Proceedings of the 40th IEEE
Conference on Decision and Control
Orlando, Florida USA, December 2001

Protocols for Media Access Control and Power Control in Wireless Networks[1] [2]

Vikas Kawadia, Swetha Narayanaswamy, Robert Rozovsky, R. S. Sreenivas
and P.R. Kumar
Department of Electrical Engineering, and
Coordinated Science Laboratory,
University of Illinois at Urbana-Champaign
USA

Abstract

We present two protocols for ad hoc wireless networks, one for the media access control problem, and the other for the power control problem.

For the media access control problem we present a protocol called SEEDEX which does not explicitly make reservations for packets, a la the IEEE 802.11 protocol, yet allows scheduling to minimize conflicts. The idea is to use known finite state machines at nodes which are driven by pseudo-random number generators. The seeds of these pseudo-random number generators are exchanged between nodes in a two hop neighborhood. A further refinement is a hybrid version, which employs SEEDEX only on the RTS-CTS handshake of the IEEE 802.11 protocol. This algorithm provides improved throughput-delay and delay jitter performance in an ns simulation.

For the power control problem we first provide a framework for conceptualizing the problem. This leads us to propose a network layer approach to power control which consists of finding the least common network wide power level at which all nodes are connected. This can be shown to maximize the throughput traffic carrying capacity of the network. We then propose a feedback algorithm COMPOW which tunes to this minimum power level adaptively. We also propose a software architecture for integrating this into the OSI protocol stack. The new idea is to introduce a parallel analog of the hierarchical OSI layers into the network layer which still allows modularity, and usability with any routing table driven routing algorithm. We also describe our implementation, which takes advantage of the port demultiplexing service provided by the transport layer.

[1]Please address all correspondence to the last author at University of Illinois, CSL, 1308 West Main St., Urbana, IL 61801-2307. Email: prkumar@uiuc.edu

[2]This material is based upon work partially supported by the Office of Naval Research under Contract No. N00014-99-1-0696, and the U.S. Army Research Office under Contract Nos. DAAD19-00-1-0466 and DAAD19-01010-465, and DARPA under contract Nos. N00014-01-1-0576 and F33615-01-C-1905. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the above agencies.

1 Media Access Control in Wireless Networks: The SEEDEX Protocol

The wireless medium distinguishes itself from the wired medium by the fact that it is a shared medium. Hence transmissions can interfere with each other. Consider the situation shown in Figure 1.

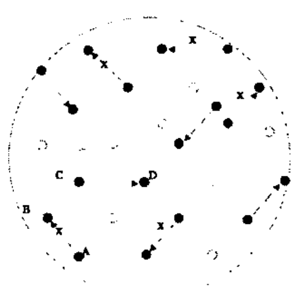

Figure 1: Collisions between simlutaneous transmissions

Node B cannot recieve a packet from node A at the same time that node C is transmitting to node D, if B is within range of C. Thus, only certain sets of simultaneous transmissions can take place in a wireless network.

The media access control problem is to schedule transmissions online in a distributed asynchronous manner so that packets reach their intended one hop neighbor recipients without collisions.

In the IEEE 802.11 protocol, see [5], an RTS-CTS handshake is used. The transmitter sends an RTS

0-7803-7061-9/01/$10.00 © 2001 IEEE

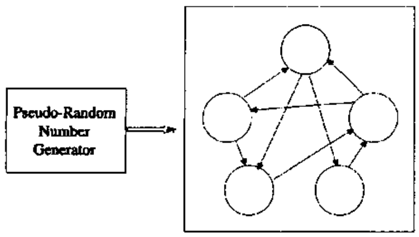

Figure 2: A finite state machine runs at each node which is driven by a pseudo-random number generator.

(request-to-send) packet when it has a DATA packet to be sent to a receiver. Upon hearing this, all nodes in the vicinity of the transmitter stay silent, while the receiver sends back a CTS (clear-to-send) packet. This also silences all nodes in the vicinity of the receiver. Upon hearing the CTS, the transmitter successfully sends its DATA packet to the receiver which then replies back with an acknowledgment.

