Medium Access Channel Sublayer

"Data Link Layer for LANs"

Can divide networks into point-to-point and broadcast. Look at broadcast networks and their protocols.

When many stations compete for a channel (e.g., broadcast channel such as an Ethernet), an algorithm must arbitrate access to the shared channel.

Need a way of insuring that when two or more stations wish to transmit, they all wait until doing so won't interfere with other transmitters. Broadcast links include LANs, satellites (WAN), etc.

LAN characteristics:

- diameter of not more than a few kilometers
- data rate starting at least several Mbps and going up to Gbps
- complete ownership by a single organization

MANs cover a city-wide area with LAN technology. For example, cable TV.

Can have higher speed, lower error rate lines with LANs than WANs.

Performance Analysis

We want to analyze performance of the communication channel to determine how well different technologies work.

There are three general approaches to analyzing systems:

1. Actually build the system and measure its performance. Drawback? Unfortunately, this can be expensive in time and effort.
2. Write a computer program that simulates the system. That is, processes can represent customers or servers, and processes move from queue to queue. For example, a "customer" process could generate arriving packets by depositing a message in a queue. The server process would then "consume" these messages, removing them from the queue and then simulating the steps required to actually send the packet.
3. Represent service times, arrival rates, etc. as a series of equations and solve the system mathematically. Drawback? Very hard to represent a system using mathematical equations without making simplifying assumptions that may invalidate the model (and thus any conclusions).

In order to be able to solve a system mathematically, one must typically make simplifying assumptions. Some common simplifications:

1. Queues can be infinitely long. Of course, no queue is infinitely long in practice, but if a fixed length queue never overflows, we can pretend that it has infinite length without compromising our model.
2. Choose a service time function that is "well behaved". Typically, the service time is chosen to vary around a mean, with small deviations from the mean. Exponential distributions are commonly used, because their mathematical representations make analytic problems solvable. Note: in real networks, traffic patterns rarely behave at all like these functions.
3. Likewise, assume customers arrive at a server's queue according to some well-behaved probability function.

Model Validity

The above simplifications may be made for computer simulations as well. Thus, when using simulations or mathematics to analyze a given system, one must always ask whether the assumptions are relevant to the system being considered. For instance:

- Real systems do not have an infinite amount of buffer space. If the system has "a lot", the assumption is probably reasonable. That is, just because the model assumes an infinite amount of buffer space doesn't mean that all that buffer space is actually used. If only a small amount of space is ever used, then the assumption is fine.
- Hot point of contention between analysts and experimentalists: Do service times and arrival times correspond to "well behaved" functions? In most cases, the answer is no. In most real systems, actual traffic patterns are quite complex and do not match simplistic functions. For example, flow control eventually forces a sender to stop generating packets. This contradicts with the "randomness" assumption of arrival functions.

Queueing Theory

Although assumptions are required, it is a powerful mathematics tool for making quantitative analyses of computer networks. Need to be aware of as we will use it in analysis. Appendix of Tanenbaum.

High level overview: Model the system as a set of customers (frames needing to be delivered) and servers (part of network that processes frames). See picture.

Queueing Systems

Characterized by five components

1. interarrival-time probability density function
2. service-time probability density function
3. number of servers
4. queueing discipline
5. amount of buffer space in the queues

Assume an infinite number of customers (i.e. what is happening in the system does not affect new arrivals)

Concentrate on infinite-buffer, single-server systems using first-come, first-served. Use the notation A/B/m where:

- A is the interarrival-time probability density
- B is the service-time probability density
- m is the number of servers

Density functions are chosen from:

- M – exponential probability density (M stands for Markov)
- D – all customers have the same value (D stands for deterministic)
- G – general (arbitrary)

Usually assume a M/M/1 model.

Handwritten Notes

Dynamic Channel Allocation Assumptions

Before looking at specific technologies, it is useful to look at general characteristics of the various technologies.

In the *station model*, stations generate one frame at a time and block until the frame has been successfully transmitted. Thus, a station cannot have multiple frames queued for transmission.

With the *single channel assumption*, all stations share one communications medium (channel), and all stations are equivalent. That is, if two or more stations wish to transmit, no station has an inherent transmission priority over the other.

When two or more stations transmit frames simultaneously, a *collision* takes place. The *collision assumption* model assumes that all stations detect a collision and the collided frames become garbled and must be retransmitted.

In the *continuous time* model, a station can transmit frames at any time.

In the *slotted time* model, a station may begin transmission of a frame only at the start of a time slot. A clock ticks at regular pulses, and a station may begin transmission only at the start of a pulse. Thus, if a station wishes to transmit, it must wait until the start of a time slot before actually attempting transmission.

