

1.204 Quiz 1

Spring 2008

Name _____

Exam guidelines:

- 1) 80 minutes are allowed to complete the quiz.
- 2) Open notes; open book.
- 3) There are 4 questions (100 points) and 7 pages (including this one) in the exam booklet.
- 4) No laptop computers, calculators, cell phones or messaging devices are allowed. Please turn off any that you have brought.
- 5) Please write legibly – you are welcome to use both sides of the paper; we can provide additional paper if necessary.

Question 1. Data modeling (25 points)

There is a set of traffic sensors that count cars in an urban area. Some are at the approaches to intersections, one sensor per lane of traffic entering the intersection. Other sensors are between intersections, one sensor per lane of traffic on the road. An intersection has several lanes entering it. Each lane is associated with a road. Each sensor has a status (working or not) and a count of the number of vehicles it has detected, in 5 second intervals, for the previous 3 minutes. Each road and each intersection has a name and a unique identifier.

Draw a data model that corresponds to this set of system rules. You only need to create one drawing that includes all the elements listed in steps a-e.

- a. **Draw a box for each entity: give each an appropriate name**
- b. **List the attributes in the box for each entity**
- c. **Indicate the primary key for each entity by placing the phrase (PK) next to its name.**
- d. **Draw all relationships between the entities in the model. Indicate foreign keys by placing the phrase (FK) next to attributes that are foreign keys.**
- e. **Indicate the cardinality of the relationship: many-many, many-one or one-one. Use crows-foot notation; if you use another notation, define it.**

To repeat: You only need to create one drawing that includes all the elements listed in steps a-e above.

Please draw your data model on this page.

2. SQL (15 points)

Assume that the data model you built in the previous question has been implemented exactly as you drew it in a relational database management system such as SQL Server. Each entity is stored as a table, and attributes, keys and relationships are as you specified them. Write the SQL query to return the list of intersection IDs that have more than 10 cars in a 5 second interval in an approaching lane.

3. Analysis of algorithms. (25 points)

a. What is the best case running time for heapify? Describe the case, and state what its running time as a function of its size n . Use the heapify code from lecture as the implementation for which you derive the bound.

b. What is the best case running time for heapsort? Describe the case, and state what its running time as a function of its size n . Use the heapsort code from lecture as the implementation for which you derive the bound.

4. Algorithm design (35 points)

When a grain elevator opens each morning, there is queue of trucks waiting to unload that have traveled from a set of farms. The grain elevator owns the trucks and pays the drivers by the hour, and thus would like to find an efficient order in which to unload the trucks. After a truck has been unloaded the driver can make another trip. (The drivers are indifferent to the order in which trucks are unloaded since they are paid by the hour.)

The grain elevator chooses to minimize

$$T = \sum_{i=1}^n t_i$$

where T = total time in system for all trucks

t_i = time required to unload truck i

n = number of trucks

The time required to unload a truck depends on its capacity and the type and size of its unloading hatches. For example, some trucks have large, automatically operated hatches; others have small, manually operated hatches.

Design an algorithm to let the grain elevator optimize its chosen objective function.

a. Write pseudocode (please define any non-Java-like symbols or conventions you use) that defines how the algorithm works.

b. What kind of algorithm is it: divide and conquer, greedy, or ad hoc? Say why, in one sentence.

c. What is the running time of the algorithm? Use $O()$, $\Omega()$ or $\Theta()$ notation as appropriate. Briefly support your answer. No derivations are needed; justify your answer informally.

MIT OpenCourseWare
<http://ocw.mit.edu>

1.204 Computer Algorithms in Systems Engineering
Spring 2010

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.