Chapter 9

Internet Control Message Protocol (ICMP)

Outline

- Types of messages
- Message format
- Error reporting
- Query
- Checksum
- ICMP package

Introduction

- IP provides unreliable and connectionless datagram delivery
- Drawbacks
  - Lack of error control mechanism
  - Lack of assistance mechanism
- Solution
  - ICMP

Lack of Error Control Mechanism

- No error-reporting or error-correcting mechanism
  - What happens if a router must discard a datagram because
    - Cannot find a router to the destination
    - Time-to-live field has a zero value
  - What happens if the final destination host must discard all fragments of a datagram
    - Because it has not received all fragments within a predefined time limit

Lack of Assistance Mechanism

- Lack of a mechanism for host and management queries
  - How to determine if a router or another host is alive?
  - How to obtain information from another host or router?

Solution

- ICMP: *Internet Control Message Protocol*
  - A network layer protocol
  - However, it message are not passed directly to the data link layer
  - The message are first encapsulated inside IP datagram before going to the lower layer

Position of ICMP in the Network Layer

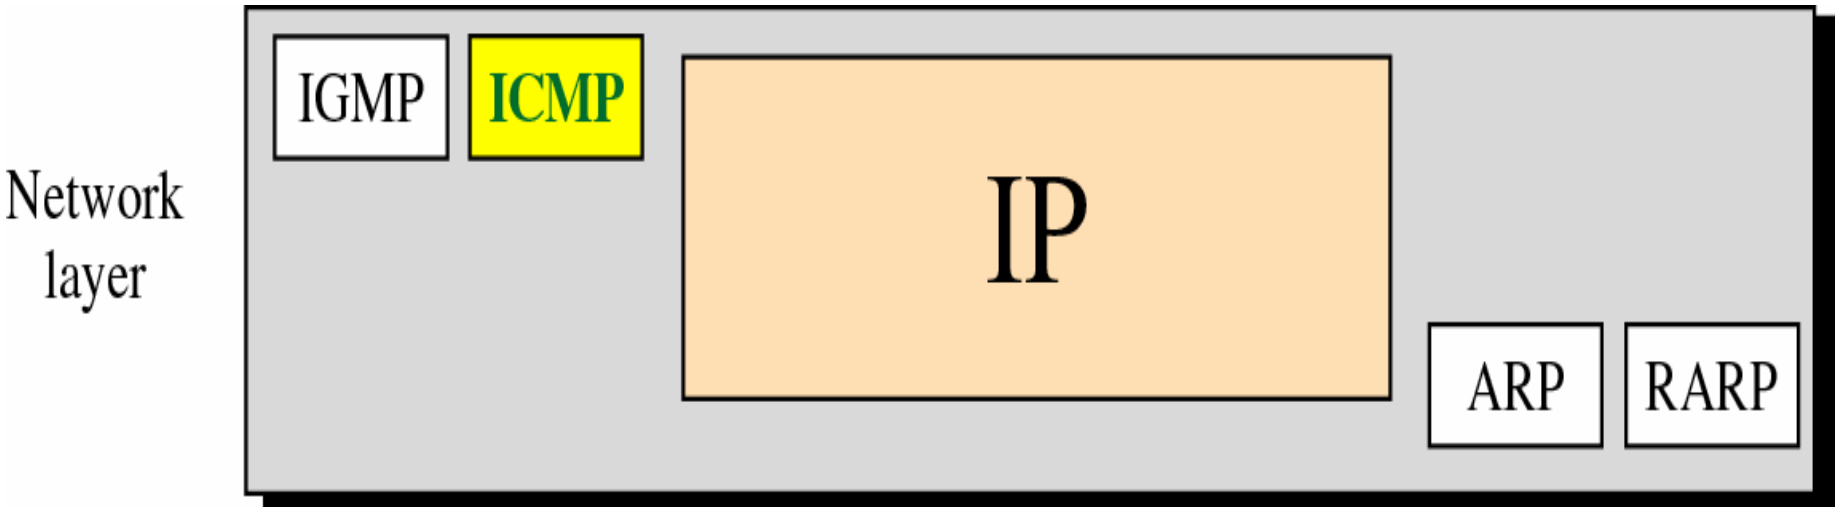

Figure 9.2

Encapsulation of ICMP Packet

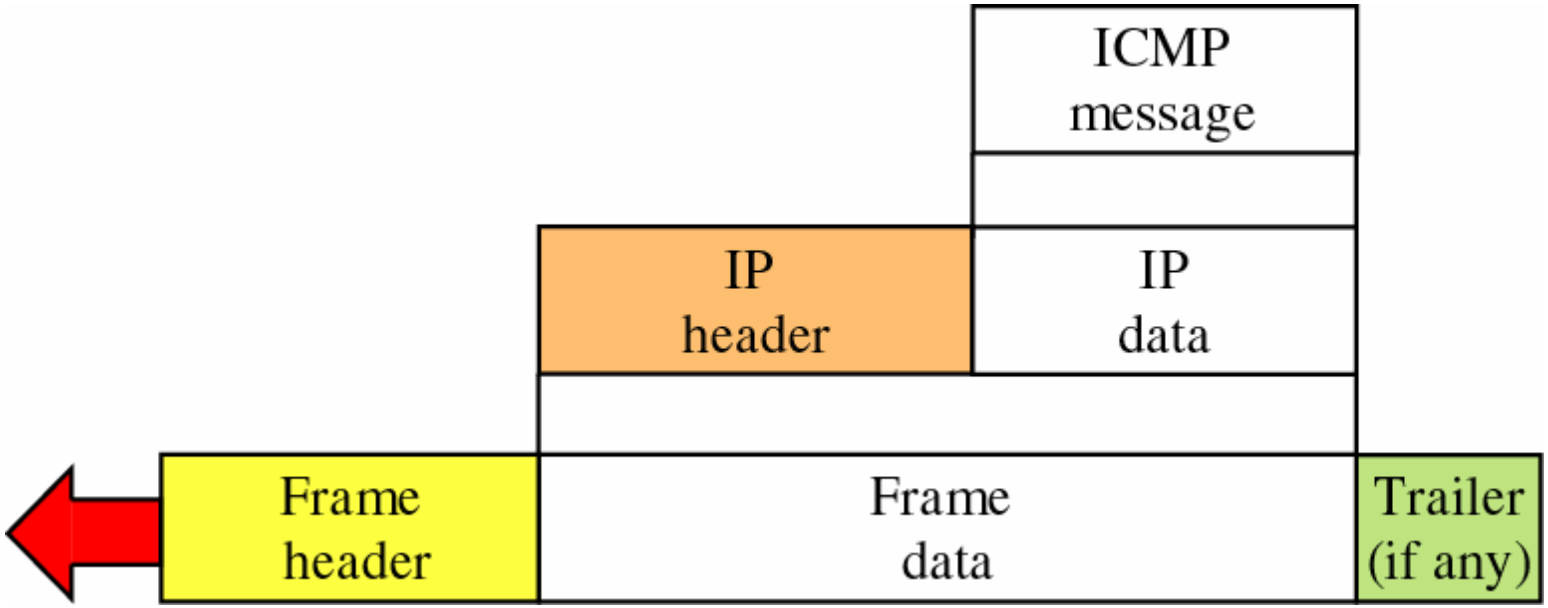

9.1

TYPES OF MESSAGES

ICMP Messages

- ICMP messages are divided into
  - *Error-reporting message*
    - Report problems that a router or a host (destination) may encounter when it processes an IP packet
  - *Query message*
    - Help a router or a network manager to get specific information from a router or another host

ICMP Messages

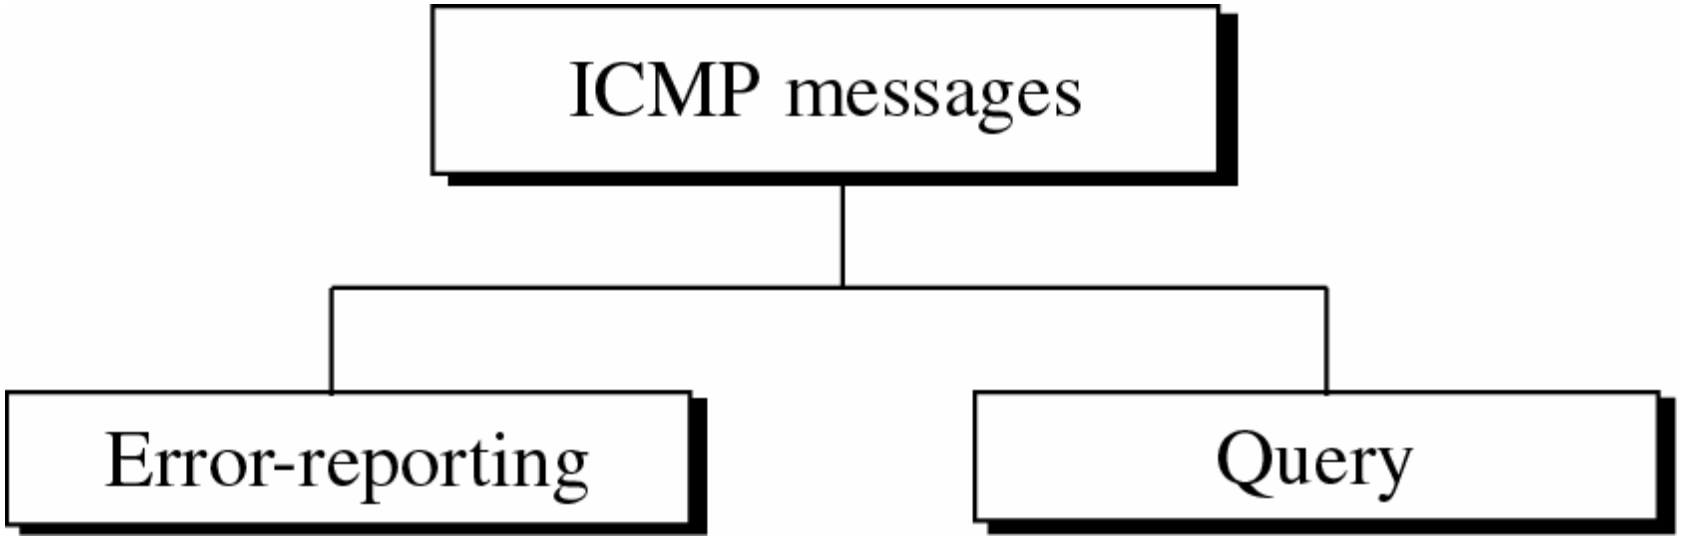

ICMP Messages

Category	Type	Message
Error-reporting message	3	Destination unreachable
	4	Source quench
	11	Time exceeded
	12	Parameter problem
	5	Redirection
Query messages	8 or 0	Echo request to reply
	13 or 14	Timestamp request or reply
	17 or 18	Address mask request or reply
	10 or 9	Router solicitation or advertisement

