

**THE UNIVERSITY OF WESTERN ONTARIO
LONDON CANADA**

**COMPUTER SCIENCE 457a
MIDTERM EXAMINATION
OCTOBER 28, 2006
2 HOURS**

NAME: __ Marking Scheme _____

STUDENT NUMBER: _____

Question

1-25. _____

26. _____

27. _____

28. _____

29. _____

30. _____

31. _____

32. _____

33. _____

34. _____

TOTAL _____

(Out of 120 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. The period for transmitting Real Time Control Protocol (RTCP) packets for a sender depends on the:
 - a. Average RTCP packet size.
 - b. Number of senders.
 - c. Session bandwidth.
 - d. All of the above.
 - e. None of the above.

2. In Internet Differentiated Services, while a Per-Hop Behaviour defines differences in performance among classes, it does not mandate any particular mechanism for achieving these performances.
 - a. True.
 - b. False.

3. Suppose that the Weighted Fair Queuing (WFQ) scheduling policy is applied to a buffer that supports three classes, and suppose the weights are .5, .3, and .2 for classes 1, 2, and 3, respectively. Also suppose that each class has a large number of packets in the buffer. Then the classes could be served in which of the following sequences:
 - a. 121312131213121312131213.....
 - b. 121312131212131213121213121312.....
 - c. 123123123123123....
 - d. Any of the above.
 - e. None of the above.

4. Multimedia data has which of the following fundamental characteristics:
 - a. It is loss intolerant and delay intolerant.
 - b. It is loss intolerant and delay tolerant.
 - c. It is loss tolerant and delay intolerant.
 - d. It is loss tolerant and delay tolerant.
 - e. None of the above.

5. Which of the following are continuous forms of media?
 - a. Text.
 - b. Audio.
 - c. Video.
 - d. Audio and video.
 - e. Audio, video, and text.

6. The same amount of bandwidth tends to be given to both luminance and chrominance signals in video transmission.
- a. True.
 - b. False.
7. How can one improve the digitization of an analog audio wave?
- a. Decrease the sampling rate and decrease the number of quantization bits.
 - b. Decrease the sampling rate and increase the number of quantization bits.
 - c. Increase the sampling rate and decrease the number of quantization bits.
 - d. Increase the sampling rate and increase the number of quantization bits.
 - e. None of the above.
8. Which of the following redundancies can be found in a digital video stream but not in a single still image?
- a. Spatial.
 - b. Temporal.
 - c. Chromatic.
 - d. All of the above.
 - e. None of the above.
9. A still image painted on the screen at 20 frames per second will not be jerky, but will flicker because the images will decay from the retina too quickly.
- a. True.
 - b. False.
10. A movie with 20 different frames per second, each painted four times in a row won't flicker, but will be jerky.
- a. True.
 - b. False.
11. Using TCP for multimedia transmission usually requires a larger playout delay than if UDP was used instead.
- a. True.
 - b. False.
12. Which of the following does the Real Time Steaming Protocol (RTSP) standard do:
- a. Define the encodings, compression, and encapsulation of audio and video over the network.
 - b. Restrict the transport protocol allowed to transport streamed media.
 - c. Specify how media players must buffer audio and video data.
 - d. All of the above.
 - e. None of the above.

13. The Real Time Protocol (RTP) is used in what part of the network?
- a. End systems.
 - b. Network routers.
 - c. Both end systems and routers.
 - d. Neither end systems nor routers.
14. Suppose you have two streams of packets, one at 100 packets per second, and another at 6000 packets per minute. Since these have the same long term average rate, they must also have the same peak rate.
- a. True.
 - b. False.
15. Suppose a Certificate Authority (CA) has Bob's certificate registered with it, binding Bob's public key to Bob. This certificate is signed with
- a. Bob's public key.
 - b. The CA's private key.
 - c. Bob's private key.
 - d. The CA's public key.
 - e. None of the above.
16. It is possible to construct a polyalphabetic cipher out of multiple monoalphabetic ciphers.
- a. True.
 - b. False.
17. If a network adapter is in promiscuous mode, it means:
- a. The network adapter is malfunctioning.
 - b. The network adapter can read all network traffic.
 - c. The network adapter is constantly writing to the network.
 - d. The network adapter is encrypting all transmissions.
 - e. None of the above.
18. Which of the following poses a potential risk in handling backup copies of data:
- a. Protecting backups from being overwritten.
 - b. Hiding all backups in a locked location on-site to prevent theft.
 - c. Encrypting backups to protect their contents.
 - d. Verifying all backups to ensure they are properly working.
 - e. None of the above.
19. The Secure Sockets Layer (SSL) provides:
- a. Encryption for messages sent by both client and server.
 - b. Server authentication.
 - c. Optional client authentication.
 - d. All of the above.
 - e. None of the above.

