EECS 507: Introduction to Embedded Systems Research
Midterm Exam

Robert Dick

25 October 2022

You have 80 minutes for the exam.

Closed book. Closed notes. No communicating with anybody except the teacher about the exam. This remains true even after you have submitted the exam. I'll tell you when it is O.K. to discuss. Someone is taking it late due to COVID-19.

There are answer length limits to control exam duration. For the sake of fairness to students who honor the limits, if you exceed the length limits, I will evaluate only the portion of the answer within the length limit.

Please use × marks for checkboxes.

Skim all the questions before starting so you can budget your time. They have different difficulties, but each is worth similar credit; that's intentional. Do the problems that are easiest for you first if you are concerned about the time limit.

Printed name ☐

Sign below to acknowledge the Engineering Honor Code: "I have neither given nor received aid on this examination, nor have I concealed a violation of the Honor Code."

Signature ☐

$\boxed{10}$ 1. Literature survey: Which one of the following methods is generally least effective for finding research papers during a literature survey? [Lecture]

Answer: Searching using a general-purpose search engine.

$\boxed{10}$ 2. Specification and modeling: In the diagram below, at time zero, only A has a token (a small circle), i.e., only A is ready to fire. The numbers near the sources of arcs indicate the numbers of tokens consumed to fire once. The numbers near the destinations of arcs indicate the number of tokens generated when the arc fires. [Jantsch and Sander '05]

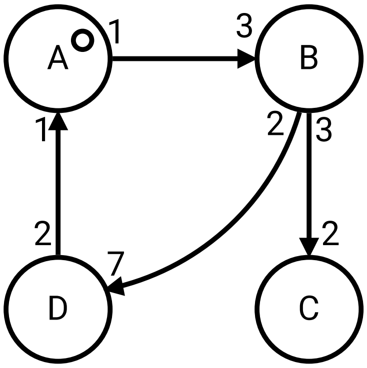

What is the maximum number of tokens that can ultimately exist in the system?
Answer: ∞.

$\boxed{10}$ 3. Optimization: If one started from a Parallel Recombinative Simulated Annealing algorithm but replaced Boltzmann trials with trials that always select the solution with lower cost, which (single) optimization approach would this result in? [Lecture]
Answer: genetic algorithm.

$\boxed{10}$ 4. Power: Voltage scaling is unlikely to bring as much benefit in the future as it has brought in the past and this is largely due to the tightening range of permissible operating voltages. [Lecture]

Use at most one sentence to describe the most important lower bound on the minimum operating voltage.
Answer: Dropping V_{DD} below V_T makes the circuit slow and generally increases energy consumption due to the increase in delay exceeding the decrease in power consumption.

Use at most one sentence to describe the most important upper bound on the maximum operating voltage.
Answer: Multiple answers including power consumption increases with V_{DD} so power constraints bound it or field strength in the very thin gate dielectric increases with V_{DD} so reliablity constraints bound it.

$\boxed{10}$ 5. Temperature: In the simple thermal model, below, use at most one short phrase to indicate the purpose served by the voltage source. [Lecture, General Reasoning]

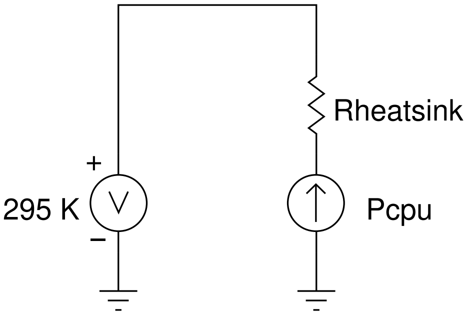

Answer: It represents the ambient temperature, in kelvin. Many students conflated power (i.e., heat) with temperature in their answers. Some gave a definition of a voltage source, which doesn't explain the purpose of that component in the circuit.

$\boxed{10}$ 6. Sensor networks: Use at most one sentence to indicate why geophones are well suited to the application in this paper? [Bonde et al. '21]
Answer: Sensitive in the appropriate frequency range.

