Network Layer

Katerina Lionta
klionta@csd.uoc.gr

Data Plane

Data Plane

- How a router forwards the packets that arrives in its incoming links to outcoming links

IPv4 Address (1/2)

- 32-bits, decimal representation **a**.**b**.**c**.**d**/**x**
- 4 Octets, **a**=**b**=**c**=**d**=8 bits
- Each network has a range of IPs that can assign to its devices
  - Called prefix of the network, /**x**
  - The prefix indicates how many bits are fixed
  - To find the number of the available IPs given a prefix
    - $2^{(32-x)}-2$
  - e.g Given the prefix 10.0.8.0/24, how many IPs are available?
    - $2^{(32-24)}-2=2^{8}-2=254$ IPs

IPv4 Address (1/2)

- There are IPs that we cannot assign to devices
  - The **network IP** which is the first IP of the network
    - identifies the network
  - The **broadcast IP** which is the last IP of the network
    - sends the messages to every device in the network
  - e.g Given the prefix 10.0.8.0/24
    - Network IP: 10.0.8.0
    - Broadcast IP: 10.0.8.255
  - **That's the reason that we subtract 2 when we compute the number of available IPs**

Forwarding Table (1/2)

- Each router has a forwarding table
- The table is completed by routing algorithms of the control plane
    - OSPF, BGP
- The entries of the table contain the prefixes with the corresponding output interface

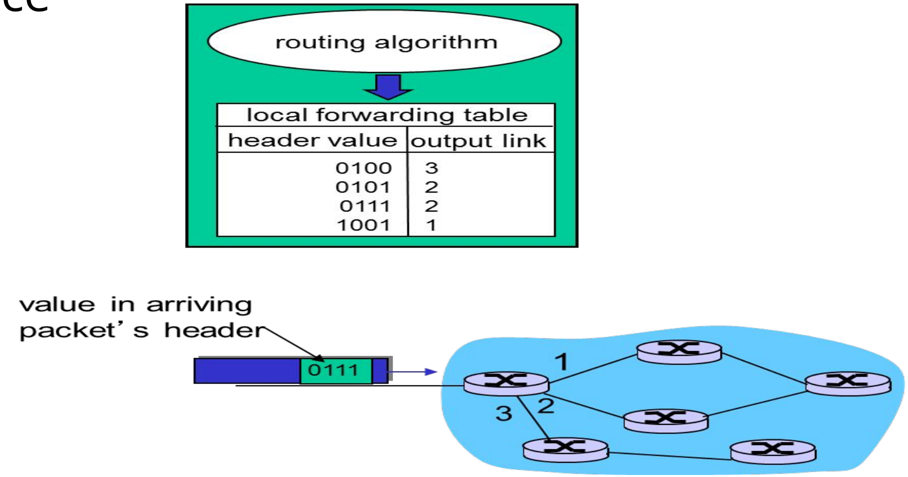

Forwarding Table (2/2)

- The router extracts the **destination IP** of each incoming packet
- Then try to match this IP with on of the prefixes in its forwarding table
  - Ẃhen the router finds the matching subnet sends the packet to the corresponding output interface
- The forwarding table follows the rule of **longest prefix matching**:
  - if the destination IP matches with multiple subnets the router selects the longer

Forwarding Table Example (1/3)

- How the router forwards the following IPs?
  - 11001000 00010111 00010001 10011010
  - 11001000 00010111 00011000 10101010

Subnet	Out Interface
11001000 00010111 00010*** ********	0
11001000 00010111 00011000 ********	1
11001000 00010111 00011*** ********	2
Otherwise	3

Forwarding Table Example (2/3)

- How the router forwards the following IPs?
  - **11001000 00010111 00010**001 10011010 **-> interface 0**

Subnet	Out Interface
11001000 00010111 00010*** ********	0
11001000 00010111 00011000 ********	1
11001000 00010111 00011*** ********	2
Otherwise	3

Forwarding Table Example (3/3)

- How the router forwards the following IPs?
    - **11001000 00010111 00011000** 10101010 **-> interface 1**

