

**THE UNIVERSITY OF WESTERN ONTARIO  
LONDON CANADA**

**COMPUTER SCIENCE 457a  
MIDTERM EXAMINATION  
OCTOBER 27, 2007  
2 HOURS**

NAME: \_\_Marking Scheme\_\_\_\_\_

STUDENT NUMBER: \_\_\_\_\_

Question

1-25. \_\_\_\_\_

26. \_\_\_\_\_

27. \_\_\_\_\_

28. \_\_\_\_\_

29. \_\_\_\_\_

30. \_\_\_\_\_

31. \_\_\_\_\_

32. \_\_\_\_\_

33. \_\_\_\_\_

34. \_\_\_\_\_

TOTAL \_\_\_\_\_

(Out of 140 marks)

There are no cheat sheets, books, or other reference materials allowed for this exam. No calculators, cell phones, or other electronic devices are permitted either.

Part I – Multiple Choice, True/False – Choose the best answer from the choices given. Circle your answer on the paper, and fill in the answer on the Scantron form. [50 marks total, 2 marks each]

1. RTSP is an out-of-band signaling protocol because:
  - a. RTSP messages are encapsulated into HTTP messages.
  - b. RTSP messages are interleaved with the media stream.
  - c. RTSP messages are sent over circuit-switched networks.
  - d. RTSP messages are sent through a separate logical connection from the media stream.
  - e. None of the above.
  
2. Which of the following approaches to recovering from packet loss incurs no redundancy overhead:
  - a. Forward Error Correction through exclusive-OR operations.
  - b. Forward Error Correction through piggybacking.
  - c. Interleaving of data.
  - d. All of the above incur no redundancy overhead.
  - e. All of the above, in fact, incur redundancy overhead.
  
3. A single leaky bucket can be used to police a flow's:
  - a. Average rate.
  - b. Burst size.
  - c. Peak rate.
  - d. All of the above.
  - e. None of the above.
  
4. If stored audio is streamed directly from a Web server to a media player, then the application must be using TCP as the underlying transport protocol.
  - a. True.
  - b. False.
  
5. When using RTP, it is possible for a sender to change encoding in the middle of a session.
  - a. True.
  - b. False.
  
6. SIP mandates that all SIP clients support particular audio and video encodings.
  - a. True.
  - b. False.

7. Suppose Alice wants to establish a SIP session with Bob. In her INVITE message, she includes the line: m=audio 48753 RTP/AVP 3 (AVP 3 denotes GSM audio). Alice has therefore indicated in this message that she wishes to send GSM audio.
- a. True.
  - b. False.
8. Referring to the previous question, Alice has indicated in her INVITE message that she will send audio to port 48753.
- a. True.
  - b. False.
9. Suppose we choose a larger value for a fixed playout delay for a real-time interactive multimedia application. This will result in:
- a. Less loss, less interactivity.
  - b. Less loss, higher interactivity.
  - c. More loss, less interactivity.
  - d. More loss, higher interactivity.
  - e. None of the above.
10. Suppose we choose a smaller value for a fixed playout delay for a real-time interactive multimedia application. This will result in.
- a. Less loss, less interactivity.
  - b. Less loss, higher interactivity.
  - c. More loss, less interactivity.
  - d. More loss, higher interactivity.
  - e. None of the above.
11. To support the Real Time Protocol (RTP) all hosts and routers must be updated.
- a. True.
  - b. False.
12. The Session Initiation Protocol (SIP) provides which of the following services:
- a. Call setup.
  - b. Determination of callee's current address.
  - c. Call management.
  - d. All of the above.
  - e. None of the above.
13. To implement RSVP, RSVP software need only be installed in end systems and not on network routers.
- a. True.
  - b. False.

