An unslotted multichannel channel-access protocol for distributed direct-sequence networks

Arvind R. Raghavan and Carl W. Baum *

Electrical and Computer Engineering Department, Clemson University, Riggs Hall, Clemson, SC 29634-0915, USA

A multichannel reservation-based channel-access protocol is investigated in this paper. The available system bandwidth is divided into distinct frequency channels. Under the protocol, one channel (the *control* channel) is used to exchange reservation messages and the remaining channels (the *traffic* channels) are used for information-bearing traffic. The performance of this scheme is compared to that of a single-channel reservation-based protocol. A simple contention-based slotted-Aloha protocol is also considered. Performance results take into account the effects of multiple-access interference on acquisition and packet errors. Results show that the reservation-based approach is advantageous under conditions of high traffic. In addition, a pacing mechanism that mitigates multiple-access interference and promotes fairness is described, and results are presented that demonstrate its effectiveness.

1. Introduction

Direct-sequence spread-spectrum technology is being used increasingly in mobile wireless networks. Most commercial applications employ centralized networks in which a base station controls the operation of all the radios in the network. In other applications, however, a central network controller may not be practical. For example, in military applications a central control unit may present a vulnerability to the network as a whole, or in disaster relief operations, it may not be feasible to install a central controller. For such applications, distributed packet radio networks must be considered.

In a distributed network, each radio must coordinate its transmissions with other radios in its vicinity so that the channel is utilized efficiently. Towards this end, it is important to design a channel-access protocol that ensures adequate throughput in a network in which the topology constantly changes as a result of radio mobility. Channel-access protocols can be broadly classified as contention-based and reservation-based. In a contention-based protocol, there is no prior coordination among the radios in the network before a transmission is initiated [1]. As a result, simultaneous transmissions are not prohibited, and significant multiple-access interference may occur. In reservation-based protocols, short transmissions are exchanged between radios to reserve use of a channel for a longer subsequent transmission. For a detailed discussion of different reservation mechanisms, see [2].

In a distributed environment, it is difficult to establish a common time reference that is precise enough to allow the use of short slots for transmission of reservation packets. In this paper, we restrict our attention to *unslotted* reservation-based protocols. We consider a multichannel protocol in which it is assumed that the available bandwidth is divided into distinct frequency bands. One of the channels is designated the *control* channel and is used exclusively for the transmission of reservation packets. The other channels carry information-bearing packets and are referred to as *traffic* channels. The idea of using separate control and data channels was first suggested in [3], where the authors describe a protocol that uses only two channels. The analysis does not take into account the bandwidth cost of the control channel. A special single-channel version of this protocol is also considered in which reservation and information packets share the same channel. For comparison purposes, we also consider a simple contention-based slotted-Aloha protocol [4].

The control channel considered in our protocol is a random access channel, and there is no coordination between radios before reservation requests are made. For source radios at unequal distances from the destination radio, the probability of successful reservation requests can be unequal, resulting in an unfair advantage to closer source radios. Moreover, unless reservation requests are overheard by radios in the vicinity, multiple-access interference will continue to pose a threat to transmissions in the traffic channels. In this paper, we describe a *pacing* mechanism that forces radios to overhear reservation requests. This pacing mechanism is a part of the channel access protocol, unlike the pacing described in [5] which can be considered to be a network layer mechanism because it provides a form of flow control as well as congestion control. The purpose of pacing as described in this paper is to control the "pace" of transmissions at the channel access layer to mitigate the interference in the traffic channels and also reduce the disadvantage suffered by distant radios.

An event-driven simulation has been used to investigate the performance of channel-access protocols in distributed direct-sequence networks. In the simulation, the effects of acquisition, modulation and coding have been

* This work was supported by the Army Research Office under grant number DAAH04-95-1-0247 and by the Office of Naval Research under grant number N00014-96-1-0896.

© Baltzer Science Publishers BV

modeled carefully. Multiple-access interference, additive white Gaussian noise and carrier phase drift are taken into account in the link-level portion of the simulation.

The remainder of this paper is organized as follows. In section 2, the models for the physical layer (the wireless link) are described. Section 3 contains a detailed description of the channel-access protocols considered in this paper. In section 4, we present simulation results, and conclusions are presented in section 5.

2. Physical layer models

2.1. Acquisition

We assume that all packet transmissions in the network are preceded by a short acquisition header in which no data modulation is present. We consider a simple non-coherent serial acquisition technique that employs passive in-phase and quadrature filters each matched to the pseudo-noise (PN) sequence present in the acquisition header. The output of the non-coherent detector is compared to a threshold to decide whether synchronization has occurred. A detailed analysis of this technique is given in [6], wherein a Gaussian approximation for the multiple-access interference is derived. This analysis takes into account the effects of an arbitrary number of interferers with unequal powers as well as thermal noise and carrier phase drift due to Doppler shifts and oscillator inaccuracies. The validity of the approximation depends only on the presence of a sufficient number of chips in the acquisition header. The effects of an automatic gain control (AGC) amplifier are modeled by normalizing the received signal power, interference power and noise power with respect to the total power. It is shown in [6] that a single threshold can be chosen that is robust with respect to a variety of channel conditions. Based on this analysis a fixed threshold is used in the simulation.