20. Suppose Bob is purchasing merchandise from Alice Inc. over the Internet. SSL permits:
- a. Bob to determine whether Alice Inc. is authorized to accept credit card purchases.
 - b. Alice Inc. to determine if Bob has a good credit history.
 - c. Bob to determine if Alice Inc. is a legitimate company.
 - d. All of the above.
 - e. None of the above.
21. SSL is intended to provide security at which layer of the Internet reference model:
- a. Physical layer.
 - b. Data link layer.
 - c. Network layer.
 - d. Transport layer.
 - e. None of the above.
22. IPsec is intended to provide security at which layer of the Internet reference model:
- a. Physical layer.
 - b. Data link layer.
 - c. Network layer.
 - d. Transport layer.
 - e. None of the above.
23. IPsec can be used with either TCP or UDP data.
- a. True.
 - b. False.
24. When using IPsec's Authentication Header (AH) protocol, the type of payload (TCP, UDP, ICMP, etc.) is specified in the AH header.
- a. True.
 - b. False.
25. With IPsec's Encapsulation Security Payload (ESP) protocol, an intruder cannot determine the type of payload (TCP, UDP, ICMP, etc.) because the payload type is encrypted.
- a. True.
 - b. False.

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

26. The following questions deal with Forward Error Correction (FEC). [8 marks total]

- a. In class we discussed two different FEC schemes for compensating for errors and loss in media streams. Briefly summarize each approach. [4 marks]

The first approach is based on XOR operations. In this approach, the stream of data packets is broken into sequences that are n packets long. An $n+1^{\text{st}}$ packet is constructed for each sequence of n packets from the XOR-ing of the n packets together. If a single packet out of this sequence of $n+1$ packets is lost, the lost packet can be reconstructed by performing a similar operation on the n remaining packets that were actually received.

The second approach is based on piggybacking ... a lower resolution stream of data is sent along with the original higher resolution stream such that the lower resolution version of chunk x arrives with the higher resolution chunk $x+1$. That way, if a packet is lost, the lower resolution counterpart can be recovered from the next message received without a total loss occurring. This approach can also be generalized to allow multiple lower resolution data chunks to accompany each higher resolution chunk to provide additional redundancy.

- b. Both approaches to FEC discussed in class increase the transmission rate of the stream by adding overhead. Does interleaving to compensate for loss or errors also increase the transmission rate? Justify your answer. [4 marks]

No. Since interleaving does not impose additional space overhead, as no redundant information is transmitted, it does not affect the bandwidth requirements of a data stream. Consequently, there is no increase in the transmission rate of the data stream either.

27. The following question parts deal with adaptive playout delays. [8 marks total]
- a. What is the general goal of using adaptive playout delays in real-time interactive multimedia? [4 marks]

The general goal of an adaptive playout delay is to minimize the playout delay of a stream of data while at the same time keeping the late loss rate acceptably low.

- b. Recall the adaptive playout strategy discussed in class for our Internet phone example (with talk spurts and such), where audio data was accumulated and transmitted every 20 msec during a talk spurt. How can two successive packets received at the destination have timestamps that differ than more than 20 msec when the two packets belong to the same talk spurt? [2 marks]

If a packet is lost, it is possible that two successive packets have timestamps differing by more than 20 msec even if the packets are in the same talk spurt. (For example, suppose there are packets A, B, and C with timestamps of 100, 120, and 140. If packet B is lost but A and C are received, the difference between the timestamps in A and C is more than 20 msec.)

- c. How can the receiver in our Internet phone example use timestamps and sequence numbers to tell if a packet is the first one in a talk spurt? Be specific. [2 marks]

Suppose that during a talk spurt that packets are generated every x msec. If the difference in timestamps between two successive packets is greater than x msec and there is no gap in the sequence numbers of the packets, then a new talk spurt has begun with the second packet in the sequence being the first in the new talk spurt.

28. The following questions deal with Quality of Service (QoS). [8 marks total]
- a. What are the differences between hard QoS requirements and soft QoS requirements? Give an example of an application with hard requirements, and an example of an application with soft requirements. [4 marks]


If an application has hard QoS requirements, the application is not functionally correct if its QoS is missed. In this case, the application provides no or negative value if it cannot meet its QoS requirements. Examples of this include flight control, patient monitoring, and nuclear reactor control.

If an application has soft requirements, the application is still functionally correct if its QoS is missed. In this case, the application still provides some or perhaps all of its potential value, even if it cannot always meet its QoS requirements. Most multimedia applications fall into this category.