14. In the Differentiated Services approach, simple functions are put in the network core, and complex functions are placed at the edge in end systems.
- a. True.
  - b. False.
15. Using public key cryptography, suppose Bob wants to send a secret message to Alice. Then Bob should:
- a. Encrypt the message with Alice's private key and send Alice the message.
  - b. Encrypt the message with Alice's public key and send Alice the message.
  - c. Encrypt the message with his private key and send Alice the message.
  - d. Encrypt the message with his public key and send Alice the message.
  - e. None of the above.
16. Using public key cryptography, suppose Bob wants to send a message to Alice and Alice wants to be sure that the message was indeed sent by Bob. Then Bob should:
- a. Encrypt the message with Alice's private key and send Alice the message.
  - b. Encrypt the message with Alice's public key and send Alice the message.
  - c. Encrypt the message with his private key and send Alice the message.
  - d. Encrypt the message with his public key and send Alice the message.
  - e. None of the above.
17. Using public key cryptography, suppose Bob wants to send a secret message to Alice and Alice wants to be sure that the message was indeed sent by Bob. Then Bob should:
- a. Encrypt the message with Alice's public key, encrypt the result with his public key and then send Alice the message.
  - b. Encrypt the message with his private key, encrypt the result with Alice's private key, and then send Alice the message.
  - c. Encrypt the message with his private key, encrypt the result with Alice's public key, and then send Alice the message.
  - d. Encrypt the message with his public key, encrypt the result with Alice's public key, and then send Alice the message.
  - e. None of the above.
18. Suppose Bob wants to send Alice a digital signature for the message  $m$ . To create the digital signature:
- a. Bob applies a hash function to  $m$  and then encrypts the result with his private key.
  - b. Bob applies a hash function to  $m$  and then encrypts the result with his public key.
  - c. Bob encrypts  $m$  with his private key and then applies a hash function to the result.
  - d. Bob applies a hash function to  $m$  and then encrypts the result with Alice's public key.
  - e. None of the above.

19. Suppose Alice receives from Bob a message  $m$  along with a digital signature for the message  $m$ . To verify that the message was not changed and that Bob indeed sent the message, Alice:
- Applies Bob's public key to the digital signature, then a de-hashing function to the result. She then compares the result of this last operation with the message  $m$ .
  - Applies Bob's public key to the digital signature, applies the hash function to  $m$ , and compares the results of the two operations.
  - Applies a de-hashing function to the digital signature and compares the result to  $m$ .
  - Any of the above.
  - None of the above.
20. Suppose Bob is purchasing merchandise from Alice Inc. over the Internet. SSL permits:
- Bob to determine whether Alice Inc. is authorized to accept credit card purchases.
  - Alice Inc. to determine if Bob has a good credit history.
  - Bob to determine if Alice Inc. is a legitimate company.
  - All of the above.
  - None of the above.
21. In IP spoofing, the attacker interchanges the source and destination addresses in the sender's IP datagram.
- True.
  - False.
22. When a public key is signed by any Certificate Authority, you know that public key can always be trusted.
- True.
  - False.
23. If everyone were to adopt IPsec, it would be able to effectively put an end to IP spoofing.
- True.
  - False.
24. Which of the following is an optional Secure Sockets Layer (SSL) service?
- Server authentication.
  - Client authentication.
  - Data encryption.
  - All of the above.
  - None of the above. All of these services are, in fact, mandatory.
25. The Secure Sockets Layer (SSL) can only be used between web browsers and web servers.
- True.
  - False.

Part II – Short/Long Answer – Complete the following questions in the space provided on the exam paper.

26. The following questions deal with delay and jitter in multimedia networks. [8 marks total]
- a. What is the difference between end-to-end delay and packet jitter? [2 marks]

End-to-end delay refers to the total time taken for a packet to be generated and sent across a network up to and including its eventual reception at its destination. (So this includes transmission delays, queuing delays, and so forth ...) Packet jitter, on the other hand, refers to the variation of packet delays between packets in the same stream.


- b. What are the causes of packet jitter in a network? [2 marks]

Packet jitter tends to be caused by fluctuations in the volume of data along the same path as the flow of packets on which jitter is being measured. If the volume of data changes drastically, this will tend to impact queuing delays at routers along this path, which will cause issues. Additionally, if your stream is being sent over a TCP connection, anything which causes retransmissions (network congestions, packet errors, and so on) will impact packet jitter as well.

- c. Why is a packet that is delayed and received after its scheduled playout time considered lost? [4 marks]

In a time-sensitive application such as a streaming or real-time interactive application, any interruption will be readily noticeable by the user; having to pause the application to wait for data is an example of one such interruption. As a result, pausing an application for this reason is typically unacceptable. Because of this, packets are given playout times, which are essentially deadlines by which the packets must be received or else the application will need to be paused to wait for the packet to arrive. So, if a packet misses its playout time, instead of pausing and waiting, the application will treat the packet as being lost to allow itself to proceed in a timely fashion, with the hope that the “loss” of the packet will be relatively minor in comparison to the pausing of the application. In the packet is eventually received it is discarded as the opportunity to make use of it has unfortunately passed. As a result, packets received after their playout times are considered lost.