This protocol thus effectively silences two neighborhoods, one of the transmitter and the other of the receiver, each and every time it has a packet to send. On small networks this could be a substantial portion of the whole network, leading to a loss in throughput. Similarly, the RTS-CTS handshake also incurs overhead. A previous study, see [1], has exhibited a throughput scaling of $\frac{c}{n^{1.7}}$ bits/sec as the number of nodes n is increased from two to twelve. This compares poorly with the optimal scaling law of $\frac{c}{\sqrt{n \log n}}$ bits/sec achievable in principle; see [4].

With the goal of reducing this overhead and thus improving efficiency, we present a protocol, called SEEDEX, introduced in [2], for the media access control problem.

The SEEDEX algorithm attempts to make reservations without, paradoxically, explicitly making them a la IEEE 802.11. The basic idea is to have a known finite state machine running at each node. These machines are driven by pseudo-random number generators, which are essentially linear congruence relations; see Figure 2. If the seed of the pseudo-random number generator at a node is known, then one knows the future state sequence of the state machine at that node. This exchanging of the seeds within a two-hop neighborhood can be done by a two-state fan-in and fan-out procedure; see Figure 3.

Thus all nodes within two hops of each other, including the so called "hidden" and "exposed" terminals know each other's states.

With such knowledge of the states of all nodes in a two-hop neighborhood of itself, a node can opportunistically schedule its transmissions to minimize conflicts. This leads to the SEEDEX protocol.

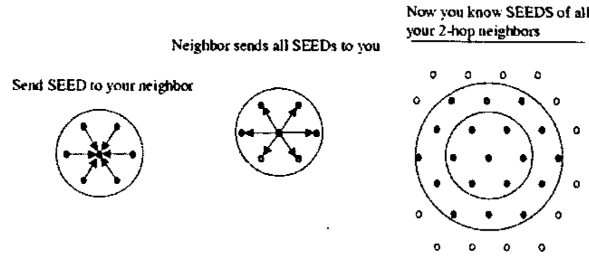

Figure 3: The seed exchange procedure.

Let us consider a particularly simple variant when each node wanders back and forth between a *Listen* and a *Permission* to Transmit state, as an i.i.d. Bernoulli process with probability p of being in the Permission to Transmit state, and a probability $(1-p)$ of being in the Listen state. When a node T has a packet to send to a neighboring node R, it waits for a time when it (T) is in the Permission to Transmit state, and its intended recipient R is in the Listen state, and then transmits the packet in the slot with a probability $\frac{\alpha}{n+1}$ where n is the number of other neighbors of R that are also in the Permission to Transmit state; see Figure 4.

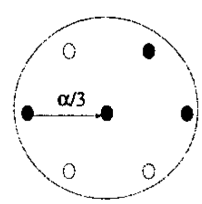

Figure 4: The node transmits with probability $\frac{\alpha}{3}$ when two other neighbors of R are also in the Permission to Transmit state.

The value of α should be low (about 1.5) in heavy traffic to minimize the number of conflicts, and high (about 2.5) in light traffic to minimize the number of idle slots, while the $\frac{1}{n+1}$ scaling ensures that the probability of one successful transmission is maximized when all the $(n+1)$ nodes do have a packet to transmit to R. One should note that this only approximates the true scenario since the other nodes may have different intended recipients; see Figure 5.

One can further pursue this idea by considering a hybrid with IEEE 802.11. One can employ the SEEDEX

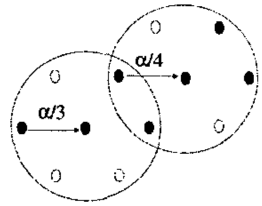

Figure 5: Two nodes with different intended recipients choose different probabilities for transmission.

protocol only for the RTS and CTS handshake packets. After this, DATA packets can proceed uninterrupted just as in IEEE 802.11. Thus, consider a situation as in Figure 6, where nodes contend using SEEDEX for the RTS and CTS slots, which alternate with each other, until a successful reservation is made, at which point a DATA packet straddling a fixed number of multiples of the RTS and CTS slot durations is thereafter transmitted. The next slot is then reserved for an ACK. Following this, another contention period for the RTS-CTS handshake ensues, again employing SEEDEX.

