

CSC553: Homework 4

Due: May 19th, 2022

This assignment is on query optimization and conflict-serializability.

1 Query Optimization

Consider a join of relation R and relation S . Suppose that schema for R is (Name, YearOfBirth, Profession) and schema for relation S is (Year, Month, Week, Description). Also, $N_R = 1000$ and $V(R, YearOfBirth) = 42$ and $N_S = 500$ and $V(S, Year) = 10$.

(i) What are the two different estimates for the size of join result of $R \bowtie_{YearOfBirth=Year} S$?

(ii) Which one is more accurate?

2 Histograms

Recall that an equi-width histogram splits the value range into X equal ranges and fills in each bucket with a count of values within each particular range. An equi-height histogram adjusts the bucket sizes in such a way that every bucket contains the exact same number of values.

Given the following data:

[1, 2, 5, 6, 8, 11, 18, 26, 34, 36, 37, 39, 43, 50, 61, 62, 66, 67, 70]

(i) Construct an equi-width histogram (with 3 buckets).

(ii) Construct an equi-height histogram (also with 3 buckets).

3 Histograms

Consider the following histogram that represents Hours column

1. What is the answer to `SELECT AVG(Hours) FROM SleepTable?`

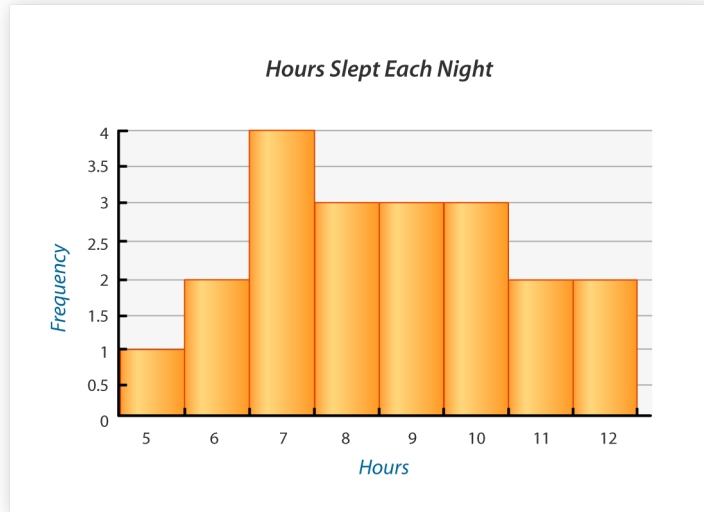


Figure 1: Histogram

2. What is the answer to `SELECT COUNT(Hours) FROM SleepTable`?
3. What is the answer to `SELECT COUNT(*) WHERE Hours = 6`?
4. What is the answer to `SELECT COUNT(*) WHERE Hours BETWEEN 6 and 11`? (`BETWEEN` is inclusive)

4 Conflict-Serializability

For each of the following schedules:

1. $R_1(A) R_2(A) R_3(B) W_1(A) R_2(C) R_2(B) W_2(B) W_1(C)$
2. $R_1(A) W_1(B) R_2(B) W_2(C) R_3(C) W_3(A)$
3. $W_3(A) R_1(A) W_1(B) R_2(B) W_2(C) R_3(C)$
4. $R_1(A) R_2(A) W_1(B) W_2(B) R_1(B) R_2(B) W_2(C) W_1(D)$
5. $R_1(A) R_2(A) R_1(B) R_2(B) R_3(A) R_4(B) W_1(A) W_2(B)$

Answer the following:

- (i) What is the precedence graph for the schedule?
- (ii) Is the schedule conflict-serializable? If so, what are the equivalent serial schedules?

(ii) Are there any serial schedules that must be equivalent but are not conflict-equivalent?