

The Internet

These exercises are intended to help you master and remember the material discussed in lectures and explored in labs. In future semesters, we may make some or all of these exercises required, but for now they remain optional. We suggest that you do them as we go over the material, but you may also want to use them to review concepts before the exam.

Please note also that some of the exercises are meant to be done with a calculator, while in exams, we just want you to be able to set up the equations correctly.

Rather than using this version directly, we suggest that you use the version without solutions to solve the problems before looking at the answers. Many studies have shown that people often trick themselves into believing that they know how to solve a problem if they are presented with the answer before they try to solve the problem themselves.

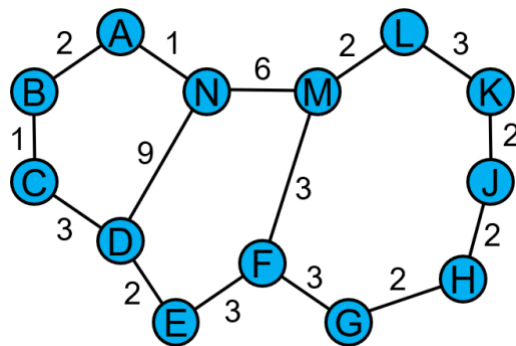
- [L5] Construct routing tables for nodes F and N in the network shown below. Each table should contain an entry for every other node in the network. For simplicity, instead of numbering the interfaces, just write the name of the node to which F (or N) should send a packet next. For example, node N should send to A in order to reach A.

Table for F

Dest	Next node
A	M
B	E
C	E
D	E
E	E
F	-
G	G
H	G
J	G
K	M
L	M
M	M
N	M

Table for N

Dest	Next node
A	A
B	A
C	A
D	A
E	A
F	M
G	M
H	M
J	M
K	M
L	M
M	M
N	-



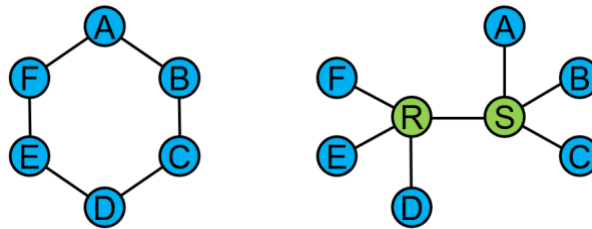
2. [L5] Compare the two network topologies below, in which six nodes (blue) are connected with or without routers (green). Briefly give two advantages of each topology over the other, keeping in mind that you are only comparing the other topologies, not “all possible or ideal” networks.

Advantages of the ring (left):

1. No routers needed (lower cost).
2. Fewer links, fewer network interfaces (lower cost)
3. No single link or node failure disconnects any pair of computers other than a failed node.
4. Up to $2\times$ link capacity from any computer to any other computer, in absence of contention.

Advantages of the hierarchical topology (right):

1. More overall capacity when subgroups (ABC, DEF) communicate amongst themselves.
2. Cross-subgroup paths do not interfere with local paths (B cannot reach D, E, or F without interfering with either A or C’s network use).



3. [L5] Consider the network graph shown below.

- A. Identify a link that can fail without affecting the operation of the network, and use your knowledge of routing to give the new route taken between a pair of nodes (you choose) that normally (before such a failure) uses the failed link. Count hops (cost 1 per link) to pick the routes.

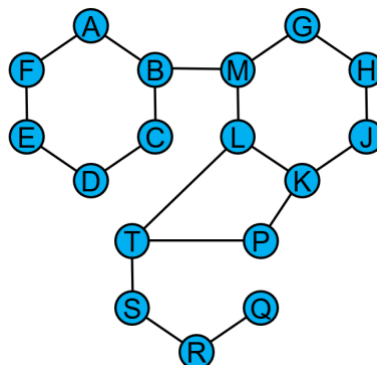
If (L,T) fails, the network remains connected. Messages from L to T, which prior to a fail would use the failed link, instead go to K, then P, and finally to T.

- B. Identify a node that can fail without affecting the operation of the network (other than the node itself, of course).

Nodes A, C, D, E, and F can each fail (by themselves) without affecting connectivity for other nodes.

- C. Identify a link that is critical—if it fails, some nodes will become disconnected—and a pair of nodes that becomes disconnected if the link fails.

