing devices in general, all need to conserve energy. In RFID systems nodes are small, battery-operated, inexpensive devices with radio receiving/transmitting and processing capabilities, integrated into the size of an ID card or smaller. These identification devices are designed for extremely low-cost, large scale applications such that the replacement of batteries is not feasible. This imposes a critical energy constraint on the communications (access) protocols used in these systems, so that the total time a node needs to be active for transmitting or receiving information should be minimized. Among existing protocols, classical random access protocols are not energy conserving while deterministic protocols lead to unacceptable delays. This paper deals with designing communications protocols with energy constraint, in which the number of time slots in which tags need to be in the active state is minimized, while the access delay meets the applications constraints. We propose three classes of protocols which combine the fairness of random access protocols with low energy requirements.

1 Introduction

Radio Frequency and Infrared Identification Devices (RFIDs and IRIDs), such as warehouse identification tags and intelligent ID cards, are examples of a new world of applications which use small, inexpensive devices for which battery conservation is a critical system parameter. For reasons of conciseness we shall refer to all of these devices simply as "tags" and to the connecting networks as People/Item IDentification NETworks - IDNETs (also referred to in literature as RFID systems).

A typical IDNET is composed of a number of interconnected base stations communicating over a shared wireless channel to a large number of small, low-cost, wireless nodes or tags. These tags usually contain some sort of microprocessor, power source in the form of a battery, capacitor or solar cell, a radio frequency receiver, possibly a transmitter, and some support logic.

The range of potential uses for RFID tags is extremely large. It is estimated that the RFID tag market will expand 20 fold to over $5 billion within the next five years [4].

A few examples of current uses for RFID tags are:

- Tags attached to the ears or worn around the necks of livestock. These tags allow for location tracking of the animals as well as specialized feeding, milking or medication schedules. The animals can be granted or denied access to specific areas based on the tags they carry.

- Smart tags used in warehouses to track inventory. They allow companies to virtually immediately know the location of any item in the warehouse as well as track boxes as they enter or leave the building.

- Numerous tag companies targeting the retail market. Prices displayed by electronic shelf labels can be automatically adjusted based on current market prices or the amount of excess inventory.

There are four fundamental characteristics of RFID wireless nodes which make RFID networks distinct from other wireless systems:

1. *Scale*: There will be potentially a very large number of wireless nodes per base station. E.g. in a warehouse application, there may be thousands of RFID nodes per base station.

2. *Cost*: Nodes must be remarkably inexpensive. Due to the number of nodes and the low cost of some of the "tagged" items, nodes cannot cost more than a few dollars, sometimes even less than a dollar.

3. *Size*: Nodes must be very small. The size of a pack of cards will be the maximum size for many applications.

4. *Traffic*: Communication is based on typically short, simple messages.

From these characteristics follow a number of important observations which further define the unique constraints of an RFID network. Most importantly, due to the large number of nodes, it is economically impossible to replace or recharge the batteries in the tags. This can be attributed to either the cost of the design or the sheer number of tags. Considering that battery life is not expected to increase more than 30% in the near future [9], the battery's energy is therefore a limited and scarce resource.

Radio frequency communication can be highly demanding of energy. A consequence of this is a limited uplink capability. Uplink transmission can typically use twice as much energy as reception [3]. It is possible to use spread spectrum modulation such that the base station indirectly provides the energy for limited uplink communication, but this requires that any uplink traffic be base station initiated. Furthermore, limited unlicensed bandwidth and simplicity of the tag means that all tags must share the same broadcast band.

For all of the above considerations, these types of systems require new access protocols which are designed around these unique constraints and provide a combination of two important factors: low delay and low energy requirements.

The allowable delay is an application dependent constraint. For example, tracking the movement of tags across the cells within a system requires updates to be performed within a short, bounded, amount of time. The system is already constrained by the speed of the shared channel and has to manage a potentially large number of tags. Therefore, it is very important for access protocols to not add significantly to the transmission delay.

Low energy consumption is the second requirement that the access protocol must satisfy. Since it is not possible to replace batteries, it is necessary to design RFID nodes that require a minimum of energy to operate. The tags need to be in the *awake* state for communication. In order to conserve battery life, the tag can enter a *sleep* state where the CPU is in a low power mode and radio reception is disabled. The ratio of energy consumed between the sleep and awake states (i.e., when the CPU operates at full energy) is typically on the order of 100 or more.

