TECHNISCHE UNIVERSITEIT DELFT
Faculteit Elektrotechniek, Wiskunde en Informatica

IN4390 – Quantitative Evaluation of Embedded Systems

January 25th, 2021, from 09:00 to 12:00

Koen Langendoen, Lydia Chen, George Iosifidis, Marco Zuniga

Question:	1	2	3	4	5	6	7	Total
Points:	10	7	8	20	15	15	15	90
Score:								

- This is an online (open book) exam
- Write your answers with a black or blue pen, not with a pencil
- **Always justify your answers, unless stated otherwise**
- Upload your answers as one file per question in PDF, PNG, or JPEG format

- The exam covers the following material:
  (a) the paper "Basic Concepts and Taxonomy of Dependable and Secure Computing" by A. Avizienis ; J.-C. Laprie ; B. Randell ; C. Landwehr
  (b) chapters 18-20,22-23 (DoE), and 30-33 (Queueing Theory) of the book "The Art of Computer Systems Performance Analysis" by R. Jain
  (c) the paper "Petri nets: Properties, analysis and applications" by T. Murata
  (d) chapters 11.2 (DTMC), and 11.3 (CTMC) of the book "Introduction to probability, statistics, and random processes" by H. Pishro-Nik
  (e) the paper "Exploring Exploration: A Tutorial Introduction to Embedded Systems Design Space Exploration" by A.D. Pimentel

Operational Laws	
Utilization law	$U = XS$
Little's law	$N = XR$
Forced-flow law	$X_k = V_k X$
Bottleneck law	$U_k = D_k X$
Operational Bounds	
Througput	$X \leq \min\left(\frac{1}{D_{max}}, \frac{N}{D+Z}\right)$
Response time	$R \geq \max\left(D, N \times D_{max} - Z\right)$
Queueing Theory M/M/1	
Utilization	$U = XS = \lambda/\mu = \rho$
Probability of n clients in the system	$P_n = \rho^n(1-\rho)$
Mean #clients in the system	$N = \rho/(1-\rho) = \lambda/(\mu - \lambda)$
Mean #clients in the queue	$N_Q = N - \rho$
Mean response time	$R = N/\lambda = 1/(\mu - \lambda)$
Mean waiting time	$W = R - S = \rho/(\mu - \lambda)$
Basic Math	
Geometric series	$\sum_{k=0}^{\infty} r^k = \frac{1}{1-r}$, for $\lvert r \rvert < 1$

ANOVA Table for One Factor Experiments

Component	Sum of Squares	%Variation	DF	Mean Square	F-Comp.	F-Table
y	SSY=$\sum y_{ij}^2$		ar			
$\bar{y}_{..}$	SS0=$ar\mu^2$		1			
y-$\bar{y}_{..}$	SST=SSY-SS0	100	ar-1			
A	SSA $= r\Sigma\ \alpha_i^2$	$100\left(\frac{\text{SSA}}{\text{SST}}\right)$	a-1	MSA $= \frac{\text{SSA}}{a-1}$	$\frac{\text{MSA}}{\text{MSE}}$	$F_{[1-\alpha; a-1, a(r-1)]}$
e	SSE=SST- SSA	$100\left(\frac{\text{SSE}}{\text{SST}}\right)$	$a(r-1)$	MSE $= \frac{\text{SSE}}{a(r-1)}$		

Question 1 [10 points]

(a) $\boxed{\text{5 points}}$ Avizienis et al. [reading list] distinguish 3 important **threats** to the dependability of embedded systems: faults, errors, and failures. Discuss the relationship between these concepts **and** illustrate them with an example from the embedded systems domain.

(b) $\boxed{\text{5 points}}$ Mention 3 **means** to address the above-mentioned threats **and** link them to your example.

Question 2 [7 points]

Here is a 2-way ANOVA analysis with 2 factors being A and B. You need to figure out the missing values to answer the questions below.

	DF	SS	MS	F
A	2	80	40	**
B	3	100	33.33	**
Interaction	**	40	6.67	0.8
Error	**	100	**	
Total	23	320		

(a) $\boxed{\text{1 point}}$ How many replicates of experiments were performed?

(b) $\boxed{\text{3 points}}$ We want to know which factors (A, B, and AB) are significant. What statistics should we compare against the computed F statistics for factor A,B, and AB, respectively? Assume the significant level is 0.05.

(c) $\boxed{\text{2 points}}$ The P-value is another way to check if factors are significant. We obtained the following P-values (0.0097, 0.56, 0.53) for factor A, B, and AB, respectively. OOPs! we actually made a mistake in one of the P values. Which factor is mistaken? Give your reason!

(d) $\boxed{\text{1 point}}$ If you want to fit a regression model, which factors shall be included? And why?

