Chalmers University of Technology
Computer Science and Engineering
Elad Michael Schiller

2012-10-26

Written exam in EDA387/DIT661 Computer Networks 2012-10-26. Exam time: 4 hours.

Means allowed: Nothing except paper, pencil, pen and English - xx dictionary.

Examiner: Elad Michael Schiller, phone: 073-6439754 and 031-7721052

Credits:	30-38	39-47	48-Max
Grade:	3	4	5
Grade (GU)	G	G	VG

1. The answer must be written in English (even for Swedish students). Use proper grammar and punctuation.
2. All answers need to be motivated, unless otherwise stated. Correct answers without motivation or with wrong motivation will not be given full credit.
3. Answer concisely, but explain all reasoning. Draw figures and diagrams when appropriate.
4. Write clearly. Unreadable or hard-to-read handwriting will not be given any credit.
5. Do not use red ink.
6. Solve only one problem per page.
7. Sort and number pages by ascending problem order.
8. Anything written on the back of the pages will be ignored.
9. Do not hand in empty pages or multiple solutions to the same problem. Clearly cross out anything written that is not part of the solution.

Question 1 IPv6

1.1 (6 points) An IPv6 node has an Ethernet-interface with MAC address 5C-26-0A-66-77-7C.
Question: Is EACH of the following a correct representation of an IPv6 address for the interface?

(a) 2001:06B0:0000:0000:5E26:0AFF:FE66:777C

(b) FF02:0000:0000:0000:0000:0001:FF66:777C

(c) FE80:0000:0000:0000:5E26:0AFF:FE66:777C

If NOT, explain clearly "why" and "what is wrong". If YES, rewrite the IPv6 address using optimal zero-compression, and then give the type and scope of the address explaining its use in IPv6 packets that will be sent to the node.

1.2 (4 points) There is neither broadcasting nor ARP in the new version of the Internet Protocol IPv6.

Task: Describe the substituting operation, its purpose, the protocols and the messages used by IPv6 nodes that are attached to an Ethernet LAN.

Hint: You are allowed to make use of the address(es) in the previous question if you would like to include example(s) in your answer.

Question 2 Congestion control

2.1 (3 points) Assume TCP is sending segments using a maximum window size (64 Kbytes) on a channel that has unbounded bandwidth and an average roundtrip time of 20 milliseconds. What is the maximum throughput? How does throughput change if the roundtrip time increases to 40 milliseconds (while bandwidth remains unbounded)? For simplicity consider that the congestion-window is very large and does not impose any limitations.

Solution: cf. slide 19; no need to assume/consider the size of segments, since the bandwidth is unbounded.

2.2 (4 points) Describe what is a network overlay and give (with justification) two example uses of network overlays, one for user applications and one for the network core.

Solution: see the slides.

Question 3 Socket API (5 points)

Please find below a diagram that depicts the interaction among clients and servers. For each item in the diagram, please provide a description. When referring to items (1), (2), … (13), please write only the name of the system calls (in the correct order). When referring to items (i), (ii), and (iii), please describe the interaction at the TCP level. When referring to items (a), (b), … (e), please give a brief description from the program perspective (no more than ten words).

