



**NTNU – Trondheim**  
Norwegian University of  
Science and Technology

# **Network Layer – Data Plane**

TTM4200 – Computer Networks

# Hint for the RATs

- **Binary Subnet Mask Structure - Example /24**

A /24 subnet mask has the first 24 bits set to 1 and the last 8 bits set to 0, shown as 11111111.11111111.11111111.00000000.

- **Decimal and Binary Relationship**

Each octet of eight 1s corresponds to decimal 255, resulting in the dotted decimal notation 255.255.255.0 for a /24 mask.

- **Network and Host Segmentation**

Routers use the subnet mask's binary form to separate network and host portions via an AND operation with IP addresses.

# individual Readiness Assurance Test

- Find the right answer **alone**.
- Choose the answer that fits **best**.
- **One** handwritten A4 page as helping material.



20:00

# **team** Readiness Assurance Test

- Find the right answer **in teams**.
- Choose the answer that fits **best**.
- **One** handwritten A4 page as helping material.



20:00

# Time for a break

*we restart at 11:20*

Join at [menti.com](https://www.menti.com) | use code .....




## Instructions

Go to

**[www.menti.com](https://www.menti.com)**

Enter the code


**3812 3495**



Or use QR code


# Networks and subnets pt. 1

- **Identify each sub-network (subnet)**, i.e., mark which devices, interfaces, and links belong to the same subnet. Assume that the switches are Layer-2, so they do not speak IP and they do not have IP addresses.




# Networks and subnets pt. 2

- Within each subnet, **how many interfaces** are there? **How many IP addresses are needed** for each subnet, if each interface only gets one IP address?




# Networks and subnets pt. 3

- Assume that none of the routers is a NAT gateway, and that NAT is **not** used. Based on your answer to the previous question, what are the **smallest possible IPv4 subnet sizes** for each subnet?



# Networks and subnets pt. 4

- If classfull addressing was still the only option (i.e. no CIDR / variable length subnet masking, only A, B, C networks), what subnet sizes would be required? Why is this none-desirable?



# Enough for today

*See you on Thursday 10:15*

# Welcome

*We start at 10:15*

Join at [menti.com](https://www.menti.com) | use code .....




