Intra-domain Routing

- Distance vector and link state algorithms

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Intra-domain Unicast Routing

- BGP gives information on the path to reach other networks
- How to route traffic within an AS?
  - Single trust domain
    - No policy restrictions on who can determine network topology
    - No policy restrictions on which links can be used
  - Desire efficient routing → shortest path
    - Make best use of the network you have available
  - Two approaches
    - Distance vector – the Routing Information Protocol (RIP)
    - Link state – Open Shortest Path First routing (OSPF)

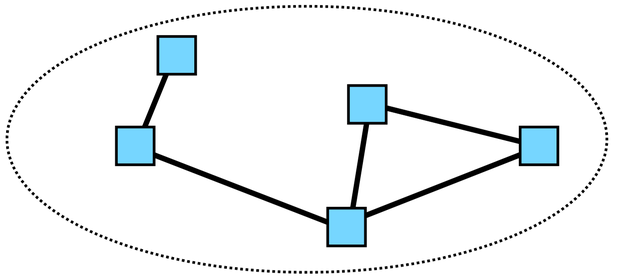

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Distance Vector Routing

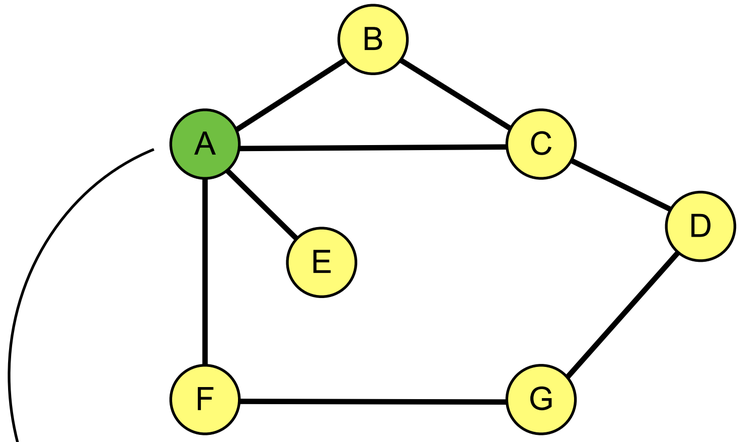

Routing Table at Node A

Destination	Distance	Next Hop
B	1	B
C	1	C
D	2	C
E	1	E
F	1	F
G	2	F

- Nodes maintain a vector containing the distance to every other node in network
  - Periodically exchanged with neighbours, eventually every node know distance and next hop to every other node
- Packets forwarded on shortest path to destination
- Not widely used in practice:
  - Slow to converge in normal operation
  - Count-to-infinity problem makes this worse on failure in networks containing loops

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Link State Routing

- Nodes know the links to their neighbours, and cost of using those links – **link state** information
  - Flood this information throughout the network on startup, and whenever a link is added or removed
    - Send address of node sending the update, the list of directly connected neighbours and costs for link usage, and message sequence number
    - Sequence numbers prevent loops
  - Each node then directly calculates shortest path to every other node, uses as routing table
- Flooding link state data from all nodes ensures all nodes know the entire topology
  - Each node uses Dijkstra's shortest-path algorithm to calculate optimal route to every other node
  - Optimal is assumed to be the shortest path, by cost

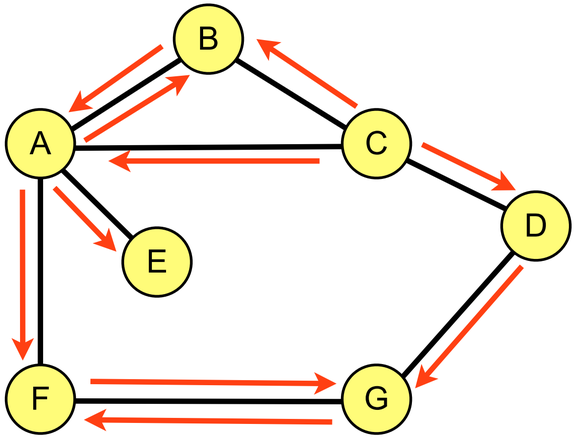

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Distance Vector vs. Link State

Distance vector routing:

- Simple to implement
- Routers only store the distance to each other node: O(n)
- Suffers from slow convergence

Link State routing:

- More complex
- Requires each router to store complete network map: O(n^2)
- Much faster convergence

Slow convergence times make distance vector routing unsuitable for large networks

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Challenges in Intra-domain Routing



- How to rapidly recover from link failures?
  - Construction work surprisingly good at breaking network cables; trawlers do similar damage to undersea cables
  - Good network designs have multiple paths from source to destination
  - Must quickly notice failure and switch to backup path
    - If the link is used for telephony or video streaming, this must happen in $<1/60^{th}$ of a second else it interrupts the speech or video playout
  - **Fast failover** – pre-calculate alternative paths to allow rapid switchover
- How to balance load across multiple paths?
  - **Equal cost multipath routing** – if two paths with the same cost, alternate traffic between them
  - Care needed: TCP uses a triple duplicate ACK to signal loss so is insensitive to small amount of packet re-ordering, but treats large amounts of reordering as loss and slows down
  - RTP applications don't care about order, as long as packets arrive before deadline

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Internet Routing

- Content distribution networks
- Inter-domain routing
- Routing security
- Intra-domain routing

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow