TECHNISCHE UNIVERSITEIT DELFT
Faculteit Elektrotechniek, Wiskunde en Informatica

IN4390 – Quantitative Evaluation of Embedded Systems

January 30th, 2020, from 09:00 to 12:00

Mitra Nasri, Koen Langendoen, Lydia Chen, Marco Zuniga

Question:	1	2	3	4	5	6	7	Total
Points:	10	10	7	18	15	15	15	90
Score:								

- This is a closed book exam
- You may use a **simple** calculator only (i.e. graphical calculators are not permitted)
- Write your answers with a black or blue pen, not with a pencil
- **Always justify your answers, unless stated otherwise**

- 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	$W = R - S = \rho/(\mu-\lambda)$

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 $= \frac{\text{SSE}}{a(r-1)}$		

Washington University in St. Louis CSE567M ©2008 Raj Jain

Question 1 [10 points]

Answer the following short questions.

(a) 6 points List one advantage and one disadvantage of

(i) analytical modeling

(ii) simulation

(iii) measurement-based evaluation techniques.

(b) 3 points List at least three sources of delay that are measured by **wall-clock time**, **system time**, and **user time**.

(c) 1 point In which situation may **wall-clock time** become smaller than **system time + user time**?

Question 2 [10 points]

A two-factor ANOVA table is obtained through a factorial design.

Two-way ANOVA: y versus, A, B

Source	DF	SS	MS	F
A	1	0.322	___	___
B	___	80.554	40.2771	4.59
Interaction	___	___	___	___
Error	12	105.327	8.7773	
Total	17	231.551		

(a) 2 points How many levels were used for factor B?

(b) 2 points How many replicates of the experiments were performed?

(c) 4 points What are the F statistics for factor A, and factor AB (interaction)?

(d) 2 points Which factors are significant given the following F Distribution Table for $\alpha = 5\%$?

$F_{1,12}$	$F_{2,12}$	$F_{3,12}$	$F_{4,12}$	$F_{5,12}$	$F_{6,12}$	$F_{7,12}$	$F_{8,12}$	$F_{9,12}$
4.75	3.89	3.49	3.26	3.11	3.00	2.91	2.85	2.80

Question 3 [7 points]

A 16-run experiment was performed to figure out what conditions are critical for running deep neural networks on personal laptop. Based on the following "sign table", answer the following questions.

Run	CPU speed [GHZ] (A)	Number of cores (B)	Cache size [Kbytes](C)	Memory size [GB](D)	Number of executors (E)
1	2.1	4	64	1	6
2	4.2	4	64	1	3
3	2.1	8	64	1	3
4	4.2	8	64	1	6
5	2.1	4	512	1	3
6	4.2	4	512	1	6
7	2.1	8	512	1	6
8	4.2	8	512	1	3
9	2.1	4	64	4	3
10	4.2	4	64	4	6
11	2.1	8	64	4	6
12	4.2	8	64	4	3
13	2.1	4	512	4	6
14	4.2	4	512	4	3
15	2.1	8	512	4	3
16	4.2	8	512	4	6

(a) [5 points] Write out the alias structure.

(b) [1 point] What kind of design is this?

(c) [1 point] What is the resolution?

Question 4 [18 points]

Consider the following reachability graph.

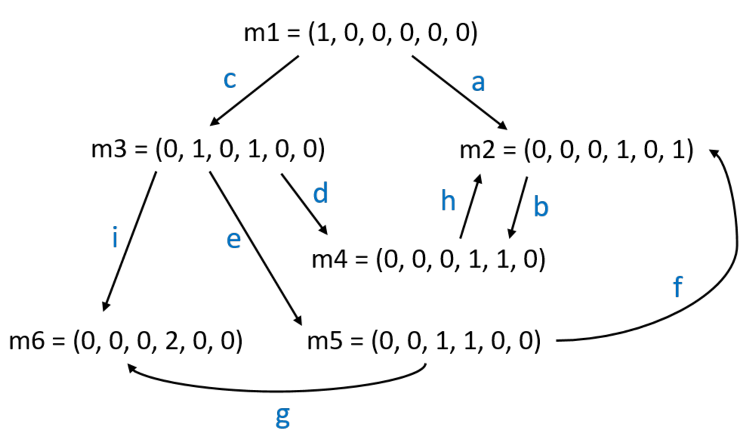

(a) [10 points] Draw the Petri net that has produced this reachability graph.

