Internetworking (2G1305) Examination
Tuesday 31-May-05 9:00-13:00

Instructor: G. Q. Maguire Jr.

- ***No help material is allowed.***
- You may answer the questions in English or Swedish.
- The questions for ***section A*** should be answered ***on the exam itself***, for the other sections the questions should ***each*** be answered on a ***separate*** page.
- **For multiple-choice questions, selecting a wrong alternative will reduce the number of points** (with a lower limit of 0 points for the problem)**.**
- The ***entire*** exam must be turned in along with your answers.

The exam is divided into three sections:

- ***Section A*** consists of multiple-choice questions. Each question is worth two points - if all correct alternatives (*regardless of how many there are*) are selected. Each missed correct alternative will reduce the score by one point. Each selected alternative that is wrong will also reduce the score by one point. The total score for each question will not be lower than zero.
- ***Section B*** consists of short questions. Each correctly answered questions is worth 2 points.
- ***Section C*** consists of essay questions where a more elaborative answer is required. A correct answer is worth four points.

The exam grades will be:

- Grade 3: at least 26 points in section A alone.
- Grade 4:
  - ♦ at least 28 points in section A and
  - ♦ at least 4 points in section B and
  - ♦ at least 4 points in section C
- Grade 5:
  - ♦ at least 30 points in section A and
  - ♦ at least 6 points in section B and
  - ♦ at least 8 points in section C

Results will be announced on the institution’s announcement board - before 2005-06-21.

Name/Your name: ____________________

Personnummer/student number: ____________________

Utbildningslinje/Your major ____________________

If you are a student from 2G1507 or another version of the course please indicate the course number: ____________________

Initials/Initials:	Section A:
Inlämnat/Handed in: :	Section B:
Antal sidor/Number of pages:	Section C:
	Betyg/Grade:

Please check off which questions you have answered:

Fråga/Questions	**Besvarad/**Answered	**Fråga är värd/** Question worth	**Rättning/** Grading	Subtotals
1	❑	2		
2	❑	2		
3	❑	2		
4	❑	2		
5	❑	2		
6	❑	2		
7	❑	2		
8	❑	2		
9	❑	2		
10	❑	2		
11	❑	2		
12	❑	2		
13	❑	2		
14	❑	2		
15	❑	2		
16	❑	2		
Total section A:				
17	❑	2		
18	❑	2		
19	❑	2		
20	❑	2		
Total section B:				
21	❑	4		
22	❑	4		
23	❑	4		
24	❑	4		
Total section C:				
			Total	

Section A: Multiple choice

1. RIP version 1 uses which of the following for its transport protocol?

❍ TCP
❍ UDP <--- correct answer

2. If we use firewalls from two *different* manufacturers in *series*, what basic security principle does this exemplify?

❍ flow control
❍ virtual circuits
❍ least privilege
❍ delayed acknowledgments
❍ defense in depth <-- answer
❍ weakest links
❍ encryption

3. Is there any checksum protection in an ARP request or reply?

❍ Yes
❍ No <-- answer
No. {As this is expected to be a link local transfer it depends on the link layer checksum.}

4. In the Resource Reservation Protocol (RSVP) does the source or the destination (receiver) determine the required resources?

❍ Source
❍ Destination <-- answer
The destination (receiver) sends the actual reservation message that establishes or updates the reservation state.

5. When an node wants to send an IP packet on an ethernet segment it uses ARP to find the destination's link address. When a node wants to send an IP packet on a PPP link does it need to transmit anything to learn the destination's ***link*** address?

❍ Yes
❍ No <-- answer
No. It does not need to know the link address of the destination, since it is a point to point connection -- hence it already knows the link address to use, therefore it simply frames the IP packet and sends it.

6. The “Version” field is the same in IPv4 and IPv6, but is this actually used to differentiate between IPv4 and IPv6 packets in the case of an ethernet (IEEE 802.3) physical and an IEEE 802.2 logical link layer.

   ❍ Yes

   ❍ No <-- answer

   No, these packets have been assigned different link types, hence the driver software knows whether to send it to the IPv4 or to the IPv6 stack, therefore there is no intermediate IP level which dispatches to the particular version of the IP stack which is needed. Hence the version field is not actually used when processing these packets (except for being part of what is transmitted and its participation in checksums).