Subnet	**Out Interface**
11001000 00010111 00010*** ********	0
11001000 00010111 00011000 ********	1
11001000 00010111 00011*** ********	2
Otherwise	3

Subnetting

- Many organizations divide their networks to smaller to:
  - facilitate the management
  - enhance the security by reducing the broadcast domains

Subnetting Example (1/2)

- A network with an IP Prefix 223.1.17.0/24
  - Separate the network in two subnets
- Solution:

Decimal IP	223	1	17	0
Binary IP	11011111	00000001	0010001	00000000

- /24= $2^{(32-24)}$-2= 254 IPs
- two subnets->254/2=127 IPs each subnet

Subnetting Example (2/2)

- subnet 1 , first 127 IPs:

from	11011111	00000001	0010001	00000001=1
to	11011111	00000001	0010001	01111110=126

Prefix 1: 223.1.17.0/25, network IP=223.1.17.0, broadcast IP=223.1.17.127

- subnet 2, next 127 IPs:

from	11011111	00000001	0010001	10000001=129
to	11011111	00000001	0010001	11111110=254

Prefix 2: 223.1.17.128/25,network IP=223.1.17.128, broadcast IP =223.1.17.255

Network Address Translation (NAT) (1/2)

- The border router of the network separates the the internal network (LAN) from the Internet (WAN)
- **NAT router**: is a border router that assigns the same source IP to the packets that are sent from internal network to the Internet
- Benefits:
  - Privacy
  - IP saving

Network Address Translation (NAT) (2/2)