Failure of the link (S,T), for example, disconnects Q, R, and S from the rest of the network.



4. [L5] When shared Ethernets were commonplace, a common failure mode occurred in which one machine would start to send signals continuously into the network. Explain why this failure mode was a big problem for shared Ethernets and why it is less important with a star topology.

When a machine fails in such a manner, it fails to obey the rules about taking turns on the network (called the medium access control, or MAC protocol). Other machines thus never get a turn, and the entire network becomes unusable, disconnecting all machines on the shared network from all other machines.

In a star topology, any given link is shared only by the failed machine and the hub. The hub may not be able to transmit to the failed machine, but otherwise, other computers in the network are unaffected by the failure.

5. [L6] New ideas in networking have historically been tested globally by implementing the idea within several ASs distributed around the world (for example, on several university campuses) and then using tunneling—a form of encapsulation—to pass through parts of the Internet that do not support the new ideas.

Imagine that UIUC has developed a new Illini Protocol (ILP) and that we are using it within the campus. Two other schools have also adopted ILP locally. ILP packets are not recognized by the Internet. Explain how encapsulation can be used to move ILP packets from UIUC through the Internet to one of the other campuses. Where should the packets be wrapped up, and using what protocol? Where should they be unwrapped and return to being ILP packets?

Let's imagine that we look into forwarding tables. The path from any UIUC machine to any machine in another campus must pass from a UIUC router out into the rest of the Internet, and must pass through one of the other campuses' routers before reaching its final destination. At these routers—at the “top” of the UIUC hierarchy—we can encapsulate each ILP packet using IP, with the IP destination address set to that of the router at the “top” of the other campus' hierarchy. When the router receives such an IP packet, it removes the IP wrapper and finds an ILP packet, which it can then deliver locally within the other campus, using the local support for ILP to reach the final destination.

Tunneling was used to test a wide variety of new ideas in networking, including multicast, quality-of-service, new generations of the IP protocol (such as IPv6), and so forth.

6. [L6] Encapsulation is not free. Packet headers must also be transmitted over the same physical channels as actual data bits, so larger headers implies longer transmission times (or lower effective throughput). Modern Ethernet, for example, maps every 8 bits into a 10-bit pattern before translating into signals. Although doing so provides several benefits, including detecting roughly 75% of bit errors, one also loses 20% of the capacity of any Ethernet link. Instead of being able to send a Gigabit in a second, for example, the real data can only be 800 Megabits.

Imagine that you are sending your data over TCP. Assume that TCP headers require 20B, IP headers require another 20B, and Ethernet headers require 20B. Assume also that each Ethernet packet can hold 1560B, which means 1500B of data plus 60B of headers. Finally, the data (including the headers) in the Ethernet packet is 8b/10b-coded. Assuming that you are sending lots of data, what is the best you can hope to achieve in terms of a fraction of the link capacity. (If your answer were 60%, for example, that would mean that data could occupy as much as 600 Mbps of a 1 Gbps link.)

Each packet contains 1500 Bytes to data. Before 8b/10b coding, each packet requires 1560 Bytes to transmit. However, the data and header are converted before being sent out, increasing the amount to $1560 \times 10 / 8 = 1950$ Bytes. Thus only the fraction $1500 / 1950 = 0.769 = 76.9\%$ of the link capacity can be used by data. This result is an upper bound: if smaller packets are sent, as is often the case in practice, the inefficiency increases, since the size of the headers and 8b/10b expansion remains the same.

7. [L6] Your friend performs an experiment with UDP, sending packets A, B, C, and D—in that order—to another computer across the Internet. The other computer records receiving C, then A, and finally D. B is never received. Explain what properties TCP provides that might improve your friend's results. In particular, what problems shown in the example are solved by TCP?

TCP provides reliable, in-order delivery. In particular, B will be retransmitted until it has been received, and then the computer at the end will reorder the packets so that the application sees A first, then B, then C, and finally D.

8. [L6] Explain an example application in which reliable delivery is not an advantage.

The examples given in class were games and teleconferencing, but any reasonable example is acceptable. In games, messages often convey world updates, such as changes in a player's position. Since the player may keep moving, sending old updates is a waste of time and resources. A similar explanation serves for teleconferencing, in which a new video frame is captured by a camera periodically (usually 10s of milliseconds). Re-transmitting old frames is again a waste of time and resources.