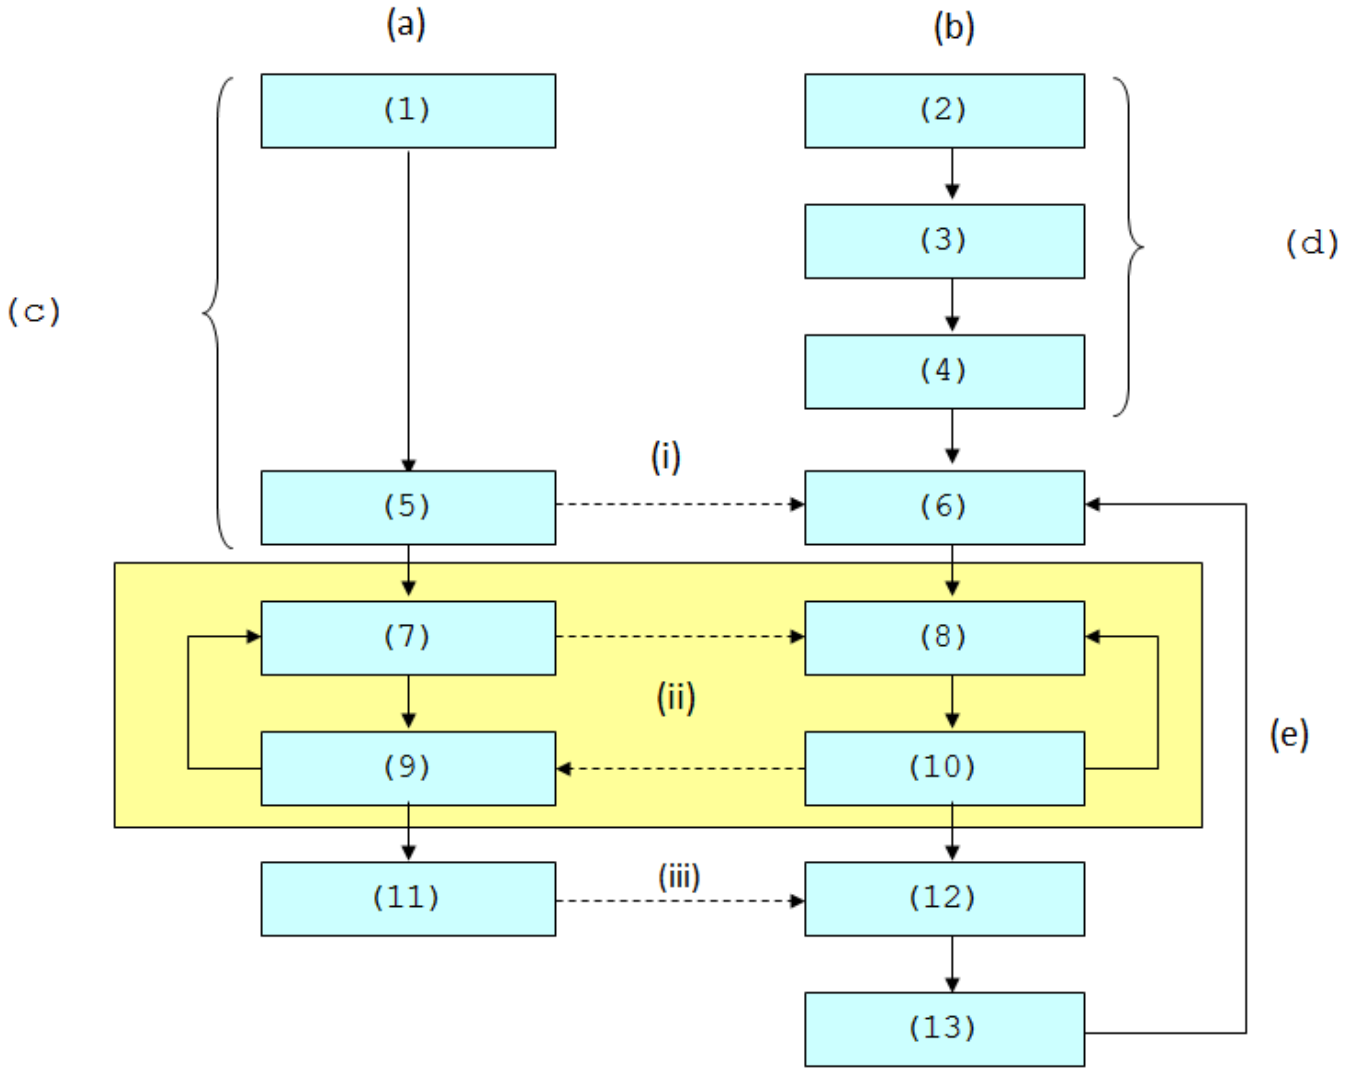

Solution: slide 16 in (36.3.a) EMS Networking API 1

Question 4 Network Security

Please take a look at the figure in which Alice is preparing a message to Bob.