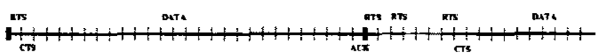

Figure 6: Using SEEDEX for the RTS-CTS handshake.

This results in a procedure with better throughput as well as delay and delay jitter (i.e., standard derivation of delay) characteristics, in an ns simulation. These are shown in Table 1 as a function of loading, for the case of three intersecting flows shown in Figure 7, in an NS simulation.

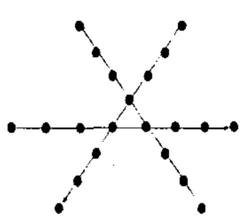

Figure 7: Three intersecting flows.

The SEEDEX protocol has been implemented at the University of Illinois. One challenge we faced was to synchronize slots on available off the shelf hardware such as Lucent's WaveLan cards or CISCO' Aironet cards which use IEEE 802.11. One may note that synchronized slots are in fact used in the Bluetooth standard. Another challenge was to overcome the carrier sensing feature built into these cards.

Table 1: Performance of SEEDEX compared with IEEE 802.11 standard : Mean delay and Standard deviation of delay (Delay jitter).

	Mean	Mean	Deviation	Deviation
Throughput	SEEDEX	802.11	SEEDEX	802.11
0.2	15.52	24.34	2.85	18.68
0.3	15.74	21.56	3.08	13.61
0.4	15.50	20.34	2.90	11.59
0.5	15.54	24.04	2.97	15.54
0.55	15.64	30.13	3.29	21.01
0.6	33.63	809.9	18.93	748.77

2 Power Control in Wireless Networks: A conceptualization, the COMPOW Protocol, and implementation

Consider now the power control problem for ad hoc networks. For some earlier approaches to power control, we refer the reader to [6], [7], and [8].

We begin with a framework for conceptualizing the issues involved. For a variety of reasons we argue that all the nodes in a network should employ a common power level (we are assuming identical nodes). First, a proper functioning of the RTS-CTS handshake requires that. For example in Figure 8, R's CTS is not powerful enough to be heard by A, and so does not silence it. Yet, if A thereafter transmits, it collides at R with T's packet.

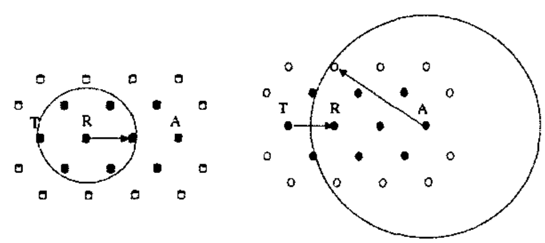

Figure 8: R should transmit at at least the same power as A in order to silence it. R's CTS is not powerful enough to be heard by A. Yet A can interfere with R if its power level is higher.

Second, link level acknowledgments are necessary in wireless networks. The wireless medium is inherently unreliable and so without such acknowledgments the transport layer would function very poorly. Also, the presence of conflicts at the MAC layer necessitates such acknowledgments, as done, for example, in IEEE 802.11. For such link level acknowledgments to reach the original transmitter T from R, R's power level must be at least equal to that of T, as argued in Figure 9.

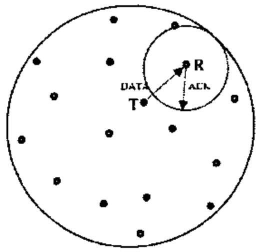

Figure 9: If R's power level is less than T's, then R's ACK may not be heard by T

For these reasons and some others, the power levels of two neighboring nodes must be equal. By transitivity, this equality continues to hold in a two-hop neighborhood, and by induction in an n-hop neighborhood and thus throughout the entire network.

