

Introduction to Databases

Exam of 29/09/2023

With Solutions

Diego Calvanese

Bachelor in Computer Science

Faculty of Computer Science

Free University of Bozen-Bolzano

<http://www.inf.unibz.it/~calvanese/teaching/exams/idb/>

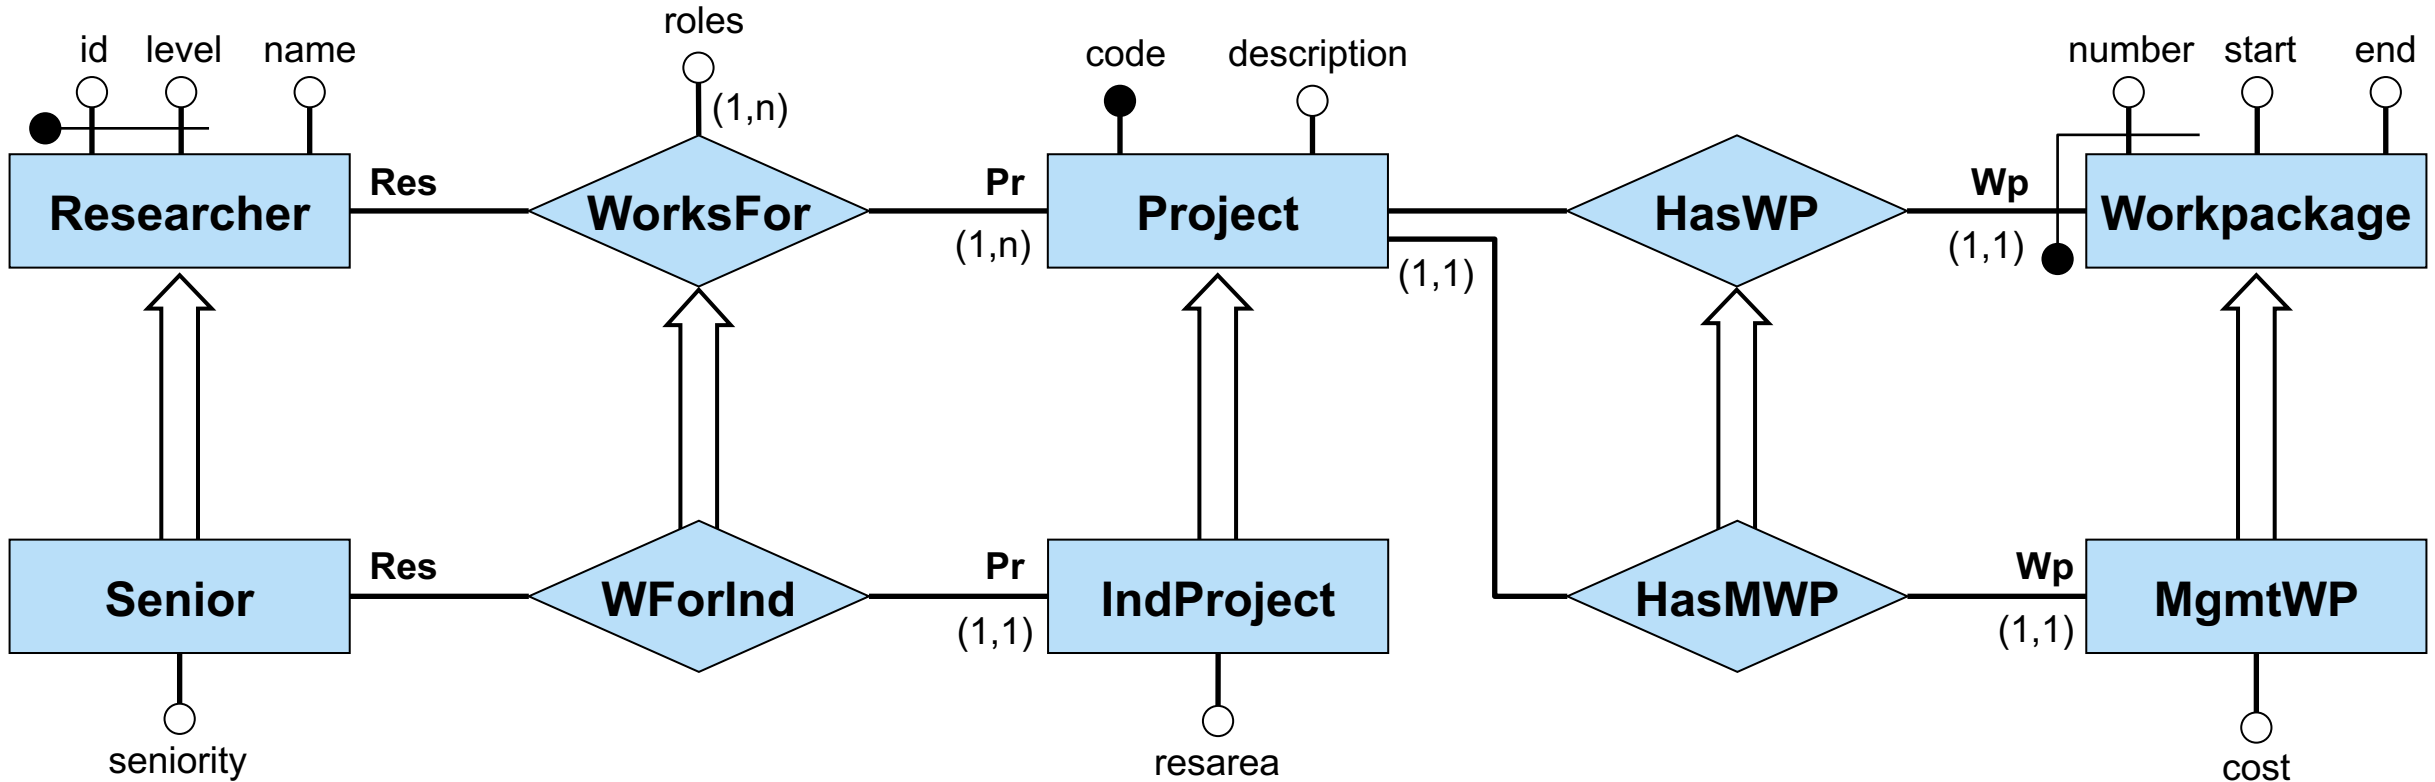
Problem 1

Design the Entity-Relationship schema of an application related to the management of research projects. Of each **researcher**, we are interested in the id and level (that together identify the researcher), the name, and the projects for which they work, with the roles (at least one) they hold in each project (e.g., responsible for publicity, budget responsible, etc.). Some researchers are **seniors** and for them we are also interested in their seniority. Of each **project**, we are interested in the code (identifier), the description (a string), and the researchers (at least one) that work for it. Some projects are **individual projects** and for them we are also interested in the research area. Note that only a single researcher participates to an individual projects, and such researcher must be senior. For each project, we are also interested in the **workpackages** that belong to it. Of each workpackage, we are interested in its number, which