Network Layer II

ELEC 1200

Routing
Link-state routing with
Dijkstra's shortest-paths algorithm

*The slides are adapted from ppt slides (in substantially unaltered form) available from "Computer Networking: A Top-Down Approach," 4th edition, by Jim Kurose and Keith Ross, Addison-Wesley, July 2007. Part of the materials are also adapted from MIT 6.02 course notes.

Layering in the Internet



Network layer

- transport segment from sending to receiving host
- on sending side encapsulates segments into datagrams
- on rcving side, delivers segments to transport layer
- network layer protocols in *every* host, router
- router examines header fields in all IP datagrams passing through it



Datagram networks

- no call setup at network layer
- routers: no state about end-to-end connections
 - no network-level concept of "connection"
- packets forwarded using destination host address
 - packets between same source-dest pair may take different paths



Router Architecture Overview

Two key router functions:

- run routing algorithms/protocol
- *forwarding* datagrams from incoming to outgoing link



Interplay between routing and forwarding



Why is Routing Hard?

- Inherently distributed problem
 - Information about links and neighbors is local to each node, but we want global reach
- Handling dynamic conditions difficult
 - Must tolerate link, switch, and network faults
 - Failures and recovery could be arbitrarily timed, messages could be lost, etc.
 - Mobility makes life even harder
- Scaling to large sizes difficult
- And on the Internet, many independent organizations must cooperate

Graph abstraction



Graph: G = (N,E)

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

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

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')

 cost could be inversely related to bandwidth

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

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 physicallyconnected 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

Link-State Routing

- Conceptually, a two-step process
- Each node *floods* information about its links to its neighbors; they re-send to their neighbors, etc.
 - This process allows each node to discover every other node and link
- Each node then runs the same local shortest path computation over its version of graph
 - If each node implements computation correctly and each node has the same graph, then resulting forwarding will work correctly
 - Each node makes *local* decision about next hop

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 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**

- 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





Dijkstra's algorithm: example

Resulting shortest-path tree from u:



Resulting forwarding table in u:

destination	link		
V	(u,v)		
×	(u,x)		
У	(u,x)		
W	(u,x)		
Z	(u,x)		

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(n²)

St	ер	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•	<u>2,u</u>	З,у			4,y
	3	uxyv 🗲		-3,y			4,y
	4	uxyvw 🔶					4,y
	5	uxyvwz 🔶					



Summary

- The network layer implements the "glue" that achieves connectivity
 - Does addressing, forwarding, and routing
- Forwarding entails a routing table lookup; the table is built using routing algorithms such as Dijkstras Link State algorithm