## Instructions

Go to

**[www.menti.com](https://www.menti.com)**

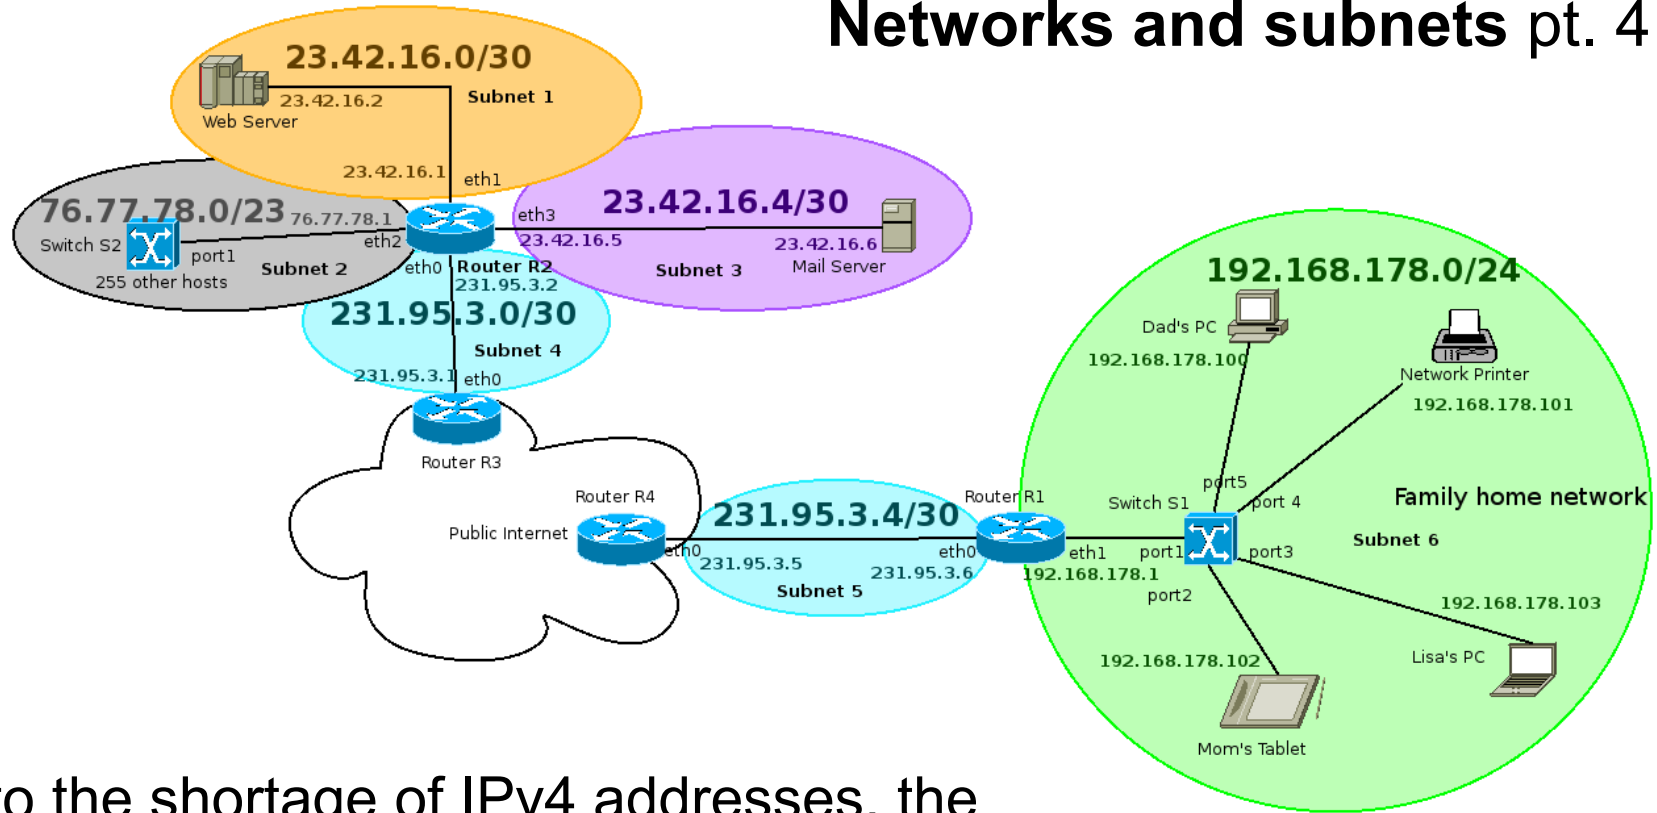
Enter the code

**3812 3495**



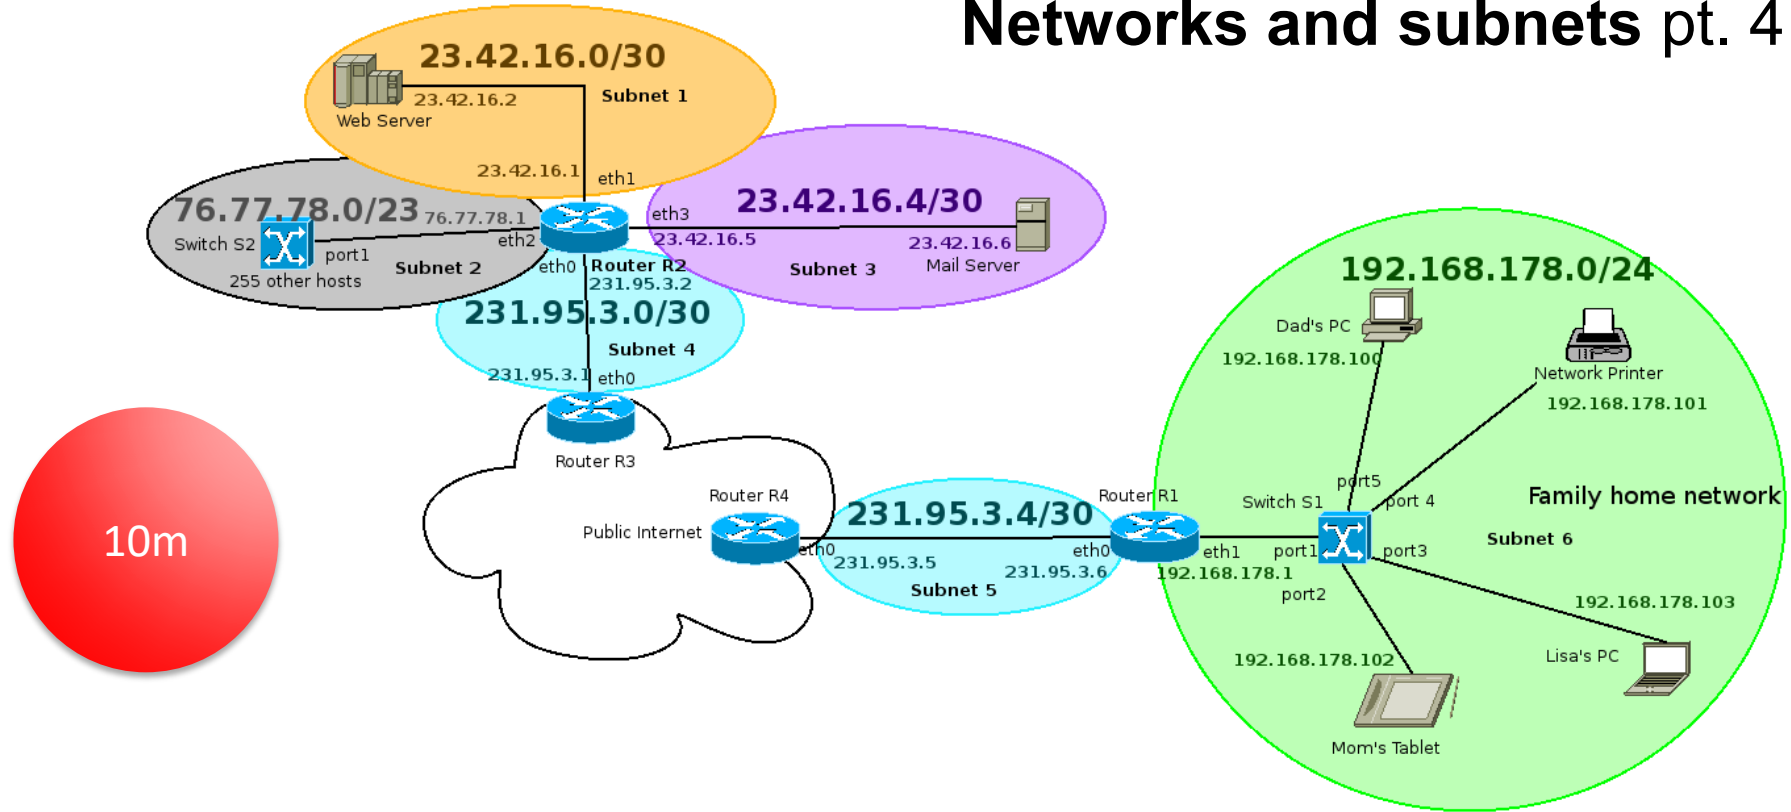
Or use QR code

# Networks and subnets pt. 4



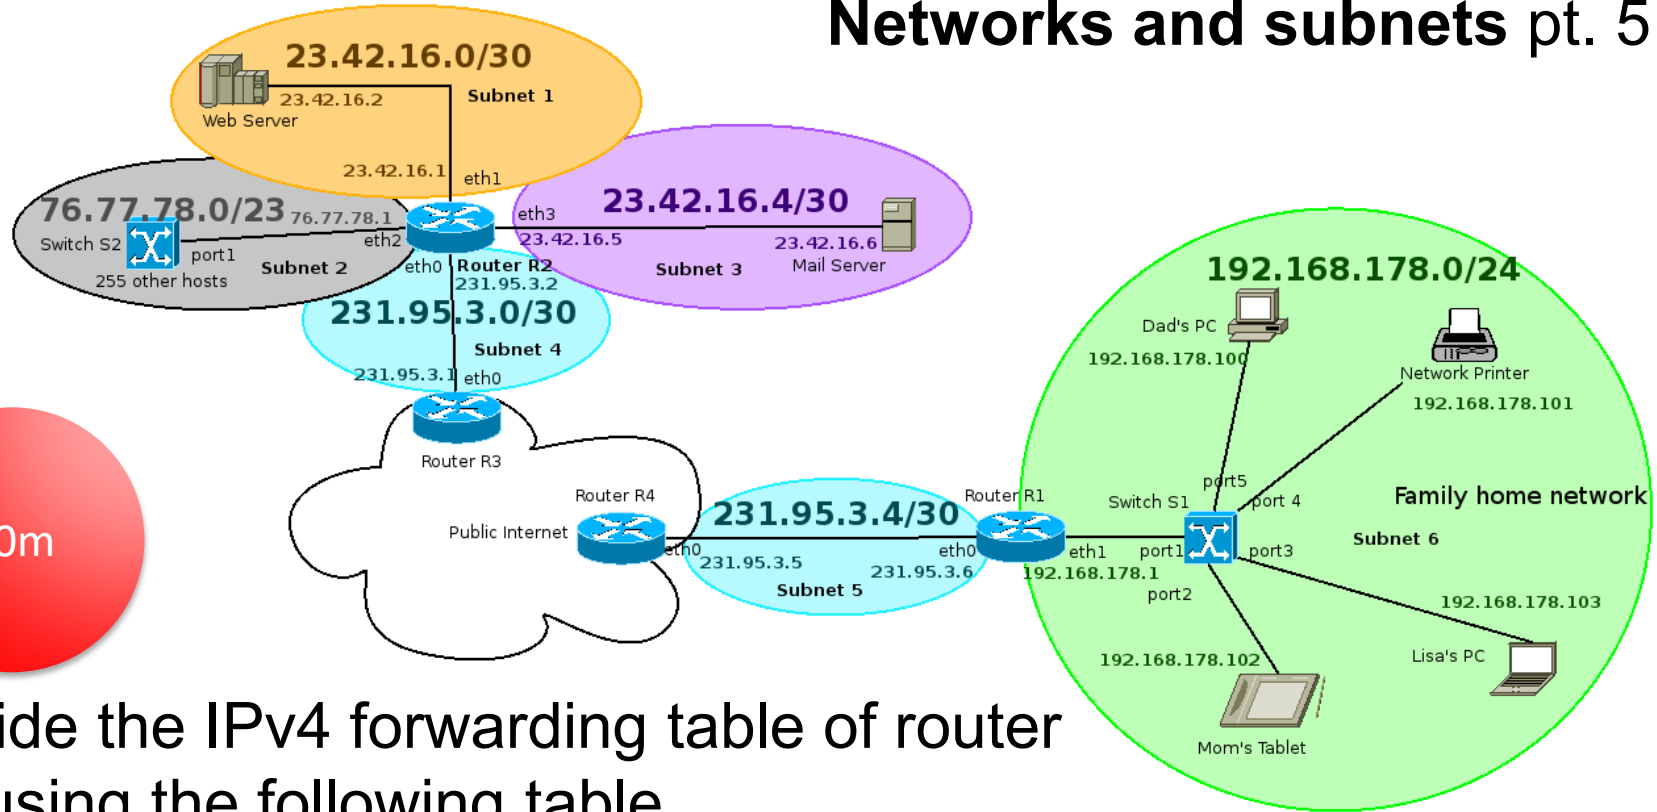

- Due to the shortage of IPv4 addresses, the family home network did not get public IPv4 addresses