The next issue: What is this common power level to be? As Figure 10 shows, if this common power level is too low, then there may not be a sufficient number of links in the network to render it connected.

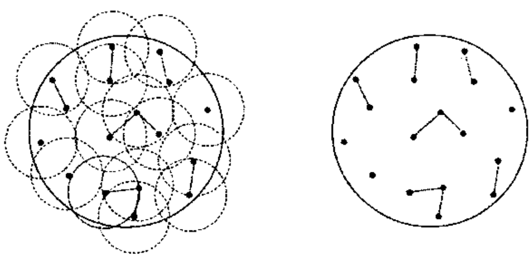

Figure 10: If the power level, and thus the range (assumed omnidirectional for illustration only) is too small, then the resulting network may not be connected.

On the other hand, if the power level is too high, then as Figure 11 shows, the network may have too many links. This increases the degree of each node (where the degree is the number of edges incident to a node). When the degree is large it simply means that there are several nodes whose transmissions can cause interference at the given node. To calculate the effect of this, consider the following simplified calculation (a more detailed analysis can be found in [4]).

Consider a network with n nodes, each capable of of transmitting at a data rate of W bits/sec, in a disk of area A square meters. Let us also suppose that all transmitters have a common range r. Further, suppose that if node R is to successfully receive a transmission from T at distance $\bar{r}$, then there can be no other transmission within a distance $(1+\Delta)\bar{r}$ from R, where $\Delta > 0$. Now consider two simultaneous successful transmissions, one from T to R at a distance $\bar{r}$, and another from T' to R' at a distance $\bar{r}'$, as shown in Figure 12.

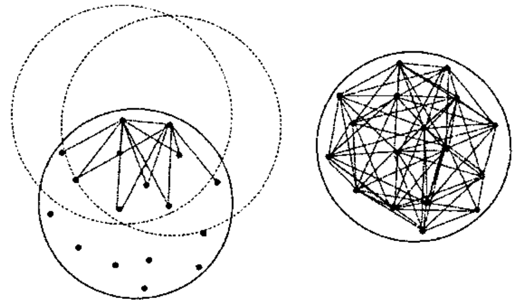

Figure 11: If the power level is too high, there can be too much interference.

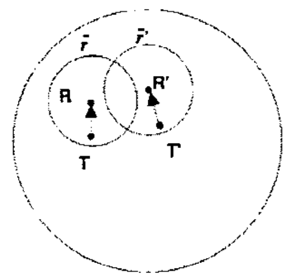

Figure 12: Two simultaneous successful transmissions.

Thus,

$$\begin{aligned} |T-R| &= \bar{r} \\ |T'-R'| &= \bar{r}' \\ |T-R'| &\leq (1+\Delta)\bar{r}' \\ |T'-R| &\leq (1+\Delta)\bar{r}, \end{aligned}$$

must hold. From this, an application of the triangle inequality shows that

$$\begin{aligned} |R-R'| &\geq |R-T'| - |R'-T'| \\ &\geq (1+\Delta)\bar{r} - \bar{r}'. \end{aligned}$$

Similarly,

$$|R-R'| \geq (1+\Delta)\bar{r}' - \bar{r}.$$

Thus

$$|R-R'| \geq \frac{\Delta(\bar{r}+\bar{r}')}{2}.$$

Hence a circle of radius $\frac{\Delta\bar{r}}{2}$ around R, and a circle radius $\frac{\Delta\bar{r}'}{2}$ around R' are disjoint, as shown in Figure 13.