is unique within the project to which the workpackage belongs, and in its start and end date. In the case where the workpackage is a **management workpackage**, we are also interested in its cost. Note that each project has exactly one management workpackage, and the start and end dates of all workpackages of a project must fall within the start and end dates of the management workpackage of that project.

Problem 1: Conceptual schema

For each instance I of the schema:

1. If $w \in \text{instances}(I, \text{Workpackage})$, $(w, sw) \in \text{instances}(I, \text{start})$, and $(w, ew) \in \text{instances}(I, \text{end})$, then $sw \leq ew$.
2. If $(\text{Project}:p, \text{Wp}:w) \in \text{instances}(I, \text{HasWP})$, $(\text{Project}:p, \text{Wp}:m) \in \text{instances}(I, \text{HasMWP})$, $(w, sw) \in \text{instances}(I, \text{start})$, $(w, ew) \in \text{instances}(I, \text{end})$, $(m, sm) \in \text{instances}(I, \text{start})$, and $(m, em) \in \text{instances}(I, \text{end})$, then $sm \leq sw$ and $ew \leq em$.
3. If $ip \in \text{instances}(I, \text{IndProject})$ and $(\text{Res}:r, \text{Pr}:ip) \in \text{instances}(I, \text{WorksFor})$, then $(\text{Res}:r, \text{Pr}:ip) \in \text{instances}(I, \text{WForInd})$.



