Medium Access Control, cont'd

1 Review from Previous Class

The MAC Layer exists to provide more reliable communication among nearby devices over the wireless channel and runs on top of the physical layer. There are a few different implementations of the MAC Layer. One notable implementation is Carrier-Sense Multiple Access (CSMA).

CSMA consists of link-level acknowledgements and time reservation packets. It uses either RTS/CTS/Data or RTS/CTS/Data/Ack sequences to avoid collision.

Gallager considers coding theory and some basic differences and issues between coding theory research and collision research. He points out some basic problems in the existing research. Gallager points out that coding theorists well take into account the channel characteristics of the communication medium but fail to consider random and bursty arrivals. On the hand, collision research takes well into account the random and bursty nature of arrivals but does not fully consider the channel.

In coding theory, most consider only the case of two senders who are sending continuously. It is able to characterize achievable rates of sending assuming very low probabilities of error. Coding technology simply doesn't exist at this point due to several issues. Firstly, there should be more than just two transmitters. Secondly, the system should be able to handle random arrivals from oustide the existing sources. Thirdly, there should be a scenario where a small and changing subset of users want to use the channel randomly.

Within collision avoidance research, Gallager points out some assumptions that are made. In the slotted system approach, there is need for time synchronization but there is no difference made between short and long packets. Furthermore, there is no consideration given to the channel since any packet recieved in an analyzed system is alays recieved.

As discussed in last class, there are a few noteworthy solutions for MAC. These include slotted ALOHA, splitting algorithms (including first come first serve) and carrier sensing. In this class, we continue to look at MAC papers.

2 Komlos, Greenburg Paper

This is a four page paper that discusses a nonadaptive algorithm for conflict resolution. That is, given a bound on k, the number of transmitting stations, they are able to show that a nonadaptive algorithm can be just as good as an adaptive algorithm.

2.1 The Problem

The general problem formation is done with formulated using n total devices with slots number $1, 2, 3, \ldots$. Some number k of the devices n have messages to send. This subset may be denoted by the set I. That is, $|I| = k$. In this system, the input arrives and the devices in set I immediately know they have something to transmit.

In the algorithm, at each step, some stations among those in I can try to transmit:

- If no devices try, then no packets are transmitted.
- If exactly one tries, the packet is successfully bcast.
- If two or more try, collision, no packet is successfully bcast.

Moreover, in any case, all stations receive feedback $(0, 1, c)$, immediately.

The cost measure is the worst-case number of slots befre every station succeeds in transmitting. In the original tree algorithms, the worst-case bound was $O(k + k \log n/k)$. But, the nodes used feedback—using $(0, 1, c)$ results from earlier slots to decide on whether to transmit in next slot.

2.2 The Proposed Nonadaptive Case

This paper considers the nonadaptive case. That is, the "algorithm" is simply a list $Q_1, Q_2, \ldots, Q_t$ of query sets, chosen ahead of time, independently of the input set I.

A bit more formally: I_j is defined to be the set of contenders that still remains to transmit, after slot j. Thus, I_0, is the set that wants to transmit at the beginning, is defined by $I_0 = I$. But then the later sets are defined recursively, by: $I_j = I_{j-1} - Q_j$ if $|I_{j-1} \cap Q_j| = 1$, and $I_j = I_{j-1}$ otherwise.

If $|I_{j-1} \cap Q_j| = \{x\}$, a single station, we say that slot j *isolates* x. That means it succeeds in letting it transmit alone. The goal is to isolate all the stations in I. Cost measure is the total length t of a sequence that always succeeds, for any I, in isolating all the stations in I.

2.3 A Fast Nonadaptive Algorithm

In this paper, they show that there is the *existence*of an algorithm t of the sequence is $O(k + k \log n/k)$. The actual algorithm is never stated. The algorithm relies on knowledge of k. However, they reasonably assert that this constraint may be remove using a standard successive doubling trick.

