

Name: _____

Student ID: _____

Signature: _____

MIDTERM EXAM 1, Fall 2006

CS 564 Introduction to Database Management Systems

Department of Computer Science

University of Wisconsin-Madison

Exam Rules:

- 1) Close book and notes, 60 minutes
- 2) Please write down your name and student ID number NOW.
- 3) Please wait until being told to start reading and working on the exam.
- 4) If you think a problem is ambiguous, write down your assumptions, argue that they are reasonable, then work on the problem using those assumptions.

Scores:

Problem 1:	out of	20 points
Problem 2:	out of	15
Problem 3:	out of	15
Problem 4:	out of	10
Problem 5:	out of	10
Problem 6:	out of	15
Problem 7:	out of	15

Total:	out of	100 points
--------	--------	------------

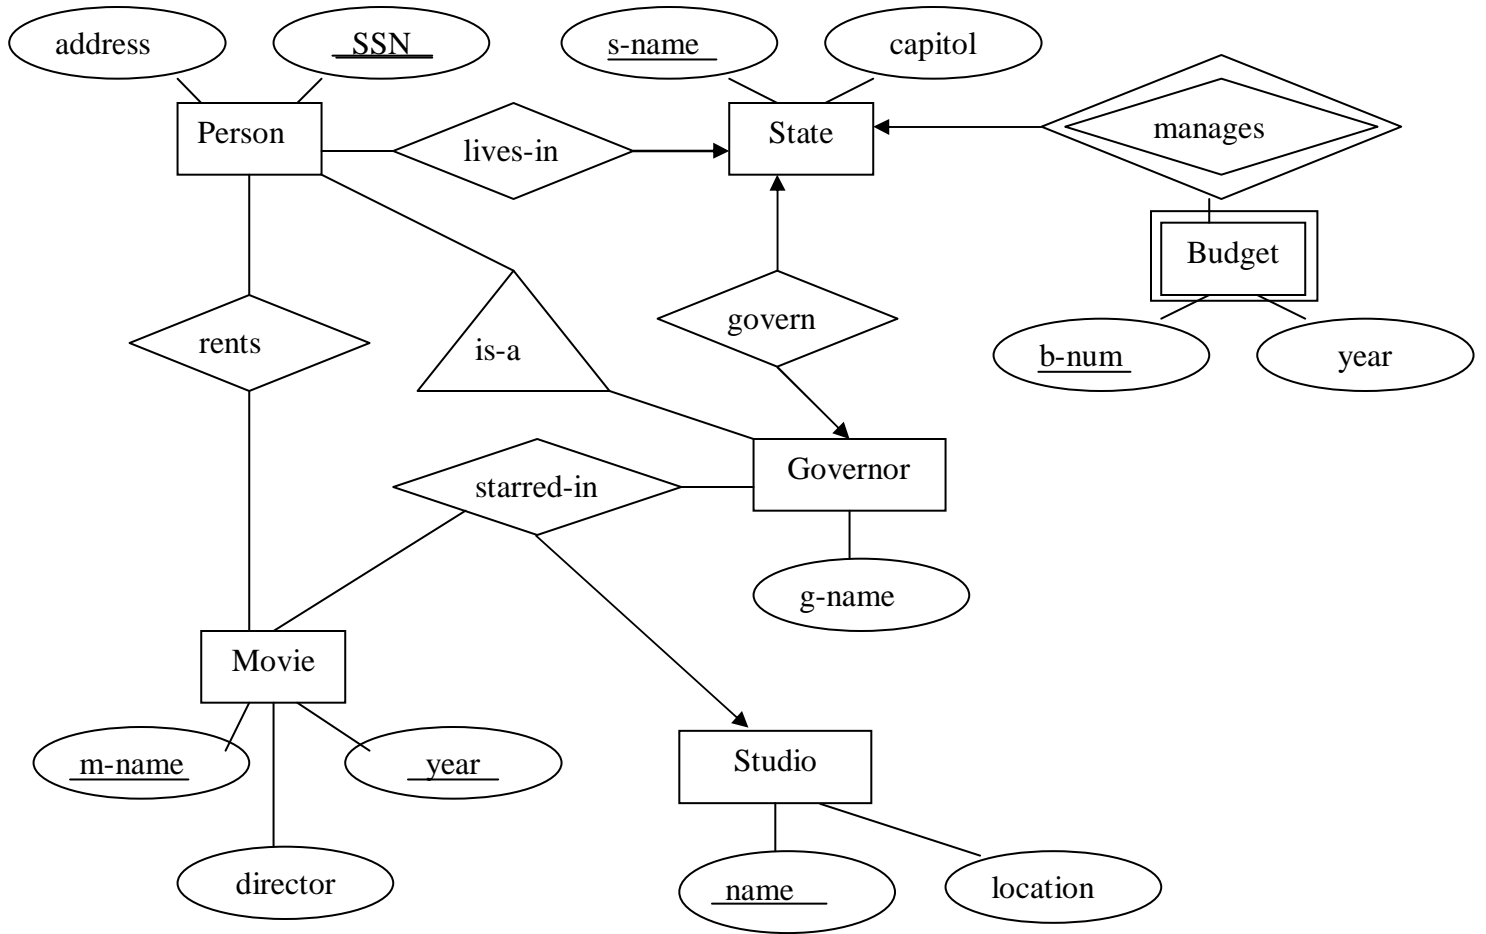
1. (20 points; 2 points each) Multiple choice questions:

For each of the following statements or questions, circle the correct answer.

1. Attributes in the ER model can only have *atomic* values.
(a) true (b) false
2. A multiway ER relationship can always be converted into a set of binary relationships.
(a) true (b) false
3. In the ER model, referential integrity constraint implies that at most one value exists in a given role.
(a) true (b) false
4. Updates to a relational database schema is painful; the main reason is because SQL does not provide good facilities to carry out such updates.
(a) true (b) false
5. Given two entity sets E and F, and a relationship R between them, when translating it is OK to combine the table of R with the table of E, if relationship R is many-one from entity set E to entity set F.
(a) true (b) false
6. If an is-a hierarchy has a total of three entity sets, then translating the entire is-a hierarchy into relational tables using the ER method would result in
(a) 3 tables (b) 4 tables
7. LRU is always a good choice for buffer replacement policy.
(a) true (b) false
8. A relational table can have at most one non-clustered index.
(a) true (b) false
9. When inserting or deleting records from a page, the slot array can never be compacted.
(a) true (b) false
10. If a value in a table is NULL, it means that that value is missing.
(a) true (b) false

2. (15 points) ER and translation to relational model.

Translate this ER diagram into a relational schema. For each relation in your schema, specify the key of that relation. In translating a subclass hierarchy, use the ER style translation.



3. (15 points) Buffer Management:

a) List at least three reasons why the database system does not want to let the OS manage space and buffer.

b) What is sequential flooding? Draw a small example to illustrate the situation.

c) What happens if a page is requested when all pages in the buffer pool are dirty?

4. (10 points) Page Management:

Assume a newly created page P. Consider the following sequence of actions:

Insert record A. Insert record B. Insert record C. Delete record B. Delete record C. Insert record D.

Draw the page **after** the above sequence of action has been executed.

5. (10 points) Record Layout:

Discuss two alternatives for organizing the fields of variable-length records, and the pros and cons of each. What happens if a field of a record grows during a modification and no longer fits into the space remaining on its page?

6. (15 points) Indexes:

Assume a file which has 950 thousands (that is, 950000) records. Assume also that we are indexing this file using a B+ tree. In this particular B+ tree, the average page has an occupancy of 100 pointers (that is, the tree's average branching factor is 100). Assume further that the amount of memory set aside for storing the index is 150 blocks, and that all 950000 records of the above file reside on disk.

Assume no duplicate search key exists, given a search key K , compute the minimal number of disk I/O needed to retrieve the record with that search key. Explain your answer.

7. (15 points) B+ Tree:

Consider the B+ tree below, where each node (except the root) must contain between 2 and 4 entries. Draw the tree after the value 16 has been inserted into the tree.