Problem 2

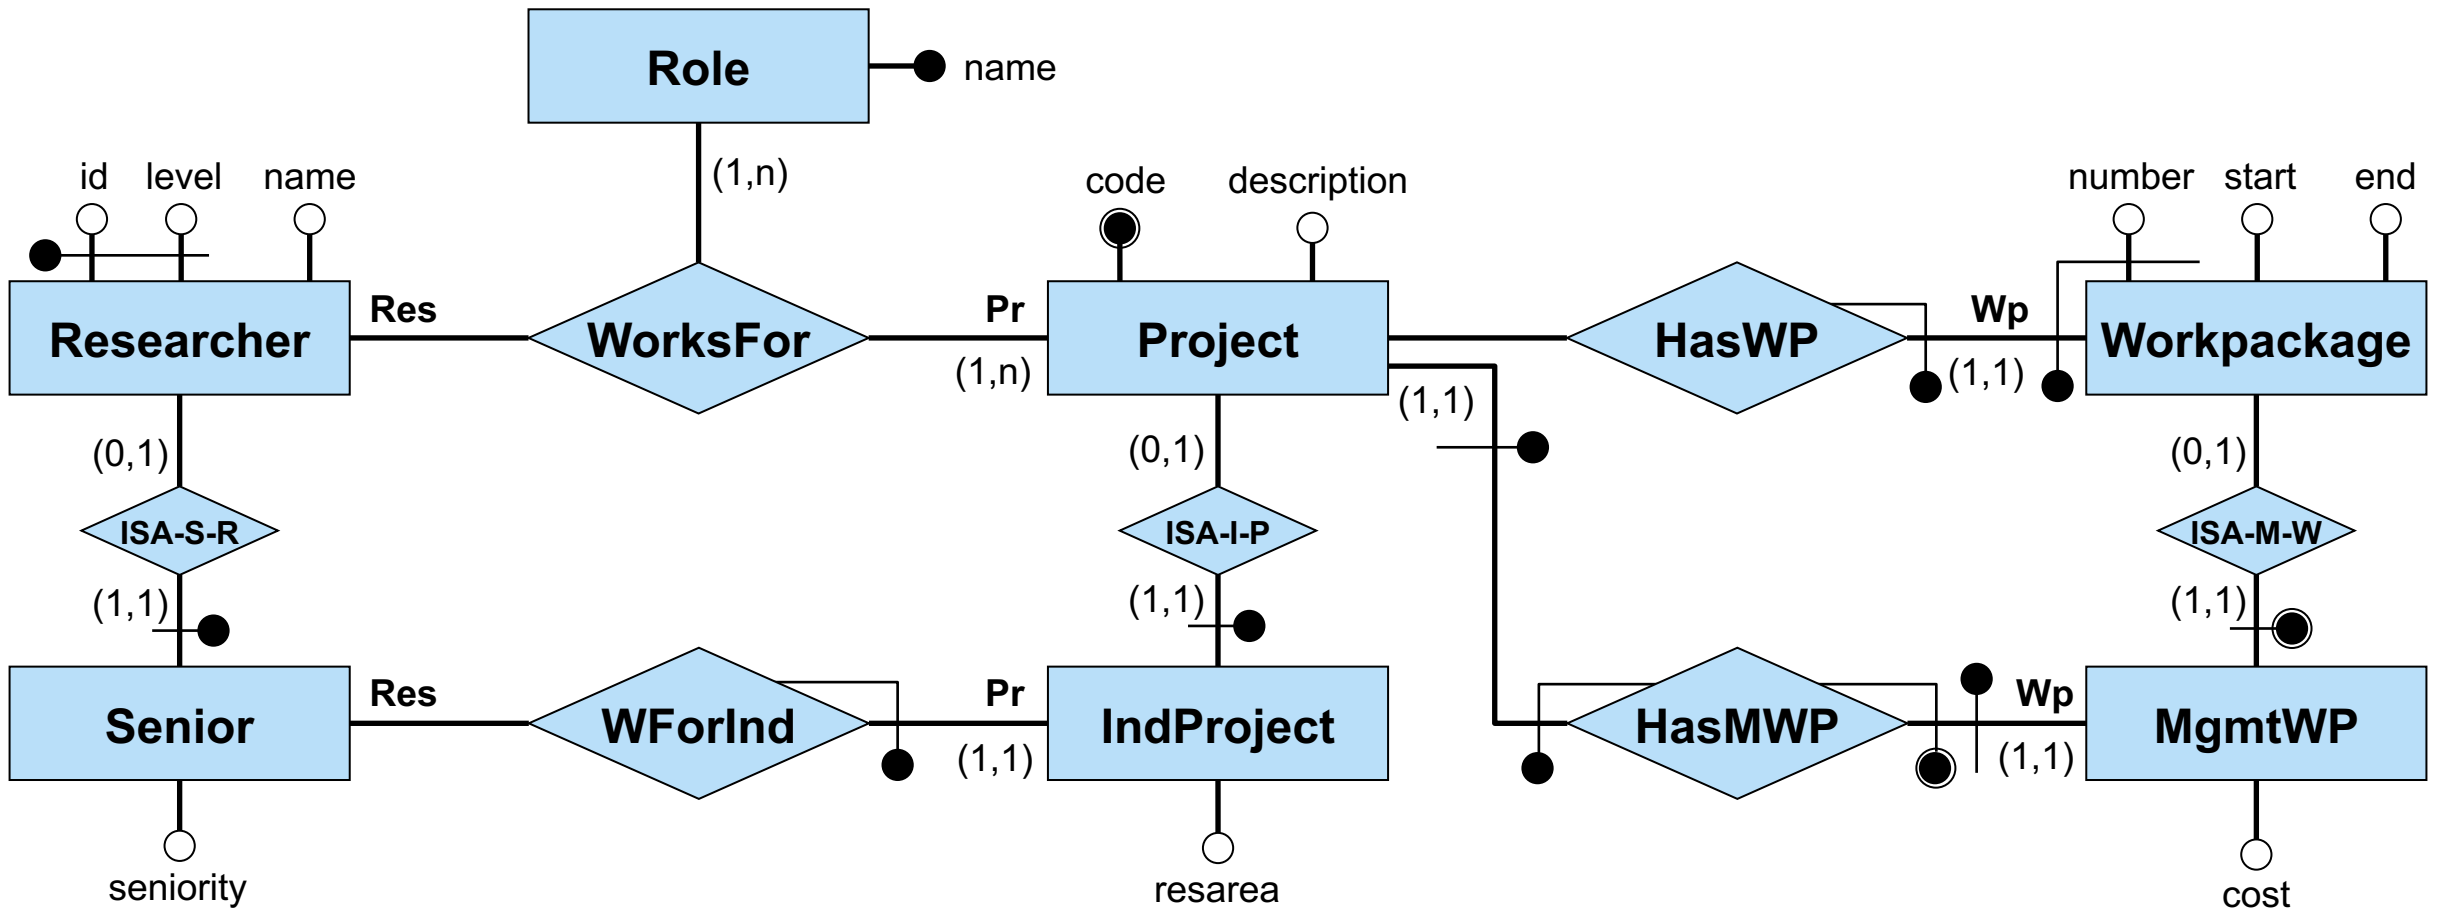
Carry out the logical design of the database, producing the complete relational schema with constraints, taking into account that when we access the information about a project, we always want to know whether it is an individual project, and if it is so, we always want to know the research area and the senior researcher that works for the project.