Current access protocols, such as those described in [2], have not been designed to meet the dual requirements of low energy consumption and low delay and are consequently not appropriate for tag systems. Random access protocols, such as Aloha, applied to the tag network system translate to base stations sending packets at random times and tags awaking at random times. The probability is very low of a tag being awake in the same slot in which the base station is transmitting to it. Therefore, due to repeated transmission attempts, the energy required and the packet delay will be quite high. Deterministic protocols, such as classical TDMA, assign each tag an individual slot in which it may receive transmissions. Although this has the advantage of a low energy requirement since each tag only needs to be awake $\frac{1}{N}$ slots, where N is the number of tags in the system, in a situation with a large number of tags and very low transmission speeds, this access protocol will take a prohibitively long time to deliver each packet.

To meet the combined energy/delay objective it is easy to see that a fundamental trade-off exists between tags waking up too frequently, leading to a high energy consumption and implying a short lifetime, and a sparse wake-up schedule, which although conserves energy, leads to unacceptably high access delays and therefore inadequate mobility and other application support.

This paper deals with the design of *communications protocols with energy constraint,* in which the fraction of timeslots in which tags need to be in the active (awake) state is minimized, and the access delay meets the applications constraints. In attempting to combine the fairness from random access protocols with low energy requirements from classical TDMA, while maintaining acceptable access delays, we present three different protocol approaches to this problem.

In the remainder of the paper we first describe the tag network model, followed by a detailed description of the protocols. In the ensuing analysis of energy consumption and access delay, we derive the system behavior for both uniformly and non-uniformly distributed traffic destinations. We follow with a detailed evaluation of the RFID protocols and conclude with a summary.

2 Network Model and Protocols Description

We consider a single cell system where a base station communicates with N tags through a radio channel of bandwidth B. The communication is packet-oriented. We assume the time to be slotted and the base station's transmissions to be synchronized to the beginnings of slots. The packet length c is constant, and exactly one packet can be transmitted during each slot. In this analytic model, we do not explicitly treat transmission errors.

We define an *access protocol* as consisting of two components: a *transmission scheduling strategy* at the base station which in each slot selects a packet for transmission from the arrival queue, and a *wake-up schedule* at each tag which determines the slots in which the tag is awake. In general, the transmission scheduling strategy can take into account different parameters: the number of packets in the queue, the packets' ages, as well as the wake-up schedules of their destinations. In the protocols discussed here the "oldest packet" criterion is generally adopted to help meet the application delay requirements. We next present and compare the following three classes of protocols for constructing efficient wake-up schedules:

- *Grouped-tag TDMA protocols*

- *Directory protocols*

- *Pseudo-random protocols*

2.1 Grouped-Tag TDMA Protocols

Classical TDMA can be adapted for use in IDNETs in the following way: we divide tags into $\left\lfloor \frac{N}{x} \right\rfloor$ disjoint groups, with the cardinality of each group differing by at most one tag, and assign (reserve) each slot of the TDMA cycle to a unique group. Compared with the one-tag-per-slot TDMA, a grouped approach increases the average energy consumption per slot by the cardinality of the groups, but decreases the average delay since there is a greater probability that a tag will be awake soon after a packet for it has arrived at the base station. Since intuitively for a given load there is an optimal group size, it is

possible in this protocol to optimize energy and delay parameters. However, the grouped-tag TDMA protocol has a number of potential drawbacks. Since tags' schedules are cyclic and independent of packets in the base station's queue (a tag does not know when the base station has a packet destined for it) tags continue to wake up cyclically. Therefore significant energy waste may occur. Secondly, it is easy to see that if the packets' destination distribution is heavily clustered, the performance of the protocol can degrade severely and the reorganization of the groups, even if possible, can be excessively time-consuming.

Many paging protocols use a variation of the Grouped-Tag TDMA protocol for energy-conservation. For example, Europe's most prevalent paging protocol, POCSAG, divides the pagers into eight groups based on their fixed ID number[8]. The small, static number of pager groups provides a limited paging delay at the cost of a fixed, minimal energy savings even in situations of very low traffic, where waking up less frequently would be advantageous.

2.2 Directory Protocols

The use of a directory for the transmission of wireless data has been introduced by Imielinski, Vishnatwan and Badrinath [5], who considered this principle for applications which do not have strict delay constraints. Their work, therefore, concentrates on the intelligent and efficient organization of the directory. The IEEE 802.11 wireless LAN draft standard also suggests the use of a directory for energy conservation purposes [1]. This system utilizes variable-sized packets on a shared, non-slotted, CSMA/CA channel such that the transmission of the directories may be deferred for extended periods of time when one of the mobile nodes is transmitting on the channel. In contrast to these works, our application requires optimal selection of the transmission set size and careful scheduling of the idle periods leading to a different protocol.

