MIT 6.S062 2017 – Mobile and Sensor Computing
Exam 1

There are 14 questions and 11 pages in this quiz booklet. To receive credit for a question, answer it according to the instructions given. *You can receive partial credit on questions.* You have **90 minutes** to answer the questions.

Don't forget to write your name on this cover sheet NOW!

Some questions may be harder than others. Attack them in the order that allows you to make the most progress. If you find a question ambiguous, be sure to write down any assumptions you make. Be neat. If we can't understand your answer, we can't give you credit!

THIS IS A CLOSED BOOK QUIZ.
YOU MAY USE TWO DOUBLE-SIDED PAGE OF NOTES.

Do not write in the boxes below

1-2 (12)	3-4 (12)	5-6 (16)	7-10 (23)	11-12 (10)	13-14 (17)	Total (90)

Name:

Location

1. *(5 points)* We studied several papers on localization. These papers described systems or methods using one or more underlying technologies:

   A. Radio frequency (RF) signals

   B. Ultrasonic signals

   C. Acoustic signals

   For each method or system below, as described in the corresponding papers, which of the above technologies are used? (More than one choice from A, B, C may apply in each case below.)

   (a) GPS:

   (b) RADAR from Microsoft:

   (c) Cricket from MIT:

   (d) Counter-sniper system from Vanderbilt (PinPtr):

2. *(1+2+2+2=7 points)* Circle **True** or **False** for the following statements about GPS.

   (a) **True** / **False** GPS provides information about time as well as position.

   (b) **True** / **False** One of the reasons why GPS consumes significant amounts of energy on mobile devices is that the weak signal-to-noise ratios of the satellite signals require computationally expensive algorithms to decode.

   (c) **True** / **False** A GPS receiver uses information about the known positions of satellites in the sky to estimate angles of arrival (AoA) to estimate its own position.

   (d) **True** / **False** Suppose you take your smartphone from the US to Europe on a flight during which your phone is off, and you turn it on when you land in Europe. In general, if the phone has Internet connectivity in Europe, it is more likely to obtain a GPS satellite fix quicker, compared to the case without any network connectivity.

Performance metrics

3. *(4 points)* In the lecture on network connectivity options for IoT systems, we studied various options. Circle **True** or **False** for each statement below.

   (a) **True** / **False** Bluetooth Low Energy (BLE) uses a time division MAC protocol, which allows the radio to save energy by sleeping except when required to be awake by the time schedule.

   (b) **True** / **False** The typical single-hop communication range of BLE communication is smaller than Wi-Fi.

   (c) **True** / **False** Cellular connectivity consumes unnecessarily high energy even when very little data needs to be sent and received because link rates are much higher (Megabits per second) than required for this use case.

   (d) **True** / **False** Cellular connectivity consumes unnecessarily high energy even when very little data needs to be sent and received because the radio remains on in one or more ”idle” states even when no data is being delivered.

4. *(8 points)* In several papers we studied, *precision* and *recall* were used to evaluate the performance of various inference, learning, or detection methods (e.g., CTrack, Glimpse). Consider a sensor-based system set up to detect the malfunction of industrial equipment in a factory. You observe the following sequence of malfunction events detected by the system:

Time slot	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Event detection by system	x	x			x	x		x	x		x		x	x
Actual event (truth)		x	x	x		x		x	x			x	x	

For this sequence of events, what are the precision and recall metrics? Explain your answer in the space below and write the numbers in the blanks.

Precision: ______________________________.

Recall: ______________________________.

BLE energy calculation

5. *(7 points)* Alyssa P. Hacker has developed a BLE-equipped sensor node to use on a bicycle. It advertises data when the bicycle is moving, consuming on average 0.5 milliAmp in the advertisement state; at all other times, the node is on standby consuming 1 microAmp. She estimates that the average use of her bicycle (i.e., advertisement state) is 1 hour per day. Ignore the energy consumed in switching between standby and advertisement states. The node uses a coin-cell battery with capacity 250 milliAmp-hours. For how long will the node function before running out of energy?

Mesh networks

6. *(2+4+3=9 points)* Ben Bitdiddle runs experiments on a Wi-Fi mesh network with three routing protocols and produces the following cumulative distribution function chart of the performance observed. The protocols are:

   A. "**Best possible**" (based on an offline determination of routes using packet loss rate data)

   B. Routing with **minimum hop count**, where a hop exists between a pair of nodes if they can communicate at a non-zero packet success rate.