27. The following question parts deal with issues in audio representation. [10 marks total]
- a. Using a diagram and an accompanying textual description, describe the process for converting an analog sound wave into a digital representation. [4 marks]



In a) we have an original analog sound wave to be digitized. The wave is sampled or read at a fixed rate in b), with each sample being an arbitrary real number. Finally, in c), the samples are quantized or rounded to one of a finite number of values, determined by the number of bits desired to represent each sample. This gives us the final digital representation of the analog wave.

- b. Suppose you have an original analog audio wave that is converted to a digital representation, and is converted back to an analog audio wave for playback to a human ear. Explain why the resulting audio wave is never exactly the same as the original audio wave. [4 marks]

Information is lost in both the sampling and quantization process, and it is not possible to have samples arbitrarily close together or of infinite resolution to compensate for this. Information is consequently going to be missing between samples and as samples are rounded during the quantization process. Since this information is lost and gone for good, reconstruction of the wave may be good, but can never be exactly the same.

- c. What can you do in the analog to digital conversion process to improve the approximation of the digital representation to the original sound wave? [2 marks]

We can increase the sampling rate or the number of quantization levels to be able to retain more of the information contained in the original wave, thereby giving us a better approximation.

28. The following questions deal with video representation and use. [10 marks total]
- a. Why do most video representations, especially for network transmission, not use RGB pixel values, and instead deal with chrominance and luminance values? [2 marks]

Ultimately, better compression can be had by using chrominance and luminance values. Since the human eye is more sensitive to brightness than colour, we can use fewer bits to represent colour, thereby allowing us to reduce the bandwidth needs of our stream. At the same time, we can still make use of the same redundancies in the video stream, perhaps in some ways better. (For example, if an entire scene has uniform illumination, the luminance values might be more consistent, allowing for even more compression this way than when an RGB scheme was used instead.)

- b. How can a layered or scalable video stream, such as that produced using MPEG 4 compression, adapt to varying amounts of available network bandwidth? How could this work with Differentiated Services where packets can be prioritized? [4 marks]

A layered or scalable video scheme can be decomposed into multiple layers containing different elements of the original video stream. For example, you might have a base layer containing the bare essentials, with multiple enhancement layers providing more detail and improving quality. As long as the base layer makes it through, you can enjoy some experience ... if bandwidth is available and enhancement layers can also make it through, then the user can have a higher quality experience. As bandwidth fluctuates, the enhancement layers can be dropped depending on resource availability or lack thereof. Differentiated Services can help with this by marking the packets with priorities, making the base layer the highest priority and enhancement layers lower priorities. When bandwidth is scarce, the base layer still makes it because it is the highest priority. When bandwidth is plentiful, the lower priority packets containing the enhancement layers can also make it through to improve video quality

- c. What are the advantages and disadvantages of constant and variable bit rate approaches to video compression as far as network quality of service is concerned? [4 marks]

Constant bit rate video has constant and predictable resource requirements, making it easier to schedule on a network. But, quality might be sacrificed or resources might be wasted to ensure this rate is met. Variable bit rate video maintains quality at a target level while not wasting resources, but resource needs are unpredictable and scheduling on a network is harder as a result.

29. The following questions deal with scheduling of packets for networks. [12 marks total]
- a. In class, we discussed non-preemptive priority queuing for scheduling packets in networks. What would be preemptive priority queuing? Does preemptive priority queuing make sense for computer networks? Explain. [4 marks]

Preemptive priority queuing would allow the transmission of one packet to be preempted or interrupted when a higher priority packet arrives to be sent out over the network. Non-preemptive priority queuing does not allow this; the higher priority packet would have to wait for the lower priority packet to complete transmission before it is scheduled. Preemptive priority queuing does not make sense for networks as there is no point in interrupting the transmission of a packet ... the partial packet that makes it onto the network will be useless and a waste of resources spent to that point. It is better to let transmission complete and move on than wasting resources in this fashion.

- b. What purpose does a “discard policy” have in a FIFO scheduling discipline? Name and briefly discuss three possible discard policies. [4 marks]

The discard policy in a FIFO discipline is used to determine which packet should be dropped when an arriving packet arrives at a full queue. Since the packet cannot be put into a full queue for transmission, something must be dropped. Policies discussed in class include tail dropping (drop the arriving packet), priority dropping (dropping packets based on some kind of priority scheme), and random dropping (selecting a packet at random to drop).

