Routing in Switched Data Networks

ITS323: Introduction to Data Communications

Sirindhorn International Institute of Technology
Thammasat University

Prepared by Steven Gordon on 4 August 2011
ITS323Y11S1L10, Steve/Courses/ITS323/Lectures/routing.tex, r1903

Contents

Routing in Switched Networks

Which path or route to take from source to destination?

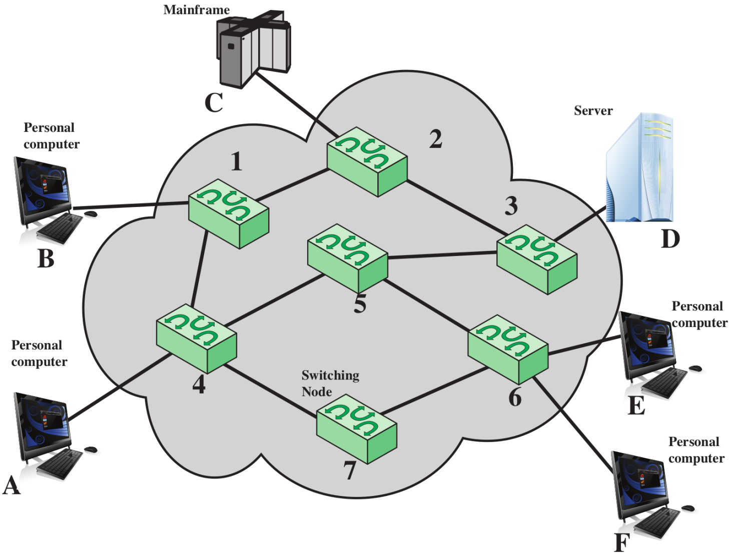

Routing in Switched Networks

- Routing is a key design issue in switched networks
- Question: What path (route) should be taken from source to destination?
- Answer: Choose the "best" path!
- What is "best", and how to choose it?
- Real networks may have 100's to 100,000+ nodes, and many possible paths
- Routing is needed in circuit-switched and packet-switched networks (we focus on packet-switched networks)

Requirements of Routing Algorithms

Correctness	path must be from intended source to intended destination
Simplicity	easy/cheap to implement
Robustness	still deliver in presence of errors or overload
Stability	path changes should not be too frequent
Optimality	choose best paths
Fairness	ensure all stations obtain equal performance
Efficiency	minimise the amount of processing and transmission overhead

Routing Terminology

Link direct connection between two nodes

Path a way between two nodes, via one or more links

Hop to traverse a link

Neighbour a node at the other end of a link

Cost value assigned to link to indicate cost of using that link

Topology arrangement of nodes and links in a network

Least-cost routing is typically used: choose a path with least cost

Example of Network Configuration

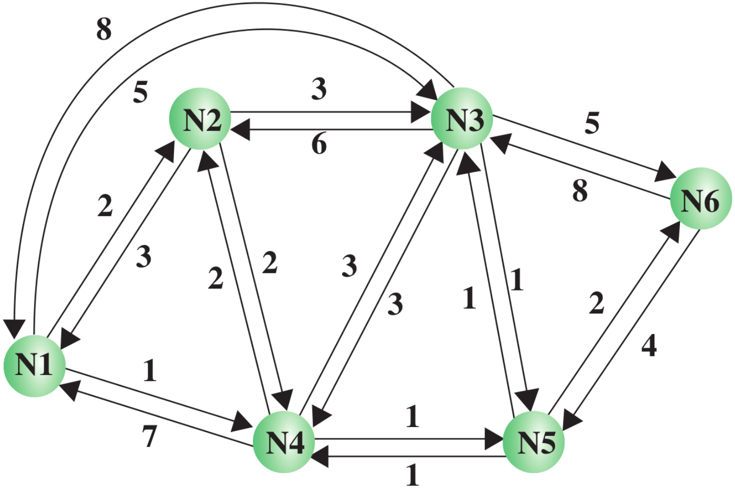

Elements of Routing Techniques

- Which performance criteria are used to select a path?
- When is a path selected?
- Which nodes are responsible for selecting path?
- Which nodes provide information about network status?
  - topology, link costs, current usage
- How often is network status information updated?

