Test of Pearl 101 — Operating Systems and Computer Networks

Pearls of Computer Science (201700139) / Introduction to BIT (201700149)
Bachelor module 1.1, EWI

October 15, 2017, 13:45–14:45

Module coordinator: Doina Bucur, Maurice van Keulen
Instructor: Pieter-Tjerk de Boer

(updated 2018-11-05, small correction in 2(d))

- You may use 1 A4 document with your own notes for this exam and a *simple* calculator.
- Scientific or graphical calculators, laptops, mobile phones, books etc. are not allowed.
  Put those in your bag now!
- Write your answers on this paper, in the provided boxes , and hand this in.
- Total number of points: 100.
  Total number of pages: 5.

Your name:

(please underline your family name (i.e., the name on your student card), so that we know how to sort)

Your student number:

Continued on next page...

1. Operating systems

5 pt (a) In the context of operating systems, a "process" is:
(choose one, no explanation needed)

A. either memory management, file management, or I/O.
B. a sequence of instructions to be executed by the processor.
C. the consequence of insufficient protection agains malicious software.
D. a sequence of events coming in from the outside world, like key presses.
E. a sequence of events for the outside world, such as data sent to a printer.
F. the sequence of actions the operating system takes when the computer is switched on.

B

5 pt (b) An operating system runs two processes simultaneously by:
(choose one, no explanation needed)

A. Installing an extra CPU core for each process.
B. Executing alternately one instruction from each process.
C. Executing one process for a while, then the other for a while.
D. At any time only executing the process the user is interacting with.
E. Putting instructions from both processes in alternate memory locations.
F. Putting instructions from the processes in the lower and upper half of memory, respectively.

C

5 pt (c) A typical situation in which an operating system "swaps", is the following:
(choose one, no explanation needed)

A. When extra memory is installed into the computer.
B. When data is being copied from a USB stick to a harddisk.
C. When two processes run which each need a lot of memory.
D. When one process gets a timeout and another is dispatched.
E. When files are being uploaded and downloaded at the same time.

C

We also accept answer D, since that is the very moment at which the swapping happens if it is needed. We aimed for answer C though, as the question asked for a situation, not for the precise moment.

8 pt (d) Which kinds of information are typically stored as metadata in the file system when a file is created?
Select **one or more** from the following list:

A. Current time
B. Size of the file
C. Size of the disk
D. Name of the file
E. Password of the user
F. Name of the computer
G. Access rights of the file
H. IP address of the computer

A,B,D,G

7 pt (e) Suppose a process is started, waits *twice* for a keypress, then does some computations and finally terminates. During the computations, it is interrupted by an operating system timeout *two times* for other processes to run.
How many times does this process pass through the "waiting" state?
(one number, no explanation needed)

5

Continued on next page...

2. Networks – protocols

5 pt (a) If a packet from protocol X is *encapsulated* in protocol Y, then:

A. protocol X is at a lower layer than protocol Y.
B. protocol X is at a higher layer than protocol Y.
C. protocol X and protocol Y are at the same layer.
D. protocol X and protocol Y do not fit in the layering model.

B

5 pt (b) Although an IPv4 address has 32 bits, allowing for 2^{32} = about 4 billion combinations, there is already a shortage with only about 1 billion computers connected to the internet. Why is this?

A. The real problem is not a shortage of addresses, but of host names.
B. Large companies have bought too many addresses and refuse to sell them.
C. Addresses are assigned systematically in entire blocks, causing some to be wasted.
D. Many old computers have already been thrown away, and their addresses can't be reused.

C

5 pt (c) When a web browser (client) talks to a web server, the TCP port number on the server side normally is 80. What can we say about the client side port number?

A. It must also be 80, since that's the port for web traffic.
B. It must not be 80, since the client is not a web server.
C. It is some random number, and different for every next connection.
D. It is some random number, but stays the same when multiple connections are made.

C

Here you see a few network packets as displayed by Wireshark:

nr.	source IP	destination IP		source/dest.port		TCP seq./ack.numbers
1	130.89.144.74	130.89.13.213	TCP	7701 > 56922	[ACK]	Seq=1000 Ack=2001 Len=90
2	130.89.4.7	130.89.13.213	TCP	7701 > 56922	[ACK]	Seq=1090 Ack=2001 Len=40
3	130.89.13.213	130.89.4.7	TCP	56922 > 7701	[ACK]	Seq=2001 Ack=1110 Len=0
4	130.89.144.74	130.89.13.213	TCP	7701 > 56922	[ACK]	Seq=1130 Ack=2001 Len=30
5	130.89.144.74	130.89.13.213	TCP	7701 > 569	[ACK]	Seq=1500 Ack=2001 Len=30
6	130.89.13.213	130.89.144.74	TCP	56922 > 7701	[ACK]	Seq=2001 Ack=1160 Len=0
7	130.89.13.213	130.89.144.74	TCP	569 > 7701	[ACK]	Seq=2001 Ack=1530 Len=0

6 pt (d) How many different TCP connections are there in this trace?
(one number, no explanation needed)

Packets 1, 4 and 6 belong together in one connection; packets 2 and 3 are from the second connection; and packets 5 and 7 from the third.