9.2

MESSAGE FORMAT

Figure 9.4

General Format of ICMP Messages

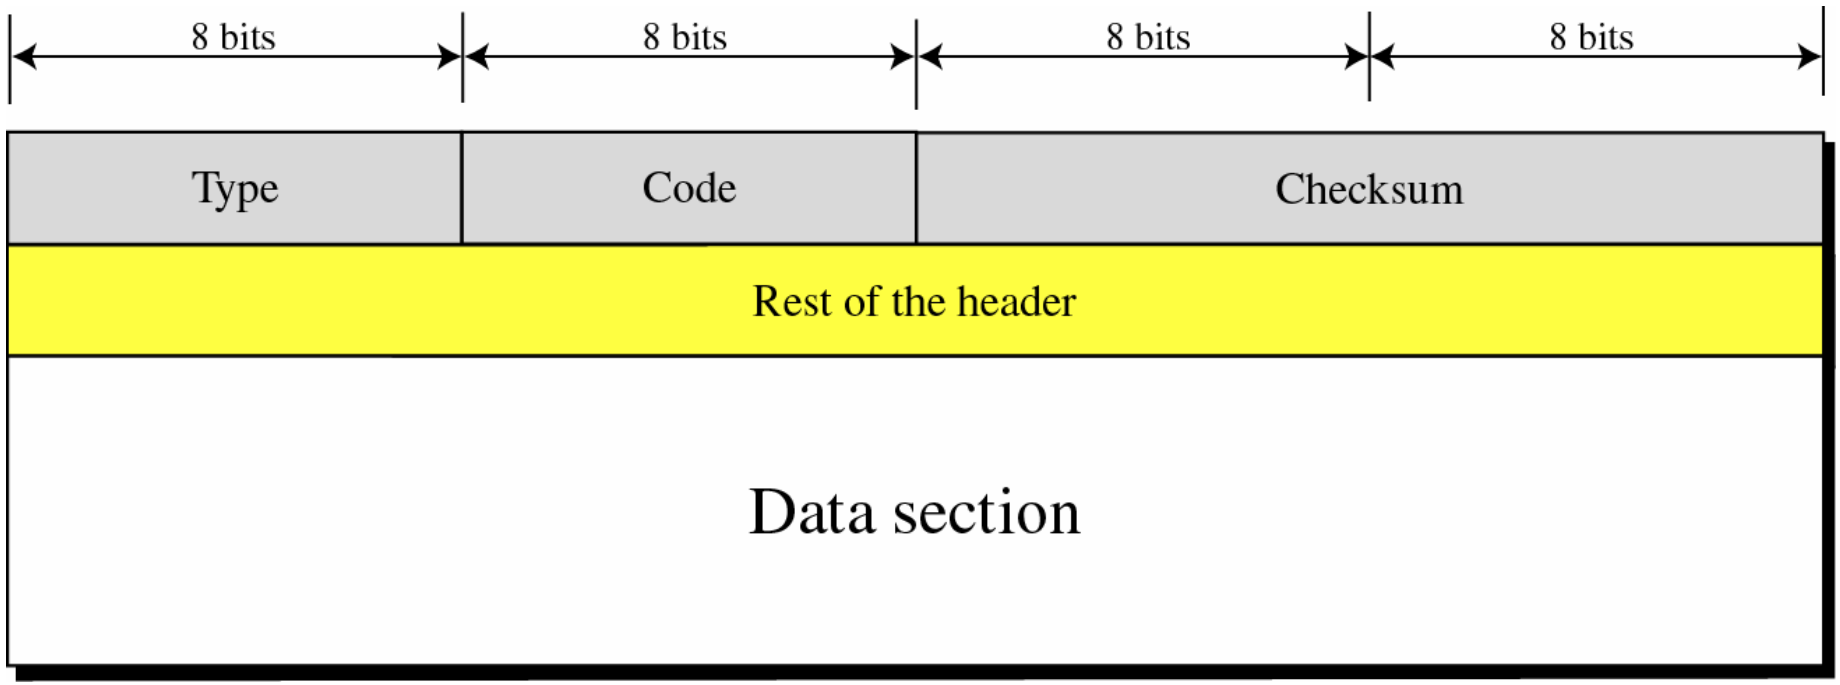

Format of ICMP Message

- 8-byte header
  - The first 4-byte are common to all
    - Type(1-byte): define the type of the message
    - Code(1-byte): specify the reason for the particular message type
    - Checksum(2-byte)
  - The rest is specific for each message type
- A variable-size data section
  - For error message
    - Carries information for finding the *original packet* that had the error
  - For query message
    - Carries extra information based on the type of the query

9.3

ERROR REPORTING

Error Reporting

- ICMP only report error
  - Does not correct error
  - Error correction is left to the higher-level protocol

- Error message are always sent to the original source
  - Because the only information available in the datagram is the *source* and *destination IP address*

Note

ICMP always reports error messages to the original source.

Figure 9.5

Error-Reporting Messages

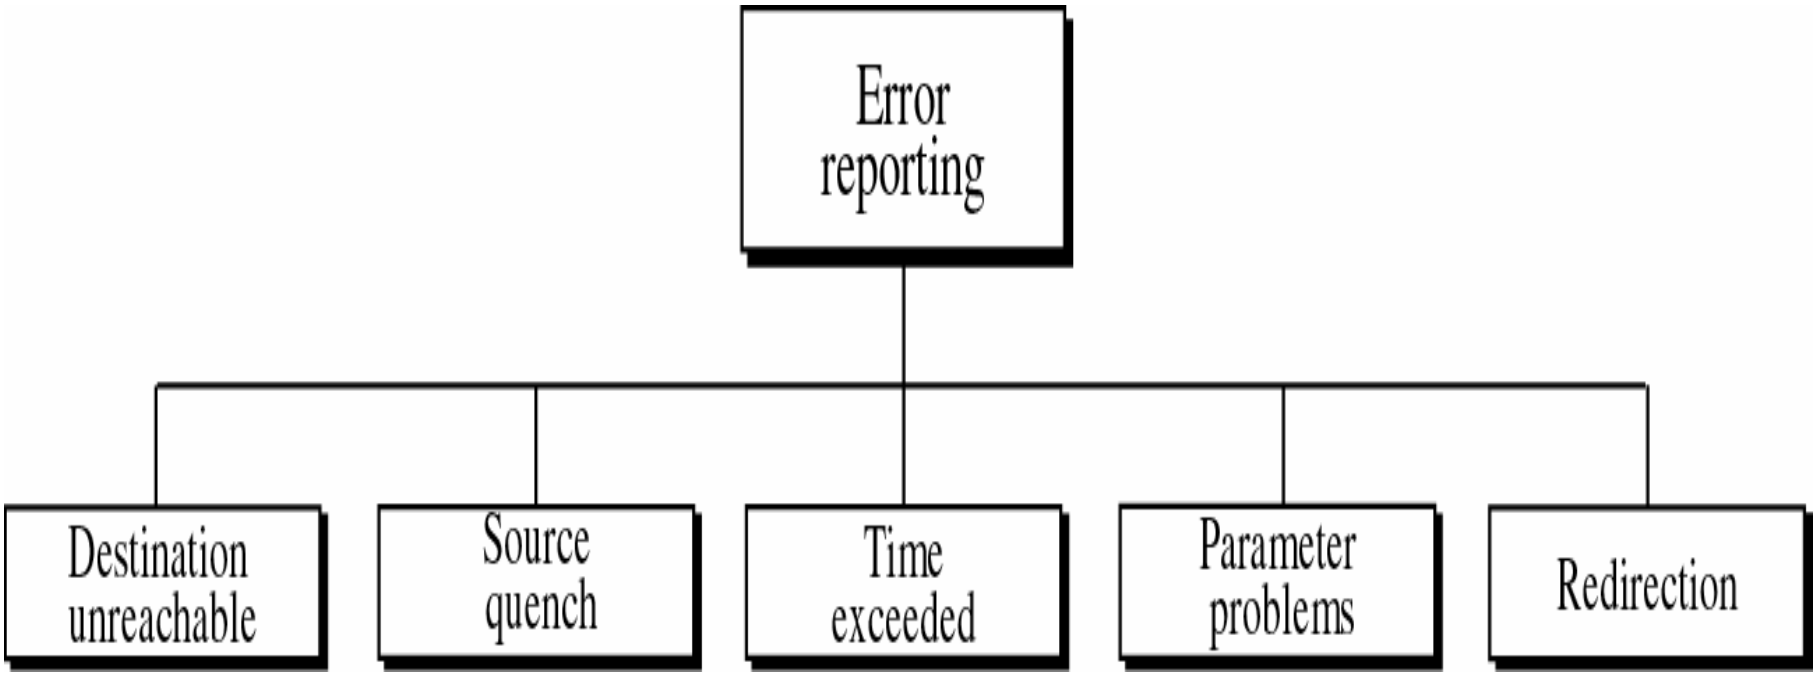

Important Points about ICMP Error Messages

- ***No ICMP error message for a datagram carrying an ICMP error message***
- ***No ICMP error message for a fragmented datagram that is not the first fragment***
- ***No ICMP error message for a datagram having a multicast address.***
- ***No ICMP error message for a datagram with a special address such as 127.0.0.0 or 0.0.0.0.***

ICMP Packet Data Section

- The data section in all error message contain includes
  - The *IP header* of the original datagram
    - Give the original source information about the datagram itself
  - The *first 8-byte of data* in that datagram
    - Provides information about the port number (UDP and TCP) and sequence number (TCP)
    - Source then can inform the upper layer protocols (TCP or UDP) about the error

