Address Resolution Protocol (ARP):

Anytime a host or a router has an IP datagram to send to another host or router, it has the logical (IP) address of the receiver. But the IP datagram must be encapsulated in a frame to be able to pass through the physical network. This means that the sender needs the physical address of the receiver. ARP accepts a logical address from the IP protocol, maps the address to the corresponding physical address and pass it to the data link layer.

❖ ARP Packet Format

- **Hardware type (2B):** Defining the type of the network on which ARP is running. Ethernet is given the type 1.
- **Protocol type (2B):** It Defines the protocol encapsulated in the ARP packet. The value of this field for the IPv4 protocol is $(0800)_{16}$.
- **Hardware length (1B):** Defining the length of the physical address in bytes. For Ethernet the value is 6.
- **Protocol length (1B):** Defining the length of the logical address in bytes. For IPv4 protocol the value is 4.
- **Operation (2B):** Defining the type of packet. Two packet types are defined: ARP request (1), ARP reply (2).
- **Sender hardware address:** This is a variable-length field defining the physical address of the sender. For example, for Ethernet this field is 6 bytes long.
- **Sender protocol address.** This is a variable-length field defining the logical (for example, IP) address of the sender. For the IPv4 protocol, this field is 4 bytes long.
- **Target hardware address:** This is a variable-length field defining the physical address of the target. For Ethernet this field is 6 bytes long. For an

ARP request message, this field is all 0s because the sender does not know the physical address of the target.

- **Target protocol address:** This is a variable-length field defining the logical (for example, IP) address of the target. For the IPv4 protocol, this field is 4 bytes long.

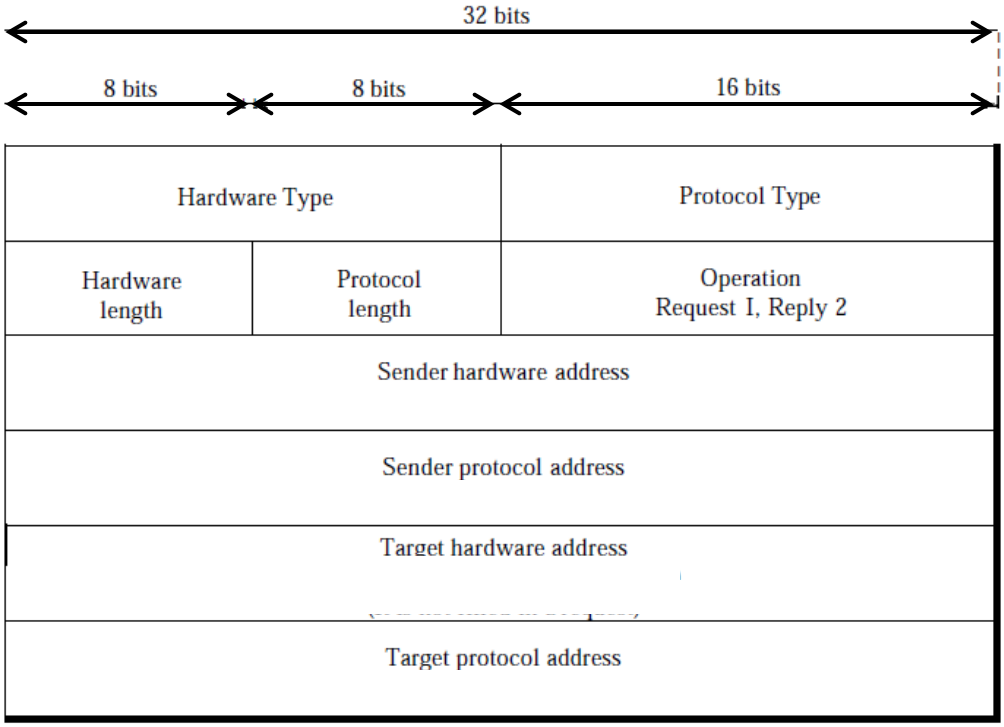

❖ ARP Operation

These are seven steps involved in an ARP process:

1. The sender knows the IP address of the target.
2. IP asks ARP to create an ARP request message, filling in the sender physical address, the sender IP address, and the target IP address. The target physical address field is filled with 0s.
3. The message is passed to the data link layer where it is encapsulated in a frame using the physical address of the sender as the source address and the physical broadcast address as the destination address.

