ICS 351: Today's plan

- OSPF, continued
- BGP
- Routing in general

link-state routing game

- 1. prepare a list of all neighbors and the links to them (HELLO protocol)
- 2. make a copy for every router in the network
- 3. distribute a copy to every router in the network (by sending it to the neighbors, and letting them distribute it)
- 4. build the network map
- 5. find the shortest path to each router
- 6. build the routing table

Flooding Link State Advertisements

- each router is responsible for distributing copies of every link-state advertisement that it gets
- if there are loops, this means each router will get multiple copies of each LSA
- so LSAs are only forwarded if they are new
- LSAs are acknowledged to the sender, and LSAs are resent if there is no acknowledgement, so that the transmission is reliable

OSPF areas

- all areas are connected to the backbone area
- all routing information is disseminated over the backbone area
- routers in OSPF play different roles, for example a backbone router is connected to the backbone, an area border router is connected to more than one area (and is usually also a backbone router), and an internal router is only connected to routers in the same area
- every area has one Designated Router which receives then rebroadcasts link-state updates
- defined in RFC 2328 (and in RFC 5340 for IPv6)

RIP compared to OSPF

- both RIP and OSPF find optimal paths
- OSPF generally finds them much more quickly
  - especially when links go down
- RIP generally sends less data (lower overhead)
- OSPF is more complex: more configurable, more code

More OSPF details

- each router has a list of all link states
- an algorithm such as Dijkstra's shortest path algorithm can be used to build a directed acyclic graph (DAG) with the router at the root, and all other networks reachable through the DAG
- multiple equal-cost paths can be used for each destination
- OSPF supports authentication among routers (null authentication is an option)
- link-state advertisements expire if they are too old

OSPF areas

- areas can be used in larger networks to minimize the amount of information exchanged among routers
- routers outside an area don't have all the link state information of routers inside the area
- areas form a 2-level hierarchy with the backbone at the root, and all the other areas below it

BGP brief summary

- Border Gateway Protocol is an Exterior Gateway Protocol (EGP) used to route between Autonomous Systems
- BGP uses a variant of distance-vector routing in which the entire path to reach a destination is distributed
- BGP might use different criteria for advertising routes and for using routes
- these criteria may be set by the network administrator to define policy
- BGPv4 is defined in *RFC 4271*

some BGP properties

- a router configured to run BGP is a BGP speaker, as opposed to other routers which might not run BGP
- BGP uses TCP for reliable transmission of data, so timeouts can be faster

Why routing matters

- the routing protocol builds the routing tables, which are essential to correct routing of packets
- incorrect routing tables can lead to packets:
- o being dropped
- o being sent in a loop
- o being sent over slow links
- o being sent over congested links
- all of these cause network "malfunction", even when the hardware is working well!

routing responsibilities

- establish routes to destination networks
- maintain routes in the face of changing configuration: link loss, router loss, new links, new routers
- be trustworthy, do not advertise routes to which we don't know how to deliver

routing possibilities

- blacklist attacking hosts/networks as close to the source as possible
  - gateways for non-transit networks can drop packets with spoofed source address not from the network
- route depending on packet (flow) type, e.g. low latency, high throughput, constant bit rate (CBR)
- smart routing, based on packet content (may conflict with net neutrality)

what routing is not:

- Ethernet switching does not use IP addresses in any way, and only connects Ethernet segments with the same network number
- Network Address Translation (NAT) allows a single IP address to be used as a "front end" for a number of systems that use TCP/IP or UDP/IP (a home “router”)
- A firewall blocks access to most TCP/UDP ports, only allowing selected ports to connect to or from the outside world
- each or all of these may be combined in the same box with a routing function, but they are logically separate