In our directory protocol, the base station waits for a group of k packets in the queue to accumulate before starting the transmission. The transmission consists of broadcasting a directory which lists the destinations of the k packets which are to be transmitted, followed by the transmission of the packets themselves. Tags listen to the directory to find out if and when packets will be sent to them. They can then schedule their wake-up slots to coincide with the broadcast of their packets. When there is no group being currently transmitted, the tags wake up periodically every v slots in order to give the base station an opportunity to start the transmission of a new group.

The choice of the parameter k depends on the load and must take into account the trade-off between the increase in the delay due to a larger k and the energy savings from more infrequent broadcasting of directories. Additionally, the parameter v should depend upon k and the load. A tag system with small value of v and a low load will have the tags waking up frequently until enough packets have accumulated at the base station, while a large value of v will incur an increase in delay before the start of group's transmission.

The directory approach solves the potential unfairness encountered by the grouped-tag TDMA when the destination distribution is heavily clustered. However, for the directory protocol, when group size k is large, the size of the directory becomes prohibitively large, and therefore the amount of energy used by the tags just to read the directory becomes a major factor.

Several improvements can be considered for reducing the energy consumption during the directory transmission. By ordering the directory according to destination IDs and sending the boundaries of the "transmission interval" (i.e., the upper and lower value of the destination IDs of the k packets) it is possible for tags to go to sleep after discovering when their packet will be transmitted or as soon as they determine that there will be no packet addressed to them in this group. The latter event occurs if either the tag does not fall within the transmission interval or the directory indicates the transmission of a packet to a higher-order ID.

2.3 Random Access Protocol

In the random access protocol the transmission scheduling strategy adopted by the base station is to randomly select one of the packets in the queue. The packet is successfully received during the current slot if and only if the destination of the packet, randomly chosen by the base station, is awake. Tags are awake with probability p. As mentioned previously, the probability of a tag being awake when a packet is sent to it is extremely low unless p is large, indicating a high energy consumption. Therefore, this protocol is considered only for the sake of a baseline comparison.

2.4 Pseudo-Random Protocols

The pseudo-random protocols are a new class of protocols [6, 7] based on deterministic (pseudo-random) schedules which can preserve the power of randomization for fairness, while providing the advantages of determinism, i.e., the base station's ability to predict tags' state in each slot. In this class of protocols all tags run the same pseudo-random number generator and determine their state (awake or asleep) at each slot based on a probability p and the stored state of the random number generator. In order to avoid a complete overlap of the wake-up schedules, the pseudo-random generator of each tag is initialized with a unique seed, which is known at the base station. Therefore, by using the same pseudo-random number generator it is possible for the base station to determine the schedules of the tags it wants to transmit to. The base station can initiate changes in the value of p as a function of the load, the number of tags, etc. Moreover, different tags can operate with different p's based upon tags' individual expected traffic rates, making the protocol appropriate for handling heterogeneous traffic patterns.

3 Analysis

To compare the performance of the various protocols proposed, we consider the behavior of a single cell tag system.

Since the time needed to successfully receive a packet, given that the tag is awake, is exactly one slot, then the slot duration is $\sigma = b + \tau$, with b denoting the packet transmission time ($\frac{c}{B}$) and τ the maximum propagation delay. The presented evaluation utilizes the following definitions:

T − the average waiting time in slots experienced by a packet in the system from arrival at the base station to successful reception at the tag.

E − the average percentage of slots in which a tag is awake

L − the average number of packets in the system

The energy measure proposed does not take into account the contribution during slots where a tag is asleep. The energy used while in this state is consumed over the lifetime of the tag, whether the tag is being used or not. Therefore only the percentage of awake slots is necessary for comparing different energy-conserving access protocols.

We assume that the packets arrive at the base station according to a Poisson process with interarrival rate λ. Each packet is addressed to a single destination selected from either a uniform or a Gaussian distribution. The latter is introduced in order to model the heterogeneous nature of the traffic which is common to many tag applications. We use the following notations:

λ^1 − the mean arrival rate expressed in packets per slot ($\lambda^1 = \lambda(b + \tau)$)