In your design you should follow the methodology adopted in the course, and you should produce:

- the restructured ER schema (possibly with external constraints),
- the direct translation into the relational model (possibly with external constraints), and
- the restructured relational schema (again with constraints).

You should motivate explicitly how the above indications affect your design.

Problem 2: Restructured conceptual schema



Problem 2: Restructured conceptual schema – External constraints

For each instance I of the schema:

1. If $w \in instances(I, \mathbf{Workpackage})$, $(w, sw) \in instances(I, \mathbf{start})$, and $(w, ew) \in instances(I, \mathbf{end})$, then $sw \leq ew$.
2. If $(Project:p, Wp:w) \in instances(I, \mathbf{HasWP})$, $(Project:p, Wp:m) \in instances(I, \mathbf{HasMWP})$, $(w, sw) \in instances(I, \mathbf{start})$, $(w, ew) \in instances(I, \mathbf{end})$, $(m, sm) \in instances(I, \mathbf{start})$, and $(m, em) \in instances(I, \mathbf{end})$, then $sm \leq sw$ and $ew \leq em$.
3. Constraint resulting from the elimination of ISA between relationships **WForInd** and **WorksFor**, and also capturing constraint 3 of the original schema:
If $(Res:s, Pr:ip) \in instances(I, \mathbf{WForInd})$ and p is the (unique) instance of **Project** such that $(IndProject:ip, Project:p) \in instances(I, \mathbf{ISA-I-P})$, then there is a unique instance r of **Researcher** such that $(Res:r, Role:ro, Pr:p) \in instances(I, \mathbf{WorksFor})$ for some instance ro of **Role**, and moreover $(Senior:s, Researcher:r) \in instances(I, \mathbf{ISA-S-R})$.
4. Constraint resulting from the elimination of ISA between relationships **HasMWP** and **HasWP**:
If $(Project:p, Wp:m) \in instances(I, \mathbf{HasMWP})$ and w is the (unique) instance of **Workpackage** such that $(MgmtWP:m, Workpackage:w) \in instances(I, \mathbf{ISA-M-W})$, then $(Project:p, Wp:w) \in instances(I, \mathbf{HasWP})$.

Problem 2: Direct translation (1/2)

Researcher(id, level, name)

Senior(id, level, seniority)

foreign key: Senior[id,level] \subseteq Researcher[id,level]

Project(code, description)

inclusion: Project[code] \subseteq WorksFor[project]

foreign key: Project[code] \subseteq HasMWP[project]

IndProject(code, resarea)

foreign key: IndProject[code] \subseteq Project[code]

foreign key: IndProject[code] \subseteq WForInd[project]

Role(name)

inclusion: Role[name] \subseteq WorksFor[role]

WorksFor(resid, reslevel, role, project)

foreign key: WorksFor[resid,reslevel] \subseteq Researcher[id,level]

foreign key: WorksFor[role] \subseteq Role[name]

foreign key: WorksFor[project] \subseteq Project[code]

Problem 2: Direct translation (2/2)

WForInd(resid, reslevel, project)

foreign key: WForInd[resid,reslevel] \subseteq Senior[id,level]

foreign key: WForInd[project] \subseteq IndProject[code]

foreign key: WForInd[resid,reslevel,project] \subseteq WorksFor[resid,reslevel,project]

Workpackage(number, project, start, end)

foreign key: Workpackage[project] \subseteq Project[code]

constraint: start \leq end

MgmtWP(number, project, cost)

foreign key: MgmtWP[number,project] \subseteq Workpackage[number,project]

foreign key: MgmtWP[project] \subseteq HasMWP[project]

HasMWP(project, project2)

foreign key: HasMWP[project] \subseteq MgmtWP[project]

foreign key: HasMWP[project2] \subseteq Project[code]

constraint: project = project2

Note that the fact that the primary key of **MgmtWP** is just the **project** attribute is a consequence of the fact that each project has a unique management workpackage.

Also, the constraint **project = project2** holding for **HasMWP** is a consequence of the constraint deriving from the inclusion between relationships **HasMWP** in **HasWP** in the original ER schema and the fact that **project** is (part of) the primary key of **MgmtWP** and of **Workpackage**.

