

The Network Layer

Antonio Carzaniga

Faculty of Informatics
Università della Svizzera italiana


November 24, 2017

- Basic network-layer architecture of a datagram network
- Introduction to forwarding
- Introduction to routing
- General architecture of a router
- Switching fabric and queuing
- Internet network-layer protocol
- The Internet protocol (IP)
- Fragmentation

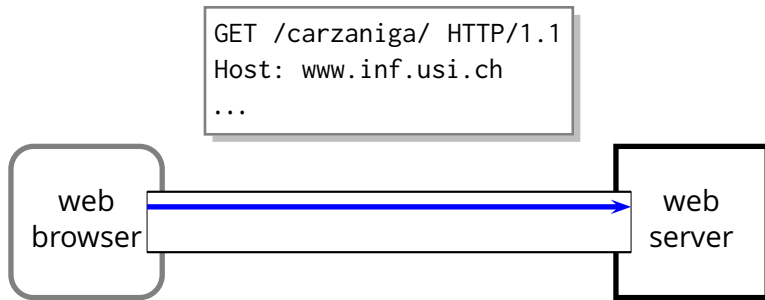
Application Level




Application Level



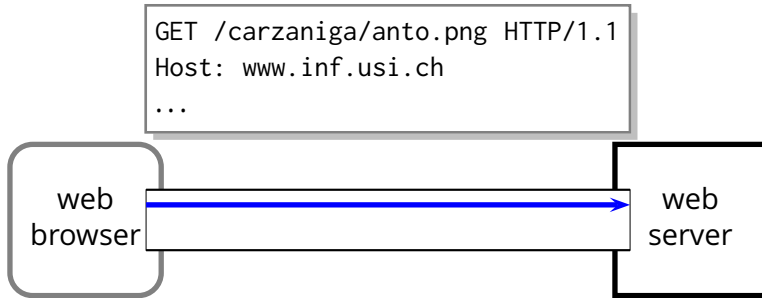
Application Level




Application Level




Application Level




Application Level




Transport Level



Transport Level




Transport Level






Network Layer




Network Layer













- Fundamental component of the network layer



- Fundamental component of the network layer
- *A node in a graph*



- Fundamental component of the network layer
- *A node in a graph*
- A finite set of input/output (physical) connections
 - ▶ a.k.a., ***interfaces*** or ***ports***

Focus: “Datagram” Networks

Focus: “Datagram” Networks

- *Packet-switched network*

Focus: “Datagram” Networks

- *Packet-switched network*

- ▶ information is transmitted in discrete units called ***datagrams***

Focus: “Datagram” Networks

- *Packet-switched network*

- ▶ information is transmitted in discrete units called ***datagrams***

- *Connectionless service*

Focus: “Datagram” Networks

■ *Packet-switched network*

- ▶ information is transmitted in discrete units called ***datagrams***

■ *Connectionless service*

- ▶ a datagram is a ***self-contained message***
- ▶ treated independently by the network
- ▶ no connection setup/tear-down phase

Focus: “Datagram” Networks

- *Packet-switched network*

- ▶ information is transmitted in discrete units called ***datagrams***

- *Connectionless service*

- ▶ a datagram is a ***self-contained message***
- ▶ treated independently by the network
- ▶ no connection setup/tear-down phase

- *“Best-effort” service*

Focus: “Datagram” Networks

■ *Packet-switched network*

- ▶ information is transmitted in discrete units called ***datagrams***

■ *Connectionless service*

- ▶ a datagram is a ***self-contained message***
- ▶ treated independently by the network
- ▶ no connection setup/tear-down phase

■ *“Best-effort” service*

- ▶ delivery guarantee: none

Focus: “Datagram” Networks

■ *Packet-switched network*

- ▶ information is transmitted in discrete units called ***datagrams***

■ *Connectionless service*

- ▶ a datagram is a ***self-contained message***
- ▶ treated independently by the network
- ▶ no connection setup/tear-down phase

■ *“Best-effort” service*

- ▶ delivery guarantee: none
- ▶ maximum latency guarantee: none

Focus: “Datagram” Networks

■ *Packet-switched network*

- ▶ information is transmitted in discrete units called ***datagrams***

■ *Connectionless service*

- ▶ a datagram is a ***self-contained message***
- ▶ treated independently by the network
- ▶ no connection setup/tear-down phase

■ *“Best-effort” service*

- ▶ delivery guarantee: none
- ▶ maximum latency guarantee: none
- ▶ bandwidth guarantee: none

