

Introduction to CS 464/564 Advanced Database Systems

Outline

- Course Outline
- CS 364 Review

Course Outline

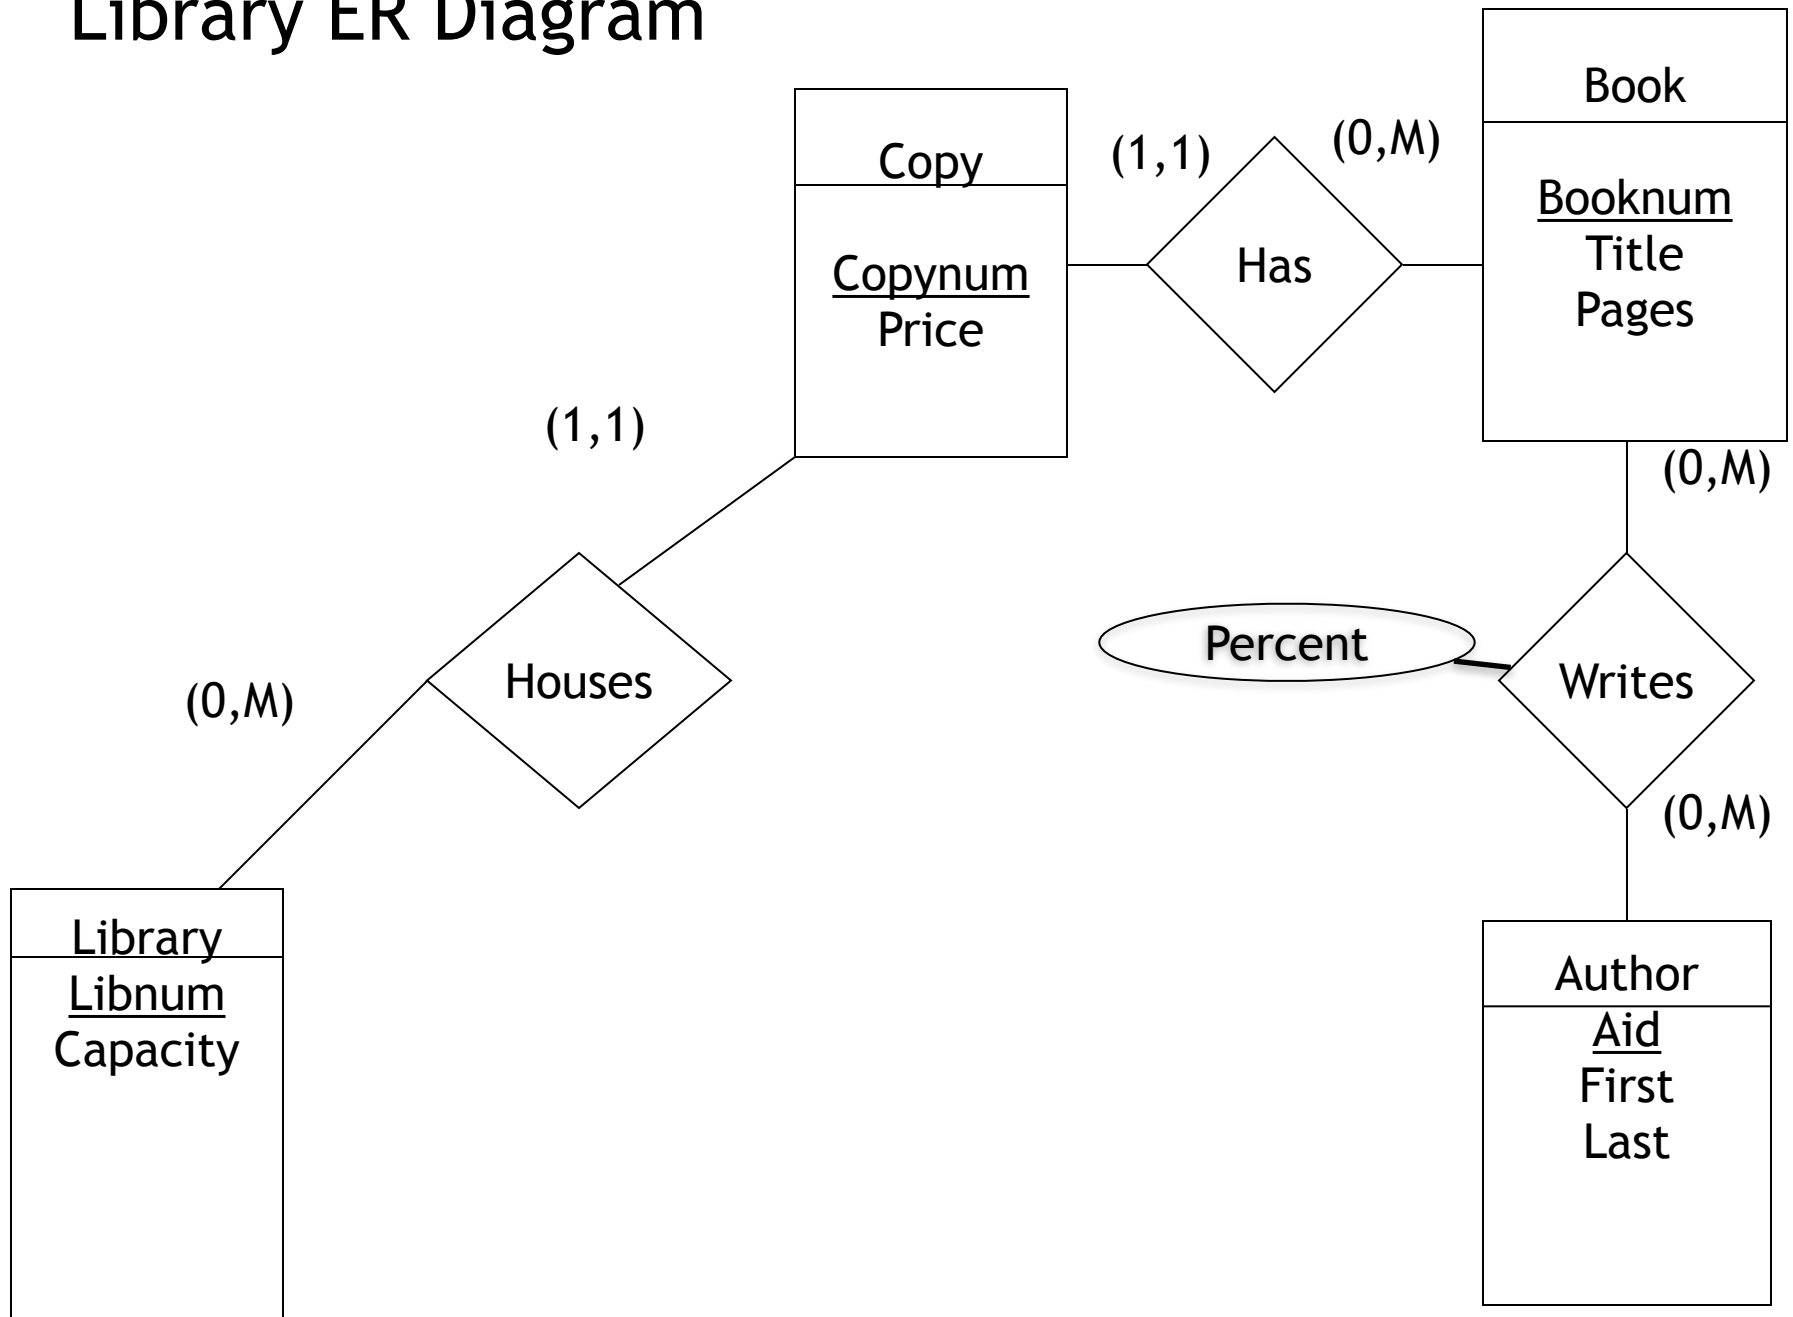
- CS 364 Review
- Postgres
- Relational Database Topics Beyond CS 364 Content
 - Stored Procedures
 - Relation Algebra and Relational Calculus
 - Normalization
 - Query Processing
 - Transaction Processing
- Other Data Models
 - NoSQL Databases
 - MongoDB
 - Neo4J
- Other Issues with Database Systems

