Chalmers University of Technology
Computer Science and Engineering
Elad Michael Schiller

2013-01-18

Written exam in EDA387/DIT663 Computer Networks 2013-08-29. 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) IPv6 Addresses
The following IPv6 addresses are given below using the optimal zero-compression.

(i) ff02::1:ff66:777c
(ii) 2001:6b0:2:2006::a01:e4
(iii) fe80::20c:f1ff:fe6c:4cd

1.1.(a) Decompress and rewrite each of the given addresses showing all hexadecimal digits.

1.1.(b) What is the type of each address?

1.1.(c) What is the scope of each address?

1.1.(d) Explain how routers and hosts will process IPv6 packets if each one of the given addresses would be the destination address of a packet.

1.2 (4 points) IPv6 Autoconfiguration

Task: Describe clearly the stateless autoconfiguration of the two most important IPv6-node's addresses. Suppose that the node is attached to an Ethernet LAN with connection to the Internet.

Note: In your answer mention in detail the addresses, the protocols, the messages, and their relevant content, which are utilized in order to achieve this type of IPv6 automatic configuration.

1.3 (10 points) Traceroute

A student at Chalmers has run the program "traceroute6" using the IPv6-enabled computer "`ju-025-16.studat.chalmers.se`" in one of the Chalmers computer rooms.

The output of running the trace is given below. Examine the output carefully and then answer the questions.

```
[student@ju-025-16 ~]$ traceroute6 basun.umdc.umu.se

traceroute to basun.umdc.umu.se (2001:6b0:e:1::f:2) from
2001:6b0:2:2f0a:20c:f1ff:fe6c:14dd, 30 hops max, 16 byte packets

 1  gw-2f0a.ch1.chalmers.se (2001:6b0:2:2f0a::1)          1.877 ms     1.007 ms   4.587 ms
 2  bb2-ch1-gw.chalmers.se (2001:6b0:2:1003::1)           8.592 ms     5.211 ms  11.683 ms
 3  chalmers2-bb2-gw.chalmers.se (2001:6b0:2:1007::2)    11.766 ms     7.413 ms   7.761 ms
 4  goteborg1-SRP2.sunet.se (2001:6b0:feed:dada::4:1)     7.835 ms     2.393 ms   5.000 ms
 5  goteborg1-POS1.sunet.se (2001:6b0:dead:beef:2::2d)    1.700 ms    22.270 ms   7.824 ms
 6  goteborg2-POS0.sunet.se (2001:6b0:dead:beef:2::a)     7.766 ms     7.509 ms   7.908 ms
 7  boras2-POS0.sunet.se (2001:6b0:dead:beef:2::d)        7.748 ms     7.422 ms  11.946 ms
 8  jonkoping1-POS3.sunet.se (2001:6b0:dead:beef:2::41)   8.247 ms     8.261 ms   8.306 ms
 9  c1sth-ge5-0-9.sunet.se (2001:6b0:dead:beef:2::fd)    11.342 ms     8.545 ms  18.963 ms
10  a1sth-ae4.sunet.se (2001:6b0:dead:beef:2::d6)         8.562 ms    32.667 ms  12.257 ms
11  (2001:6b0:dead:beef:2::2e2)                          17.827 ms    17.872 ms  17.843 ms
12  basun.ipv6.umu.se (2001:6b0:e:1::f:2)                17.761 ms    18.546 ms  17.729 ms
```

1.3.(a) Explain the purpose(s) of running "`traceroute`" when it is used by a network engineer.

1.3.(b) Explain the operation performed by "`traceroute`" when running on a host. In your answer please include what TCP/IP protocols and messages are employed by "`traceroute`".

1.3.(c) Explain the information (not value but the meaning) displayed within each hop and clearly describe how "`traceroute`" could find each part of the information (`name, address, time1, time2, time3`).

1.3.(d) What is the IPv6 prefix and default router used by the `studat` computer?

1.3.(e) What are the hostname, canonical name and IPv6 address of the target computer?

1.3.(f) How many IPv6 routers outside the Chalmers network are there in the path to the target computer during this trace?

Question 2 Congestion control

2.1 (3 points) Explain the main limitation that TCP poses for applications that need bandwidth/timing guarantees (e.g. real-time streaming applications).

2.2 (4 points) Explain the notion "TCP-friendly congestion control" and how it can be implemented.

Question 3 (6 points) **Socket API:** `select()`

Each of the following parts of a program contains a flaw. Identify and describe the flaw in a few short sentences or points. You do not have to correct the flaw; you should just find and describe it! (Note: you're not looking for, e.g., syntax errors. Find conceptual flaws in the program.)

Hint: The program uses `select()` and they are supposed to be non-blocking. Consider which operations can actually block the processes that execute these programs.

