University of Glasgow

Tuesday, 6th May 2025
14:00 - 15:30 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 worth a total of 60 marks.

1. **(a)** A web browser, running on a client machine, connects to a web server for the first time and retrieves an HTTPS page over a TCP connection. Starting with the first packet of the TCP connection until the HTTP response is received, describe what packets are sent by each of the client and server, and outline what is their purpose. [6]

(b) In the scenario described in Question 1(a), the first part of the requested data is received by the client approximately three round-trip times after the initial TCP packet was sent. Describe three ways in which the protocols could change in order to reduce the time taken to retrieve the data to approximately two round-trips, and explain why each of them is problematic and unusable in practice [6]

(c) The process of peer-to-peer connection establishment in the presence of NAT devices proceeds in three stages: binding discovery, candidate exchanges, and then probing for connectivity.

Assume two peers, Alice and Bob, are trying to establish such a peer-to-peer connection between their devices. Both are using laptops that have active Ethernet and WiFi network interfaces. Alice is on a dual-stack network with IPv4 and IPv6 available on both interfaces. Bob has a dual stack Ethernet, with IPv4 and IPv6 enabled, but has access to only IPv4 on WiFi. Both Alice and Bob have access to a STUN server that is accessible over IPv4 that can be used for binding discovery, a relay server also accessible over IPv4, and a signalling channel over which they can exchange candidates.

Explain how many candidate addresses each of Alice and Bob will have available after the binding discovery phase. State, with justification, what is the number of combinations of candidate address pairs that will need to be probed for connectivity in the worst case. Describe any assumptions you make [8]

2. **(a)** Modern DNS resolvers can be accessed by sending DNS queries using TLS (called "DoT" for short) or using HTTPS ("DoH"), in either case running over a TCP connection. Explain how the two approaches work and discuss why DoH has proven controversial while DoT has not. [6]

(b) At the top-level of the DNS hierarchy are the root servers. While the physical servers that implement the DNS root are geographically widely distributed, with hundreds of replicas located around the world, technical limitations of the DNS protocol are such that there can be only 13 root server operators. For historical reasons, most of the root server operators are US-based organisations.

Discuss whether this concentration of root server operators into a single jurisdiction is a concern, considering both technical and non-technical factors. [6]

(c) With the aid of an example, explain what is meant by ossification of network protocols and why it is a concern. Explain what is GREASE, and discuss why the combination of pervasive encryption and GREASE is thought to help avoid protocol ossification. [8]

3. **(a)** The IPv4 address space was fully allocated some years ago and it is now common to use network address translation (NAT) to allow a number of devices on a private network to share a single public IPv4 address. If you are a network operator, and have been allocated

the public IPv4 address range 192.0.2.0/24 for your network and customers, how many customers can you support, assuming that each customer is allocated only a single IPv4 address? [4]

(b) Assuming they cannot switch to IPv6 or increase the size of their IPv4 address allocation, describe how the network operator described in Question 3(a) could organise their network to support more customers. Is there a limit to the number of customers that operator can support, given only the 192.0.2.0/24 address range? [6]

(c) As the Internet has grown, a number of organisations have deployed private networks on a global scale. These include Internet companies such as Apple, Facebook, Google, Netflix, and Microsoft, as well as cloud computing infrastructure providers and content distribution networks such as Amazon Web Services, Akamai, and Cloudflare, amongst others.

The experience of these companies is that the latencies they measure across their private networks are often significantly lower than the latencies measured between similar sites across the public Internet. That is, the latency of intradomain paths is lower than that of the interdomain paths selected by BGP, despite both sets of paths being between data centres in the same cities. For example, the latency measured between the Amazon data centres in São Paulo, Brazil, and Dublin, Ireland, is often lower when measured through Amazon's private network than when measured across the public Internet.

Considering both the operation of Internet routing and other factors, discuss why the latencies measured in private networks are lower than those measured across the public Internet. [10]

END OF QUESTION PAPER