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: Thursday, February 23, 2006, 18.30-21.30 in M129

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*: during the course 4 homework assignments have been distributed among the students and placed on the Web site. For each homework assignment each student has received a grade between 1 and 10. The deadline for the latest homework assignment is December 25, 2005. The average of the 4 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 three questions (A, B and C), 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	6	total
A	3	3	5	3	3	3	**20**
B	6	8	6				**20**
C	5	5	5	5			**20**

If N is the total number of credits you get for questions A, B and C (out of 60), then your grade for the written exam will be calculated according to the following formula:

grade = $1 + 9N/60$.

Good luck!

QUESTION A: GSM network dimensioning problem

A mobile operator of a GSM network wants to determine how many base stations are needed to satisfy its customers' Quality of Service (QoS) demands. To this end, the operator smartest wants to determine the maximum size of a cell for the call blocking probability is small enough. Voice telephone calls are generated with rate 3 calls per minute *per square kilometer* (i.e., km^2), and the call duration is a gamma distribution with mean 2 minutes. Assume that each voice call requires a single 14.4 kbps channel to the nearest base station, and that each cell can support only 4 channels in parallel.

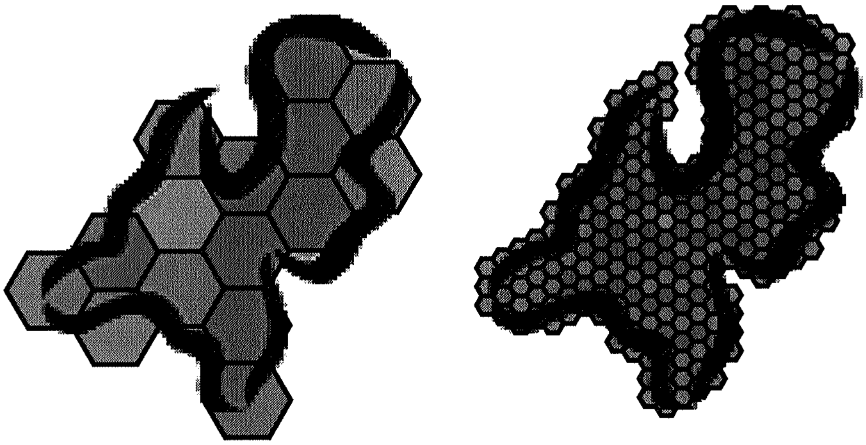

Figure 1. Illustration of GSM network dimensioning problem.

To make a proper decision on the number of base stations to be placed to offer good QoS to its customers, the operator wants to understand the impact of the cell size (in km^2) and the call blocking probability.

A.1 Formulate a simple model description for the problem. Be precise, motivate your assumptions and clearly define any notation!

A.2 Calculate the call blocking probability for cell size values 0.25 km^2, 1 km^2 and 5 km^2.

Now suppose the service provider wants to offer a new *additional* service to its customers: video conferencing. Each video conferencing call requires a 57.6 kbps connection (4 parallel channels). Video conferencing calls arrival according to a Poisson arrival process with rate 0.1 call per minute per km^2, and the mean conference call duration is 20 minutes. A conference call attempt is blocked when there are less than 4 lines available.

A.3 Determine the blocking probability for *both* the voice telephone calls and the video conferencing calls if the cell size is 1 km^2 if the new service is added.

A few basics...

A.4 Erlang's blocking formula is known to have some *insensitivity property*. What *exactly* does that mean? Be precise!

A.5 Poisson processes occur naturally in the modelling in human-initiated events, such as telephone-call initiation moments and Web-session initiation moments. Why is that? Be precise.

A.6 What is the relation between a Poisson *process* and a Poisson *distribution*? And what is the difference between them? Be precise!

QUESTION B: Measurement procedures

B.1 Explain the impact of the latency-bandwidth product of a network connection on the bandwidth that is achievable by a sliding-window protocol.