   C. Routing with the **ETX** metric.

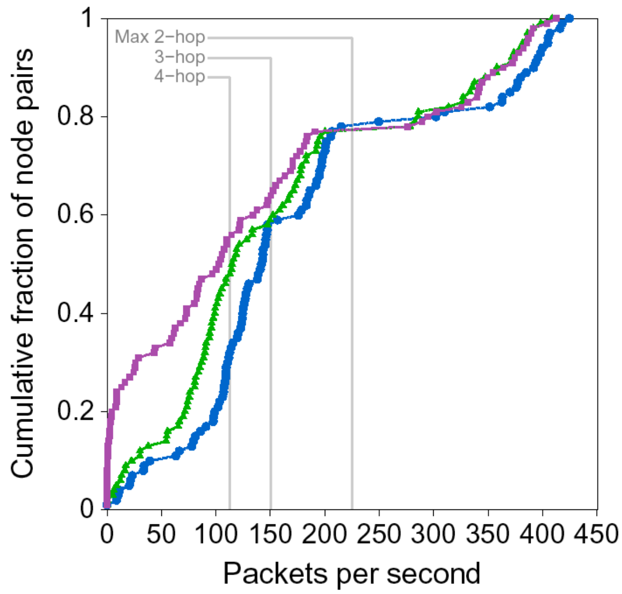

(a) In the picture above, **clearly** label each protocol (A, B, C) next to the corresponding CDF curve.

(b) Circle **True** or **False** for these statements:

   i. **True** / **False** At least 20% of the node pairs shown in the chart communicate directly with each other with the ETX protocol.

   ii. **True** / **False** Approximately 80% of the node pairs shown in the chart in each of the protocols do not communicate directly with each other.

   iii. **True** / **False** This chart shows that the reason why ETX is a much better protocol than minimum-hop count is that the median throughput improvement is substantial.

   iv. **True** / **False** The gap between one of the curves and the other two on the right top corner is because of routing protocol overhead, which reduces the usable bandwidth of the network.

(c) On one of the curves, nearly 20% of the node pairs have a throughput close to 0. Why is that?

Pothole Patrol

7. *(6 points)* Looking at the three plots of X and Z acceleration data below, which, if any, do you believe correspond to a pothole, and why:

   Plot 1:

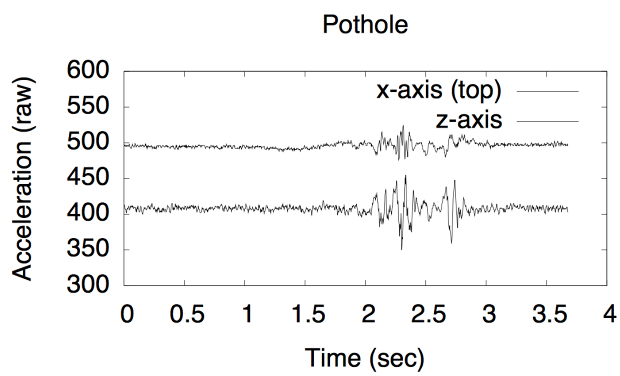

   Plot 2:

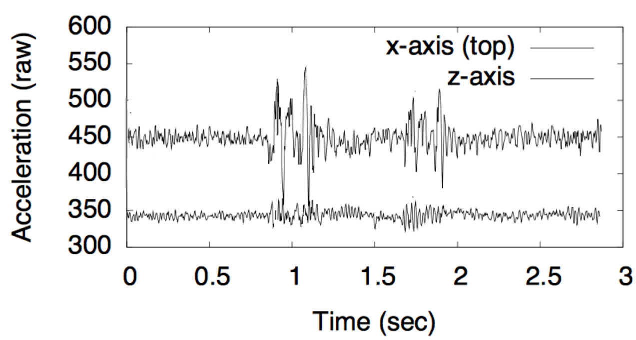

   Plot 3:

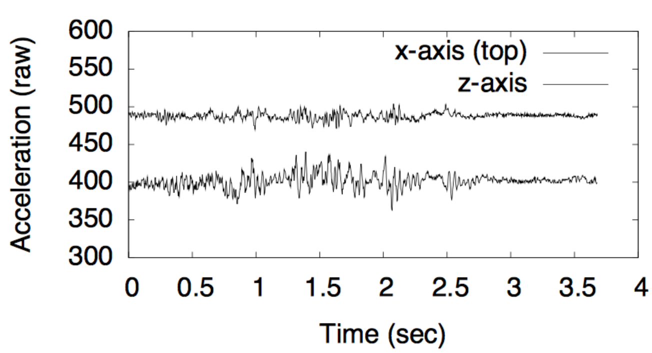

   Write your answer (including a brief justification) in the space below:

TinyDB and TAG

8. *(2+3+4=9 points)* Consider the two aggregation trees below. Estimate the sum of the total bytes sent over the radio by all nodes in the network if we are computing a single epoch's worth of data using the TAG in-network aggregation protocol. Assume all numbers (sensor values, counts, etc) require 4 bytes each, and that the root node (0) does not need to transmit any bytes over the radio, and that only data needed to compute the aggregate is transmitted (i.e., there are no node IDs/addresses, epoch numbers, etc).