but is using a /24 from the private IPv4 address space. The router R1 is running a **NAT gateway**.

# Networks and subnets pt. 4



- What is the **broadcast IPv4 address** in each of the subnets shown in the figure?

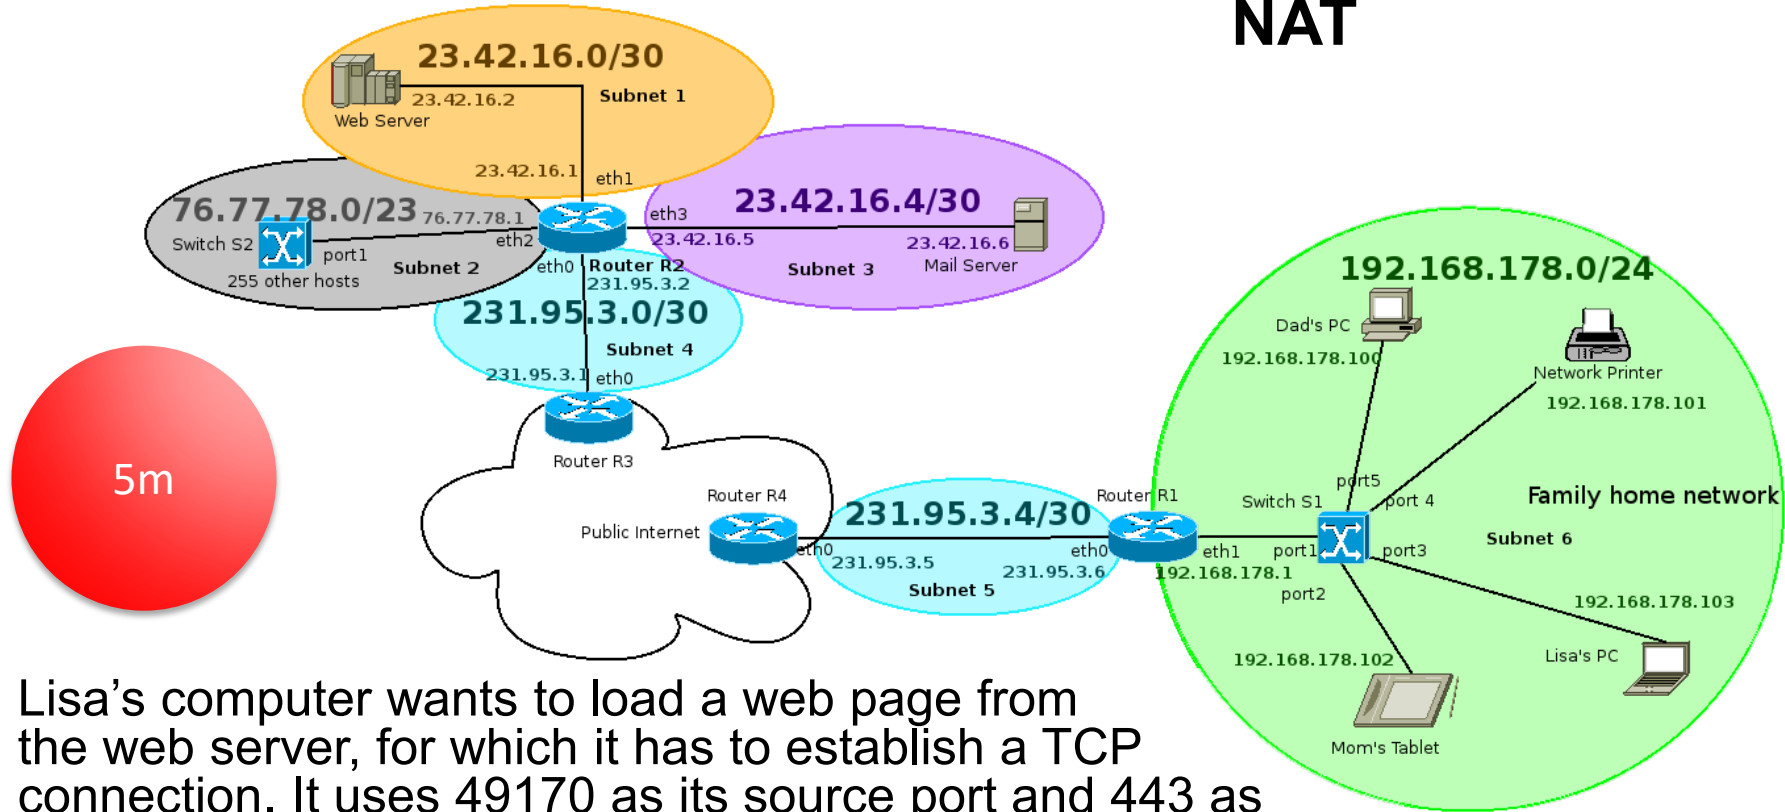
# Networks and subnets pt. 5



- Provide the IPv4 forwarding table of router R2, using the following table.