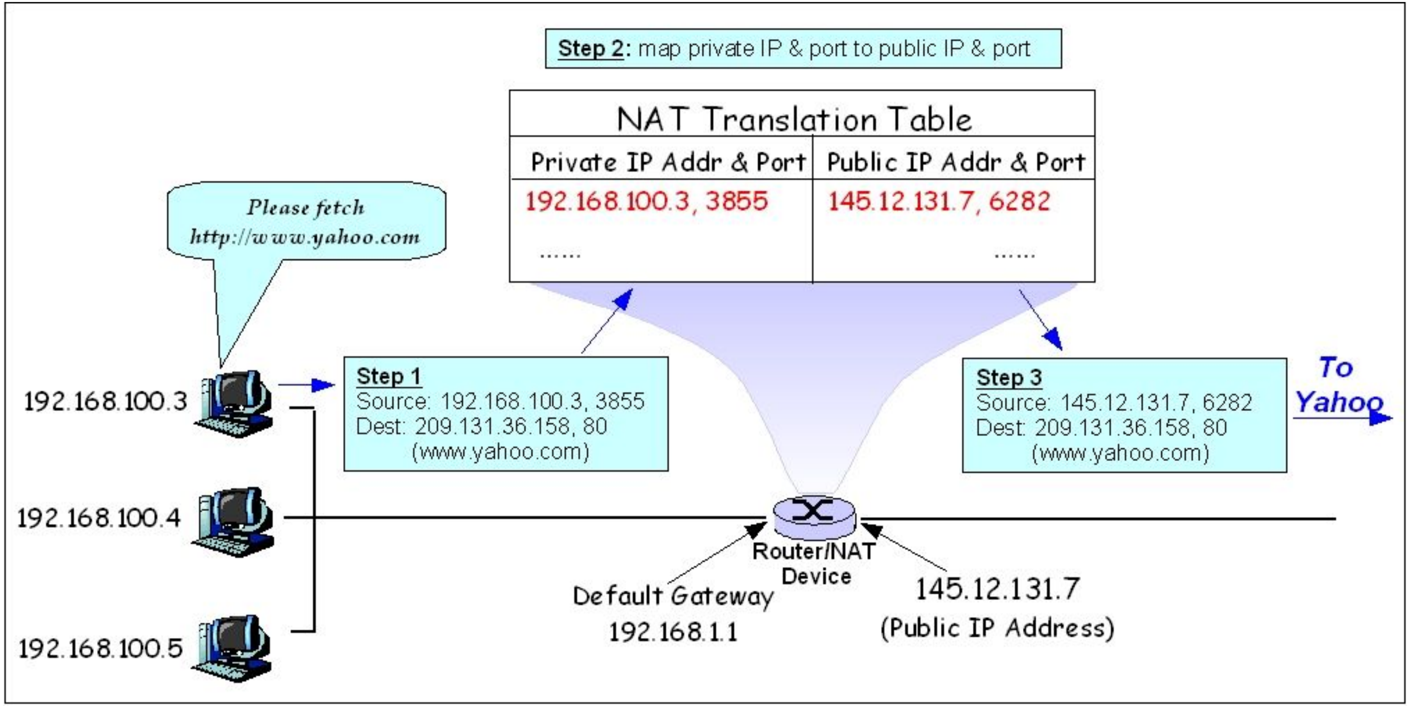

NAT Example

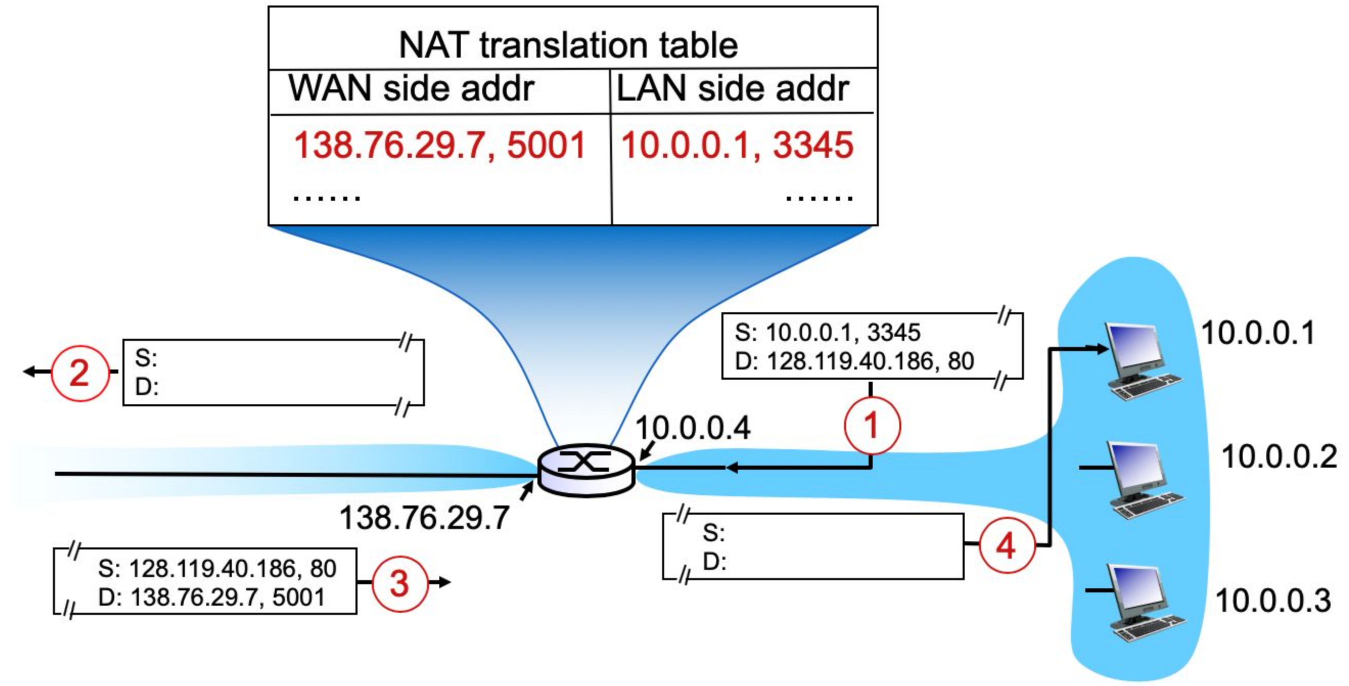
NAT translation table
WAN side addr
LAN side addr
138.76.29.7, 5001
10.0.0.1, 3345
……
……
S: 10.0.0.1, 3345
D: 128.119.40.186, 80
1
10.0.0.4
10.0.0.1
10.0.0.2
10.0.0.3
2
S:
D:
138.76.29.7
S: 128.119.40.186, 80
D: 138.76.29.7, 5001
3
S:
D:
4