Contents of Data Field for Error Messages

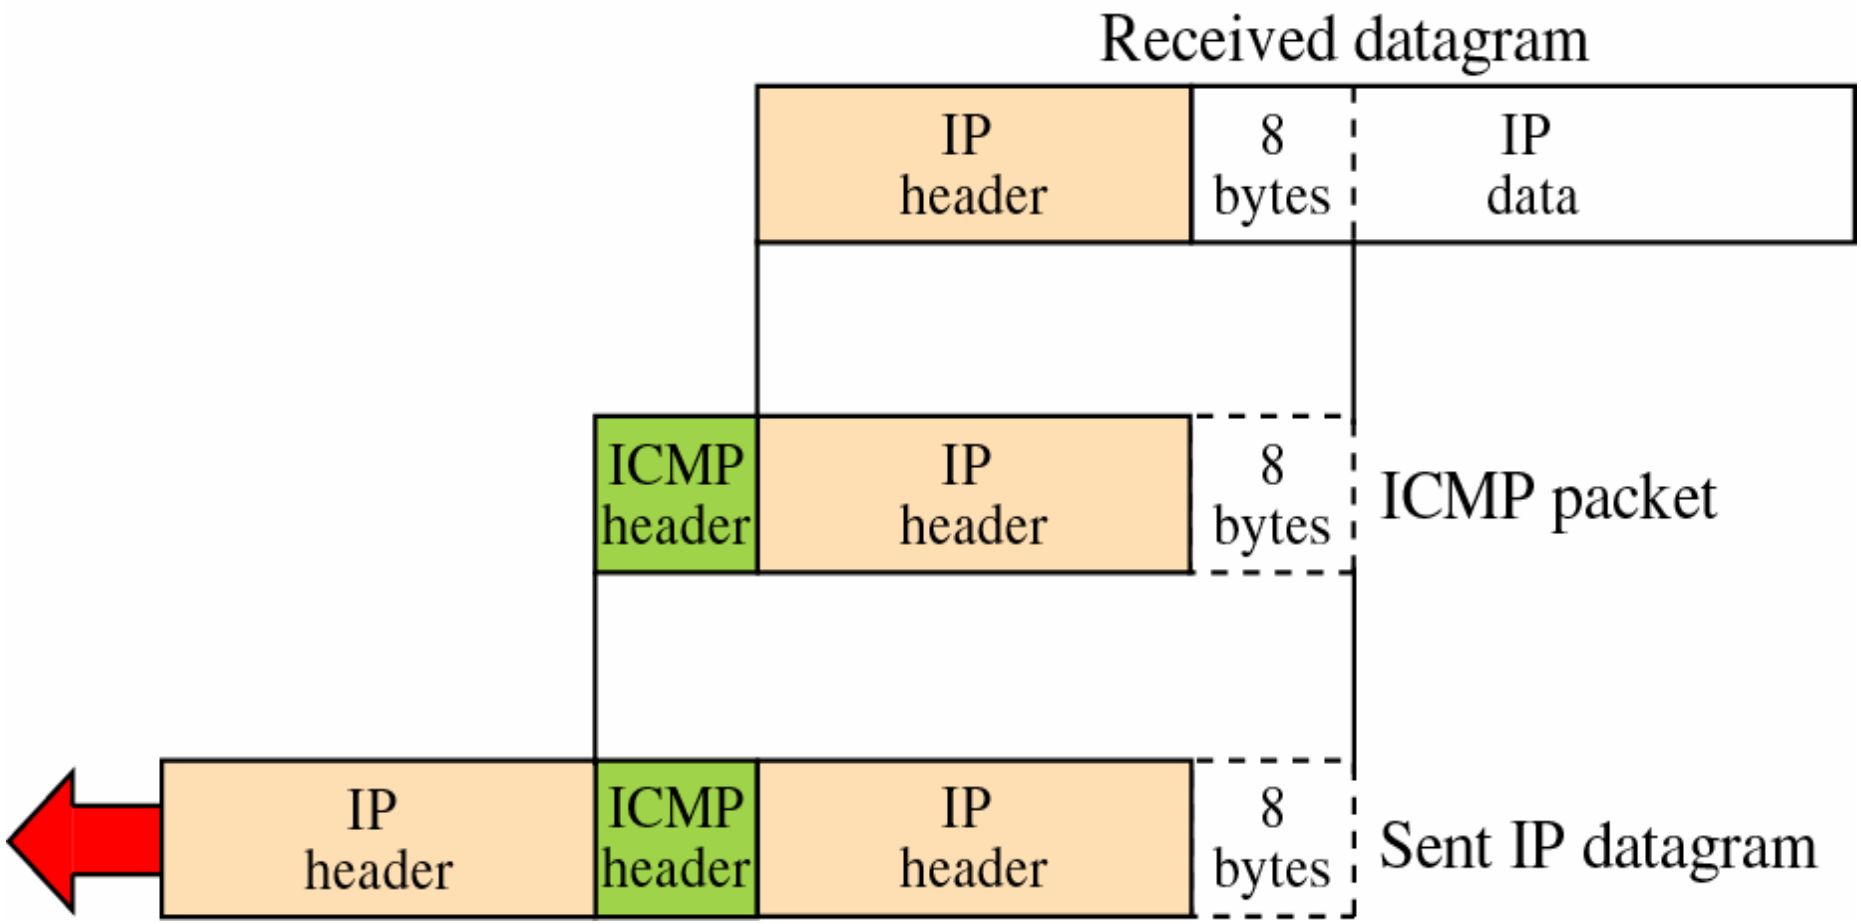

Destination Unreachable

- When a router cannot route a datagram or a host cannot deliver a datagram
  - Discard the datagram
  - The router or host sends a destination unreachable message back to the source host

Figure 9.7

Destination-Unreachable Format

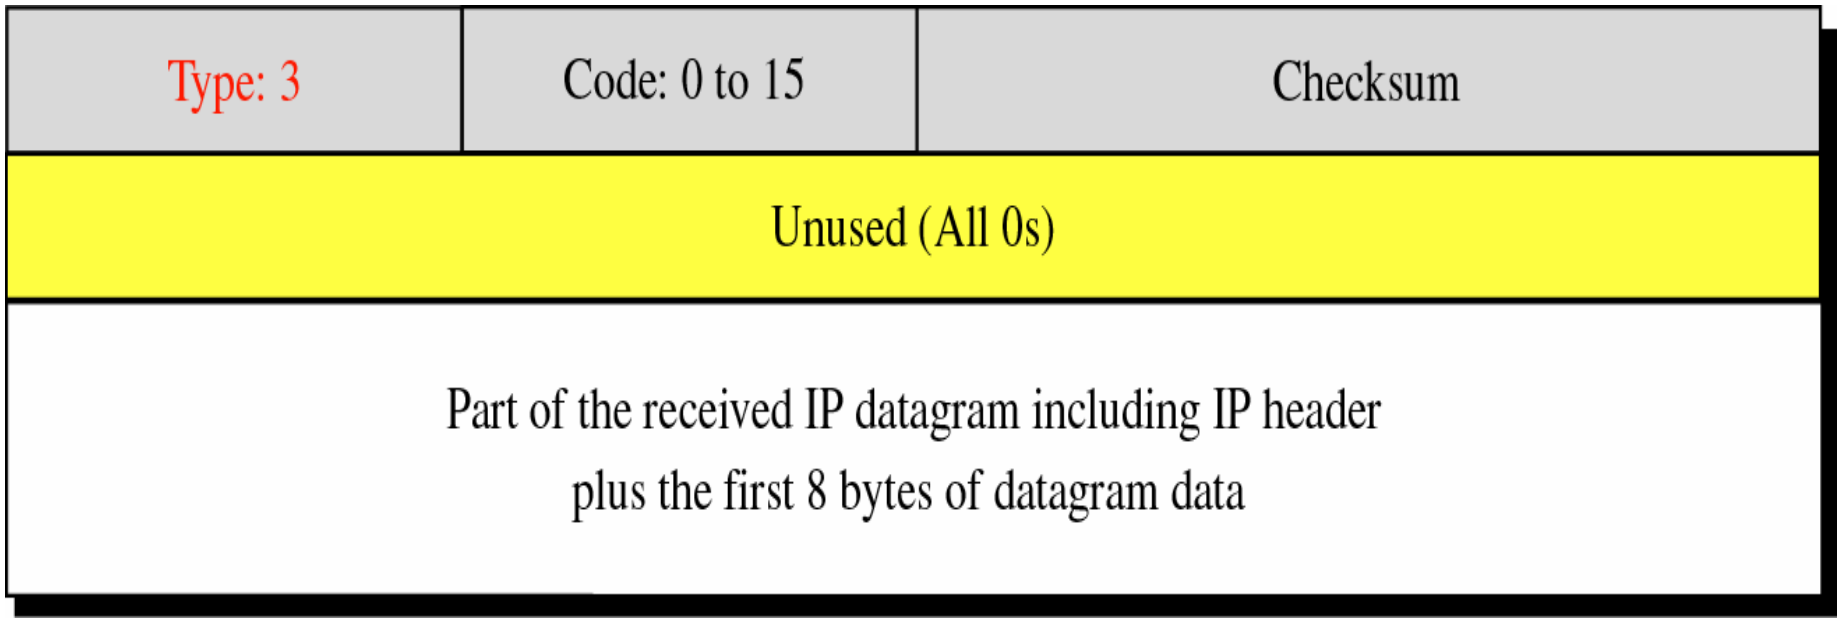

Code Field for Destination-Unreachable

- Code 0: the network is unreachable
  - Possibly due to hardware failure
  - Can only be generated by a *router*
- Code 1: the host is unreachable
  - Possibly due to hardware failure
  - Can only be generated by a *router*
- Code 2: the protocol is unreachable
  - Delivery to the upper layer protocol (TCP, UDP) is failed
  - Can only be generated by a *destination host*

Code Field for Destination-Unreachable (Cont.)

- Code 3: the port is unreachable
  - Can only be generated by a *destination host*
- Code 4: fragmentation is required, but the DF (do not fragment) field has been sent
  - The sender specifies no fragmentation
  - But the router is impossible without fragmentation
- Code 5: source routing cannot be accomplished
  - One or more routers defined in the source routing cannot be visited

Code Field for Destination-Unreachable (Cont.)

