Network Layer – Part II

Computer Networks, Winter 2019/2020

Lecturers: Prof. Xiaoming Fu, Dr. Yali Yuan

Assistant: Yachao Shao, MSc
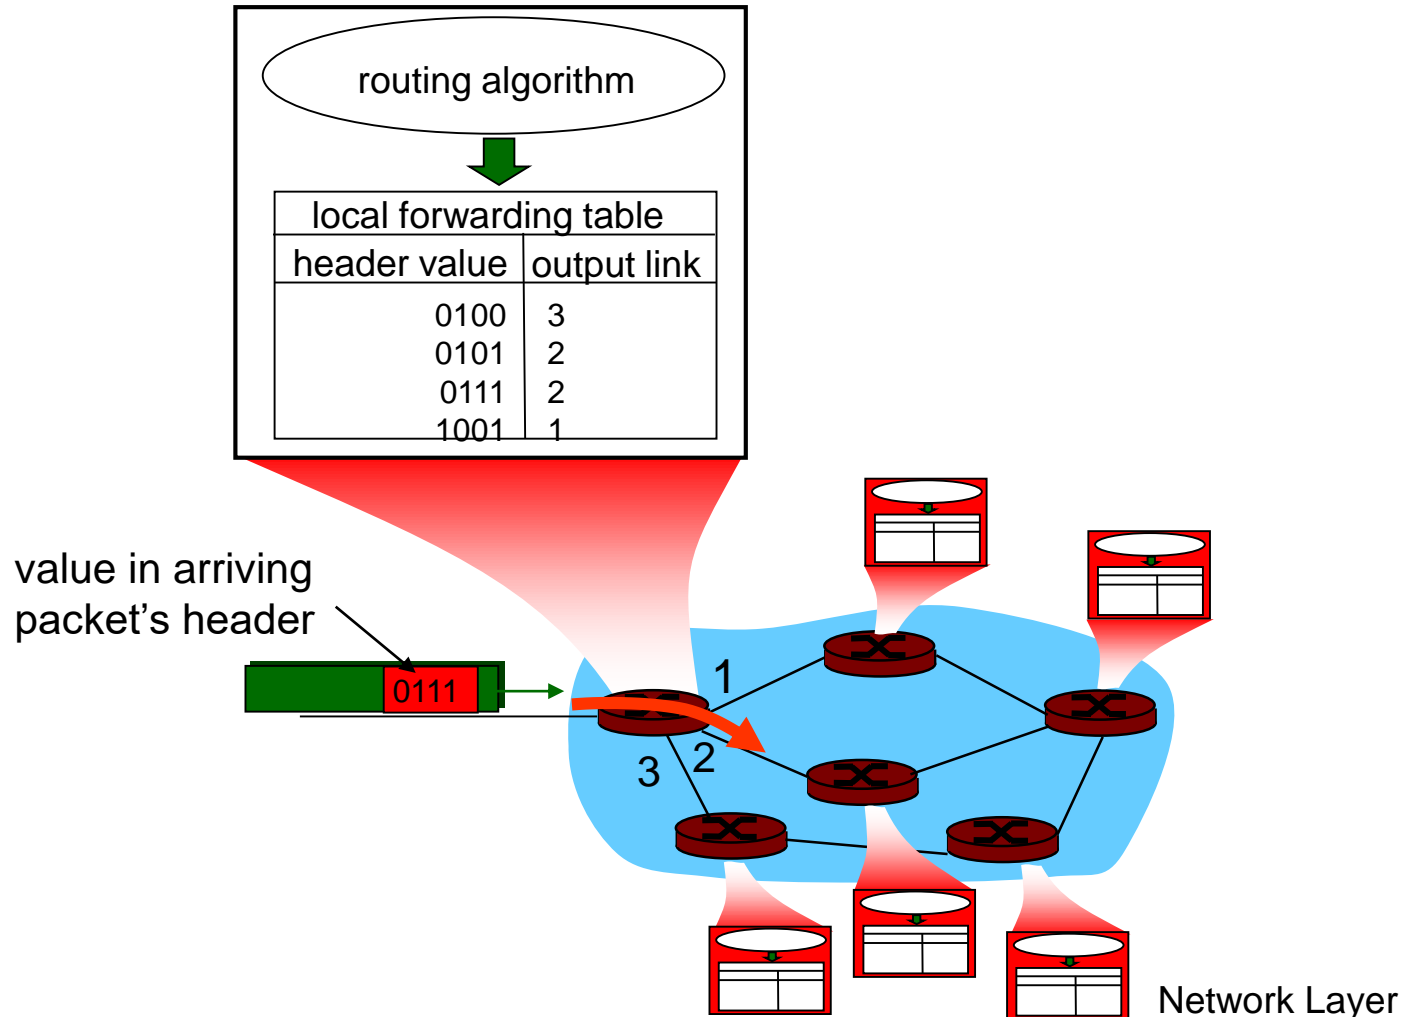
Final Exam

- Date: 10th, February, 2020

- Room: MN08

- Duration: 14:00-16:00

Network Layer II

o 3.4 Routing algorithms

   o Link state

   o Distance Vector

   o Hierarchical routing

o 3.5 Routing protocols

   o Routing Information Protocol (RIP)

