CS 7270: Internet Applications

Midterm exam

October 2, 2007

Duration: 80 minutes

Name: _____

Part	**Maximum**	**Obtained**
I	25	
II	25	
III	25	
IV	25	

Final grade:

Problem 1 (25 points)

We discussed the traffic classification problem during the first week of classes. Here, you are asked to think independently about a specific instance of this problem. You will not find the answer in any research paper or at your notes. Think of the following as a small research problem that you need to approach within 20 minutes.

Specifically, suppose that you want to design a traffic classification scheme that aims to map each new flow to one of the following three classes of traffic: (1) voice-over-IP (voice), (2) video streaming (video), and (3) peer-to-peer file transfers (p2p). The rest of the traffic can be grouped as "other".

The constraints:

- you cannot inspect the payload of the packets,
- you cannot make assumptions about the protocol or port numbers that each application uses,
- the method needs to classify each new flow in real-time, after observing the first few (say 10-100) packets of the session,
- the method should be effective when deployed at an edge router, connecting say a university to the Internet.

How would you design such a classifier? Which are the key features that you would associate with each of the three previous applications? You do not need to give a complete design or a rigorous algorithm. Focus on the key ideas and explain why they are important.

Problem 2 (25 points)

We discussed voice-over-IP applications during the second week of classes. Here, you are asked to think independently about a specific instance of this problem. You will not find the answer in any research paper or at your notes. Think of the following as a small research problem that you need to approach within 20 minutes.

Specifically, suppose that you want to design a conferencing application (only voice, no video or text) for large groups of people. Each pair of participants has a different network delay. Further, some users are behind "challenged" network links, i.e., low bandwidth, high jitter, or frequent packet losses.

Describe an architecture that would work well for such a conferencing application. The main constraint is that, at least for most users, the bandwidth requirement to participate at the call should not increase with the number of participants N.

Further, the quality of the overall call (determined by codec selection, loss recovery method, jitter compensation) should not be limited by the "most challenged" user, i.e., the user that is behind the link with the lowest bandwidth or highest packet loss rate.

Finally, you should try to provide each user with the best quality that her/his network link allows.

Problem 3 (25 points)

We discussed peer-to-peer applications during the fifth week of classes. Here, you are asked to think independently about a specific instance of this problem. You will not find the answer in any research paper or at your notes. Think of the following as a small research problem that you need to approach within 20 minutes.

Specifically, suppose that you want to design a peer-to-peer file distribution application that works well in "challenged networks". Such networks are lossy and they only provide intermittent connectivity (think of mobile wireless devices in fast-moving vehicles). In the networks we consider, it is unlikely that any two peers A and B will stay connected for more than few seconds.

The files are relatively large and they may require several minutes to be transferred. You can assume that there are many, perhaps thousands, of users in this peer-to-peer system. Further, you can assume that users are completely altruistic and they are interested to minimize the download time of other users, even if they don't get anything out of it (for instance, think of soldiers in a battleground sharing a file that shows a video of the next target).

Describe the key features that such a peer-to-peer system should have to perform well. You can borrow ideas from existing systems, but you may also need to come up with novel techniques. You do not need to describe a complete algorithm. Focus on the key ideas and explain why they are important.

Problem 4 (25 points - the only wrong answer is to not give an answer)

Suppose that we are in 2017. Try to imagine 3-4 new Internet applications that people will be using then, and describe them at a high level. What do you think "the Web" will become by then?