

This is a written exam for the course “*Performance Analysis of Communication Networks*”

Lecturers: prof.dr. R.D. van der Mei and dr. T. Kielmann

Date and location of exam: Wednesday, February 17, 2010

Rules for the exam:

1. Allowed material: This is an open book exam. For answering the questions, you are allowed to use all kinds of written material like textbooks, printouts of the lecture slides, your own notes, etc. You are allowed to bring your laptop for looking up electronic versions of course reading material, but electronic communication during the exam is strongly prohibited.
2. Language disclaimer: You are kindly asked to answer the questions using the English language. However, if it helps clarifying your answers, you may use some Dutch here and there. Doing so, will not affect your result.
3. Calculation of end grade for the course: the end grade for the course is built up in two parts: homework assignments and a written exam.
 - *Homework assignments*: The average of the three grades counts to 50% of the final grade.
 - *Written exam*: for this written exam you get a grade between 1 and 10. This grade will count for the remaining 50% of the final grade.
 - *Final grade*: the final grade is calculated as the average of the grade for the written exam on the one hand, and the average homework grade on the other hand, with the restriction that the grade for the written exam must be at least 4.0.
4. Credits: This written exam consists of four questions (A, B, C and D), each of which consists of a number of sub-questions. The maximum number of credits you can get is distributed as follows amongst the sub-questions:

	1	2	3	4	5	total
A	4	4	4	4		16
B	3	6	4	3		16
C	3	3	3	4	3	16

Good luck!

A. Capacity planning for Video-on Demand for CableCom

Cable TV company CableCom plans to offer Video-on-Demand (VoD) services, allowing their subscribers to watch videos upon request *at any time*. This is fundamentally different from the current situation, where CableCom only offers standard cable TV services: for each TV channel pre-scheduled TV programs (see your TV guide) are simply broadcast to all customers at specific times. CableCom has installed two types of servers: a signalling server and a video server. The process of setting up a VoD session consists of two phases. First, the client send a request to the signalling server to set up a connection between the video server and the client (phase I). Once a connection between the video server and the client has been established, the video server immediately starts to send the video traffic stream to the client TV (phase II). See Figure 1.




Figure 1: CableCom's Video-on-Demand service.

CableCom expects that VoD will be a commercial success, and therefore, wants to plan the capacity of its infrastructure properly and timely anticipate on performance problems when the number of users increases in the near future. In this context, CableCom wants to make sure that the signalling server is well-dimensioned so that the connection set-up phase does not take too long. Moreover, they want to make sure that the amount of network bandwidth is sufficiently large.

We make the following simplifying assumptions:

Regarding Phase I: The amount of time involved in processing a connection set-up request by the signalling server is exponentially distributed with mean 1 second. The signalling server handles connection set-up requests in the order of arrival, and can handle one request at a time (single-threaded). The network latency and bandwidth consumption involved in establishing a connection are negligible.

Regarding Phase II: Once the connection is established the video server will generate a traffic stream over the CATV network trunk at a constant rate of 5 Megabits per second for the duration of a movie. The duration of a movie has a gamma distribution with mean 2 hours. The CATV network trunk is shared by all clients of CableCom and its capacity is 1 Gigabit per second. When there is insufficient bandwidth available for running a VoD session over a newly established connection, the session is blocked and the connection is terminated. The time involved in terminating a connection is negligible.

- A.1 Formulate a performance model that encompasses *both* the delay involved in setting up a connection and the blocking of VoD sessions. Define the relevant notation and the performance metrics. Motivate your assumptions and be precise!
- A.2 What is the expected time it takes to set up a connection between the video server and the client?
- A.3 CableCom wants to deliver good service to its VoD customers, and requires that the average duration of the connection establishment phase (including both waiting time and processing time) is less than 5 seconds. What is maximum number of connection set-up requests per minute that can be handled while meeting this constraint?

- A.4 How many VoD sessions can the CATV network trunk handle simultaneously? Give an expression for the session blocking probability.

B: Sliding-window protocols

- B.1 For a sliding-window protocol, the achievable bandwidth between a pair of sender and receiver depends on the performance characteristics of the network in between. On which characteristics precisely?
- B.2 For a pair of sender and receiver, what is the optimal size of both send window and receive window to maximize achievable bandwidth while minimizing the utilization of resources? What is the consequence if the send window is either bigger or smaller than the receive window?
- B.3 What happens if the timeout value (for waiting for acknowledgments) is either too high or too low, for a given connection?
- B.4 Which additional mechanism has been added to TCP to shorten the time wasted because of lost packets?

C: Analysis of a simple file server

Consider a single-threaded file server that handles incoming requests one-at-a-time and in the order of arrival. The required amount of processing time per request, denoted by B , has the following probability distribution: For $x > 0$,

$$\Pr\{B > x\} = \frac{1}{(1+x)^2}.$$

- C.1 Define a simple model for the file server, and introduce the proper notation. Be precise.
- C.2 Is the probability distribution of B light-tailed or heavy tailed? Motivate your answer.
- C.3 Show that the mean request processing time equals 1.

Hint: the mean of B can be expressed as $E[B] = \int_{t=0}^{\infty} \Pr\{B > t\} dt$.

- C.4 The response time of a request is defined as the waiting time plus the processing time. What is the mean response time of an arbitrary request? Give an intuitive explanation for this peculiar result.

Suppose now instead of handing requests one-at-a-time, the file server handles the requests simultaneously in a processor sharing fashion.

- C.5 What is in that case the mean time it takes for an arbitrary request to be completed? This result is fundamentally different from the result in C.4. Give a clear intuitive explanation for this.