4.1 (4 points) For each rectangle from top left to bottom right, (1), (2), (3) and (4), please explain the steps that Alice takes.

4.2 (4 points) Please draw the Bob's process when processing Alice's message. Please mark your rectangles from top left to bottom right using (5), (6), (7), ….

4.3 (4 points) Please explain the steps that Bob must take when processing Alice's message (for each rectangle from top left to bottom right).

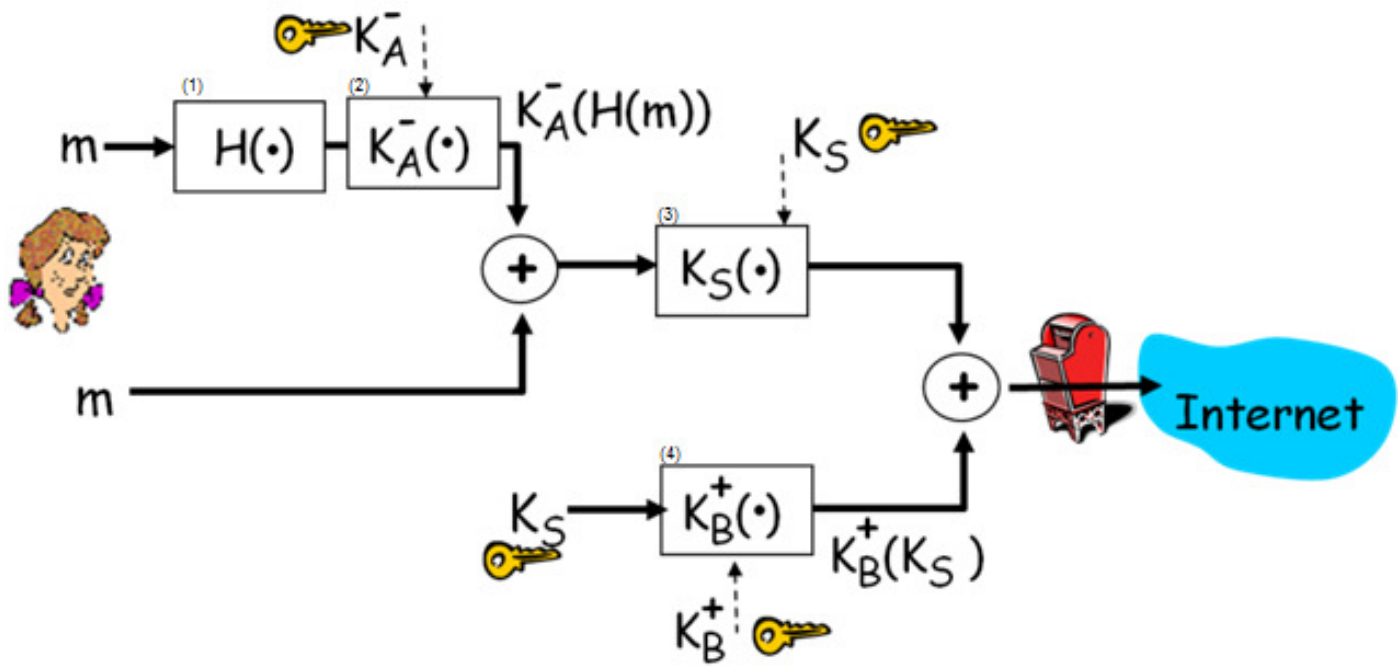

Solution: slide 62 to 69 in (42) EMS Network Security (short)

Question 5 (5 points)

Below please find a leader election algorithm. The question is whether this algorithm is self-stabilizing. In case you believe that the answer is positive, please prove your answer. Otherwise, please give a starting configuration, *c*, such that every fair execution that starts from *c* does not satisfies the task of self-stabilizing leader election.