- Code 6: the destination network is unknown
  - In code 0: the router knows that the destination network exists, but it is unreachable at the moment
  - In code 6: the router has no information about the destination network
- Code 7: the destination host is unknown
  - In code 1: the router knows that the destination host exists, but it is unreachable at the moment
  - In code 7: the router is unaware of the existence of the destination host

Code Field for Destination-Unreachable (Cont.)

- Code 8: the source host is isolated
- Code 9: communication with the destination network is administratively prohibited
- Code 10: communication with the destination host is administratively prohibited
- Code 11: the network is unreachable for the specified type of service
- Code 12: the host is unreachable for the specified type of service

Code Field for Destination-Unreachable (Cont.)

- Code 13: the host is unreachable because the administrator has put a filter on it
- Code 14: the host is unreachable because the host precedence is violated
  - The requested precedence is not permitted for the destination
- Code 15: the host is unreachable because its precedence was cut off

Note

Destination-unreachable messages with codes 2 or 3 can be created only by the destination host.

Other destination-unreachable messages can be created only by routers.

Error Reporting

- Even if a router does not report a destination-unreachable message
  - Does not mean that the datagram has been delivered
  - For example, in a Ethernet network, there is no way that a router knows a packet has been delivered to the destination or the next router
    - Ethernet does not provide an acknowledge mechanism

Note

A router cannot detect all problems that prevent the delivery of a packet.

Source Quench

- In IP, there is no communication between the source host, the router, and the destination
- As a result, no *flow control* in IP
  - The source never knows if it is producing datagram faster than can be
    - Forwarded by router
    - Processed by the destination host
- Problem
  - Congestion in routers or the destination host

Source Quench (Cont.)

- Source-quench message thus adds a kind of flow control to the IP
  - Inform the source that the datagram has been discarded
  - Warm the source that
    - There is congestion somewhere in the path
    - The source should slow down (quench) the sending process

Figure 9.8

Source-Quench Format

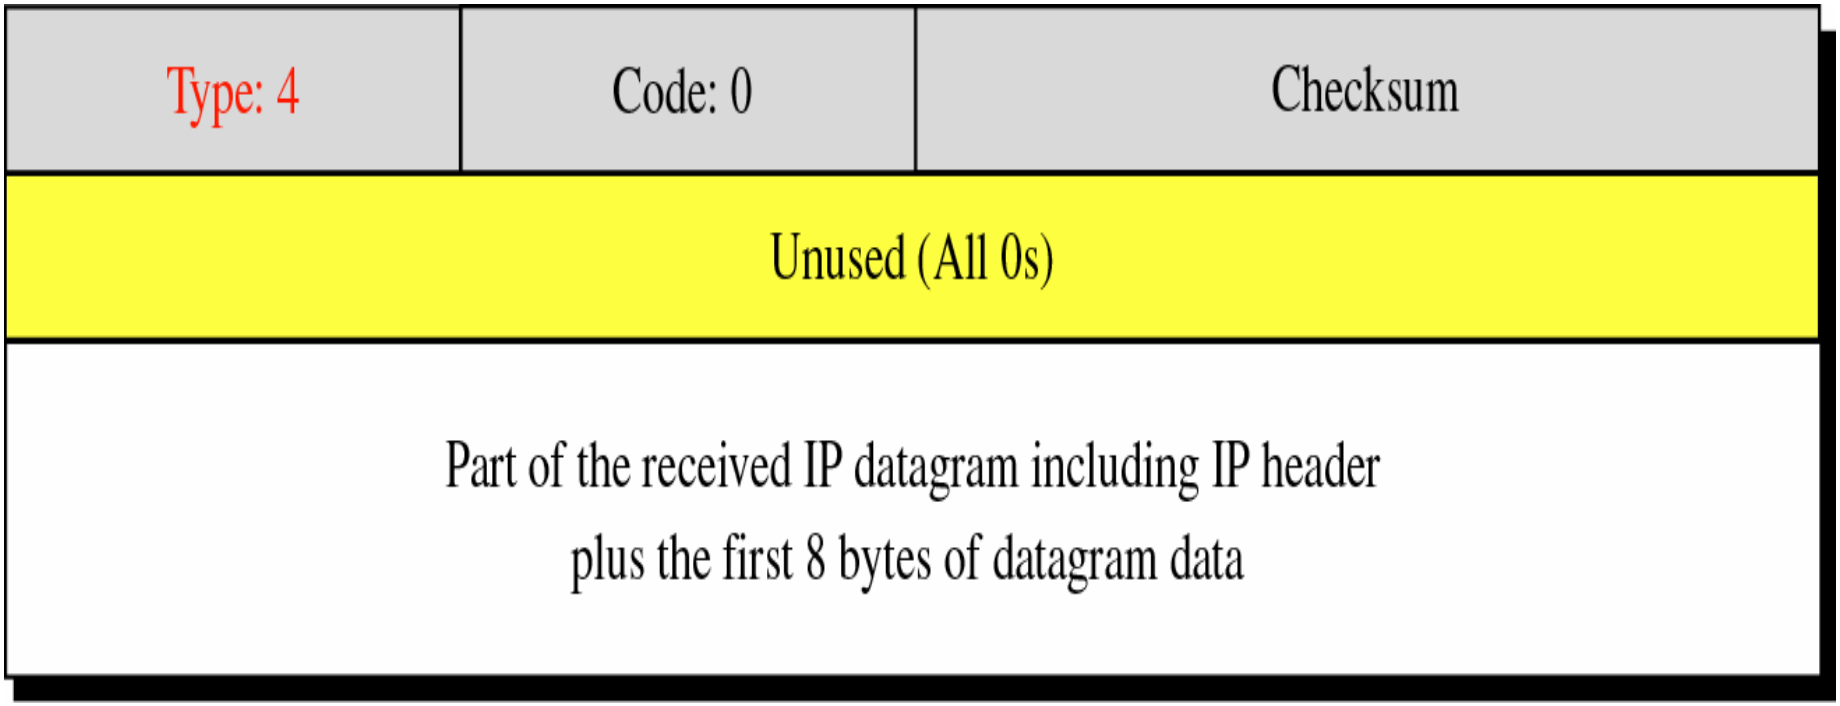

Note about the Source Quench

- One source-quench message should be sent *for each datagram* that is discarded due to congestion
- There is no mechanism to tell the source that the congestion has been relieved and the source can speed up its sending rate
  - The source continue to slower that rate until no more source-quench message are received

Note about the Source Quench (Cont.)

- The congestion can be created either by one-to-one or many-to-one communication
  - In one-to-one: source quench is helpful
  - In many-to-one: may be useless
    - Each source sends datagram at a different rate
    - The router or the destination has no clue which source is responsible for the congestion
    - Thus, it may drop a datagram from a very slow source

Time Exceeded: Two Situations

- The packet travel in a loop or a cycle
  - Caused by errors in the routing table
  - Finally, *time-to-live* value is 0
  - The router discards the datagram and sends time-exceeded message
- When all fragments that make up a message do not arrive at the destination within a certain time limit
  - When the first datagram arrives at the destination, it starts a timer
  - When the timer expires and all the fragments are not arrived
  - The destination discards all the fragments and sends a time-exceeded message

Note

Whenever a router receives a datagram with a time-to-live value of zero, it discards the datagram and sends a time-exceeded message to the original source.

Note

When the final destination does not receive all of the fragments in a set time, it discards the received fragments and sends a time-exceeded message to the original source.

Time-Exceeded Message Format

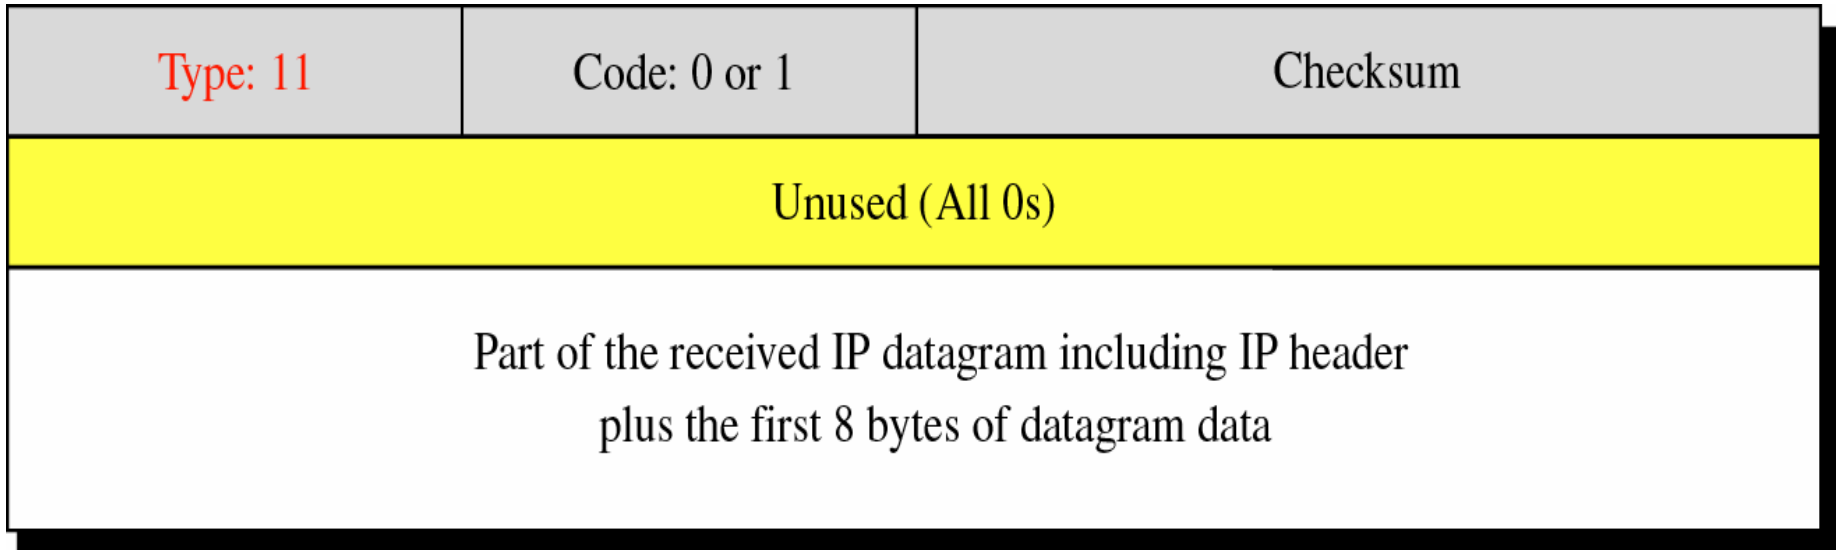

Code 0: Time to live is zero

Code 1: Fragmentations are not arrived with a set time

Note

In a time-exceeded message, code 0 is used only by routers to show that the value of the time-to-live field is zero. Code 1 is used only by the destination host to show that not all of the fragments have arrived within a set time.

