

UG Exam Performance Feedback

First Year

2016/2017 Semester 2

COMP18112 Fundamentals of Distributed Systems

Alvaro Fernandes
Kung-Kiu Lau

Comments Question 1

Question 1 was taken by 170/248 (69.5%) of the students. The average percent mark was ~57%, i.e., about 8% below the planned-for degree of difficulty, with a rather high standard deviation of 22.

In terms of subparts, parts (c.iv), (d) and (e) were the only ones below the overall average for the question.

In part (a), many students forgot to focus on what the question was specifically asking, i.e., an example of a company website pushing to the company's customers a role in the performance of a specific business process of that company. General answers (e.g., e-commerce websites like Amazon) fail to meet the requirements. A good example of a good answer is, as cited in the lecture notes, Dell's (now widely emulated) approach of pushing to customers the process of configuring a product prior to ordering it.

In part (b), the main reason for loss of marks was a failure to explicitly contrast how the in-cloud application delivers transparencies better than the corresponding off-cloud one.

Parts (c.i) and (c.ii) were generally well-answered, though some students, again, failed to read the questions carefully (e.g., failing to concentrate on contrasting message semantics).

Parts (c.iii) and (c.iv) revealed a faulty understanding of the notion of early- and late-binding RMI. These are better understood as static and dynamic, in the sense that with early-binding the desired responder is fixed in the code the requester uses, whereas with late-binding the desired responder is not fixed in the code, instead the requester makes a decision immediately prior to the call (e.g., by consulting a service directory) and writes out the desired responder before the request is sent.

In part (d), the most common cause for loss of marks was forgetting that the question was about what to optimize for, i.e., what is the dominant cost before and after. Moving the system to a distributed setting changes the dominant cost from between-memory data movements to communication costs between components.

In part (e), the most common cause for loss of marks was a failure to take the analogy strictly. So, sending (i.e., writing a message) is blocking because one does not (normally) interleave the writing/pinning with some other action, and likewise, receiving is blocking for analogous reasons. Too many students overcomplicated this in quite unnecessary ways.

Question 2

Question 2 was taken by 147/248 (59.3%) of the students. The average percent mark was ~67%, i.e., a little above the planned-for degree of difficulty, but with a rather high standard deviation of 24.

In terms of subparts, parts (a), (c)(iii) and (f) were the only ones below the overall average for the question.

In part (a), the most common cause for loss of marks was a failure to explicitly link the loss of performance in parallel processing to the potential for contention (e.g., congestion on the interconnect) and overheads (e.g., splitting, spawning and merging too many subtasks).

In part (c)(iii), the most common cause for loss of marks was forgetting that split, spawn and merge occur for *both* mappers *and* reducers.

In part (f), the most common cause for loss of marks was ignoring the fact that in the scenario in hand, the bespoke interconnect was seen as having one single administrator, in which case, change is manageable *and* managed, therefore the axioms lose force and performance becomes more predictable.

Question 3

Overall the performance is satisfactory.

Parts a) and b): Most students show they understand cache and its benefits for DNS and IMAP.

Part c): Most students do not define atomic transaction correctly, although they give reasonably correct examples.

Part d): Most students incorrectly say it is stateful, but correctly explain why stateful is not desirable.

UG Exam Performance Feedback

First Year

2016/2017 Semester 2

Part e): It is not necessary to show all Coffman conditions to hold, just one is enough to make deadlock possible. Some students only explain Coffman conditions in general; they should be applied specifically to this example.

Part f): Most students correctly say it is inconsistent but many do not explain the reason clearly: the key thing is to show a contradiction.

Part g): Most students give the wrong definition for critical section, confusing it with atomic transaction.
