

Name: _____

1. (1 pt.)

- **Read all material carefully.**
- You may refer to your books, papers, and notes during this test.
- No computer or network access of any kind is allowed (or needed).
- Write, and draw, carefully. Ambiguous or cryptic answers receive zero credit.
- Use the conventions used in class and the textbook for all material.
- Some questions refer to the database suggested by Figure 1 on the last page. You may detach that page for easy reference. Return it without reattaching.
- Reminder: In all query-writing questions, your answers should work on *all* database instances, not only the instance of Figure 1.

Write your name in the space provided above.

2. (14 pts.) Write a SQL query to find the batteries with the maximum reduction in measured capacity within a calendar month. Use SQL:1999 *window functions* (recall the AQuery paper¹) where appropriate, in addition to the standard SQL described in the textbook and in class. Briefly justify the correctness of your query.

Note that, over time, the measured capacity of any battery may both increase (perhaps due to refresh-charging) and decrease (perhaps due to age or poor use); there is no monotonicity in this regard.

¹Alberto Lerner and Dennis Shasha, “AQuery: Query Language for Ordered Data, Optimization Techniques, and Experiments,” in *Proceedings of the 29th International Conference on Very Large Data Bases (VLDB)* (Berlin, Germany, 2003).

[additional space for answering the earlier question]

3. (15 pts.) Repeat Question 2 (including justification of correctness) without using window functions, instead using only the SQL features described in class (and the textbook) prior to the AQuery paper.

[additional space for answering the earlier question]

4. (15 pts.) Render the query of Question 3 in extended bag algebra. Explain why the resulting algebraic query is correct.

[additional space for answering the earlier question]

5. (15 pts.) Prove or disprove: It is possible to express the following query using only the relational algebra operators of selection, projection, cross product, union, and difference: Names of all battery manufacturers M such that all batteries manufactured by M (in our database) have recorded a measured capacity (at least once) that is 1500 or lower.

[additional space for answering the earlier question]

Scratch page

Material here will not be graded.

You may detach this page. Return it without reattaching.

Batteries

tag varchar(50)	model varchar(50)	buy_date date	price float	manuf_date date	color varchar(50)	notes varchar(50)
x1	Tenergy AB	2009-01-23	2.20	2008-06-01	blue	heavy
x2	Tenergy AB	2009-01-23	2.20	2008-06-01	blue	light
pq	Tenergy AB	2010-10-03	2.25	2009-06-01	blue	check
pq2	Tenergy AB	2010-10-30	2.50	2009-06-01	blue	

Charges

tag varchar(50)	model varchar(50)	charger varchar(50)	date date	mAh float	method varchar(50)
x1	Tenergy AB	maha-101	2012-01-22	1883	charge
x1	Tenergy AB	maha-101	2012-02-12	1983	refresh

Models

m_id varchar(50)	manuf varchar(50)	model varchar(50)	rated_mAh float	notes varchar(50)
Tenergy AB	Tenergy	Essential	2500	
Amazon B	Amazon	Basics	2200	OEM unknown

For notational convenience in relational algebra, we may abbreviate as follows:

Batteries(tag, model, buy_date, price, manuf_date, color, notes) $B(T, M, B, P, F, C, N)$

Charges(tag, model, charger, date, mAh, method) $C(T, M, C, D, A, E)$

Models(m_id, manuf, model, rated_mAh, notes) $M(I, F, M, A, N)$

Figure 1: A battery database.