Elements of Routing Techniques

Performance Criteria	Network Information Source
Number of hops	None
Cost	Local
Delay	Adjacent node
Throughput	Nodes along route
	All nodes
Decision Time	
Packet (datagram)	**Network Information Update Timing**
Session (virtual circuit)	Continuous
	Periodic
Decision Place	Major load change
Each node (distributed)	Topology change
Central node (centralized)	
Originating node (source)	

Contents

Strategy 1: Fixed Routing

- Use a single permanent route for each source to destination pair
- Routes are determined using a least cost algorithm, e.g. Dijkstra, Bellman-Ford algorithms
- Route is fixed
  - At least until a change in network topology (node/link added/deleted)
  - Hence cannot respond to traffic changes (e.g. overload in one portion of the network)
- No difference between routing for datagrams and virtual circuits
- Advantage is simplicity
  - You assign the routes at the start, and then nothing to do
- Disadvantage is lack of flexibility
  - When the network is operating, changes in load may mean better routes than initially selected should be used

Fixed Routing Example

How many least-cost paths are there? What are they?

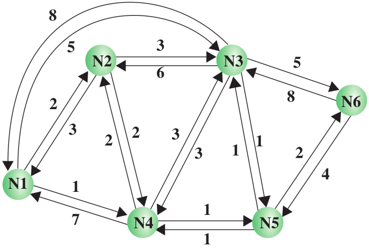

Routing Tables

- A node determines least-cost paths to all possible destinations
- No need to store entire path; store only next node in path (and optionally cost of path)
- Path information stored in *routing table* (or directory)

Destination	*Next*	*Path Cost*
Node_1	Node_x	c_1
Node_2	Node_y	c_2
$\vdots$		

- Routing table may be stored on central node or distributed amongst each node
- Separation of routing and forwarding:

  Routing: strategies, protocols and algorithms are used to create routing table

  Forwarding: routing table used to determine where to send the data to next

Centralised Routing Table Example

Routing table stored on one node

		From Node					
		1	2	3	4	5	6
To Node	1	—	1	5	2	4	5
	2	2	—	5	2	4	5
	3	4	3	—	5	3	5
	4	4	4	5	—	4	5
	5	4	4	5	5	—	5
	6	4	4	5	5	6	—

Distributed Routing Tables Example

Each node has its own routing table

Node 1 Directory

Destination	Next Node
2	2
3	4
4	4
5	4
6	4

Node 2 Directory

Destination	Next Node
1	1
3	3
4	4
5	4
6	4

Node 3 Directory

Destination	Next Node
1	5
2	5
4	5
5	5
6	5

Node 4 Directory

Destination	Next Node
1	2
2	2
3	5
5	5
6	5

Node 5 Directory

Destination	Next Node
1	4
2	4
3	3
4	4
6	6

Node 6 Directory

Destination	Next Node
1	5
2	5
3	5
4	5
5	5

Fixed Routing Summary

- When is a decision made for a route? *At network startup*
- Which node chooses the route? *Centralised or distributed*
- Where does the network information come from? *All nodes*
- When is the network information updated? *Never*
- In practice, only used for small, stable networks (10s of nodes)

Strategy 2: Flooding

- Instead of choosing a route before sending the data, just send the data to everyone
  - A copy of the original packet is sent to all neighbours of the source
  - Each node that receives the packet, forwards a copy of the packet to all of its neighbours
- Advantages:
  - All possible routes are tried; at least one packet will take minimum hop route, e.g. setup a virtual circuit
  - All nodes are visited, e.g. distributing network status (topology) information
  - Simple
- Disadvantages:
  - Inefficient: need to send many copies of packet to get one packet from source to destination
  - Using hop limit and/or selective flooding, packet may not reach destination

Flooding Extensions

1. Don't send back to the node that just sent you the packet
2. Only forward packet once: nodes remember which packets they have forwarded (based on sequence number and source/destination addresses); do not forward a packet if you have previously forwarded that same packet
3. Duplicate detection: each packet has a sequence number, so if destination receives multiple copies of the same packet, it can discard the duplicates
4. Hop Limit: include a "hop counter" in the packet; decrement the counter each time the packet is forwarded—if it is 0, then discard the packet
5. Selective Flooding: send to selection of neighbours. E.g. random, round-robin, probability-based