Parameter Problem

- Occurred when a router or a destination discovers an ambiguous or missing value in any field of the datagram
- Code 0: there is an error or ambiguity in one of the *header* fields
  - Pointer field points to the byte within the problem
- Code 1: the required part of an *option* is missing

Parameter-Problem Message Format

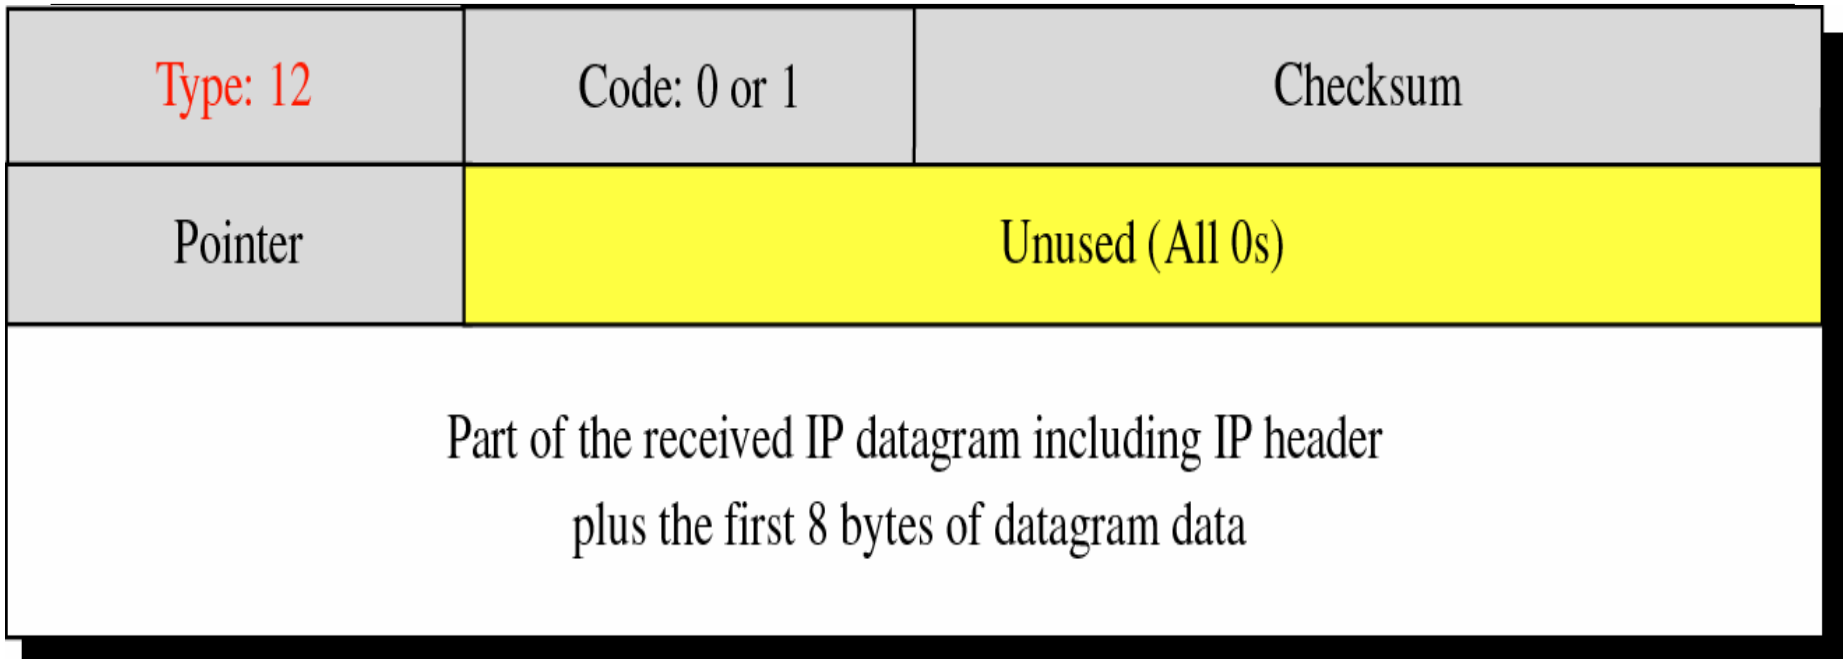

Code 0: Main header problem

Code 1: Problem in the option field

Note

A parameter-problem message can be created by a router or the destination host.

Redirection

- Routing table is updated dynamically
- However, for efficiency, hosts do not take part in the routing update process
  - There are terrible number of hosts
  - Host thus use *static routing*
    - Usually knows only one IP address of the router, the default router

Redirection (Cont.)

- Thus, the host may send a datagram to the wrong router

- Solution
  - A router can send a *redirection message* to the host

Figure 9.11

Redirection Concept

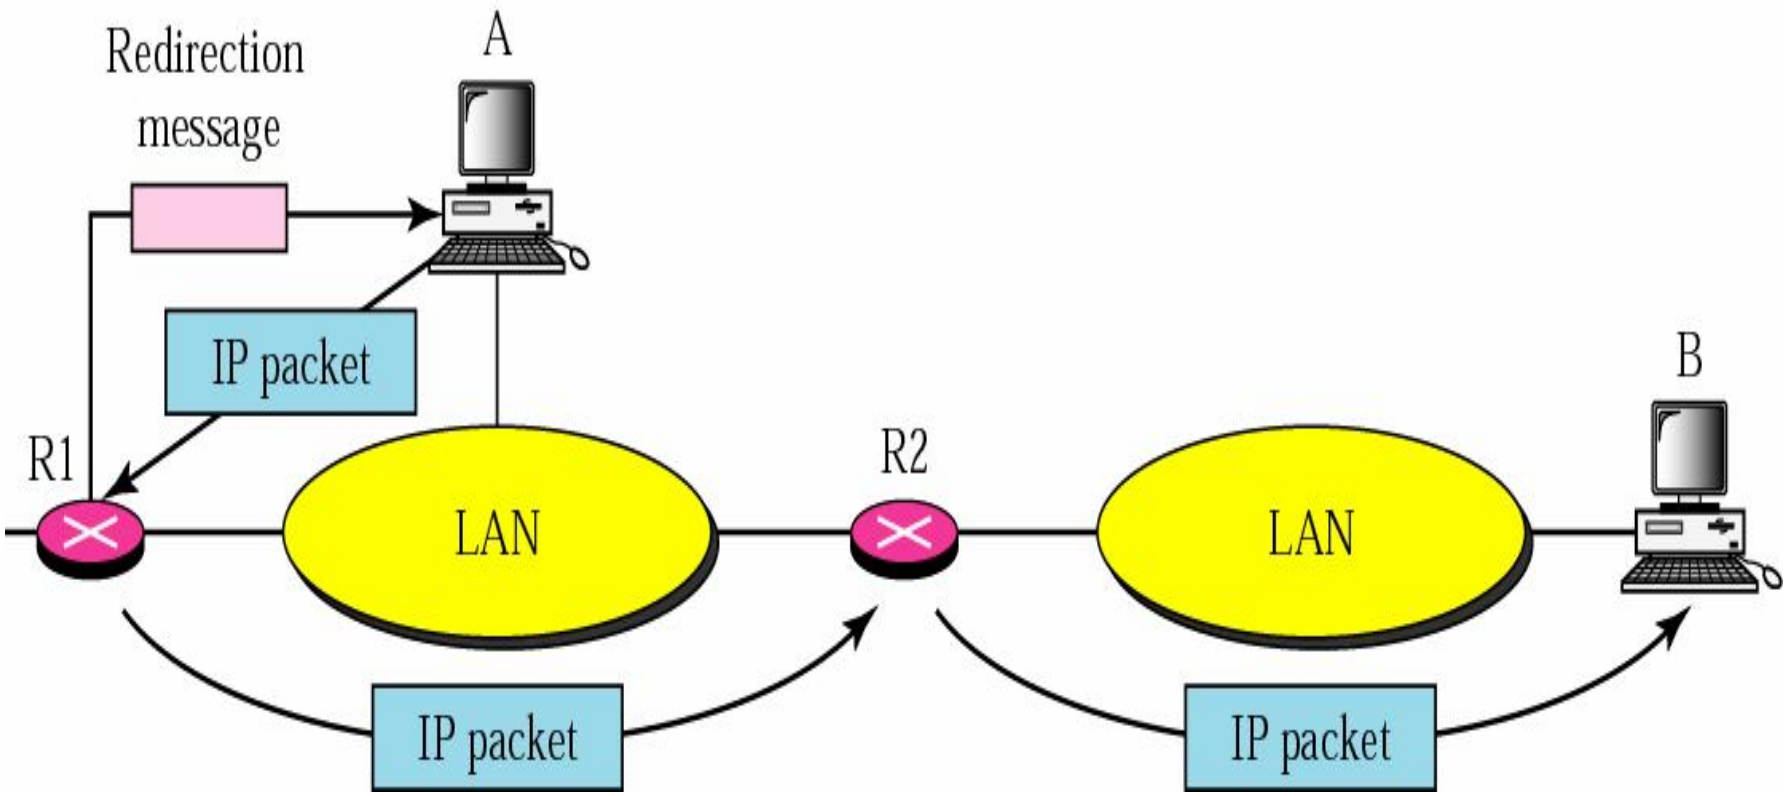

Note

A host usually starts with a small routing table that is gradually augmented and updated. One of the tools to accomplish this is the redirection message.