NAT Example Solution

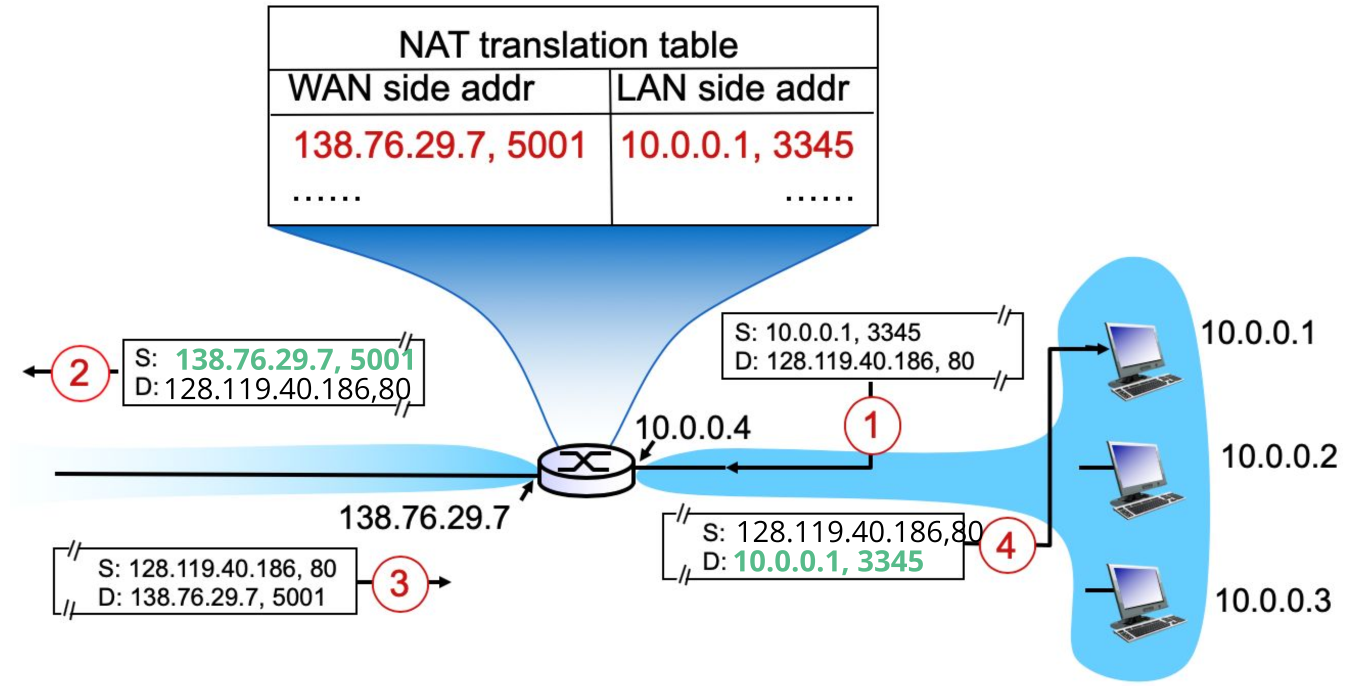

Control Plane

Control Plane

- How a packet is routed among the routers

Routing Algorithms: Dijkstra (1/2)

- Centralized routing algorithm
- Each router in the network knows the entire topology
  - Link state algorithm
- Compute the minimum distance from one router to every other router in the network
  - find the minimum cost paths from one node to any other node

Routing Algorithms: Dijkstra (2/2)

- Find the minimum distance from A to eve router
- D: Distance
- p: previous node

