Assignment 4: WiFi and Fairness
Due the week of October 3/4, 2012

The version of Ethernet we all use today is standardized by the Institute of Electrical and Electronics Engineers (IEEE). The IEEE network standards committees are all assigned numeric codes which then become the official identifiers for the protocols they produce. Ethernet, for example, is officially known as the IEEE 802.3 standard.

This week, we will spend some time discussing two other members of the IEEE 802. family of protocols including

- 802.11 (the standards for WiFi), and

- 802.6 (a standard for slightly-larger-than-local area networks, otherwise known as MANs (Metropolitan Area Networks) — think of your cable company and cable modems).

In the current market, Ethernet and WiFi are clearly the dominant local area network protocols. These two protocols are also very similar in important ways. In particular, they both are broadcast networks and they both depend on backoff and retransmissions to deal with packet collisions. Finally, one very important characteristic they both share is that they are currently widely deployed

The other protocols we will consider this week are not widely used. One is a new proposal that might someday be incorporated to improve the efficiency of wireless networks. The other, IEEE 802.6 is essentially defunct. It was expected to be widely used when it was standardized, but failed commercially for both technical and economic reasons.

Despite its commercial failure, IEEE 802.6 is interesting to study for several reasons. While Ethernet and WiFi depend on randomized competition to allocate bandwidth among transmitters, the 802.6 standard was based on reservation mechanisms that could provide more predictable guarantees of service than randomized mechanisms. It also used a different network topology than Ethernet (a directed linear buses). In addiion, 802.6 is one of many examples of a slotted network. Most importantly, studying this network standard will allow us to examine the relationship between a network's efficiency, its "normalized channel propagation delay", and its fairness. We will examine the issue of fairness in several ways this week.

To prepare for this week, you should first read §2.7 of the text which describes a variety of wireless networking technologies. They provide the most details (and we will focus most of our attention on) WiFi networks, but they also discuss Bluetooth and cellular networks.

As usual, I have some readings beyond the text book. These readings will both provide some additional background on the IEEE standards (both 802.11 and 802.6) and provide information about some other LAN proposals.

For more details on wireless networks, I ask you to read pp. 116-120 of

- B. P. Crow, I. Widjaja, J. G. Kim, and P. T. Sakai. IEEE 802.11 Wireless Local Area Networks. *IEEE Communications Magazine*, pages 116–126, September 1997

The next paper:

- Abeysundara, Bandula W. and Ahmed E. Kamal, High-Speed Local Area Networks and Their Performance: A Survey. *ACM Computing Surveys*, 23:2 (June 1991), pp. 221-264.

discusses several local area network proposals including DQDB, the basis for IEEE 802.6/ You don't need to read all of this paper. First read §2 and then skip ahead to read the paper's discussion of DQDB (§3.2.6). Next, skip to the conclusion of the discussion of bus based networks (§3.2.9).

I'd like you to read the first two sections of a paper that describes a proposal for improving the DBDQ network:

- Kabetepe, Mete and Kenneth S. Vastola, The Fair Distributed Queue (FDQ) Protocol for High-Speed Metropolitan-Area Networks, *IEEE/ACM Transactions on Networking.* 4:3 (June 1996), pp. 331-339.

Finally, please read the first three sections of

- Heusse, Martin, Franck Rousseau, Romaric Guiller, and Andrzej Duda, Idle Sense: An Optimal Access Method for High Throughput and Fairness in Rate Diverse Wireless LANs, *SIGCOMM 2005*, August 2005.

This paper describes an alternative to the binary exponential backoff algorithm used by both WiFi networks and Ethernet.

Exercises

1. Consider the description of the hidden node and exposed node problems given on page 137 of our text. According to the text, "802.11 addresses these problems by using CSMA/CA."

   In fact, it is not clear that 802.11 addresses the exposed node problem. Given a set of network stations like those shown in Figure 2.31, if C hears B send an RTS to A but does not hear any CTS from A, C can deduce that either it is out of range to interfere with A or B's transmission failed for some other reason and conclude that it should be safe for C to send to D even if B is sending to A at the same time because A will only hear B and D will only hear C. While this argument implies that the data can be received correctly, what about the acknowledgments? If C is still sending after B finishes, its transmission will interfere with the acknowledgment A sends to B. B will then believe that it has to retransmit its frame.

   (a) Can this actually happen in the 802.11 protocol? Precisely explain how this could happen or why it could not happen.

   (b) In part (a) we were assuming that C heard B's RTS. If B sends an RTS to A at exactly the same time that C sends and RTS to D, C will not hear B's RTS (because stations don't listen while sending). Would either B or C or both complete their transmissions successfully in this case?