Focus: “Datagram” Networks

■ *Packet-switched network*

- ▶ information is transmitted in discrete units called ***datagrams***

■ *Connectionless service*

- ▶ a datagram is a ***self-contained message***
- ▶ treated independently by the network
- ▶ no connection setup/tear-down phase

■ *“Best-effort” service*

- ▶ delivery guarantee: none
- ▶ maximum latency guarantee: none
- ▶ bandwidth guarantee: none
- ▶ in-order delivery guarantee: none

■ *Packet-switched network*

- ▶ information is transmitted in discrete units called ***datagrams***


■ *Connectionless service*

- ▶ a datagram is a ***self-contained message***
- ▶ treated independently by the network
- ▶ no connection setup/tear-down phase


■ *“Best-effort” service*

- ▶ delivery guarantee: none
- ▶ maximum latency guarantee: none
- ▶ bandwidth guarantee: none
- ▶ in-order delivery guarantee: none
- ▶ congestion indication: none


Datagram Network




Datagram Network




Datagram Network




Datagram Network




Datagram Network




Datagram Network



Datagram Network




Datagram Network




- Potentially *multiple paths* for the same source/destination

Datagram Network




- Potentially *multiple paths* for the same source/destination

Datagram Network




- Potentially *multiple paths* for the same source/destination

Datagram Network




- Potentially *multiple paths* for the same source/destination

Datagram Network




- Potentially *multiple paths* for the same source/destination

Datagram Network




- Potentially *multiple paths* for the same source/destination

Datagram Network




- Potentially *multiple paths* for the same source/destination

Datagram Network




- Potentially *multiple paths* for the same source/destination

Datagram Network




- Potentially *multiple paths* for the same source/destination
- Potentially *asymmetric paths*

Datagram Network




- Potentially *multiple paths* for the same source/destination
- Potentially *asymmetric paths*

Datagram Network




- Potentially *multiple paths* for the same source/destination
- Potentially *asymmetric paths*

Datagram Network




- Potentially *multiple paths* for the same source/destination
- Potentially *asymmetric paths*

Datagram Network




- Potentially *multiple paths* for the same source/destination
- Potentially *asymmetric paths*


Datagram Network




- Potentially *multiple paths* for the same source/destination
- Potentially *asymmetric paths*



- A sends a datagram to *B*

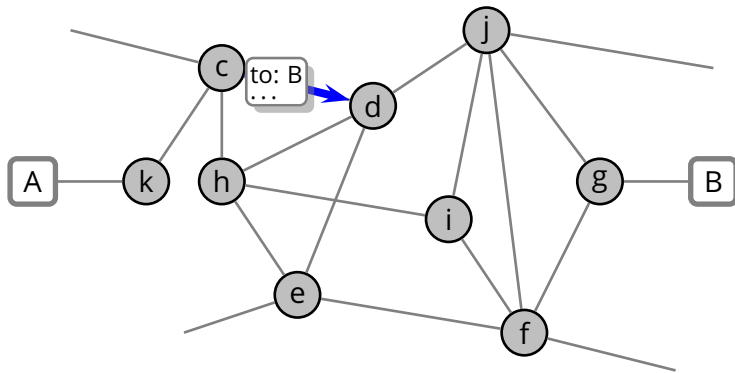


- A sends a datagram to *B*
- The datagram is **forwarded** towards *B*




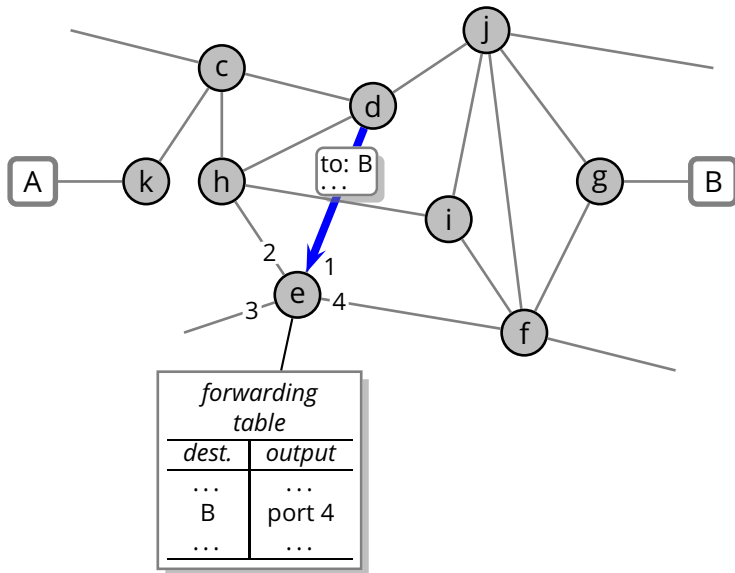
- A sends a datagram to *B*
- The datagram is **forwarded** towards *B*