   o Open Shortest Path First (OSPF)

   o Border Gateway Protocol (BGP)

Interplay between routing and forwarding

routing algorithm

local forwarding table	
header value	output link
0100	3
0101	2
0111	2
1001	1

value in arriving
packet's header

0111

1

3 2

Graph abstraction

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) }

Remark: Graph abstraction is useful in other network contexts

Example: P2P, where N is set of peers and E is set of TCP connections

Graph abstraction: costs

- c(x,x') = cost of link (x,x')

    - e.g., c(w,z) = 5

- cost could always be 1, or inversely related to bandwidth, or inversely related to congestion

Cost of path $(x_1, x_2, x_3,\ldots, x_p)$ = $c(x_1,x_2)$ + $c(x_2,x_3)$ + $\ldots$ + $c(x_{p-1},x_p)$

Question: What's the least-cost path between u and z ?

Routing algorithm: algorithm that finds least-cost path

Routing Algorithm classification

Global or decentralized information?

- Global:
  - all routers have complete topology, link cost info
  - "link state" algorithms

- Decentralized:
  - router knows physically-connected neighbors, link costs to neighbors
  - iterative process of computation, exchange of info with neighbors
  - "distance vector" algorithms

Static or dynamic?

- Static:
  - routes change slowly over time

- Dynamic:
  - routes change more quickly
    - periodic update
    - in response to link cost changes

Network Layer II

o 3.4 Routing algorithms

    o Link state

    o Distance Vector

    o Hierarchical routing

o 3.5 Routing protocols

    o Routing Information Protocol (RIP)

    o Open Shortest Path First (OSPF)

    o Border Gateway Protocol (BGP)

A Link-State Routing Algorithm

Dijkstra's algorithm

- net 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 dest.'s

Notation

- $c(x,y)$: link cost from node x to y; = ∞ if not direct neighbors
- $D(v)$: current value of cost of path from source to dest. v
- $p(v)$: predecessor node along path from source to v
- N': set of nodes whose least cost path is definitively known

Dijsktra's Algorithm

```
1  Initialization:
2    N' = {u}
3    for all nodes v
4      if v adjacent to u
5          then D(v) = c(u,v)
6      else D(v) = ∞
7
8  Loop
9    find w not in N' such that D(w) is a minimum
10   add w to N'
11   update D(v) for all v adjacent to w and not in N' :
12       D(v) = min( D(v), D(w) + c(w,v) )
13   /* new cost to v is either old cost to v or known
14     shortest path cost to w plus cost from w to v */
15 until all nodes in N'
```

Dijkstra's algorithm: example

| Step	N'	D(v) p(v)	D(w) p(w)	D(x) p(x)	D(y) p(y)	D(z) p(z)
0	u	7,u	③,u	5,u	∞	∞
1	uw	6,w		⑤,u	11,w	∞
2	uwx	⑥,w			11,w	14,x
3	uwxv				⑩,v	14,x
4	uwxvy					⑫,y
5	uwxvyz					

notes:

❖ construct shortest path tree by tracing predecessor nodes

❖ ties can exist (can be broken arbitrarily)

Dijkstra's algorithm: example

Step	N'	D(v),p(v)	D(w),p(w)	D(x),p(x)	D(y),p(y)	D(z),p(z)
0	u	2,u	5,u	1,u	∞	∞
1	ux	2,u	4,x		2,x	∞
2	uxy	2,u	3,y			4,y
3	uxyv		3,y			4,y
4	uxyvw					4,y
5	uxyvwz					

Dijkstra's algorithm: example (2)

Resulting shortest-path tree from u:

Resulting forwarding table in u:

destination	link
v	(u,v)
x	(u,x)
y	(u,x)
w	(u,x)
z	(u,x)

