SUBJECT

WIRELESS NETWORKS

SESSION 6 Cellular Systems Medium Access Control (MAC)"

SESSION 6

Cellular Systems “Data Link Control Protocols “Cellular Concepts and Basics

: "Medium Access Control (MAC)"

- A summary or tutorial about the essentials or basic concepts of a mobile phone and cellular telecommunications systems and cellular technology.

IN THIS SECTION

- Cellular concepts
- Cellular multiple access schemes
- TDD & FDD duplex
- Cellphone electronics
- Cellular network architecture
- Cellphone registration
- Handover / handoff

Cellular systems are widely used today and cellular technology needs to offer very efficient use of the available frequency spectrum. With billions of mobile phones in use around the globe today, it is necessary to re-use the available frequencies many times over without mutual interference of one cell phone to another. It is this concept of frequency re-use that is at the very heart of cellular technology. However the infrastructure technology needed to support it is not simple, and it required a significant investment to bring the first cellular networks on line.

Early schemes for radio telephones schemes used a single central transmitter to cover a wide area. These radio telephone systems suffered from the limited number of channels that were available. Often the waiting lists for connection were many times greater than the number of people that were actually connected. In view of these limitations this form of radio communications technology did not take off in a big way. Equipment was large and these radio communications systems were not convenient to use or carry around.

The need for a spectrum efficient system

To illustrate the need for efficient spectrum usage for a radio communications system, take the example where each user is allocated a channel. While more effective systems are now in use, the example will take the case of an analogue system. Each channel needs to have a bandwidth of around 25 kHz to enable sufficient audio quality to be carried as well as enabling there to be a guard band between adjacent signals to ensure there are no undue levels of interference. Using this concept it is only possible to accommodate 40 users in a frequency band 1 MHz wide. Even of 100 MHz were allocated to the system this would only enable 4000 users to have access to the system. Today cellular systems have millions of subscribers and therefore a far more efficient method of using the available spectrum is needed.

Cell system for frequency re-use

The method that is employed is to enable the frequencies to be re-used. Any radio transmitter will only have a certain coverage area. Beyond this the signal level will fall to a limited below which it cannot be used and will not cause significant interference to users associated with a different radio transmitter. This means that it is possible to re-use a channel once outside the range of the radio transmitter. The same is also true in the reverse direction for the receiver, where it will only be able to receive signals over a given range. In this way it is possible to arrange split up an area into several smaller regions, each covered by a different transmitter / receiver station.

These regions are conveniently known as cells, and give rise to the name of a "cellular" technology used today. Diagrammatically these cells are often shown as hexagonal shapes that conveniently fit together. In reality this is not the case. They have irregular boundaries because of the terrain over which they travel. Hills, buildings and other objects all cause the signal to be attenuated and diminish differently in each direction.

It is also very difficult to define the exact edge of a cell. The signal strength gradually reduces and towards the edge of the cell performance will fall. As the mobiles themselves will have different levels of sensitivity, this adds a further greying of the edge of the cell. Therefore it is never possible to have a sharp cut-off between cells. In some areas they may overlap, whereas in others there will be a "hole" in coverage.

Cell clusters

When devising the infrastructure technology of a cellular system, the interference between adjacent channels is reduced by allocating different frequency bands or channels to adjacent cells so that their coverage can overlap slightly without causing interference. In this way cells can be grouped together in what is termed a cluster.

Often these clusters contain seven cells, but other configurations are also possible. Seven is a convenient number, but there are a number of conflicting requirements that need to be balanced when choosing the number of cells in a cluster for a cellular system:

- Limiting interference levels
- Number of channels that can be allocated to each cell site

It is necessary to limit the interference between cells having the same frequency. The topology of the cell configuration has a large impact on this. The larger the number of cells in the cluster, the greater the distance between cells sharing the same frequencies.

In the ideal world it might be good to choose a large number of cells to be in each cluster. Unfortunately there are only a limited number of channels available. This means that the larger the number of cells in a cluster, the smaller the number available to each cell, and this reduces the capacity.

This means that there is a balance that needs to be made between the number of cells in a cluster, and the interference levels, and the capacity that is required.

Cell size

Even though the number of cells in a cluster in a cellular system can help govern the number of users that can be accommodated, by making all the cells smaller it is possible to increase the overall capacity of the cellular system. However a greater number of transmitter receiver or base stations are required if cells are made smaller and this increases the cost to the operator. Accordingly in areas where there are more users, small low power base stations are installed.