3

9 pt (e) One packet in this trace has been lost. What must have been the contents of this missing packet?

source IP = **130.89.144.74**

destination IP = **130.89.13.213**

Seq = **1090**

Ack = **2001**

Len = **40**

This packet is missing, as there's a sequence number gap between packets 1 and 4.
We gave partial scores for several other, also somewhat plausible answers.

Continued on next page...

3. Networks – delay

Consider a network consisting of an endhost A, two routers B and C, and an endhost D. The only path from A to D is via B and C. The link from A to B is 2 megabit/s, from B to C 1 megabit/s, and from C to D 2 megabit/s.

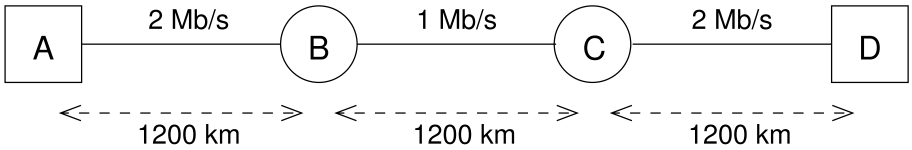

We assume that the computation time needed by routers B and C to decide where to send the packet, is negligible. We also assume that each cable is 1200 km long, and the signals travel over it at 200 000 km/s.

An application on host A generates 3 packets of 2000 bits each (incl. headers), at times $t_1 = 0$ ms, $t_2 = 10$ ms and $t_3 = 11$ ms. There is no other traffic in this network.

15 pt (a) Calculate the transmission and propagation delays for one packet on each of the links, or indicate why it is negligible:

Transmission delay on link A–B:

$$\frac{\textbf{2000 bits}}{\mathbf{2 \cdot 10^6}\ \textbf{bits/s}} = \mathbf{0.001\ s = 1\ ms}$$

Propagation delay on link A–B:

$$\frac{\mathbf{1200 \cdot 10^3\ m}}{\mathbf{2 \cdot 10^8\ m/s}} = \mathbf{0.006\ s = 6\ ms}$$

Transmission delay on link B–C:

similar calculation, 2 ms

Propagation delay on link B–C:

same, 6 ms

Transmission delay on link C–D:

1 ms

Propagation delay on link C–D:

same, 6 ms

25 pt
(b), (c), (d)
together

(b) At what time will the first packet have arrived completely at host D? Show your calculation.
If the packet incurs any queueing delays, clearly indicate at which nodes and how much.

Adding all transmission and propagation delays to t_1 gives t = 22 ms as the arrival time.
Since this is the first packet, there is nothing in front of it in the queues.

Max. 8 pt for (b).

(c) At what time will the second packet have arrived completely at host D? Show your calculation.
If the packet incurs any queueing delays, clearly indicate at which nodes and how much.

We will not have any queueing delay, due to the 10 ms interval, which is (much) longer than any of the transmission times of the first packet.
So it takes again 22 ms, starting from t_2 = 10 ms, so it arrives at t = 32 ms.

Max. 6 pt for (c).

(d) At what time will the third packet have arrived completely at host D? Show your calculation.
If the packet incurs any queueing delays, clearly indicate at which nodes and how much.

Check where each packet is when, with a text or table or so.

Transmission of packet 2 at A is from t = 10 to t = 11 ms. Propagation to B takes 6 ms, so its last bit arrives at B at t = 11 + 6 = 17 ms. Transmission at B is from t = 17 to t = 19 ms, arrival at C at t = 25 ms, transmission by C from t = 25 till t = 26 ms, and arrival of last bit at D at t = 32 ms.

Transmission of packet 3 at A is from t = 11 to t = 12 ms; its last bit arrives at B at t = 18 ms. But then B's transmitter is still busy with packet 2, so the third packet is queued. Its transmission can start at t = 19 ms (so we have 1 ms of queueing delay), and takes until t = 21 ms. Arrival of last bit at C is at t = 27 ms; by that time, C's transmitter is not busy with the second packet anymore, so no queue, transmission starts immediately, takes until t = 28 ms; after propagation the packet arrives at D at t = 34 ms.

Max. 11 pt for (d).

End of this exam.