

CSC553: Homework 2

Due: April 24th, 2022

This assignment is on relational algebra, logical plan equivalence, and indexes.

1 Relational Algebra

The following questions pertain to the following database specification:

Plays(Play#, Title, Director, Year, Cost)

Artists(Artist#, LName, FName, Gender, Birthdate)

Roles(Play#, Artist#, Character)

Write relational algebra expressions for the following English queries:

1. Find the directors of the play titled “Romeo and Juliet”.
2. Find the titles of all plays in which the director is also an artist.
3. Find the artists who have played two characters in the same play; show the title of each such play, first name and last name of the artist and the two characters.
4. Find the titles of the plays in which the artists are all of the same gender.
5. Find the play with the least cost.

2 RA Equivalence

Give examples to show that: (i) Projection cannot be pushed below set union.

(ii) Duplicate elimination cannot be pushed below projection.

(iii) Starting with an expression $\Pi_L(R(a, b, c) \bowtie S(b, c, d, e))$, push projection down as far as it can go if L is (a) $b + c \rightarrow x, c + d \rightarrow y$. (b) $a, b, a + d \rightarrow z$.

3 B+-trees

3.1

What are the minimum number of keys and pointers in B+-tree (i) interior node, and (ii) leaves, when:

- (i) $d = 10$, (ii) $d = 11$.

3.2

Give the state of the B+-tree as shown below: (i) Perform the following inserts:




Figure 1: Btree

1, 14, 15, 16. For each insert draw the final tree only.

(ii) Perform the following deletes: 23, 23 and higher. For each delete draw the final tree only.

(iii) Consider the original tree again. How many pages will be read from disk to answer each of the following queries. Assume that the buffer pool is empty before each query executes, and that it is large enough to hold the entire database. Assume also that 2 tuples from the base relation fit on one page in the data file and that all pages are full:

- Look up all tuples with search key value of 41 if the index is clustered.
- Look up all tuples with search key value of 41 if the index is unclustered.
- Look up all tuples with search key value in the range $[25, 45]$ if the index is clustered.
- Look up all tuples with search key value in the range $[25, 45]$ if the index is unclustered. Assume the worst-case scenario in terms of how the keys are stored in pages.