# Dijkstra's algorithm, discussion

Algorithm complexity: n nodes

o each iteration: need to check all nodes, w, not in N

o n(n+1)/2 comparisons: $O(n^2)$

o more efficient implementations possible: $O(n*\log(n))$

Oscillations possible:

o e.g., link cost = amount of carried traffic

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 II

o 4.4 Routing algorithms

    o Link state

    o Distance Vector

    o Hierarchical routing

o 4.5 Routing protocols

    o Routing Information Protocol (RIP)

    o Open Shortest Path First (OSPF)

    o Border Gateway Protocol (BGP)

Bellman–Ford algorithm

o Finds shortest paths in weighted, directed graph for given source

o Approach: relax all edges repeatedly until stable ($|V| - 1$ times)

Distance vector algorithm

Bellman-Ford equation (dynamic programming)

let

$d_x(y)$:= cost of least-cost path from x to y

then

$$d_x(y) = \min_v \{c(x,v) + d_v(y) \}$$

cost from neighbor v to destination y

cost to neighbor v

min taken over all neighbors v of x

Bellman-Ford example

Clearly, $d_v(z) = 5$, $d_x(z) = 3$, $d_w(z) = 3$

B-F equation says:

$$d_u(z) = \min \{ c(u,v) + d_v(z),$$
$$c(u,x) + d_x(z),$$
$$c(u,w) + d_w(z) \}$$
$$= \min \{2 + 5,$$
$$1 + 3,$$
$$5 + 3\} = 4$$

Node that achieves minimum is next
hop in shortest path ➔ forwarding table

Distance Vector Algorithm

○ $D_x(y)$ = estimate of least cost from x to y

○ Node x knows cost to each neighbor v: $c(x,v)$

○ Node x maintains distance vector $\mathbf{D}_x = [D_x(y): y \in N]$

○ Node x also maintains its neighbors' distance vectors

  ○ For each neighbor v, x maintains $\mathbf{D}_v = [D_v(y): y \in N]$

Distance vector algorithm – Basic Idea

o From time-to-time, each node sends its own distance vector estimate to neighbors

o Asynchronous

o When x receives new DV estimate from neighbor, it updates its own DV using BF equation

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

o Under minor, natural conditions, the estimate Dx(y) converge to the actual least cost dx(y)

# Distance Vector Algorithm (cont'd)

Iterative, asynchronous:
each local iteration caused by:

○ local link cost change

○ DV update message from neighbor

Distributed:

○ each node notifies neighbors *only* when its DV changes

  ○ neighbors then notify their neighbors if necessary

Each node:

wait for (change in local link cost or msg from neighbor)

↓

recompute estimates

↓

if DV to any dest has changed, *notify* neighbors

$$D_x(y) = \min\{c(x,y) + D_y(y), c(x,z) + D_z(y)\}$$
$$= \min\{2+0, 7+1\} = 2$$

$$D_x(z) = \min\{c(x,y) + D_y(z), c(x,z) + D_z(z)\}$$
$$= \min\{2+1, 7+0\} = 3$$

node x table

cost to

	x	y	z
x	0	2	7
y	∞	∞	∞
z	∞	∞	∞

from

cost to

	x	y	z
x	0	2	3
y	2	0	1
z	7	1	0

from

node y table

cost to

	x	y	z
x	∞	∞	∞
y	2	0	1
z	∞	∞	∞

from

node z table

cost to

	x	y	z
x	∞	∞	∞
y	∞	∞	∞
z	7	1	0

from

time

$$D_x(y) = \min\{c(x,y) + D_y(y), c(x,z) + D_z(y)\}$$
$$= \min\{2+0, 7+1\} = 2$$

$$D_x(z) = \min\{c(x,y) + D_y(z), c(x,z) + D_z(z)\}$$
$$= \min\{2+1, 7+0\} = 3$$

node x table

cost to

|	x	y	z
x	0	2	7
y	∞	∞	∞
z	∞	∞	∞

from

cost to

	x	y	z
x	0	2	3
y	2	0	1
z	7	1	0

from

cost to

	x	y	z
x	0	2	3
y	2	0	1
z	3	1	0

from

node y table

cost to

	x	y	z
x	∞	∞	∞
y	2	0	1
z	∞	∞	∞

from

cost to

	x	y	z
x	0	2	7
y	2	0	1
z	7	1	0

from

cost to

	x	y	z
x	0	2	3
y	2	0	1
z	3	1	0

from

node z table

cost to

	x	y	z
x	∞	∞	∞
y	∞	∞	∞
z	7	1	0

from

cost to

	x	y	z
x	0	2	7
y	2	0	1
z	3	1	0

from

cost to

|	x	y	z
x	0	2	3
y	2	0	1
z	3	1	0

from

time

Distance Vector: link cost changes

Link cost changes:

o node detects local link cost change

o updates routing info, recalculates distance vector

o if DV changes, notify neighbors

"good news travels fast"

At time t_0, y detects the link-cost change, updates its DV, and informs its neighbors.