The following program accepts new connections using the `listenfd` socket. The first byte sent by a client is expected to be an 8 bit ID.

- You may assume that the `handle_*_error()` methods do something sensible.
- The helper method `register_client(client, id)` verifies the client ID is acceptable and if that is the case, enters the client into a global list. Otherwise it closes the connection.
- The method `add_client_sockets_to_readfds()` properly adds all active clients in the global list to the readfds. It returns the largest socket number it encounters.
- `handle_registered_clients()` handles clients that are ready to send data according to `readfds`, and removes clients that close their associated connections from the global list. No data is ever sent to the clients, the program only receives and processes data sent to it.

```
/* includes, declarations, etc. */
int main() {
      int listenfd = -1;
/* initialization code, such as setting up a listening socket on listenfd,
   has been omitted - this is not the error you're looking for */
      while( 1 ) {
            fd_set readfds; // initialize read set
            FD_ZERO( &readfds );
            int maxfd = add_client_sockets_to_readfds( &readfds );

            FD_SET( listenfd, &readfds );
            if( listenfd > maxfd ) maxfd = listenfd;

            int ret = select( maxfd+1, &readfds, 0, 0, 0 ); // call select

            if( -1 == ret ) handle_select_error();

            // is there a new client waiting?
            if( FD_ISSET( listenfd, &readfds ) ) {
                  sockaddr_in clientAddr;
                  socklen_t clientAddrLen = sizeof(clientAddr);

                  int client = accept( listenfd,
                        (sockaddr*)&clientAddr,
                        &clientAddrLen
                  );

                  if( -1 == client ) handle_accept_error();

                  // receive 8bit client ID
                  unsigned char id;
                  int ret = recv( client, &id, sizeof(id), 0 );

                  if( 0 == ret ) {
                        close( client );
                        continue;
                  }

                  if( -1 == ret ) handle_recv_error();

                  // register client
                  register_client( client, id );
            }

            handle_registered_clients(&readfds);//handle registered clients
      }
      return 0;
}
```

Question 4 Network Security

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

4.1 (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), ...

4.2 (4 points) Please explain the steps that Bob must take when processing Alice's message (for each rectangle from top left to bottom right).

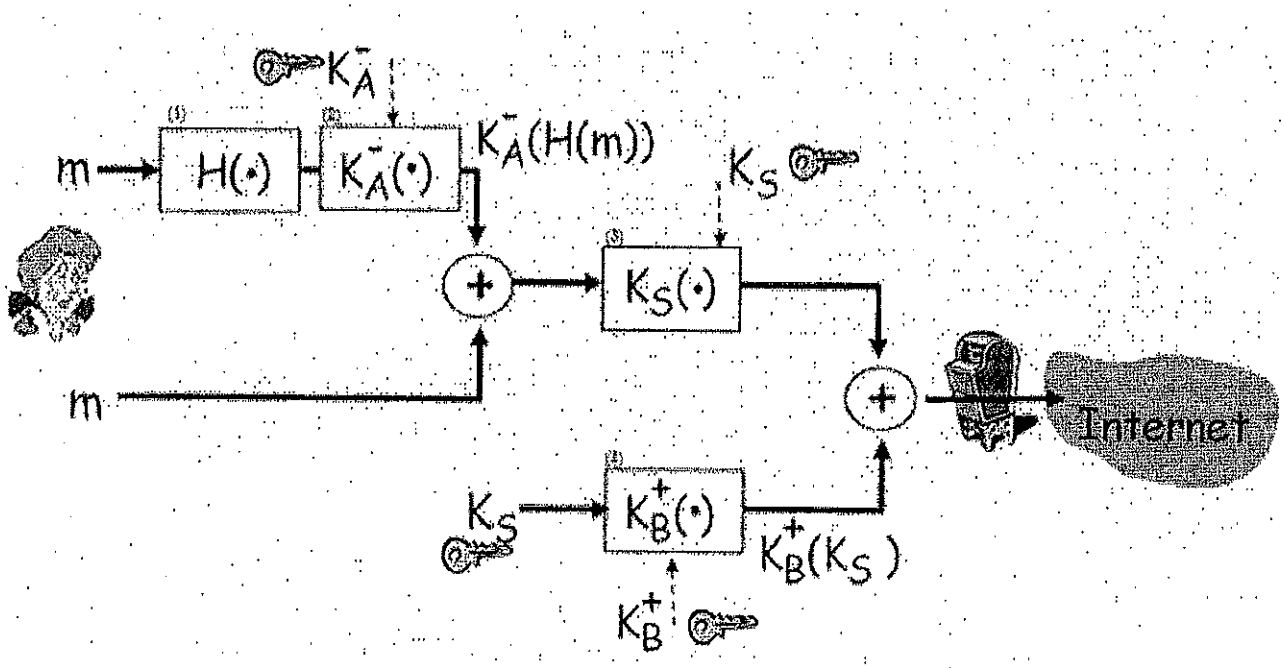

Question 5 (4 points)