- c. Describe how the weighted fair queuing (WFQ) approach to network scheduling works. How can a weighted fair queuing approach be made to act the same as a work conserving round robin approach? [4 marks]

Weighted fair queuing is a generalized work conserving round robin approach, in which each scheduling class gets a weighted amount of service in each cycle; each class may get different amounts of service compared to the other classes. This allows for some scheduling prioritization without starvation issues, and is generally accepted to be a good approach to use. This can act exactly like a work conserving round robin approach by having all the class weights equal and set to one.

30. What are the differences between a brute-force attack, a ciphertext-only attack, a known-plaintext attack, and a chosen-plaintext attack? [8 marks total]

Brute force: This attack simply attempts all possibilities to try and breach the cryptosystem, and is therefore very time consuming to do.

Ciphertext-only attack: In this case, only the encrypted ciphertext is available, so statistics and other information to decrypt the intercepted ciphertext.

Known-plaintext attack: In this case, not only is the ciphertext available, but at least some of the original plaintext is also available, potentially making the attack easier.

Chosen-plaintext attack: In this case, an intruder can choose the plaintext message and receive the ciphertext form. So, not only do they have access to the plaintext, they can select it themselves to suit their attack better.

31. What is the purpose behind the next header fields that are part of the headers and trailers of the Authentication Header (AH) and Encapsulation Security Payload (ESP) protocols that are part of IPsec? What advantage is there to having this field encrypted in the Encapsulation Security Payload (ESP) protocol? [6 marks total]

IP packets have a protocol field that indicates which higher level protocol entity should process the packet once it has been received. (For example, this is used to mark TCP packets and UDP packets, so that they can be processed and delivered correctly when they are received.) When IPsec is used, the protocol field must be used to indicate which IPsec protocol is being used (AH or ESP), and so the original protocol field in the packet must be stored elsewhere to ensure it can be eventually delivered after it has been received and processed by the AH or ESP protocol handlers on the receiving end. The next header fields in the headers and trailers of these protocols are used to contain this information.

By encrypting this field in the ESP protocol, someone intercepting the packet cannot determine if the packet ultimately is a UDP, TCP, ICMP, or other type of packet. In theory, this information might have been useful to the attacker, to guess at the type of communication going on, to deduce the packet structure of the original packet, and so on. Encrypting this information is therefore a good idea.

The following questions deal with trusted intermediaries for security. [12 marks total]

- a. What is a Key Distribution Center (KDC)? How does it work? Why can it be a critical part of network security infrastructure? [3 marks].

A KDC is a trusted intermediary for symmetric key cryptography. Each user shares a secret key with the KDC, allowing them to create secure sessions between themselves and the KDC. Through this mechanism, users can talk to the KDC to create session keys and required authentication information to communicate with other users of the KDC. This is a critical part of infrastructure, as it must be present to allow users to create secure sessions between one another across the network.

- b. What is a Certificate Authority (CA)? How does it work? Why can it be a critical part of network security infrastructure? [3 marks]

A CA is a trusted intermediary for public key cryptography. Its purpose is to bind a public key to a particular entity, vouching for the fact that a certain public key does belong to that entity. It does this by issuing certificates signed with its own private key that contain the public key and identifying information for a particular entity. This certificate can then be used whenever another party needs the public key for that entity to communicate securely with it. It can play a critical role because, without a CA, there is no trustworthy way of linking a public key with a particular entity, which is a serious security problem.

- c. Consider the KDC and CA servers. Suppose a KDC goes down. What is the impact on the ability of parties to communicate securely; that is, who can and cannot communicate? Justify your answer. Suppose now that a CA goes down. What is the impact of this failure instead? [6 marks]

In the case of a KDC going down, no new session keys can be generated. People with existing tickets and keys can use them until they expire, but nothing new can be created, causing a serious impact on the ability to communicate. If a CA goes down, the only thing affected is the ability to generate new certificates. Otherwise, communication can proceed as normal using existing issued certificates. As a result, a CA going down has a relatively minor impact on the ability to communicate ... there is only an issue if you need a new certificate issued by the CA.

32. The following questions deal with firewalls and intrusion detection. [12 marks total]

- a. What are the two main approaches to firewall policies? Explain each approach briefly. Identify which approach you believe provides better security, and justify your answer. [6 marks]

One approach is default permit. In this case, all packets are allowed through except for packets are explicitly denied. The other approach is default deny. In this case, all packets are denied by default, except for those that are explicitly allowed in. Generally speaking, a default deny policy provides better security as you do not need to identify bad things and block them ... everything is blocked already by default, and only good things are allowed in.