The different types of cells are given different names according to their size and function:

- ***Macro cells:*** Macro cells are large cells that are usually used for remote or sparsely populated areas. These may be 10 km or possibly more in diameter.
- ***Micro cells:*** Micro cells are those that are normally found in densely populated areas which may have a diameter of around 1 km.
- ***Pico cells:*** Picocells are generally used for covering very small areas such as particular areas of buildings, or possibly tunnels where coverage from a larger cell in the cellular system is not possible. Obviously for the small cells, the power levels used by the base stations are much lower and the antennas are not position to cover wide areas. In this way the coverage is minimised and the interference to adjacent cells is reduced.
- ***Selective cells:*** Sometimes cells termed selective cells may be used where full 360 degree coverage is not required. They may be used to fill in a hole in the coverage in the cellular system, or to address a problem such as the entrance to a tunnel etc.
- ***Umbrella cells:*** Another type of cells known as an umbrella cell is sometimes used in instances such as those where a heavily used road crosses an area where there are microcells. Under normal circumstances this would result in a large number of handovers as people driving along the road would quickly cross the microcells. An umbrella cell would take in the coverage of the microcells (but use different channels to those allocated to the microcells). However it would enable those people moving along the road to be handled by the umbrella cell and experience fewer handovers than if they had to pass from one microcell to the next.

Infrastructure technology

Although the illustrations used here to describe the basic infrastructure technology used for cellular systems refers to the original first generation systems, it serves to provide an overview of the basic cellular concepts that form the cornerstones of today's cellular technology. New techniques are being used, but the basic concepts employed are still in in use.

DLC (data link control) is the service provided by the Data Link layer of function defined in the Open Systems Interconnection (OSI) model for network communication. The Data Link layer is responsible for providing reliable data transfer across one physical link (or telecommunications path) within the network. Some of its primary functions include defining frames, performing error detection or ECC on those frames, and performing flow control (to prevent a fast sender from overwhelming a slow receiver).

Many point-to-point protocols exist at the Data Link layer including High-level Data Link Control (HDLC), Synchronous Data Link Control (SDLC), Link Access Procedure Balanced (LAPB), and Advanced Data Communications Control Procedure (ADCCP). All of these protocols are very similar in nature and are found in older networks (such as X.25 networks). In the Internet, one of two point-to-point protocols are used at this layer: Serial Line Internet Protocol (SLIP) or Point-to-Point Protocol (PPP) with PPP being the newer, approved standard. All of these protocols are used in point-to-point connections such as those on metropolitan area network (MAN) or wide area network (WAN) backbones or when we dial our Internet service provider (ISP) from home using a modem.

In local area networks (LANs) where connections are multipoint rather than point-to-point and require more line-sharing management, the Data Link layer is divided into two sublayers:

the Logical Link Control layer and theMedia Access Control layer. The Logical Link Control layer protocol performs many of the same functions as the point-to-point data link control protocols described above. The Media Access Control (MAC) layer protocols support methods of sharing the line among a number of computers. Among the most widely used MAC protocols are Ethernet (IEEE 802.3), Token Bus (IEEE 802.4), and token ring (IEEE 802.5) and their derivatives.

Medium Access Control (MAC)

The Medium Access Control (MAC) protocol is used to provide the data link layer of the Ethernet LAN system. The MAC protocol encapsulates a SDU (payload data) by adding a 14 byte header (Protocol Control Information (PCI)) before the data and appending an integrity checksum, The checksum is a 4-byte (32-bit) Cyclic Redundancy Check (CRC) after the data. The entire frame is preceded by a small idle period (the minimum inter-frame gap, 9.6 microsecond (µS)) and a 8 byte preamble (including the start of frame delimiter).

Preamble

The purpose of the idle time before transmission starts is to allow a small time interval for the receiver electronics in each of the nodes to settle after completion of the previous frame. A node starts transmission by sending an 8 byte (64 bit) preamble sequence. This consists of 62 alternating 1's and 0's followed by the pattern 11. Strictly speaking the last byte which finished with the '11' is known as the "Start of Frame Delimiter". When encoded using Manchester encoding, at 10 Mbps, the 62 alternating bits produce a 10 MHz square wave (one complete cycle each bit period).

