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

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.

> **Solution:**
>
> (i) **Positive**: (1) Provides the best insight into the effects of different parameters and their interaction, and (2) it can be done before the system is built, (3) it is usually fast
> **Negative**: these models (1) are rarely accurate, (2) usually need many simplifying assumptions, and (3) their correctness depend on the quality and correctness of these simplified assumptions
>
> (ii) **Positive**: (1) provides a full control of simulation model, parameters, level of detail (high flexibility), and (2) Can be done before the system is built.
> **Negative**: these models (1) may still include simplifying assumptions, (2) simulations might be inaccurate in comparison to the actual performance
>
> (iii) **Positive**: (1) it is the true to the actual performance and is the most convincing evaluation
> **Negative**: has high costs since it requires a full implementation of the system

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

> **Solution:** (1) Preemptions and their overhead such as context switch overhead, (2) I/O delays, and (3) suspensions

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

> **Solution:** This can happen, for example, when there are multiple processors in the system and the program runs in parallel on those processors.

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?

Solution: MSB = SSB/DFB, so DFB = SSB/MSB = 80.554/40.2771 = 2. Then factor B had 3 levels.

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

Solution: DFE = ab(r-1), so r = 12/(2×3) + 1 = 3 replications.

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

Solution: FA = MSA/MSE, with MSA=SSA/DFA = 0.322/1. Thus FA = 0.322/8.7773 = 0.036.

For factor AB we have to do a bit more work. Knowing that SSE = SST - SSA - SSB - SSAB, we can compute SSAB = SST - SSA - SSB - SSE = 45.348. Nest we compute MSAB = SSAB/DFAB = 45.38 / (1×2) = 22.67, and finally arrive at FAB = MSAB/MSE = 22.67/8.7773 = 2.58.

(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

Solution: FA=0.036 < F1,12=4.75, Factor A is not significant

FB=4.59 > F2,12=3.89, **Factor B** is significant

FAB=2.58 < F2,12=3.89, Factor AB is not significant

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.

Solution: Key is finding the generator: I=ABCDE, which follows from "seeing" E=ABCD (or any other alias).

The complete alias structure is then as follows: A=BCDE, B=ACDE, C=ABDE, D=ABCE, E=ABCD, AB=CDE, AC=BDE, AD=BCE, AE=BCD, BC=ADE, BD=ACE, BE=ACD, CD=ABE, CE=ABD, DE=ABC

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

Solution: One half factorial design, 2^{5-1}.

(c) 1 point What is the resolution?

Solution: The resolution follows from the generator: 5

Question 4 [18 points]

Consider the following reachability graph.

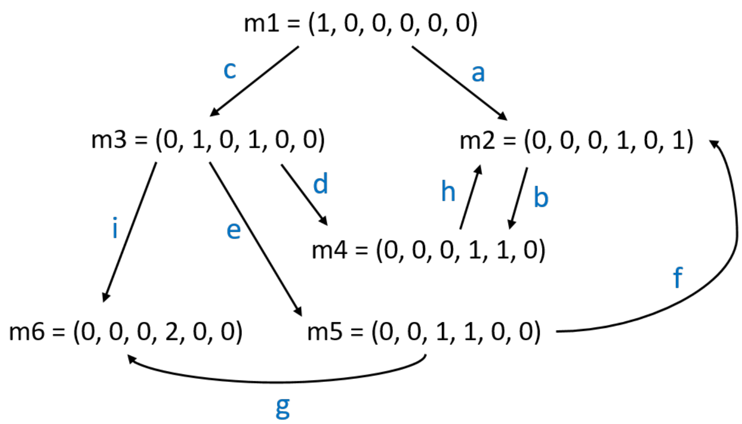

(a) 10 points Draw the Petri net that has produced this reachability graph.

