CSc 460 — Database Design
Spring 2026 (McCann)

http://www2.cs.arizona.edu/classes/cs460/spring26/

Homework #4

(115 points)

Due Date: April 9th, 2026, at the beginning of class

Directions

1. **This is an INDIVIDUAL assignment; do your own work! Submitting answers created by other people or generated by AIs (e.g., ChatGPT) are a few examples of NOT doing your own work.**
2. Write complete answers to each of the following questions, in accordance with the given directions. Create your solutions as a PDF document such that each question is on a separate page; all parts of a multi–part question may be on the same page. Show your work, when appropriate, for possible partial credit.
3. If you have questions about any aspect of this assignment, help is available from the class staff via piazza.com and our office hours.
4. When your answers are ready to be turned in, do so on gradescope.com. Be sure to assign pages to problems after you upload your PDF. Need help? Visit https://help.gradescope.com/ and search for "Submitting an Assignment."
5. Remember that you can use at most one late day on a homework assignment, because we will be distributing solutions after that time.
6. Start early! Getting help is much easier n days before the due date/time than it will be n hours before.

Connolly/Begg Chapter 5: Relational Algebra and Relational Calculus

1. (10 points) In class we covered several varieties of join — the inner joins (theta, equi, and natural), and the outer (left, right, and full).
   (a) What distinguishes inner joins from outer joins?
   (b) How are the three types of inner join different?
   (c) How are the three types of outer join different?
2. (10 points) Question 5.10 on page 140 of our Connolly/Begg text provides a few tuple relational calculus queries. Rewrite queries (a) and (d) in relational algebra, using the σ, π, etc., notation and syntax. (The queries are based on the Hotel schema on page 118.)
3. (10 points) Again using the Hotel schema from page 118, write relational algebra queries (again using σ, π, etc.) to answer the following questions.
   (a) List the names and cities of all of the hotels.
   (b) List the guest names and addresses of all of the guests currently staying at the Relational Hotel. ("Currently" means that today's date (`CURRENT_DATE`) is within the guest's dateFrom and dateTo range.)
4. (10 points) Using the Library schema on page 141, write relational algebra queries (yes, still using σ, π, etc.) to answer the following questions.
   (a) Which borrowers have the book "The ABCs of DBs" checked out today? (Use a similar range idea as above.)
   (b) Which borrowers have books that are more than a week overdue?

(Continued . . .)

Connolly/Begg Chapter 6: SQL: Data Manipulation

5. (5 points) Why does SQL allow us to express conditions in both the WHERE clause and in the HAVING clause? That is, why have HAVING at all?

6. (15 points) Using the Hotel schema from page 118, write SQL queries to answer the following questions.

   (a) What are the names and addresses of all guests from Tucson in the Guest table, in ascending lexicographic (alphabetical) order by guest name?

   (b) How many unique guests hold bookings for April? (See examples 6.14 (p. 160) and 6.35 (p. 178) for queries that use dates. You may use comparison operators (`=`, `<=`, etc.) with dates.)

   (c) What are the types, prices and (for the currently occupied rooms) guestNos of all of the rooms at the Relational Hotel? (As above, use (`CURRENT_DATE`) for today's date.)

7. (10 points) Using the Library schema on page 141, write SQL queries to answer the following questions.

   (a) List the copyNo values of the copies of "The ABCs of DBs" that are not currently checked out.

   (b) The library has multiple copies of popular books. Which library patrons have, at any time, checked out copies of books for which the library has multiple copies?

Connolly/Begg Chapter 7: SQL: Data Definition

8. (10 points) Write commands to create the following items based on the Projects schema defined on page 140 of the text.

   (a) Create (that is, write the CREATE TABLE command for) the Employee relation with the additional constraints that 'sex' is limited to the single characters 'M', 'F', and '-', and 'position' is limited to 'dept mgr', 'team mgr', 'associate', and 'intern'. Studying examples provided in Chapter 7 may be helpful.

   (b) Create a view that shows the full names of the employees and the names of the projects on which they are working. Give the view more meaningful field names than those of the base tables.

Connolly/Begg Chapter 8: Advanced SQL

9. (5 points) Give two advantages and two disadvantages of the use of triggers in SQL.

Connolly/Begg Chapter 14: Normalization

10. (10 points) The parts of this question refer to this table of dentist and patient information:

staffNo	dentistName	patNo	patName	apptdate	appttime	procedureNo
S1011	Tony Smith	P100	Gillian White	2026-04-06	10:00	S15
S1011	Tony Smith	P105	Jill Bell	2026-04-06	11:00	S15
S1024	Helen Pearson	P108	Ian MacKay	2026-04-06	10:00	S10
S1024	Helen Pearson	P108	Ian MacKay	2026-04-08	14:00	S10
S1032	Robin Plevin	P105	Jill Bell	2026-04-08	16:30	S15
S1032	Robin Plevin	P110	John Walker	2026-04-09	18:00	S13

   (a) List the functional dependencies that are present in this table. When creating your list, which assumptions, if any, did you make about the attributes and data values?

   (b) When the data in the table is changed, update anomalies may occur. Provide an example of a deletion anomaly and an example of an update anomaly that could occur on this table.

(Continued . . .)

Questions on Topics Not Well–Covered in the Text

11. (5 points) Consider our SPJ schema from the handout I gave you several weeks ago (and available from the class web page), and the following view definition:

```
create view suppliedparts ("Part Number")
as select distinct p.pno
     from p, spj
    where p.pno = spj.pno;
```

    Imagine that a DBMS is storing this as a materialized view. Which specific kinds of changes to the base relations (p and spj) would require the view to be re–materialized, and why?

12. (5 points) Consider this schema: $R(A, B, C, D, E, F, G, H, I)$. Also consider the FDs $AE \rightarrow F, B \rightarrow A, BF \rightarrow G, BH \rightarrow I, DE \rightarrow C$. Compute both BDE^+ and $ABCD^+$.

13. (10 points) Consider the schema $S(J, K, L, M, N)$ and the FDs $J \rightarrow L$ and $KL \rightarrow N$. Using only Armstrong's Axioms, prove that the FD $JKM \rightarrow JN$ holds.