Question 3 [8 points]

You want to know the significant parameters that affect the image recognition software on your cell phone. You use DoE to figure that out. You consider the following parameters 4 parameters: batch size (A), learning rate (B), and memory size (C), and image sizes (D).

(a) $\boxed{\text{2 points}}$ You start with the 2k full factorial design. Write down the sign table for only factor A, B, C, and D.

(b) $\boxed{\text{1 point}}$ With 3 replication per configuration, how many experiments do you need to run?

(c) $\boxed{\text{2 points}}$ However, it's too time consuming to run the full factorial design. And, you decide to run a half factorial design by following the generator function of I=ABCD . Based on this, write down the sign table for only factor A, B, C, and D.?

(d) $\boxed{\text{2 points}}$ Tell us what kind of design this is and its resolution?

(e) $\boxed{\text{1 point}}$ Is this a good generator function, compared to I=ABD? And, why?

Question 4 [20 points]

(a) $\boxed{\text{5 points}}$ Is the following Petri Net 6-bounded? Please justify your answer.

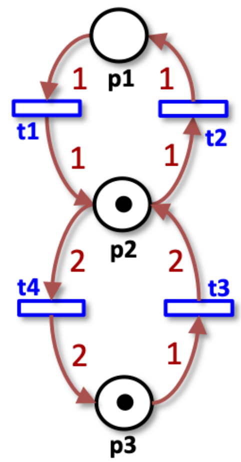

(b) Consider the following incidence matrix:

$$A = \begin{bmatrix} -2 & 1 & 1 \\ 1 & -1 & 0 \\ 1 & 0 & -1 \\ 0 & -3 & 1 \end{bmatrix}$$

i. $\boxed{\text{3 points}}$ Draw the Petri Net that is associated with incidence matrix A.

ii. $\boxed{\text{1 point}}$ Is there any initial marking with at least 4 tokens in total, for which some of the transitions cannot be fired in any firing sequence? Justify your answer with a proof or an example.

iii. $\boxed{\text{1 point}}$ If the initial marking m_0 is as follows:

$$m_0 = \begin{bmatrix} 3 \\ 1 \\ 1 \\ 0 \end{bmatrix}$$

and t_1 fires, calculate analytically, i.e., using algebraic equations, what the obtained marking m_1 will be. Please include the calculations in your answer.

(c) Consider the following Petri Net:

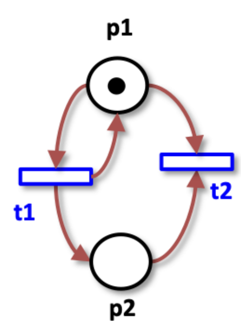

i. $\boxed{\text{2 points}}$ What is the purpose (or, usage) of a coverability tree?

ii. $\boxed{\text{3 points}}$ Construct the coverability tree for the above Petri Net .

(d) $\boxed{\text{5 points}}$ Reason the maximum number of jobs that can be executed concurrently in the system that the following Petri Net models.

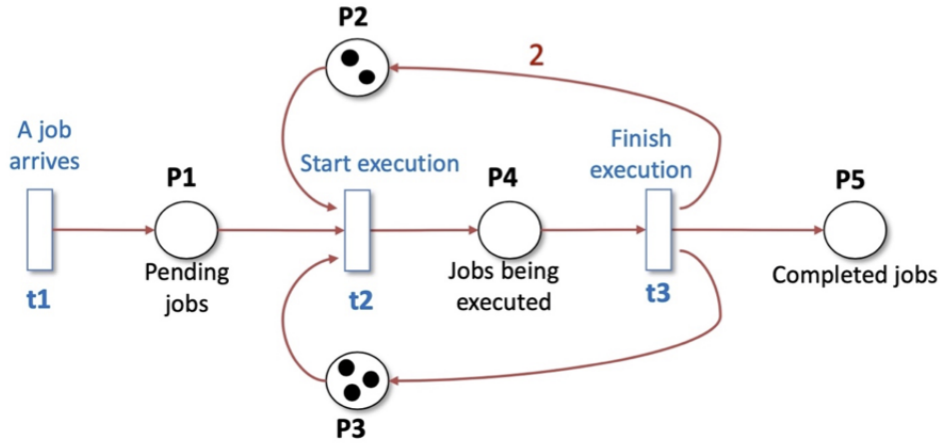

Question 5 [15 points]

Consider the Discrete Markov Chain below with state-space $S = \{A, B, C\}$:

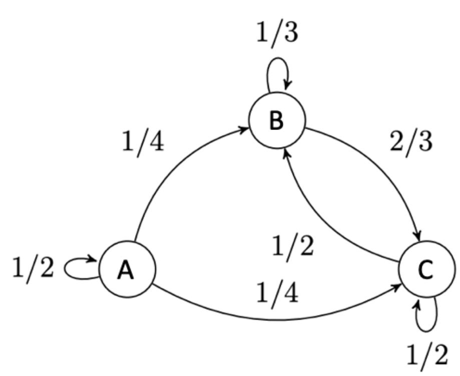

Assume that $X_t = X_0, X_1, \cdots, X_n, \cdots$ denotes the state of the system at time t, and the initial distribution is $\pi_0 = [1, 0, 0]$.