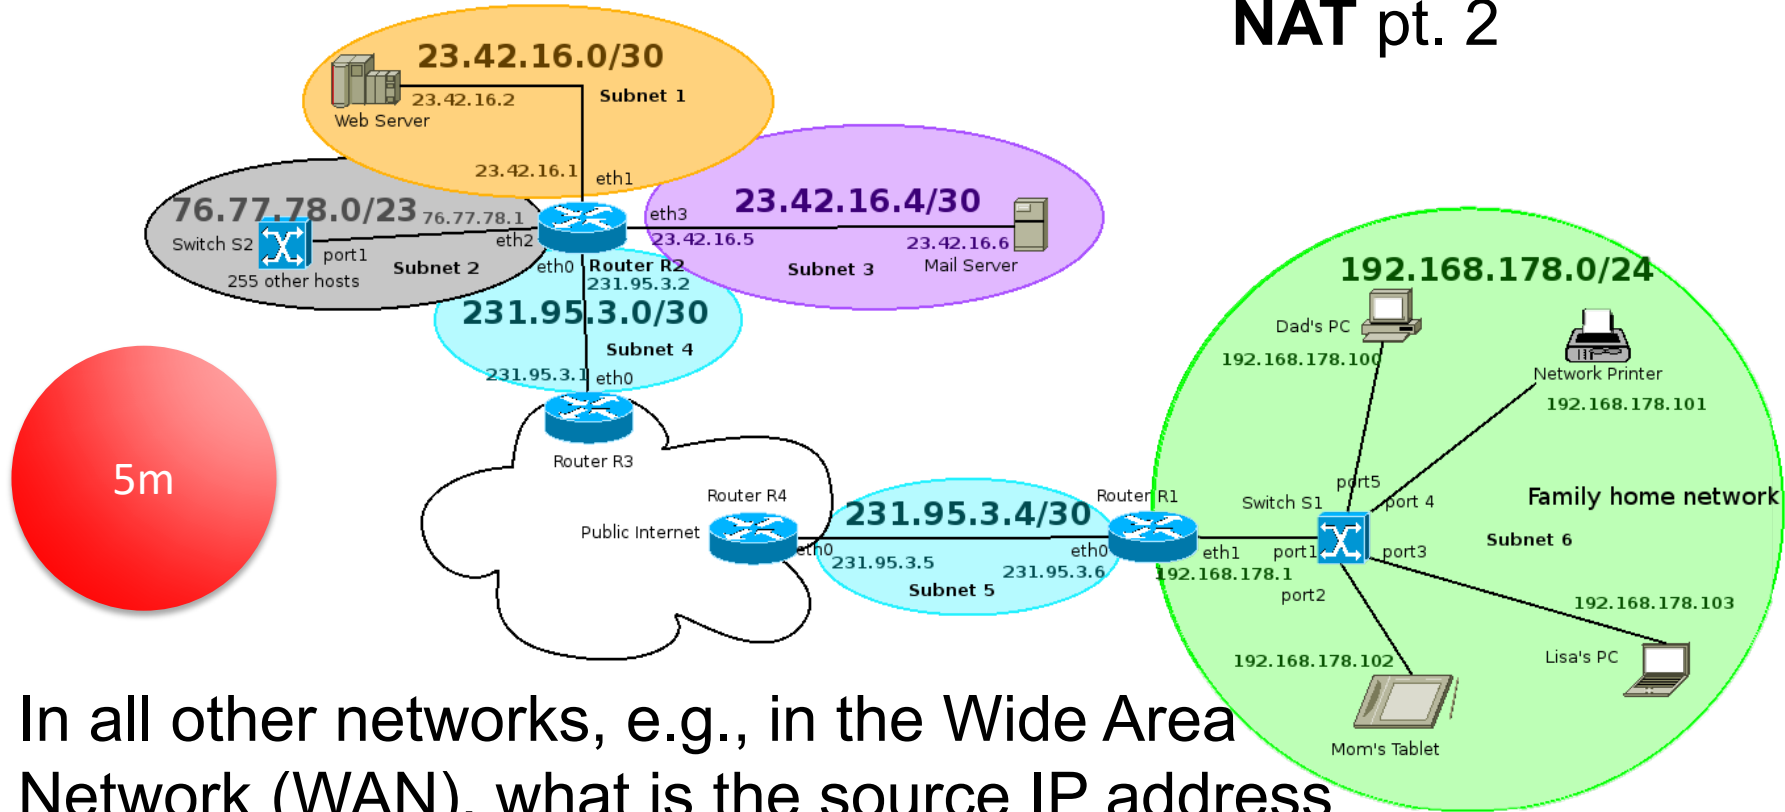
prefix/mask	next hop IP addr	iface

## NAT



- Lisa's computer wants to load a web page from the web server, for which it has to establish a TCP connection. It uses 49170 as its source port and 443 as destination port, and IPv4. Inside the home network (LAN), **what is the source IP address and the destination IP address** of the packet containing the initial TCP SYN?

## NAT pt. 2

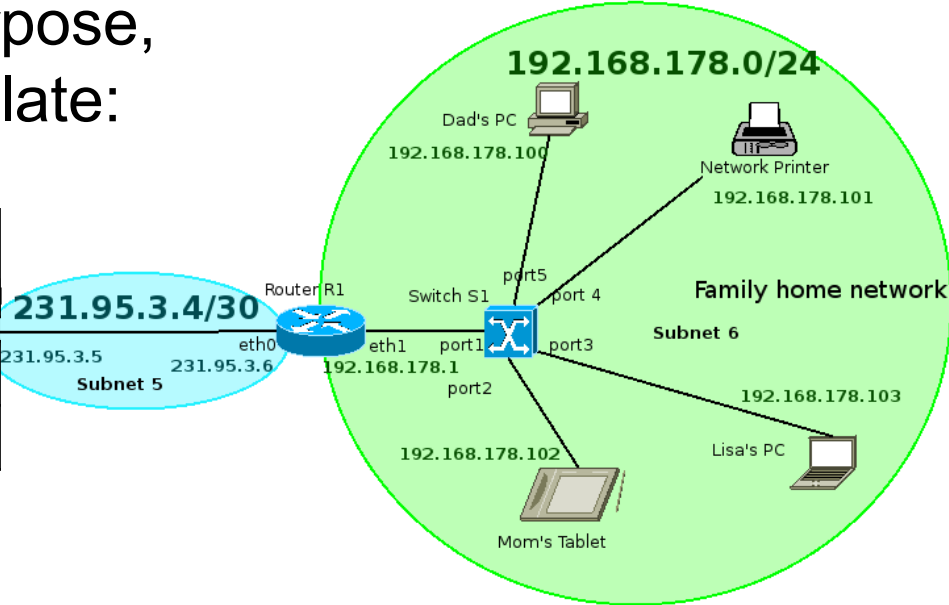


- In all other networks, e.g., in the Wide Area Network (WAN), what is the source IP address and the destination IP address of the packet containing the initial TCP SYN?