Figure 9.13

Redirection Message Format

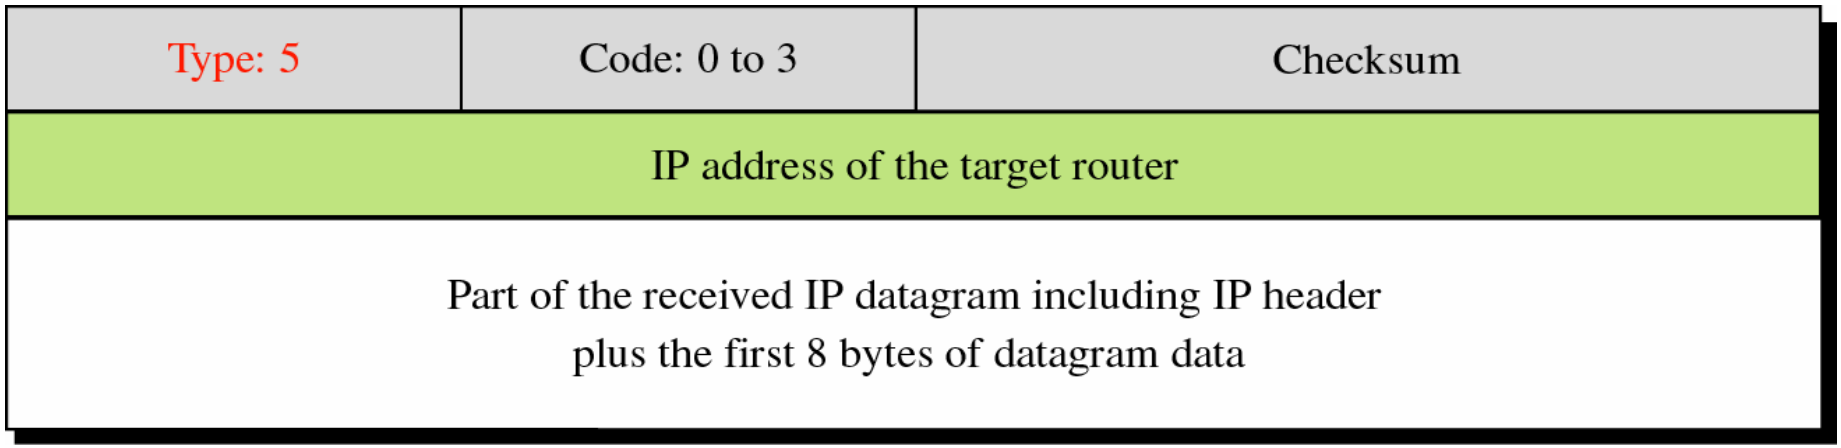

Code 0: Redirection for a network-specific route
Code 1: Redirection for a host-specific route
Code 2: Redirection for a network-specific route based on a specified type of service
Code 3: Redirection for a host-specific route based on a specified type of service

Note

A redirection message is sent from a router to a host on the same local network.

9.4

QUERY

Query

- ICMP can also diagnose some network problems
  - Accomplished by the query message
  - A group of four different pair of messages

Figure 9-13

Query Messages

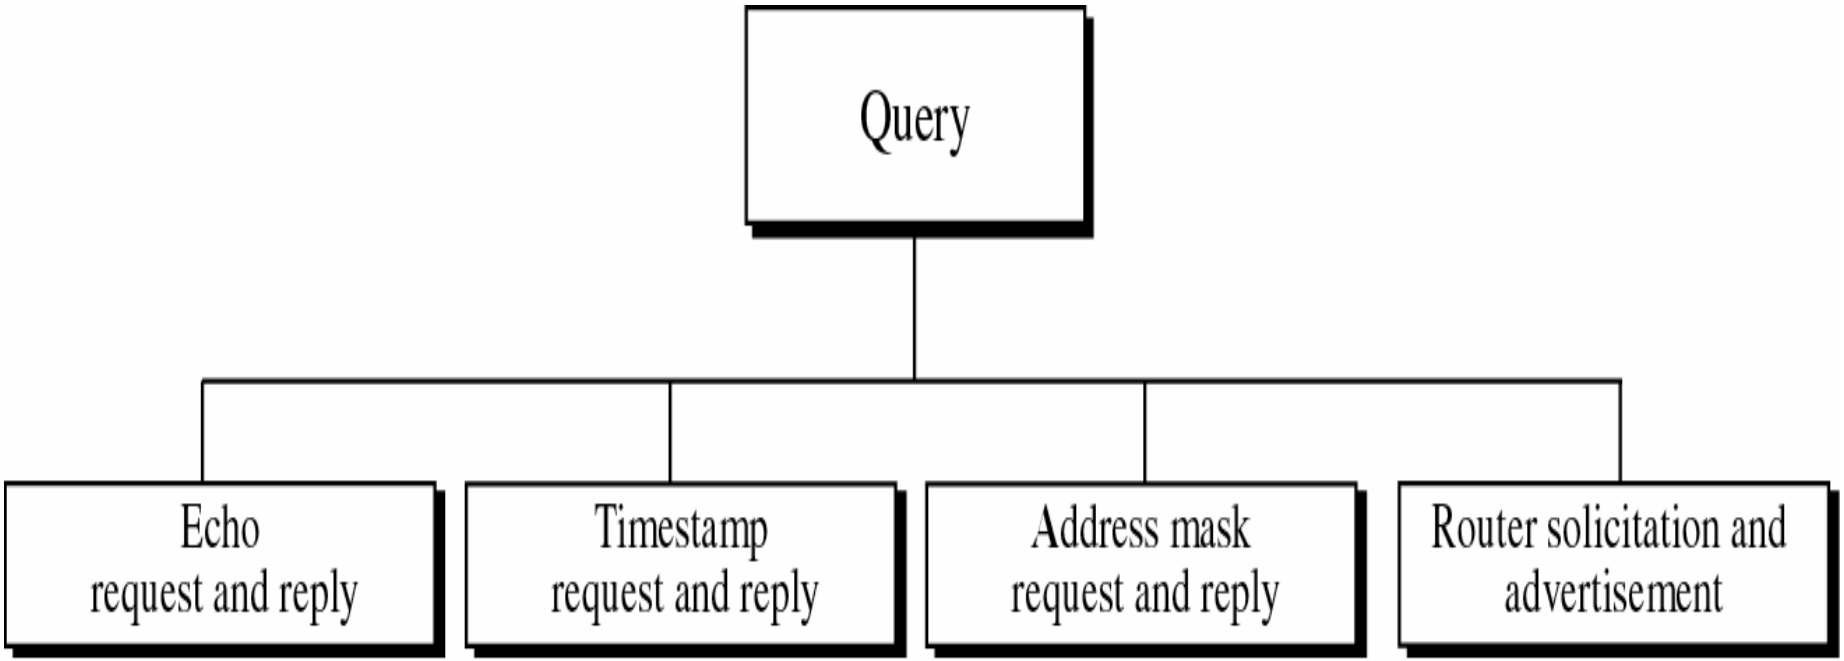

Echo Request and Reply

- Determine whether two systems (hosts or routers) can communicate with each other
  - Determine if there is communication at the IP level
    - Because ICMP are encapsulated in IP datagram
  - Also be used by a host to see if another host is reachable
    - At the user level, this is done by *ping* command

Note

An echo-request message can be sent by a host or router.

An echo-reply message is sent by the host or router which receives an echo-request message.

Note

Echo-request and echo-reply messages can be used by network managers to check the operation of the IP protocol.

Note

Echo-request and echo-reply messages can test the reachability of a host. This is usually done by invoking the ping command.

Figure 9.14

Echo-Request and Echo-Reply Message Format

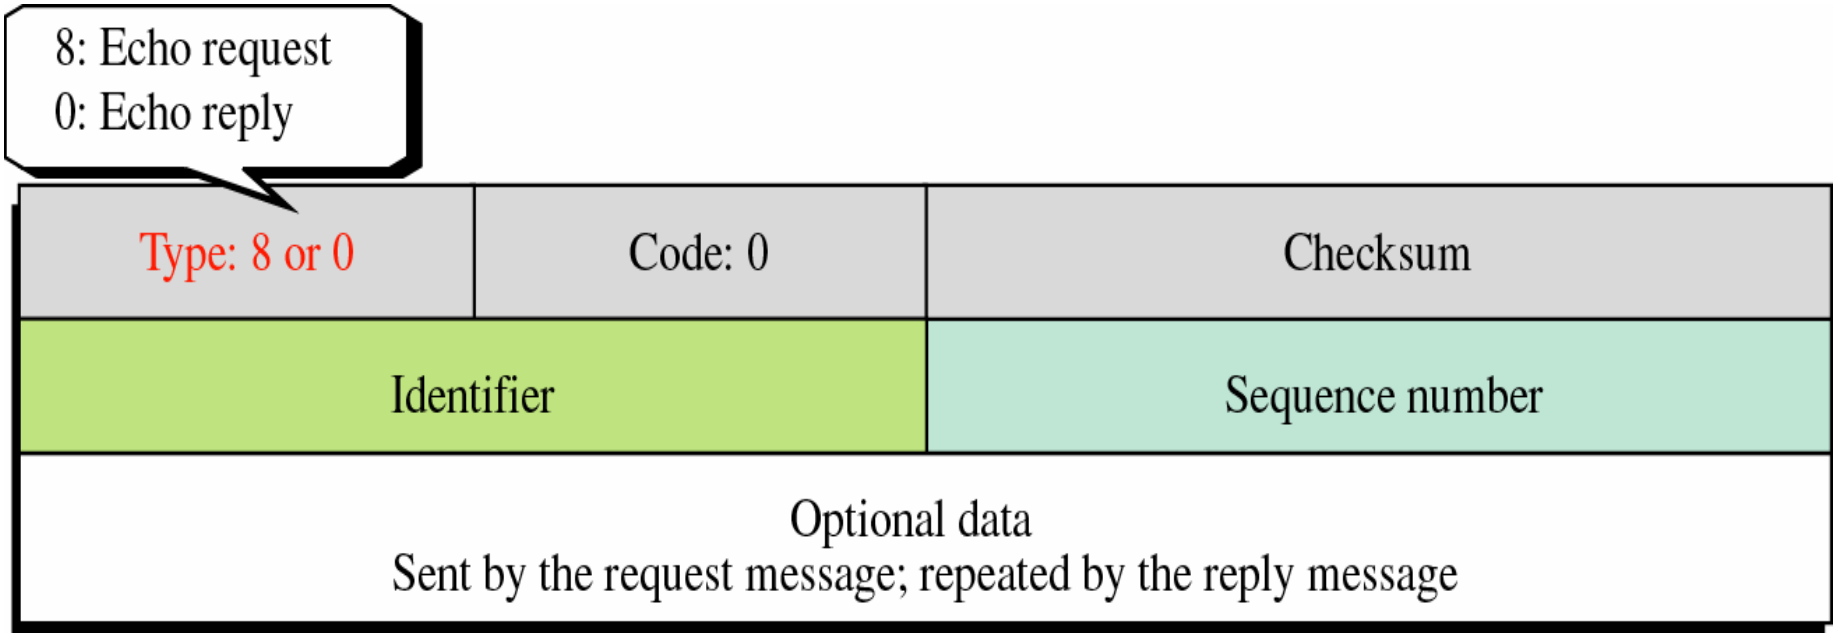

- Identifier and sequence number
  - Are not formally defined by the protocol
  - Can be used arbitrarily by the sender