Let X_i denote the sum of the squares of the in-phase and quadrature filter outputs, sampled at the ith chip of the header. If this output X_i exceeds a threshold η, a *hit* is declared. If the detector declares a hit at least one chip interval before the acquisition header is received, a *false alarm* occurs. A finite amount of time is required for the tracking stage to detect the occurrence of a false alarm so that the receiver can return to the acquisition mode. If the end of the acquisition header arrives during this *false-alarm processing interval*, the desired signal is not acquired. Conversely, if the detector fails to declare a hit until after the end of the acquisition header arrives, a *miss* occurs, and the desired signal is not acquired. In the simulation, the length of the false-alarm processing interval is chosen to be equal to the header duration. Thus, successful acquisition of the desired signal is possible only if no false alarms occur for the $M-1$ samples immediately preceding the sample corresponding to the actual signal delay. In addition, if a false alarm occurs at least M samples prior to the correct sample, it has no effect on the acquisition of the desired signal.

Now, acquisition occurs if and only if $X_M > \eta$ and $X_j < \eta$, $1 \leqslant j \leqslant M-1$. It is shown in [6] that X_j, $j = 1, \ldots, M-1$, are approximately central chi-square random variables, X_M is approximately a non-central chi-square random variable, and $X_1, \ldots, X_M$ are mutually independent. Let $N_0/2$ denote the power spectral density of the thermal noise, let P_0 denote the power of the signal that is being acquired, and let P_k, $k = 1, \ldots, K$, denote the powers of the interfering transmissions after normalization by the AGC. Let T_c denote the chip duration and let τ_k denote the fractional chip offset of the kth interfering transmission. The density of $X_1, \ldots, X_{M-1}$ is given by

$$f_\mathrm{u}(x_j) = \frac{1}{2\sigma_\mathrm{u}^2} \exp\left(-\frac{x_i}{2\sigma_\mathrm{u}^2}\right), \quad x_j \geqslant 0, \tag{2.1}$$

where

$$\sigma_\mathrm{u}^2 = \frac{M}{2} N_0 T_\mathrm{c} + \frac{M}{2} \sum_{k=0}^{K} P_k \left[\tau_k^2 + (T_\mathrm{c} - \tau_k)^2\right], \tag{2.2}$$

and the density of X_M is given by

$$f_M(x_M) = \frac{1}{2\sigma_\mathrm{u}^2} \exp\left(-\frac{x_M + \mu^2}{2\sigma_m^2}\right) I_0\left(\frac{\mu\sqrt{x}_M}{\sigma_m^2}\right), \quad x_M \geqslant 0, \tag{2.3}$$

where

$$\sigma_m^2 = \frac{M}{2} N_0 T_\mathrm{c} + \frac{M}{2} \sum_{k=1}^{K} P_k \left[\tau_k^2 + (T_\mathrm{c} - \tau_k)^2\right], \tag{2.4}$$

and

$$\mu^2 = P_0 M^2 T_\mathrm{c}^2, \tag{2.5}$$

and $I_0(\cdot)$ is the modified Bessel function of the first kind and zero order. The probability of false alarm can be expressed as

$$\begin{aligned} P_\mathrm{FA} &= \mathrm{P}(X_{M-1} < \eta, \ldots, X_1 < \eta) \\ &= \left[\int_0^\eta f_\mathrm{u}(x)\,\mathrm{d}x\right]^{M-1} \\ &= \left[1 - \exp\left(-\frac{\eta}{2\sigma_\mathrm{u}^2}\right)\right]^{M-1}. \end{aligned} \tag{2.6}$$

In the simulation, we make certain "worst-case" assumptions. Any interfering transmission present for any fraction of the acquisition header duration is assumed to be present for the entire duration of the header. Also, the fractional chip offsets are set to the value that results in the greatest variance; i.e., $\tau_k = 0$, $k = 1, \ldots, K$.

For each header received when a radio is in the acquisition mode, the non-central chi-square detector output X_M is simulated. If X_M exceeds the threshold η, then the probability of false alarm is calculated using (2.6), and a Bernoulli random variable with parameter P_FA is generated to decide whether a false alarm occurred. If X_M exceeds the threshold and a false alarm does not occur, then the packet is successfully acquired.

