15-744 Computer Networks — Spring 2017
Homework 1
Due by 2/20/2017, 10:30am

(to be submitted in the beginning of class in hard copy)

Name:

A Short Answer

1. Why is Internet an internetwork (rather than one massive, homogeneous network)? Give at least two reasons.

> **Solution:**
>
> - This allows the use of different link-layer technologies (Ethernet, Wi-Fi, LTE, etc.).
> - Network management is simplified because each "real-world entity" (e.g., CMU) can control its own network.
> - Fault tolerance is improved because many failures can be confined to a single network.

2. Explain how the NAT on the router in your home network violates fate sharing.

> **Solution:** If your router is rebooted, you will lose all of your connections to the outside world even though neither your machine nor the outside hosts to which you were connected failed.

3. Please list two more goals that you think are important but have not been considered in the "Design" paper.

> **Solution:**
>
> - Security
> - Quality of service
> - Support for mobility
> - Easy to deploy new services/protocols
> - (Any reasonable answers will be accepted)

4. Why does DNS service use UDP, instead of TCP? Doesn't DNS need reliable transfer?

Solution:

- UDP is faster, since it doesn't have to build a connection before sending a request. It also reduces the burden of a DNS server because no connection state needs to be maintained.
- DNS can still implement reliability features in the application layer. In the simplest case, it could use the stop-and-wait approach, since a DNS request is usually so small that fits in a single UDP packet.

5. The "4D" paper presents an architecture with a centralized controller, where the "Onix" paper implements a distributed control plane. Discuss the tradeoffs in the different design.

Solution:

- Pros for a single (centralized) controller: Easy to implement; Easy to update policies; The control node has global view of the network thus it's easy to make optimal decisions.
- Pros for a distributed control plane: Avoids single point of failure; Scales well to large networks; React faster to change.

B Intra-domain routing

6. Link-state and distance-vector are two popular algorithms for intra-domain routing. Why can't they be used in inter-domain routing?

> **Solution:**
>
> - These algorithms don't scale well.
> - Inter-domain routing has to taken into account many metrics that are not considered in these two algorithms, such as business relationships. As a result, the inter-domain routing algorithm has different rules in deciding what to publish and which route to select.

7. Consider a simple network topology as shown in the figure below. The number on each line represents the cost of the link.

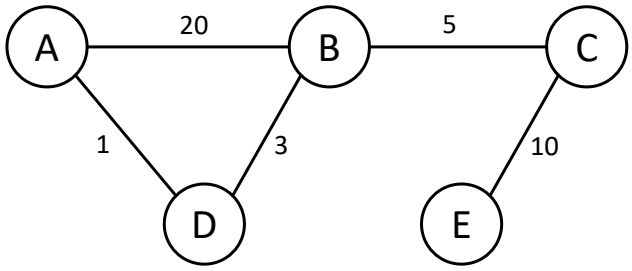

(a) Complete the table below for routing table on node A when the network is stable.

Destination	Cost	Next hop
B	4	D
C	9	D
D	1	D
E	19	D

(b) Using the basic distance vector algorithm, what will happen when the link between B and D is down (the cost between them becomes infinity)? Assume each node broadcasts its routing information every t seconds, how long does it take the routing tables on different nodes to become stable again?

> **Solution:** The above situation will have the "count to infinity" problem and it will take a while for the network to become stable again. Eventually, A will use B as the next hop to reach B, C, and D, and D will need to go through A to reach the other part of network.
>
> From the perspective of A and D, The average cost to reach B will increase by 1 every t second. A initially has a cost of 4 to reach B, and the network won't be stable until this cost reaches 20. So it will take about $16t$ to $17t$ seconds for the network to stabilize.

(c) One partial solution to the above problem is *split horizon with poison reverse (see Wikipedia link)*. Identify a scenario using the above network topology when this approach would fail. How can you solve the problem then?