# NAT pt. 3

- Show the NAT table of R1 after forwarding the TCP SYN packet of that connection. For this purpose, use the following table template:

source IP	source port	outside src port	dest. IP	dest. port



# Time for a break

*we restart at 11:20*

Join at [menti.com](https://www.menti.com) | use code .....




## Instructions

Go to

**[www.menti.com](https://www.menti.com)**

Enter the code

**3812 3495**



Or use QR code

# Network layer: "data plane" roadmap

- Network layer: overview
  - data plane
  - control plane
- What's inside a router
  - input ports, switching, output ports
  - buffer management, scheduling
- IP: the Internet Protocol
  - datagram format
  - addressing
  - network address translation
  - IPv6
- Generalized Forwarding, SDN
  - match+action
  - OpenFlow: match+action in action
- Middleboxes



# Network layer: "data plane" roadmap


- Network layer: overview
  - data plane
  - control plane
- What's inside a router
  - input ports, switching, output ports
  - buffer management, scheduling
- IP: the Internet Protocol
  - datagram format
  - addressing
  - network address translation
  - IPv6
- Generalized Forwarding, SDN
  - Match+action
  - OpenFlow: match+action in action
- Middleboxes



# Generalized forwarding: match plus action


*Review:* each router contains a **forwarding table** (aka: **flow table**)

- **"match plus action"** abstraction: match bits in arriving packet, take action
  - *destination-based forwarding*: forward based on dest. IP address
  - *generalized forwarding*:
    - many header fields can determine action
    - many action possible: drop/copy/modify/log packet



# Flow table abstraction

- **flow**: defined by header field values (in link-, network-, transport-layer fields)
- **generalized forwarding**: simple packet-handling rules
  - **match**: pattern values in packet header fields
  - **actions**: for matched packet: drop, forward, modify, matched packet or send matched packet to controller
  - **priority**: disambiguate overlapping patterns
  - **counters**: #bytes and #packets



# OpenFlow abstraction

- **match+action**: abstraction unifies different kinds of devices

## Router

- *match*: longest destination IP prefix
- *action*: forward out a link

## Switch

- *match*: destination MAC address
- *action*: forward or flood

## Firewall

- *match*: IP addresses and TCP/UDP port numbers
- *action*: permit or deny

## NAT

- *match*: IP address and port
- *action*: rewrite address and port

# Generalized forwarding: summary

- “match plus action” abstraction: match bits in arriving packet header(s) in any layers, take action
  - matching over many fields (link-, network-, transport-layer)
  - local actions: drop, forward, modify, or send matched packet to controller
  - “program” *network-wide* behaviors
- simple form of “network programmability”
  - programmable, per-packet “processing”
  - *historical roots*: active networking
  - *today*: more generalized programming: P4 (see [p4.org](http://p4.org)).

# Network layer: "data plane" roadmap

- Network layer: overview
- What's inside a router
- IP: the Internet Protocol
- Generalized Forwarding
- **Middleboxes**
  - middlebox functions
  - evolution, architectural principles of the Internet



# Middleboxes


Middlebox (RFC 3234)

“any intermediary box performing functions apart from normal, standard functions of an IP router on the data path between a source host and destination host”

# Middleboxes everywhere!

**NAT:** home,  
cellular,  
institutional

**Application-specific:**  
service  
providers,  
institutional,  
CDN



**Firewalls, IDS:** corporate,  
institutional, service  
providers, ISPs

**Load balancers:**  
corporate, service  
provider, data center,  
mobile nets

**Caches:** service  
provider, mobile, CDNs


# Middleboxes

- initially: proprietary (closed) hardware solutions
- move towards “whitebox” hardware implementing open API
  - move away from proprietary hardware solutions
  - programmable local actions via match+action
  - move towards innovation/differentiation in software
- SDN: (logically) centralized control and configuration management often in private/public cloud
- network functions virtualization (NFV): programmable services over white box networking, computation, storage

# The IP hourglass

## Internet's "thin waist":

- *one* network layer protocol: IP
- *must* be implemented by every (billions) of Internet-connected devices




*many* protocols in physical, link, transport, and application layers

# The IP hourglass, at middle age

Internet's middle age  
"love handles"?

- middleboxes, —————  
operating inside the  
network



# Architectural Principles of the Internet

## RFC 1958

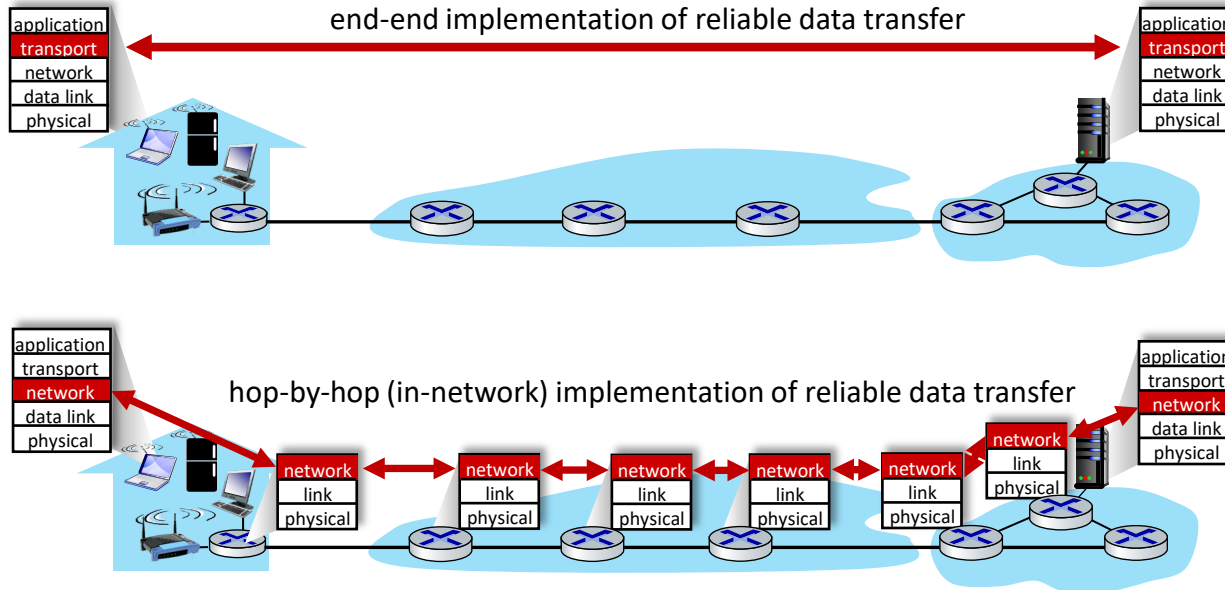
“Many members of the Internet community would argue that there is no architecture, but only a tradition, which was not written down for the first 25 years (or at least not by the IAB). However, in very general terms, the community believes that **the goal is connectivity, the tool is the Internet Protocol, and the intelligence is end to end rather than hidden in the network.**”