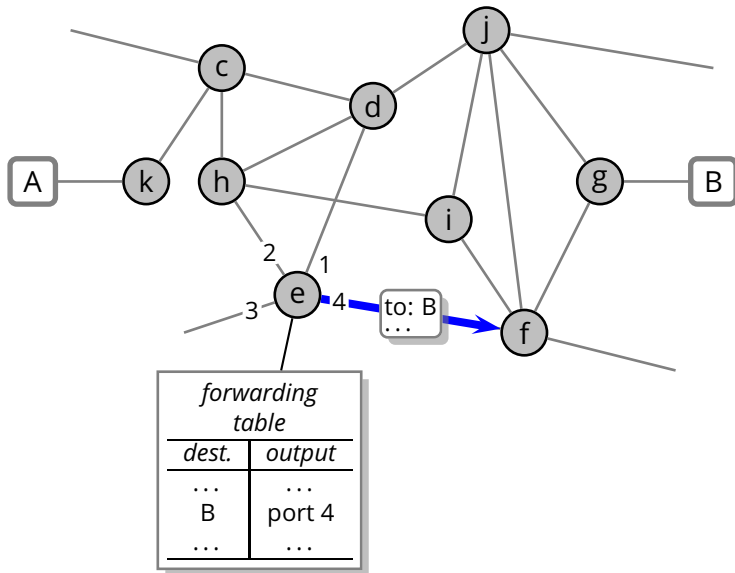
Forwarding

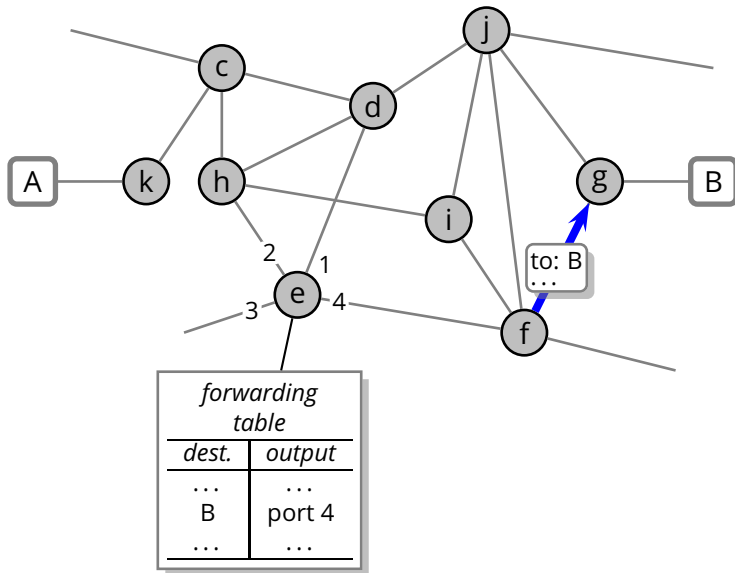



Forwarding











- *Input:* datagram destination

- *Input:* datagram destination
- *Output:* output port


- *Input*: datagram destination
- *Output*: output port
- Simple design: “forwarding table”


- *Input*: datagram destination
- *Output*: output port
- Simple design: “forwarding table”
- Issues

- *Input*: datagram destination
- *Output*: output port
- Simple design: “forwarding table”
- Issues
 - ▶ how big is the forwarding table?

- *Input*: datagram destination
- *Output*: output port
- Simple design: “forwarding table”
- Issues
 - ▶ how big is the forwarding table?
 - ▶ how fast does the router have to forward datagrams?

