

This homework is due at 11:59:59 PM on December 5, 2022 and is worth 3% of your grade.

Name: _____

NUID (with leading zeros): _____

| Problem | Possible | Score |
|----------------|-----------------|--------------|
| 1 | 20 | |
| 2 | 20 | |
| 3 | 30 | |
| 4 | 30 | |
| Total | 100 | |

- 1a. Imagine a file is being shared via BitTorrent, and the file is divided into 8 BitTorrent pieces $\{A, B, C, D, E, F, G, H\}$. Now, suppose that you and two other leechers $\{l_1, l_2\}$ are in the system, and that everyone is connected to everyone else. The bitfields are as follows:

| Leecher | Bitfield [ABCDEFGH] |
|---------|---------------------|
| you | [11100100] |
| l_1 | [00000101] |
| l_2 | [00011110] |

Assume that each leecher can send and receive two pieces in each round. Given that you must upload something to each l_1 and l_2 in each round in order for them to upload something to you, which bitfield should you advertise to l_1 and l_2 in the next round in order to ensure that they stay interested in you for the longest amount of time? *Note that you can lie about your bitfield, and can reveal different bitfields to l_1 and l_2 .* (10 pts)

- 1b. Assuming ideal conditions for you, what is the minimum number of rounds until you have downloaded the entire file? (5 pts)


- 1c. Why does network address translation (NAT) make life difficult for peer-to-peer applications (such as BitTorrent), but does not generally affect client-server applications (such as Web servers)? (5 pts)

2a. Briefly explain how unstructured P2P networks are organized. Why is it sometimes difficult to search for data/files on unstructured P2P networks? (5 pts)

2b. Briefly describe how a consistent hashing algorithm works. What is the key difference between a consistent hashing algorithm and a typical hashing algorithm? (5 pts)

2c. Consider the Pastry/Tapestry structured P2P overlay. Explain why it takes a maximum of $\log_b N$ hops to route a message from any node to any other node in the worst case, where N is the number of hosts in the P2P system, and b is the base of the ID space (e.g., binary, hexadecimal). (10 pts)

3. Suppose a peer-to-peer structured overlay network uses eight digit, binary node IDs. The below figure shows an example of this overlay. Nodes *A*, *B*, and *C* are labeled, and their respective node IDs are shown. Assume that the remaining node ID-space is well-populated with other nodes (even though they aren't shown in the figure).



- 3a. What is the minimum number of hops it could take for a message to travel from node *A* to node *B*? Show an example path of this length (you may make up node IDs to complete the path if necessary). (4 pts)
- 3b. What is the maximum number of hops it could take for a message to travel from node *A* to node *B*? Show an example path of this length (you may make up node IDs to complete the path if necessary). (4 pts)
- 3c. What is the minimum number of hops it could take for a message to travel from node *A* to node *C*? Show an example path of this length (you may make up node IDs to complete the path if necessary). (4 pts)
- 3d. What is the maximum number of hops it could take for a message to travel from node *A* to node *C*? Show an example path of this length (you may make up node IDs to complete the path if necessary). (4 pts)
- 3e. What is the maximum number of peers that could be part of this P2P overlay network? (4 pts)

3f. What is the size of the routing table of each peer in this P2P overlay network, in terms of number of entries (i.e. not in bits or bytes)? (4 pts)

3g. Complete the example routing table for node A. The first two rows are given.

0*****: <IP address>


1*****: <Next row>

10*****: <Next row>

11*****: <IP address>

(6 pts)

4. Suppose a peer-to-peer structured overlay network uses eight digit, hexadecimal node IDs. The below figure shows an example of this overlay. Nodes *A*, *B*, and *C* are labeled, and their respective node IDs are shown. Assume that the remaining node ID-space is well-populated with other nodes (even though they aren't shown in the figure).



- 4a. What is the minimum number of hops it could take for a message to travel from node *A* to node *B*? Show an example path of this length (you may make up node IDs to complete the path if necessary). (4 pts)
- 4b. What is the maximum number of hops it could take for a message to travel from node *A* to node *B*? Show an example path of this length (you may make up node IDs to complete the path if necessary). (4 pts)
- 4c. What is the minimum number of hops it could take for a message to travel from node *A* to node *C*? Show an example path of this length (you may make up node IDs to complete the path if necessary). (4 pts)
- 4d. What is the maximum number of hops it could take for a message to travel from node *A* to node *C*? Show an example path of this length (you may make up node IDs to complete the path if necessary). (4 pts)
- 4e. What is the maximum number of peers that could be part of this P2P overlay network? (4 pts)

4f. What is the size of the routing table of each peer in this P2P overlay network, in terms of number of entries (i.e. not in bits or bytes)? (4 pts)

4g. Suppose that node *A* sends a message to ID *FEFA9FBC*, but a node with this ID does not exist in the P2P overlay. What happens to the message? (6 pts)