Three cornerstone beliefs:

- simple connectivity
- IP protocol: that narrow waist
- intelligence, complexity at network edge

# The end-end argument

- some network functionality (e.g., reliable data transfer, congestion) can be implemented **in network**, or at **network edge**




# The end-end argument

- some network functionality (e.g., reliable data transfer, congestion) can be implemented **in network**, or at **network edge**

“The function in question can completely and correctly be implemented only with the knowledge and help of the application standing at the end points of the communication system. Therefore, providing that questioned function as a feature of the communication system itself is not possible. (Sometimes an incomplete version of the function provided by the communication system may be useful as a performance enhancement.)


We call this line of reasoning against low-level function implementation the “end-to-end argument.”

# Where's the intelligence?




## 20<sup>th</sup> century phone net:

- intelligence/computing at network switches



## Internet (pre-2005)

- intelligence, computing at edge



## Internet (post-2005)

- programmable network devices
- intelligence, computing, massive application-level infrastructure at edge

# Chapter 4: done!

- Network layer: overview
- What's inside a router
- IP: the Internet Protocol
- *Generalized Forwarding, SDN*
- Middleboxes



*Question:* how are forwarding tables (destination-based forwarding) or flow tables (generalized forwarding) computed?

*Answer:* by the control plane (next chapter)

# Readings for next week

## Book pages

- 408–425 (full Sec. 5.1 & 5.2)
- 453–455 (5.6 ICMP)

## Key concepts

### Routing Algorithms

- Graph,  $G=(N,E)$
- Nodes, Edges, Neighbours, Cost and Paths
- Link-state vs. Distance-vector

# Network layer control plane: our goals

- understand principles behind network control plane:
  - traditional routing algorithms
  - ~~SDN controllers~~
  - ~~network management, configuration~~
- ~~instantiation, implementation in the Internet:~~
  - ~~OSPF, BGP~~
  - ~~OpenFlow, ODL and ONOS controllers~~
  - Internet Control Message Protocol: ICMP
  - ~~SNMP, YANG/NETCONF~~

# Network layer: "control plane" roadmap

- **introduction**
- routing protocols
  - link state
  - distance vector
- intra-ISP routing: OSPF
- routing among ISPs: BGP
- SDN control plane
- Internet Control Message Protocol



- network management, configuration
  - SNMP
  - NETCONF/YANG

# Network-layer functions

- **forwarding**: move packets from router's input to appropriate router output
- **routing**: determine route taken by packets from source to destination

*data plane*


*control plane*

## Two approaches to structuring network control plane:

- per-router control (traditional)
- logically centralized control (software defined networking)

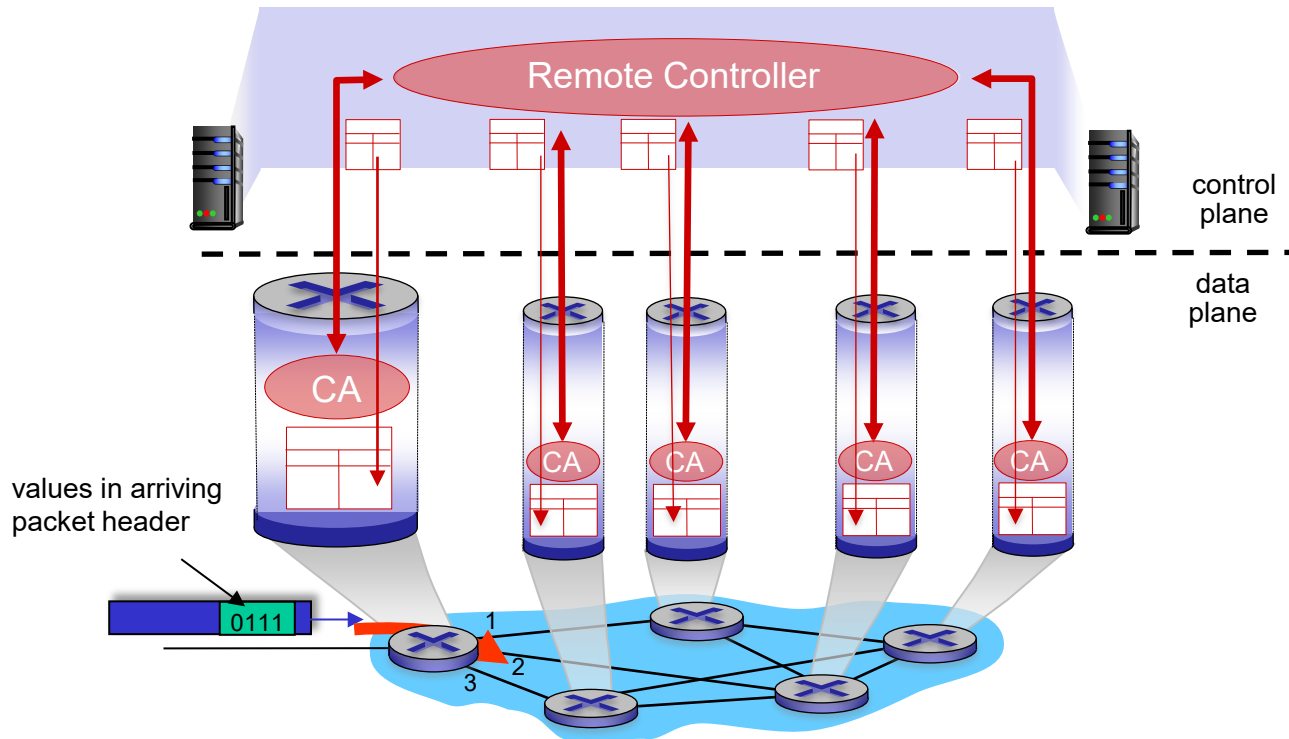
# Per-router control plane

Individual routing algorithm components *in each and every router* interact in the control plane



# Software-Defined Networking (SDN) control plane

Remote controller computes, installs forwarding tables in routers



# Network layer: "control plane" roadmap

- introduction
- routing protocols
  - link state
  - distance vector
- intra-ISP routing: OSPF
- routing among ISPs: BGP
- SDN control plane
- Internet Control Message Protocol




- network management, configuration
  - SNMP
  - NETCONF/YANG


# Routing protocols

**Routing protocol goal:** determine “good” paths (equivalently, routes), from sending hosts to receiving host, through network of routers

- **path:** sequence of routers packets traverse from given initial source host to final destination host
- **“good”:** least “cost”, “fastest”, “least congested”
- routing: a “top-10” networking challenge!