7. When a destination host receives only some of the fragments from an IP packet what happens?

   ❍ There will be a IP checksum error and the packet will be discarded

   ❍ There will be a reassembly timeout and the packet will be discarded <-- answer

   ❍ There will be a IP checksum error and the packet will be sent up the stack

   ❍ There will be a reassembly timeout and the packet will be sent up the stack

   ❍ There will be both an IP checksum error ***and*** a reassembly timeout, then the packet will be sent up the stack

   (the second choice): There will be a timeout (i.e., time-to-live equals 0) during IP reassembly, after which the destination host will send an ICMP time exceeded message (type 11) with code = 1. The destination host will discard all the received fragments. Note: the source is able to distinguish this from a transit timeout, because this later has code = 0.

8. PPP provides the following advantages as compared to SLIP:

   ❍ PPP supports multiple protocols, while SLIP only supports IP <-- answer

   ❍ PPP uses a fixed IP address, manually configured by the user

   ❍ PPP provides TCP and IP header compression <-- answer

   ❍ PPP can negotiate many data-link options <-- answer

   ❍ PPP provides a file transfer checksum

   - support for multiple protocols on a single serial line (not just IP datagrams)
   - a cyclic redundancy check on every frame
   - dynamic negotiation of the IP address for each end (using the IP network control protocol)
   - TCP and IP header compression similar to CSLIP
   - a link protocol for negotiating many data-link options

9. DHCP always assigns a client an IP address for an ***indefinite*** period of time

   ❍ Yes

   ❍ No <-- answer

   No: A DHCP server leases an IP address to a client for a period of time. By only leasng the IP address it is easier to allocate (and reallocate) addresses for clients.

10. Explain the purpose of a “hold down” when doing routing updates

❍ When a route is removed, no update of this route is accepted for some period of time - to give everyone a chance to remove the route. <-- answer

❍ When a route is removed, all updates of this route are processed immediately - to give everyone a chance to insert the correct route.

11. IPv6 fragments contain which of the following:

❍ Fragment offset <-- answer

❍ Do not fragment flag

❍ More fragments flag <-- answer

12. Which of the following are timers used in conjunction with ***TCP*** :

❍ Retransmission timer <-- answer

❍ Persist timer <-- answer

❍ Fragment reassembly timer

❍ Keepalive timer <-- answer

❍ 3 MSL timer

● Retransmission timer: TCP requires an acknowledgment for each byte of sent data. The acknowledgments are cumulative, i.e. one ACK acknowledges all data up to this sequence number. The retransmission timer is used to detect loss of segments (either the actual data segment or the segment containing the ACK). The value of the retransmission timer, RTO (retransmission timeout) is adjusted according to the current network situation based upon the measured round trip times. To keep up with wide fluctuations in the RTT the RTO calculation keeps track of the mean deviation in the RTT measurements, in addition to the smoothed RTT estimator.

● Persist timer: In TCP the receiver performs flow control by specifying the amount of data it is willing to accept. If all the receiver’s buffers are filled it will announce a window size of zero, thus stopping the sender from transmitting any further data until it receives a non-zero window size advertisement. If the ack segment containing this advertisment is lost a deadlock occurs: the receiver waits for data, the sender waits for the window to open. To prevent this deadlock, the sender uses the persist timer which causes it to query the receiver periodically, to find out if the window has been increased (window probes).

● Keepalive timer: the keepalive timer is not part of the TCP specification but is provided by many implementations. It’s intention is to allow a host to find out if an (idle) connection is still up or if the other end has crashed/rebooted. The use of keepalives is controversial as they can cause connections to be dropped during transient failures, consume unnecessary bandwidth and cost money on networks that charge by the packet.

- 2MSL timer: every implementation must choose a value for the *maximum segment lifetime* (MSL) - the maximum amount of time any segment can exist in the network before being discarded. When TCP performs an active close and sends the final ACK, that connection must stay in the TIME_WAIT state for twice the MSL. This lets TCP resend the final ACK in case this ACK is lost. Additionally, while a TCP connection is in the 2MSL wait, the socket pair defining that connection cannot be reused. Any delayed segments arriving for a connection in the 2MSL wait are discarded. Therefore when establishing a connection no delayed segments from an earlier incarnation of this connection can be misinterpreted as being part of the new connection.

