

**THE UNIVERSITY OF WESTERN ONTARIO
LONDON CANADA**

**COMPUTER SCIENCE 457a
MIDTERM EXAMINATION
OCTOBER 29, 2005
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. Which of the following methods of recovering from packet loss do not incur any redundancy overhead:
 - a. Forward Error Correction (FEC) using exclusive ORs.
 - b. Forward Error Correction (FEC) using piggybacking.
 - c. Interleaving.
 - d. All of the above.
 - e. None of the above.

2. A differentiated services router distinguishes between packet flows in implementing different Per-Hop Behaviours by using the packets':
 - a. Source IP address, destination IP address, and packet markings.
 - b. Source and destination IP addresses.
 - c. Source and/or destination port numbers.
 - d. Packet markings alone.
 - e. None of the above.

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, .2, and .3 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. 12131213121213121312131213121312.....
 - c. 123123123123123....
 - d. Any 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 can use any underlying transport protocol it chooses.
 - a. True.
 - b. False.

5. Sequence numbers are not necessary for the adaptive playout strategy for real-time interactive multimedia applications discussed in class.
 - a. True.
 - b. False.

6. When using RTP, it is possible for a sender to change encodings in the middle of a session.
 - a. True.
 - b. False.

7. Suppose a flow gets assigned a certain weight in the weighted fair-queuing discipline. At this point, the minimum throughput for the flow is fixed, and is independent of the weights assigned to the other flows.
- a. True.
 - b. False.
8. In RSVP, to maintain reservations in routers, receivers send refresh messages.
- a. True.
 - b. False.
9. Suppose an RTP session has a separate video and audio stream for each session. The audio and video streams in each session can be synchronized using only the timestamp fields from the RTP packet headers.
- a. True.
 - b. False.
10. SIP mandates that all SIP clients support particular audio and video encodings.
- a. True.
 - b. False.
11. 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.
12. 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.
13. The Real Time Protocol (RTP) is intended to provide a real-time service that can provide timely delivery of data to provide quality of service guarantees.
- a. True.
 - b. False.
14. The theft of laptop computers, as discussed during class, is a serious concern at the University of Western Ontario.
- a. True.
 - b. False.

15. Suppose Bob wants to send a secret message to Alice using public key cryptography. 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 public key and send Alice the message.
 - d. Encrypt the message with his private 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. In this case, Bob should:
- a. Encrypt the message with his private key and send Alice the message.
 - b. Encrypt the message with his public key and send Alice the message.
 - c. Encrypt the message with Alice's private key and Alice the message..
 - d. Encrypt the message with Alice's 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. In this case, Bob should:
- a. Encrypt the message with Alice's public key, encrypt the result with his public key and then send the message to Alice.
 - b. Encrypt the message with his private key, encrypt the result with Alice's private key, and then send the message to Alice.
 - c. Encrypt the message with his private key, encrypt the result with Alice's public key, and then send the message to Alice.
 - d. Encrypt the message with his public key, encrypt the result with Alice's public key, and then send the message to Alice.
 - e. None of the above.
18. In IP spoofing, the attacker interchanges the source and destination addresses in the IP packets it sends.
- a. True.
 - b. False.
19. Two parties can use public-key encryption to agree on a shared one-time symmetric session key.
- a. True.
 - b. False.
20. The ROT13 algorithm is the same as a Caesar cipher with the key $k=13$.
- a. True.
 - b. False.

21. In public key cryptography, if a public key can decrypt a message, that message must have been encrypted with:
- a. A symmetric key.
 - b. A session key.
 - c. The same public key.
 - d. The corresponding private key.
 - e. None of the above.
22. In public key cryptography, if a private key can decrypt a message, that message must have been encrypted with:
- a. A symmetric key.
 - b. A session key.
 - c. The corresponding public key.
 - d. The same private key.
 - e. None of the above.
23. When a public key is signed by any Certificate Authority, you know that public key can always be trusted.
- a. True.
 - b. False.
24. The Secure Sockets Layer (SSL) can be used with either TCP or UDP sockets.
- a. True.
 - b. False.
25. Intermediate routers process Authentication Header (AH) and Encapsulation Security Payload (ESP) protocol messages as they would regular data messages.
- 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 scheduling and scheduling policies. [8 marks total]
- In class, we discussed non-preemptive priority queuing for networks. What would be preemptive priority queuing in a network? Does preemptive priority queuing make sense for computer networks? Explain. [4 marks]

Preemptive priority queuing in a network would amount to allowing the interruption of the transmission of a lower priority packet if a higher priority packet became available for sending, so that the higher priority packet could be sent out immediately.