# Graph abstraction: link costs



$c_{a,b}$ : cost of *direct* link connecting  $a$  and  $b$

e.g.,  $c_{w,z} = 5$ ,  $c_{u,z} = \infty$

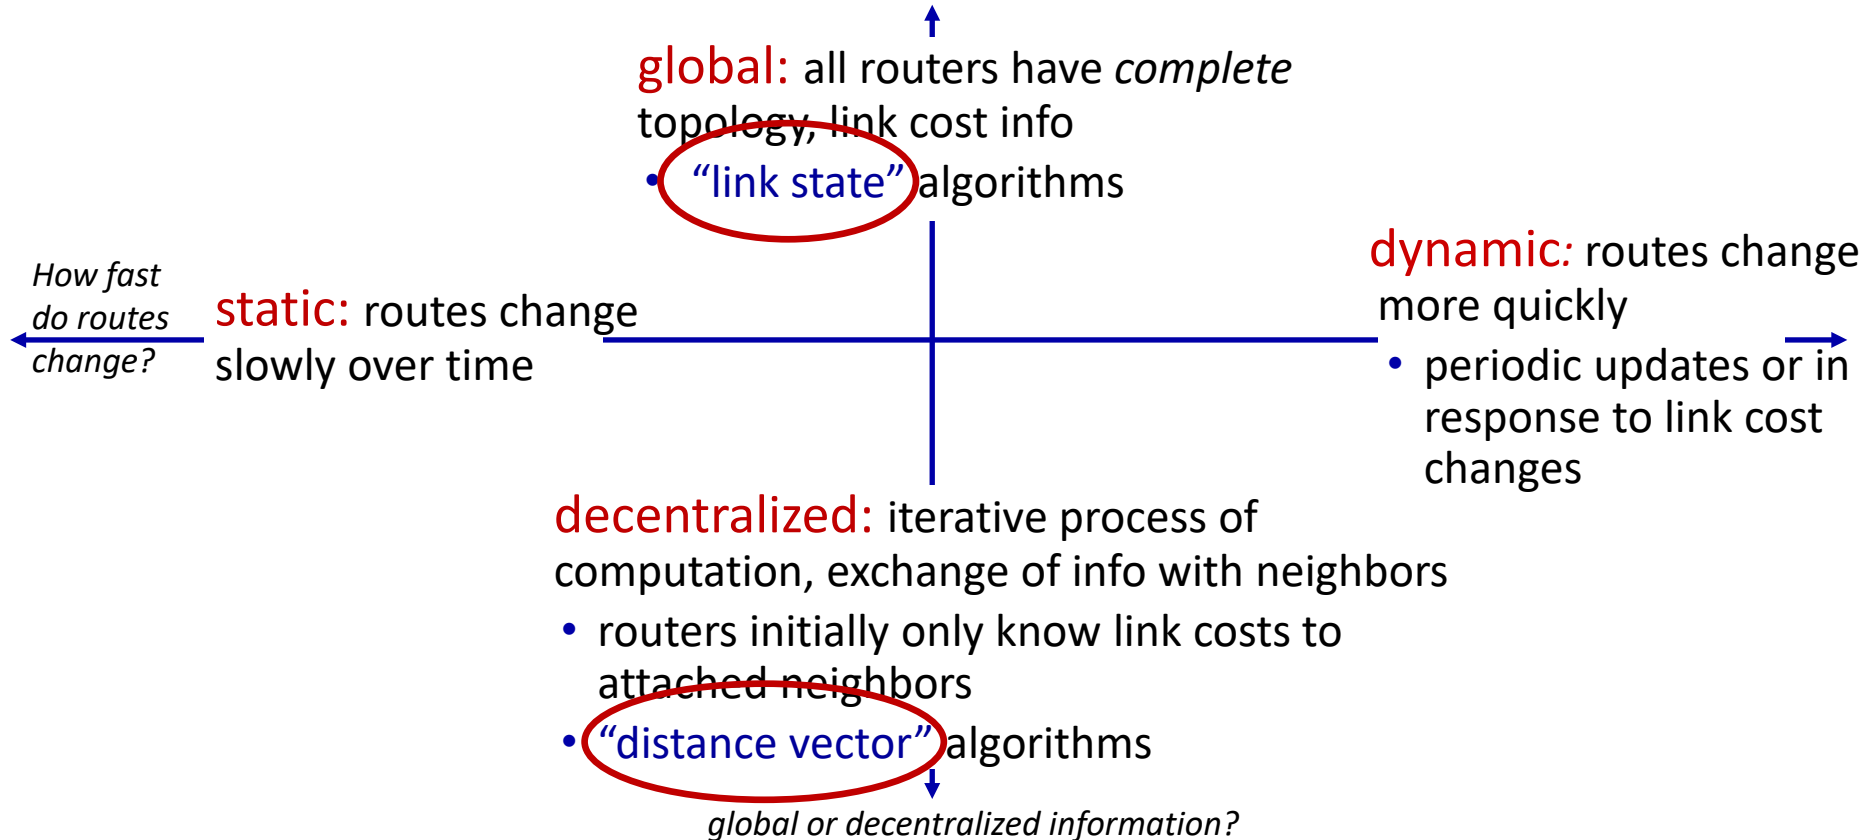
cost defined by network operator:  
could always be 1, or inversely related  
to bandwidth, or inversely related to  
congestion

graph:  $G = (N, E)$

$N$ : set of routers =  $\{ u, v, w, x, y, z \}$

$E$ : set of links =  $\{ (u,v), (u,x), (v,x), (v,w), (x,w), (x,y), (w,y), (w,z), (y,z) \}$

# Routing algorithm classification



# Network layer: "control plane" roadmap

- introduction
- routing protocols
  - link state
  - distance vector
- intra-ISP routing: OSPF
- routing among ISPs: BGP
- SDN control plane
- Internet Control Message Protocol



- network management, configuration
  - SNMP
  - NETCONF/YANG

# Dijkstra's link-state routing algorithm

- **centralized**: network topology, link costs known to *all* nodes
  - accomplished via "link state broadcast"
  - all nodes have same info
- computes least cost paths from one node ("source") to all other nodes
  - gives *forwarding table* for that node
- **iterative**: after  $k$  iterations, know least cost path to  $k$  destinations

## notation

- $c_{x,y}$ : direct link cost from node  $x$  to  $y$ ;  $= \infty$  if not direct neighbors
- $D(v)$ : *current* estimate of cost of least-cost-path from source to destination  $v$
- $p(v)$ : predecessor node along path from source to  $v$
- $N'$ : set of nodes whose least-cost-path *definitively* known

# Dijkstra's algorithm: discussion

**algorithm complexity:**  $n$  nodes


- each of  $n$  iteration: need to check all nodes,  $w$ , not in  $N$
- $n(n+1)/2$  comparisons:  $O(n^2)$  complexity
- more efficient implementations possible:  $O(n \log n)$