Definition: A list of queries is (α, k, n)-*universal* if it isolates all but at most $(1-\alpha)k$ of the members of I for every I with $|I| \leq k$.

The overall goal can be restated as constructing a $(1, k, n)$-universal list of queries, of length $O(k + k \log n/k)$.

Suppose we choose a sequence of query sets $Q_1, Q_2, \ldots, Q_p$ independently, with each set chosen uniformly at random from the set of all subsets of the n stations of size exactly n/k. Moreover, the number of sets $p = k/2$.

Lemma 1: For any I, and any $Q_1, \ldots, Q_{j-1}$, the probability that Q_j isolates some member of the input set I is $> 1/(2e^2)$.

Proof: The case where $k = n$ is an easy boundary case. ($|Q_j| = 1$, and at least half of the entire set of n processes are still trying to transmit. So the chance is at least $1/2$ of choosing one of these.) That is we can say that: $2 \leq k \leq n/2$.

The first expression gives the probability of isolating some member of the still-active set I_{j-1}. Let x denotes $|I_{j-1}|$. The denominator, $nchoose(n/k)$, is simply the total number of ways to choose Q_j. The numerator represents the number of choices of Q_j that successfully isolate some member of I. In the numerator, the first factor x gives the number of choices for the station to isolate. Suppose that is in Q_j The second factor, $((n-x)choose(n/k-1))$, is the number of ways to choose the remaining $n/k - 1$ elements of Q_j from among the non-transmitting stations. Those can be stations that already succeeded in transmitting, or stations that weren't in the input set I to begin with.

Expanding the binomial coefficients into factorials, using the fact that each fraction in the resulting expansion is $\leq 1/2$ and the fact that for $u \leq 1/2$ we get $1 - u > e^{-2u}$. The result is a lower bound that is a harmonic series. Use an integral approximation for the resulting harmonic series, and finally, use the approximation expression we have seen before, for $(1 - 1/n)^{n-1}$.

Regardless of what has happened before, the probability of the next query set isolating some station is bounded from below. Now use this repeatedly, for the whole series of $p = k/2$ queries.

Lemma 2: For some constants b and c, the probability that $Q_1, Q_2, \ldots, Q_p$ isolates at least ck members of the input set I is $> 1 - 1/e^{bk}$.

This leads to Theorem 1.

Theorem 1: For every k, n, $2 \leq k \leq n$ (with the two other assumptions about k and n), there exists a (c, k, n)-universal list of query sets of length $O(k + k \log n/k)$, where c is the constant in Lemma 2.

Proof: The idea is to use enough queries so that, for every I, the probability that I is a "bad input set" (does not get enough stations isolated) is strictly less than $1/(nchoosek)$. This probability is taken over the random choices of the query sets, as usual. Then the expected number of bad input sets is the number of input sets I, which is $(nchoosek)$, times this probability. This implies that the expected number of bad input sets is < 1.

If the expected number of bad input sets is less than one (expectation taken over the query lists), then there must be some particular query list (but we don't know what it is, only that it exists) that guarantees that there are no bad input sets.

To get the probability that each I is a bad input set to be sufficiently low ($< 1/(nchoosek)$, we use Lemma 2 repeatedly, on m batches of queries, each of length $p = k/2$.

This upper-bounds the probability that every batch fails to isolate enough stations, by e^{-bmk}. By choosing m as indicated in the proof, approximately $log(nchoosek)/bk$, we get this probability to be $< 1/(nchoosek)$, as needed.

And the length of the list isn't too great: $t = mp$ gives $log(nchoosek)/bk$ times $k/2$, which is $O(k + klog(n/k))$. (Use Stirling's formula to show $(nchoosek) \leq (en/k)^k$. Then take logs, getting $\log nchoosek \leq k(\log e + \log(n/k))$, which is equal to $k + k\log(n/k)$.)

Finally, we concatenate (c, k_i, n)-universal lists, in a geometric progression of lengths, to get a $(1, k, n)$-universal list.