2.2. Packet error probability

We assume that the radios employ binary phase shift keying with coherent demodulation and that the packets are encoded using a convolutional code. The probability of packet error for a direct-sequence spread-spectrum signal with convolutional coding cannot be computed directly, so bounds must be used, or the probability must be obtained by simulation. A bit-by-bit simulation to obtain packet error probabilities is computationally intensive from the perspective of a packet radio network simulation, so we resort to the use of an upper bound. An upper bound based on the use of hard-decision decoding is given by

$$P_{\mathrm{E}} \leqslant 1 - \left[1 - P_{\mathrm{u}}(\rho)\right]^{L}, \tag{2.7}$$

where P_{E} is the packet error probability, L is the number of bits in a packet, ρ is the probability of bit error, and $P_{\mathrm{u}}(\rho)$ is the first-event error probability for the convolutional code [7]. The probability $P_{\mathrm{u}}(\rho)$ is difficult to evaluate, so the Van de Meerberg bound on the first-event probability is used. This bound is given by [7]

$$\begin{aligned} P_{\mathrm{u}}(\rho) \leqslant \frac{\Gamma_{n_0}}{2}\big[&T(D) + T(-D) \\ &+ D[T(D) - T(-D)]\big]_{D=2\sqrt{\rho}}, \end{aligned} \tag{2.8}$$

where

$$\Gamma_{n_0} = \binom{2n_0 - 1}{n_0} 2^{-2n_0}, \tag{2.9}$$

n_0 is half the free distance of the code, and $T(D)$ is the generating or transfer function of the code. We assume that the system uses the NASA standard, rate $1/2$, constraint length 7 code. For this code, n_0 is 5 and $T(D)$ is given in [8]. The bit error probability ρ is a function of the chip offsets and relative phase of the interfering signals, which are likely to vary over the duration of the packet. It is therefore desirable to use an upper bound on the bit error probability. If all the signals are received with equal power, the worst-case bit error probability for coherent, binary signaling with random signature sequences is obtained when the interfering signals have synchronous phase and chip offsets. This upper bound can be approximated by

$$\rho \approx Q\left(\left[\frac{N_0}{2P_0NT_{\mathrm{c}}} + \frac{\sum_{k=1}^{K} P_k}{NP_0}\right]^{-1/2}\right), \tag{2.10}$$

where N is the number of chips per channel bit (processing gain) [7]. If the interfering signals are not received with equal power, the value for ρ computed using (2.10) still tends to provide a pessimistic estimate of the bit error probability.

In the simulation, it is possible for interfering transmissions to begin or end during the reception of a packet. To estimate the worst-case bit error probability, the maximum interference power present at any point during the packet reception is used for the sum of interference powers in (2.10). A pessimistic estimate of the packet error probability is obtained by substituting (2.10) and (2.8) in (2.7), and this probability is used to generate a Bernoulli random variable to decide whether the packet decodes correctly.

3. Channel-access protocols

3.1. The RTS-CTS protocol

In this section, we describe the multichannel reservation-based protocol which is the focus of this paper. Each radio in the network has a unique spreading code (PN sequence) assigned to it and each radio in the network maintains a table with the spreading code assignments for all the other radios in the network. All point-to-point transmissions are made using a receiver-directed code; i.e., the spreading code assigned to the destination radio is used. All broadcast transmissions are made using a common spreading code known to all the radios in the network. We assume that all radios are capable of monitoring a channel for incoming transmissions on both the receiver-directed code and the common code simultaneously. The radios are capable of transmitting and receiving on multiple channels, but at any given time, only one channel can be used. The radios we consider are half-duplex; i.e., they can transmit and receive, but cannot do both simultaneously.

A radio that has a packet to transmit first sends out a Request-To-Send packet (RTS) to the destination radio. This transmission is made on the control channel using the receiver-directed code. If the destination radio successfully acquires and decodes the RTS, it transmits a Clear-To-Send packet (CTS) on the control channel using the common code. The CTS contains the identity of the radio transmitting it (the destination radio), the identity of the radio it is being transmitted to (the source radio) and the traffic channel the source radio should use for the subsequent information-bearing packet. If the CTS is received by other nearby radios that are monitoring the control channel, they refrain from transmitting on the traffic channel specified in the CTS while the traffic channel is being used by the source-destination pair. If the CTS is received at the source radio, the source radio transmits the information packet on the traffic channel specified in the CTS. This transmission is made using a predetermined offset of the receiver-directed code (the *reserved* receiver-directed code) to avoid collisions with RTS packets which may be directed to the destination radio using the regular receiver-directed code. If the information packet is successfully acquired by the destination radio, it transmits an acknowledgment packet (ACK) or negative-acknowledgment packet (NACK) depending on whether the packet decodes correctly or not. The ACK or NACK is sent on the same traffic channel as the information packet using a reserved receiver-directed code.

Each radio maintains a list of traffic channels on which it is allowed to transmit, and this list is specified in the RTS that is sent to the destination radio. The destination radio compares this list to its own list of open traffic channels, and chooses a channel randomly from the channels that are common to both the lists. If there are no channels in common, then the destination radio refrains from sending a CTS.

The protocol described in the preceding paragraph can be adapted for the special case in which there is only one channel available. In this case, the control channel and the traffic channel are the same, and all the packets are transmitted on the same channel.