- *Input*: datagram destination
- *Output*: output port
- Simple design: “forwarding table”
- Issues
 - ▶ how big is the forwarding table?
 - ▶ how fast does the router have to forward datagrams?
 - ▶ how does the router build and maintain the forwarding table?







router *k*

A	2	...
B	1	


Router Functions



Router Functions




Router Functions

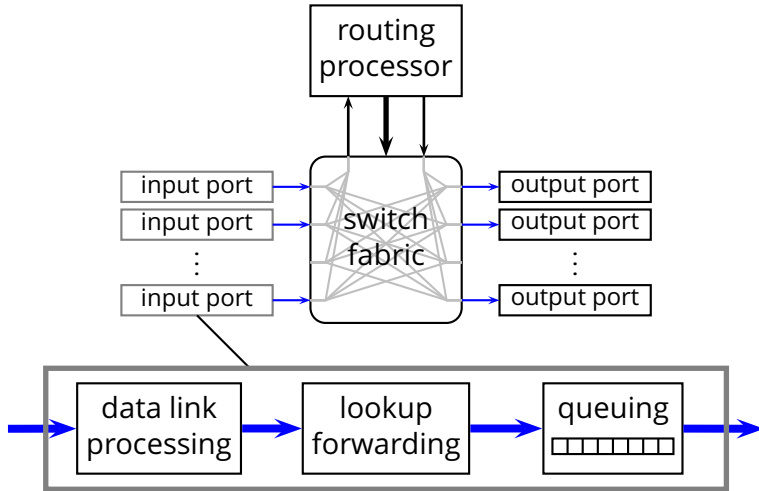


Anatomy of a Router


Anatomy of a Router




Anatomy of a Router



Anatomy of a Router



Anatomy of a Router



- Where does queuing occur?

- Where does queuing occur?
- Input ports
 - ▶ queuing may occur here if the switching fabric is slower than the aggregate speed of all the input lines. I.e., $R_S < nR_{in}$

- Where does queuing occur?
- Input ports
 - ▶ queuing may occur here if the switching fabric is slower than the aggregate speed of all the input lines. I.e., $R_S < nR_{in}$
- Output ports
 - ▶ queuing may occur here because of the limited throughput of the output link. I.e., $R_{out} < \min(R_S, nR_{in})$

- What happens when packets queue up in a router?

- What happens when packets queue up in a router?
- *Scheduling*: deciding which packets to process

- What happens when packets queue up in a router?
- *Scheduling*: deciding which packets to process
 - ▶ *first-come-first-served*

- What happens when packets queue up in a router?
- *Scheduling*: deciding which packets to process
 - ▶ *first-come-first-served*
 - ▶ *weighted fair queuing*: the router tries to be balance traffic evenly among the different end-to-end connections. Essential to implement *quality-of-service guarantees*

- What happens when packets queue up in a router?
- *Scheduling*: deciding which packets to process
 - ▶ *first-come-first-served*
 - ▶ *weighted fair queuing*: the router tries to balance traffic evenly among the different end-to-end connections. Essential to implement *quality-of-service guarantees*
- Deciding when to drop packets, and which packets to drop

- What happens when packets queue up in a router?
- *Scheduling*: deciding which packets to process
 - ▶ *first-come-first-served*
 - ▶ *weighted fair queuing*: the router tries to be balance traffic evenly among the different end-to-end connections. Essential to implement *quality-of-service guarantees*
- Deciding when to drop packets, and which packets to drop
 - ▶ *drop tail*: drop arriving packets when queues are full

- What happens when packets queue up in a router?
- *Scheduling*: deciding which packets to process
 - ▶ *first-come-first-served*
 - ▶ *weighted fair queuing*: the router tries to balance traffic evenly among the different end-to-end connections. Essential to implement *quality-of-service guarantees*
- Deciding when to drop packets, and which packets to drop
 - ▶ *drop tail*: drop arriving packets when queues are full
 - ▶ *active queue management*: a set of policies and algorithms to decide when and how to drop or mark packets in the attempt to prevent congestion

Internet Network Layer

- *Routing*: defining paths and compiling forwarding tables

- *Routing*: defining paths and compiling forwarding tables
 - ▶ RIP
 - ▶ OSPF
 - ▶ BGP

- *Routing*: defining paths and compiling forwarding tables
 - ▶ RIP
 - ▶ OSPF
 - ▶ BGP