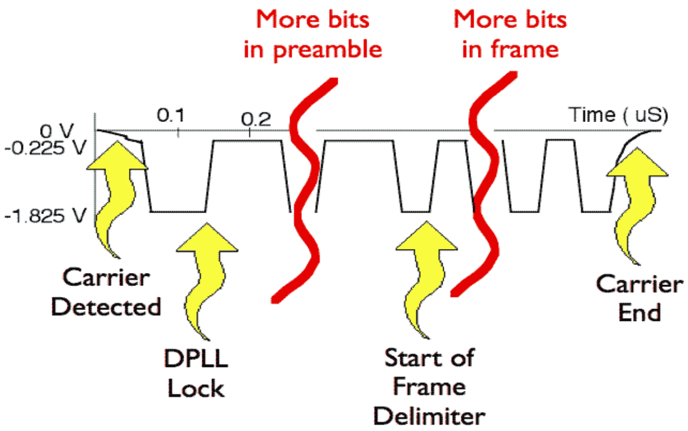

The purpose of the preamble is to allow time for the receiver in each node to achieve lock of the receiver Digital Phase Lock Loop which is used to synchronise the receive data clock to the transmit data clock. At the point when the first bit of the preamble is received, each receiver may be in an arbitrary state (i.e. have an arbitrary phase for its local clock). During the course of the preamble it learns the correct phase, but in so doing it may miss (or gain) a number of bits. A special pattern (11), is therefore used to mark the last two bits of the preamble. When this is received, the Ethernet receive interface starts collecting the bits into bytes for processing by the MAC layer. It also confirms the polarity of the transition representing a '1' bit to the receiver (as a check in case this has been inverted).

Header

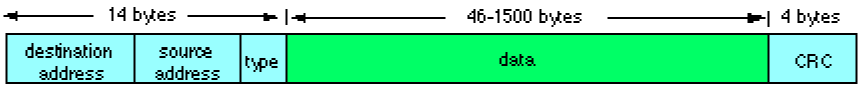

MAC encapsulation of a packet of data

The header consists of three parts:

- A 6-byte destination address, which specifies either a single recipient node (unicast mode), a group of recipient nodes (multicast mode), or the set of all recipient nodes (broadcast mode).
- A 6-byte source address, which is set to the sender's globally unique node address. This may be used by the network layer

protocol to identify the sender, but usually other mechanisms are used (e.g.arp). Its main function is to allow address learning which may be used to configure the filter tables in a bridge.

- A 2-byte type field, which provides a Service Access Point (SAP) to identify the type of protocol being carried (e.g. the values 0x0800 is used to identify the IP network protocol, other values are used to indicate other network layer protocols). In the case of IEEE 802.3 LLC, this may also be used to indicate the length of the data part. Th type field is also be used to indicate when a Tag field is added to a frame.

CRC

The final field in an Ethernet MAC frame is called a Cyclic Redundancy Check (sometimes also known as a Frame Check Sequence). A 32-bit CRC provides error detection in the case where line errors (or transmission collisions in Ethernet) result in corruption of the MAC frame. Any frame with an invalid CRC is discarded by the MAC receiver without further processing. The MAC protocol does not provide any indication that a frame has been discarded due to an invalid CRC.

The link layer CRC therefore protects the frame from corruption while being transmitted over the physical mediuym (cable). A new CRC is added if the packet is forwarded by the router on another Ethernet link. While the packet is being processed by the router the packet data is not protected by the CRC. Router processing errors must be detected by network or transport-layer checksums.

Inter Frame Gap

After transmission of each frame, a transmitter must wait for a period of 9.6 microseconds (at 10 Mbps) to allow the signal to propagate through the receiver electronics at the destination. This period of time is known as the Inter-Frame Gap (IFG). While every transmitter must wait for this time between sending frames, receivers do not necessarily see a "silent" period of 9.6 microseconds. The way in which repeaters operate is such that they may reduce the IFG between the frames which they regenerate.