13. A host using Mobile IP must tunnel packets to the corresponding host via its home agent when sending packets under which of the following conditions

❍ When using IPSEC
❍ When there is ingress filtering<-- answer
❍ When the mobile host is in the home network
❍ When the mobile host is in the same network as the corresponding host

14. An application layer gateway can be used to forward which of the following kinds of traffic through a NAT/Firewall:

❍ RTP<-- answer
❍ IPSEC<-- answer

15. SCTP provides which of the following:

❍ A reliable message oriented service<-- answer
❍ multihoming with load balancing
❍ multihoming without load balancing<-- answer
❍ a variable number of streams<-- answer
❍ a fixed number of streams
❍ The sender selects its primary own address
❍ The receiver selects the primary address for the other endpoint<-- answer
❍ An association is allocated resources when a COOKIE ECHO is received<-- answer
❍ An association is allocated resources when a COOKIE ACK is received
❍ An association is allocated resources when an INIT ACK is received

16. The BGP keepalive generates approximately how much traffic:

❍ 5 bps<-- answer
❍ 50 bps
❍ 50 kbps
❍ 50 Mbps

Section B: Short Answers

17. Many sites allow anonymous FTP access, but request that the user enter their e-mail address as the password. Explain how the site can check to see if the domain name of this e-mail address corresponds to the domain from which the user is connecting to this FTP server.

    The server can do a DNS reverse lookup (also called pointer query) to convert the incoming IP address (i.e., the user's source address) into a string name and compare this to the domain name string which the user enters as part of their e-mail address. If the user's mail system supports the VRFY command, the site could also check if the whole e-mail address corresponds to a real user {but it can't know if this corresponds to the user who is logging in!}

18. How many IPv4 addresses are associated with a single ethernet interface? Describe how this/these address(es) is (are) used.

    if the interface is down: 0 addresses
    If the interface is up:
    - at least 2 addresses - a host specific IP address and the link local broadcast address;
    - n addresses -- as there can be multiple secondary addresses, multiple multicast addresses, and the link local broadcast address; and
    - 2^{32} addresses if the interface is in promiscuous mode

19. When is the Nagle algorithm disabled? Given an example of this.

    As the Nagle algorithm is used to avoid tinygrams (i.e. packets with a very small amount of data, e.g. keystrokes from an interactive session) it must be disabled when we actually ***do*** want to quickly receive a small amount of data -- such as X windows mouse events, etc.

20. Describe the purpose of anycast in IPv6.

    Anycast is used to find an instance of a specific service (for example, router, DHCP server, etc.). This enables a node to send a packet to a generic address to get a specific service from the "nearest" instance. This puts the burden of determining which instance to deliver it to on the routing system.

Section C: Essay Answers

21. If a host is participating in a multicast group, why does it delay responding to an IGMP query?

    Since it is a *multicast* group, there may be more than one host which is a member of this group on a given link. Use of a random delay prevents all of the hosts from trying to respond at the same time. While delaying its IGMP response the host can listen to hear if another node reports membership in a group it is interested in, if so its response is cancelled; if not, then it will respond after waiting for its random delay.

22. Karn’s algorithm is a solution for the TCP retransmission ambiguity problem. Explain what this problem is.

    When a packet is transmitted and a timeout occurs, the RTO is backed off using TCP’s exponential backoff. The packet is retransmitted with the longer RTO and an acknowledgment is received. *The problem now is that TCP doesn’t know if the ACK is for the original or the retransmitted packet.*
    Karn’s Algorithm specifies that, when timeout and retransmission occur, the backed off RTO has to be used and the RTT estimators *must not* be updated until an acknowledgment is received for a segment that was not retransmitted -- this removes the ambiguity.

23. Describe the Security Parameter Index (SPI) used in the Authentication and ESP protocols in IPsec and IPv6.

    Authentication and encryption require that the sender(s) and receiver(s) agree on a key, on an authentication or encryption algorithm, and on a set of parameters such as the lifetime of the key or the details of the algorithm’s utilization. This set of data forms a security association between the sender(s) and the receiver(s). Each authenticated and/or encrypted packet contains a parameter called the Security Parameter Index (SPI) which is normally negotiated as part of the key exchange procedure. This SPI is used to link packets with the context of a security association. The SPI contained in the packet enables the receive to determine which security association (and hence which algorithms and keys) it should use with this packet.

24. Explain the steps in a DNS lookup of the name “www.kth.se” assuming that you are doing this from a host named “fred.it.kth.se”. Comment on the performance of DNS. {You should assume that the host “fred” is a UNIX/linux workstation, which is *not* using “yellow pages” or a similar service. In addition, assume that the local machine does *not* know the address corresponding to “www.kth.se”.}

    The query will begin with the local machine (fred.it.kth.se), if “fred” knows the address, then it returns the answer (according to the stated assumptions it does not know the answer), then it will use the DNS server (defined in /etc/resolv.conf -- note that on some UNIX machines the local search order may be different), if this DNS knows then it will answer, otherwise it will recursively go up the DNS tree until it finds an answer or it reaches a root DNS server (which by definition must know how to find the top level domain). In this case, it will find an answer before reaching a root level server, because the two machines share a common postfix. The key to the performance of DNS is that this recursive lookup will benefit from the caching of answers to earlier requests.