- IP

- *Routing*: defining paths and compiling forwarding tables

- ▶ RIP
- ▶ OSPF
- ▶ BGP

- IP

- ▶ addressing
- ▶ datagram format
- ▶ fragmentation and packet handling

- *Routing*: defining paths and compiling forwarding tables

- ▶ RIP
- ▶ OSPF
- ▶ BGP

- IP

- ▶ addressing
- ▶ datagram format
- ▶ fragmentation and packet handling

- ICMP

■ *Routing*: defining paths and compiling forwarding tables

- ▶ RIP
- ▶ OSPF
- ▶ BGP


■ IP

- ▶ addressing
- ▶ datagram format
- ▶ fragmentation and packet handling


■ ICMP

- ▶ error reporting
- ▶ signaling

IPv4 Datagram Format




IPv4 Datagram Format



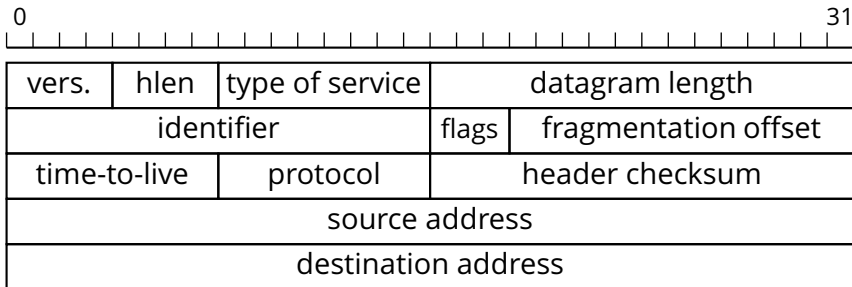
IPv4 Datagram Format



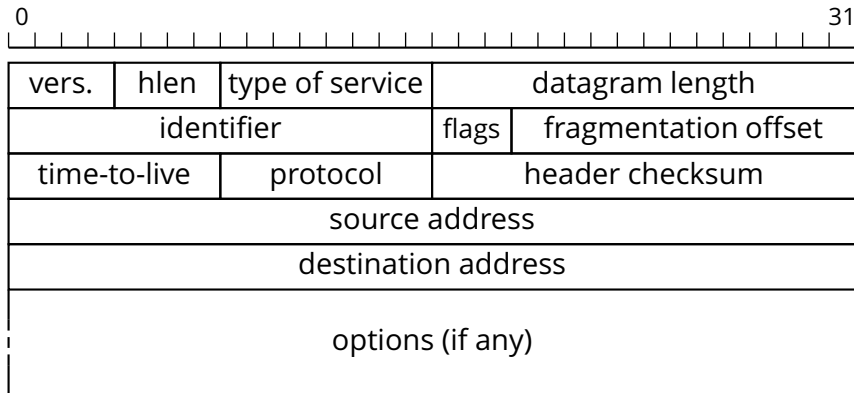
IPv4 Datagram Format



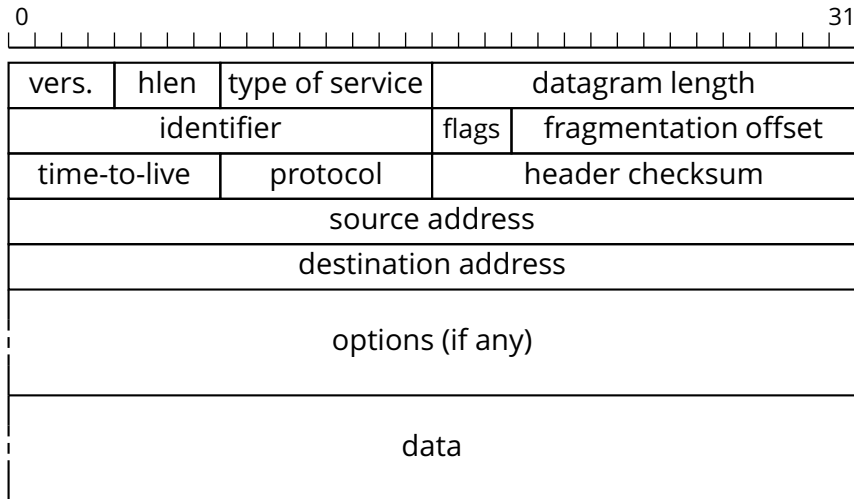
IPv4 Datagram Format



IPv4 Datagram Format




IPv4 Datagram Format




Fragmentation


Fragmentation




Fragmentation




Fragmentation



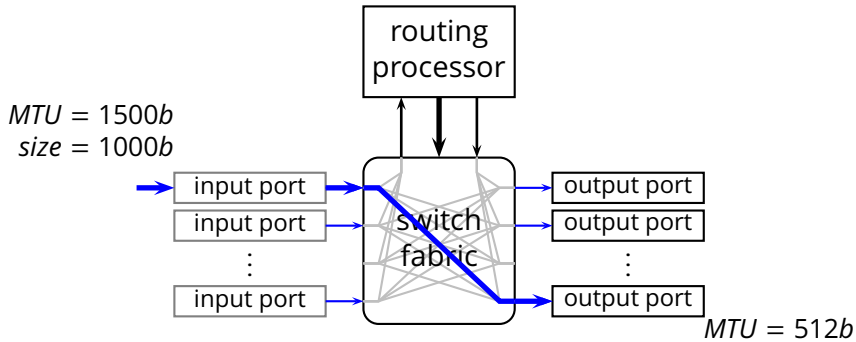
Fragmentation




Fragmentation




Fragmentation






- How does the router handle cases where the size of an input datagram exceeds the maximum transmission unit (MTU) of the output link?




