CS 336 Lecture 15– 3/13/98

Announcements

- Exercise 5 is now available.

Token Ring Reliability

1. Reliability is a much bigger issue for token rings than it is for Ethernet.
   - A token ring has many active components that could act a "single points of failure".
   - There are several ways in which "transient" problems could lead to permanent problems.
2. The basis of the token ring's mechanisms from recovering from "transient" errors is the assignment of the role of "monitor" to one of the ring's stations.
   - For example, one of the monitor's jobs is to timeout if it does not see a token or packet for longer than the maximum transmission time plus the ring's propagation time.
   - This job can't safely be shared by all stations since if two stations timed-out nearly together a ring might end up having several free tokens.
3. To ensure reliability even if the monitor fails, the assignment of the monitor task to a station is not fixed. When necessary (for example when the machines on the network are first turned on or when the previous monitor crashes or is turned off), the stations of the ring "hold an election" to choose a monitor.
   - When a station realized that a new monitor needs to be elected (we will discuss how it does this later), it sends out special "election" packets (the text calls it a claim token) containing its own hardware address.
   - While sending these packets, it listens for incoming packets.
     - If it receives a packet from another station, it compares its own address to the address in the packet.
       * If its address is smaller, it stops sending packets of its own and instead forwards what it receives.
       * If its address is larger, it just gobbles up the incoming packet and keeps sending its own.
     - If it receives one of its own election packets back, it claims victory and starts acting like a monitor:
       * First it sends out a message to make sure that all the other stations know that the election is over.
       * When this packet returns the new monitor sends out a free token.
4. Of course, given that the monitor is very important to the rings ability to handle failures, we need a way to handle the failure of the monitor itself.

   This can be handled by having station's timeout if they don't see a free token for a very long time.

   Alternately:
   - Each packet contains a bit called the monitor bit.
   - When a station sends a message it stores zero in the monitor bit.
   - When a message passes through the monitor, the monitor sents the packet's monitor bit to 1.
   - If a station ever see one of its own packet's return without having the monitor bit sent, it suspect trouble and starts an election for a new monitor.
5. The monitor bit also solves another problem. If the monitor ever sees a packet with the monitor bit set, it will realize that the packet has been around the ring at least twice. Whichever station was suppposed to remove the packet from the ring has failed! In this case, the monitor removes the packet from the ring and sends a free token. What order should the monitor bit and the free token bit fall to make this possible? Why aren't they in this order according to the readings?

High Speed LANs

1. There are two basic measures of the performance of a medium access control protocol:

   Delay measures the difference between the amount of time required to deliver a packet using a given access control protocol and the time required if the sending station had exclusive access to the network.

   Efficiency (or utilization) measures the amount of data delivered per unit time compared to the amount that could be delivered in the same time if medium access control was not an issue.
2. I have already pointed out that the ratio of the time it takes a packet to propagate through a network to the time it takes to transmit a packet has a significant effect on the efficiency of most protocols.
   - For unclear reasons, this ratio is frequently called "a".
   - The value of a can be thought of as the number of average packets required to just fill the network with bits. In fact, it should be noted that "a" becomes large when either the transmission rate increases, the packet size decreases or the network size increases.

- At high load, the main factor producing inefficiency in a token ring is the time the token spends moving from one station to another. If this time is measured in bits rather than microseconds, then the larger the value of "a", the more bit transmission times get wasted as the token moves from one station to another.
  - In the paper by Stallings I included in the readings, he gives a bound on the efficiency of token ring.
  - Basically, assuming that a ring contains N stations, that the packet size used in the calculation of "a" is the size of the maximal number of bits a station on the ring can send before passing the token, that all stations have data to send and that the time to send such a packet is one unit, then the time spent sending a packet is 1 plus the time it takes the token to move between two stations on average, a/N. So, the efficiency is at most

$$\frac{1}{1+a/N}$$

  - At low load, the average delay introduced waiting for a token to arrive is half the propagation time of the ring regardless of the transmission rate. As the transmission rate increases, however, the delay increases as a fraction of the total time required to transmit and deliver a packet.
- As "a" gets larger, the efficiency of contention protocols such as CSMA/CD also suffers.
  - Basically, if we again assume our units of time are such that the average packet transmission time is 1, then the time wasted by each collision is proportional to the x time required for a to propagate throughtout the network, a. If "C" is the expected number of collisions, then the time to successfully send a single packet will be $1 + kaC$ for some constant k. Therefore, the efficiency of the protocol will behaves as

$$\frac{1}{1+kac}$$

  - One would expect the "C" in this formula to increase as the number of stations increased. Thus, unlike token ring, CSMA/CD gets worse as either "a" or the number of stations increases.

3. The problems associated with large "a" values are typically caused by increases in transmission speed. Therefore, LANs in which "a" is large are considered *High Speed LANs* even though a large "a" value may actually be the result of an increase in network size.
4. A great deal of effort has been devoted to the design of protocols that operate efficiently on high speed LANs (a.k.a. HSLNs).
   - A fair amount of work has sought to produce bus networks that could compete with ring networks.
     - The reliability advantages of a passive bus helped motivate this work.
   - Most of this work has assumed that the physical layer of the networks involved would be based on optical fiber. As a result, many of the designs are based on buses with uni-directional data flow.
5. We will first discuss collision free protocols for uni-directional bus networks. Then we will look at two alternate approaches to the design of ring protocols: slotted rings and insertion rings. Finally, we will look at a slotted bus protocols, DBDQ.