1. write *id* to r_i
2. for $m := 1$ to n do $lr_m := \text{read}(r_m)$
3. Leader := (id == *maximum* $\{lr_m.\text{id} \mid 1 \leq m \leq \text{n}\}$)
4. (* if *Leader* == True then *act _like_a_leader()* *)

Solution: This solution is based on slide 5 in (41.2) EMS Self-* methods 4 and post-lecture questions (4).

Consider a configuration, c, in which the program counter of all processors is just after the execution of line 1. Suppose that for all p_i in P, we have that r_i=*minimum*($\{id_i \mid 1 \leq m \leq n\}$) in c. The rest of the answer is completed by considering all executions of the above algorithm that starts from c. The answer should show that after the algorithm execution, for all p_i in P, we have that $Leader_i == False$.

Question 6

Below please find self-stabilizing Dijkstra's algorithm for token circulation.

6.1 (3 points) Please define its safe configuration, c_{safe}.

Solution: At least two possible answers:

A: For all p_i in P, $x_i=x_1$.

B: For all p_i in P, $x_i=x_1$ OR ($x_n := (x_1+1) \bmod (n+1)$ and there is a j in [2, n-1], such that: for all k in $[2, j]$ $x_k=x_1$ and for all l in $[j, n-1]$ $x_l=x_n$).

6.2 (2 points) Say whether your definition of c_{safe} describes the entire set of possible configurations that a legal execution can consider.

Solution A does not where as B does.

Reminder: A legal execution is an execution in which the system behaves according to the system requirements. The definition of a safe configuration considers the set of legal executions. We say that configuration *c* is safe with respect to the set legal execution set, *LE*, if every system execution that starts from *c* is in *LE*. The definition of safe configuration, with respect to a given algorithm, must consider all variables in the system, since we are required to guarantee the satisfaction of the task requirement in the execution that follows.

You are welcome to explain and clarify your answer, but there is **no need to provide a formal proof**.

01 P_1: **do** forever **if** $x_1 = x_n$ **then** $x_1 := (x_1+1) \bmod (n+1)$

02 P_i $(i \neq 1)$: **do** forever **if** $x_i \neq x_{i-1}$ **then** $x_i := x_{i-1}$

Question 7 (4 points)

Above please find self-stabilizing Dijkstra's algorithm for token circulation. Please prove that for every configuration there exists at least one integer j such that for every i: xi ≠j (Lemma 2.3). Hint: the proof uses a very famous principle from discrete mathematics.

Solution: This solution is based on slide 19 in (39.1) EMS Self-* Methods 2 and post-lecture questions (3).

The system has exactly *n* variables. Each variable can encode *n+1* different value. Thus, by the pigeonhole principle, at any time, there is at least one value, j, that is not encoded by the system.

Question 8

Take a look at the self-stabilizing maximal matching algorithm that appears below. We assume the existence of a central daemon. Given a configuration *c*, we say that a processor p_i is:

- **matched** in *c*, if p_i has a neighbor p_j, such that $pointer_i = j$ and $pointer_j = i$.
- **single** in *c*, if $pointer_i = null$ and every neighbor of p_i is matched.
- **waiting** in *c*, if p_i, has a neighbor p_j such that $pointer_i = j$ and $pointer_j = null$.
- **free** in *c*, if $pointer_i = null$ and there exists a neighbor p_j, such that p_j is not matched.
- **chaining** in *c*, if there exists a neighbor p_j for which $pointer_i = j$ and $pointer_j = k$, $k \neq i$.

We define the variant function *VF(c)* as one that returns a vector (*m+s*, *w*, *f*, *c*), where *m* stands for matched, *s* stands for single, *w* stands for waiting, *f* stands for free, and *c* stands for chaining.

Program for P_i:

01 **do** forever

02 **if** $pointer_i = null$ **and** $(\exists P_j \in N(i) \mid pointer_j = i)$ **then** $pointer_i = j$

03 **if** $pointer_i = null$ **and** $(\forall P_j \in N(i) \mid pointer_j \neq i)$ **and** $(\exists P_j \in N(i) \mid pointer_j = null)$ **then** $pointer_i = j$

05 **if** $pointer_i = j$ **and** $pointer_j = k$ **and** $k \neq i$ **then** $pointer_i = null$

06 **od**

Solutions: see slides 8 and 9 in (41.3) EMS Self-* methods 5

8.1 (4 points) Please use the value of *VF(c)* to define the safe configuration, c. **Hint**: it is a vector that includes values that are either *0* or *n*, where n is the number of nodes in the system. For that value show that: (1) $pointer_i = j$ implies that $pointer_j = i$, and (2) if $pointer_i = null$ then there is no neighbor p_j, such that if $pointer_j = null$.