- How does the router handle cases where the size of an input datagram exceeds the maximum transmission unit (MTU) of the output link?
- The datagram is **fragmented**

Fragmentation




input datagram


Fragmentation




Fragmentation




Fragmentation



- The *destination* reassembles fragmented datagrams



- The *destination* reassembles fragmented datagrams
 - ▶ push complexity out of the network
 - ▶ a datagram may have to be fragmented further along the path



- The *destination* reassembles fragmented datagrams
 - ▶ push complexity out of the network
 - ▶ a datagram may have to be fragmented further along the path
- Requirements
 - ▶ destination must recognize two fragments of the same original datagram
 - ▶ destination must see if and when all the fragments have been received
 - ▶ intermediate routers must be able to fragment a datagram to whatever level necessary

Fragmentation

- Initial (non-fragmented) datagram format (*datasize* = 1000)

- Initial (non-fragmented) datagram format (*datasize* = 1000)
 - ▶ sender host assigns a 16-bit *identifier* to the datagram (e.g., 789)

- Initial (non-fragmented) datagram format (*datasize* = 1000)
 - ▶ sender host assigns a 16-bit *identifier* to the datagram (e.g., 789)
 - ▶ the *fragment offset* is set to 0, indicating that this packet contains data starting at position 0 of the original datagram
 - ▶ *fragment offset* is actually the offset in units of 8 bytes (remember it's only 13 bits...)

- Initial (non-fragmented) datagram format ($datasize = 1000$)
 - ▶ sender host assigns a 16-bit *identifier* to the datagram (e.g., 789)
 - ▶ the *fragment offset* is set to 0, indicating that this packet contains data starting at position 0 of the original datagram
 - ▶ *fragment offset* is actually the offset in units of 8 bytes (remember it's only 13 bits...)
 - ▶ the “*more fragments*” flag is set to 0, indicating that no (more) fragments have been sent

- Initial (non-fragmented) datagram format ($datasize = 1000$)
 - ▶ sender host assigns a 16-bit *identifier* to the datagram (e.g., 789)
 - ▶ the *fragment offset* is set to 0, indicating that this packet contains data starting at position 0 of the original datagram
 - ▶ *fragment offset* is actually the offset in units of 8 bytes (remember it's only 13 bits...)
 - ▶ the “*more fragments*” flag is set to 0, indicating that no (more) fragments have been sent

<i>identifier</i>	<i>fragment offset</i>	<i>more fragments</i>	<i>header length</i>	<i>total length</i>
789	0	0	20	1020

Fragmentation

- Fragmentation to an MTU of 512

- Fragmentation to an MTU of 512
 - ▶ sender must split the datagram into 3 fragments:

- Fragmentation to an MTU of 512

- ▶ sender must split the datagram into 3 fragments:

<i>identifier</i>	<i>fragment offset</i>	<i>more fragments</i>	<i>header length</i>	<i>total length</i>
789	0	1	20	508

■ Fragmentation to an MTU of 512

- ▶ sender must split the datagram into 3 fragments:

<i>identifier</i>	<i>fragment offset</i>	<i>more fragments</i>	<i>header length</i>	<i>total length</i>
789	0	1	20	508

<i>identifier</i>	<i>fragment offset</i>	<i>more fragments</i>	<i>header length</i>	<i>total length</i>
789	61	1	20	508

■ Fragmentation to an MTU of 512

- ▶ sender must split the datagram into 3 fragments:

<i>identifier</i>	<i>fragment offset</i>	<i>more fragments</i>	<i>header length</i>	<i>total length</i>
789	0	1	20	508

<i>identifier</i>	<i>fragment offset</i>	<i>more fragments</i>	<i>header length</i>	<i>total length</i>
789	61	1	20	508

<i>identifier</i>	<i>fragment offset</i>	<i>more fragments</i>	<i>header length</i>	<i>total length</i>
789	122	0	20	44