μ^1 − the mean service rate expressed in packets per slot

ρ^1 − the utilization factor $\frac{\lambda^1}{\mu^1}$

3.1 Analysis of the Grouped-Tag TDMA

The service time for packets whose destinations are in different groups is independent. Therefore, for each group i, we can associate a different queueing system whose Poisson traffic stream has rate λ_i^1. Every time the queue is empty, the server goes "on vacation" for one TDMA cycle ($m = \lfloor \frac{N}{x} \rfloor$ slots). Otherwise the service time of a packet is constant and equals 1. Thus, the average waiting time in the queue for a packet addressed to group i is equivalent to that of an $M/D/1$ queueing system with vacations [2]:

$$W_i = \frac{\rho_i^1}{2\mu_i^1(1 - \rho_i^1)} + \frac{m}{2} \quad (1)$$

$$= \frac{m}{2(1 - \lambda_i^1 m)} \quad (2)$$

Let T_x denote the average number of slots, as a function of the parameter x, during which a packet is waiting in the system. Then:

$$T_x = \sum_{i=1}^{m} \frac{\lambda_i^1}{\lambda^1}(W_i + 1) \quad (3)$$

$$= 1 + \sum_{i=1}^{m} \frac{\lambda_i^1 m}{2\lambda^1(1 - \lambda_i^1 m)}. \quad (4)$$

It remains to calculate the λ_i^1. For sake of presentation we compute these values under the assumption $\frac{N}{m}$ is integer. However, it is easy to see that the approach can be easily extended to the non-integer case. The group interarrival rate parameters depend on the destination distribution utilized. Under a uniform destination distribution the percentage of the global traffic dedicated to each group is proportional to the number of tags in that group. Thus,

$$\lambda_i^1 = \frac{\lambda^1}{m} \quad (5)$$

We next consider the case of when the packets destinations are chosen according to a Gaussian distribution with a density function f_x, average μ and variance σ. $\lambda_i^1 = \lambda^1 p_i$, where p_i denotes the probability of choosing a packet in the i-th group. p_i is the area of the region bounded above by f_x and the interval associated to the i-th group, normalized by the area of the *usable* part of the Gaussian (i.e., the one corresponding to the interval $[1...N]$). So

$$p_i = \frac{F_x(\frac{iN}{m}) - F_x(\frac{(i-1)N}{m})}{F_x(N) - F_x(0)} \quad (6)$$

where

$$F_x(y) = \phi(\frac{y - \mu}{\sigma}) \quad (7)$$

and

$$\phi(z) = \int_{-\infty}^{z} \frac{e^{\frac{-t^2}{2}}}{\sqrt{2\pi}} dt \quad (8)$$

To compute the average energy consumption we observe that each tag is awake exactly once in a grouped-tag TDMA cycle. Thus:

$$E = \frac{1}{\lfloor \frac{N}{x} \rfloor} \quad (9)$$

3.2 Analysis of the Directory Protocol

We analyze an improved version of the directory protocol, where the directory is ordered by increasing *id* of packets' destinations and the first two records of the directory contain the extreme values of the transmission interval. Let $a = \lfloor \frac{c}{\lceil \log N \rceil} \rfloor$ be the number of different records which fit into a slot. Then, the service time of a group is $k + k'$, where k' indicates the time needed for transmitting the directory and is given by $\lceil \frac{k+2}{a} \rceil$. When a cycle ends and no completed groups are in the queue, the tags go to sleep for an idle interval of v slots. Thus, a server can start processing a new group if and only if there is at least one completed group in the queue and either an idle or a service cycle ended in the previous slot. We compute the average waiting time T of a packet in the system as:

$$T = W_p + T_g - W_s \quad (10)$$

$$\cong \frac{k - 1}{2\lambda^1} + T_g - (\frac{k - 1}{2}) \quad (11)$$

$-W_p$ is the average waiting time of a packet in the queue before its group is completed and is $\cong \frac{k-1}{2\lambda^1}$

$-T_g$ is the average waiting time of a group in the system

$-W_s$ is the average time between the successful reception of a packet and the completion of its group's transmission. $W_s = \frac{k-1}{2}$

To compute T_g we construct the following analytic model. We observe the system at the regeneration points embedded at the beginning of each slot and model the system as a discrete time, infinite Markov chain $M = <S_{<i,j,h>}, P_{<i,j,h>,<i',j',h'>}>$. Each state $S_{<i,j,h>}$ denotes the number i of completed groups in the system, the number j of non-grouped packets in the queue, and the index h of the current slot in either the idle ($0 \leq h \leq v - 1$) or the service ($v \leq h \leq v + k + k' - 1$) cycle. Note that within either a service or idle cycle we have only to increase the position of the slot and keep track of the arrivals, while during the last slot of a cycle we also have to decide the nature (busy/idle) of the next interval and possibly indicate the exit of a group from the system. We obtain the following transition probabilities:

Case $h \in [0...v-1) \bigcup [v...v+k+k'-1)$
$$P_{<i,j,h>,<i',j'h'>} =$$
$$\begin{cases} 0 & (h' \neq h+1) \vee (i' < i) \\ & \vee ((i' = i) \wedge (j' < j)) \\ A_{j'-j+(i'-i)k} & otherwise \end{cases}$$

Case $h = v - 1$
$$P_{<i,j,v-1>,<i',j',h'>} =$$
$$\begin{cases} 0 & ((h' \neq 0) \wedge (i' = 0)) \\ & \vee ((h' \neq v) \wedge (i' \geq 1)) \\ & \vee (i' < i) \vee ((i' = i) \wedge (j' < j)) \\ A_{j'-j+(i'-i)k} & otherwise \end{cases}$$

Case $h = v + k + k' - 1$
$$P_{<i,j,v+k+k'-1>,<i',j',h'>} =$$
$$\begin{cases} 0 & ((h' \neq 0) \wedge (i' = 0)) \\ & \vee ((h' \neq v) \wedge (i' \geq 1)) \\ & \vee (i' < i-1) \vee ((i' = i-1) \\ & \wedge (j' < j)) \\ A_{j'-j+(i'-i)k+k} & otherwise \end{cases}$$

$-A_r$ is the probability of r messages originated by a Poisson process during a slot and is given by:

$$A_r = e^{-\lambda^1} \frac{(\lambda^1)^r}{r!} \qquad (12)$$

The steady state probability $\pi_{<i,j,h>}$ of being in the $<i,j,h>$ state is obtained by solving the following system of equations:

$$\begin{cases} \Pi = P^T \Pi \\ \sum_i \sum_j \sum_h \Pi = 1 \end{cases} \qquad (13)$$

Let us denote $\pi_g(i)$ as the probability of having i completed groups in the system:

$$\pi_g(i) = \sum_j \sum_h \pi_{<i,j,h>} \qquad (14)$$

The average number of completed groups in the system L_g can then be given by:

$$L_g = \sum_i i \pi_g(i). \qquad (15)$$

Finally, applying Little's theorem we can compute the average waiting time of a group in the system:

$$T_g = \frac{L_g k}{\lambda^1} \qquad (16)$$

The energy consumption E can be computed as follows:

$$E = Pr_i E_i + Pr_b E_b \qquad (17)$$

$-Pr_i$ is the percentage of slots belonging to idle cycles and is given by:

$$Pr_i = \sum_i \sum_{j=0}^{k-1} \sum_{h=0}^{v-1} \pi_{<i,j,h>}. \qquad (18)$$

$-Pr_b$ is the percentage of busy slots, $Pr_b = 1 - Pr_i$.

$-E_i$ is the average energy consumption per tag per slot during an idle period and is equal to $\frac{1}{v}$.

$-E_b$ is the average energy consumption per tag per slot during a service cycle.

$$E_b = \frac{E_{dir} + \frac{k}{N}}{k + k'}, \qquad (19)$$

where E_{dir} indicates the average energy consumption during the directory transmission.

Let us define $p_{i,h}$ as the probability that the i-th tag is awake during the h-th slot of the directory. $p_{i,h}$ expresses the probability that the i-th tag is in the transmission interval and at least $l = a(h-1) - 2$ packets have been chosen in $\{1...i\}$. Then:

$$E_{dir} = \frac{1}{N} \sum_{i=1}^{N} \left[1 + \sum_{h=2}^{k'} p_{i,h} \right] \qquad (20)$$

When using a uniform destination distribution

$$p_{i,h} = \frac{\sum_{j=l}^{k-1} i^j (N-i)^{k-j} \binom{k}{j} + i^k - (i-1)^k}{N^k} \qquad (21)$$

We now assume that the packet destination is chosen according to a Gaussian distribution of average μ and variance σ.

The probability p_i of choosing the i-th tag as the destination is given by:

$$p_i = \frac{F_x(i) - F_x(i-1)}{F_x(N) - F_x(0)} \qquad (22)$$

Then the probability p_s associated to a sequence $s = i_1, ..., i_k$ is

$$p_s = \prod_{j=1}^{k} p_{i_j} \qquad (23)$$

Let us call $S(j, k, N)$ the set of distinct sequences having j elements in $1, ..., i$ and including i in the transmission interval, then

$$p_{i,h} = \sum_{j=l}^{k} \sum_{s \in S(j,k,N)} p_s \qquad (24)$$

It remains to compute $\sum_{s \in S(j,k,N)} p_s$. Let $s < h, j, i >$ be the h-th sequence we get when we enumerate all the possible ways to choose j elements out of a set of cardinality i.
Case $j \neq k$:

