OSPF Overview

Vermont State University

Peter Chapin

Autonomous Systems

- An *autonomous system* is a group of hosts together with their connecting network infrastructure that is under the control of a single administrative mind.
  - Examples include The Vermont State Colleges (VSC), UVM, etc.
- The Internet can be considered a group of interconnected autonomous systems.
  - Some autonomous systems are extensive (e.g., an ISP and its customers)
  - Some are less large (the VSC)
  - All are “significant” in size.
- Every AS gets an autonomous system number (the VSC’s is 54257)

Interior vs Exterior

- An “interior gateway protocol” is a routing protocol used inside an autonomous system.
  - It can deal with fine details…
  - … but it doesn’t necessarily scale to global sizes
- An “exterior gateway protocol” is a routing protocol used to connect autonomous systems.
  - It is concerned with large-scale routing on a worldwide scale…
  - … but lets each AS worry about its internal structure.

Routing Protocols

- Examples of interior gateway protocols
  - Router Information Protocol (RIP)
  - Enhanced Interior Gateway Routing Protocol (EIGRP)
    - Cisco proprietary
    - And yes, there was an “Interior Gateway Routing Protocol” (IGRP), but it’s obsolete.
  - **Open Shortest Path First (OSPF)**
    - It is defined by RFC-2328 (i.e., it is open). That RFC describes version 2, which is the most current for IPv4 only. Version 3 covers IPv6 and is defined by RFC-5340.
- Examples of exterior gateway protocols
  - Exterior Gateway Protocol
    - Yes, that is its name. Obsolete.
  - Border Gateway Protocol (BGP)

OSPF at the Vermont State Colleges

- Tom Maguire, the network engineer at the VSC central office:
  - “We switched from Cisco’s EIGRP to a more vendor-neutral OSPF with a Single Area design.”
  - “Cost values are very important with redundant links at our locations.”

Routing Protocol?

- Two parts…
  - … A network protocol by which routers communicate with each other to trade information they know about the network.
  - … A decision-making procedure (algorithm) by which routers decide how to forward packets to distant parts of the network.
- Routing Table
  - For every destination network, which interface should be used?
  - Incoming packets are copied to the appropriate interface depending on the destination address.

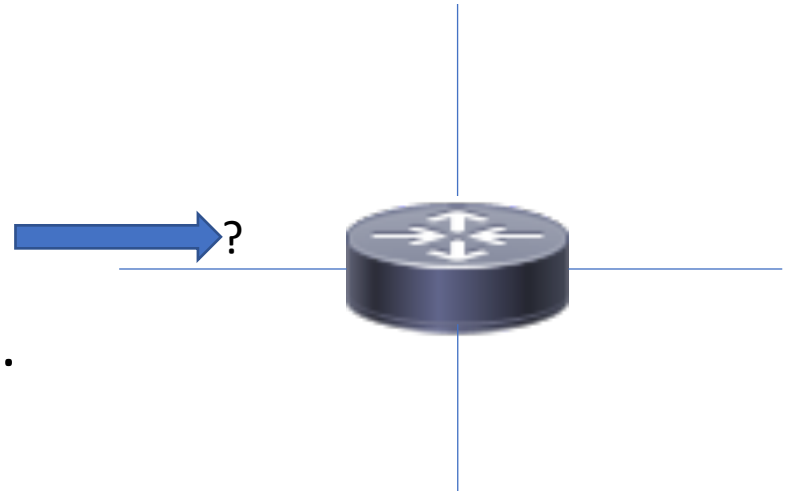

OSPF Overview

- Link-state protocol (not distance-vector)
  - Routers exchange information about links they know about.
  - Eventually, all routers in an *area* will learn about all links in that area.
- Routers build a graph representing the area in which they are located
  - Compute the *shortest path tree* in that graph rooted at themselves using Dijkstra's Algorithm.
  - Using the shortest path, compute a routing table to all known links (in the area).
  - Path distance is based on "cost" values assigned by the administrator.
- If a change occurs, the graphs are rebuilt, etc.

