

CMSC 724, Spring 2016: Homework 1

Due Friday February 12, 2016, 6pm.

The homework is to be done by yourself. Be concise in your answers – answers to most questions should be at most one paragraph.

1. Give an example that illustrates that IMS violates the desirable property of *physical data independence*.
2. Briefly explain the notion of “currency indicators” in CODASYL.
3. Briefly explain what problems the *normal forms* (as a concept) were designed to address/solve, and how they solve those problems. What are the key tradeoffs different normal forms allow us explore?
4. What were the problems with Normalization theory that E/R model provides an effective solution for?
5. What was the key motivation behind Object-oriented databases? Why did they not succeed?
6. Briefly explain the notion of UDFs and their key advantages.
7. Briefly explain why I/O processes are typically needed if the database server is a single multi-threaded process. Is there something the OS could provide to address this problem?
8. Explain one scenario where it makes sense to assign multiple DBMS workers to a single SQL query, and discuss the pros and cons of doing so. There is a brief discussion towards the end of Section 2.3 (in the Architecture paper).
9. List three of the responsibilities of the query rewrite module in a database system.
10. Given three reasons why databases need fine-grained control over the buffer management.
11. Briefly explain why “Steal/No Force” is preferable, and what kind of problems it raises for guaranteeing ACID properties.
12. What trade-offs do Isolation Levels allow one to explore?
13. What are the three major concurrency control techniques?
14. Clearly explain the problem of “double-buffering” when running a database system on top of an OS. How does *mmap* help solve that problem?
15. Briefly explain the LRU-k buffer replacement policy, and how it helps circumvent some of the concerns with using the more standard LRU for database operations. There is a brief discussion in Section 5.3 of the “Architecture” paper.
16. What are some of the issues that come up if a database is being used on top of a RAID storage device?