Homework 1

This assignment is due in class on **Tuesday, January 28**. It covers readings and lectures through Thursday, January 23. Late homeworks will not be accepted.

Write your name, your e-mail address, and the date on the paper that you hand in.

1. MASS-MARKET BUSINESSES AND NETWORK EFFECTS (10 points)

A. ***(6 points)*** Explain briefly the difference between a *mass-market* business and a business that exhibits *network effects*.

B. ***(4 points)*** Give two examples of mass-market businesses that were discussed in class, one of which relies much more heavily on network effects than the other; identify which of the two relies much more heavily on network effects.

2. INTERNET ARCHITECTURE AND PROTOCOLS (10 points)

A. ***(2 points)*** What type of number must your browser translate the name **`www.microsoft.com`** into so that you can view a page on the Microsoft website?

B. ***(4 points)*** If the correct translation of **`www.microsoft.com`** is not already known to your browser, which Internet service is used to look it up? Briefly explain the lookup process.

C. ***(4 points)*** For one point each, identify the Internet layer that is most closely associated with each of:
(i) Retransmission of dropped packets
(ii) Congestion control
(iii) Interdomain routing
(iv) Ethernet connection

3. OSPF ROUTING (22 points)

Open Shortest Path First (OSPF) routing is used to determine best routes within a single administrative domain. Because the routers involved are in the same domain, they can work together towards a common, simple goal, which is choosing the *lowest-cost path* available. In addition, because domains are relatively small compared to the Internet, it is possible to share more information about the network without causing congestion. Thus, OSPF works by sending information about the network topology throughout the network so that each node can calculate the lowest-cost path to reach a destination.

For example, consider a node ***A*** that has two neighbors:

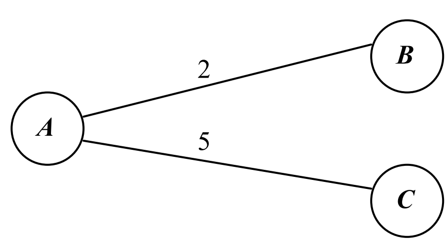

A would then broadcast the following *link-state packet* (LSP):

Link	Weight
A→B	2
A→C	5

Suppose ***A*** receives the following LSPs from other nodes:

Link	Weight
B→A	2
B→C	1
B→D	3

Link	Weight
C→A	5
C→B	1
C→D	4

Link	Weight
D→B	3
D→C	4

From this information, we can deduce the topology of the network, calculate the shortest paths to the other nodes, and construct a forwarding table. Here is what the whole network looks like, based on the above:

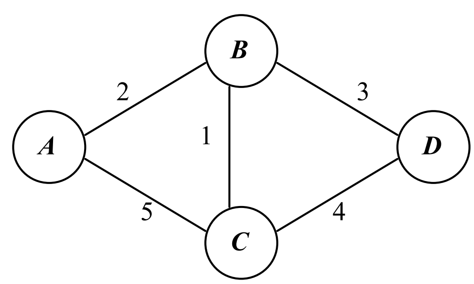

The shortest paths (paths of minimum total weight) from ***A*** to the other nodes are:

Destination	Path	Total Cost
A→B	***A→B***	2
A→C	***A→B→C***	3
A→D	***A→B→D***	5

Then the *forwarding table* for ***A***, which indicates where next to send packets for a given destination, is:

Destination	Next Hop
B	***B***
C	***B***
D	***B***

When ***A*** receives updated LSPs from new or existing nodes on the network, it recalculates its forwarding table to reflect the new shortest paths.

Now consider a new network in which node ***A*** has three neighbors:

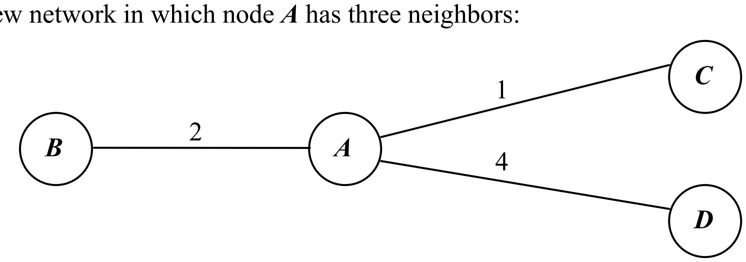

and ***A*** receives the following LSPs:

Link	Weight
B→A	2
B→E	2

Link	Weight
C→A	1
C→E	3

Link	Weight
D→A	4

Link	Weight
E→B	2
E→C	3

A *(4 points).* Give the forwarding table for ***A***.

B *(5 points).* Assume that a new node, ***F***, joins the network. ***A*** receives the following updated LSPs:

Link	Weight
C→A	1
C→E	3
C→F	1

Link	Weight
D→A	4
D→F	1

Link	Weight
E→B	2
E→C	3
E→F	3

Link	Weight
F→C	1
F→D	1
F→E	3

Give the new forwarding table for ***A*** (including the new entry for destination ***F***).