Under high traffic conditions, it is possible that a radio will transmit its own reservation packets repeatedly without listening to the control channel to obtain information about the availability of traffic channels. This may result in an uncoordinated use of the traffic channels, thereby defeating the purpose of the reservation packets. To prevent this situation, each radio has a wait state wherein it refrains from transmitting an RTS for a certain period of time. Each radio enters this wait state after *any* transmission attempt. The radio also enters the wait state after a reception attempt *unless* (i) it intercepts a CTS which still leaves at least one traffic channel available, or (ii) it receives an RTS to which it cannot respond (because no traffic channel from the list sent in the RTS is available). We refer to the use of such a wait state as *pacing*, and the silent period is called the *pacing time*.

Ideally, the intent of pacing is to have each radio wait for a time equal to the time taken for a complete transmission cycle so that any reservation made after such a wait will not result in the selection of a traffic channel that is already in use. However, since the pacing time will itself become a part of the transmission cycle at every radio, it is not possible for all radios to wait for a period equal to the transmission cycle. One method to mitigate this problem is to randomize the pacing time at each radio, so that the pacing time is uniformly distributed over an interval.

Pacing is also used in one other situation. If a radio is required to refrain from transmitting for a certain duration because a CTS it receives renders the last traffic channel unavailable, then the radio waits for an additional duration equal to the pacing time after the first traffic channel becomes available. The intent here is to avoid contention with other radios in the same situation. I.e., multiple radios may become aware that a traffic channel is available and initiate the reservation process at the same time, resulting in an uncoordinated use of the traffic channels. Note that in this case, it is evident that the pacing time should be randomized to be effective.

3.2. The slotted-Aloha protocol

Performance of the slotted-Aloha protocol has also been studied for comparison purposes. We assume that the radios using this protocol employ only receiver-directed codes for all transmissions. All radios that have packets to transmit do so at the beginning of the slot. All idle radios monitor the channel for incoming transmissions on the spreading code assigned to them. Any receiver that successfully acquires and decodes a packet transmits an ACK back to the source radio. The slot length is set equal to the sum of the packet length, length of the ACK, and twice the maximum propagation delay of the network. Slot synchronism may be maintained, at least in theory, by use of stable clocks and by addition of a short guard-time at the end of each slot.

4. Simulation results and discussion

To model the attenuation of radio signals, we use a simple expression for path loss given by

$$a = \left(\frac{\lambda}{4\pi r}\right)^b.$$

Here, a is the attenuation factor, λ is the wavelength of the radio signal, r is the distance between the transmitter and the receiver, and b is the path-loss index which is typically a value between 2 and 4. To study the performance of the channel access protocols, we consider the network topology shown in figure 1, in which eight radios have been placed at the vertices of a regular octagon. The diameter of this network has been set to ensure a high probability of acquisition and low probability of packet error for an interference-free transmission between radios that are diametrically across from each other. Thus, in the absence of any multiple-access interference, the network can be considered to be fully connected. It is assumed that a packet is discarded after one unsuccessful transmission attempt; i.e., there are no retransmissions. Packet generation at each source radio follows the Poisson distribution.

The performance measure considered here is the network throughput which is obtained by simulation for various values of the packet generation rate. We use a normalized generation rate λ defined as the ratio of the minimum cycle time to the mean of the inter-arrival time of the Poisson

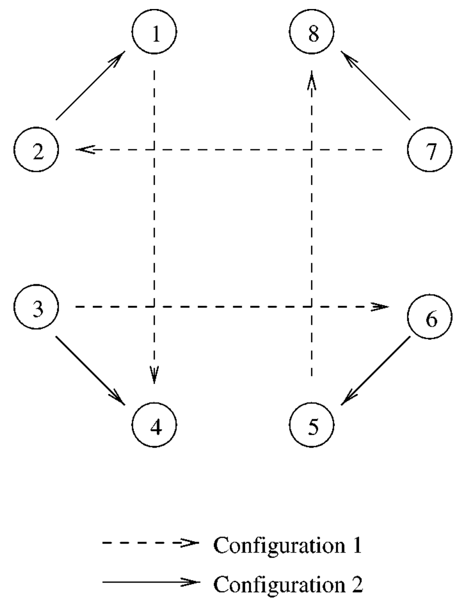

Figure 1. Octagonal network topology.

Table 1
Parameters used in the simulation.

Parameter	Value
Lowest carrier frequency	1 GHz
Total system bandwidth	4 MHz
Path-loss index	3.0
Data rate	200 kbps
Acquisition header length (M)	25 bits
Reservation packet (RTS, CTS) length	100 bits
Information packet length	1000 bits
Acknowledgement packet length	100 bits

arrival process. Here, the minimum cycle time is the time for one successful RTS-CTS transmission cycle (consisting of the RTS, CTS, information packet, and ACK, but no pacing). Throughput is defined as the product of the normalized generation rate and the probability of a successful transmission cycle at that rate.