> **Solution:** It will fail when the link between C and E is down, and A, B, and D form a loop to cause the "count to infinity" problem. The solutions can be to either

- Use other completely different protocols (e.g. link state),
- Set "infinity" in the network to be a smaller number (e.g. 16) so that the network can stabilize more quickly, or
- Each node not only advertises next hop, but also complete path to reach different destinations, which then becomes the path vector protocol.

C Inter-domain routing

8. The following figure shows 4 AS paths in the network.

   - Path 1: $1 \to 3 \to 5 \to 9$
   - Path 2: $1 \to 3 \to 6 \to 7 \to 4 \to 2$
   - Path 3: $9 \to 5 \to 4 \to 2$
   - Path 4: $2 \to 4 \to 8$

   In this network, AS 4 is the only tier-1 ISP. There is a copy of the figure without paths at the end of the question which you can use as a working area.

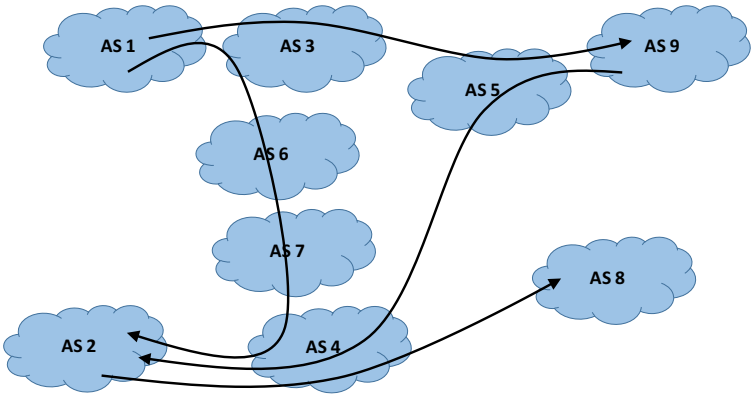

(a) Assume in the network there is only customer-provider relationship. Please mark the edges in the figure below the customer-provider relationship (customer $\to$ provider).

Solution: See figure below. The arrow between AS3 and AS5 can be either way.

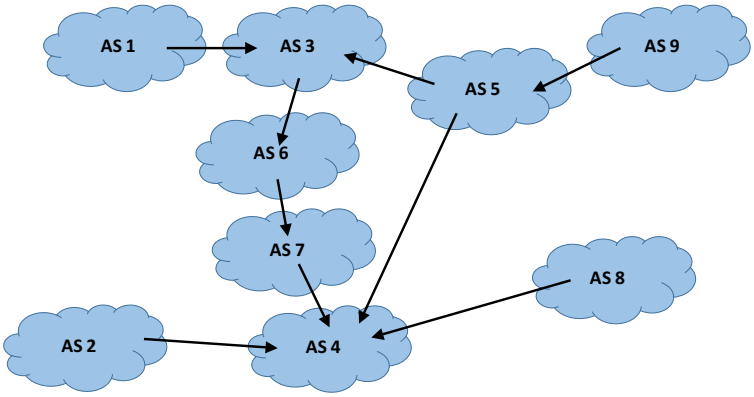

(b) Assume there is also peering relationship in the network, please identify one possible peering relationship.

Solution: AS 3 and 5.

(c) Explain why in the network, AS 4 is most likely the top provider.

Solution: AS4 has the highest degree.

(d) Topologically, there is a shorter path between AS 1 and AS 2 through $1 \to 3 \to 5 \to 4 \to 2$. Do you think, under the existing customer-provider relationship, BGP should use this path, why or why not?

Solution: No, because it violates the valley-free property.

(e) Topologically, there are two paths between AS9 and AS 8, $9 \to 5 \to 4 \to 8$, and $9 \to 5 \to 3 \to 6 \to 7 \to 4 \to 8$. Are they both valid under valley-free BGP? If they are not, please explain why. If they are, please explain which is in the best position to enforce one path, not another and how does it enforce?

Solution: Yes, they are both valid. AS 5 is in the best position. For instance, It can enforce $9 \to 5 \to 4 \to 8$ by exporting the route to AS 9 only to AS 4, not AS3.