At time t_1, z receives the update from y and updates its table. It computes a new least cost to x and sends its neighbors its DV.

At time t_2, y receives z's update and updates its distance table. y's least costs do not change and hence y does *not* send any message to z.

Distance Vector: link cost changes

Link cost changes:

- good news travels fast

- bad news travels slow - "count to infinity" problem!

- 44 iterations before algorithm stabilizes: see textbook

Poisoned reverse:

- If Z routes through Y to get to X :

  - Z tells Y its (Z's) distance to X is infinite (so Y won't route to X via Z)

Comparison of LS and DV algorithms

Message complexity

o <u>LS:</u> with n nodes, E links, O(nE) msgs sent

o <u>DV:</u> exchange between neighbors only

  o convergence time varies

Speed of Convergence

o <u>LS:</u> O(n²) algorithm requires O(nE) msgs

  o may have oscillations

o <u>DV:</u> convergence time varies

  o may be routing loops

  o count-to-infinity problem

Robustness: what happens if router malfunctions?

LS:

  o node can advertise incorrect *link* cost

  o each node computes only its *own* table

DV:

  o DV node can advertise incorrect *path* cost

  o each node's table used by others

  • error propagate through network

Network Layer II

- o 3.4 Routing algorithms
  - o Link state
  - o Distance Vector
  - o Hierarchical routing
- o 3.5 Routing protocols
  - o Routing Information Protocol (RIP)
 - o Open Shortest Path First (OSPF)
  - o Border Gateway Protocol (BGP)

Hierarchical Routing

o Our routing study thus far - idealization

    o all routers identical

    o network "flat"

    o … not true in practice

o Scale: with 200 million destinations:

    o can't store all dest's in routing tables

    o routing table exchange would swamp links

o Administrative autonomy

    o internet = network of networks

    o each network admin may want to control routing in its own network

Hierarchical Routing

- Aggregate routers into regions
  - Autonomous Systems (AS)
- Routers in same AS run same routing protocol
  - "intra-AS" routing protocol
  - routers in different AS can run different intra-AS routing protocol
- Gateway router
  - Direct link to router in another AS

Interconnected ASes

- forwarding table configured by both intra- and inter-AS routing algorithm
    - intra-AS sets entries for internal dests
    - inter-AS & intra-As sets entries for external dests

Inter-AS tasks

- suppose router in AS1 receives datagram destined outside of AS1:

  - router should forward packet to gateway router, but which one?

AS1 must:

1. learn which dests are reachable through AS2, which through AS3

2. propagate this reachability info to all routers in AS1

Job of inter-AS routing!

Network Layer II

o 3.4 Routing algorithms

   o Link state

   o Distance Vector

   o Hierarchical routing

o 3.5 Routing protocols

   o Routing Information Protocol (RIP)

   o Open Shortest Path First (OSPF)

   o Border Gateway Protocol (BGP)

Routing Protocols

o Intra-AS routing aka Interior Gateway Protocols (IGP)

   o Routing Information Protocol (RIP)

   o Open Shortest Path First (OSPF)

   o Enhanced Interior Gateway Routing Protocol (EIGRP) (Cisco proprietary for decades, until 2016)

o Inter-AS routing

   o Border Gateway Protocol (BGP)

   o Exterior Gateway Protocol (EGP)

Network Layer II

o 3.4 Routing algorithms

    o Link state

    o Distance Vector

    o Hierarchical routing

o 3.5 Routing protocols

    o Routing Information Protocol (RIP)

   o Open Shortest Path First (OSPF)

    o Border Gateway Protocol (BGP)

Routing Information Protocol (RIP)

o distance vector algorithm

o included in BSD-UNIX Distribution in 1982

o distance metric: # of hops (max = 15 hops)

From router A to subnets:

destination	hops
u	1
v	2
w	2
x	3
y	3
z	2

Network Layer II

○ 3.4 Routing algorithms

  ○ Link state

  ○ Distance Vector

  ○ Hierarchical routing

○ 3.5 Routing protocols

  ○ Routing Information Protocol (RIP)

  ○ Open Shortest Path First (OSPF)

  ○ Border Gateway Protocol (BGP)

Open Shortest Path First (OSPF)

o "open": publicly available

o uses Link State algorithm
  - o LS packet dissemination
  - o topology map at each node
  - o route computation using Dijkstra's algorithm

o OSPF advertisement carries one entry per neighbor router

o advertisements disseminated to entire AS (via flooding)
  - o carried in OSPF messages directly over IP (rather than TCP or UDP

OSPF "advanced" features (not in RIP)

o Security: all OSPF messages authenticated (to prevent malicious intrusion)

■ multiple same-cost paths allowed (only one path in RIP)

■ for each link, multiple cost metrics for different TOS (e.g., satellite link cost set low for best effort ToS; high for real-time ToS)

o Integrated uni- and multicast support:

   o Multicast OSPF (MOSPF) uses same topology data base as OSPF

o Hierarchical OSPF in large domains.