(b) [4 points] In a scenario in which a **deadlock** has happened in this Petri net, which of the transitions **could NOT** have been fired? (mention all of them)

(c) [4 points] Is there any **livelock** in this Petri net? If your answer is **no**, remove one transition (and two arcs) to create a livelock in this Petri net. If your answer is **yes**, determine which places are involved in the livelock and provide a sequence of markings that reaches to the livelock.

Question 5 [15 points]

At any given day, an embedded system is either working or being repaired. If it's working today, then there is a 95% chance it will be working tomorrow. If it is being repaired today, there is a 40% chance that it will be working tomorrow. However, if the system has been broken for four days, the overall system is replaced with a new one.

(a) [5 points] What fraction of time is the system working?

(b) [3 points] Consider the same probability for the operational part (95%), but now assume that the repair shop guarantees that it will take only one day to repair the system. What fraction of time is the system working?

(c) [3 points] Consider the original probabilities for the operational (95%) and repair (40%) parts, but now assume that the repair shop gives no guarantees. That is, the system could be repaired in one day or in an infinite number of days. What fraction of time is the system working?

(d) [4 points] For all the cases above, the fraction of time that the system is working can be generalized into a single equation that is solely a function of the transition probabilities and the number of states. Derive that equation and explain the connection with parts (a)-(c).

Question 6 [15 points]

In a gas station there is one gas pump. Cars arrive at the gas station according to a Poisson process. The arrival rate is 20 cars per hour. An arriving car finding n cars at the station immediately leaves with probability $q_n = n/4$, and joins the queue with probability $1 - q_n$, $n = 0, 1, 2, 3, 4$. Cars are served in order of arrival. The service time (i.e. the time needed for pumping and paying) is exponential. The mean service time is 3 minutes.

(a) [3 points] Draw the state transition diagram for the cars taking gas at the station.

(b) [5 points] Determine the stationary distribution of the number of cars taking gas.

(c) [3 points] Determine the mean number of cars taking gas.

(d) [2 points] Determine the mean time spent at the station (waiting time plus service time) of the cars taking gas (tricky!).

(e) [2 points] Determine the mean time spent at the station of **all** arriving cars.

Question 7 [15 points]

When designing an embedded system multiple factors should be taken into account. Consider the case of selecting the "right" Arduino model for your next robot project. The table below gives an overview of the specifications of a range of models differing in speed (higher = better), amount of memory (more = better) and power consumption (lower voltage = better):

model	processor	operating voltage	CPU speed	SRAM
101	Intel Curie	5.5 V	32 MHz	24 kB
Gemma	ATtiny85	3.3 V	8 MHz	0.5 kB
LilyPad	ATmega168V	2.7 V	8 MHz	1 kB
Mega 2560	ATmega2560	5 V	16 MHz	8 kB
Micro	ATmega32U4	5 V	16 MHz	2.5 kB
Pro	ATmega168	3.3 V	8 MHz	1 kB
Pro Mini	ATmega328P	3.3 V	8 MHz	2 kB
Uno	ATmega328P	5 V	16 MHz	2 kB
Esplora	ATmega32U4	5 V	16 MHz	2.5 kB
Leonardo	ATmega32U4	5 V	16 MHz	2.5 kB
Mini	ATmega328P	5 V	16 MHz	2 kB
Nano	ATmega168	5 V	16 MHz	1 kB

(a) [2 points] Provide the definition of Pareto dominance for the Arduino case. That is, define when model X outperforms model Y.

(b) [3 points] Enumerate the models that are part of the Pareto set (front).

(c) [4 points] One way of handling a multi-objective design space is scalarization, which effectively projects the Pareto front onto a one-dimensional evaluation criterion. In the Arduino case one can, for example, apply a weighted sum with factors $W_{speed} = 1, W_{memory} = 2, W_{power} = -3$. What is then the best (= maximum value) model(s) with this scalarization?

(d) [6 points] A second approach to multi-objective optimization is to apply a genetic algorithm. In the environmental selection step the candidate set of "new" solutions needs to be shrunk. To keep a diverse set of solutions the hypervolume indicator is known to be a good approach. Which three points would remain from the Pareto set enumerated in (b)?

Hint: as computing 3D hypervolumes is complex, it is enough to outline the procedure, and make an educated guess (explain your reasoning).