(a) $\boxed{\text{1 point}}$ $P(X_5 = B \mid X_4 = C, X_3 = B, X_2 = A, X_1 = A) = ?$

(b) $\boxed{\text{1 point}}$ $P(X_3 = B, X_2 = B \mid X_1 = A) = ?$

(c) $\boxed{\text{4 points}}$ $P(X_2 = B) = ?$

(d) $\boxed{\text{1 point}}$ Explain why the chain is reducible.

(e) $\boxed{\text{1 point}}$ In general, with reducible chains, we cannot guarantee that the steady-state is always the same (independently of the initial state). Explain why that is the case.

(f) $\boxed{\text{1 point}}$ How many classes does the chain have? Explain why.

(g) $\boxed{\text{1 point}}$ Are the classes recurrent or transient? Explain why.

(h) $\boxed{\text{5 points}}$ Even though the above chain is reducible, we can still obtain the steady-state for any initial state. Explain why that is the case and calculate the steady-state.

Question 6 [15 points]

Consider the following interactive system:

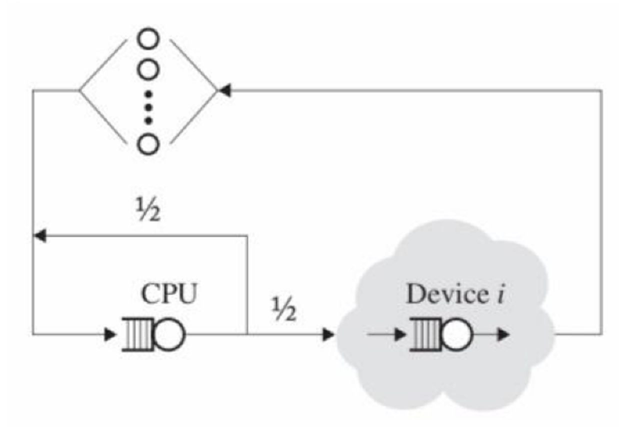

in which users generate requests that need somewhat complex computations as well as post processing in the cloud that involves a Device i amongst others. Through an extensive measuring campaign the following facts were uncovered:

mean user think time = 30 seconds
mean service time at Device i = 15 ms
utilization of Device i = 0.3
#visits-per-request to Device i = 10
#jobs in the cloud subsystem = 31
mean total time in the system (including think time) = 50 seconds

(a) $\boxed{\text{3 points}}$ Compute the overall throughput of the system. Show your calculations and mention which operational law(s) you apply.

(b) $\boxed{\text{3 points}}$ What is the average number of users that is eagerly waiting for the results of their request? (Mention the operational law(s) that you apply).

(c) $\boxed{\text{2 points}}$ What is the mean number of requests in the processing part (CPU + queue) of the system?

Consider the processing part of the system in which a user request may be processed multiple times as jobs only exit the system with a certain probability, say p:

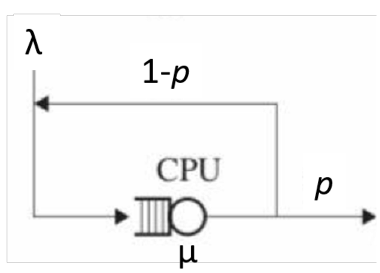

(d) $\boxed{\text{5 points}}$ Reason that the loop-back shown above is equivalent to an M/M/1 system with service rate $p\mu$.

(e) $\boxed{\text{2 points}}$ What is the CPU utilization in the complete interactive system?

Question 7 [15 points]

A small telephone switchboard has four lines. Calls during the busiest part of the day arrive in a Poisson process with rate 80 per hour and the lengths of the calls are independently exponentially distributed with mean length 3 minutes. Any call that is made when all the lines are busy is lost.

(a) $\boxed{\text{3 points}}$ Draw the state transition diagram for the calls connected by the switchboard.

(b) $\boxed{\text{5 points}}$ Determine the stationary distribution of the number of ongoing calls.

(c) $\boxed{\text{2 points}}$ What proportion of calls are lost?

(d) $\boxed{\text{4 points}}$ It is proposed to increase the size of the switchboard to six lines. What will be the proportion of lost calls in this case?

(e) $\boxed{\text{1 point}}$ Will the alternative of splitting the district into two parts, each being served by a 3-line switchboard reduce the fraction of lost calls even further?