

COMS3200

COMS3200 – Week 9 Routing

School of Information Technology and Electrical Engineering
The University of Queensland

COMS3200

Outline

- Quick Quiz
- Routing
 - Distance Vector Routing
 - Link State Routing
 - Dijkstra's algorithm
 - Broadcast routing
 - Multicast routing
- Routing in mobile or dynamic networks
 - Delivery of packets to Mobile Hosts
 - Introduction to routing in Ad-hoc networks

2

COMS3200

This Week's Learning Objectives

- Understand and be able to reproduce the operation of simple routing algorithms
- Understand broadcast and multicast routing
- Understand routing protocols supporting mobile hosts and dynamically changing networks

3

COMS3200

Quick Quiz

- A message of 3000 bytes is to be transmitted using IP over two networks, the first with an MTU of 2000 bytes, the second with an MTU of 1020 bytes. How many fragments are expected at the destination and what will be the values of the following IP header fields in each fragment?
 - MF
 - Fragment offset
 - Packet length
- What's the purpose of the Protocol field in an IP header?

4

COMS3200

Routing

- Consider datagram network
 - Have datagram packet with destination address
 - **HOW** do we get it to its destination?
- More precisely...
 - How does a router determine which output port is the best choice for that packet?

5

COMS3200

How to determine the best output port for a packet?

- Simple answer...
 - Consult a [routing table](#) (or forwarding table)
 - Simple Example:

IP routing table	Destination	Gateway	Genmask	Iface
	192.168.118.0	*	255.255.255.0	eth0
	127.0.0.0	*	255.0.0.0	lo
	default	gateway118.slabb	0.0.0.0	eth0

6

COMS3200

Forwarding vs Routing

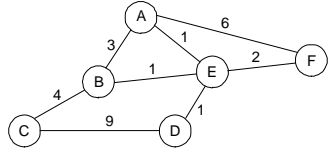
- **Forwarding**
 - Sending packet in a direction determined by routing table
 - Sometimes called switching
- **Routing**
 - Process by which routing tables are built

T-362 7

COMS3200

Network as a Graph

- Network represented as a **Graph**



```

graph TD
    A ---|3| B
    A ---|1| E
    A ---|6| F
    B ---|4| C
    B ---|1| E
    C ---|9| D
    D ---|1| E
    E ---|2| F
    
```

- We'll consider
 - nodes to be routers
 - edges to be network links, weighted by some "cost" (e.g. delay, \$ cost, congestion...)
- Problem: Find lowest cost path between two nodes
 - Path cost = sum of link costs

8

COMS3200

Static Approach to Routing

- Calculate all shortest paths and store them permanently in each node
- What's wrong with this approach?
 - Link or node failure
 - Can't handle dynamic changes
 - Addition
 - Link weight changes
- Practically, network routing is **dynamic** – we run **routing protocols** between the nodes

9

COMS3200

Routing Protocols

- Distributed, not centralised
 - What's wrong with centralisation?
- Distribution can cause problems
 - at one instant, two routers may have different ideas of shortest paths in the network
 - can cause routing loops

10

COMS3200

Routing Algorithm Classification

<p>Global or decentralized information?</p> <p>Global:</p> <ul style="list-style-type: none"> ● all routers have complete topology, link cost info ● "link state" algorithms <p>Decentralized:</p> <ul style="list-style-type: none"> ● router knows physically-connected neighbors, link costs to neighbors ● iterative process of computation, exchange of info with neighbors ● "distance vector" algorithms 	<p>Static or dynamic?</p> <p>Static:</p> <ul style="list-style-type: none"> ● routes change slowly over time <p>Dynamic:</p> <ul style="list-style-type: none"> ● routes change more quickly <ul style="list-style-type: none"> ■ periodic update ■ in response to link cost changes
---	--

11

COMS3200

Intradomain Routing

- Within a **domain**
 - Domain = internetwork where routers are under same administrative control (e.g. University campus)
 - Small to mid-sized network
 - NOT the Internet (interdomain routing)
- Protocols called **Interior Gateway Protocols (IGPs)**
 - Distance Vector Routing
 - Link State Routing

12

The Optimality Principle

If J is on optimal path from E to K, then optimal path from J to K also falls along the same route

(a) A network (b) A sink tree for router B

T-365 13

Distance Vector Routing Summary

- Each node maintains a set of triples
 - (Destination, Cost, NextHop)
- Updates exchanged with directly connected neighbors
 - periodically (on the order of several seconds)
 - whenever table changes (called *triggered* update)
 - Each update is a list of pairs:
 - (Destination, Cost)
- Update local table if receive a "better" route
 - smaller cost
 - came from next-hop
- Refresh existing routes; delete if they time out

14

Example

- Assume link weights are all 1
 - Example table for Node B

Destination	Cost	NextHop
A	1	A
C	1	C
D	2	C
E	2	A
F	2	A
G	3	A

- How do we generate this table?

