CS 168 Introduction to the Internet: Architecture and Protocols
Fall 2022 Sylvia Ratnasamy

Discussion 7: Midterm Prep

1 Link-State Routing

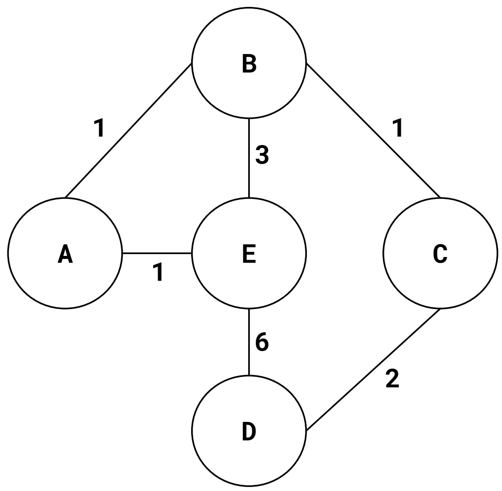

(1) E begins to calculate its shortest paths to all other routers. Fill in the following table that contains routing information for E. Columns represent destination routers. You may use any shortest path algorithm (i.e. Dijkstra's) to calculate shortest paths.

	A	B	C	D
Shortest path (i.e. E → A → B)	E → A	E → A → B	E → A → B → C	E → A → B → C → D
Path cost	1	2	3	5

(2) Assuming all other routers have also computed their shortest routes, what path does a packet take from E to D?

Solution: E → A → B → C → D

(3) Assume that the cost of the link between C and D suddenly increases from 2 to 20, and routing has not yet re-converged. Is it possible that packets now enter a loop?

Solution: Yes. As an example, C and D will both update their routing information immediately, but the topology change will not immediately reach B. Therefore, if B tries to send a packet to D, the packet will be sent to C, which will send the packet back to B, and so on and so forth.

2 IP Header

IPv4 and IPv6 have very different headers. In particular, IPv6 dropped many of the fields that IPv4 headers had. Why did IPv6 drop the following fields from its header?

(1) Checksum:

Solution: It is sufficient to offload this responsibility to the end-host. Corrupt destinations/source addresses result in hosts ignoring traffic. Corruptions in the TTL field may cause the packet to fail to reach the destination, or potentially travel longer than necessary.

(2) Header Length:

Solution: The IPv6 packet header is constant length, there is no need to include it. The "Next Header" field tells you about any additional information needed in order to consider options and such things.

(3) Fragmentation:

Solution: Again, it is sufficient to offload this to the end hosts. Hosts can proactively send small enough packets by using Path MTU discovery, a technique for exploring the MTU across a given path. Otherwise, routers will send a "packet too big" message back to the sending host. This reduces the work to handle reassembly and only adds (at most an RTT) delay for sending the first packet in a stream. This also reduces a whole class of side-channel attacks that exploit the fragmentation mechanism in IPv4 (host buffers, ID fields, etc.).

(4) IPv6 also hosts much larger address fields (128 bits). Is that enough bits, or might some future IPvN need to expand these fields further?

Solution: 128 bits is plenty for the potential future of the human race. $2^{128} \approx 10^{38}$ addresses.

It has been (over)estimated that there will be 50 billion internet connected devices by the year 2020. Lets continue that trend to the year 2100 with an estimate of 50 trillion internet connected devices on Earth. Lets also suppose that humanity has conquered many solar systems, and controls 1000 planets, each with as many devices as Earth.

This means there would be about $1000 \cdot 50 \cdot 10^{12} = 5 \cdot 10^{16}$ internet connected devices on the inter(planetary)net.

$100\% \cdot \frac{5 \cdot 10^{16}}{10^{38}} \approx 5 \cdot 10^{-20}\%$ address space use. IPv6 gives us more than enough addresses.

3 Putting it All Together

Host A in AS 1 wants to communicate with host B in AS 2. This problem will walk through many of the various steps necessary for this to happen.

3.1 BGP

Consider the network topology below. Assume all ASes are following Gao-Rexford policies. Arrows point from providers to customers (e.g. AS X is a provider to AS 1).

Solid lines denote peering (e.g. AS 1 is peered with AS W)

Dotted lines connect individual hosts to their ISP's AS (e.g. A's ISP is AS 1)

Gao-Rexford Rules

- Export policy: To customers export all routes. To peers and providers only export routes to customers.
- Selection policy: pick routes in the following order: Customers Peers Providers.

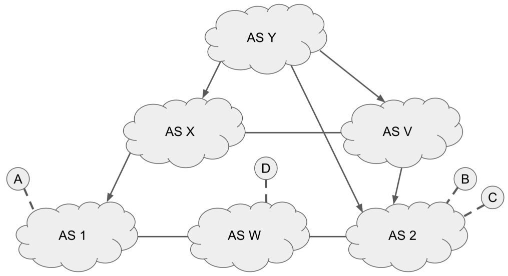

(1) What path does AS1 use to reach AS2?

