University of Glasgow

Thursday, 19 May 2022
09:30-11:00 BST
Duration: 1 hour 30 minutes
Additional time: 30 minutes
Timed exam – fixed start time

DEGREES of MSc, MSci, MEng, BEng, BSc, MA and MA (Social Sciences)

NETWORKED SYSTEMS (H)
COMPSCI 4012

Answer all 3 questions

This examination paper is an open book, online assessment and is worth a total of 60 marks.

1. **(a)** There have been numerous studies measuring the sizes of the packets that traverse of the Internet. Three common findings are the occurrence of a very large number of packets sized approximately 40 bytes, the occurrence of a large number of packets of approximately 1500 bytes in size, and relatively few packets of other sizes. Explain the reason for each finding. [5]

(b) A client connected to a WiFi network opens a TCP connection, sends 4000 bytes of data to a server, then closes the connection. The server accepts the connection and repeatedly calls `recv()` on that connection, providing a 4000 byte buffer to read into each time, until the connection is closed. How often would you expect the server to need to call `recv()` in order to read all the data on the connection? How many bytes of data would you expect the call(s) to the `recv()` function to return? Explain your answers, including some discussion around whether you'd always expect the same answer if the scenario was repeated on multiple occasions, and what might cause differences in behaviour? [5]

(c) A TCP connection is used to transmit a file containing 1024 kilobytes of data. The connection traverses a network that has a path maximum transmission unit (MTU) large enough to allow each TCP segment to deliver one kilobyte of payload data, plus any necessary header data. If no packets are lost, and the initial congestion window is 1 kilobyte, what will be the duration of the TCP connection, measured in multiples of the network round-trip time? Explain your working. [6]

(d) You are running a voice-over-IP phone call, where the sender transmits one UDP packet containing speech data every 20 ms. When sending over an otherwise idle IP network, the UDP packets arrive every 20ms, matching the timing with which they are sent, after some delay depending on the propagation time of the packets across the network. You start a large file downloading using a TCP connection, traversing the same network path as the voice-over-IP call. Discuss how this affects the voice-over-IP call. [4]

2. **(a)** In the original design of the Internet, packet forwarding was based solely on the destination IP address of each packet, and the choice of transport protocol was solely a matter for the end systems. In practice, this is no longer the case, and it is difficult for end systems to use a transport protocol other than TCP or UDP (this is why, for example, QUIC is built to run on top of UDP). Explain what prevents the use of other transport protocols in the current Internet. Discuss the trade-offs around whether it is desirable to prevent the use of new transport protocols in this manner. State whether you think such blocking is overall helpful or harmful for the Internet, and justify your conclusion. [12]

(b) You are building a networked application, and decide that it might benefit from using the QUIC transport protocol. When testing the application, you find that QUIC improves performance for some clients, but that other clients fail to establish QUIC connections to your servers, due to the presence of firewalls that block the UDP traffic on which QUIC relies. Explain how you would work around this problem. Discuss whether the effort needed to implement this work-around is worthwhile. [8]

3. **(a)** Many popular websites use content distribution networks (CDNs) to improve their service. CDNs are designed to improve scalability and serve requests quickly by directing them to a

data centre located near to the user making the request. Some CDNs do this by using the DNS to direct requests, others use a technique known as anycast routing. Briefly explain how these two approaches work. Discuss which would be most appropriate for a CDN that is handling content where the number and geographic location of the users making the requests is changing rapidly, and briefly explain why the other approach is not suitable.

[12]

(b) CDNs often advertise that they provide protection from distributed denial of service (DDoS) attacks for websites. Discuss what properties of a CDN might make it well-suited to providing such a service. [2]

(c) A common use of CDNs is to support streaming video services, such as those offered by Netflix or the BBC iPlayer. Consider the case where you are watching streaming video of a live sporting event on such a service. Due to a fault, your Internet connection drops out for several seconds before reconnecting. This causes the video play-back to stall.

Some time later you are watching a pre-recorded movie when your Internet connection fails in the same way, dropping out for a few seconds before reconnecting. This time, however, the movie continues playing, seemingly unaffected. Discuss why this difference in behaviour occurs.

[6]

END OF QUESTION PAPER