Example Network

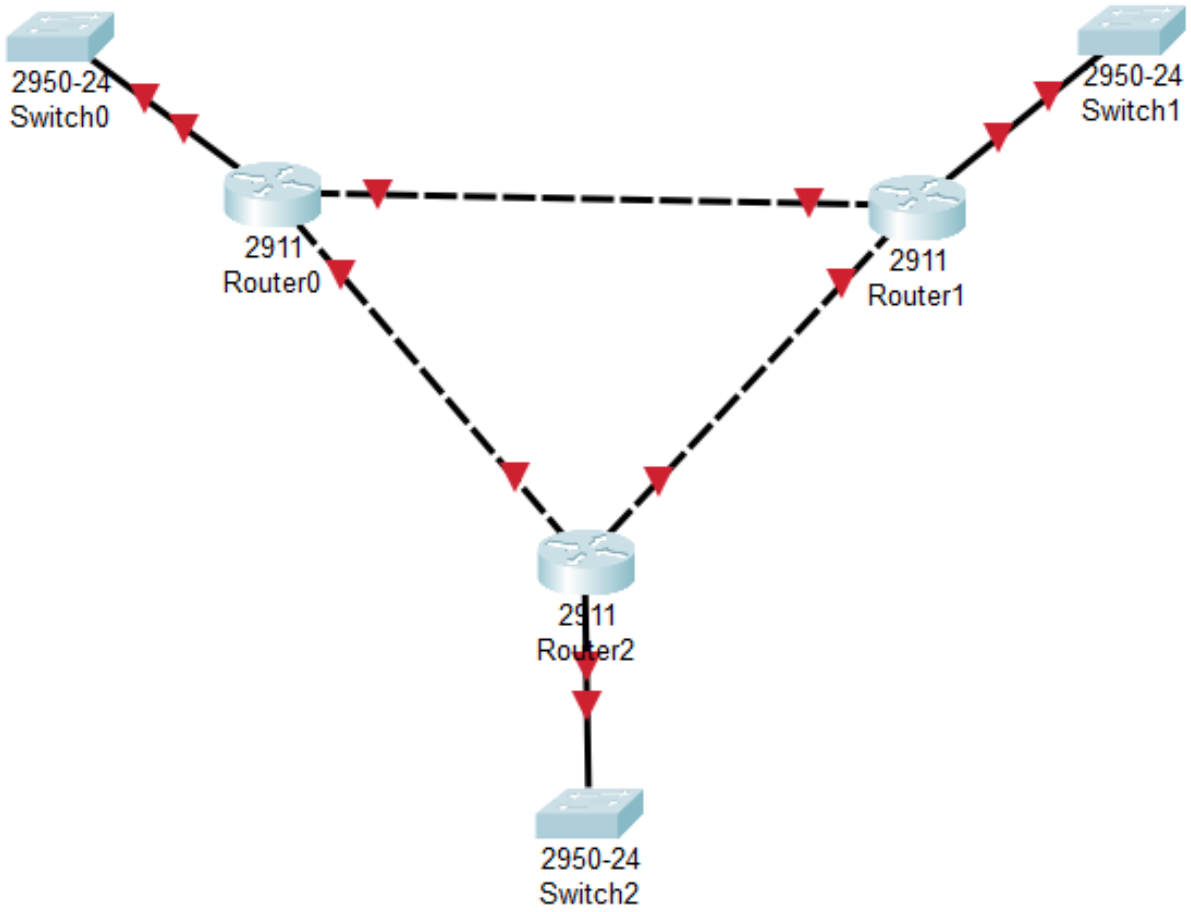

Example Network Graph

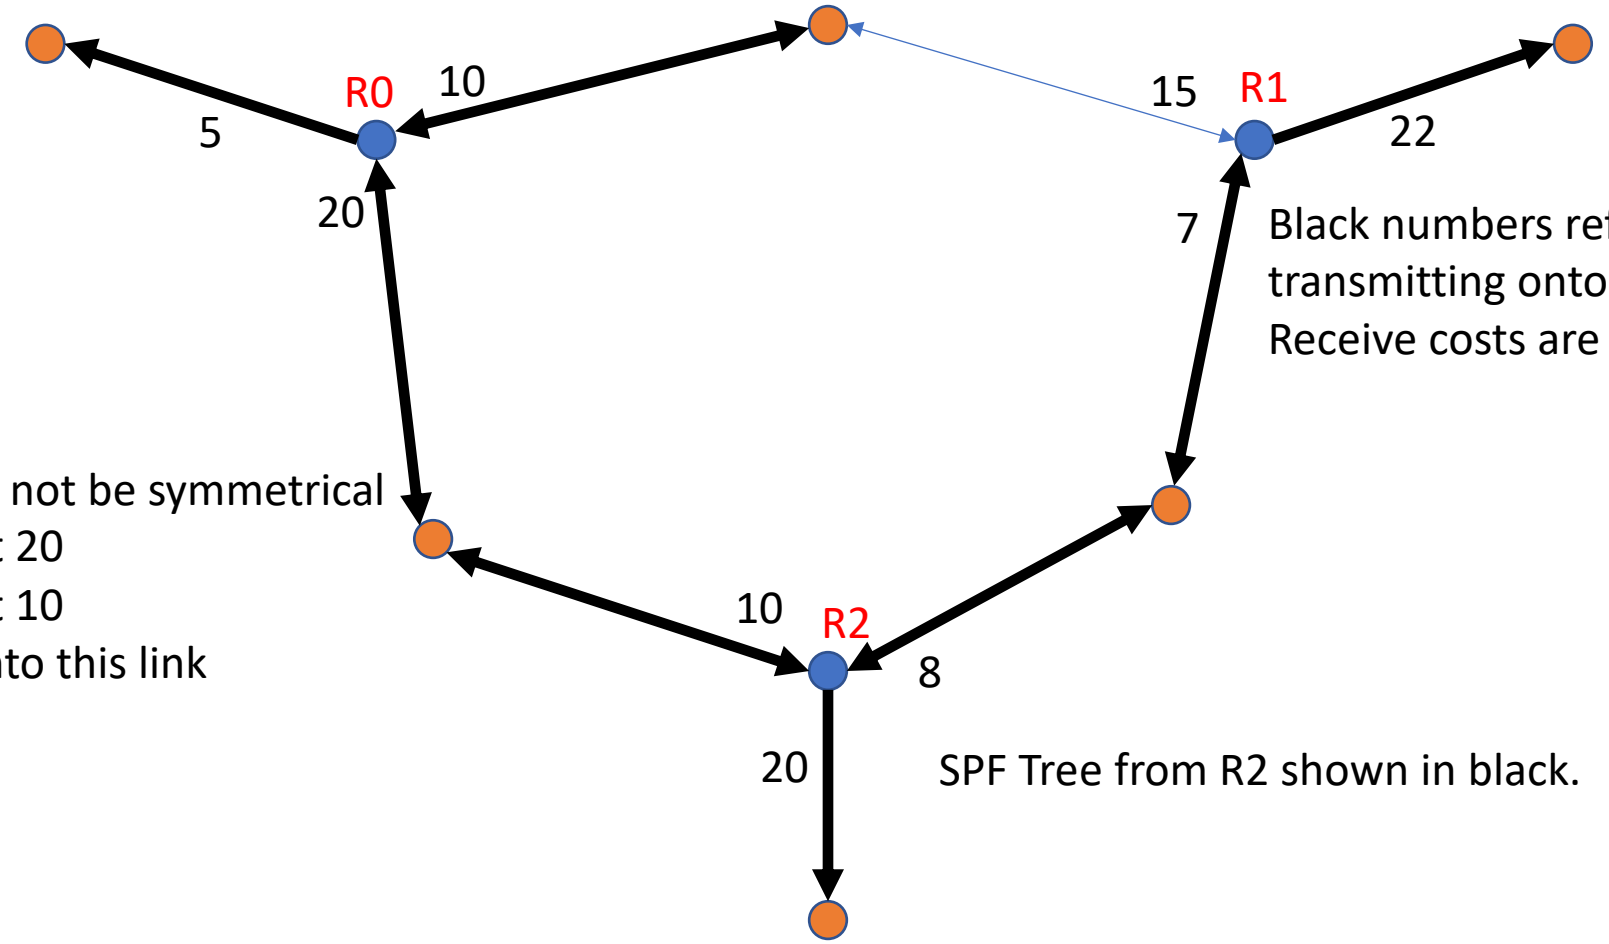

Costs need not be symmetrical
R0 has cost 20
R2 has cost 10
… to get onto this link

Black numbers reflect cost of transmitting onto a link.
Receive costs are always zero with OSPF

SPF Tree from R2 shown in black.

Some Notes

- Once OSPF has decided on a route, it is committed to that route until there is a change in the link state somewhere (i.e., a link goes down, costs change)
  - There is no accounting for congestion (unless it gets reflected in link cost).
  - No attempt at load balancing (e.g., splitting traffic 60/40 based on cost weighting)
- However…
  - A change in link state causes routers to flood the change. Then, all routers recompute the graph, re-execute Dijkstra's Algorithm, and recompute routes.
  - The network adapts to such changes "quickly" with no manual reconfiguration.

Single Area

- In *Single Area* OSPF…
  - All routers know about all other links
  - All routers build the same graph
  - Each router computes its own shortest path tree (rooted at itself)
  - Each router computes its own routing table based on its shortest path tree

Multi-Area OSPF

- In *Multi-Area* OSPF…
  - Routers are grouped into multiple “areas,” each with an *area number*.
  - Some routers have interfaces in several areas. They are called *area border routers*.
  - Routing is done in two levels. First, send the packet to the right area (i.e., to an area border router). Then, route the packet inside that area.
  - Routers only build graphs and trees for their area and do not “know” about links in other areas.
  - This cuts down on the information that must be flooded over the network and improves scalability.
  - Area 0 is special. It is the *backbone area*. All other areas must connect to it.