University of Glasgow

Monday, 29th April 2024
09:30 - 11:00 BST
Duration: 1 hour 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)** TCP connections proceed in two phases. First, the connection is established using a three-way handshake (SYN, SYN+ACK, ACK) where the client and server agree to communicate and synchronise their initial sequence numbers. Then, after that handshake has completed, data is sent on the connection. Consider a modified version of TCP that overlaps these two phases, where the client provides some data as an additional parameter to the `connect()` function which is sent in the TCP segment that has the SYN bit set, with the response being returned in the SYN+ACK segment. State what would be the benefit of this idea and discuss why it is not feasible in practise. [5]

(b) When TCP starts to send data, it uses the slow start algorithm to determine an appropriate congestion window size then switches to the congestion avoidance mode to adapt the congestion window to ongoing network conditions. It is sometimes suggested that TCP should remember the congestion window it used when sending data to a particular destination, and start new connections to that destination with the saved congestion window rather than using slow start. Briefly discuss the potential benefits and risks of this suggestion. [5]

(c) Recent versions of HTTP include a server push feature. This allows a web server to send unsolicited data to the client, responding as if the server had received a request to retrieve a particular URL without such a request being submitted. This server push feature is claimed to improve performance in certain cases, although experiments have shown that the actual performance impact is mixed. Discuss why and how server push can improve performance, giving an example to illustrate your argument, and stating by what metric are improvements measured. Explain in what circumstances server push might not help, and might even hurt, performance. Finally, state whether you think HTTP server push is likely to be a net benefit overall, justifying your answer. [10]

2. **(a)** The network is in the middle of a transition from IPv4 to IPv6, with current estimates showing that almost half of the traffic is using IPv6 and that this fraction is growing. The main change between the two IP versions is that IPv6 uses 128 bit addresses compared to the 32 bit addresses used in IPv4. Do you think 128 bits was the right size for IPv6 addresses? Discuss the trade-offs that would have been considered when selecting the size of the IPv6 address.

[6]

(b) People use domain names (e.g., `google.com`) to identify sites, with software applications performing a DNS lookup to convert these into IP addresses prior to use. The time it takes to perform such a DNS lookup can vary significantly. Explain what causes this variation.

[4]

(c) DNS queries can be sent to a resolver using one of four transport protocols: UDP, TCP, TLS, or HTTPS. The use of the latter transport, DNS over HTTPS, has proven controversial, with claims that it helps users evade legitimate access control mechanisms. Explain why such criticism has been levelled at DNS over HTTPS but not at DNS running over the other transport protocols. [10]

3. **(a)** When sending multimedia data, such as video conferencing traffic or live TV broadcasts, across the Internet, the amount of data that must be sent in a given time interval varies,

exhibiting bursty behaviour between fixed upper and lower bounds. Briefly explain why the traffic pattern is bursty and why there are constraints on the minimum and maximum amount of data that needs to be sent in a given time period. [2]

(b) You are participating in an interactive video conference with the video data being sent in UDP packets at a rate that matches the needs of the video. While doing this, you download a very large file over a TCP connection. Assume the TCP connection uses the Cubic congestion control algorithm, and the network capacity is enough to comfortably transmit the video data alone without packet loss. State what happens to the quality of the video call while the TCP download is running and explain why this happens. [8]

(c) Interactive video conferencing applications typically send their data in UDP. Explain why these applications use UDP rather than TCP. Discuss in what circumstances such applications would be able to effectively use TCP with Cubic congestion control as a transport. [10]

END OF QUESTION PAPER