$$\sum_{s \in S(j,k,N)} p_s = \binom{k}{j} \sum_{h_1=1}^{i^j} p_{s<h_1,j,i>}$$
$$\cdot \sum_{h_2=1}^{i^{k-j}} p_{s<h_2,k-j,N-i>} \qquad (25)$$

$$= \binom{k}{j} \sum_{h_1=1}^{i^j} \prod_{m_1=0}^{j-1} p'_{\lceil \frac{h_1}{i^{m_1}} \rceil \bmod i}$$
$$\cdot \sum_{h_2=1}^{i^{k-j}} \prod_{m_2=0}^{k-j-1} p'_{i+\lceil \frac{h_2}{(N-i)^{m_2}} \rceil \bmod (N-i)} \qquad (26)$$

$$p'_u = \begin{cases} p_i & u = 0 \\ p_N & u = i \\ p_u & otherwise \end{cases} \qquad (27)$$

Case $j = k$:

$$\sum_{s \in S(k,k,N)} p_s = \sum_{h_1=1}^{i^k} \prod_{m_1=0}^{k-1} p'_{\lceil \frac{h_1}{i^{m_1}} \rceil \bmod i} \qquad (28)$$
$$- \sum_{h_2=1}^{(i-1)^k} \prod_{m_2=0}^{k-1} p''_{\lceil \frac{h_2}{(i-1)^{m_2}} \rceil \bmod (i-1)} \qquad (29)$$

where

$$p''_u = \begin{cases} p_{i-1} & u = 0 \\ p_u & otherwise \end{cases} \qquad (30)$$

3.3 Analysis of the Random Access Protocol

We evaluate this protocol's performance by introducing a discrete time infinite Markov chain embedded at the beginning of a slot $M = < S_i, P_{i,j} >$. The states S_i represent the number of packets in the system.

Thus:

$$P_{0,j} = A_j \qquad (31)$$
$$P_{i,j} = \begin{cases} 0 & j < i-1 \\ pA_0 & j = i-1 \\ pA_{j-i+1} + (1-p)A_{j-i} & otherwise \end{cases} \qquad (32)$$

Let Π denote the steady-state probabilities given by the solution of the following linear equations:

$$\begin{cases} \Pi = P^T \Pi \\ \sum_i \Pi = 1 \end{cases} \qquad (33)$$

π_i expresses the probability that there are i packets in the system. Then, L can be computed as follows:

$$L = \sum_i i\pi_i \qquad (34)$$

and, using the Little's result:

$$T = \frac{L}{\lambda^1} \qquad (35)$$

Finally, from the definition of energy consumption given in the previous section it follows that $E = p$.

3.4 Approximate Analysis of the Pseudo-Random Protocol

We analyze the protocol under the assumption that all tags use the same awake probability parameter p. To obtain an exact description of the system behavior through a Markov chain, the states should include the number of packets addressed to each tag, since the probability of successfully transmitting depends on the number of unique packet destinations in the queue. It is therefore impossible to analyze systems of a realistic size. We solve this problem by making use of Stern's independence assumption [10], stating that at the beginning of a time slot each message draws a new destination from a given uniform distribution. The Markov chain describing the system is then represented by the total number of packets at the beginning of a slot. The transition probability matrix P is given similar to the random access case, by

$$P_{0,j} = A_j \qquad (36)$$
$$P_{i,j} = \begin{cases} 0 & j < i-1 \\ \sum_{k=1}^{min(i,N)} R_{i,k}(1-(1-p)^k)A_0 & j = i-1 \\ \sum_{k=1}^{min(i,N)} R_{i,k}(1-(1-p)^k)A_{j-i+1}+ \\ \sum_{k=1}^{min(i,N)} R_{i,k}(1-p)^k A_{j-i} & otherwise \end{cases} \qquad (37)$$

$-R_{i,k}$ is the probability that i packets are addressed to k different destinations and is given by

$$R_{i,k} \quad = \quad \frac{\begin{pmatrix} N \\ k \end{pmatrix} \sum_{j=0}^{k-1} (-1)^j \begin{pmatrix} k \\ j \end{pmatrix} (k-j)^i}{N^i} \quad (38)$$

We note that the difference between the pseudo-random sequences and the random access protocol is reflected in the different definition of successful reception. It is now sufficient that at least one destination is awake during the current slot, instead of only the destination of the packet chosen by the base station.