Solution: AS 1 - AS X - AS V - AS 2

(2) AS Y would export routes for which hosts to AS X?

Solution: Hosts B and C. AS Y does not export a route to host D because it never learns about host D.

(3) AS 2 is a customer of both AS V and AS Y. What is the general term for having multiple providers in this fashion?

Solution: Multihoming

3.2 Address Allocation

Assume that AS X has the prefix 1.2.240.0/16 and plans to give some portion of this address space to its customers (AS 1).

(1) AS 1 wants to be able to support up to 1,000 hosts in its AS. What is the maximum prefix length AS 1 could use?

Solution: 22. To support 1,000 hosts, we will need at least 2^{10} addresses available $(32-10=22)$

(2) Assume AS X has another customer AS to which it allocates the block 1.2.248.0/24 (not pictured). How many hosts can this AS support?

Solution: 254. There are 2^{32-24} addresses in this block, but the first and last are reserved.

(3) Will any of these allocation decisions for AS X's customers affect the routing tables for AS Y (AS X's provider)?

Solution: No, either way, AS Y will have one route entry for AS X's customers using its overall prefix.

3.3 BGP Selection

Suppose that the AS topology is the same as question 2.1 except Assume ASes X and V no longer peer, and that AS W becomes another provider to poth AS1 and AS2.

Suppose that the internal network topology of AS1 is as follows

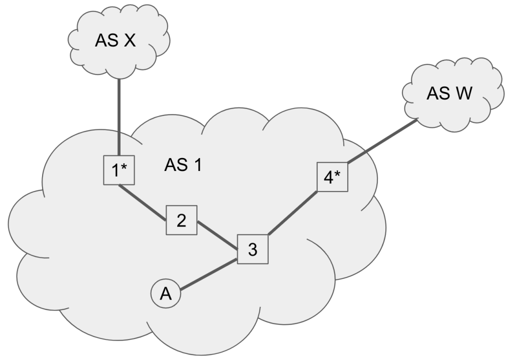

The goal of this question is to test your understanding on BGP and its connection to routing algorithms.

Note that the *LocalPrefs* represent preferences for AS paths. The paths listed in the table have been abbreviated to the last hop AS. **The table contains exports for destination B.**

Routers	*LocalPref*	*ASPATH*	*IGP path*	*MED*
1	ASX:5, ASW:1	-	-	-
2	ASX:5, ASW:1	-	-	-
3	ASX:3, ASW:2	-	-	-
4	ASX:2, ASW:3	-	-	-

(1) Suppose that host A is connected to 3 and wants to send a packet to B. Which path does the packet take?

Solution: Based on router 3's local pref, we would pick a path to ASX, therefore the packet would go from $A \rightarrow 3 \rightarrow 2 \rightarrow 1$. If A was connected to 4, then A would send a packet from $A \rightarrow 4$.

(2) Suppose now that all routers have the same local pref for each autonomous service and A is still connected to router 3. Given the ASPATH's below which path would A's packet to B take? **NOTE**: the ASPATH routing state has not converged yet.

Routers	*LocalPref*	*ASPATH*	*IGP path*	*MED*
1	ASX:5, ASW:5	ASX-ASY-AS2	-	-
2	ASX:5, ASW:5	ASX-ASY-AS2	-	-
3	ASX:5, ASW:5	ASW-AS2	-	-
4	ASX:5, ASW:5	ASW-AS2	-	-

Solution: Packet would go through ASW-AS2.

(3) Now suppose that all routers have the same local pref, the ASPATHs have converged, and each link has a cost of 1. Which AS does the packet go through?

Solution: It goes through ASW because 4 is the closest router according to the forwarding table produced by distance vector.

3.4 Learning Switches in AS2

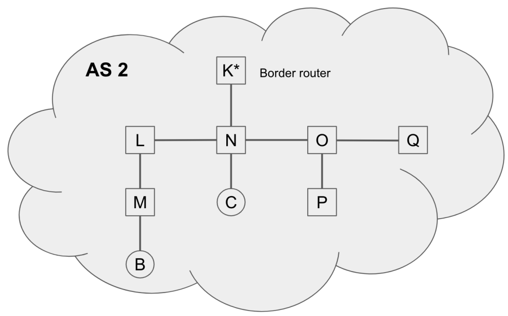

For this question, K is AS2's single border router connecting to all other ASes from part 1. Assume that internally, AS2 is running the learning switch algorithm for IGP. All routers in the network start with empty forwarding tables.

(1) K forwards A to B, which routers receive this packet?

Solution: All routers receive this packet since routers that don't have the destination address of the packet in their forwarding table flood that packet

(2) C sends a packet to B, which routers receive this packet?

Solution: All routers by the same reasoning as the previous question

(3) B responds to A, which switches receive this packet?

Solution: M $\rightarrow L \rightarrow N \rightarrow K$