In *carrier sense* systems, a station can sense if a channel is being used before it attempts transmission. That is, it can listen to the channel to determine if another station is currently transmitting.

With *no carrier sense*, stations cannot sense the channel before transmission starts.

ALOHA Protocols

Back in 1970, the University of Hawaii built a network out of radios that broadcast signals. Basic idea:

1. Anyone may transmit whenever they want. (Continuous time model.)
2. Each radio detects collisions by listening to its own signal. A collision is detected when a sender doesn't receive the signal that it just sent.
3. After a collision, wait a random amount of time and transmit the same frame again. This technique is known as *backoff*.

With the queueing theory, we can now examine its efficiency:

1. Assume that all frames are the same size. That is, the service time is deterministic (when no collision takes place).
2. When a station sends, it blocks until transmission is successful. A station cannot try to send more than one packet at a time. Note that this means that if a packet's transmission is delayed due to collisions, the station will also delay generating additional packets to send.
3. Assume new frames are generated according to a Poisson distribution, with a mean of S frames per frame time.

   Note:

   (a) For $S > 1$, load exceeds channel's capacity and many collisions occur.

   (b) for S close to 0, few collisions occur.
4. Assume that the probability of n transmission attempts per frame time is Poisson distributed with a mean of G per frame. G is called the *offered load* and includes both new transmissions and retransmissions.

   Note: $G = \lambda t$, where t is one frame time, and $G \geq S$.

ALOHA Analysis

The probability that n frames will be generated during a given frame time is given by the Poisson distribution:

$P_n(t) = \frac{(\lambda t)^n e^{-\lambda t}}{n!} = \frac{G^n e^{-G}}{n!}$

When will transmission be successful? Answer: If there are no collisions. There are two cases. First, if another station starts transmission after we do (within time G), or if we collide with another frame already in transmission (within time G before we start transmitting).

Thus, the probability if no other traffic being generated the entire vulnerable period is

$P_0(2G) = e^{-2G}$ and

Throughput $= GP_0(2G) = Ge^{-2G}$

When is throughput maximized?

Answer: when $G = 0.5$; the throughput at this value is only 18%!

The above method is called *pure ALOHA*. An alternative method called *slotted ALOHA* behaves as follows:

1. Only allow transmission to start at the beginning of a slot time. Need a means of synchronization. Each slot has duration equal to the maximum frame size. When a station has a new frame to send, it must wait until the start of the next slot time before attempting transmission.

2. Because transmissions always start and stop on slot boundaries, the collision interval drops from two to one frame time. Result:

   Throughput $= Ge^{-G}$, which doubles the max throughput to 37% at $G = 1$

Intuitively, the throughput is poor because of the following:

- At low offered loads, there are few collisions, but many empty slots. Thus bandwidth is frequently unused and wasted.
- At higher loads, few empty slots, but many collisions, effectively wasting the full slots.
- Throughput is maximized somewhere in the middle.

CSMA Protocols

Under what conditions do collisions take place?

1. If a station begins transmission when the channel is already busy. Obviously, if one transmits when the channel is already in use, a collision will happen. Note: Pure Aloha has this problem, which is one reason its performance is poor. We can avoid these types of collisions by sensing the channel before we attempt transmission, and send only if it is idle.
2. If two stations begin transmission at (roughly) the same time, each thinking that the channel is idle. Because of propagation delay times, one station may think the channel is idle, when in fact another station has already begun transmission, but its signals have not reached the first node.

Collision avoidance protocols all attempt to reduce one or more of the above factors.

Carrier Sense Protocols

Problem with ALOHA is that frames are blindly sent—bound to be collisions.

Stations listen for a transmission before trying to send data—carrier sense. Only send if channel is idle.

- 1-Persistent CSMA (Carrier Sense Multiple Access). Sense channel, if idle then send, if busy wait until idle and then send. 1-persistent because it sends with probability of one when senses channel is idle. Collisions?

  Effect of propagation delay (takes time for signal to propagate from one channel to another). Even if zero could have collisions.

- nonpersistent CSMA—less greedy. Sense channel. If idle then send. If busy then wait random amount of time before repeating the same routine. Collisions go away? No, can still send at the same time.

- p-persistent CSMA—applies to slotted channels. If channel idle then send with probability *p*, with probability *q=(1-p)*, it defers to the next time slot. Delay for random time if channel busy.

  Note: p = 1 implies we transmit immediately, p = .1, we transmit with probability .1.