The steady-state probabilities π_i are the solutions of the following set of equations

$$\begin{cases} \Pi = P^T \Pi \\ \sum_i \Pi = 1 \end{cases} \quad (39)$$

We can now compute the average number L of packets in the system

$$L \quad = \quad \sum_i i\pi_i \quad (40)$$

and, using the Little's result:

$$T \quad = \quad \frac{L}{\lambda^1} \quad (41)$$

Finally, as in the previous section, the average energy consumption is obviously $E = p$.

4 Performance

For validation purposes we simulated the reception of $15,000$ packets for each of the protocols in a system of $N = 1000$ tags. Interarrival rate parameters were equal to 0.05, 0.2 and 0.5. We considered the cases of uniform and heterogeneous destination distributions. The latter was simulated in two cases, the first being a Gaussian with mean $\frac{N}{2}$ and variance $10N$, and the second a Gaussian with mean $\frac{N}{2}$ and variance N. To validate the approximate analytical models described in the previous section, we compared computationally derived results of the analysis with the results of the simulations. The infinite Markov chain of the pseudo-random protocol was approximated by truncating the transition probability matrix after the first 100 states while up to 800 states were considered for the directory protocol with the limitation $k \in \{1, \ldots, 5\}$ and $v \in \{1, \ldots, 10\}$. Our simulator counted in each slot i the number of awake tags E_i and the delay D_i incurred to each received packet. These quantities were then used to compute the


Figure 1: Classical Access Protocols

Figure 2: Grouped TDMA Protocol, Uniform Destination Distribution, Optimal group size (varying interarrival rates with a sampling of analysis results depicted in discrete points for $\lambda = 0.05$)

average access delay T ($\sum_i D_i / 15,000$packets) and the average energy consumption E ($\sum_i E_i/(N * \text{simulation length})$, where the length of the simulation was the amount of time to generate and send 15,000 packets with the appropriate Poisson arrival distribution). The average delay for the grouped-tag TDMA sequences computed with the analysis was consistently around half a slot greater than that obtained through the simulations, since using slotted arrivals "shifts" the arrivals of packets from within a certain slot to the end of it. In all other cases the comparison of the simulations with the analytic results showed an excellent approximation with error under 2%. A sampling of the analytical results, for each of the protocols for the case of $\lambda = 0.05$, is displayed as discrete points in Figs. 2–4.

The simulation results are shown in Figs. 1–6. Fig 1 depicts the energy vs. delay behavior of the random access and classical TDMA protocols. Notice that it is necessary to have

Figure 3: Directory Protocol, Uniform Destination Distribution, Optimal K (varying interarrival rates with a sampling of analysis results depicted in discrete points for $\lambda = 0.05$)


Figure 6: Energy conserving protocols, Gaussian destination distribution ($\mu = \frac{N}{2}, \sigma^2 = 10N$)


Figure 4: Probabilistic Sequences Protocol, Uniform Destination Distribution (varying interarrival rates with a sampling of analysis results depicted in discrete points for $\lambda = 0.05$)


Figure 7: Energy conserving protocols, Gaussian destination distribution ($\mu = \frac{N}{2}, \sigma^2 = N$)


Figure 5: Energy conserving protocols, Uniform destination distribution

a much coarser scale in order to show the results for these two protocols. The random selection criteria adopted by the base station in the random access protocol makes unlikely a synchronization of the two schedules and leads to poor performance of the protocol, which is stable only for high values of p and therefore a high energy consumption ($\geq \lambda$ is a necessary condition). The classical TDMA shows on the other hand a very good energy consumption (0.001 for $N = 1000$) but an extremely long delay (≥ 500 slots) since the cycle time increases as the number N of tags.

Fig.2,3,4 compare the performance of the three protocols proposed using a uniform destination distribution. Fig.2 shows the trade-off between energy and delay for optimized values of group size using the grouped-tag TDMA protocol. The choice of a larger size decreases the average delay but increases the energy consumption per tag per slot. We also note that the performance is better for low loads because there are fewer packets available for transmission for the same slot in a cycle. The

performance of the directory protocol depends on the choice of the parameters k and v (Fig.3). By waiting for a larger number of packets we reduce the energy required for reading the directory but heavily increase the average delay experienced by a packet due to having to wait for its group to be completed before it can be transmitted. The length of this waiting time decreases as λ increases. Given a value of k, a higher v results in a lower energy consumption but an increase in delay before the start of next group's transmission. Fig.4 shows the energy vs. delay graph for the pseudo-random protocol given a uniform destination distribution. The performance achieved is slightly worse than that of the grouped-tag TDMA although this difference decreases as λ increases. The simulation results for each protocol under a traffic load of $\lambda = 0.2$ is reproduced in Figure 5 for ease of direct comparison.