3 IEEE 802.11

3.1 Overview

IEEE standard for wireless MAC layer (and accompanying physical layer).

802.11 has two pieces: Centralized PCF (Point Coordination Function) and Distributed DCF (Distributed Coordination Function)

PCF offers both asynchronous and time-bounded service. Time-bounded services is good for supporting voice/video. Needs an AP to coordinate. It does so using a centralized polling protocol, polling all the devices it is servicing, round-robin, to see if they have something to send.

DCF offers just asynchronous service (no time bound guarantees). Works for both cellular and ad hoc settings. DCF, in turn, has two pieces:

A version of CSMA/CA, and an RTS/CTS reservation exchange. RTS/CTS is often turned off, because of the overhead.

3.2 Basic facilities

- Time synchronization

  If there's a base station (AP), it transmits (periodically) Beacon Frames, which contain the value of the AP's clock at the moment of the actual transmission.
  Receiving stations reset their clocks accordingly.

  Beacon Frames are supposed to be sent at particular scheduled periodic "target beacon transmission times".
  However, that doesn't always succeed, because of congestion-induced delays.
  The AP does the best it can, trying to approximate the scheduled times.

  What if there is no base station (ad hoc mode)?
  Then all the devices will try to send beacons, at the target times.
  The characteristics of the broadcast medium will ensure that (most of the time), one of the broadcasts will win over the others.
  Then everyone else sets their clocks according to the received beacon.

  Sometimes, in ad hoc mode, no beacon will succeed in getting through at some target time.
  Then the beacon for that time is lost.
  But things should recover next time—a pretty robust algorithm.

- Physical carrier sensing

  802.11 does physical carrier-sensing (compares level of energy on the channel with usual "noise" level).
  (Not the same as "busy-tones"—just listening.)
  All the supported physical layer technologies have some type of carrier-sensing capability.

- Link-layer Acks

  Provides link-layer Acks for data transmissions.
  Uses absence of Acks (after a few attempts to transmit data) as a way of detecting collisions.

3.3 802.11 RTS/CTS Exchange

A sequence of RTS/CTS/Data/Ack is used. Virtual carrier sensing is used. That is, all stations receiving either RTS or CTS will set their Virtual Carrier Sense indicator (NAV—Network Allocation Vector, which remembers the duration of time it must remain silent) for the entire duration of the proposed data transmission + Ack.

This reduces number of collisions due to hidden terminals and reduces window of bad simultaneous attempts to just the short interval of the RTS/CTS handshake. The cost is just the time for these transmissions, not the time for the whole data message.

Overhead: Consumes channel resources. May be worth it if data is very large, compared to the RTS itself. Otherwise, it is often turned off, because of the overhead.

4 MACA and MACAW

Bharghavan: MACAW: A Media Access Protocol for Wireless LANs, 1994

4.1 Overview

Paper describes two Media Access protocols for single-channel wireless LANs. MACAW is an improvement on top of the MACA protocol.

MACA: RTS/CTS/Data packet exchange, binary exponential backoff

MACAW: An improvement. RTS/CTS/Data/Ack message exchange, different backoff strategy.

They compare using simulations, and also using tests on their own (Xerox PARC) wireless network.

4.2 Xerox PARC's Radio Network

This system consisted of base stations in the ceiling and portable devices (called "Pads"). All wireless communication between base stations and pads. The system uses a "near-field" radio technology that eliminates multipath effects, so it's suitable for indoor wireless LAN. Everyone uses same signal strength with a range 3-4 meters.

In this sytem both collision and interference is observed. A collision is defined as when a receiver in reception range of two transmitters, can't receive signal cleanly from either. An interference is defined as a receiver in range of one transmitter, slightly out of range of another. But can't receive the closer one's signal because of interference from the further one's signal.

All of this leads to techniques involving the RTS/CTS exchange and the backoff timer presented in the MACA/MACAW paper. However, it should generally be understood that there are better ways to control the physical medium.