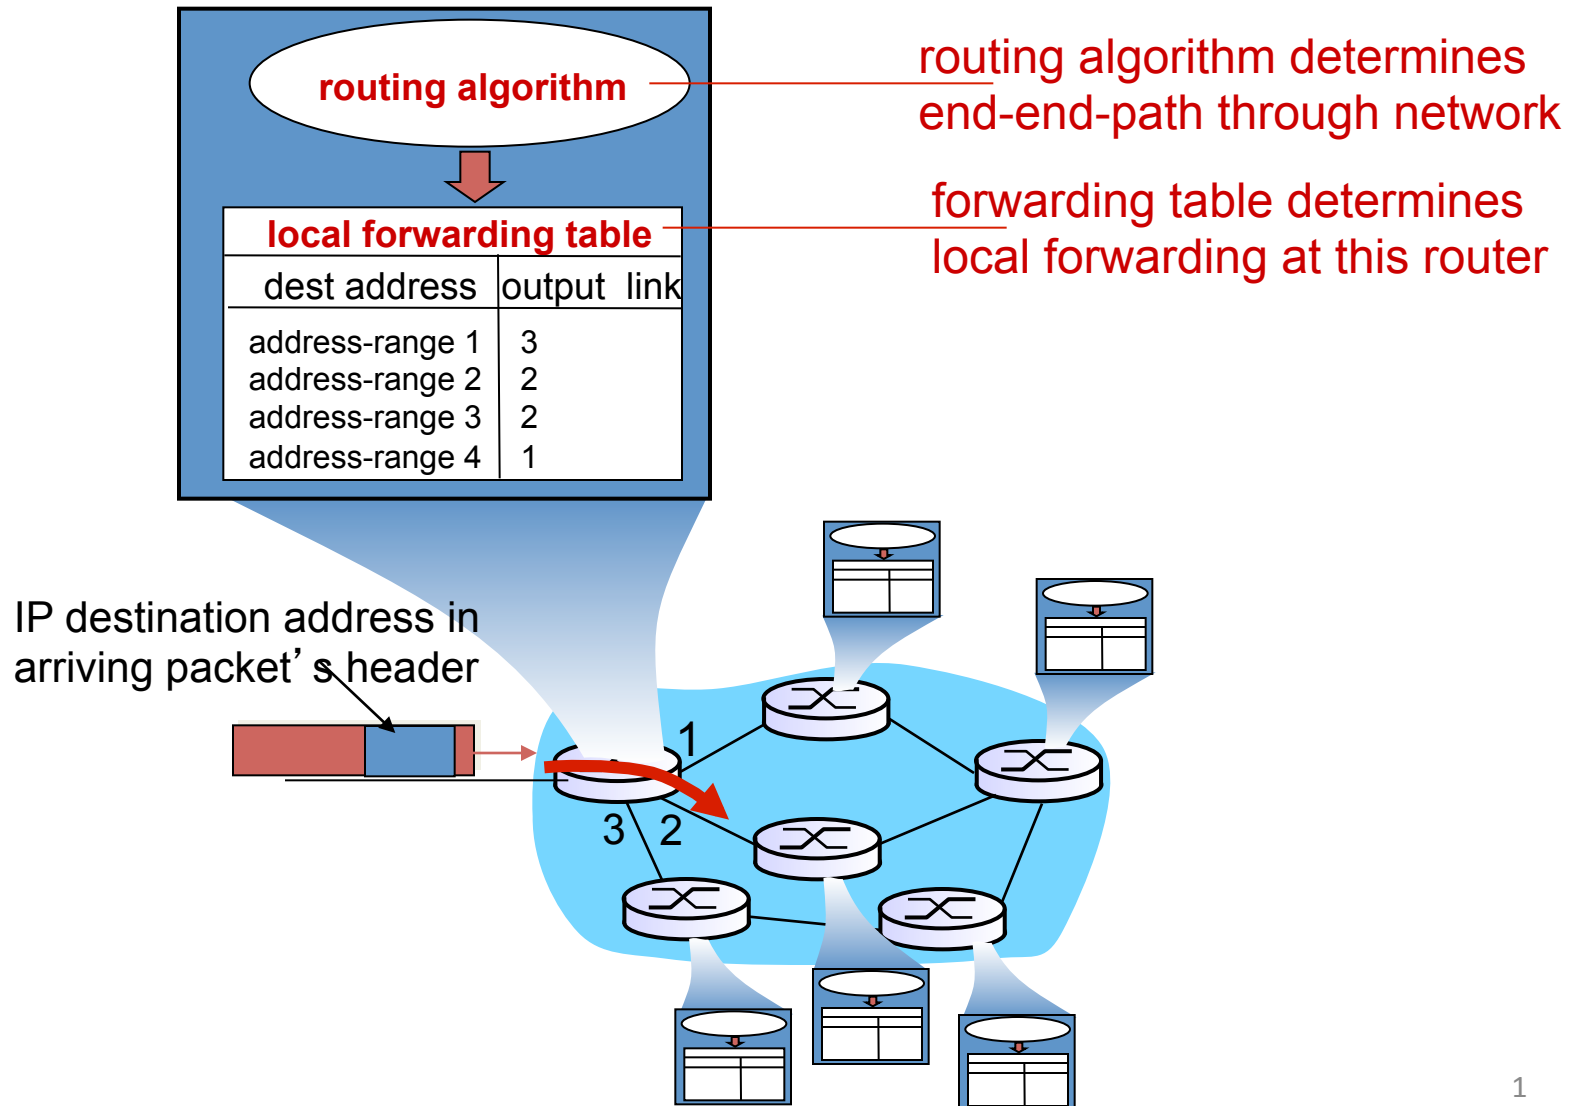



Interplay between routing, forwarding



Graph abstraction




graph: $G = (N, E)$

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

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

aside: graph abstraction is useful in other network contexts, e.g., P2P, where N is set of peers and E is set of TCP connections

Graph abstraction: costs



$c(x, x') = \text{cost of link } (x, x')$
e.g., $c(w, z) = 5$

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

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

key question: what is the least-cost path between u and z ?
routing algorithm: algorithm that finds that least cost path

Routing algorithm classification

Q: 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

Q: static or dynamic?

static:

- routes change slowly over time

dynamic:

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

Select a correct statement

- A. In a global routing algorithm, all routers have complete topology, link cost info
- B. In a decentralized routing algorithm, router knows physical-connected neighbors, link costs to neighbors
- C. Link costs might be changed over the time
- D. A and B
- E. A, B and C

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; $= \infty$ 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 definitively known

Dijkstra'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 $D(v) = \infty$
7

8 ***Loop***

15 ***until all nodes in N'***