Problem 2: Direct translation – External constraints

1. Has been specified as **tuple constraint** of the **Workpackage** relation.
2. **CHECK constraints** on relation **Workpackage**:
 - `start >= (SELECT WP.start FROM Workpackage WP, MgmtWP M WHERE WP.number = M.number AND WP.project = M.project)`
 - `end <= (SELECT WP.end FROM Workpackage WP, MgmtWP M WHERE WP.number = M.number AND WP.project = M.project)`
3. The constraint resulting from the elimination of ISA between relationships **WForInd** and **WorksFor** has already been expressed as a foreign key on **WForInd**. The additional constraint 3 of the original conceptual schema can be expressed through a **CHECK constraint**:

```
CHECK (SELECT WF.resid, WF.reslevel, WF.project
        FROM WorksFor WF, IndProject IP WHERE WF.project = IP.code)
        IN (SELECT resid, reslevel, project FROM WForInd)
```
4. The constraint resulting from the elimination of ISA between relationships **HasMWP** and **HasWP** has already been expressed through the **foreign key** between **MgmtWP** and **Workpackage**, together with the fact that the project code is part of the primary key of **Workpackage**.

Problem 2: Restructuring of the relational schema

When we access the information about a project, we always want to know whether it is an individual project, and if it is so, we always want to know the research area and the senior researcher that works for the project.

We take into account the above indication in the following way:

- We merge relation **IndProject** into **Project**, making the attribute **resarea** nullable, and using it as a flag to distinguish the instances of **IndProject**.
- We also merge relation **WForInd** into **Project**, making the two attributes **resid** and **reslevel** nullable, and constraining them to be null exactly when attribute **resarea** is null.

Additionally, we remove the redundant attribute **project2** from **HasMWP**, at which point the relation **HasMWP**, becomes redundant, since the information about the project of an instance of **MgmtWP** is already contained in **MgmtWP**.

Problem 2: Restructured relational schema

We specify here only the relations with their constraints that have been changed with respect to the schema obtained through the direct translation.

The relations **IndProject**, **WForInd**, and **HasMWP** have been removed.

Project(code, description, resarea*, resid*, reslevel*)

inclusion: **Project**[code] \subseteq **WorksFor**[project]

foreign key: **Project**[code] \subseteq **HasMWP**[project]

foreign key: **Project**[resid,reslevel] \subseteq **Senior**[id,level]

foreign key: **Project**[resid,reslevel,code] \subseteq **WorksFor**[resid,reslevel,project]

The constraints 1, 2, and 4 stay unchanged, while the additional constraint 3 of the original conceptual schema can now be expressed through a **CHECK constraint**:

```
CHECK (SELECT WF.resid, WF.reslevel, WF.project
        FROM WorksFor WF, Project P
        WHERE WF.project = P.code AND IS NOT NULL P.resarea)
IN (SELECT resid, reslevel, code FROM Project)
```

Problem 3

Consider a database containing the tables **Worker** (ssn, age, department), where null values are allowed for the attributes age and department, and **Purchase** (ssn, item, year), where null values are not allowed. The first table stores the ssn (primary key), age, and department of a set of workers. Each tuple in the second table represents a purchase made by a worker for an item in a certain year. Note that there is nothing that prevents a purchase p from having a “twin”, i.e., a purchase done by the same worker of the same item in the same year as those of p . We know that the database satisfies the foreign key constraint from **Purchase**[ssn] to **Worker**[ssn].

1. Express in **SQL** a query that, calculates how many purchases in the **Purchase** table have no twins.
2. Express in **relational algebra** a query that, computes the ssn of workers who have done at least one purchase from 2020 onwards, and for whom if the department is not known, the age is known.