VF(c) = (n,0,0,0)

8.2 (6 points) Suppose that *c* is safe. Let a_i be a step of processor p_i that is taken immediately after *c*. Moreover, let *c'* be configuration that immediately follows by a_i. Show that c' is safe, i.e., the closure property. **Hint**: consider the case that a_i includes the execution of either line 02, 03 or 04. For each of these three cases, show that *VF(c) = VF(c').*

As simple as the hint.

Question 9 Multiple Choice questions (9 points)

Instructions: Select the single correct choice (labeled W, X, Y or Z) among the available options. Each correct answer will give you **one** point. **Zero** points are given for a blank answer, or an **incorrect** answer.

Please write down the **question letter and the full text of the answer** so we can avoid confusion!

9.1 Over a stream oriented connection, such as a TCP connection, a peer sends a 120-bytes (octets) long data stream. The receiving side attempts to `recv()` 64 bytes:

```
ssize_t ret = recv( sock, bufferPointer, 64, 0 );
```

Which of the following is a possible outcome of this operation?
[W] The recipient receives 64 bytes (ret = 64), but the entire 120 byte long data stream is kept and can be received with `recv()` as a whole at a later time.

[X] The recipient receives 64 bytes (ret = 64), and 56 bytes are discarded.
[Y] The recipient receives 60 bytes (ret = 60), and 60 bytes can be received with another call to `recv()` at a later time. (Note: the third parameter for recv() specifies how much data we can receive at most. It's perfectly possible for recv() to receive less than that!)
[Z] The recipient receives 68 bytes (ret = 68), and 52 bytes can be received with another call `recv()` at a later time.

9.2 In relation to blocking and non-blocking modes set on sockets used in TCP/IP communications, which of the following statements is *not* true?
[W] In blocking mode, `send()` will wait (i.e., block) until the remote peer has successfully received all data that is to be sent to it.
[X] In non-blocking mode, `recv()` will return an error if no data is available to be received.
[Y] In non-blocking mode, `connect()` may return an error condition (i.e., it may return -1), but the connection can still succeed at a later time.
[Z] In blocking mode, `recv()` may receive and return less data than what was requested of it.

9.3 The `select()` method can be used to detect events affecting one or more sockets. Sockets can be added to three sets, depending on what events are to be detected. One of the sets is the `readfds`.

```
int select(int nfds, fd_set *readfds, fd_set *writefds,
           fd_set *exceptfds, struct timeval *timeout);
```

Which of the following is an event that is never detectable using `select()` and the `readfds`?
[W] Data is ready to be received from the socket. [X] The connection associated with the socket has closed.
[Y] A non-blocking attempt to establish a connection has completed (but the peer has not yet sent any data and/or closed the connection). (This is detected with the writefds, the 2nd parameter to select(). Which makes sense: when the connection has been established, it's possible for us to send() (=write) data to it. Reading would require that the remote peer has sent something, which is not necessarily the case when the connection is first established.)
[Z] A new connection is ready to be accepted from the socket.

9.4 What is true about query of type A? Choose the correct answer.
[W] Answer to an A-type query may provide more than one IP address mapped for one host name but with different preference values.
[X] Answer to an A-type query will provide only one-to-one mapping of host name and IP address.
[Y] Answer to an A-type query may provide more than one IP address mapped for one host name.
[Z] Answer to an A-type query may provide more than one host name mapped for one IP address.

9.5 When an authoritative DNS server replies with an answer, it issues a TTL value for each RR. Choose the correct statement.
[W] The authority's administrator specifies only one TTL value for all RRs in the database of the authoritative DNS server.
[X] TTL value will be used to limit the number of hops across the DNS servers before answer reaches the client.