- b. Can the problem of providing QoS guarantees be solved simply by “throwing enough bandwidth” at the problem, that is, by upgrading all link capacities so that bandwidth limitations are no longer a concern? Explain your answer. [4 marks]

Unfortunately, the answer is no. First of all, you will never have enough bandwidth to be able to solve all of your needs. You will be stuck in an endless cycle of increasing bandwidth in an attempt to keep everyone happy. Most importantly though, having ample bandwidth does not necessarily help anything if that bandwidth isn't properly managed. Applications can still compete too much and starve each other out, causing shortages of bandwidth, but also causing impacts on delay or delay jitter. If flows are not isolated from one another, even having lots of bandwidth does not necessarily mean anything, because they can still negatively affect one another. While throwing bandwidth at the problem might help somewhat in the short term, it is ultimately not going to work in the long term.

29. Explain how the leaky token bucket policing mechanism works. Use a diagram in your answer. [8 marks total]



In this scheme, a packet can only be transmitted when there is a token available in the bucket; when it is transmitted, the token is removed from the bucket. Tokens fill the bucket at a rate of r tokens/second, unless the bucket is already full to its capacity of b tokens. This limits the long-term average rate of transmission to r packets/second ... if this amount is exceeded for an extended period, the bucket will drain and there will be no tokens for transmission available. This also permits a burst size of b tokens, because the transmitter can burn through all of the available tokens without having to pause or wait to continue. In the end, over a time period of length t , the number of transmitted packets is limited to $r*t + b$.

30. In a group of N people, suppose each person wants to communicate with each of the $N-1$ other people in the group using symmetric key encryption. All communication between any two people, i and j , is visible to all other people, but no other person should be able to decode their communication. How many keys are required in the system as a whole? Now suppose that public key encryption is used instead. How many keys are required in this case? [6 marks total]

With N people total, we would need $N*(N-1)/2$ symmetric keys. We need one symmetric key between each pair of people in the group of N . How many pairs are there? There are N choices for the first person, and $N-1$ choices for the second in forming pairs. Also, ordering in the pair doesn't matter, so we have $N*(N-1)/2$ pairs, and consequently $N*(N-1)/2$ keys required.

For public key cryptography, we simply need $2*N$ keys. Each person would need their own public and private keys, so we would have 2 keys per person, or $2*N$ keys.

31. The following questions deal with message digests. [6 marks total]
- a. In what way does a message digest provide a better message integrity check than a regular checksum such as the Internet checksum? [3 marks]

A message digest provides a better message integrity check than a regular checksum because checksums are relatively easy to fool, and it is not difficult to modify a message so that it will still have the same checksum computed for it. Message digests are stronger, having typically been generated by a one way hash function that is difficult to reverse or predict. It is exceedingly difficult to modify a message in a way that produces the same message digest, although nothing prevents the modification of the digest as well.

- b. In what way does a public key encrypted message digest provide a better digital signature than using the public key encrypted message? [3 marks]

A public key encrypted message digest is a better digital signature because the digest is potentially much smaller than the original message, and so the encryption activity both to sign the digest and verify the signature will be much more efficient. Since the digest is mathematically tied to the original message, signing the digest is pretty much as strong as signing the whole message, cryptographically speaking. (Assuming you have a good digesting function, that is ...)

32. The following questions deal with authentication. [10 marks total]
- a. What is the purpose of authentication? [2 marks].

The purpose of authentication, in essence, is to have a reliable method of establishing the identity of someone else.

- b. What are the three main ways of authenticating an individual discussed in class? Provide an example of each method. [6 marks]

Something you own (basically an object in your possession that identifies who you are), such as a physical key, security card or badge, or a one-time password generator.

Something you know (basically a secret piece of information that only you should be in possession of), such as a password, cryptographic key, or the correct answer to a challenge-response test.

Something you are (typically some kind of biometric measurement), such as facial features, fingerprint, retina scan, voice print, and so on, plus a liveness test.

- c. Why is it usually unwise to ask someone for their public key directly when trying to authenticate them? [2 marks]

It is too easy for an imposter to offer you a valid public key that just happens to be their own. (Or at the very least, it does not belong to the individual the imposter is claiming it to be from.) This opens you up for “man-in-the-middle” style of attacks, which was demonstrated in the class notes.

33. The following questions deal with types of security threats. [8 marks total]
- a. What is meant by the term “mapping”, as it pertains to network security threats? Provide one possible defense that can be used to prevent it, or at least limit its effectiveness. [2 marks]

Mapping basically refers to determining the valid addresses and ports within a network, in essence to determine what exactly is there before attempting to launch an attack. A good way to defend against this is to protect your network with a default deny policy ... it will make it difficult for an outsider to determine very much about your network.