- b. What are the differences between stateless and stateful packet filtering for network firewalls? [2 marks]

A stateless packet filter, as the name implies, maintains no state and therefore filters on a packet-by-packet basis, regardless of what has occurred in earlier packets, or what will occur in later packets. Stateful packet filters, however, maintain state on connections and sessions, allowing filtering decisions to be made based on the status of the connection or session. (For example, by not allowing a connection to enter a state that does not makes sense.)

- c. How does packet filtering for firewalls, as discussed above, compare to the inspection and filtering capabilities of an Intrusion Detection System (IDS)? [4 marks]

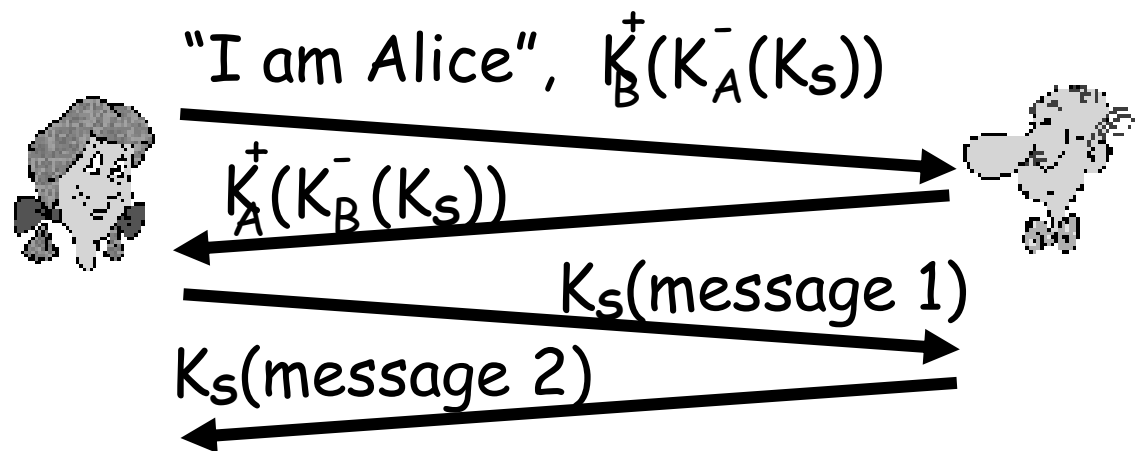
Packet filtering, as described above, only examines IP headers or TCP/UDP headers for individual packets, or packets within the same connection or session. While this gives a fair bit of filtering power, it does not allow filtering based on packet contents.

An IDS, however, allows for deep packet inspection, where packet contents can be examined and used for decisions being made. Furthermore, an IDS allows for decisions to be based on packets from multiple connections or sessions. (For example, this allows the detection of port scanning attempts, and other similar attacks.)

33. The following questions deal with secret, public key, and hybrid approaches to cryptography for computer networks. [12 marks]
- a. Compare and contrast the secret and public key approaches to cryptography in terms of their cryptographic strength, convenience, and performance. [6 marks]

With proper keys and a solid algorithm, both approaches have comparable cryptographic strength. Public key methods are more convenient, as there does not need to be a shared secret key between two entities for them to be able to communicate, making key distribution more manageable and convenient. Secret key approaches, however, tend to perform better than public key approaches, particularly on large quantities of data.

- b. Consider the hybrid secret-public key authentication protocol given below:



Explain why Alice and Bob have mutually authenticated each other, and why the session key  $K_S$  selected by Alice can be trusted for future communication. (You can assume the usage of sufficient nonces to ensure message freshness.) [6 marks]

Alice has authenticated herself to Bob because of her use of her private key in encrypting the session key. If Alice's public key decrypts this session key correctly, then Alice must have sent it, as only her private key could have produced it. As a result, Alice has been authenticated. Likewise, Bob authenticates himself as he sends the session key back to Alice encrypted using his private key. Since his public key can decrypt this session key, he must have sent this message. (Furthermore, only Bob has the private key to decrypt the message received from Alice in the first place; likewise only Alice has the private key necessary to decrypt the message Bob sent, thereby adding another level of protection.) Thus, Alice and Bob have authenticated each other. The session key can be safely used since the use of Bob's and Alice's public keys in encrypting messages prevents anyone else other than Bob or Alice with the appropriate corresponding public keys from decrypting the messages to recover the keys.

This page has been left intentionally blank. Use it as additional workspace or extra space for answers if necessary.