Byte Order

It is important to realise that nearly all serial communications systems transmit the least significant bit of each byte first at the physical layer. Ethernet supports broadcast, unicast, and multicast addresses. The appearance of a multicast address on the cable (in this case an IP multicast address, with group set to the bit pattern 0xxx xxxx xxxx xxxx xxxx xxxx) is therefore as shown below (bits transmitted from left to right):

```
0                        23 IP Multicast Address Group  47
|                        | <--------------------------->|
1000 0000 0000 0000 0111 1010 xxxx xxx0 xxxx xxxx xxxx xxxx
|                                 |
Multicast Bit                         0 = Internet Multicast
                                  1 = Assigned for other uses
```

However, when the same frame is stored in the memory of a computer, the bits are ordered such that the least significant bit of each byte is stored in the right most position (the bits are transmitted right-to-left within bytes, bytes transmitted left-to-right):

```
0                        23                          47
|                        |                           |
0000 0001 0000 0000 0101 1110 0xxx xxxx xxxx xxxx xxxx xxxx
       |                 <-------------------------->
       Multicast Bit         IP Multicast Address Group
```

CSMA /CD

The Carrier Sense Multiple Access (CSMA) with Collision Detection (CD) protocol is used to control access to the shared Ethernet medium. A switched network (e.g. Fast Ethernet) may use a full duplex mode giving access to the full link speed when used between directly connected NICs, Switch to NIC cables, or Switch to Switch cables.

Receiver Processing Algorithm

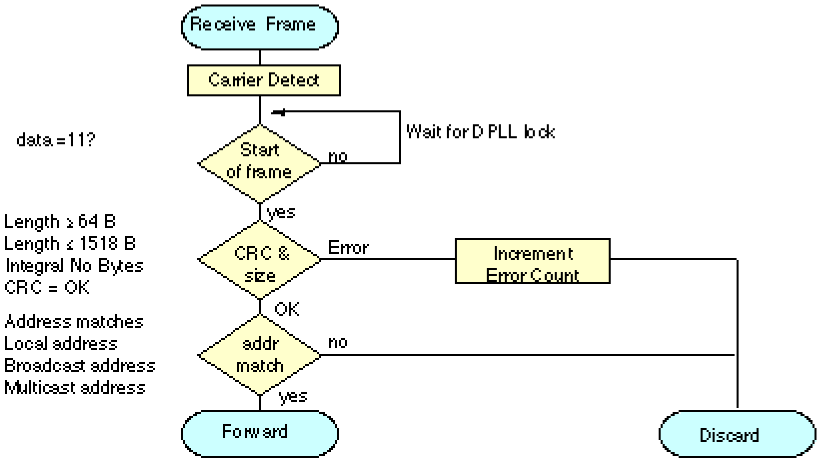

Runt Frame

Any frame which is received and which is less than 64 bytes is illegal, and is called a "runt". In most cases, such frames arise from a collision, and while they indicate an illegal reception, they may be observed on correctly functioning networks. A receiver must discard all runt frames.

Giant Frame

Any frame which is received and which is greater than the maximum frame size, is called a "giant". In theory, the jabber control circuit in the transceiver should prevent any node from generating such a frame, but certain failures in the physical layer may also give rise to over-sized Ethernet frames. Like runts, giants are discarded by an Ethernet receiver.

Jumbo Frame

Some modern Gigabit Ethernet NICs support frames that are larger than the traditional 1500 bytes specified by the IEEE. This new mode requires support by both ends of the link to support Jumbo Frames.Path MTU Discovery is required for a router to utilise this feature, since there is no other way for a router to determine that all systems on the end-to-end path will support these larger sized frames.

A Misaligned Frame

Any frame which does not contain an integral number of received bytes (bytes) is also illegal. A receiver has no way of knowing which bits are legal, and how to compute the CRC-32 of the frame. Such frames are therefore also discarded by the Ethernet receiver.

Other Issues

The Ethernet standard dictates a minimum size of frame, which requires at least 46 bytes of data to be present in every MAC frame. If the network layer wishes to send less than 46 bytes of data the MAC protocol adds sufficient number of zero bytes (0x00, is also known as null padding characters) to satisfy this requirement. The maximum size of data which may be carried in a MAC frame using Ethernet is 1500 bytes (this is known as the MTU in IP).

A protocol known as the "Address Resolution Protocol" (arp) is used to identify the MAC source address of remote computers when IP is used over an Ethernet LAN.

Exception to the Rule

An extension to Ethernet, known as IEEE 802.1p allows for frames to carry a tag. The tag value adds an extra level of PCI to the Ethernet frame header. This increases the size of the total MAC frame when the tag is used. A side effect of this is that NICs and network devices designed to support this extension require a modification to the jabber detection circuit.