For simplicity of exposition, suppose that $\bar{r} = \bar{r}' = r$ (see [4] for a more general analysis). Then note that

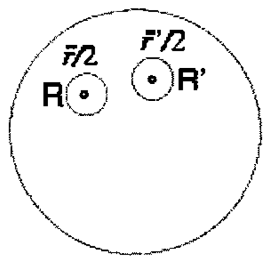

Figure 13: Circles around receivers are disjoint. Hence every successful transmission consumes area.

even when such circles are located near the periphery of the domain, at least a quarter of the area, or $\frac{\pi\Delta^2 r^2}{16}$ must lie within the domain. Since the available area in the domain is only A, it implies that there can only be $\frac{16A}{\pi\Delta^2 r^2}$ simultaneous transmissions, each of W bits/sec.

Now suppose that the network carries traffic at a rate λ bits/sec from each of the n nodes to their destinations, which are at an average distance $\bar{L}$ away. Then such sessions must, on average, involve $\frac{\bar{L}}{r}$ hops. Thus each node must transmit on average at $\frac{\bar{L}\lambda}{r}$ bits/sec, to carry the relaying burden.

Hence the average throughput of λ per node can be supported by the network only if

$$\frac{n\bar{L}\lambda}{r} \leq \frac{16AW}{\pi\Delta^2 r^2},$$

or

$$\lambda \leq \frac{c}{rn}$$

for some constant c.

This suggests that the common range r must be made as small as possible. However, as we have seen earlier, too small a value for r can render the network disconnected. So the question that arises is: How small a value of r will still guarantee connectivity? This question has been answered in [3]. Consider n nodes located at random (i.i.d. uniformly distributed) in a disk of area A. If all nodes have a common range $r(n)$, then

$$\lim_{n\to+\infty} \text{Prob(Network is connected)} = 1$$

if and only if

$$r(n) = \sqrt{\frac{A(\log n + k(n))}{\pi n}} \text{ where } k(n) \to +\infty.$$

From this it follows that the maximal throughput that the network can support per node is

$$\lambda(n) = O\left(\frac{1}{\sqrt{n \log n}}\right) \text{ bits/sec.}$$

In practice the number of nodes in the network may be unknown, the area of the domain may be unknown, the range may not be omnidirectional, etc. Hence instead of dealing with the range r, one wants to deal directly with the power level P, noting that when power levels are identical, then one node is in the range of the other if and only if the other is in its range, without requiring omnidirectionality.

Thus one can formulate the network-wide power control problem as follows: What is the smallest common power level P which ensures that the network is connected?

What we further require is a distributed asynchronous algorithm which converges to such a power level. Note also that since nodes are mobile, this algorithm will have to be adapt constantly to its environment, and thus be running continuously.

We propose the following network-wide feedback strategy to converge to the smallest power level at which the network is connected. Let $\{P_{min} = P_0, P_1, \ldots, P_J = P_{max}\}$ be the discrete set of power levels available at each node, arranged in the order $P_i \leq P_{i+1}$. For each level P, denote by $\mathcal{R}(P)$ the set of nodes which are connected to a distinguished node, when all nodes use the common power level, and the connections are allowed to be multi-hop. Thus $\mathcal{R}(P)$ is the "reachable set" at a common power level P.

Note that $\mathcal{R}(P_{max})$ is the maximal reachable set. The feedback strategy is simply this: Set the current power level to the smallest level at which the network is connected, that is

$$P(t) = \{ P_i \text{ such that } \mathcal{R}(P_{i-1}) \subsetneq \mathcal{R}(P_{\max})$$

Here t denotes an instant of some set of sampling times at which the power levels are changed. We call this the COMPOW (Common Power) power control algorithm. (We can also modify the algorithm so that the power level is k level higher than that above, which will provide k levels of robustness).

The next issue that arises is how one may obtain $\mathcal{R}(P)$ in practice. Fortunately this is readily obtained from any routing table driven routing algorithm, by simply examining the number of entries in the routing table.

There is yet another architectural, and important, issue that needs to be confronted: Where in the current OSI layers should power control be situated for the case of wireless networks?