4. Every host or router receives the frame. Because the frame contains a broadcast destination address, all stations remove the message and pass it to ARP. All machines except the one targeted drop the packet. The target machine recognizes the IP address.
5. The target machine replies with an ARP reply message that contains its physical address. This message is unicast.
6. The sender receives the reply message. It now knows the physical address of the target machine.
7. The IP datagram, which carries data for the target machine, is now encapsulated in a frame and is unicast to the destination.

❖ ARP Cases

- **Case 1: The sender is a host and wants to send a packet to another host on the same network. In this case, the logical address that must be mapped to a physical address is the destination IP address in the datagram header.**
- **Case 2: The sender is a host and wants to send a packet to another host on another network. In this case, the host looks at its routing table and finds the IP address of the next hop (router) for this destination. If it does not have a routing table, it looks for the IP address of the default router. The IP address of the router becomes the logical address that must be mapped to a physical address.**
- **Case 3: The sender is a router that has received a datagram destined for a host on another network. It checks its routing table and finds the IP address of the next router. The IP address of the next router becomes the logical address that must be mapped to a physical address.**
- **Case 4: The sender is a router that has received a datagram destined for a host in the same network. The destination IP address of the datagram becomes the logical address that must be mapped to a physical address.**

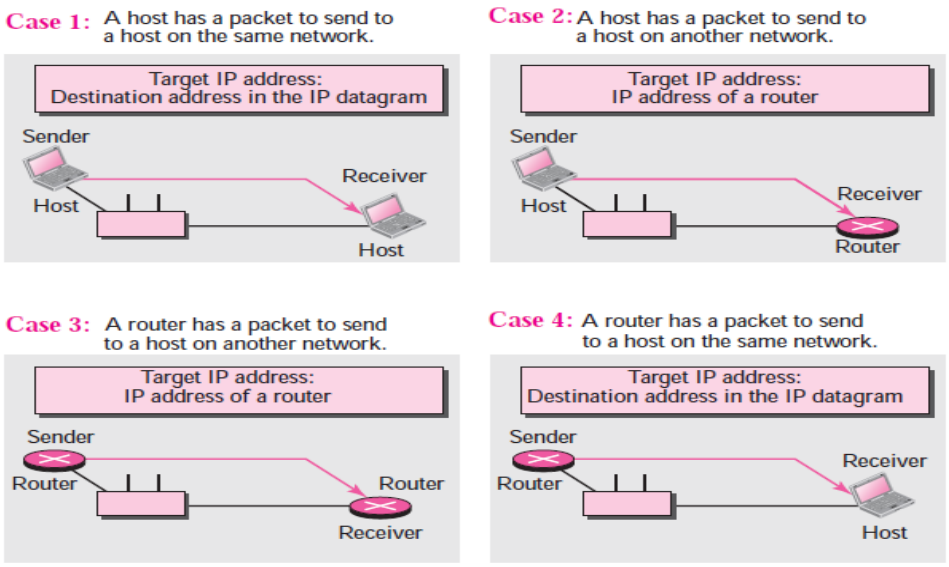

Example: A host with IP address 130.23.43.20 and physical address B2:34:55:10:22:10 has a packet to send to another host with IP address 130.23.43.25 and physical address A4:6E:F4:59:83:AB (which is unknown to the first host). The two hosts are on the same Ethernet network. Show the ARP request and reply packets encapsulated in Ethernet frames.

Solution : ARP request and reply packets are shown below. Note that the ARP data field in this case is 28 bytes, and that the individual addresses do not fit in the 4-byte boundary. That is why we do not show the regular 4-byte boundaries for these addresses. Also note that the IP addresses are shown in hexadecimal.