CS 364 Review

Problem Statement

- Create a database to store information about authors, books, copies of books and libraries. Associated with each author is a unique author aid, a first name and a last name. Associated with each book is a unique book number, a title and the number of pages in the book. Authors can write multiple books and a book can have multiple authors. For each author of a book the percent of the book the written by the author is recorded in the database. Associated with each copy is a unique copy number and a price. Each copy is a copy of exactly one book. Associated with each library is a unique libnum and a capacity. The capacity is the number of copies that can be stored in the library. Each copy is stored in exactly one library.

Library ER Diagram



Problems

- You should be able to implement the library database
 - You can use any relational database software you want.
 - Write create tables statements for the ER diagram on the previous slide.
 - Assume the data type for the aid, booknum, copynum, libnum, pages, capacity and percent is integer. The data type for price is decimal(8,2). The data type for title, first and last is TEXT or varchar.
 - Load data from the files author.txt, book.txt, writes.txt, library.txt and copies.txt into the database,

Query Problems Problems

- Find the number of pages in *War and Peace*
- Find the titles of books written by Anton Chekhov
- Find the price of each copy of *Middlemarch*
- Find the names of the authors of books with more than 600 pages.
- Find the number of books written by each author.
- Find the average copy price for each book.
- Find the library with the largest capacity.
- Find the unused capacity for each library.
- Find the libraries with copies of books by every author
- Find the libraries with copies of books by every author that has at least one book in the database

SQL Keywords

- create table
 - constraints
 - primary key
 - unique
 - foreign key
 - not null
 - check clause
- alter table
 - add column
 - drop column
- drop table
- create view
- drop view
- select
 - distinct
 - aggregate functions
 - count
 - sum
 - avg
 - min
 - max
 - natural join
 - join on
 - left join
 - right join
 - full join

SQL Keywords

- Select
 - in
 - exists
 - and
 - or
 - not
 - all
 - group by
 - having
 - order by
- Intersect
- Union
- Except
- update
- delete
- insert
- grant
- revoke