CS 336 Lecture 12– 3/6/98

Announcements

- Exercise 4 is due today.
- For Monday and/or Wednesday, read §6.3 of Shay and the paper by Bux which I have added to the binders of reserve readings in the library. I have also added a paper by Stallings which you should read for Wednesday or Friday.

Aloha

1. Ethernet is just one example of the general class of medium access control protocols called contention protocols.
2. To appreciate the design decisions that distinguish Ethernet, it is useful to consider alternate contention protocols.
3. As a form of homage, it is traditional to start discussions of contention protocols with a consideration of the first such protocol implemented, the Aloha protocol. Consideration of Aloha is useful to illustrate the impact of the choices between slotted and unslotted and between using and not using carrier sensing.
4. The Aloha protocol was designed as part of a project at the University of Hawaii. It provided data transmission between computers on several of the Hawaiian Islands using radio transmissions.
   - Communications was typically between remote stations and a central sited named Menehune or vice versa.
   - All message to the Menehune were sent using the same frequency.
   - When it received a message intact, the Menehune would broadcast an ack on a distinct outgoing frequency.
   - The outgoing frequency was also used for messages from the central site to remote computers.
   - All stations listened for message on this second frequency.
5. Unlike Ethernet, Aloha stations did not check to make sure no one else was sending before they began transmitting.
   - I'm not sure exactly why Aloha was designed this way, but a plausible explanation would be the possibility that two stations on opposite sides of the Menehune might be out of range of each other's transmissions while both being within range of the central Menehune.
6. An Aloha station would start transmitting as soon as a new message was presented for delivery and complete the transmission without listening for other, interfering transmissions. Then, it would listen for an ack. If no ack arrived it would wait a random time and try again.
7. While I will not try to present a detailed analysis of the efficiency of any of the variations of the Aloha protocol some consideration of what such an analysis would look like can help you appreciate the motivation for the variations.
8. Even imprecise and incomplete analyses require simplifying assumptions. For this discussion, we will assume that all packets are of equal length, L, an require TR_F time for transmission.
9. Both the efficiency of the overall system and the delay experienced by any single station will depend upon the expected number of transmission of a given packet required before it is transmitted without colliding with another station.

   In particular, the formula for the efficiency of Aloha should look like

   $$\frac{TR_F}{TR_F + CT}$$

   Where CT is the time stations spend in contention before one station wins and gets its packet delivered.

   Given our experience deriving formulas for expected values, it is clear that the correct for CT will involve the probability, p, that a given attempt to transmit a packet fails due to a collision. The bigger this probability, the bigger CT will be.
10. All other things being equal, the probability of a packet suffering a collision will increase with the time it takes to transmit the packet.
    - Collisions in Aloha need not be "head on". If any part of one packet overlaps with any part of another packet enough will be garbled to invalidate the CRC (or other checksum).
    - So, any station that starts a transmission within TR_F of the beginning of the transmission of a given packet will collide with the "given packet". The bigger TR_F ("all other things being equal"), the bigger the chance of a collision.
11. Things are actually even a bit worse than you might think. A given transmission will obviously collide with any transmission started within TR_F after its start time. It will also collide with any transmission started within TR_F before its start time.

Basically, any two stations that get the urge to transmit within $2TR_F$ of one another will collide.

12. Slotted Aloha is a variation of "Pure" Aloha intended to improve performance by reducing the length of the window in time during which stations that start transmissions will collide.

    - The idea is to divide time into slots each of which are just a bit longer than the packet transmission time and then, only allow stations to begin transmission of a packet at the beginning of a slot.
    - Using this scheme, two transmissions can collide only if they begin at the start of the same slot.
    - A station will start a transmission at the beginning of a slot only if it gets the urge to transmit sometime during the preceding slot.

      Accordingly, stations in slotted Aloha collide only if the become ready to transmit within TR_F of one another. Stations in pure Aloha collide if they become ready to transmit within twice this time period.
    - Not surprisingly, a complete analysis predicts that slotted Aloha's maximum efficiency (of about 36%) will be twice that of pure Aloha.

13. The key idea behind the difference between pure and slotted Aloha is that shrinking the "window of vulnerability" during which a packet may suffer a collision around the beginning of its transmission is a way to improve efficiency.

14. Generally, the window of vulnerability can be cut even further by making stations listen before they send a packet.

    - Like "taking turns", the simple idea of listening gets a very fancy name in the world of networking. It is called CSMA for *carrier sense multiple access.*
    - The idea is very simple. If a station listens before sending and refrains from sending if it hears another transmission in progress, then two stations can only collide with one another if they both become ready to transmit within the time it takes for a message to propagate through the network.
    - Note that the degree to which this helps depends on the degree to which the time to propagate through the network is less than the time it takes to send a packet.

15. CSMA (as narrowly proposed/described above), decreases the probability of collisions, but it does not reduce the cost of those collisions that do occur.

    - In Aloha, when two stations collide they both send their full packets before they have any chance to notice that they have been wasting time.
    - Using CSMA only implies making stations listen before they start sending. Although it would seem to make sense to keep listening so that we might notice collisions faster, simple CSMA does not do this.

16. As soon as the nature of a network makes listening possible, it is natural to go beyond using CSMA and also incorporate early *collision detection* by having stations continue to listen while they send and abort as soon as a collision is detected.

    - the resulting scheme is called CSMA/CD where the CD stands for collision detection.

17. Note that although adding the initials CD to a protocol certainly suggests we have added collision detection to the protocol, this is not really the case. Even pure Aloha incorporates collision detection. It just takes a long time (until a station times out waiting for an ack). Thus, what CD really stands for is **early** collisions detection.