Hierarchical OSPF

Hierarchical OSPF

o two-level hierarchy: local area, backbone.

  o Link-state advertisements only in area

  o each node has detailed area topology; only know direction (shortest path) to nets in other areas.

o area border routers: "summarize" distances to nets in own area, advertise to other Area Border routers.

o backbone routers: run OSPF routing limited to backbone.

o boundary routers: connect to other AS's.

Network Layer II

o 3.4 Routing algorithms

- o Link state
- o Distance Vector
- o Hierarchical routing

o 3.5 Routing protocols

- o Routing Information Protocol (RIP)
- o Open Shortest Path First (OSPF)
- o Border Gateway Protocol (BGP)

Inter-AS routing: BGP

o BGP (Border Gateway Protocol): the de facto standard

o BGP provides each AS means to:

- eBGP: obtain subnet reachability information from neighboring ASes

- iBGP: propagate reachability information to all AS-internal routers.

- determine "good" routes to other networks based on reachability information and *policy*

o allows subnet to advertise its existence to rest of Internet: "I am here"

eBGP, iBGP connections

AS 1

AS 2

AS 3

— — — eBGP connectivity

- - - - - - iBGP connectivity

gateway routers run both eBGP and iBGP protools

BGP basics

- **BGP session:** two BGP routers ("peers") exchange BGP messages over semi-permanent TCP connection:
  - advertising *paths* to different destination network prefixes (BGP is a "path vector" protocol)
- when AS3 gateway router 3a advertises path AS3,X to AS2 gateway router 2c:
  - AS3 *promises* to AS2 it will forward datagrams towards X


BGP advertisement:
AS3, X

BGP basics

- **BGP session:** two BGP routers ("peers") exchange BGP messages over semi-permanent TCP connection:
  - advertising *paths* to different destination network prefixes (BGP is a "path vector" protocol)
- when AS3 gateway router 3a advertises path AS3,X to AS2 gateway router 2c:
  - AS3 *promises* to AS2 it will forward datagrams towards X

AS 3

AS 1

AS 2

*BGP advertisement:*
AS3, X

Path attributes and BGP routes

- advertised prefix includes BGP attributes
  - prefix + attributes = "route"
- two important attributes:
  - AS-PATH: list of ASes through which prefix advertisement has passed
  - NEXT-HOP: indicates specific internal-AS router to next-hop AS
- *Policy-based routing:*
  - gateway receiving route advertisement uses *import policy* to accept/decline path (e.g., never route through AS Y).
  - AS policy also determines whether to *advertise* path to other other neighboring ASes

BGP path advertisement

- AS2 router 2c receives path advertisement AS3,X (via eBGP) from AS3 router 3a

- Based on AS2 policy, AS2 router 2c accepts path AS3,X, propagates (via iBGP) to all AS2 routers

- Based on AS2 policy, AS2 router 2a advertises (via eBGP) path AS2, AS3, X   to AS1 router 1c