  The primary advantage of p-persistent protocols is that they reduce the number of collisions under heavy load. The primary disadvantage is that they increase the average delay before a station transmits a frame. Under low loads, the increased delay reduces efficiency.

Protocol Efficiency

The following table gives maximum efficiency percentages for some of the protocols we have studied so far:

Protocol	Channel Utilization
pure ALOHA	18%
slotted ALOHA	37%
1-persistent	55%
.5-persistent	70%
0.01-persistent	99%

CSMA/CD

Another way to reduce the number of collisions is to abort collisions as soon as they are detected. CSMA networks with Collision Detect (CSMA/CD) do just that. How long does it take to detect collisions:

- at least twice the propagation delay, or 2τ (in worst case, τ for the signal to travel from one end of cable to the other, another τ for the collision indication to travel back)
- we'll call this interval the *contention period*
- what does this say about building broadcast networks that span large distances? (increase propagation delay) The longer the physical distance, the longer the contention period. For satellite communication, the contention period is .5 seconds, much longer than the transmission time for a single frame!
- small frames? (could send entire frame before detecting a collision)–pad out the frames in the standard

IEEE 802 Protocol Standards for LANs

The IEEE has produced a set of LAN protocols known as the IEEE 802 protocols. These protocols have been adopted by ANSI and ISO:

- 802.2: logical link standard (device driver interface)
- 802.3: CSMA/CD
- 802.4: token bus
- 802.5: token ring
- 802.11: wireless LAN

Ethernet is a specific product implementing (or nearly so) the IEEE standard. Interesting to note that having an Ethernet port on a machine has become a standard (certainly for workstations).

The 802.3 protocol is described as follows:

- 1-persistent CSMA/CD LAN (always sends if sense as idle)
- its history is as follows:
  1. started with ALOHA
  2. continued at Xerox, where Metcalf & Boggs produced a 3 Mbps LAN version
  3. Xerox, DEC, and Intel standardized a 10Mbps version
  4. IEEE standardized a 10Mbps version (with slight differences from the Xerox standard
- the maximum length of a segment of cable is 500 meters
- segments can be separated by *repeaters*, devices that regenerate or "amplify" signals (not frames); a single repeater can join multiple segments
- maximum distance between two stations 2.5 km, maximum number of repeaters along any path: 4 (why these limits?)
- uses Manchester encoding

802.3 Frame Layout

At the Medium Access (MAC) sublayer, frames consist of the following:

1. Frames begin with 56-bit preamble consisting of alternating 1s and 0s. Purpose? Similar to start bit in RS-232-A. It allows the receiver to detect the start of a frame.
2. Start of the frame designated by the byte "10101011". That is, two consecutive 1 bits flag the end of the preamble.
3. 48-bit destination address (16-bit for lower speed version).
4. 48-bit source address (16-bit for lower speed version).
5. 16-bit data length field; maximum data size 1500 bytes.
6. 32-bit checksum field. The checksum is a number dependent on every bit in the frame. The sender computes a checksum and appends it to the frame. The receiver also recalculates the checksum, comparing result with value stored in frame. If the two checksums differ, some of the bits in the frame must have changed and the packet is discarded as having errors.

There are two types of addresses:

1. *Unicast* addresses start with a high-order bit of 0. Unicast addresses refer to a single machine, and every Ethernet address in the world is guaranteed to be unique (e.g., address is sort of like a serial number).
2. *Multicast* (group) addresses start with a high-order bit of 1. Multicast addresses refer to a set of one or more stations.

   A *broadcast* address (all 1's) is a special case of multicasting. All machines process broadcast frames.

   The management of multicast addresses (e.g., entering or leaving a group) must be managed by some outside mechanism (e.g, higher-layer software).

Binary Exponential Backoff

802.3 (Ethernet) LANs use a *binary exponential backoff* algorithm to reduce collisions:

1. It is 1-persistent. When a station has a frame to send, and the channel is idle, transmission proceeds immediately.
2. When a collision occurs, the sender generates a noise burst to insure that all stations recognize the condition and aborts transmission.
3. After the collision, wait 0 or 1 contention periods before attempting transmission again. Station has 50-50 probability of waiting 0 or 1 contention period. The idea is that if two stations collide, one will transmit in the first interval, while the other one will wait until the second interval. Once the second interval begins, however, the station will sense that the channel is already busy and will not transmit.
4. If another collision occurs: Randomly wait 0, 1, 2, or 3 slot times before attempting transmission again.
5. In general, wait between 0 and $(2^r - 1)$ times, where r is the number of collisions encountered.
6. Finally, freeze interval at 1023 slot times after 10 attempts, and give up altogether after 16 attempts.