15

Distance Vector Routing Table Generation

- Initially, each node sets
 - cost of 1 to directly connected nodes
 - cost of infinity (∞) to all other nodes
- Then, each node sends a message to its directly connected neighbours containing its list of distances
- Continues until **convergence**
 - steady state

16

Example / Exercise

- Generate distance vector routing table for node A of previous example

17

Routing Updates

- Two circumstances
 - Periodic
 - Each node sends update message periodically, even if no change
 - Triggered
 - When node receives update from neighbour that causes routing table change, then send update to neighbours
- Link failure detection
 - Test link regularly (control packet)
 - Non-receipt of expected periodic update

18

COMS3200

Routing Update Example 1

- F detects that link to G has failed
- F sets distance to G to infinity and sends update to A
- A sets distance to G to infinity since it uses F to reach G
- A receives periodic update from C with 2-hop path to G
- A sets distance to G to 3 and sends update to F
- F decides it can reach G in 4 hops via A

19

COMS3200

Routing Update Example 2

- link from A to E fails
- A advertises distance of infinity to E
- B and C advertise a distance of 2 to E
- B decides it can reach E in 3 hops; advertises this to A
- A decides it can reach E in 4 hops; advertises this to C
- C decides that it can reach E in 5 hops...
- Count to infinity problem...

20

COMS3200

Loop-Breaking Heuristics

- Set infinity to some "small" number
 - e.g. 16
- Split horizon
 - Don't send what you learnt from your neighbour back to your neighbour
- Split horizon with poison reverse
- Split horizon techniques only work for loop of 2 nodes, need more drastic techniques for larger loops

21

COMS3200

Routing Information Protocol (RIP)

- Classic example of a distance-vector routing protocol
 - Included in BSD-UNIX Distribution in 1982
 - routed daemon
 - Periodic updates every 30 seconds
- Most common distance-vector protocol today is **IGRP**
 - Interior Gateway Routing Protocol
 - Cisco proprietary

22

COMS3200

Link State Routing

- Strategy
 - Send to all nodes (not just neighbours) information (i.e. state) about directly connected links (not entire routing table)
- Steps – each router must
 1. Discover its neighbours and learn their network addresses
 2. Measure the delay/cost to each of its neighbours
 3. Construct a packet telling all it has just learned
 4. Send this packet to all other routers
 5. Compute the shortest path to every other router

23

COMS3200

Learning about the Neighbors

(a) Nine routers and a broadcast LAN, (b) A graph model of (a)

HELLO packet is sent to neighbours
N- artificial node representing LAN

T-374 24

Measuring Line Cost

Each router has to get an estimate of the delay to each neighbour (ECHO packet)

25

Building Link State Packets

(a) A subnet

(b) The link state packets for all routers

26

Link State Packet

- Link State Packet (LSP), contains
 - id of the node that created the LSP
 - cost of link to each directly connected neighbor
 - sequence number (SEQNO)
 - time-to-live (TTL) for this packet
- How are packets distributed reliably?

27

Distributing the Link State Packets

The packet buffer for router B

Source	Seq.	Age	Send flags			ACK flags			Data
			A	C	F	A	C	F	
A	21	60	0	1	1	1	0	0	
F	21	60	1	1	0	0	0	1	
E	21	59	0	1	0	1	0	1	
C	20	60	1	0	1	0	1	0	
D	21	59	1	0	0	0	1	1	

28

Link State Routing Summary

- Store most recent LSP from each node
- Forward LSP to all nodes but one that sent it (flooding)
- Generate new LSP periodically
 - increment SEQNO
- Start SEQNO at 0 when reboot
- Decrement TTL of each stored LSP
 - discard when TTL=0

29


Link State (cont)

- Reliable flooding
 - Making sure all nodes get a copy of link-state information from all other nodes
- Reliability through
 - acknowledgements/retransmissions
 - sequence numbers ensure only store/forward latest information
 - discard LSPs from node with SEQNO <= current stored SEQNO for that node
 - store and forward LSPs with SEQNO > current
 - don't forward to node from which received