C *(13 points).* Assume that the link between ***A*** and ***C*** fails and must be withdrawn. Assume further that the link between ***A*** and ***D*** becomes less congested, and so both ***A*** and ***D*** recalculate the cost of the ***A↔D*** link to be 1.

(i) Give the updated LSPs that ***A***, ***C***, and ***D*** send *(3 points)*.
(ii) Give the updated forwarding table for ***C*** *(5 points)*.
(iii) Give the updated forwarding table for ***D*** *(5 points)*.

4. BUSINESS ON THE INTERNET (12 points)

In **Trust and Risk in Internet Commerce**, Camp identifies the following stages of a commercial transaction: account acquisition, browsing or discovery, price negotiation, payment, merchandise delivery, dispute resolution, and collection and final settlement. For simplicity, restrict your attention in this question to business-to-consumer transactions (*e.g.*, those on Amazon) or consumer-to-consumer transactions (*e.g.*, those on eBay), and ignore business-to-business transactions. Choose three of these stages and, for each, state whether you think that the use of the Internet enables a unique type of consumer experience in that stage. If you think that it does, briefly explain what is unique about the consumer experience and what property of the Internet enables this unique experience. If you think that it does not, give another example of a communication system or environment in which a consumer can have a similar experience. *(You will get 4 points for each correct answer.)*

5. INFORMATION ECONOMY (12 points)

Recall that information goods often have high fixed costs and low marginal costs; furthermore, it is often the case that anyone (*i.e.*, not only the creator or authorized manufacturer) can reproduce them. Briefly describe three techniques that can be used to build effective information businesses despite this fundamental fact about information goods. *(You will get 4 points for each correct answer.)*

6. BGP ROUTING (34 points)

The *Border Gateway Protocol (BGP)* is used to determine best routes between sources and destinations that are not in the same *autonomous system (AS)*, *i.e.*, administrative domain. Because paths often traverse networks managed by separate companies, route choices are often determined by a combination of factors, possibly including specific local preferences at each router, business relationships with neighbors, and path length. BGP accommodates these complex routing policies that are not uniform throughout the Internet.

Because BGP performs *interdomain routing* (*i.e.*, between domains), it is primarily concerned with reaching an AS containing a specific host destination, not with reaching the host itself. Once the packets are routed to the AS containing the destination, the job of getting them to the specific host is left to *intradomain routing* (*i.e.*, within a domain). For example, to reach the host 192.1.12.14 from some far-away AS, it is enough to know that 1) this host is contained within some AS, say #1702, and 2) the far-away AS has a route to AS1702. The routers within AS1702 can then route the packet to host 192.1.12.14.

For the purposes of interdomain routing, think of an AS as one node in a network. The nodes connected to it are its *neighbors*. Routers in neighboring ASs are physically connected, so that packets can be passed between them.

Here is how information about hosts gets transmitted through the network:

1. An AS "advertises" to its neighbors the hosts contained within it.
2. When an AS receives information about a destination from its neighbors, it looks at its "best" path to that destination. If there is no path to that destination or the newly received path is better, it records the new path as the "best" path.
3. Deciding the "best" path involves considering many factors and is done differently by each AS.
4. Whenever an AS's "best" path changes because of an update, it announces the change to its neighbors if its routing policies permit.

Consider the following hypothetical example.

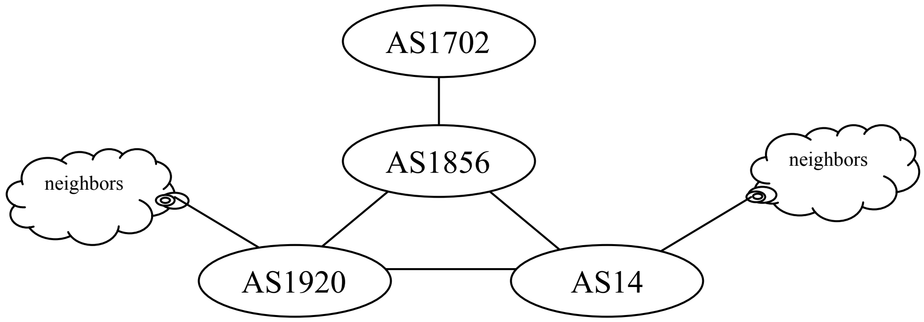

1. AS1702 announces the destinations 192.1.12.0 through 192.1.12.255.
2. Its neighbor AS1856 receives this information for the first time and updates its table to reflect that it has a direct connection to AS1702. It then announces that information to its neighbors, AS1920, AS14, and AS1702. (AS1702 already has a path to itself; so it ignores this new information.)
3. AS1920 and AS14 update their tables to reflect the path to AS1856 and then AS1702. These ASs advertise this path to their neighbors.
4. Because AS14 and AS1920 are connected, AS14 receives a path through AS1920 in addition to the one it already has through AS1856. Suppose its policies tell it to prefer paths through AS1920. It will then change its table to reflect the path AS14, AS1920, AS1856, AS1720 and will advertise this to its neighbors. (Note that cyclic paths are not allowed; so AS1856 can't choose the path AS1856, AS14, AS1920, AS1856, AS1702.)
5. This process continues as ASs throughout the network decide their best paths to AS1702. Notice that routes can keep changing as long as new information is received from neighboring ASs.