D Router Design

9. The iSLIP scheduling algorithm matches input and output ports for a given time slot. Consider the router state shown in the figure below, and answer the following questions.

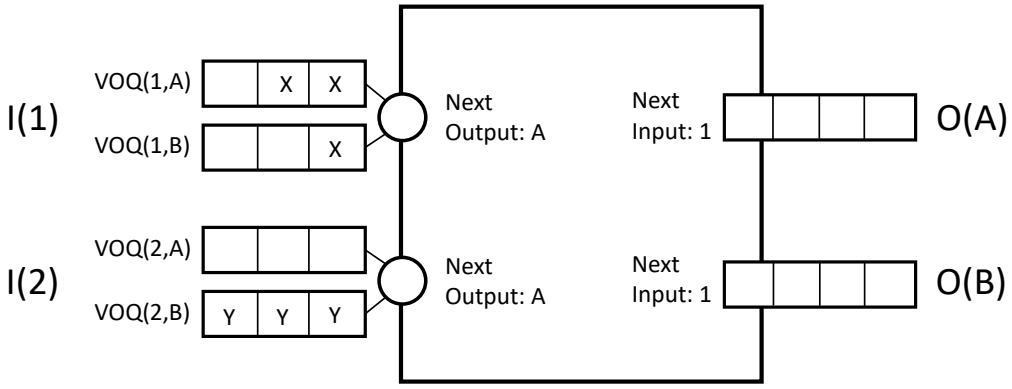

(a) For the next time slot, how many iterations does the iSLIP algorithm take to complete?

Solution: 2

(b) After the next time slot, what are the pointer values at the input and output ports?

Solution:
I(1): next output is B
I(2): next output is A
O(A): next input is 2
O(B): next input is 1

(c) Please draw on the figure what the output packet sequence will be like once all input packets have been sent.

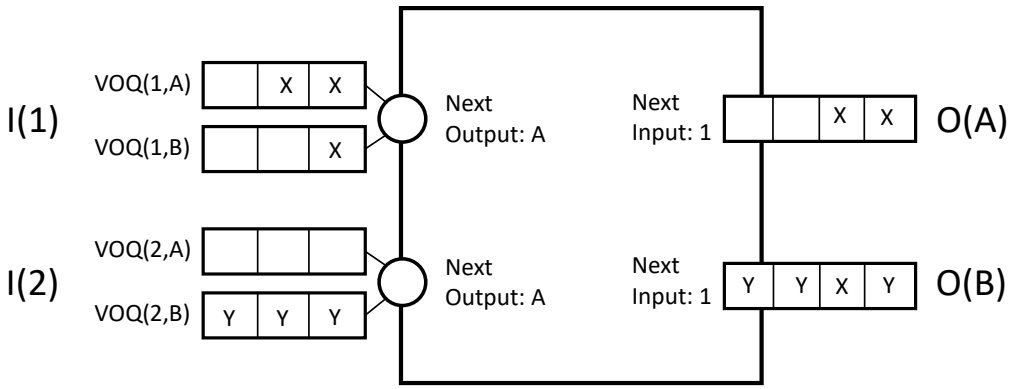

(d) For a router with N output ports, what is the maximum number of iterations iSLIP could take to complete?

Solution: It would take N iterations at maximum.

E Basic Tools

10. Take a look at the man pages for `traceroute` to answer the following questions.

    Perform a `traceroute` on www.berkeley.edu at three different hours of the day (submit the times as part of the homework).

    1. Find the average and standard deviation of the round-trip delays at each of the three hours.

    2. Find the number of routers in the path at each of the three hours. Did the paths change during any of the hours?

    3. Try to identify the number of ISP networks the `traceroute` packets pass through from source to destination. Routers with similar names and/or similar IP addresses should be considered as part of the same ISP. In your experiments, do the largest delays occur at peering interfaces between adjacent ISPs?

    4. Repeat the above for a destination on a continent different then the source. Compare the intra- and inter-continent results.

    5. What kind of problem do you expect to be able to solve using `traceroute`?

       Solution: It can be used to localize reachability failure points or bottlenecks.