2. At the very end of their discussion of collision avoidance in 802.11 (p. 139), Peterson and Davie mention that in the event that a station does not receive a CTS frame in response to a RTS, it waits an amount of time determined "by the same exponential backoff algorithm used on the Ethernet." This suggests that there must be a "MACA capture effect" similar to the Ethernet capture effect explored by several of last week's problems. In fact, the designers of 802.11 modified the backoff algorithm to avoid the capture effect. Briefly explain the change(s) they made.

3. Even with the capture effect, the Ethernet backoff algorithm ensures that if several computers connected to an Ethernet are all continuously trying to send packets, they will each, over a sufficiently long period of time, be able to send roughly the same number of packets. In some

sense, therefore, the collision detection and retransmission process used by both Ethernet and the 802.11 protocols is "fair". Fairness, however, is a slippery concept.

(a) Consider the data in figure 3-3 of the paper "Measured Capacity of an Ethernet: Myths and Reality" by Boggs, Mogul, and Kent. Suppose that exactly two hosts, A and B, are both continuously sending 1536 byte packets as quickly as they can. Approximately how many megabytes per second will each host be able to transmit? (Yes, this is meant to be a trivial question.)

(b) In their paper, Boggs, Mogul, and Kent include data from experiments in which they used a "bimodal" distribution of packet lengths. That is, instead of having the computers all transmit packets of a single length, they used packets of two different length, 64 bytes and 1536 bytes. Suppose that we modify the scenario described in part (a) so that host A continues to constantly transmit 1536 byte packets as quickly as possible, but host B instead constantly transmits 64 byte packets. How many megabytes per second will each host be able to transmit? Is this fair?

(c) Suppose that two stations, A and B, are constantly trying to send maximum length packets to two access points on an otherwise idle 802.11g network at 54Mbps. Assuming that interrupt processing leads to transmissions without collisions (as in the 2 host case on the network described by Boggs et.al.), how many megabytes per second will each host be able to transmit? (Yes, this is also meant to be a trivial question.)

(d) As Peterson and Davie explain in §2.7.1 (on page 136 under the heading "Physical Properties"), the 802.11 wireless standard supports several different transmission rates including 2 Mbps, 11Mbps, and 54Mbps. In fact, most WiFi products support all of these data rates. One advantage of this support is that if environmental conditions make it impossible for two WiFi nodes to communicate at one of the higher data rates, they can switch to one of the slower rates to increase transmission reliability.

Suppose that two stations, A and B, are constantly trying to send maximum length packets to two access points on an otherwise idle 802.11 network but that while A is able to transmit at 54Mbps, B has been forced to lower its transmission rate to 2Mbps. Still assuming that interrupt processing leads to transmissions without collisions, how many megabytes per second will each host be able to transmit? Is this fair?

4. In their description of the "Idle Sense" alternative to the exponential backoff algorithm used by 802.11 networks, Heusse et.al., provide an analysis leading to a formula for the expected number of idle slots a station should observe between collisions and or transmission on the network. This analysis includes the derivation for a formula for P_e^{OPT}, the optimal value for the probability with which an station that has data to sends attempts a transmission in a given slot.

In the original Ethernet paper, Metcalfe and Boggs state that "We assume that a queued station attempts to transmit in the current slot with probability $\frac{1}{Q}$ ...; this is known to be the optimum statistical decision rule..." In other words, Metcalfe and Boggs claim that $P_e^{OPT} = \frac{1}{Q}$ where Q is the number of stations with packets to transmit.

(a) First, let's assume that Metcalfe and Boggs are right (since they are certainly more famous than Heusse et.al.). Assuming that $P_e^{OPT} = \frac{1}{Q}$ derive an expression for the expected number of consecutive idle slots that will be observed on an Ethernet.

(b) Compute (and if possible plot) the actual values predicted by your answer to (a) for reasonable values of Q (1-50?). You may want to write a little program or script to do this.

(c) Explain why the formula for P_e^{OPT} given by Heusse et.al. is different from $\frac{1}{Q}$. How do their assumptions differ from those made by Metcalfe and Boggs? In particular, under what conditions would the two papers agree on the correct value for P_e^{OPT}?

5. The authors of the FDQ paper claim to have designed their protocol in response to an inherent unfairness in the original DQDB protocol. Give a specific example of how such an unfair allocation of bus bandwidth could occur in the original DQDB protocol. Try to construct a scenario involving as few active stations as possible. Quantify the degree of unfairness present in terms of the parameters that describe the network (transmission rate, propagation time, packet size, etc.). That is, if possible give a formula for the ratio of the fraction of the bandwidth received by the most favored station to that received by the least favored.

6. The authors of the FDQ paper also claim to have eliminated such unfairness in the design of their protocol. Unfortunately, in this world little comes for free. Demonstrate that the FDQ protocol sometimes has to sacrifice efficiency for fairness. As for the last question, construct a scenario in which network bandwidth is wasted by FDQ using as few active stations as possible. Again, quantify the amount of bandwidth wasted.