30

COMS3200

LSP Flooding



31

COMS3200

LSP Generation

- Two circumstances
 - Periodically
 - Not as often as for DV since LSPs reliable
 - Change in topology
 - directly connected link or neighbour down
- New LSPs
 - increment sequence number
 - sequence numbers do NOT wrap
- When reboot, start SEQNO at 0

32

COMS3200

LSP Time-to-live

- LSPs also carry time-to-live (TTL)
 - Ensures old information removed
- TTL decremented
 - before forwarding
 - periodically
- When TTL reaches 0
 - LSP reflooded as an indication to delete that LSP from other nodes
 - LSP discarded

33

COMS3200

How does this help routing?

- Steps
 - Every node receives a copy of the LSP of every other node
 - Compute network map (topology)
 - Decide on best route to each destination
- Dijkstra's shortest path algorithm

34

COMS3200

Dijkstra's Algorithm

- N = set of nodes in the graph
- $I(i, j)$ = non-negative cost for edge (i, j)
 - infinity if no edge connecting i, j
- s = this node
- M = the set of nodes incorporated so far
- $C(n)$ = cost of the path from s to node n

```

M = {s}
for each n in N - {s}
    C(n) = I(s, n)
while (N != M)
    M = M ∪ {w} such that C(w) is the minimum for
        all w in (N - M)
    for each n in (N - M)
        C(n) = MIN(C(n), C(w) + I(w, n))
    
```

35

COMS3200

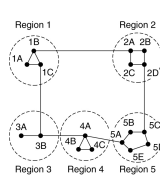
Now in Plain(er) English...

- Start with node s
 - permanently label s with $[0, s]$
 - tentatively label others $(\infty, -)$
 - make node s the working node, n_w
- Repeat until all nodes permanently labelled
 - For each tentatively labelled node (n) next to n_w calculate
 - d = cost to n_w + distance from n to n_w
 - if $d < n$'s tentative distance then tentatively relabel to (d, n_w)
 - Of all tentatively labelled nodes, select the one with the least cost, make its label permanent and make it the working node n_w

36

Hierarchical Routing

In large networks routing has to be hierarchical



Dest.	Line	Hops
1A	--	--
1B	1B	1
1C	1C	1
2A	1B	2
2B	1B	3
2C	1B	3
2D	1B	4
3A	1C	3
3B	1C	2
4A	1C	3
4B	1C	4
4C	1C	4
5A	1C	4
5B	1C	5
5C	1B	5
5D	1C	6
5E	1C	5

Dest.	Line	Hops
1A	--	--
1B	1B	1
1C	1C	1
2	1B	2
3	1C	2
4	1C	3
5	1C	4

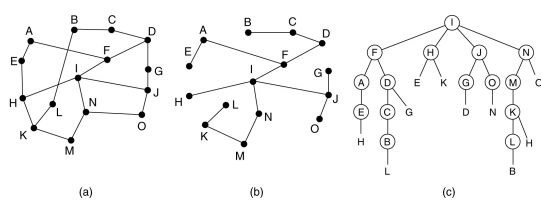
T-380 (a) (b) (c) 43

Broadcast Routing

- Broadcast – sending to all destinations simultaneously
- Various protocols:
 - Flooding
 - Multi-destination routing
 - list of destinations included
 - Sending using spanning tree
 - Reverse path forwarding

44

Broadcast – Reverse path forwarding



Reverse path forwarding (a) A subnet (b) a Sink tree for I (c) The tree built by reverse path forwarding

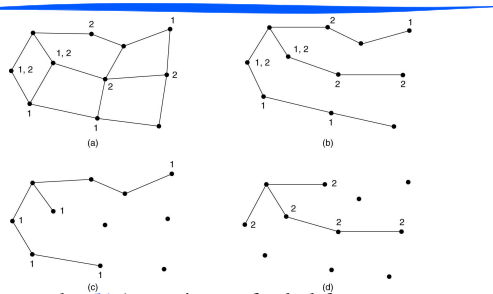
T-381 45

Multicast Routing

- Multicast – sending to a group
- Requires group management
 - Creation and deletion of groups
 - Joining and leaving groups
- Multicast routing
 - Spanning tree covering all routers is calculated
 - Spanning tree is pruned to remove nodes which do not belong to the group
- Alternative approach: core-based trees

46

Multicast Routing

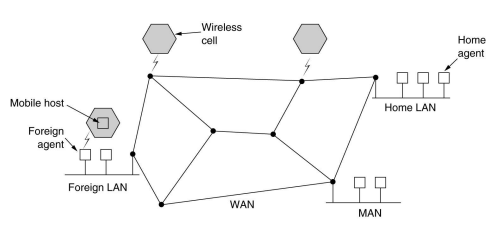


(a) A network. (b) A spanning tree for the leftmost router. (c) A multicast tree for group 1. (d) A multicast tree for group 2.