Preemptive priority queuing does not make sense for computer networks, because partially sent packets that are interrupted are useless to the recipient, and consume network bandwidth needlessly in transit. It is better to finish the packet first, and then send the higher priority data.

- 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]

In WFQ, you have multiple classes of packets, each with own queue. The scheduler cyclically scans each class queue, giving each class a weighted amount of service in each cycle, allowing certain classes to be able send more packets, or send packets at a higher rate. Each class may get different amounts of service compared to the other classes, depending on the weights. WFQ can act like a round robin approach if the weight for each queue is set to be 1.

27. The following question parts deal with multimedia data and applications. [8 marks total]
- a. In streaming multimedia applications, what requirement is placed on all data yet to be transmitted? [2 marks]

In a streaming application, the requirement placed on all data to be transmitted is that it must be received in time for playout.

- b. When we say that audio and video are continuous forms of media, what do we mean? [2 marks]

We mean that audio and video have a time-wise component; in other words they are not a snapshot in time like a still image. Instead, they flow in a continuous fashion in time.

- c. Compare and contrast constant and variable bit rate compression for audio and video data. [4 marks]

Constant bit rate compression varies the quality and ratio of compression to maintain a consistent bit rate through a data stream, while variable bit rate compression varies the bit rate of the compressed stream to maintain a consistent level of quality. Consequently, constant bit rate compression had consistent and predictable resource requirements, but either sacrifices quality at times or wastes space at other times. Variable bit rate compression has quality maintained without wasting resources, but at the cost of predictable resource consumption.

28. The following questions deal with comparing competing approaches to multimedia communications. [8 marks total]
- Compare and contrast the strengths and weaknesses of using TCP versus UDP for the transmission of multimedia data. [4 marks]

UDP is not affected by congestion control and does not use timeouts, which could potentially allow more data through and can keep the playout delay of multimedia data lower. This comes at the cost of unreliable transfer, meaning such mechanisms (if needed) must be built into the application (but multimedia data is resistant to some loss anyways ...) UDP data can have issues at firewalls. TCP must respect the rules of congestion control and retransmissions resulting in more variability and fluctuations in transfer rates, which in turn requires a larger playout delay in compensation. Data transfer is reliable, but at this cost. TCP data, however, can usually traverse firewalls more easily.

- Compare and contrast the strengths and weaknesses of Integrated Services versus Differentiated Services for the next generation Internet. [4 marks]

Integrated services has scalability issues because it manages resources and maintains flows on a per-flow basis. Since Differentiated services deals with aggregates of flows into classes, it scales more easily. Integrated services lacks flexibility in defining service models, but has flexibility in that it can tailor service to particular flows. Differentiated services has more flexibility in terms of defining new service models, but cannot tailor service to particular flows. Both integrated services and differentiated services suffer from the fact that they need widespread adoption in order to be successful.

29. The following questions deal with working with audio data. [8 marks total]
- a. 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. What can you do in the conversion process to improve the approximation of the digital representation to the original sound wave? [4 marks]

In doing this conversion the audio wave must be sampled at a fixed interval (which can lose the continuity of the analog wave) and then quantized into a fixed number of bits (which can lose accuracy of representation of the wave). Since these two things must be done in the conversion, we will always lose information and therefore the resulting wave is never quite the same as the original. To improve the approximation, we can increase the sampling rate and increase the number of bits used to store each sample during quantization.

- b. How does Differential Pulse Code Modulation (DPCM) improve over regular Pulse Code Modulation (PCM)? How does Adaptive Differential Pulse Code Modulation (ADPCM) improve over Differential Pulse Code Modulation (DPCM)? [4 marks]

DPCM exploits the fact that, for most audio signals, the range of differences between successive samples is much less than the range of possible samples. Only the digitized difference between samples is encoded, which tends to require fewer bits than a comparable PCM signal. ADPCM goes further by recognizing that additional bandwidth savings (or improved quality) can be obtained by varying the number of bits used for the difference signal. Fewer bits encode smaller difference values than larger values.

30. Give three reasons why network security tends to be difficult in practice. [3 marks total]
- Open and interoperable protocols, while desirable, tend to work against security.
 - Security is often sacrificed in return for gains in performance and scalability.
 - Providing good security is expensive, so it can be difficult to get resources to support it.
 - People tend to see security as a barrier to getting useful work done, and resist it.
 - Information on circumventing security is widely available, as are software tools.
 - Some people see circumventing security as a challenge and enjoy doing it.
 - Most systems and networks were not designed with any security concerns in mind.