We first consider the selection of a channel-access protocol for a network, given a fixed available bandwidth and data rate. The single-channel RTS-CTS protocol and the slotted-Aloha protocol use all the available bandwidth, and the bandwidth is modeled as equal to the chip rate $1/T_c$. For the multichannel protocol, the available bandwidth is divided equally among all the channels used. This results in a system with a chip rate equal to $1/LT_c$, where L is the number of channels. To maintain the same data rate as the higher chip-rate systems, the processing gain is reduced by the same factor L. It is assumed that the time duration of the acquisition header is the same for all the three systems. As an example of a multichannel protocol we consider a protocol that uses 4 channels. The processing gain for the system that uses the single channel protocol is set to 20, and thus the processing gain for the 4-channel system is set to 5. The parameters used for the simulation are shown in table 1. To compare the performance of these three protocols, we consider two configurations of source-destination pairs, Configuration 1 and Configuration 2 as marked by the arrows in figure 1.

A plot of throughput for Configuration 1 is shown in figure 2. In this configuration, the distance between each transmitter and the intended receiver is large compared to the distance between the receiver and the nearest interfering transmitter, thereby accentuating the near–far problem. As the packet generation rate increases, the number of collisions in the slotted-Aloha system increases and, therefore, the performance deteriorates. However, the performance of both RTS-CTS protocols is not adversely affected at high packet generation rates because of the inherent collision-avoidance mechanism. This observation is supported by the plots in figures 3 and 4, from which it can be seen that, for slotted-Aloha, the probability of not acquiring and the probability of packet error are extremely high at high packet generation rates. Comparing the two systems using the RTS-CTS protocol, it is evident that the reduced processing gain in the multichannel case results in a higher probability of packet error, but the acquisition performance of the multichannel protocol is slightly better because of

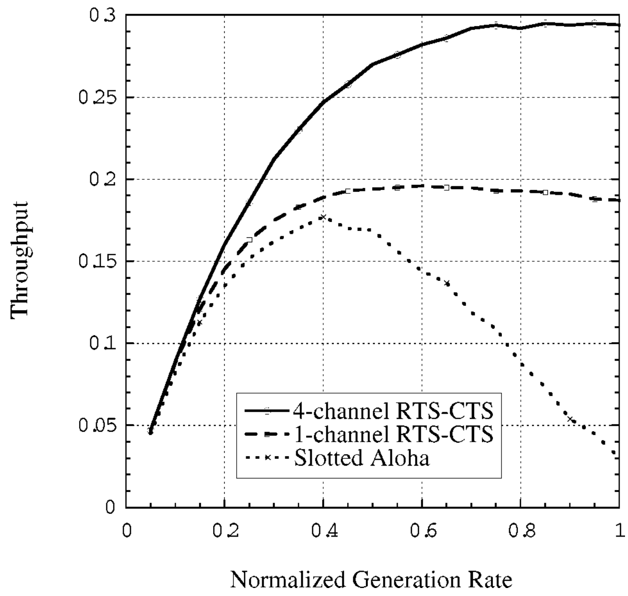

Figure 2. Throughput performance for Configuration 1.

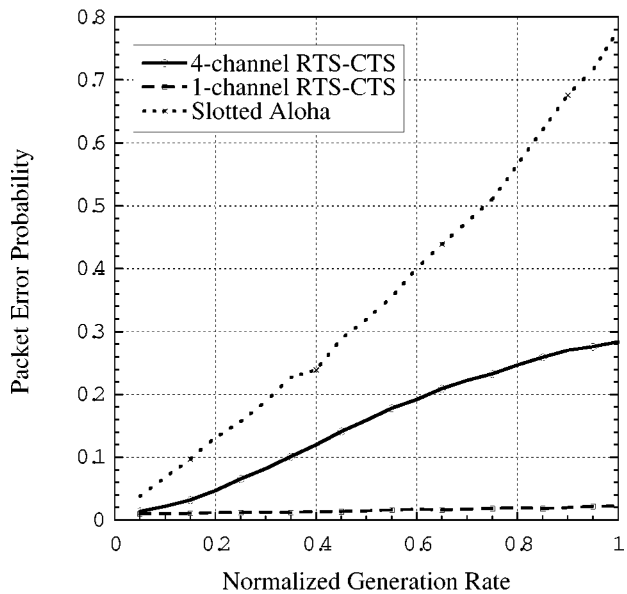

Figure 3. Packet error probability for Configuration 1.

the reduced "load" on the control channel. It is also evident that the additional channels ensure higher throughput in the multichannel case.

We consider Configuration 2 next. In this configuration, the near–far problem is not present and the receivers are relatively isolated, ensuring that the multiple-access interference is not severe. A plot of the throughput for this configuration is shown in figure 5. The performance of the system using slotted-Aloha is closely matched by that of the systems using the RTS-CTS protocols.