$\boxed{10}$ 7. Wireless communication: Use at most threes phrases to indicate the three most important characteristics of LPWANs that differentiate them from most prior communication technologies. [Raza et al. '17]
Answer: Long ranges of many km, low power to the point of enabling multi-year battery lifespans with compact batteries, inexpensive, and low data rate.

$\boxed{10}$ 8. Reliability: Consider a system containing three software components: A, B, and C. The system functions correctly if C is functioning correctly. A and B are non-essential, but enable some improvement in functionality. A and B communicate with each other to control each other's rates, and therefore computational loads. If either one of these two fails, the other will overload the processing resources of the system and prevent C from executing. This failure mode is in addition to that in which C simply fails on its own. If both A and B fail, the overload problem won't occur. The fault processes for the different components are independent. They have the following fault probabilities A: 0.1, B: 0.2, and C: 0.1. [Lecture]

What is the probability of system failure?
Answer: 0.334

One might also follow the slow but general approach presented in class.

A works	B works	C works	System works	Probability
0	0	0	0	$0.1 \cdot 0.2 \cdot 0.1 = 0.002$
0	0	1	1	$0.1 \cdot 0.2 \cdot 0.9 = 0.018$
0	1	0	0	$0.1 \cdot 0.8 \cdot 0.1 = 0.008$
0	1	1	0	$0.1 \cdot 0.8 \cdot 0.9 = 0.072$
1	0	0	0	$0.9 \cdot 0.2 \cdot 0.1 = 0.018$
1	0	1	0	$0.9 \cdot 0.2 \cdot 0.9 = 0.162$
1	1	0	0	$0.9 \cdot 0.8 \cdot 0.1 = 0.072$
1	1	1	1	$0.9 \cdot 0.8 \cdot 0.9 = 0.648$

$1 - 0.018 - 0.648 = 0.334$.

Alternatively, we can to do a little more thinking and less calculation. For example, we can treat the system as if it contains two components, the A/B group and C. It fails if either A/B or C fails. That yields the following results. A/B fails with probability $0.1 \cdot 0.8 + 0.9 \cdot 0.2 = 0.08 + 0.18 = 0.26$ and C fails with probability 0.1. We cannot just sum these two to get the system failure probability because C may fail when AB fails. Instead, $1 - (1 - 0.26) \cdot (1 - 0.1) = 1 - 0.74 \cdot 0.9 = 1 - 0.666 = 0.334$. This way requires only one pen-and-paper multiplication operation.

$\boxed{10}$ 9. Autonomous vehicles: Inference computation, e.g., LiDAR analysis and computer vision, accounts for a substantial percentage of locomotion energy consumption in self-driving cars. What would you expect this percentage to be in scaled-down self-driving vehicles roughly 1/10 the size and mass of a normal automobile? [General reasoning, Dick et al. '22]
Answer: Much more than 1/3. Depending on assumptions and shown work, answers anywhere near 80% were accepted. Those who grasped that the computational problems remained similar but the locomotion energy reduced greatly and recalled that the computational energy for the full-sized car is around 1/3 the total energy received full credit. Other conclusions based on reasonable assumptions and deductions also received credit.

$\boxed{10}$ 10. Memory hierarchy: At most one sentence to indicate why scratchpad memories are more widely used (and presumably easier to use effectively) in embedded applications than in general-purpose computing applications. [Banakar et al. '02]
Answer: The memory access patterns in many embedded applications are more predictable than many general-purpose computing applications, enabling the compile-time knowledge needed for using scratchpad memories. Another valid answer would be that scratchpads are often more energy efficient than caches, and energy efficiency is often a larger concern in embedded systems than in general-purpose computing. Arguments that they are especially useful in real-time systems due to greater predictibility than caches were also taken. Many students thought the question asked what scratchpads are or why they are good. Nope. It asked for a why they are more common (or better) for embedded applications than for general-purpose computing.