Timestamp Request and Reply

- Two goals
  - Determine the *round-trip time* need for an IP datagram
  - *Synchronize* the clocks in two machines

Timestamp-Request and Timestamp-Reply Message Format

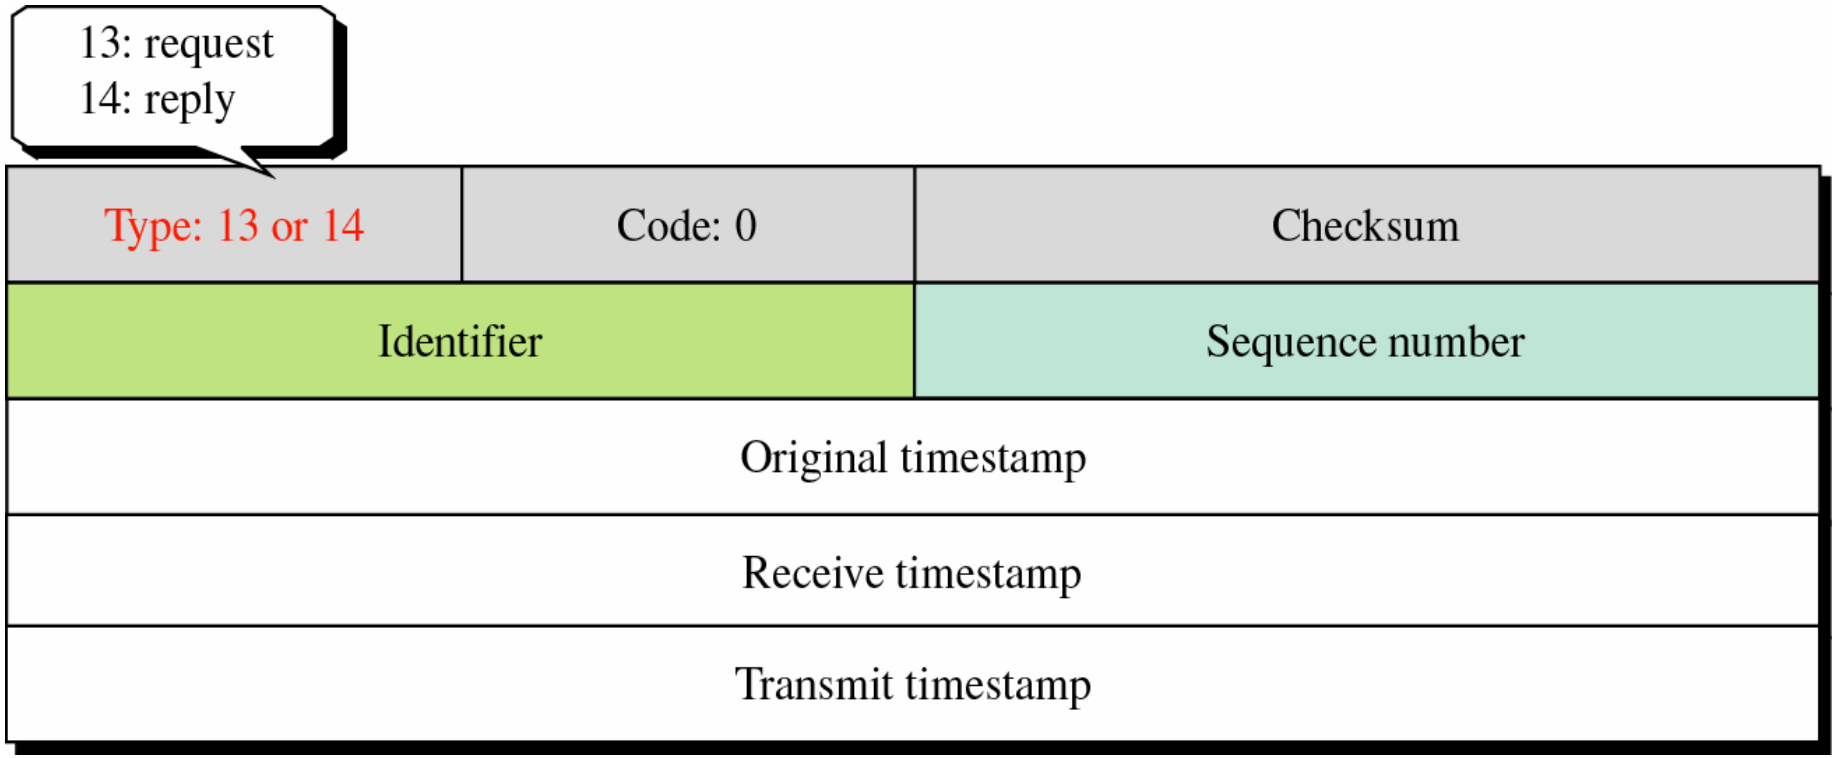

Timestamp

- Original timestamp
  - Filled by the *sender* at departure time
  - Appear in both request and reply messages
- Receive timestamp
  - Filled by the *receiver* at receiving time
  - Appear only at the reply message and fill zero in request message
- Transmit timestamp
  - Filled by the *receiver* when the reply message departs
  - Appear only at the reply message and fill zero in request message

One-Way and Round-Trip Time

- ***Sending time*** = value of receive timestamp - value of original timestamp
- ***Receiving time*** = time the packet returned - value of transmit timestamp
- ***Round-trip time*** = sending time + receiving time

One-Way and Round-Trip Time (Cont.)

- Note that
  - Sending time and receiving time are accurate only if the clocks in the source and destination are synchronized
  - The round-trip time is correct even if the two clocks are not synchronized
    - See the following next slides

Note

Timestamp-request and timestamp-reply messages can be used to calculate the round-trip time between a source and a destination machine even if their clocks are not synchronized.

Example

- Given
  - Value of original timestamp: 46
  - Value of receive timestamp: 59
  - Value of transmit timestamp: 60
  - Time the packet arrived: 67
- Derive
  - Sending time = 59 - 46 = 13 milliseconds
  - Receiving time = 67 - 60 = 7 milliseconds
  - Round-trip time = 13 + 7 = 20 milliseconds

Example (Cont.)

- To show that the round-trip time is independent of time difference
  - Following above example, assume that the time difference is 3
    - Receiving node's clock = 3 + sending node's clock
  - Sending time = (56+3) – 46 = 10 + 3
  - Receiving time = 67 – (57+3) = 10 – 3
  - RRT = (10 + 3) + (10-3) = 20

Synchronization

- *Timestamp request* and *timestamp reply* messages can also by used to synchronize the clocks in two machines

- Time difference = receive timestamp - (original timestamp field + one-way time duration)

Example

- Assume the one-way time duration can be obtained by dividing the round-trip time duration by two
  - Note that, this assumption may be wrong, depends on the network condition

- Thus, time difference = 59 - (46 + 10) = 3

Note

The timestamp-request and timestamp-reply messages can be used to synchronize two clocks in two machines if the exact one-way time duration is known.

Address-Mask Request and Reply

- The IP address of a host contains
  - A *network address*
  - *Subnet address* if subneted
  - *Host identifier*

- A host may know its full IP address, but does not know its network, subnetwork address, and its host identifier

Address-Mask Request and Reply (Cont.)

- Masking is needed for diskless station at start-up time
  - It first ask its IP address using the RARP protocol when it boots
  - Then, it use the address-mask request and reply to find out its mask

Address-Mask Request and Reply (Cont.)

- To obtain its mask, a host sends an address-mask-request message to a router
  - If it knows the router's address, send the request directly to the router
  - If it does not know, it broadcasts the message

Figure 9.16

Mask-Request and Mask-Reply Message Format

Router Solicitation and Advertisement

- A host needs to know the address of routers connected to its network
  - Send router-solicitation message by broadcasting or multicasting
  - The router receiving the message can then send the router-advertisement message
- A router may also periodically send router-advertisement message
  - Even if no hosts has solicited

Router Solicitation and Advertisement (Cont.)

