

# PGT Exam Performance Feedback

## 2017/2018 Semester 1

**Comments** 53 students took the exam. It proved to be a very hard one, with an average of only 42.2%.

Q1: This turned out to be a not very easy question, yielding an average of 56.8%

a.i/ii - There were many altogether inappropriate answers. With regards to sensible ones, the main frequent error was failing to make explicit whether set or bag semantics was being assumed.

b.i - The most frequent error was a failure to use: an iterator for every referred relation, and a join condition.

b.ii - The most frequent errors were: thinking that a nested query was needed to compute the minimum, failing to write a proper group by clause (which must contain all the attributes in the select clause that are not being aggregated), and instead of writing the predicate on price as a conjunct on the where clause (i.e., on every tuple, as intended) doing it instead in a having clause (where it only applies to the min prices, which was not intended).

Q2: This turned out to be an extremely hard question, yielding an average of only 29.8%

a - There were many errors. The most frequent was to assume that the buffer needs to hold the product (or, even more unjustifiably, the sum) of the cardinalities.

b - Almost no one answered this question as required. The question hinges on the fact that the hash join will block until the left operand is completed.

c - Here, while a fair number of students found the right join order, far fewer reproduced the run of the algorithm, and fewer still did the calculations.

Q3: This turned out to be a close to average-difficulty question, yielding an average of 60.4%

a - No major issues on this part.

b - Almost everyone did very well on this part.

c - This was the hardest part. No systematic errors worth pointing out.

Q4: This, again, turned out to be extremely hard, with an average of only 31.1%

a - There were no systematic errors, just poor recall.

b - Again, very poor recall. One common mistake was not understanding what was meant by 's' and 'p' tables in the question. Many students only showed one of 's' or 'p' rather than both, as required.

Q5: This, again, turned out to be extremely hard, with an average of only 32.6%

a - There were no systematic errors, just poor recall.

b - Same as above.

c - Again, very poor recall. One common mistake was forgetting that with join one must tag the emitted result with the left/right input it came from in the mapper and route the attributes correspondingly. Another mistake was to try to map a join as if it were a group by (i.e., directly for the entire query, rather than as a node in a query plan).