I. Now consider the following AS graph, and consider routes to a destination in AS1. Here are the individual AS's routing policies (note that *shortest path* here means a path through the smallest possible number of ASs):

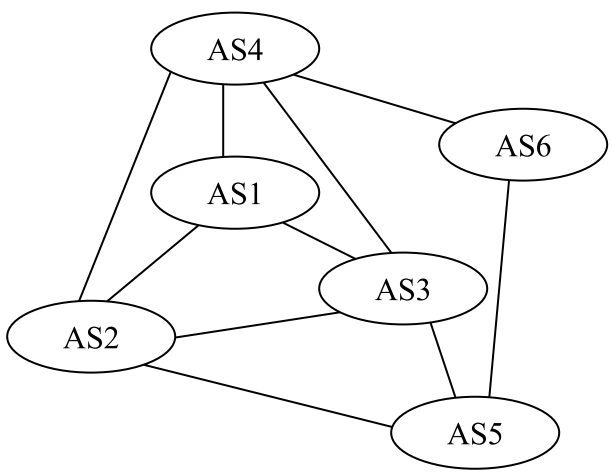

1. AS2 will not use the direct link to AS1 unless absolutely necessary. It prefers paths through AS4, then paths through AS3, then paths through AS5.
2. AS3 will not use the direct link to AS1 unless absolutely necessary. It rejects all paths through AS4.
3. AS4 always chooses the shortest available path.
4. AS5 always prefers paths through AS6 and otherwise chooses the shortest available path.
5. AS6 always chooses the shortest available path.

A *(10 points)*. Assume that the process begins with AS1 advertising its destinations. At some point, the ASs will have paths to AS1 that will no longer change (*i.e.,* the above graph should *converge* to a solution). Write down these "best" AS paths to AS1 for each AS in the network. *(You will get 2 points for each correct path for AS2–6.)*

II. Business relationships often determine routing policies. We can model the Internet using two types of relationships: *customer-provider* relationships and *peer-peer* relationships. In a customer-provider relationship, one AS, the *customer*, purchases service from another AS, the *provider*. The customer usually pays the provider to carry traffic and provide accessibility to other networks, while the provider advertises its ability to reach the customer so that the customer is accessible to the network. In a peer-peer relationship, two ASs form a mutual agreement to carry each other's traffic and share routes, often providing shortcuts to longer routes that would have to go through a provider.

These relationships affect routing policies in the following ways. ASs will advertise all their routes to their customers and will advertise routes through customers to all neighbors. (That is part of what customers pay for.) On the other hand, ASs will not advertise routes through a peer or a provider to any neighbor other than a customer. In addition, routes through customers are preferred first, then routes through peers, and finally routes through providers.

In the graph below, assume that AS11 and AS12 are peers, that AS13 is a customer of AS11, and that AS14 is a customer of AS12. Assume that the destination AS is AS14.

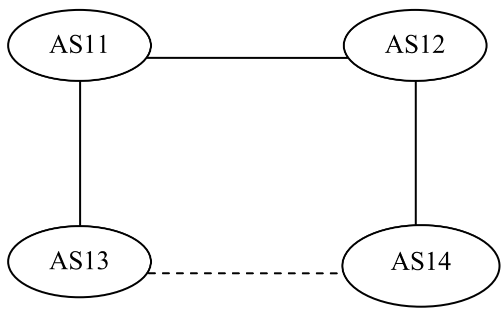

B ***(6 points).*** Assume the link between AS13 and AS14 is not present, that AS14 is advertising itself to its neighbors, and so on. What are the AS paths from each AS to AS14? *(You will get 2 points for each path.)*

Now suppose that AS14 is nervous about its link to AS12 and decides to purchase a link from AS13, *i.e.*, that it becomes a customer of AS13. However, it tells AS13 (through a special agreement) that this link is a backup link and is to be given lower preference than any other path AS13 might have to AS14. Thus there is no change to the AS paths from Question B, because this is taken into account when AS13 chooses and advertises a "best route" to AS14.

C ***(6 points).*** When a link fails, the corresponding route is withdrawn and should be replaced by another one. Suppose the link between AS12 and AS14 fails; so AS12 withdraws its route to AS14. AS13 has a backup link that it will advertise as long as there is no other route to AS14. Given this change, what are the new AS paths for each AS to AS14 once the routes stabilize? *(You will get 2 points for each path.)*

D ***(6 points).*** Suppose that the link between AS12 and AS14 is restored. Keeping in mind the preference rules for business relationships, what are the AS paths for each AS to AS14 now? *(You will get 2 points for each path.)*

E ***(6 points).*** AS14 complains to AS13 that it still receives traffic through the backup link which was not supposed to be used as long as other routes are available. AS13 reboots its routers, which temporarily disables the link from AS13 to AS14. What are the AS paths after AS13 comes back on-line? Note that AS13 becomes reconnected to the network after the AS13–AS14 path has been withdrawn. *(You will get 2 points for each path.)*