Problem 3: Solution (1/2)

Worker (ssn, age, department)

Purchase (ssn, item, year)

1. Express in **SQL** a query that, calculates how many purchases in the **Purchase** table have no twins.

```
SELECT COUNT(*)  
FROM (SELECT ssn, item, year  
      FROM Purchase  
      GROUP BY ssn, item, year  
      HAVING COUNT(*) = 1)
```

Problem 3: Solution (2/2)

Worker (ssn, age, department)

Purchase (ssn, item, year)

- Express in **relational algebra** a query that, computes the ssn of workers who have done at least one purchase after 2020 and for whom if the department is not known, the age is known.

PROJ_{ssn} ((**SEL**_{(department IS NOT NULL) OR (age IS NOT NULL)} Worker) **JOIN**
(**SEL**_{year >= 2020} Purchase))

Problem 4

Consider the relational schema S consisting of two relations $P(X, Y)$ and $Q(Z, W)$, where:

1. P has X as primary key and Y as additional key;
2. the default value for attribute Y of P is 1;
3. Q has Z as primary key;
4. P and Q have mutual foreign keys with delete policies as follows:

```
P(X, Y) UNIQUE Y, DEFAULT Y=1,  
        FOREIGN KEY P[Y] ⊆ Q[Z] ON DELETE SET DEFAULT  
Q(Z, W) FOREIGN KEY Q[W] ⊆ P[X] ON DELETE CASCADE
```

Answer the following question: Is there a database D consistent with S such that the deletion from the database D of a tuple t of P causes an error?

If the answer is negative, motivate your answer.

If the answer is positive, show such a database D and the corresponding tuple t and describe the reason for the error.

Problem 4: Solution

P		Q	
<u>X</u>	Y	<u>Z</u>	W
7	4	4	2
2	5	5	7
3	1	1	2

Consider the above database D , and assume that we delete tuple $(7,4)$ from P . Then, the **DELETE CASCADE** policy for $Q[W] \subseteq P[X]$ (and since there can be no other tuple t in P such that $t[X] = 7$) forces to delete tuple $(5,7)$ from Q . Then, the **DELETE SET DEFAULT** policy for $P[Y] \subseteq Q[Z]$ forces to set to 1 (which is the default value for $P[Y]$) the value 5 in tuple $(2,5)$.

In this way, we obtain the two tuples $(2,1)$ and $(3,1)$ in P , which causes a violation of the key constraint **UNIQUE Y** of P .

Variant of Problem 4 with Solution

Consider now a variant of Problem 4 where Condition 2 now states that the default value for attribute **Y** (instead of **X**) of **P** is 1. Then the solution of the problem is as follows.

Consider an arbitrary database D consistent with S and a tuple t of **P** that is deleted.

- If there is no tuple s in **Q** such that $s[\mathbf{W}] = t[\mathbf{X}]$, there is no further change to make.
- If instead there are tuples s_1, \dots, s_n in **Q** such that $s_1[\mathbf{W}] = \dots = s_n[\mathbf{W}] = t[\mathbf{X}]$, then the **DELETE CASCADE** policy for **Q[W] ⊆ P[X]** (and the fact that **X** is a key of **P**, hence there can be no other tuple t' in **P** such that $t'[\mathbf{X}] = t[\mathbf{X}]$) forces to delete s_1, \dots, s_n from **Q**.
- If for every i with $1 \leq i \leq n$, there is no tuple t' in **P** such that $t'[\mathbf{Y}] = s_i[\mathbf{Z}]$, there is no further change to make.
- If instead all the tuples t_j in **P** such that $t_j[\mathbf{Y}] = s_i[\mathbf{Z}]$, for some i with $1 \leq i \leq n$, are t_1, \dots, t_m , then the **DELETE SET DEFAULT** policy for **P[Y] ⊆ Q[Z]** (and the fact that **Z** is a key of **Q**) forces to set all of $t_1[\mathbf{Y}], \dots, t_m[\mathbf{Y}]$ to **NULL** (which is the default value for **P[Y]**). None of the constraints is violated by this (in particular the key constraint on **P[Y]** is not violated).
- No further change needs to be made, and the resulting database is consistent with S .

Hence, there is no database D consistent with S such that the deletion of a tuple t of **P** causes an error.