To study the effects of pacing, we consider three different traffic generation scenarios. For these scenarios, we consider a protocol that uses three channels (one control channel and two traffic channels). The first scenario is shown in figure 6. In this configuration, the four source

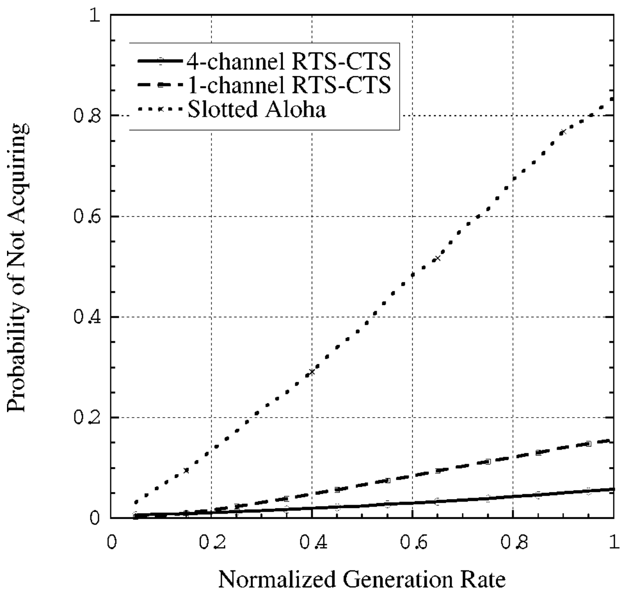

Figure 4. Probability of not acquiring for Configuration 1.

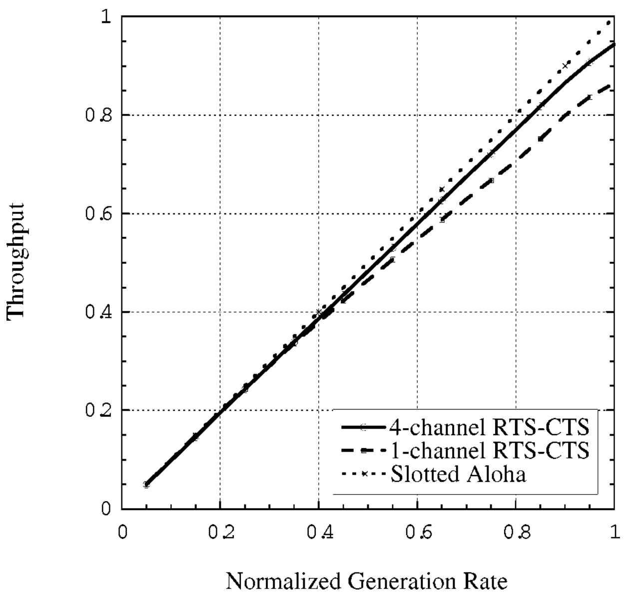

Figure 5. Throughput performance for Configuration 2.

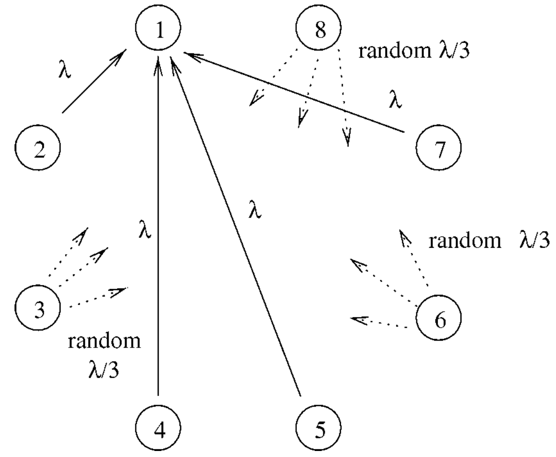

Figure 6. The network topology for the four-sources, one-destination scenario.

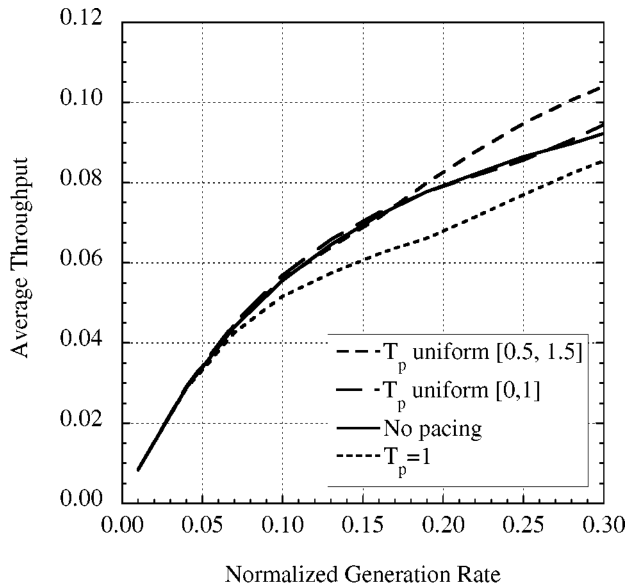

Figure 7. Average throughput for the four-sources, one-destination scenario.