B.2 A common approach to improve the achievable bandwidth is to use multiple TCP connections between a sender-receiver pair in parallel. Compare the behaviour in terms of achievable bandwidth of

a) one single TCP connection with a send window size W_1
b) n parallel TCP connections each with a send window size W_n, with $W_1 = n * W_n$

For answering the question it is sufficient to consider the steady state of the connection(s), thus neglecting the start-up phase.

B.3 Explain how the following characteristics of a network path from A to B can be measured:

(a) roundtrip delay
(b) one-way delay
(c) one-way delay variation (jitter)

C. Response time analysis of cash dispensers

A popular bank, FastCash, wants to offer cash dispenser services to its customers. To this end, the system is equipped with a front-end (FE) server and two database (DB) servers: an authentication DB server and a balance DB server containing information about the customers' balance. The service works as follows. The customer (1) inserts a credit card, (2) types in a 4-digit ID number, (3) types in how much money (in euros) he/she wants to have, and then presses the OK button. This cash retrieval request is then immediately forwarded to the FE server. The FE server then pre-processes the request, and forwards the request to the authentication DB server (step 1), which checks whether the 4-digit ID number is correct. The response of this request is sent back to the FE server (step 2). Subsequently, the FE server processes the response and forwards the request to the balance DB server (step 3), which checks whether the customer's balance is sufficient, and then returns a response to the FE server (step 4), which processes the response, makes a decision on whether or not the requested amount of cash is paid out or not and terminates the transaction.

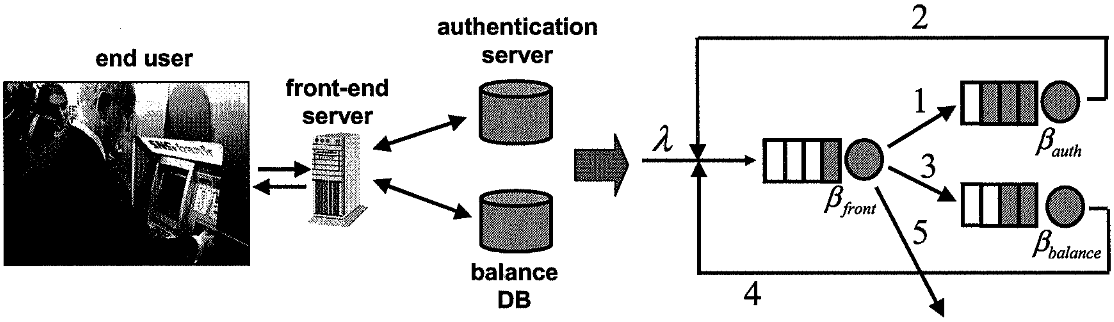

Figure 2. Processing phases of the cash dispenser service.

It is important to notice that during each customer transaction the FE server processes the request *three* times in total (namely: at the pre-processing stage, after the authentication phase and after the balance checkup phase), whereas both the authentication DB and the balance DB only process the request *once*.

Each of the three servers can handle a single request at a time, and processes the requests on a First-In-First-Out (FIFO) basis. The processing times of *each* request at the FE server, the authentication server and the balance DB are independent and exponentially distributed with means β_{FE}, β_{auth} and $\beta_{balance}$, respectively. The clients' session initiation moments occur according to a Poisson process with rate λ. Our focus is on the processing times of the servers, and therefore, the network delay is assumed to be negligible. The total sojourn time of a request is this system represents the total response time of a combined cash retrieval request. FastCash now wants to be able to predict the total response time experienced by the client.

C.1 Formulate the model as an open queueing network (see Figure 2). Define the proper variables and motivate your assumptions. Be clear!

C.2 Determine the average load at each of the three servers.

C.3 Use the theory of product-form networks to give an expression for the joint probability distribution of the number of request present at each of the queues. Be precise.

C.4 Give an expression for the expected total sojourn time in the system. Motivate how you get this answer.