Below please find the spanning-tree construction algorithm. A variant of that algorithm is proposed. In this version every processor repeatedly checks whether the value of the *dist* variable of its parent in the tree is smaller than the value of its own *dist* variable. Processor p_i does not execute lines 8 to 16 of the code when the above condition holds. Is the proposed algorithm a self-stabilizing spanning tree construction algorithm? Prove your answer.

```
01 Root: do forever
02         for m := 1 to δ do write r_im := ⟨0,0⟩
03      od
04 Other: do forever
05         for m := 1 to δ do write lr_mi := read(r_mi)
06         FirstFound := false
07         dist := 1 + min{lr_mi.dis | 1 ≤ m ≤ δ }
08         for m := 1 to δ
09         do
10               if not FirstFound and lr_mi.dis = dist -1
11                     write r_im := ⟨1,dist⟩
12                     FirstFound := true
13               else
14                     write r_im := ⟨0,dist⟩
15         od
16      od
```

Question 6 (1 point)

Write the definition of: the term *set of legal executions*.

Question 7 (1 point)

Write the definition of the term *safe configuration*.

Question 8 (4 points)

Below please find a self-stabilizing leader election algorithm for general communication networks. Define a safe configuration for that self-stabilizing algorithm.

```
01 do forever
02     ⟨candidate, distance⟩ = ⟨ ID(i), 0 ⟩
03     forall P_j ∈ N(i) do
04         begin
05             ⟨leader[j],dis[j]⟩ := read⟨ leader_j dis_i ⟩
06             if (dis[j] < N) and ((leader[j] < candidate) or
07                 ((leader[j] = candidate) and (dis[j] < distance))) then
08                     ⟨candidate, distance⟩ := ⟨ leader[j],dis[j] + 1⟩
09         end
10     write ⟨leader_i,dis_i⟩ := ⟨candidate, distance⟩
11 od
```

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 the number of points that is written next to the question title. **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!

Assume that a host has no ARP entries before a certain application generates IP datagrams to destination address 129.16.211.118. The host has Internet connection using an Ethernet NIC attached to an Ethernet switch. The following is part of the host's IP configuration.

```
C:\>ipconfig /all
Windows IP Configuration

Ethernet adapter Local Area Connection:

      Physical Address . . . . . . . . . : 00-1B-77-D3-20-B9
      DHCP Enabled . . . . . . . . . . . : Yes
      Autoconfiguration Enabled  . . . . : Yes
      IP Address . . . . . . . . . . . . : 129.16.212.119
      Subnet Mask  . . . . . . . . . . . : 255.255.252.0
      Lease Obtained . . . . . . . . . . : the 02 October 2012 13:32:18
      Lease Expires. . . . . . . . . . . : the 02 October 2012 21:32:18
      Default Gateway . . . . . . . . . : 129.16.213.23
      DHCP Server  . . . . . . . . . . . : 129.16.213.24
      DNS Servers . . . . . . . . . . . : 129.16.1.53
                                          : 129.16.2.53
```

9.1 (1 point) Please examine the above output carefully and then choose the correct statement.

[W] The host and destination have IP addresses belonging to the same subnet.

[X] The host will broadcast an ARP request inquiring the MAC address of the destination before sending the packets.

[Y] The IP packets with the destination IP address will be encapsulated in frames with the destination MAC address of the default gateway before delivering the packets.

[Z] The host cannot send any packet because the default gateway is on a different subnet.

9.2 (1 point) Please refer to the above output and assume that the DHCP server uses default renewal and rebind times for its lease of addresses. Choose the correct statement.

[W] The DHCP server is on a different subnet and therefore default gateway will act as relay agent (address-helper) in order to exchange DHCP messages between the host and the server.

[X] If the host would not be succeeded to renew the lease in time it should, one hour before 21:32:18, broadcast the DHCP-message REQUEST in order to rebind the lease.

[Y] The host will receive the DHCP message "LEASE" from the DHCP server after 17:32:18 in order to renew the lease.

[Z] The host may continue using the leased address after 21:32:18 for four hours (50% of the lease time) before requesting a new lease.

9.3 (4 points) 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.

9.4 (1 point) 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.

9.5 (1 point) DNS messages are mostly short with predefined format and maximum size. Choose the correct statement.

[W] DNS always uses UDP for all types of transactions.

[X] DNS encapsulates the messages directly in IP packets.

[Y] DNS primarily uses UDP segments to carry requests and replies.

[Z] DNS always uses TCP in order to deliver DNS information reliably.

9.6 (1 point) Which of the following is NOT a type of the DNS Resource Records?

[W] PTR

[X] TTL

[Y] AAAA

[Z] NS