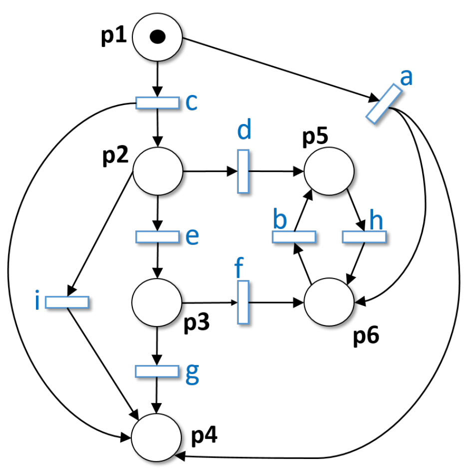

Solution:

(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)

Solution: In a deadlock scenario, a, d, f, b, h could not have been fired because as soon as they fire, the PN ends up in a livelock (or stays in the livelock).

(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.

Solution: There is a livelock that involves p_5 and p_6. It can be reached from any of the following markings: $\langle m_1, m_2 \rangle$, $\langle m_1, m_3, m_4 \rangle$ and $\langle m_1, m_3, m_5, m_2 \rangle$.

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.

Solution: All the chains are irreducible and aperiodic (the aperiodicity is due to the self-loop in the working state). Thus we can use the steady-state formulas.

From now on we assume that for all the parts in this question, the working state is denoted as S_0, and the repair states as S_1 (first day in the repair shop), S_2 (second day in the repair shop), etc.

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

Solution:

$$P = \begin{pmatrix} 0.95 & 0.05 & 0.00 & 0.00 & 0.00 \\ 0.40 & 0.00 & 0.60 & 0.00 & 0.00 \\ 0.40 & 0.00 & 0.00 & 0.60 & 0.00 \\ 0.40 & 0.00 & 0.00 & 0.00 & 0.60 \\ 1.00 & 0.00 & 0.00 & 0.00 & 0.00 \end{pmatrix}$$

For all $i > 0$, the equations from PI×P=PI follow the pattern: $S_i = S_0 \times 0.05 \times 0.6^{(i-1)}$
Utilizing the above pattern and $\sum S_i = 1$, we obtain $S_0 = 0.902$

(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?

Solution:

$$P = \begin{pmatrix} 0.95 & 0.05 \\ 1.00 & 0.00 \end{pmatrix}$$

After solving the system, $S_0 = 0.952$

(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?

Solution: This is exactly the same problem we solved in class but with a different wording

$$P = \begin{pmatrix} 0.95 & 0.05 \\ 0.40 & 0.60 \end{pmatrix}$$

After solving the system, $S_0 = 0.89$

(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).

Solution: The pattern given for S_i in Part-(a) is the hint for the solution.
Following some algebra: $S_0 = 1/(1 + 0.05 \times \sum_{i=1}^{N} 0.6^{(i-1)})$
One can then show that

- N=4 in the above equation gives the solution for Part-(a)
- N=1 in the above equation gives the solution for Part-(b)

- N=infinity in the above equation gives the solution for Part-(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.

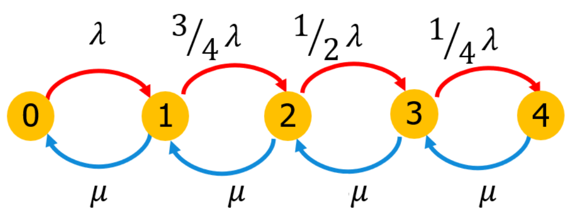

Solution:

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

Solution: $\lambda = 20/60 = 1/3$ and $\mu = 1/3$, so $\rho = \lambda/\mu = 1$ (which makes for easy calculations).

$\lambda P_0 = \mu P_1 \quad P_1 = \rho P_0 = P_0$
$\frac{3}{4}\lambda P_1 = \mu P_2 \quad P_2 = \frac{3}{4}\rho P_1 = \frac{3}{4} P_0$
$\frac{1}{2}\lambda P_2 = \mu P_3 \quad P_3 = \frac{1}{2}\rho P_2 = \frac{3}{8} P_0$
$\frac{1}{4}\lambda P_3 = \mu P_4 \quad P_4 = \frac{1}{4}\rho P_3 = \frac{3}{32} P_0$

Factoring in that $\sum_{n=1}^{4} P_n = 1$ yields the following expression:
$P_0 = \frac{1}{1+1+3/4+3/8+3/32} = 32/103$.