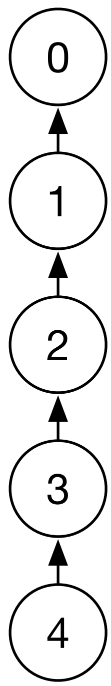

Tree 1

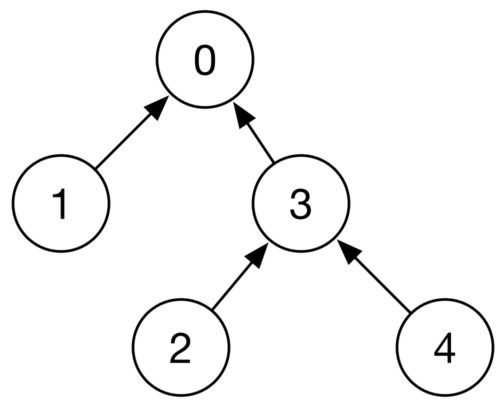

Tree 2

Write your answer in the table below:

COUNT	Total Bytes	AVERAGE	Total Bytes	MEDIAN	Total Bytes
Tree 1	______________	Tree 1	______________	Tree 1	______________
Tree 2	______________	Tree 2	______________	Tree 2	______________

Inertial sensing and activity recognition

9. *(4 points)* Louis Reasoner proposes a low-power method to estimate the velocity of a moving car traveling without turns using only a commodity smartphone-quality accelerometer and no GPS. He affixes an accelerometer rigidly to the vehicle and develops a method to estimate the matrix that transforms the axes of the accelerometer to the axes of the car. You may assume this computation is correct. Then, by integrating over time the acceleration time-series data obtained in the direction of movement of the car, he estimates the car's velocity at each point in time t, starting from a car at rest at $t = -0$:

$$v(t) = \int_{t'=0}^{t} a(t')dt'.$$

Would Louis's method work? If no, why not? If yes, how would its performance compare to GPS?

10. *(4 points)* Consider the problem of recognizing the activity a person is engaged in, using 2 two-axis accelerometers each in the same orientation on each ankle of the person. If you run a classifier over 1-second windows of accelerometer data sampled at 512 Hz, which single feature would you expect to be most effective at differentiating running from walking?

Choose the best answer (Circle One).

(a) Mean acceleration value of X or Y channel of one of the accelerometers
(b) FFT energy of the X or Y channel of one of the accelerometers
(c) Correlation between the X and Y channels on one accelerometer
(d) Correlation between the X or Y channels on the pair of accelerometers

Explain your answer in one sentence:

Map matching and map inference

11. *(4 points)* In the paper "Map inference in the face of noise and disparity", to find road centerlines, the authors develop a new gray-scale skeletonization method. This method iterates through binary skeletonization starting from the highest density level and proceeding downwards in density levels. Why does it produce a better result than using binary skeletonization once at a low density level?

12. *(6 points)* The CTrack system for map matching produces a map-matched trajectory from cellular radio data. It makes its inference using two different Hidden Markov Models. For each hidden Markov Model used in CTrack, specify the observed state (i.e., the input) and the hidden state (i.e., what is being estimated).

    First HMM Observed State:

    First HMM Hidden (Inferred) State:

    Second HMM Observed State:

    Second HMM Hidden (Inferred) State:

Shooter localization

13. *(10 points)* Consider a simplified deployment of a shooter localization sensor network similar to the PinPtr system we studied in class, as shown in the picture below. Four acoustic sensors, A through D, are located as shown. Assume that the gunshot comes from exactly one of the nine circled locations (including possibly the location of the sensors) at time 0. Sound travels at **1 foot per millisecond**. The times at which each sensor detects the acoustic muzzle blast (note: not the shockwave) are shown near each sensor. The time at which a sensor detects the muzzle blast could be ± 0.2 ms from the true time at which the sound wave would have reached it owing to errors in different components of the system.

Using the consistency method from the PinPtr system, estimate the **two most likely locations** of the shooter. **Show them in the picture by putting a square around the location and labeling them 1 (for the most likely location) and 2 (for the second-most likely location).** Explain your answer.

You may find the following numbers useful: $\sqrt{2} = 1.41$, $\sqrt{3} = 1.73$, $\sqrt{5} = 2.24$, $\sqrt{8} = 2.83$.

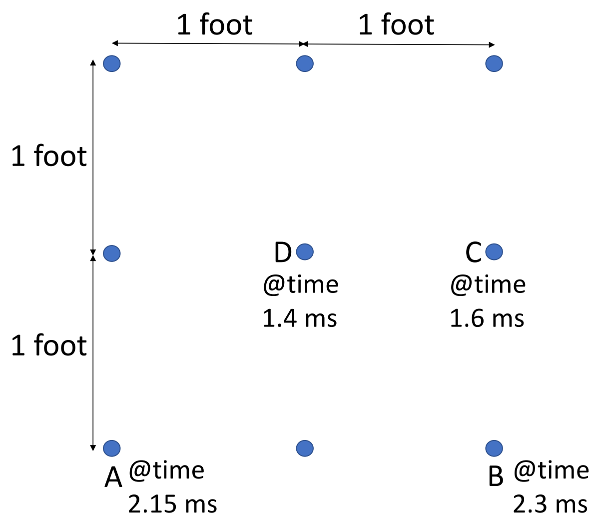

Glimpse

14. *(3+4=7 points)* The Glimpse system for continuous object recognition uses an *active cache*.

    (a) What task does the mobile client side of Glimpse perform on the active cache?

    (b) What benefit does running the above task on the active cache provide?

End of quiz!