BGP path advertisement

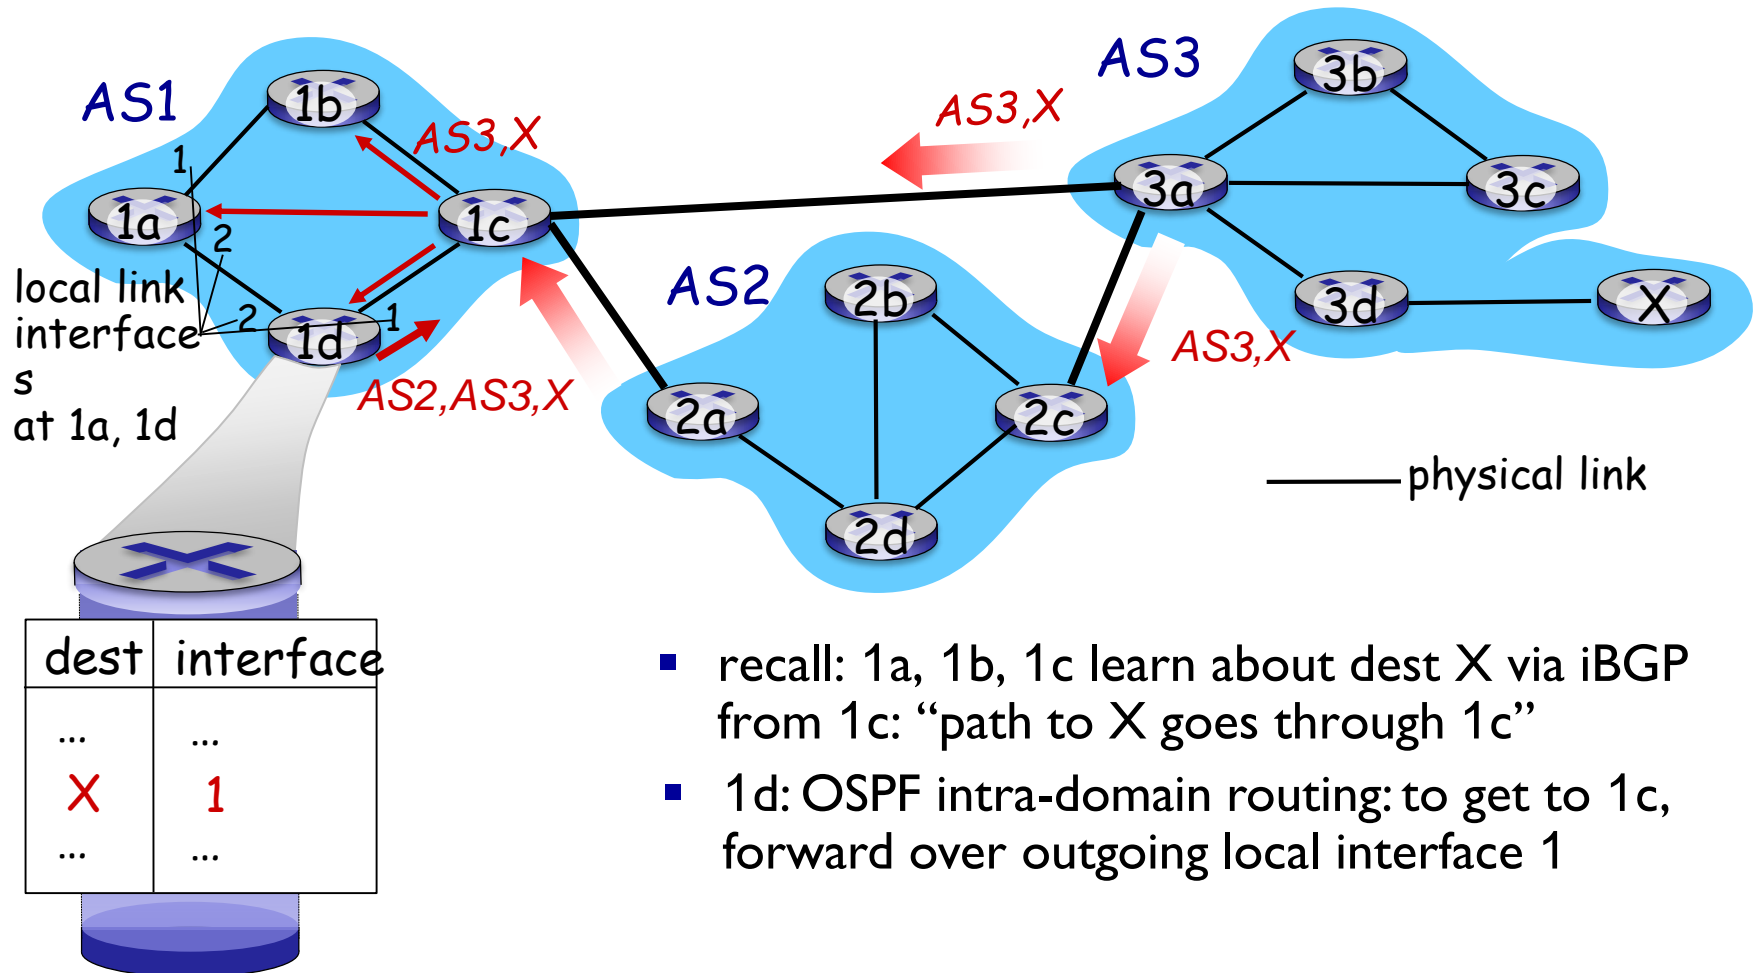
gateway router may learn about multiple paths to destination:

- AS1 gateway router 1c learns path *AS2,AS3,X* from 2a
- AS1 gateway router 1c learns path *AS3,X* from 3a
- Based on policy, AS1 gateway router 1c chooses path *AS3,X, and advertises path within AS1 via iBGP*

BGP messages

- BGP messages exchanged between peers over TCP connection
- BGP messages:
  - OPEN: opens TCP connection to remote BGP peer and authenticates sending BGP peer
  - UPDATE: advertises new path (or withdraws old)
  - KEEPALIVE: keeps connection alive in absence of UPDATES; also ACKs OPEN request
 - NOTIFICATION: reports errors in previous msg; also used to close connection

BGP, OSPF, forwarding table entries

Q: how does router set forwarding table entry to distant prefix?