[Y] TTL value will be used by the cache to specify how long time the entry of RR may be reused before it is removed.
[Z] Cache-only and recursion DNS servers are allowed to choose their own TTL values for the entries they cache from answers.

9.6 Suppose that an IPv4 host A with address 129.16.84.83/22, will send packets with destination B 129.16.83.84/22. What will be done by host A before? Choose one of the following.
[W] Will search its ARP table for an entry for host B in order to map the IP address to MAC-address.
[X] Sends an ARP request with the IP address of host B as target, so that host A will get a mapping to the MAC-address of host B from the ARP-table of its default gateway.
[Y] Sends an ARP request with the IP address of its default gateway as target, so that host A will get a mapping to the MAC-address of the default gateway.
[Z] Sends an ARP request with the IP address of host B as target which will be forwarded by its default gateway, so that host A will get back an ARP-reply from host B through the gateway.

9.7 If an ARP request arrives, the software in the receiving host examines the address of the target. Choose the correct statement.
[W] If the receiving host is not the target, it will refresh the sender's address pair in its ARP-table if the mapping is already present.
[X] If the receiving host is not the target, it will add the sender's address pair to its ARP-table if the mapping is not already present.
[Y] If the receiving host is the target, it will add the sender's address pair to its ARP-table if the mapping is not already present.
[Z] All hosts are allowed to add as well as refresh their entries in ARP-table when any ARP reply arrives regardless of the destination.

A PC-user issues the command `<ipconfig -all>` in order to show the IPv4 configuration of the computer. Examine the output carefully and then answer the questions 9.8 and 9.9 that are given below.

```
.....................................................................
Windows IP Configuration
   Host Name . . . . . . . . . . . . : CSE-GUEST
   Primary Dns Suffix  . . . . . . . : net.chalmers.se
   Node Type . . . . . . . . . . . . : Hybrid
   IP Routing Enabled. . . . . . . . : No
   WINS Proxy Enabled. . . . . . . . : No
   DNS Suffix Search List. . . . . . : net.chalmers.se

Ethernet adapter Wired LAN:   Connection-specific DNS Suffix  . :
   Description . . . . . . . . . . . : Intel(R) 82579LM Gigabit Network
Connection
   Physical Address. . . . . . . . . : 5C-26-0A-66-77-7C
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   IPv4 Address. . . . . . . . . . . : 129.16.22.150
   Subnet Mask . . . . . . . . . . . : 255.255.0.0
   Lease Obtained. . . . . . . . . . : the 02 October 2012 13:32:18
   Lease Expires . . . . . . . . . . : the 02 October 2012 21:32:18
   Default Gateway . . . . . . . . . : 129.16.1.4
   DHCP Server . . . . . . . . . . . : 129.16.4.84
   DNS Servers . . . . . . . . . . . : 129.16.1.53
                                       129.16.2.53
   NetBIOS over Tcpip. . . . . . . . : Disabled
```

9.8 Please refer to the above output and then choose the correct statement.
[W] Since auto-configuration is enabled, the host is allowed to reuse this release even if the user reboots the computer.
[X] If the DHCP server uses default rebind and renewal times for this release, the host will be allowed to reuse this release next day (the 03 October).
[Y] If the DHCP server uses default rebind and renewal times for this release, the host should stop using it and try to renew the release on the same day at 17:32:18.
[Z] If the DHCP server uses default rebind and renewal times for this release, the host should stop using it if the host does not succeed to rebind the release before 21:32:18 on the same day.

8.9 Please refer the above output and choose the correct statement.
[W] The host will broadcast the DHCP-message REQUEST in order to renew the release.
[X] The host will broadcast the DHCP-message REQUEST in order to rebind the release.
[Y] The host will receive the DHCP-message RELEASE from the DHCP-server in order to renew the release.
[Z] The host will send the DHCP-message REQUEST to the IP address 129.16.4.84 in order to rebind the release.