Fig.6, and 7 compare protocol performance for the case of heterogeneous traffic. The directory protocol is hardly affected by the destination distribution. The only slight difference is due to the reduced number of tags which must be awake during the directory slots. Therefore, the results for this protocol are not displayed in the graphs. Instead, we show how the performance of the grouped-tag TDMA degrades rapidly under heterogeneous traffic. Packets belonging to a group with a high probability of traffic are severely delayed due to the high volume of localized destinations. Since the number of packets in the queue addressed to the same group increases as either the load increases or the Gaussian distribution becomes narrower, the performance of the protocol decreases in both these cases. The case of $\lambda = 0.5$, with a destination distribution of $Gaussian(\mu = \frac{N}{2}, \sigma^2 = N)$, is completely unstable and therefore not presented in Fig.7. It can be seen that the heterogeneous traffic affects the behavior of the pseudo-random protocol, since a higher concentration of packets decreases the probability of successful transmission. However, this effect can only be seen for extremely high loads (or very narrow Gaussian destination distributions) and is much weaker than the one noticed for the grouped-tag TDMA protocol. For a wide Gaussian destination distribution ($(Gaussian(\mu = \frac{N}{2}, \sigma^2 = 10N))$ the pseudo-random protocol outperforms the grouped-tag TDMA starting from $\lambda = 0.2$, while for a narrower Gaussian ($Gaussian(\mu = \frac{N}{2}, \sigma^2 = N)$) the pseudo-random protocol is much more effective.

5 Conclusions

In this paper, we addressed the problem of wireless access protocols which include an energy constraint. We showed that classical multiple access protocols such as TDMA and Aloha are either not energy-conserving or lead to unacceptable delays. Therefore, we have proposed three energy-conserving protocols: grouped-tag TDMA, directory and pseudo-random. Careful selection of protocol parameters addresses the goal of minimizing the energy required for reception of packets while meeting the application delay constraints. Both analytical and simulation results were presented and described.

All three protocols perform very well for various loads. In particular, the grouped-tag TDMA protocol achieves the best performance for either a low traffic load or a uniform destination distribution. When the destination distribution becomes more realistically clustered, other protocols such as the directory or pseudo-random protocols need to be adopted.

The pseudo-random protocol consistently out-performs the directory protocol in both energy and delay for all loads and most destination distributions. However, for cases where the destination distribution becomes heavily clustered, it is apparent that the directory protocol is the best solution, unless dynamic handling of the protocol parameters is introduced. For typical RFID applications, the pseudo-random protocol appears therefore to be capable of providing a simple and effective access method which includes energy-constraint.

References

[1] IEEE, "Wireless LAN Media Access Control (MAC) and Physical Layer (PHY) Specification", IEEE 802.11 Draft Version 4.0, May 1996.

[2] Bertsekas, D. and Gallager, R. *Data Networks*, Prentice Hall, 1992.

[3] Harris, E.P., Depp, S.W., Pence,E., Kirkpatrick, S., Sri-Jayantha, M., and Troutman, R.R. Technology directions for portable computers. *Proceedings of the IEEE*, 83(4):636–58, April 1995.

[4] Dunn, D. RFID Engine Chugs Onto Scene. *Electronic Buyer's News*, February 5, 1996.

[5] Imielinski, T., Vishnatwan, S. and Badrinath, B. R. Data on Air. *IEEE Transactions of Data and Knowledge Engineering*, July 1996.

[6] Chlamtac, I., Petrioli, C. and Redi, J. An Energy-Conserving Access Protocol for Wireless Communication. *IEEE International Conference on Communications (ICC97)*, Montreal(Canada), June 1997, pp. 1059–1062.

[7] Chlamtac, I., Petrioli, C. and Redi, J. Extensions to the Pseudo-Random Class of Energy-conserving Access Protocols *Second IEEE International Workshop on Wireless Factory Communication Systems*, Barcelona, October 1997, pp. 11–16.

[8] British Telecom, "A Standard Code for Radiopaging," Report of the post Office Code Standardization Advisory Group (POCSAG), June 1987 and November 1980.

[9] Powers, R.A. Batteries for low power electronics. *Proceedings of the IEEE*, 83(4):687–93, April 1995.

[10] Stern, T. E. Packet Scheduling protocols in multi-beam communication satellites. *Proceedings of the International Symposium on Information Theory*, Italy, June 1979, pp. 130–131.