Result? Low delay at light loads, yet we avoid collisions under heavy loads.

Also, note that there are no acknowledgments; a sender has no way of knowing that a frame was successfully delivered.

802.3 Analysis

Assuming a constant retransmission probability (rather than backoff), the channel efficiency of 802.3 LANs is given by:

Efficiency $= \frac{P}{P+\frac{2\tau}{A}}$,

where P is the mean time to transmit a frame, and A is the probability that some station transmits during each slot.

What happens as τ increases? The efficiency drops.

An alternative from of the same formula:

Efficiency $= \frac{1}{1+\frac{2BLe}{cF}}$,

where: B = bandwidth, L = cable length, c is signal propagation rate, e = contention slots per frame, and F = mean frame length.

What does this say about networks with large BL products? They are inefficient. Fiber optics, satellites, and WANs in general have high bandwidth-delay products.

Switched 802.3 LANs

Can connect hosts to a hub switch. Advantage is that stations use the same network interface card, but they can run at higher speeds.

Fast Ethernet

See Fig. 4-21. Different approaches result in 100Mbps.

100BASE-T4 split data on 3 streams each with effective data rate of 33 1/3Mbps. Use four twisted pairs. One pair used for in-bound, one for out-bound and two are bidirectional. Use ternary encoding (8B6T).

Shared hub (all incoming lines logically connected) vs. Switched Hub (each incoming frame is buffered on a plug-in line card). 100Base-Fx cables must be buffered because they are too long for Ethernet collision detection algorithm.

Gigabit Ethernet

Use same Ethernet frame, but go 10 times faster again—1 Gbps.

All connections are point-to-point as shown in Fig 4-22.

Normal mode is to use full-duplex (two-way communication) with a central switch to buffer all communication.

In half-duplex mode, machines are connected to a hub, which does no buffering—collisions are possible. Does reduce radius to ensure collisions will be detected. Either pad frames to be longer or use *frame bursting* to combine multiple frames.

Cabling options shown in Fig 4-23.

Also added a flow control mechanism so that one end can cause the other end to PAUSE for units of minimum frame time.

IEEE Standard 802.2: Logical Link Control

On top of MAC layer for DLL equivalent (Fig 4-24).

Service options:

- Unreliable datagram service
- acknowledged datagram service
- reliable connection-oriented service

IEEE Standard 802.11: Wireless LANs

802.11 has three permitted transmission techniques.

Two problems (see Fig 4-26):

- hidden station problem
- exposed station problem

Consequently 802.11 does not use CSMA/CD as does Ethernet.

Distributed Coordination Function (DCF)

Uses CSMA/CA—CSMA with Collision Avoidance with two methods:

1. sender senses channel and if idle sends entire frame—does not sense while transmitting. If channel is busy then use binary exponential backoff to determine when to transmit if a collision.
2. use MACAW (Multiple Access with Collision Avoidance for Wireless). See Fig 4-27. In example, stations C and D hear RTS and CTS frames and know to keep quiet. Indicated by Network Allocation Vector (NAV).

   Frames can also be fragmented as shown in Fig 4-28. Smaller fragments are less suscepible to transmission error, which is more common in wireless transmission. Short interframe spacing allows additional fragments to be sent.

Point Coordination Function (PCF)

Alternate approach is to use an infrastructure of base stations (access points (APs)). Wireless nodes select an AP via *scanning*:

1. Node sends a `Probe` frame.
2. All APs within range reply with a `Probe Response` frame.
3. Node selects one of the APs and sends that AP an `Association Request` frame.
4. The AP replies with an `Association Response` frame.

Base station polls clients on a periodic basis to see if they have any frames to send.

Can allow stations to go to sleep to save power.

Can have PCF and DCF coexist.

802.11 Frame Structure

Frame format shown in Fig 4-30.

- Frame control—11 different fields.
- Duration—how long frame and ack will occupy the channel (for NAV mechanism)
- Addresses—4, source/dest addresses of hosts and base stations
- Seq—for fragments
- Data—upto 2312 bytes, but generally frames limited to max Ethernet size of 1500 bytes for compatibility.
- Checksum

Management frames use Subtype field for RTS, CTS and ACK.

802.16 Broadband Wireless

Wireless MAN.

Bluetooth

wireless standard for interconnecting computing and communications devices using short-range, low-power, inexpensive wireless radios.

Basic unit is a *piconet*—piconets can be combined into a *scatternet*. See Fig 4-35. Master-slave communication with master using time division multiplexing (TDM) to control which slave communicates.

13 applications with profiles defined. Shown in Fig 4-36.