Nodes	D(A), p(A)	D(B), p(B)	D(C), P(C)	D(D), p(D)
A	-	5, A	∞	8, A
AB	-	-	7, B	8, A
ABC	-	-	-	8, A
ABCD	-	-	-	-

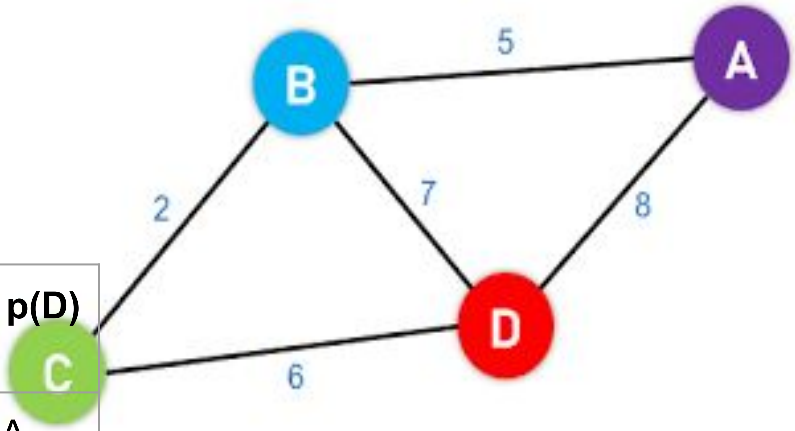

Routing Algorithms: Distance Vector (1/2)

- Decentralized Algorithm
- The routers know only their neighbors
- Repetitive algorithm
- The information of the cost of each link is distributed in the network via messages
- Based on Bellman-Ford equation:
  - $d_x(y)=min_v\{c(x,v)+d_v(y)\}$

Routing Algorithms: Distance Vector (2/2)