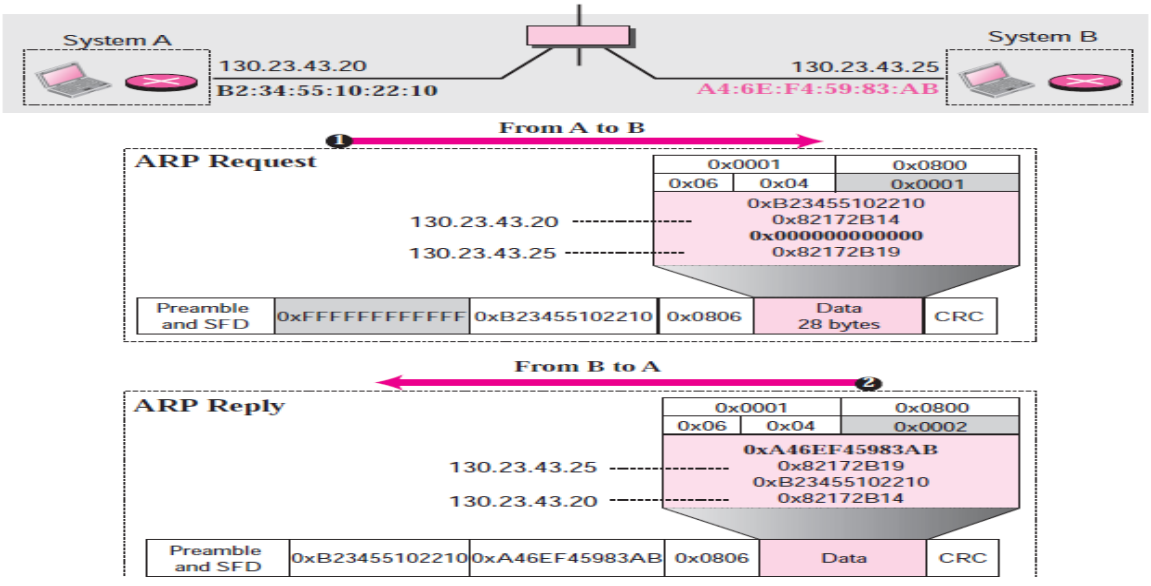

Internet Control Message Protocol (ICMP)

IP which handles most of the network layer job lacks the mechanisms for error reporting or managing host queries. ICMP has been designed to compensate for the above two deficiencies. It is a companion to the IP protocol. ICMP itself is a network layer protocol. However, its messages are not passed directly to the data link layer as would be expected. Instead, the messages are first encapsulated inside IP datagrams before going to the lower layer.

According to the purposes ICMP was designed to serve; ICMP messages are generally divided into two types:

- **Error-reporting messages:** report problems that a router or a host (destination) may encounter when it processes an IP packet.
- **Query messages:** which occur in pairs, help a host or a network manager get specific information from a router or another host.

❖ ICMP Message Format

An ICMP message has an 8-byte header and a variable-size data section. Although the general format of the header is different for each message type, the first 4 bytes are common to all. The following is a brief description for the fields of an ICMP message:

- **Type (1B):** It defines the type of the message whether it is an error-reporting or query message.
- **Code (1B):** This field specifies the reason for the particular message type.
- **Checksum (2B):** It is used for error detection for the entire ICMP packet (header and data portions) in a mechanism similar to that used by IPv4.
- **Rest of the header (variable):** It is specific for each message type.
- **Data (variable):** In error messages carries information for finding the original packet that had the error. In query messages, the data section carries extra information based on the type of the query.

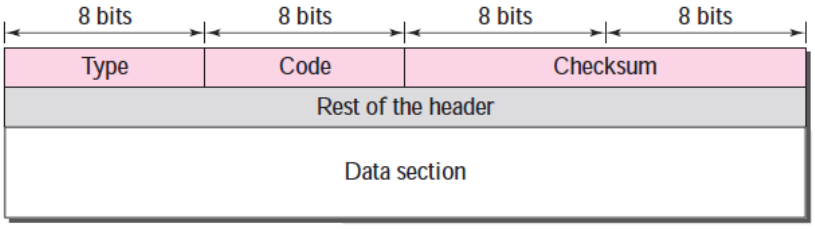

❖ Error Reporting Messages:

ICMP was designed, in part, to compensate error reporting which not a concern of IP. However, ICMP does not correct errors, it simply reports them. Error correction is left to the higher-level protocols. Error messages are always sent to the original source because the only information available in the datagram about the route is the source and destination IP addresses. ICMP uses the source IP address to send the error message to the source (originator) of the datagram.

➢ Encapsulation

All error messages contain a data section that includes the IP header of the original datagram plus the first 8 bytes of data in that datagram. The original datagram header is added to give the original source, which receives the error message, information about the datagram itself. The 8 bytes of data are included because the first 8 bytes provide information about the port numbers (UDP and TCP) and sequence number (TCP). This information is needed so the source can inform the protocols (TCP or UDP) about the error. ICMP forms an error packet, which is then encapsulated in an IP datagram.