radios, each at a different distance from the destination, transmit packets to the destination radio at the same generation rate. The three other radios in the network each transmit packets to random destinations at one-third this generation rate. The pacing time, normalized by the minimum cycle time, is denoted by T_p. The throughput averaged over the four source radios is shown in figure 7. The ratio of the maximum throughput at a radio (achieved at the closest source radio) to the minimum throughput at a radio (achieved at the farthest source radio) is shown in figure 8. These figures indicate that use of randomized pacing not only results in an increase in the average throughput of the network, but also reduces the discrepancy between the throughput of radios at unequal distances from the destination radio.

In the second scenario, every alternate radio along the edges of the octagon (four radios in all) in figure 1 generates traffic for a destination chosen at random from the other seven radios. The resulting throughput of the network averaged over the traffic-generating radios is shown in figure 9. In the final traffic generation scenario, all the radios in the octagonal topology generate packets bound to random destinations. The average throughput resulting in this situation is shown in figure 10. In both cases, the use of randomized pacing results in increased throughput for moderate values of the normalized generation rate. At higher generation rates, the average cycle time is greater when pacing is used, and as a result, the throughput of the network saturates.

Note that when T_p is uniformly distributed in the interval $[0.5, 1.5]$, the resulting performance in the first case shown in figure 7 is superior to that when T_p is uniformly distributed in the interval $[0, 1]$. This is because, using the $[0.5, 1.5]$ interval, each radio is silent for a longer period of

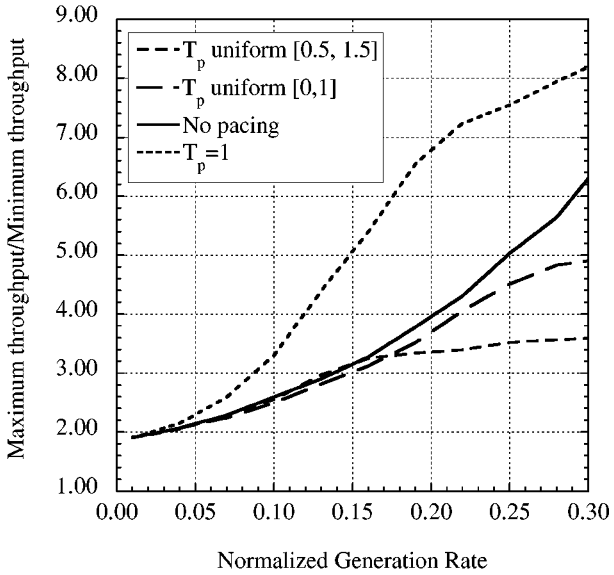

Figure 8. Throughput discrepancy for the four-sources, one-destination scenario.

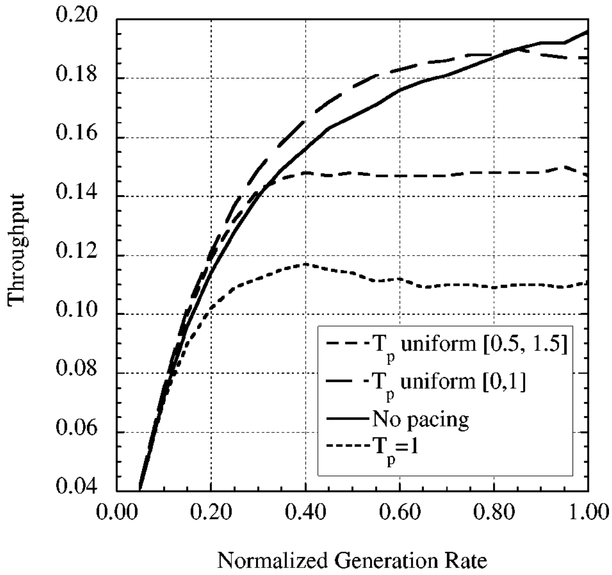

Figure 10. Average throughput for the octagonal network with eight sources and random destinations.

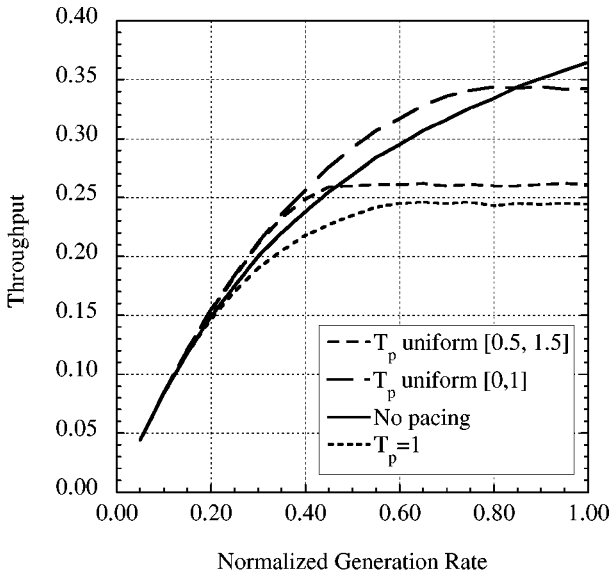

Figure 9. Average throughput for the octagonal network with four sources and random destinations.