AS1

AS3

AS2

1b

1a

1c

1d

AS3,X

AS3,X

AS3,X

AS2,AS3,X

local link interfaces at 1a, 1d

1

2

2

1

2b

2a

2c

2d

3b

3a

3c

3d

X

——— physical link

| dest	interface
...	...
X	1
...	...

- recall: 1a, 1b, 1c learn about dest X via iBGP from 1c: "path to X goes through 1c"

- 1d: OSPF intra-domain routing: to get to 1c, forward over outgoing local interface 1

BGP, OSPF, forwarding table entries

Q: how does router set forwarding table entry to distant prefix?
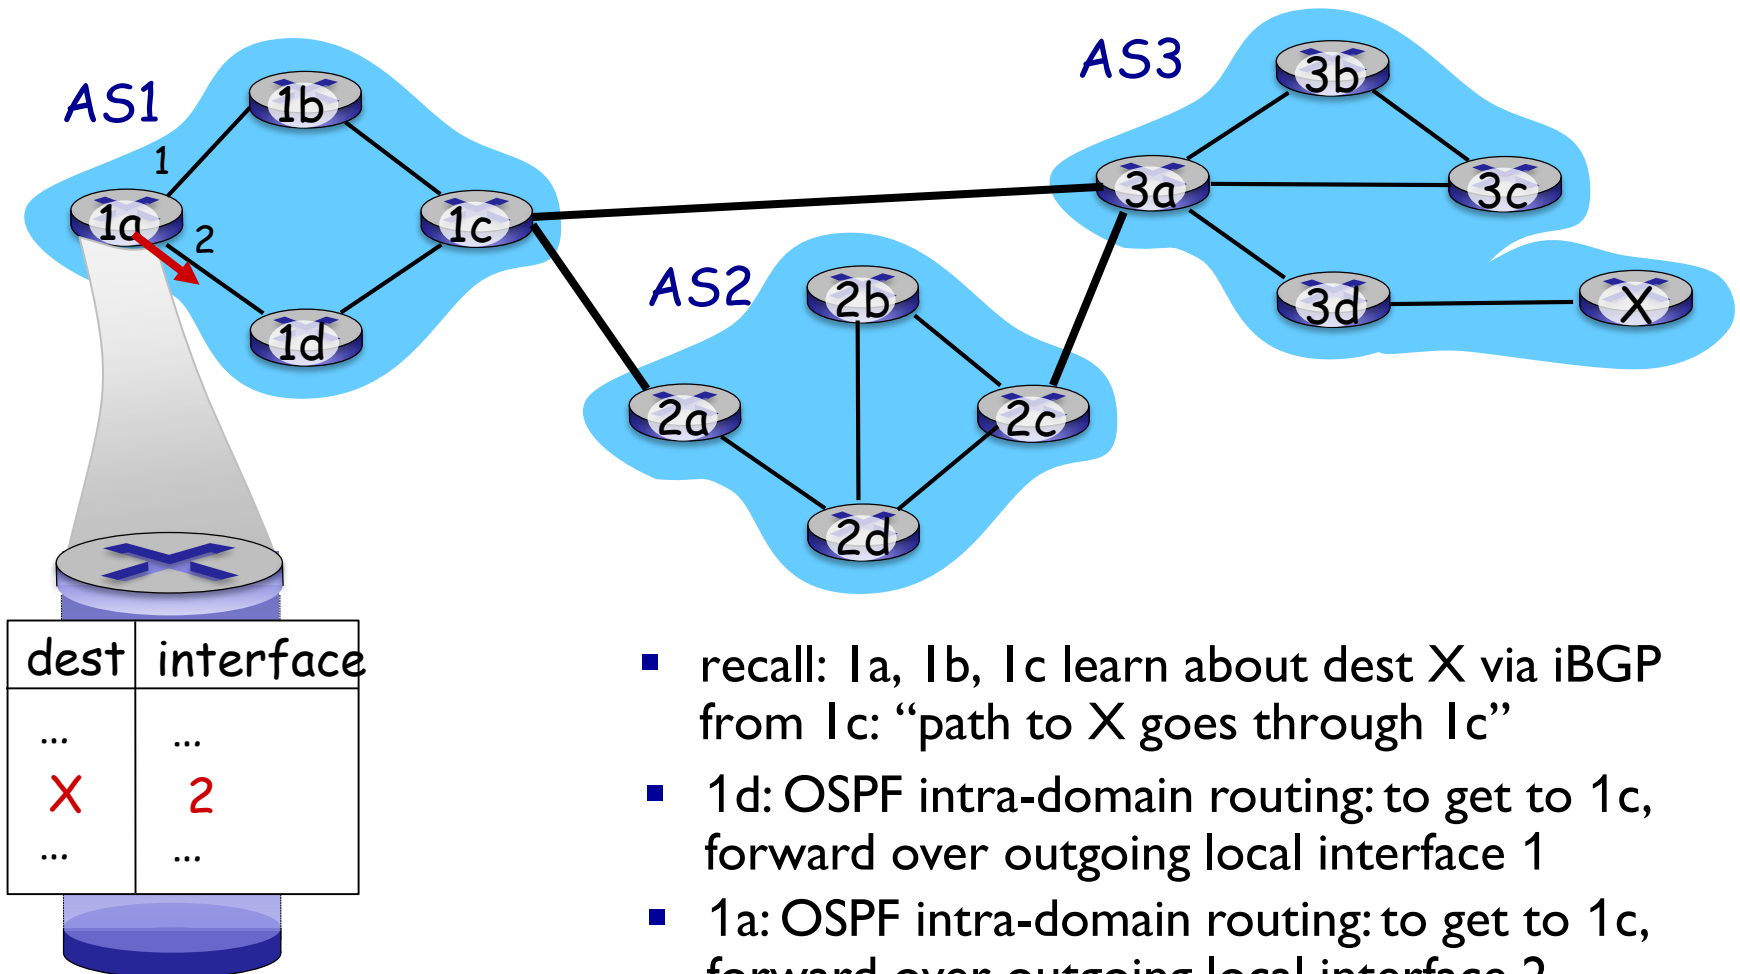
dest	interface
...	...
X	2
...	...