- Note that, in a router-advertisement message
  - Contain not only its own presence
  - But also the presence of all routers on the network of which it is aware

Figure 9.17

Router Solicitation Message Format

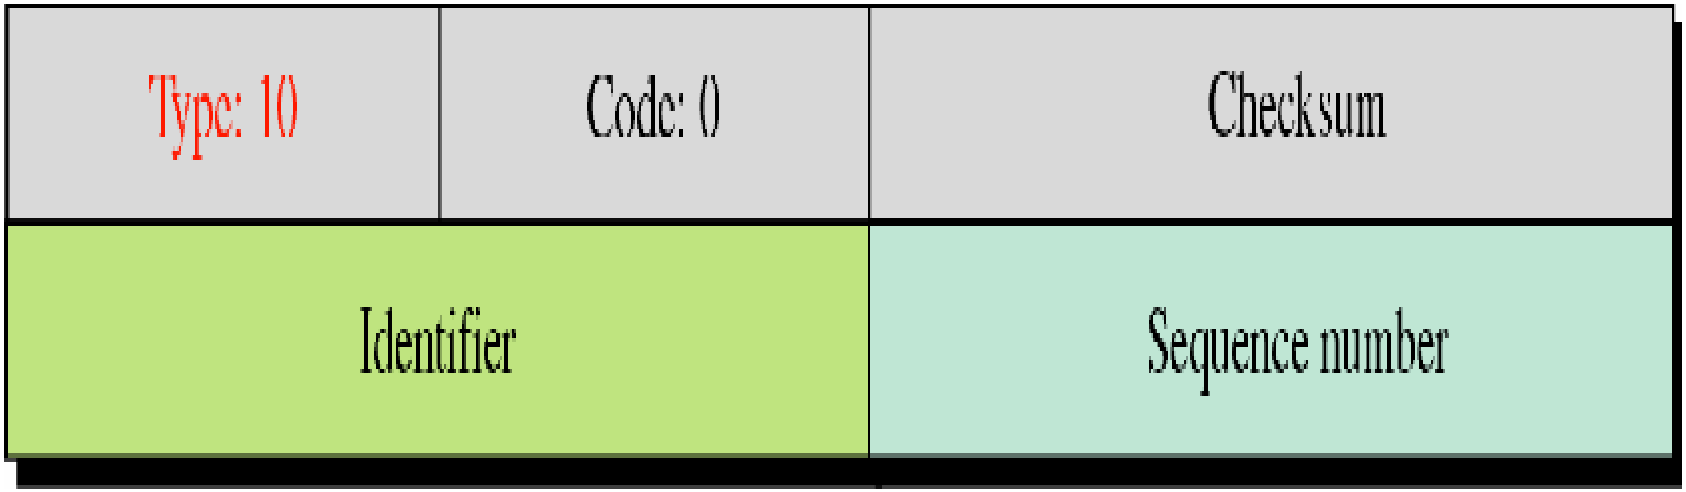

Figure 9-18

Router Advertisement Message Format

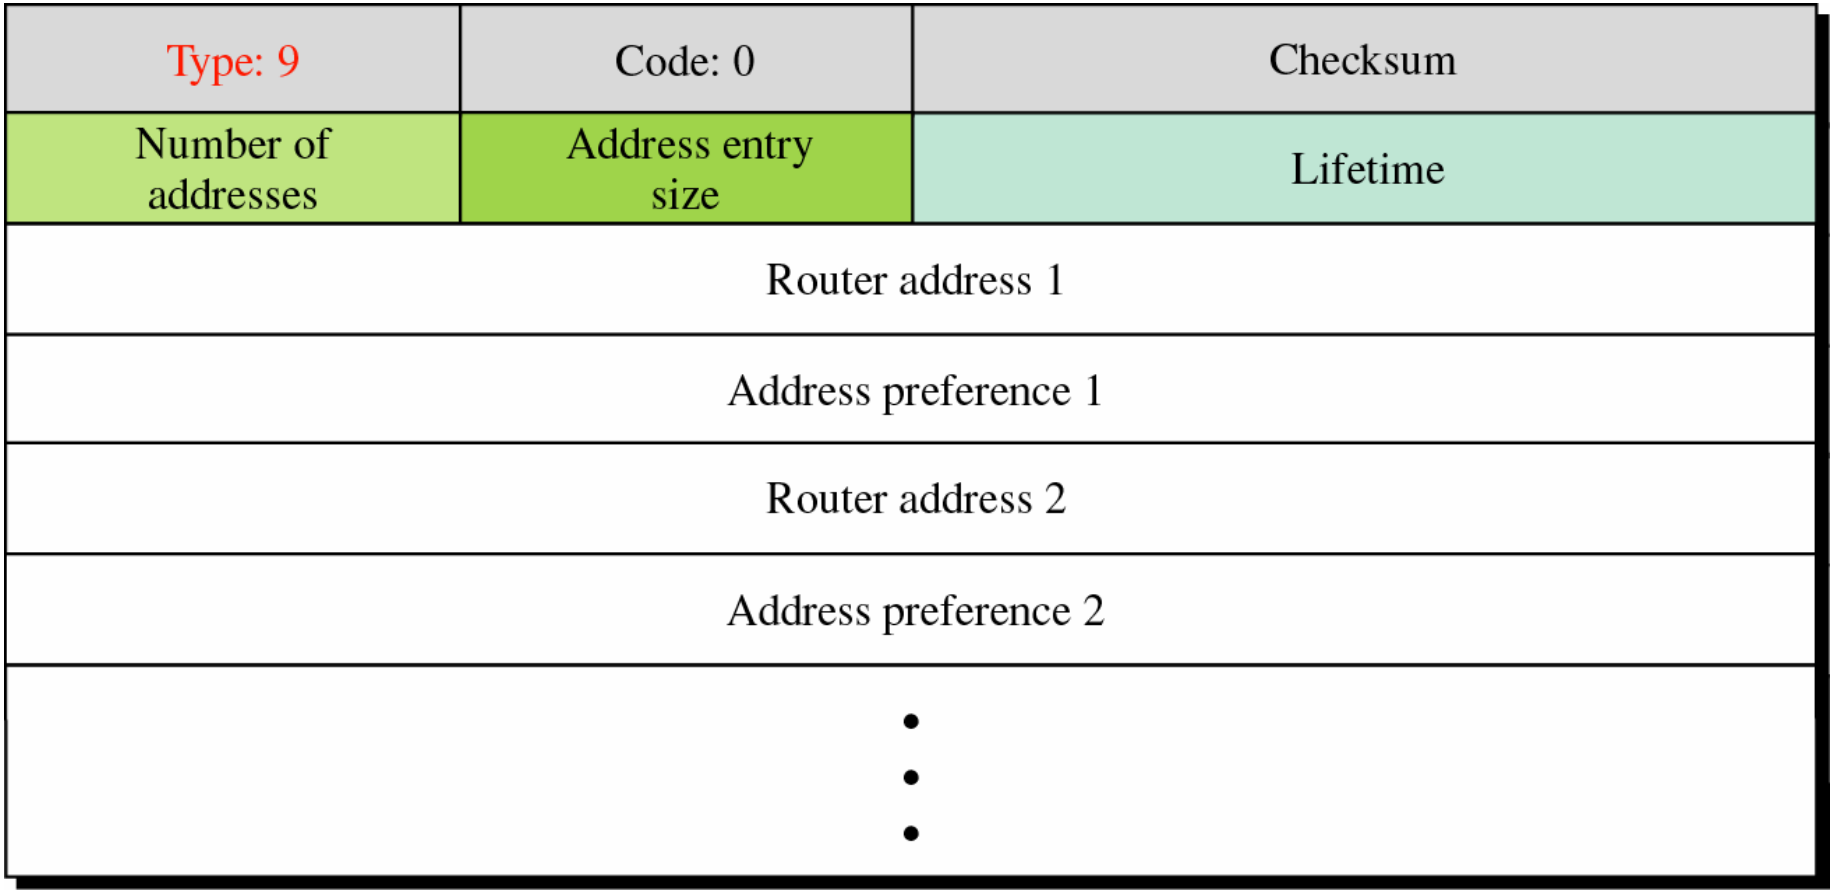

Router Advertisement Message Format

- Lifetime
  - The number of seconds that the entries are considered to be valid
- Address preference level
  - The ranking of the router and used to select a router as the *default router*
  - If zero: the router is considered as the default router
  - If 80000000_{16}, the router should never be selected as the default router

9.5

CHECKSUM

Checksum

- In ICMP, the checksum is calculated over the entire message
- Checksum calculation
  - The checksum field is set to zero
  - The sum of all the 16-bit words (header and data) is calculated
  - The sum is complemented to get the checksum
  - The checksum is stored in the checksum field

Checksum (Cont.)

- Checksum testing
  - The sum of all words (header and data) is calculated
  - The sum is complemented
  - If the result is 16 *0s*, the message is accepted
    - Otherwise, it is rejected

Figure 9-19

Example of Checksum Calculation

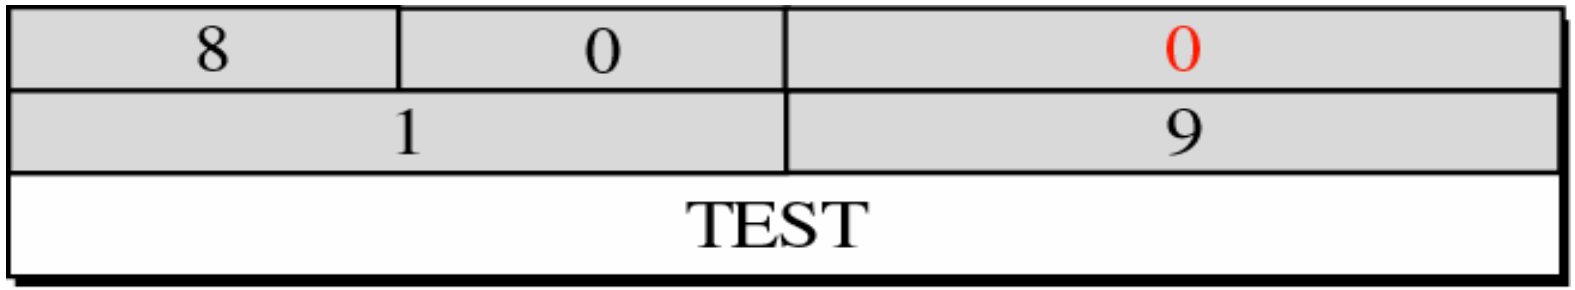

8 and 0	→	00001000 00000000
0	→	00000000 00000000
1	→	00000000 00000001
9	→	00000000 00001001
T & E	→	01010100 01000101
S & T	→	01010011 01010100
Sum	→	10101111 10100011
Checksum	→	01010000 01011100

9.6

ICMP PACKAGE

ICMP Package

- Input module: handle all received ICMP message
  - Invoked when an ICMP message is received
  - If the received packet is a request or solicitation
    - Create a reply or an advertisement and sends it out
- Output module: create request, solicitation, or error message requested by a higher level (TCP/UDP) or the IP protocol

Figure 9.20

ICMP Package