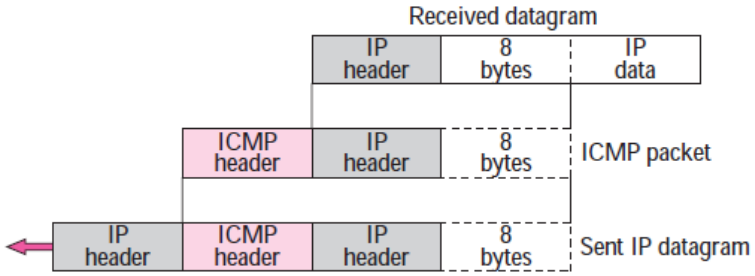

➢ Exceptions:

The following are some exceptions where ICMP error reporting messages cannot be generated:

- No ICMP error message will be generated in response to a datagram carrying an ICMP error message.
- No ICMP error message will be generated for a fragmented datagram that is not the first fragment.
- No ICMP error message will be generated for a datagram having a multicast address.
- No ICMP error message will be generated for a datagram having a special address such as 127.0.0.0 or 0.0.0.0.

➢ Types of Errors:

There are five types of errors are handled by error reporting messages these are:

1. Destination unreachable

When a router cannot route a datagram or a host cannot deliver a datagram, the datagram is discarded and the router or the host sends a destination-unreachable message back to the source host that initiated the datagram. Note that destination-unreachable messages can be created by either a router or the destination host. A Destination unreachable message has a value (3) in type field.

2. Source quench

IP does not have a flow control mechanism embedded in the protocol. The lack of flow control can create congestion in routers or the destination host. This is the case where routers or the destination host has buffers been overwhelmed with datagrams more than the limit they can forward or process. . In this case, the router or the host has no choice but to discard some of the datagrams. The source-quench message in ICMP was designed to add a kind of flow control to the IP. When a router or host discards a datagram due to congestion, it sends a source-quench message to the sender of the datagram. This message has two purposes. First, it informs the source that the datagram has been discarded. Second, it warns the source that there is congestion somewhere in the path and that the source should slow down (quench) the sending process. A source quench ICMP message has the type field set to (4).

3. Time exceeded

The time-exceeded message is generated in two cases. first; If there are errors in one or more routing tables, a packet can travel in a loop or a cycle, going from one router to the next or visiting a series of routers endlessly. In this case each router forwarding the datagram decrements the time to live field value by 1. When the time-to-live value reaches 0, after decrementing, the router discards the datagram. However, when the datagram is discarded, a time-exceeded message must be sent by the router to the original source. Second, a time-exceeded message is also generated when not all fragments that make up a message arrive at the destination host within a certain time limit. The type field in a time exceeded message is set to (11).

4. Parameter problems

Any ambiguity in the header part of a datagram can create serious problems as the datagram travels through the Internet. If a router or the destination host discovers an ambiguous or missing value in any field of the datagram, it discards the datagram and sends a parameter-problem message back to the source. The type field in a this message is set to (12).

5. Redirection

When a host comes up, its routing table has a limited number of entries. It usually knows the IP address of only one router, the default router. For this reason, the host may send a datagram, which is destined for another network, to the wrong router. In this case, the router that receives the datagram will forward the datagram to the correct router. However, to update the routing table of the host, it sends a redirection message to the host. A redirection ICMP message has the type field set to (5).

This concept of redirection is shown in the figure below. Host A wants to send a datagram to host B. Router R2 is obviously the most efficient routing choice, but host A did not choose router R2. The datagram goes to R1 instead. R1, after consulting its table, finds that the packet should have gone to R2. It sends the packet to R2 and, at the same time, sends a redirection message to host A. Host A's routing table can now be updated.