Dijkstra'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 $D(v) = \infty$

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			y	z
3	uwxv				10,v	14,x
4	uwxvy					12,y
5	uwxvyz					

A. $y = 11, w ; z = \infty$


B. $y = 10, v ; z = 12, y$

C. $y = 11, w ; z = 14, x$

D. $y = 10, w ; z = \infty$

notes:

- ❖ construct shortest path tree by tracing predecessor nodes
- ❖ ties can exist (can be broken arbitrarily)




Dijkstra's algorithm: another 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					


Shortest path from u to w according to above result

- A. 5 through v
- B. 4 through x
- C. 3 through x
- D. 3 through y



Dijkstra's algorithm: example (2)

resulting shortest-path tree from u:



resulting forwarding table in u:

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

A. $a = (u,v)$; $b = (u,x)$

B. $a = (u,u)$; $b = (u,u)$

C. $a = (u,x)$; $b = (u,v)$


Dijkstra's algorithm, discussion

algorithm complexity: n nodes


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

oscillations possible:


- e.g., support link cost equals 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

Assignment 5 – Reliable Communication

- Will be published today at 3 PM
- Deadline 7AM Wednesday April 8, 2015
- Need to install miniset as soon as possible
 - <http://mininet.org/download/> option 2
- Explanation and Demo will be in next lecture

Next lecture

- Explanation and Demo for Assignment 5
- The Distance-Vector (DV) Routing Algorithm
 - 4.5.2
- Hierarchical Routing
 - 4.5.3