From the Proceedings of The International Conference on Communications (ICC-97), June 8–12, 1997.

An Energy-Conserving Access Protocol for Wireless Communication

Imrich Chlamtac[a], Chiara Petrioli[b], Jason Redi[a]
[a] Boston University, Boston, MA, U.S.A., {*chlamtac,redi*}*@bu.edu*
[b] Università di Roma “La Sapienza”, Italy, *petrioli@dsi.uniroma1.it*

Abstract

A myriad of applications such as radio frequency identification (RFID) and smart card networks are emerging in which nodes are designed for extremely low-cost, large scale applications such that the replacement of batteries is not feasible. Energy conservation therefore becomes a major constraint. Classical access protocols are either not energy conserving or lead to unacceptable delays. In this paper, we propose a communication protocol which meets the energy constraints while yielding low access delays.

1 Introduction

A typical identification network (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, also called RFID “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 and is expected to expand to over $5 billion [2] by the year 2000. Some examples of current uses for RFID tags are: tags attached to the ears or worn around the necks of livestock for use in location tracking, smart tags used in warehouses to automatically and quickly track inventory, and the numerous tag companies targeting the retail market for electronic shelf labels and automatic purchasing.

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

1. *Scale*: There may be a very large number, perhaps thousands, of wireless nodes per base station.
2. *Cost*: Due to the large number of extremely inexpensive items to “tag”, nodes must be extremely inexpensive, often on the order of only a few dollars.
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, as in all mobile computing applications, the battery’s energy is a limited and scarce resource, which is not expected to increase in potential more than 30% in the near future [4]. Especially demanding of energy is any uplink transmission as it 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. Low energy consumption is the second requirement that the access protocol must satisfy. The large number of nodes makes it economically impossible to replace or recharge the batteries. Therefore tags must be designed so as to require a minimum of energy for operation. 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. In contrast with this, the *awake* state, in which the CPU operates at full energy and the radio frequency circuitry is active, can typically use 100 times as much energy.

No protocols are available today whose objective is to reduce the necessary power required while minimizing delay. Current access protocols, such as those described in [1], have not been designed with these objectives in mind 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 of a tag being awake in the same slot in which the base station is transmitting to it, is extremely low. 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.

This paper deals with the design of a *communications protocol which operates under an 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. We present a new approach to this problem, a *pseudo-random protocol*, which combines the fairness from random access protocols with the low energy requirements of classical TDMA. Considering an analytical model as well as simulations of the system behavior for the cases of uniformly and non-uniformly distributed traffic destinations, we show that in addition to fairness, and low energy requirements the proposed protocol provides low access delays.

2 Network Model and Protocol 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 system 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 class of protocols presented in this paper the "oldest packet" criterion is generally adopted to help meet the application delay requirements.

2.1 Pseudo-Random Protocols

A pseudo-random protocol uses a deterministic (pseudo-random) schedule 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 protocol 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 identical wake-up schedules, the pseudo-random generator of each tag is initialized using a unique seed, also known to the base station. Therefore, it is possible for the base station to determine which packets in its queue have destinations which are awake, and then transmit the oldest of these packets. The base station can also initiate changes to the value of p as a function of the load, the number of tags, etc. Lastly, we observe that different tags can operate with different p's based upon the tags' individual expected traffic rates, making the protocol appropriate for handling heterogeneous traffic patterns.

3 Analysis

To evaluate the performance of the pseudo-random protocol, 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 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 evaluating the performance of an energy-saving protocol.

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 notation:

$\lambda-$ the mean arrival rate expressed in packets per slot ($\lambda = \lambda'(b + \tau)$).

3.1 Approximate Analysis

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 [5], 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 by

$$P_{0,j} = A_j$$

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

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

$$A_r = e^{-\lambda}\frac{(\lambda)^r}{r!}.$$

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

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

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

$$\begin{cases} \Pi = P^T \Pi \\ \sum_i \Pi = 1. \end{cases}$$

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

$$L = \sum_i i\pi_i$$

and, using the Little's result:

$$T = \frac{L}{\lambda}.$$

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

4 Performance

For validation purposes we simulated the reception of $15,000$ packets for the proposed protocol in a system of $N = 1000$ tags with interarrival rate parameters equal to 0.05, 0.2 and 0.5. We considered the cases of uniform and heterogeneous destination distributions. The latter was simulated using a Gaussian with mean $\frac{N}{2}$ and variance N. To validate the approximate analytical model 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. We verified the values of both the average waiting time T in