- b. What is meant by the term “port scanning”? How can one carry out a “port scan” in a way that is more difficult to detect by intrusion detection systems? [2 marks]

Port scanning, in essence, is a type of mapping technique in which an attacker attempts to establish TCP connections with the ports on a machine to see which ports have services listening on them. Scanning ports sequentially from a single address is usually easy to detect ... breaking the scanning to come from multiple sources or to hit ports in a non-sequential fashion over a longer period of time is harder to detect. Also, limiting scans to only certain port numbers also makes the scan difficult to detect.

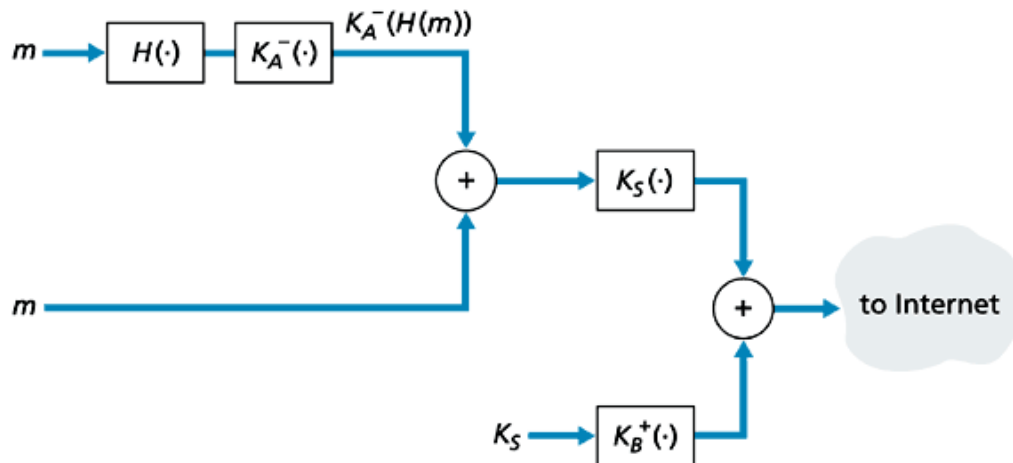
- c. What is meant by the term “packet sniffing”? Provide an example of a countermeasure that can be used to help stop it. [2 marks]

Packet sniffing generally refers to the process of intercepting packets intended for another destination, and reading their contents. Typically this involves using a system on a broadcast network that has its network interface set to promiscuous mode. To help prevent such things, one should lock down promiscuous mode access to systems, ensure only one system is on every network segment, and encrypt sensitive information.

- d. What is meant by the term “IP spoofing”? Provide an example of a countermeasure that can be used to help stop it. [2 marks]

IP spoofing refers to the creating of IP packets with forged source address fields, thereby fooling the recipient into believing that they originated elsewhere. The best strategy against this is to ensure routers don't forward on packets from impossible sources, but this is difficult to mandate for entire Internet.

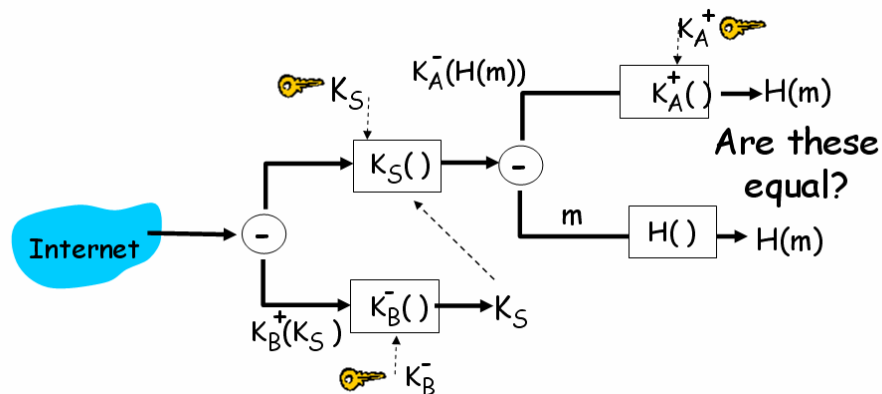
34. Suppose Alice sends a message m to Bob as shown in the diagram below. [8 marks]



a. In sending the message m in this fashion, what is provided in terms of message confidentiality, authentication, and integrity? Justify your answer. [4 marks]

Message confidentiality is provided because ultimately everything is encrypted using a symmetric session key, and only Bob can access this key using his private key. Authentication is provided because a message digest for the message has been signed using Alice's private key; the use of her public key to decrypt the digest will verify its origin. Integrity is also ensured because of the signed message digest ... if the message has been tampered with, it will no longer match the digest, and an attacker cannot generate a new digest and re-sign it because they do not have Alice's private key.

b. Construct a similar diagram showing the steps Bob will need to take to make use of the message m in its original form. If authentication and integrity checks are possible, be sure to include these in your answer. [4 marks]



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