CHAPTER 2

Relational Model

Solutions to Practice Exercises

2.1 **a.** $\Pi_{person\_name}\ ((employee\ \bowtie\ manages)$
$\bowtie_{(manager\_name = employee2.person\_name \wedge employee.street = employee2.street}$
$\wedge\ employee.city = employee2.city)(\rho_{employee2}\ (employee)))$

b. The following solutions assume that all people work for exactly one company. If one allows people to appear in the database (e.g. in *employee*) but not appear in *works*, the problem is more complicated. We give solutions for this more realistic case later.

$\Pi_{person\_name}\ (\sigma_{company\_name \neq}$ "First Bank Corporation"$(works))$

If people may not work for any company:

$\Pi_{person\_name}(employee)\ -\ \Pi_{person\_name}$
$(\sigma_{(company\_name =}$ "First Bank Corporation"$_)(works))$

c. $\Pi_{person\_name}\ (works)\ -\ (\Pi_{works.person\_name}\ (works$
$\bowtie_{(works.salary \leq works2.salary \wedge works2.company\_name =}$ "Small Bank Corporation"$_)$
$\rho_{works2}(works)))$

2.2 **a.** The left outer theta join of $r(R)$ and $s(S)$ $(r\ \mathbin{⟕}_\theta\ s)$ can be defined as
$(r \bowtie_\theta s) \cup ((r - \Pi_R(r \bowtie_\theta s)) \times (null, null, \ldots, null))$
The tuple of nulls is of size equal to the number of attributes in S.

b. The right outer theta join of $r(R)$ and $s(S)$ $(r\ ⟖_\theta\ s)$ can be defined as
$(r \bowtie_\theta s) \cup ((null, null, \ldots, null) \times (s - \Pi_S(r \bowtie_\theta s)))$
The tuple of nulls is of size equal to the number of attributes in R.

c. The full outer theta join of $r(R)$ and $s(S)$ $(r \mathbin{⟗}_\theta s)$ can be defined as
$(r \bowtie_\theta s) \cup ((null, null, \ldots, null) \times (s - \Pi_S(r \bowtie_\theta s))) \cup$
$((r - \Pi_R(r \bowtie_\theta s)) \times (null, null, \ldots, null))$
The first tuple of nulls is of size equal to the number of attributes in R, and the second one is of size equal to the number of attributes in S.

2.3 **a.**
$$employee \leftarrow \Pi_{person\_name, street, ''Newtown''}(\sigma_{person\_name=\text{“Jones”}}(employee)) \cup (employee - \sigma_{person\_name=\text{“Jones”}}(employee))$$

b. The update syntax allows reference to a single relation only. Since this update requires access to both the relation to be updated (*works*) and the *manages* relation, we must use several steps. First we identify the tuples of *works* to be updated and store them in a temporary relation (t_1). Then we create a temporary relation containing the new tuples (t_2). Finally, we delete the tuples in t_1, from *works* and insert the tuples of t_2.

$$t_1 \leftarrow \Pi_{works.person\_name, company\_name, salary}(\sigma_{works.person\_name = manager\_name}(works \times manages))$$
$$t_2 \leftarrow \Pi_{person\_name, company\_name, 1.1 * salary}(t_1)$$
$$works \leftarrow (works - t_1) \cup t_2$$