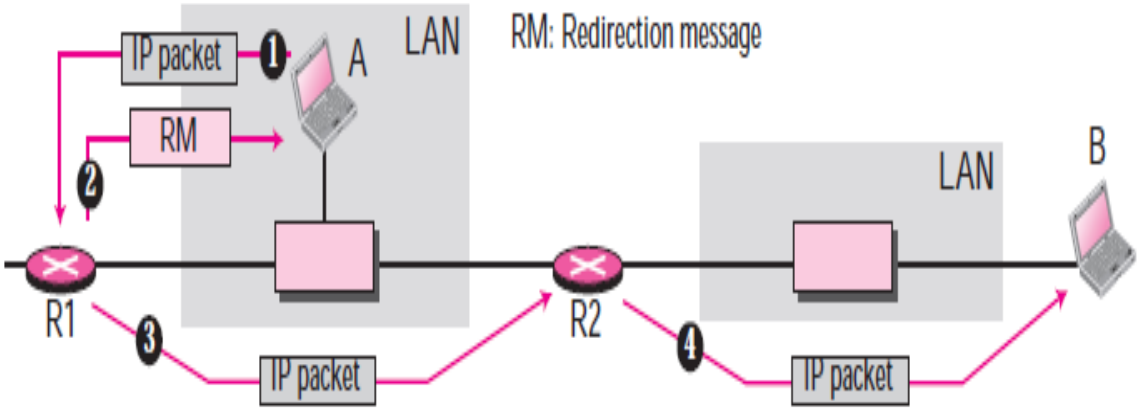

❖ Query messages

In addition to error reporting, ICMP can diagnose some network problems. This is accomplished through the query messages, a group of four different pairs of messages as will be seen in later sections of this lecture.

➢ Encapsulation

An ICMP query message is encapsulated in an IP packet, which in turn encapsulated in a data link layer frame. However, in this case, no bytes of the original IP are included in the message.

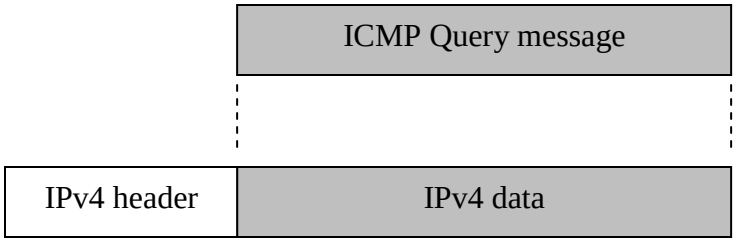

➢ Types of Queries

As mentioned earlier; there are four different pairs of query messages. These are:

1. Echo Request and Reply

The echo-request and echo-reply messages are designed for diagnostic purposes. Network managers and users utilize this pair of messages to identify network problems. The combination of echo-request and echo-reply messages finds whether two systems (hosts or routers) can communicate with each other. The echo-request and echo-reply messages can be used to determine if there is communication at the IP level. Because ICMP messages are encapsulated in IP datagrams, the receipt of an echo-reply message by the machine that sent the echo request is proof that the IP protocols in the sender and receiver are communicating with each other using the IP datagram. Also, it is proof that the intermediate routers are receiving, processing, and forwarding IP datagrams. Most systems provide a version of the ping command that can create a series (instead of just one) of echo-request and echo-reply messages, providing statistical information. The type field in the echo request message is set to (8) while it is set to (0) in its reply message.

2. Timestamp Request and Reply

Two machines (hosts or routers) can use the timestamp request and timestamp reply messages to determine the round-trip time needed for an IP datagram to travel between them. It can also be used to synchronize the clocks in two machines. The type field in the timestamp request message is set to (13) while it is set to (14) in its reply message.

3. Address-Mask Request and Reply

A host may know its IP address, but it may not know the corresponding mask. To obtain its mask, a host sends an address-mask-request message to a router on the LAN. If the host knows the address of the router, it sends the request directly to the router. If it does not know, it broadcasts the message. The router receiving the address-mask-request message responds with an address-mask-reply message, providing the necessary mask for the host. The type field in the address-mask request message is set to (17) while it is set to (18) in its reply message.

4. Router Solicitation and Advertisement

As we discussed in the redirection message section, a host that wants to send data to a host on another network needs to know the address of routers connected to its own network. Also, the host must know if the routers are alive and functioning. The router-solicitation and router-advertisement messages can help in this situation. A host can broadcast (or multicast) a router-solicitation message. The router or routers that receive the solicitation message broadcast their routing information using the router-advertisement message. A router can also periodically send router-advertisement messages even if no host has solicited. Note that when a router sends out an advertisement, it announces not only its own presence but also the presence of all routers on the network of which it is aware. The type field is set to (10) in the solicitation message while it is set to (9) in the advertising message.