31. The following question parts deal with principles of cryptography. [8 marks total]
- a. For public key cryptography to be secure, why is it important that it is difficult to deduce the private key when given the public key? [4 marks]

The public key is circulated to everyone in the general public. If one could deduce a private key given a public key, it would be simple for an attacker to retrieve someone's public key, generate the private key, and pose as the individual in question. This would effectively break the cryptosystem.

- b. For a message digest to be useful, why is it important that the function used to create it be a one-way function that is difficult to reverse or predict? [4 marks]

Suppose one could reverse or predict the results of the digest function, and this function is used to produce a digest for a message. In this case, an attacker could more easily generate an altered message that results in the same digest as the original message, making the alteration undetectable, even if the digest is digitally signed. (The digest matches, so there is no reason to attempt to reproduce and re-encrypt it.)

32. The following questions deal with failures and their impact on security. [10 marks total]
- a. What is meant by the term “central point of failure”, as discussed in class? Why is that having a “central point of failure” is particularly bad for the security of a network? [4 marks].

A “central point of failure” refers to having a single critical system in your network that, if made unavailable, seriously disrupts the services provided within your network. If that system were to fail, it could have dire consequences. Having such a central point of failure is bad because it makes your network particularly open to a denial of service attack. If the central point of failure is forced to fail or go offline in such an attack, this could greatly disrupt the functioning of the network.

- b. Consider Key Distribution Centre (KDC) and Certification Authority (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]

If the KDC goes down, no one will be able to open new channels of communication to other entities, as the KDC is required to negotiate session keys between individuals. Those that already have session keys can continue to use them without any problems whatsoever for their current sessions. If a CA goes down, this only inhibits the ability to issue new certificates. Any existing certificates can still be validated and used to get public keys safely to open new communication sessions. So long as the certificate has already been granted, communication can begin and proceed as normal.

33. The following questions deal with principles of security. [9 marks total]
- a. What is meant by the following statement, as it pertains to network security: “Do unto yourself before someone does unto you.” Why is this a very important idea for network security? [3 marks]

In essence, this means that you should probe your network and discover your own vulnerabilities before someone else does. This is important because if someone else discovers your weaknesses first they can exploit them until you notice the problem and fix it. This could result in very serious problems.


- b. What is meant by the following statement: “Security through obscurity doesn’t work.” Why is this the case? [3 marks]

The statement means that hiding or keeping secret the security-related details of your system and network will not ultimately work in the end. This is because things that are hidden can be found, and things that are meant to be kept secret can be divulged, either accidentally or maliciously. Either way, if your security depended on these things being hidden or secret, and they are now out in the open, your security is severely compromised.

- c. What is meant by the following statement: “Technology alone won’t make you safe.” Why is this the case? [3 marks]

This statement basically means that security cannot be provided through technological means alone. The best hardware and software will not protect you on its own. This is because security is very much a people or social problem, and must be treated in this fashion ... more often than not it is people with bad security habits that cause more security problems than anything else.

34. Consider the key distribution center protocol shown below. [8 marks total]



- a. Why doesn't Alice need to explicitly and directly authenticate Bob before using the session key R1? [4 marks]

In this case, Alice communicates with the KDC to get a session key R1 to communicate with Bob. Since she uses the key she shares with the KDC to talk to the KDC, and this key is only shared between her and the KDC, she can trust that the data she gets back is from the real KDC. Included in this is a token that is encrypted using the key shared between Bob and the KDC. Since only Bob has this key besides the KDC, she can trust that only Bob can decrypt this token and get R1. Since only Bob can get R1 besides Alice (and only Alice has the shared secret key with the KDC to decode the response from the KDC), this implies that if Alice uses R1, only Bob can respond. Consequently Bob is implicitly authenticated.

- b. Why doesn't Bob need to explicitly and directly authenticate Alice before using the session key R1? [4 marks]

When Bob receives the token from Alice, he knows it originated from the KDC, since it is encrypted using the secret key he shares only with the KDC. Since Bob trusts the KDC, he can trust that R1 is a valid session key. Knowing the protocol used, Bob knows that Alice is the only other individual that could have this key, since it would have been encrypted with the key she shares with the KDC. Consequently, Alice is implicitly authenticated, both by the token from the KDC Alice forwards to Bob, and through Alice's use of R1.

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