

EECS 507: Introduction to Embedded Systems Research  
Midterm Exam  
29 April 2021

Robert Dick

You have until 100 minutes for the exam, and a 20 minutes grace period to upload it to Gradescope (or, if special arrangements were made, to email it to me). If you email, you should not consider it submitted until you receive confirmation from me. It must be to me by 12:30pm.

Open book. Open notes. Open internet. 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.

If you can print and scan/photograph, do that and upload to Gradescope instead of using the forms. If you cannot print and scan/photograph, then use the forms to the degree possible and submit the exam to me via both Gradescope and email. Only those using forms need also email.

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

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.

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

- 10 1. What is the primary argument made in R. P. Dick, L. Shang, M. Wolf, and S.-W. Yang, “Embedded Intelligence in the Internet-of-Things,” *IEEE Design & Test of Computers*, Dec. 2019 to support the claim that edge devices will carry out sophisticated analysis locally instead of transmitting raw data to the cloud for analysis? Use at most two sentences.

Soln: The compute energy per bit is orders of magnitude lower than the communication energy per bit for wireless edge devices.

- 10 2. Referring to L. Yang, R. P. Dick, H. Lekatsas, and S. Chakradhar, “High-performance operating system controlled on-line memory compression,” *ACM Trans. Embedded Computing Systems*, vol. 9, no. 4, pp. 30:1–30:28, Mar. 2010 and S. Han, X. Liu, H. Mao, J. Pu, A. Pedram, M. A. Horowitz, and W. J. Dally, “EIE: Efficient inference engine on compressed deep neural network,” in *Proc. Int. Symp. Computer Architecture*, June 2016, under what circumstances would using the PBPM compression algorithm instead of the weight indexing approach in EIE result in improved energy efficiency? Use at most three sentences and be careful to focus on the most important factors.

Soln: If weights frequently shared many high-order bits, such an approach might outperform straight-forward indexing, although increase in match detection latency might overwhelm the improvement. Many other reasonable answers were also accepted.

- 10 3. As a result of process variation, different chips and memristors of the sort described in P. M. Sheridan, F. Cai, C. Du, W. Ma, Z. Zhang, and W. D. Lu, “Sparse coding with memristor networks,” *Nature Nanotechnology*, vol. 12, Aug. 2017 may have different response magnitudes to the same training stimuli. Will this prevent different chips from producing similar output results when used as Sheridan et al. describe? Why? Use at most three sentences.

Soln: No. Each chip instance would be trained independently, so the training process could compensate for instance-to-instance variation.

- 10 4. In Y. Zhu, A. Samajdar, M. Mattina, and P. Whatmough, “Euphrates: Algorithm-SoC co-design for low-power mobile continuous vision,” arXiv, Tech. Rep., Apr. 2018, what is the most important advantage and disadvantage of increasing the  $d$  parameter for motion vector calculation? Use at most two sentences.

Soln: The advantage is finding accurate matches with larger frame-to-frame block motion. The disadvantage is higher computational cost for each motion vector computation.

- 10 5. What role did percolation theory play in E. Ronen, A. Shamir, A.-O. Weingarten, and C. O’Flynn, “IoT goes nuclear: Creating a ZigBee chain reaction,” in *Proc. Symp. on Security and Privacy*, May 2017? Use at most two sentences.

Soln: Percolation theory enabled the authors to estimate the effect of changing wireless device density on virus propagation.

- 10 6. What specific IMU-based method did Bhatt, Minnehan, Radhakrishnan, and Thomas, and use to estimate velocity? What is typically the main source of error when using this method? Use at most two sentences.

Soln: They integrated acceleration to estimate velocity. This is prone to drift error, which typically increases with time. Many reasonable answers were accepted.

- 10 7. What was the main question Davies and Nguyen were attempting to answer? Use at most two sentences.

Soln: How do image-distorting weather effects change the accuracy of Digital Foveation enhanced object detection in traffic monitoring scenarios? Many answers indicating compar-

ison of training a system incorporating Digital Foveation on un-weathered and weathered images were also accepted.

- 10 8. Evaluate the failure probability of a quadcopter system containing the following three subsystems.
1. The sensing subsystem, for which at least two sensors must be operational.
  2. The processing / analysis subsystem, for which at least one of the two CPU cores must be operational.
  3. The actuation subsystem, within which there are four propellor motors with two able to rotate propellers clockwise and two able to rotate propellers counter-clockwise, and for which at least one clockwise and one counter-clockwise motor must be operational.

The component fault probabilities follow, and are uncorrelated.

Component	Fault probability
LiDAR sensor	0.01
Camera sensor	0.4
Sonar sensor	0.2
CPU core	0.001
Motor	0.005

Soln:

	State	Probability
	000	0.0008
	001	0.0032
	010	0.0012
Stage 1	011	0.0048
	100	0.0792
	101	0.3168
	110	0.1188
	111	0.4752

1-0.9156 for first stage.

0.999999 for second stage.

Really two motor stages, each with 0.999975 probability

Answer:  $1 - 0.915533 = 0.084467$

- 10 9. Redemption: This question counts for 10 points on the final exam and also gives you the opportunity to redeem your score for the most similar midterm question. Give the Pareto-rank of each of the following solutions, given that price, power consumption, and delay are all costs.

Solution	Price (\$)	Power (mW)	Delay (ms)	Pareto-rank
1	1	199	11	Soln: 3
2	18	320	26	Soln: 2
3	24	225	27	Soln: 2
4	10	489	37	Soln: 2

- 10** 10. **Redemption:** This question counts for 10 points on the final exam and also gives you the opportunity to redeem your score for the most similar midterm question. Has optimally solving the traveling salesman problem been proven to require exponential worst-case time in terms of problem instance size for a deterministic computer? No.    Yes.

Soln: No.