time on an average. However, the resulting longer average cycle time causes the throughput to saturate at a smaller value of the normalized generation rate, as can be seen in figures 9 and 10. Here, both the intervals considered for randomization of the pacing time have the same *length*. Additional simulations have shown that intervals of smaller length are not as effective. Note that in all three situations, the throughput resulting from the use of a constant pacing time ($T_\mathrm{p} = 1$) is extremely poor as expected.

5. Conclusions

The performance of a reservation-based multichannel channel-access protocol has been investigated and compared to a single-channel reservation-based protocol and a simple contention-based protocol. Performance has been determined by a simulation of a direct-sequence spread-spectrum packet radio network that uses detailed models for the physical layer. It has been shown that the multichannel reservation-based protocol is robust and provides adequate throughput even under high traffic loads. A pacing mechanism that mitigates multiple-access interference has been described. It has been shown that the pacing mechanism is also effective in promoting fairness, and results in increased throughput under conditions of low or moderate traffic. It is also noteworthy that under high traffic conditions, the protocol does not become unusable, but allows the throughput to saturate.

Acknowledgments

We would like to thank Dr. H.B. Russell for his assistance in simulation development. We would also like to thank G.S. Hosangadi and Dr. D.L. Noneaker for their helpful tips on modeling the acquisition stage. We also acknowledge the contributions of Dr. J.H. Gass, Dr. M.B. Pursley and C.S. Wilkins to the packet error probability model.

References

[1] A.S. Tanenbaum, *Computer Networks* (Prentice-Hall, Englewood Cliffs, NJ, 1981) pp. 249–323.

[2] V. Bharghavan, A. Demers, S. Shenker and L. Zhang, MACAW: a media access protocol for wireless lan's, in: *Proceedings of ACM SIGCOMM '94* (1994).

[3] F. Tobagi and L. Klienrock, Packet switching in radio channels: Part III – Polling and (dynamic) split-channel reservation multiple access, IEEE Trans. Commun. COM-24 (August 1976) 832–845.

[4] C. Namislo, Analysis of mobile radio slotted ALOHA network, IEEE J. Select. Areas Commun. 2 (July 1984) 563–574.
[5] N. Gower and J. Jubin, Congestion control using pacing in a packet radio network, in: *Proceedings of MILCOM '82* (1982).
[6] A.R. Raghavan, C.W. Baum and D.L. Noneaker, Analysis of a serial acquisition scheme for unslotted distributed direct-sequence packet radio networks, in: *ICC '98 Conference Record* (1998) pp. 35.5.1–35.5.5.
[7] M.B. Pursley and D.J. Taipale, Error probabilities for spread-spectrum packet radio with convolutional codes and Viterbi decoding, IEEE Trans. Commun. 35 (January 1987) 1–12.
[8] I.M. Onyszchuk, On the performance of convolutional codes, Ph.D. thesis, California Institute of Technology (1990).

Arvind R. Raghavan graduated from the Karnataka Regional Engineering College, Surathkal, India, with a B.E. degree in electronics and telecommunications in July 1992. He received the M.S. degree in electrical engineering from Clemson University, Clemson, USA, in August 1996. He is currently a graduate research assistant in the wireless communications group at Clemson University, where he is working towards a doctoral degree in electrical engineering. His research interests include applications of error control coding, analysis and simulation of direct-sequence CDMA systems, medium-access protocols and routing protocols for distributed ("ad hoc") networks.

Carl W. Baum received the B.S. degree in electrical engineering (with highest distinction) from the University of California, Los Angeles, in 1987, and he received the M.S. and Ph.D. degrees in electrical engineering from the University of Illinois, Urbana-Champaign, in 1989 and 1992, respectively. He was a University of Illinois Fellow from 1987 to 1988. From 1987 to 1992, he held a graduate research assistantship at the Coordinated Science Laboratory, University of Illinois, Urbana-Champaign. In August 1992, he joined Clemson University, where he is currently an Associate Professor in Electrical Engineering. His current research interests are in the general areas of detection theory and communications theory with emphasis on spread-spectrum communications. Dr. Baum received the Robert Chien Memorial Award from the University of Illinois for excellence in research in the field of electrical and computer engineering in 1992. In 1996, he was named recipient of the IEEE Browder J. Thompson Prize Award for the outstanding paper by authors under thirty years of age. He is a member of Tau Beta Pi, Eta Kappa Nu, and Phi Eta Sigma.