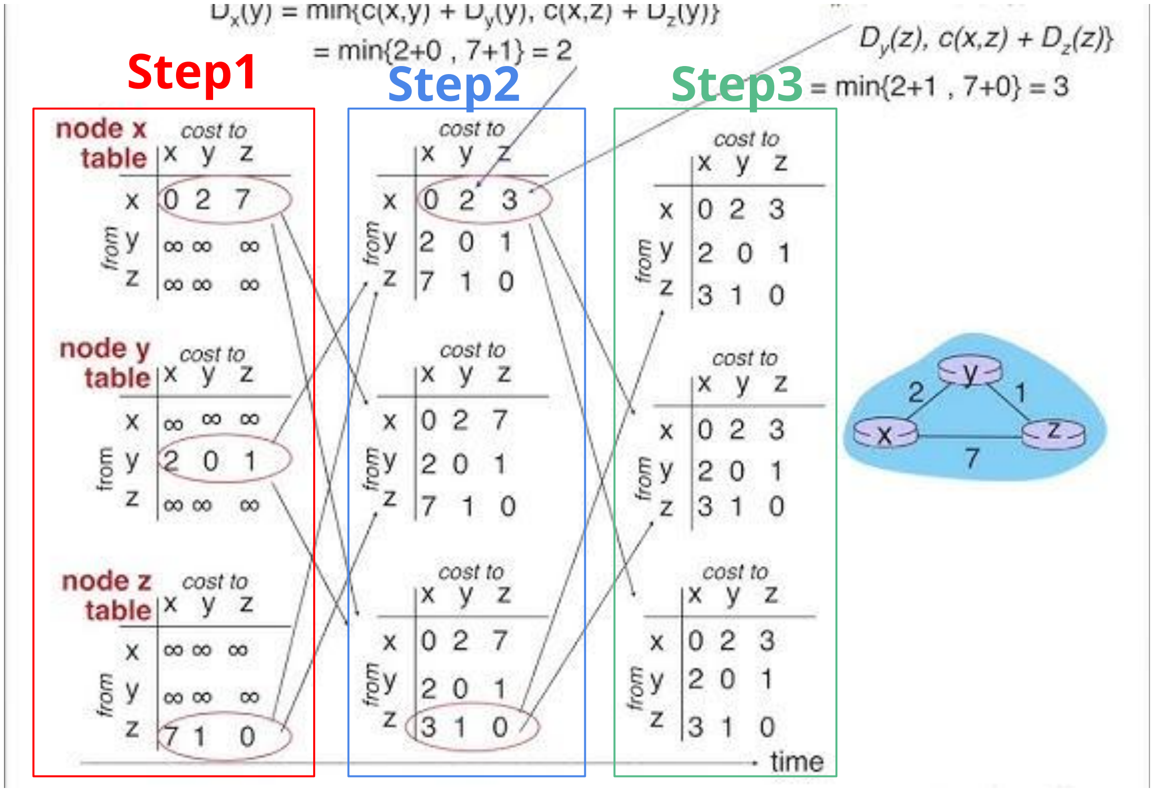

Autonomous Systems (ASes)

- A group of routers which operate under the same management
- Each AS has a unique number identifier
  - Autonomous System Number (ASN)
- The routers of each AS shares a common prefix
- Each organization (Facebook, Google, Amazon etc) has one or more ASes in different locations

AS Routing Algorithms

- Intra-autonomous system routing
  - OSPF
  - Based on Dijkstra routing algorithm
  - Determines how the traffic flows **inside the AS**
- Inter-autonomous system routing
  - BGP
  - Based on distance vector algorithm
  - Determines how the traffic flows among **different ASes**

Border Gateway Protocol (BGP) (1/2)

- Each router in the network can:
    - Determines the best paths to different networks
    - Receives information for paths from its neighbors for different networks
- A protocol that determines the best paths to reach other ASes
- **Internal router**: a router that is connected with routers of the AS that belongs to
- **Gateway router**: a router that is connected with routers from other ASes
- The routers that run BGP are connected with TCP connections
    - **eBGP**: A TCP connection between two routers from different ASes
    - **iBGP**: A TCP connection between two routers from the same AS

BOrder Gateway Protocol (BGP) (2/2)

- BGP attributes:
    - **AS-PATH**: The ASes that the AS must traverse to reach the desired prefix
    - **NEXT-HOP**: The IP of the border router in the next AS in the Path
- The AS-PATH that each router in AS1 knows to reach the router D in AS3 is:
    - AS3 AS1
- The NEXT-HOP is the IP of the router I (in green)

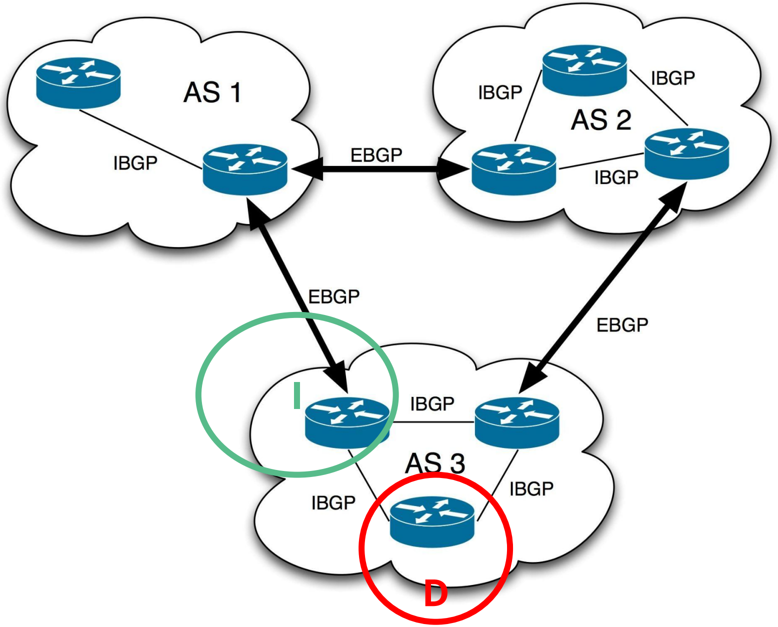

Thank You !