

5. What does ORM stand for? What problem does ORM solve? (4)

6. List four items needed for an ORM to connect to a DBMS. (4)

7. How is lazy fetching applied with respect to ORM? (4)

8. What does it mean for an object to be detached? (4)

9. Define all parts of the acronym ACID. (8)

For questions 10 and 11, use the following schedule:

$$R_1(x)R_2(x)W_1(x)R_2(y)R_3(x)R_3(y)W_2(y)$$

10. For this schedule, draw the serialization graph. Is the schedule serializable? (4)

11. Is the schedule cascadeless? Why or why not? (4)

For questions 12 through 14, consider the following query from the joke rating database we have used in class. For your convenience, a copy of the table description is included at the end of the exam.

```
SELECT username
FROM users u, jokes j
WHERE u.userid = j.author and Content = 'R';
```

12. Assuming there are 100 jokes and the content ratings are 'G', 'PG', 'PG13', 'R' and 'NC17', how many jokes would the DBMS predict to satisfy the predicate Content = 'R'? Why? (6)

13. Assuming there are 400 users, how many usernames would the DBMS predict to be returned by this query? (6)

14. What is the size of the result set (in bytes)? (6)

For questions 15 through 16, consider the set of 2d points of the form (x, x^2) where $0 \leq x \leq 9$.

15. Find the two nearest neighbors to the point $(10, 10)$. (6)

16. Use the k-means algorithm to form 2 clusters of the data. Use the points $(0, 0)$ and $(9, 81)$ as the starting points. (6)

17. Describe collaborative filtering. (6)

18. For association, what is the formula for support and confidence? (4)

19. For classification, how is the best split determined? (4)

20. For information retrieval, what is the formula for recall and precision? (4)

21. For Google, what is page rank? (4)