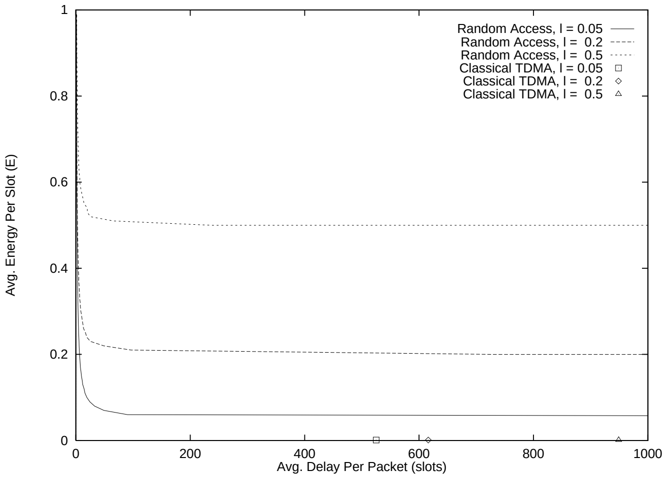

Figure 1: Classical Access Protocols

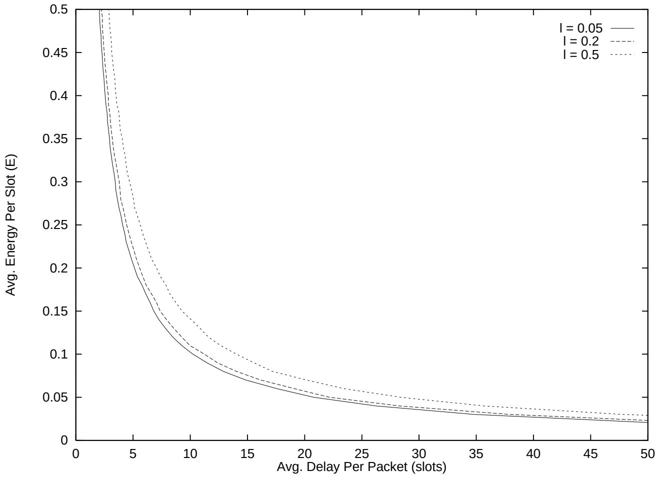

Figure 2: Pseudo-Random Protocol, Uniform Destination Distribution

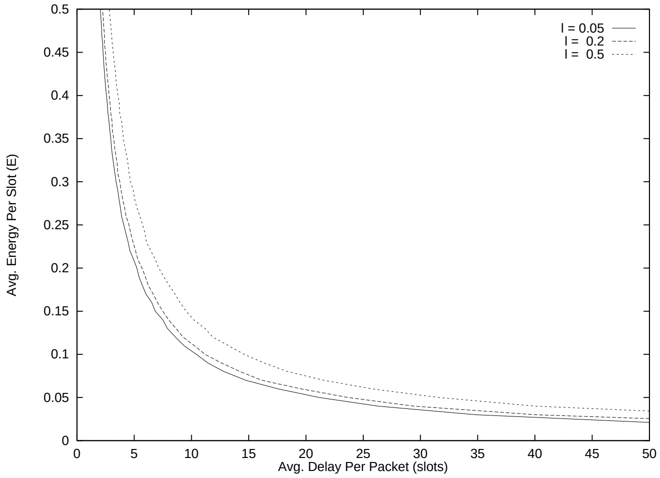

Figure 3: Pseudo-Random Protocol, Gaussian destination distribution ($\mu = \frac{N}{2}, \sigma^2 = N$)

the system and the average energy consumption E. For both these performance measures the comparison of the simulations with the analytic results showed an excellent approximation with error under 2%.

The simulation results are shown in Figs. 1–3. Fig 1 depicts the energy vs. delay behavior of the random access and classical TDMA protocols. Notice that it is necessary to have a much coarser scale in order to show the results for these two protocols. The random selection criterion adopted by the base station in the random access protocol makes unlikely a synchronization of the two schedules and leads to extremely poor performance of the protocol. It is stable only for high values of p and therefore a high energy consumption ($\geq \lambda$ is a necessary condition). On the other hand, the classical TDMA shows 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 shows the energy vs. delay graph for the pseudo-random protocol given a uniform destination distribution. The performance of the pseudo-random protocol is a dramatic improvement over the classical protocols. We note that the difference between the pseudo-random and the random access protocols reflects the different conditions for successful reception. The former requires at least one destination to be awake during the current slot, while in the latter the packet is randomly selected by the base station. Also, it is easy to see that the classical TDMA protocol performs even worse and becomes rapidly unstable when the destination distribution is clustered. Fig.3 depicts the pseudo-random protocol performance for the case of heterogeneous traffic. Non-uniform 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 barely significant.

5 Conclusions

In this paper, we addressed the problem of wireless access protocols which include an energy constraint. Since classical multiple access protocols do not satisfy the dual requirement of low energy and acceptable delays, we have proposed a new protocol class and developed analytical and simulation models for its evaluation. We showed that the pseudo-random protocol performs extremely well for various loads and its behavior is only slightly affected by the use of heterogeneous destination distributions.

References

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

[2] Dunn, D. RFID Engine Chugs Onto Scene. *Electronic Buyer's News*, February 5, 1996, pp 14.

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

[5] 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.