**message complexity:**


- each router must *broadcast* its link state information to other  $n$  routers
- efficient (and interesting!) broadcast algorithms:  $O(n)$  link crossings to disseminate a broadcast message from one source
- each router's message crosses  $O(n)$  links: overall message complexity:  $O(n^2)$

# Dijkstra's algorithm: oscillations possible


- when link costs depend on traffic volume, **route oscillations** possible
- sample scenario:
  - routing to destination a, traffic entering at d, c, e with rates 1, e (<1), 1
  - link costs are directional, and volume-dependent




initially



given these costs,  
find new routing....  
resulting in new costs



given these costs,  
find new routing....  
resulting in new costs



given these costs,  
find new routing....  
resulting in new costs

# Network layer: "control plane" roadmap

- introduction
- routing protocols
  - link state
  - **distance vector**
- intra-ISP routing: OSPF
- routing among ISPs: BGP
- SDN control plane
- Internet Control Message Protocol



- network management, configuration
  - SNMP
  - NETCONF/YANG

# Distance vector algorithm

Based on *Bellman-Ford* (BF) equation (dynamic programming):

Bellman-Ford equation

Let  $D_x(y)$ : cost of least-cost path from  $x$  to  $y$ .

Then:

$$D_x(y) = \min_v \{ c_{x,v} + D_v(y) \}$$


$\min$  taken over all neighbors  $v$  of  $x$

direct cost of link from  $x$  to  $v$

$v$ 's estimated least-cost-path cost to  $y$

# Bellman-Ford Example

Suppose that  $u$ 's neighboring nodes,  $x, v, w$ , know that for destination  $z$ :



Bellman-Ford equation says:

$$\begin{aligned} D_u(z) &= \min \{ c_{u,v} + D_v(z), \\ &\quad c_{u,x} + D_x(z), \\ &\quad c_{u,w} + D_w(z) \} \\ &= \min \{ 2 + 5, \\ &\quad 1 + 3, \\ &\quad 5 + 3 \} = 4 \end{aligned}$$

*node achieving minimum ( $x$ ) is next hop on estimated least-cost path to destination ( $z$ )*

# Distance vector algorithm

## key idea:

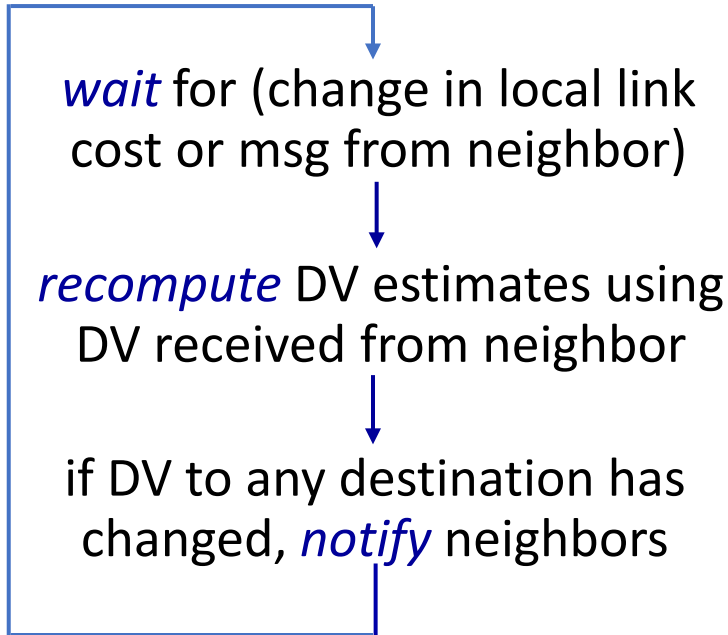
- from time-to-time, each node sends its own distance vector estimate to neighbors
- when  $x$  receives new DV estimate from any neighbor, it updates its own DV using B-F equation:

$$D_x(y) \leftarrow \min_v \{c_{x,v} + D_v(y)\} \text{ for each node } y \in N$$

- under minor, natural conditions, the estimate  $D_x(y)$  converge to the actual least cost  $d_x(y)$

# Distance vector algorithm:

each node:



**iterative, asynchronous:** each local iteration caused by:

- local link cost change
- DV update message from neighbor


**distributed, self-stopping:** each node notifies neighbors *only* when its DV changes

- neighbors then notify their neighbors – *only if necessary*
- no notification received, no actions taken!

# Distance vector: link cost changes

## link cost changes:

- node detects local link cost change
- updates routing info, recalculates local DV
- if DV changes, notify neighbors



$t_0$ : y detects link-cost change, updates its DV, informs its neighbors.

“good news travels fast”


$t_1$ : z receives update from y, updates its table, computes new least cost to x, sends its neighbors its DV.

$t_2$ : y receives z’s update, updates its distance table. y’s least costs do *not* change, so y does *not* send a message to z.

# Distance vector: link cost changes

## link cost changes:

- node detects local link cost change
- **"bad news travel slow"** – count-to-infinity problem:
  - y sees direct link to x has new cost 60, but z has said it has a path at cost of 5. So y computes "my new cost to x will be 6, via z); notifies z of new cost of 6 to x.
  - z learns that path to x via y has new cost 6, so z computes "my new cost to x will be 7 via y), notifies y of new cost of 7 to x.
  - y learns that path to x via z has new cost 7, so y computes "my new cost to x will be 8 via y), notifies z of new cost of 8 to x.
  - z learns that path to x via y has new cost 8, so z computes "my new cost to x will be 9 via y), notifies y of new cost of 9 to x.
  - ...
- see text for solutions. *Distributed algorithms are tricky!*



# Comparison of LS and DV algorithms

## message complexity

LS:  $n$  routers,  $O(n^2)$  messages sent

DV: exchange between neighbors;  
convergence time varies

## speed of convergence

LS:  $O(n^2)$  algorithm,  $O(n^2)$  messages

- may have oscillations

DV: convergence time varies

- may have routing loops
- count-to-infinity problem

robustness: what happens if router malfunctions, or is compromised?

LS:

- router can advertise incorrect *link* cost
- each router computes only its *own* table

DV:

- DV router can advertise incorrect *path* cost ("I have a *really* low cost path to everywhere"): black-holing
- each router's table used by others: errors propagate through network

# Readings for next week

## Book pages

- 408–425 (full Sec. 5.1 & 5.2)
- 453–455 (5.6 ICMP)

## Key concepts

### Routing Algorithms

- Graph,  $G=(N,E)$
- Nodes, Edges, Neighbours, Cost and Paths
- Link-state vs. Distance-vector

# Short Survey

Join at [menti.com](https://www.menti.com) | use code .....




## Instructions

Go to

[www.menti.com](https://www.menti.com)

Enter the code

**3812 3495**



Or use QR code

