

6 Databases (TGG)

We assume that for each base table R in a relational database we have two update operations : $\text{insert}(R, t)$ which inserts tuple t into table R if t does not violate any of the constraints declared for R (fails otherwise), and $\text{delete}(R, p)$ which deletes all records in R satisfying predicate p (and fails if this would violate referential integrity constraints). Update operations are combined in programs to define transactions with ACID guarantees.

Suppose that we have defined a view $V = Q(R_1, R_2, \dots, R_n)$, where the R_i indicate the base tables used in query Q . The designers of a new database system want to allow users to update directly such a view. That is, if we have an update of the form $U = \text{insert}(V, t)$ or $U = \text{delete}(V, p)$, then the database system must automatically generate a transaction T_U over the tables R_i such that for any database instance DB this diagram commutes:

$$\begin{array}{ccc} DB & \xrightarrow{T_U} & DB' \\ \downarrow Q & & \downarrow Q \\ V & \xrightarrow{U} & V' \end{array}$$

In other words, applying the update U directly to a view (as if it were a base table) produces the same result as applying T_U to the database and then evaluating the view query.

A major problem with this approach is that there may be multiple distinct solutions for T_U . We explore this now.

- (a) Explain the difference between a *database query* and a *database view*. [2 marks]
- (b) Let $V = \pi_X(R)$ be a view for some base table R and some subset X of R 's attributes Y . How could this be translated into the desired transaction T_U ? Discuss any problems with ambiguity that may arise. [5 marks]
- (c) Let $V = \sigma_q(R)$ be a view for some base table R and predicate q . How could this be translated into the desired transaction T_U ? Discuss any problems with ambiguity that may arise. [5 marks]
- (d) In the design of a database schema it was discovered that a relation R violated Boyce-Codd normal form, and so it was replaced by two base tables R_1 and R_2 resulting from the standard decomposition process. Suppose users attempt to reconstruct the original relation using the view $V = R_1 \bowtie R_2$. Discuss the problems that might arise now in the construction of transaction T_U for updates to V . [8 marks]