

Sai University

Data Base Management Systems

Mid-Semester Examination, Semester A, 2023–2024

Date : 4 October, 2023

Marks : 20

Duration : 90 minutes

Weightage : 20%

- Let r be a relation with attribute set R and s be a relation with attribute set S such that $S \subseteq R$. We can define a division operation in the relational algebra as follows: $r \div s$ is the largest relation q such that $q \times s \subseteq r$.
 - Let $r(A, B, X, Y)$ and $s(X, Y)$ be the relations displayed below. Compute $r \div s$.

r			
A	B	X	Y
a	b	w	x
c	d	y	z
a	b	y	z
e	f	w	x
c	d	w	x

s	
X	Y
y	z
w	x

(2 marks)

- Show that

$$r \div s = \Pi_{R-S}(r) - \Pi_{R-S}((\Pi_{R-S}(r) \times s) - \Pi_{R-S,S}(r))$$

where $A - B$ denotes set difference (the set of elements in A that are not in B).

(4 marks)

- An election database includes the following two tables:

$Candidates(CandidateID, CandidateName)$

$Votes(CandidateId, PollingBoothID, NumberOfVotes)$

The second table is updated booth by booth. Whenever counting is completed in a polling booth, an entry is added to the second table for each candidate who has got a nonzero number of votes, listing out the number of votes he got at that booth. For candidates who poll zero votes at a booth, no entry is recorded in $Votes$.

We need to periodically compute the current status of *all* candidates by “joining” these two tables to create a table of the form

$Status(CandidateID, CandidateName, TotalVotesTillNow)$

- Explain why a natural join is not adequate for this purpose. (2 marks)
- What kind of join should we use to ensure that *every* candidate is included in the status table, including those who have yet to poll any votes across all booths? (2 marks)
- Write an SQL query to compute the $Status$ table. (4 marks)

3. A database of movies currently showing in the city has a table of the form *Playing*(*Title*, *Theatre*, *Location*), where *Title* is the name of the movie, *Theatre* is the name of the cinema theatre and *Location* is the place where the theatre is located (a mall or multiplex).

We assume that no pair of movies currently running have the same title and no two cinema theatres have the same name. A cinema theatre may be showing different movies in different time slots, but no location has two cinema theatres showing the same movie.

- (a) What functional dependencies can you infer from these constraints? (2 marks)
- (b) Compute a BCNF decomposition of the table *Playing*. (2 marks)
- (c) Explain whether your BCNF decomposition is dependency preserving. (2 marks)
-