As such the distribution becomes:
$P_0 = 32/103$
$P_1 = 32/103$
$P_2 = 24/103$
$P_3 = 12/103$
$P_4 = 3/103$

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

Solution: We need to compute $N = \sum_{n=1}^{4} n \times P_n = 128/103 = 1.24$

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

Solution: The tricky part is that some of the cars leave immediately, so we cannot apply Little's law using the arrival rate λ. We need to compute the effective arrival rate for the cars taking gas:

$\lambda_{gas} = \lambda P_0 + \frac{3}{4}\lambda P_1 + \frac{1}{2}\lambda P_2 + \frac{1}{4}\lambda P_3 = \frac{71}{103}\lambda$

Now we can apply Little's law:

$R = N/\lambda_{gas} = N \times \frac{103}{71} \times 3 = \frac{128}{103} \times \frac{309}{71} = 5.4\,\text{min}$

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

Solution: The cars taking **no** gas don't wait, so when using using Little's Law with the overall arrival rate λ we compute the average response time of all cars **arriving** at the station:

$R = N/\lambda = 3 \times N = 384/103 = 3.7\,\text{min}$

Alternatively one can reason that the response time from (d) is diluted with all cars leaving immediately, so

$R = (1 - \lambda_{gas}) \times 0 + \lambda_{gas} \times R_{gas} = \frac{71}{103} \times (\frac{128}{103} \times \frac{309}{71}) = \frac{128 \times 3}{103} = 3.7\,\text{min}$

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.

> **Solution:** X has to better in one, and better or equal in all other dimensions than Y, with better being higher speed, more memory, or lower operating voltage.

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

> **Solution:** Note that all processors running at a certain voltage run at the same speed, which makes for easier comparison. We have a set of just 4 elements: {**LilyPad**=(2.7 V, 8 MHz, 0.5 kB), **Pro Mini**=(3.3 V, 8 MHz, 2 kB), **Mega 2560**=(5 V, 16 MHz, 8 kB), **101**=(5.5 V, 32 MHz, 24 kB)}.

(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?

> **Solution:**
>
> As the Pareto dominance relation is preserved*+ under scalarization, we only need to consider the points on the Pareto front: {**LilyPad**=1.9, **Pro Mini**=2.1, **Mega 2560**=17, **101**= 63.5}. Thus the **101** is the best.
>
> *When larger (smaller) is better and all weights are positive it holds that if a point P dominates point Q in the original design space, then that will also be the case after scalarization with Scal(P) being larger (smaller) than Scal(Q).
>
> +In the Arduino case the negative weight transforms the voltage into a "bigger is better metric": $-3 \times V = 3 \times -V$.

(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).

Solution: We have to make a number of choices in order to compute the hypervolumes of the four options (we need to drop 1 of the 4 Pareto points).

First, it's handy if all metrics have “bigger is better”, so lets reverse the operating voltage axis as in RV = 10 - V. That gives us the set: { **LilyPad**=(7.3 RV, 8 MHz, 0.5 kB), **Pro Mini**=(6.7 RV, 8 MHz, 2 kB), **Mega 2560**=(5 RV, 16 MHz, 8 kB), **101**=(4.5 RV, 32 MHz, 24 kB) }.

Second, we need to decide on a reference set R. Let's simply pick the origin: R=(0,0,0).

Third, we need to compute the volumes that get lost when dropping a point. Note that the differences in coordinates between **LilyPad** and **Pro Mini** are small compared to the other coordinate differences. Dropping **LilyPad** reduces the hypervolume with $0.6 \times 8 \times 0.5 = 2.4$, while dropping **Pro Mini** reduces the hypervolume with $1.7 \times 8 \times 1.5 = 20.4$. Thus we “conclude” that dropping **LilyPad** has the least effect, hence, the remaining set would be {**Pro Mini**, **Mega 2560**, **101**}.