One should note that, a priori, the problem of power control cuts across many of the carefully delineated layers. For example, power control is a link layer problem, as for example in cellular systems. However, as seen earlier, it is also a network level problem since network-wide connectivity depends on the choice of the power level. Additionally, power control impacts on the routing problem since different choices of power levels give rise to different sets of links.

Conceptually, suppose that the network layer has a routing table for each power level. Then, a routing table at each power level P is built by sending routing control packets at that power level. The table built using the current power level $RT_{P(t)}$ is the primary routing table, and is consulted for all DATA packets.

Our solution uses structural parallelism in the network layer, and thus preserves modularity, which has been a governing principle in network protocol design. *Parallel modularity*, that is, parallel running of several independent non-interacting routing algorithms, is equally amenable to the plug and play approach of the serial hierarchy of the OSI layer.

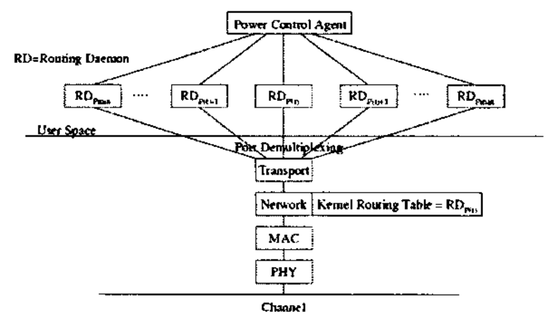

Figure 14: Implementation of COMPOW illustrating Parallel Modularity.

We implement the above concept as follows. A routing daemon corresponding to each power level is run at different ports; see Figure 14. (For example in the CISCO's Aironet 350 cards, there are six power levels available).

The multiplexing of route packets for the correct power level is automatically taken care of by the transport layer port demultiplexing. The power control agent decides the current network power level and the routing table of the corresponding routing daemon is copied to the kernel routing table. This will ensure that DATA packets follow the appropriate power levels, while each routing daemon maintains the routing table appropriate for its own power level. Finally, the implementation is completed by ensuring that the physical setting of the power level is done at the level of the device driver. For this a power level field is required in the packets.

Thus, we have provided a conceptual foundation, a mathematically sound solution, and an architecturally clean implementation which is compatible with current implementations of the protocol stack is software. This protocol has been implemented at the University of Illinois.

References

[1] P. Gupta, "Design and Performance Analysis of Wireless Network," Ph.D. Thesis, Department of Electrical and Computer Engineering, University of Illinois, Urbana-Champaign, August 2000.

[2] P. R. Kumar, "New Technological Vistas for Systems and Control: The Example of Wireless Networks," *IEEE Control Systems Magazine*, Special Section on Networks and Control, pp. 24–37, vol. 21, no. 1, Feb. 2001.

[3] P. Gupta and P. R. Kumar. Critical power for asymptotic connectivity in wireless networks. In W. M. McEneany, G. Yin, and Q. Zhang, editors, *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming*. Birkhauser, Boston, MA, 1998. ISBN 0-8176-4078-9.

[4] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, IT-46(2):388–404, Mach 2000.

[5] IEEE Protocol 802.11. Draft standard for wireless LAN: Medium access control (MAC) and physical layer (PHY) specifications. IEEE, July 1996.

[6] V. Rodoplu and T. H. Meng. Minimum energy mobile wireless networks. *IEEE Journal on Selected Areas in Communications*. vol. 17, issue 8, pp. 1333–1344, Aug. 1999.

[7] R. Ramanathan and R. Rosales-Hain. Topology control of multihop wireless networks using transmit power adjustment. Proceedings of *INFOCOM 2000: Nineteenth Annual Joint Conference of the IEEE Computer and Communication Societies*. vol 2, pp. 404–412, 2000.

[8] R. Wattendorfer, L. Li, P. Bahl, and Y-M. Wang. Distributed topology control for power efficient operation in multihop wireless ad hoc networks. To appear in *INFOCOM 2000: Twentieth Annual Joint Conference of the IEEE Computer and Communication Societies*, Alaska, April 2001.