Flooding Example

Destination is node 6; Hop limit is 3

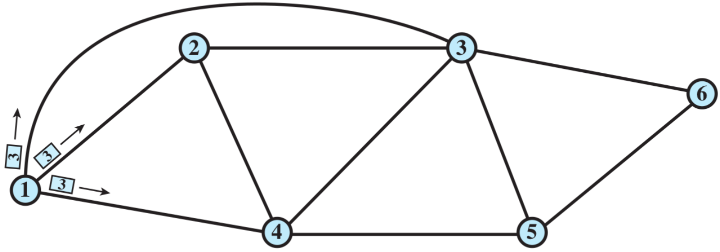

Flooding Example

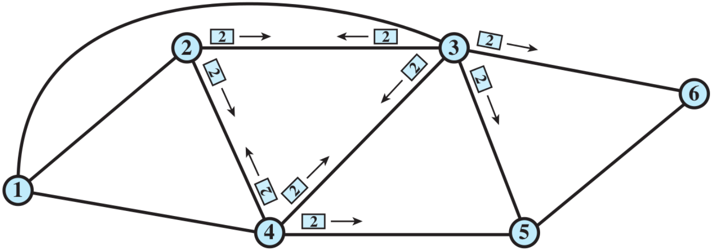

Flooding Example

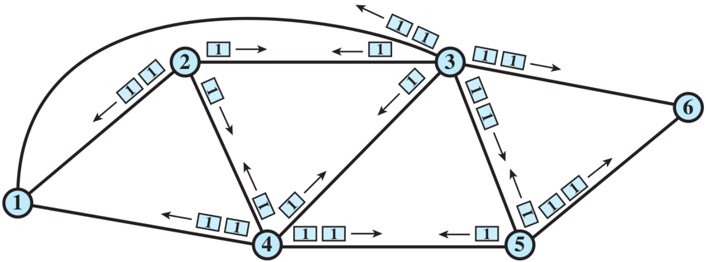

Flooding Example

How many packets transmitted in previous example? How much data is delivered to destination? What if hop limit was 2?

Selective Flooding Examples

Strategy 3: Adaptive Routing

- Use a least-cost routing algorithm to determine a route, and adapt the route as network conditions change
- Used in almost all packet switching networks, e.g. the Internet
- Requires network status information from:
  1. Local to node: route to output link that has shortest queue (rarely used)
  2. Adjacent nodes: delay/link status, then least-cost routing
  3. All nodes: similar to option 2

Characteristics of Adaptive Routing

Advantages

- Improved performance: potentially can select the most suited paths; balance amount of traffic across the network

Disadvantages

- Decisions more complex (complex algorithms needed to select the best path)
- Tradeoff between quality of network information and overhead
  - The more information required for routing decisions, and the more often updates are delivered, then the more overhead in the network
  - Reacting too quickly can cause oscillation
  - Reacting too slowly means information may be irrelevant

Contents

Routing Protocols

- A routing protocol is used by nodes to automatically determine the routes in the network
- A routing protocol specifies:
  - Routing algorithm for determining least-cost routes: e.g. Dijkstra, Bellman-Ford or variants
  - Routing information to be exchanged between nodes
  - Formats of messages used to deliver routing information
  - Rules as to when to send routing messages and what to do upon receiving them
  - Metrics to be used in routing algorithm (hop count, bandwidth, . . .)
  - Optionally, default values of specific parameters may be given
- Real routing protocols include: OSPF, RIP, BGP, IGRP, EIGRP, PNNI, IS-IS, DSDV, AODV, . . .

Example: Link State Routing

- Example: link state routing protocol that uses Dijkstra's algorithm to determine the least–cost routes
- Aim: Each node learns the topology of the network, then calculates the least–cost route from itself to every other node using Dijkstra's algorithm
- Steps for each node:
  1. Record the state of its own links (e.g. source/destination, metric)
  2. Send the state of its own links to every other node by flooding a link state packet
     - identity of the current node
     - list of links that the current node has, including their costs
     - sequence number (used by the flooding protocol)
     - hop count or age (used by the flooding protocol)
  3. Form a shortest path tree from itself to every other node
  4. Build a routing table based on the shortest path tree

Example: Link State Routing

What is the link state packet created by node 1? What is the routing table created by node 1?

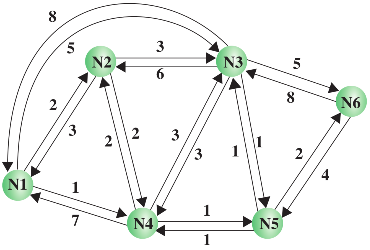

Example: Link State Routing

Summary: Concepts

- Communication networks are formed by connecting devices across multiple links
- Switching is the method of delivering data between source and destination across multiple links
  - Stations or end-user devices act as sources and destinations of data
  - Switches connect the links and forward data between source and destination
- Circuit and Packet Switching techniques determine how to deliver data across one or more paths between source and destination
- Routing determines what path to take between source and destination
- There are different routing metrics, strategies, algorithms and protocols available

Summary: Practice

- Circuit switching was developed for traditional telephone networks and is still used today in those (and other) networks
- Packet switching was developed to be more efficient for delivering computer generated data over networks
- Packet switching is the concept used in the Internet and in almost all new Wide Area Networks (WANs) today
- Adaptive routing strategies are needed for almost all WANs
- Dijkstra and Bellman Ford are two of the most common algorithms for determining the shortest path between source and destination
- The trade-offs between the different routing protocols often depend on the size of networks, the amount of traffic and the rate at which the network changes