T-384 47

Routing for Mobile Hosts

A WAN to which LANs, MANs, and wireless cells are attached



48

COMS3200

Routing for Mobile Hosts

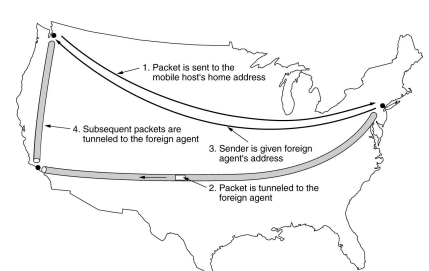
- All hosts have a permanent home location (IP address)
- Each small network has a home agent and foreign agent
 - Home agent keeps track of mobile hosts which are away from home
 - Foreign agent deals with network visitors
 - mobile hosts have to register with foreign agent

49

COMS3200

Routing for Mobile Hosts

Packet routing for mobile users



50

COMS3200

Routing in Ad Hoc Networks (MANETs)

Possibilities when the routers are mobile:

- Military vehicles on battlefield
 - No infrastructure
- A fleet of ships at sea
 - All moving all the time
- Emergency works at earthquake
 - The infrastructure destroyed
- A gathering of people with notebook computers
 - In an area lacking 802.11

51

COMS3200

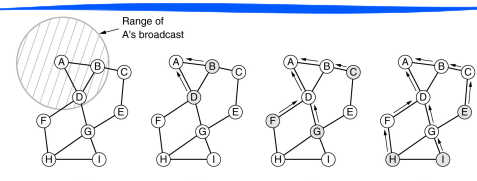
Routing in Ad Hoc Networks (MANETs)

- Very dynamic network topology
 - Nodes (hosts and routers) may come and go
- A variety of routing protocols
 - Best known: AODV (Ad hoc On-demand Distance Vector routing)
 - Takes into account limited bandwidth and low battery life
- AODV has to discover routes and maintain them

52

COMS3200

Route Discovery



(a) Range of A's broadcast
 (b) After B and D have received A's broadcast
 (c) After C, F, and G have received A's broadcast
 (d) After E, H, and I have received A's broadcast

Shaded nodes are new recipients. Arrows show possible reverse routes.

53

COMS3200

Route Discovery

Source address	Request ID	Destination address	Source sequence #	Dest. sequence #	Hop count
----------------	------------	---------------------	-------------------	------------------	-----------

Format of a ROUTE REQUEST packet

54

COMS3200

Route Discovery

Source address	Destination address	Destination sequence #	Hop count	Lifetime
----------------	---------------------	------------------------	-----------	----------

Format of a ROUTE REPLY packet

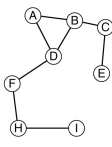
55

COMS3200

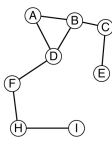
Route Maintenance

Each node periodically broadcasts HELLO to learn which nodes are no longer there

Dest.	Next hop	Distance	Active neighbors	Other fields
A	A	1	F, G	
B	B	1	F, G	
C	B	2	F	
E	G	2		
F	F	1	A, B	
G	G	1	A, B	
H	F	2	A, B	
I	G	2	A, B	



(a)



(b)

(a) D's routing table before G goes down
 (b) The graph after G has gone down

56

COMS3200

Routing in dynamic networks

More about Ad-hoc routing
 (and also Routing in Peer-to-Peer)
 in **COMS4200/7200**

57

COMS3200

Readings

- 5th ed: pp 362-392 (Routing); pp 474-484 (OSPF, BGP) ← less important
 - (4th ed: pp 350-380 (Routing) 454-461)
 - IGRP: http://docwiki.cisco.com/wiki/Interior_Gateway_Routing_Protocol
- Next week: Multimedia
- Read: Tanenbaum 5th ed.
 - 7.4 – 7.4.1, 7.4.3, 7.4.4, 7.4.5,
 - 7.5 – 7.5.1, 7.5.2, 7.5.3, 7.5.4(BitTorrent)
- (4th ed: 7.4. Does not cover material in 7.5)

58