- recall: 1a, 1b, 1c learn about dest X via iBGP from 1c: "path to X goes through 1c"
- 1d: OSPF intra-domain routing: to get to 1c, forward over outgoing local interface 1
- 1a: OSPF intra-domain routing: to get to 1c, forward over outgoing local interface 2

BGP route selection

- router may learn about more than one route to destination AS, selects route based on:
    1. local preference value attribute: policy decision
    2. shortest AS-PATH
    3. closest NEXT-HOP router: hot potato routing
    4. additional criteria

Hot Potato Routing

- 2d learns (via iBGP) it can route to X via 2a or 2c
- *hot potato routing:* choose local gateway that has least intra-domain cost (e.g., 2d chooses 2a, even though more AS hops to *X*): don't worry about inter-domain cost!

BGP routing policy

legend: provider network

customer network:

- A,B,C are provider networks
- X,W,Y are customer networks
- X is dual-homed: attached to two networks
  - X does not want to route from B via X to C
  - .. so X will not advertise to B a route to C

BGP routing policy (2)

legend:

provider network

customer network:

o A advertises path AW to B

o B advertises path BAW to X

o Should B advertise path BAW to C?

   o No way! B gets no "revenue" for routing CBAW since neither W nor C are B's customers

   o B wants to force C to route to w via A

   o B wants to route *only* to/from its customers!

Why different Intra- and Inter-AS routing?

- Policy
  - Inter-AS: admin wants control over how its traffic routed, who routes through its net.
  - Intra-AS: single admin, so no policy decisions needed

- Scale
  - hierarchical routing saves table size, reduced update traffic

- Performance
  - Intra-AS: can focus on performance
  - Inter-AS: policy may dominate over performance

Traffic engineering: difficult traditional routing

_Q:_ what if network operator wants u-to-z traffic to flow along _uvw_z, x-to-z traffic to flow _xwyz_?

_A:_ need to define link weights so traffic routing algorithm computes routes accordingly (or need a new routing algorithm)!

Link weights are only control "knobs": wrong!

Traffic engineering: difficult

Q: what if w wants to route blue and red traffic differently?

A: can't do it (with destination based forwarding, and LS, DV routing)

Software defined networking (SDN)

4. *programmable control applications*

routing

access control

. . .

load balance

3. *control plane functions external to data-plane switches*

Remote Controller

control plane

- -

data plane

CA

CA

CA

CA

CA

2. *control, data plane separation*

1: *generalized" flow-based" forwarding (e.g., OpenFlow)*

Thank you

Any questions?