

CS 448/688 Assignment #1 Solution

Notes:

1. There are typically several possible correct solutions for each question (but there are many more possible incorrect solutions).
2. For simplicity, throughout this solution, the symbol ∞ indicates JOIN.

1.

Ex 4.5.2

- a) $\pi_{ename} (\sigma_{aname='Boeing'} (Employees \infty Certified \infty Aircrafts))$
- b) $\{ t[ename] \mid e \in Employees \wedge \exists c (c \in Certified \wedge c[eid] = e[eid] \wedge \exists a (a \in Aircrafts \wedge a[aid] = c[aid] \wedge a[aname] = 'Boeing')) \}$
- c)

<i>Employees</i>	<i>eid</i>	<i>ename</i>	<i>salary</i>
	<u>_EID</u>	P.ename	

<i>Aircraft</i>	<i>aid</i>	<i>aname</i>	<i>cruisingrange</i>
	<u>_AID</u>	'Boeing'	

<i>Certified</i>	<i>eid</i>	<i>aid</i>
	<u>_EID</u>	<u>_AID</u>

Ex 4.5.4

- a) People have interpreted this in different ways. Some (including the book's authors) have interpreted it as requesting flights that can be flown by some pilot, in which case a solution is as follows:

$$\pi_{flno} (\sigma_{cruisingrange \geq distance \wedge salary > 100,000} (Employees \infty Certified \infty Aircrafts \infty Flights))$$

The proper answers use division, since you want all the high-paid pilots:

$$\pi_{flno, eid, ename, salary} (\sigma_{cruisingrange \geq distance} ((Employees \infty Certified \infty Aircrafts) \times Flight)) \div \sigma_{salary > 100,000} (Employees)$$

(Note that since Flight has no common attributes with the other relations, natural join will provide the same result as cross-product.)

- b) For the first form, we get the following:
 $\{ f[flno] \mid f \in Flights \wedge \exists a, c, e (a \in Aircraft \wedge c \in Certified \wedge e \in Employees \wedge a[cruisingrange] \geq f[distance] \wedge e[salary] > 10000 \wedge a[aid] = c[aid] \wedge e[eid] = c[eid]) \}$

For the second form, we get the following:

$$\{ f[flno] \mid f \in Flights \wedge \forall e (e \in Employees \wedge e[salary] > 100,000 \wedge \exists c (c \in Certified \wedge c[eid] = e[eid] \wedge \exists a (a \in Aircraft \wedge c[aid] = a[aid] \wedge f[distance] \leq a[cruisingrange])) \}$$

c)

<i>Employees</i>	<i>eid</i>	<i>ename</i>	<i>salary</i>
	<u>_EID1</u>		>100,000
	<u>_EID2</u>		>100,000

<i>Certified</i>	<i>eid</i>	<i>aid</i>
	<u>_EID1</u>	<u>_AID</u>

<i>Aircraft</i>	<i>aid</i>	<i>aname</i>	<i>cruisingrange</i>
	<u>_AID</u>		<u>_CR</u>

<i>Flight</i>	<i>fno</i>	<i>from</i>	<i>to</i>	<i>distance</i>	<i>departs</i>	<i>arrives</i>
	P.G.			<u>_Dist</u>		

<i>Conditions</i>
<u>_CR</u> ≥ <u>_Dist</u>
COUNT. <u>_EID1</u> = COUNT. <u>_EID2</u>

Ex 4.5.6

- a) $\pi_{eid}(\text{Employee}) - \pi_{E1.eid}(\rho(E1, \text{Employees}) \bowtie_{E1.salary < E2.salary} \rho(E2, \text{Employees}))$
 b) $\{e1[eid] \mid e1 \in \text{Employees} \wedge \forall e2 (e2 \in \text{Employees} \Rightarrow e2[salary] \leq e1[salary])\}$
 or
 $\{e1[eid] \mid e1 \in \text{Employees} \wedge \neg \exists e2 (e2 \in \text{Employees} \wedge e2[salary] > e1[salary])\}$

c)

<i>Employees</i>	<i>eid</i>	<i>ename</i>	<i>salary</i>
	<u>P.</u>		<u>_S1</u>
			<u>_S2</u>

<i>Conditions</i>
<u>S1</u> = MAX. <u>S2</u>

Ex 4.5.8

- a) This query requires us to count the number of aircraft each pilot is certified for; however, the COUNT operation is not provided in the relational algebra. Thus this query is not expressible in relational algebra.
 b) Same as part a), we don't have the required COUNT operation in tuple relational calculus, so this query is not expressible in tuple relational calculus.
 c) This is possible to express in QBE, but it is tricky to compare the results of aggregated values. One method is to use an intermediate relation (as described in Section 6.9 in the text):

<i>Certified</i>	<i>eid</i>	<i>aid</i>
	<u>G._EID</u>	<u>_AID</u>

<i>Counts</i>	<i>eid</i>	<i>num</i>
<u>I.</u>	<u>_EID</u>	COUNT. <u>_AID</u>

This defines a new relation and inserts corresponding tuples. Then we can write:

<i>Counts</i>	<i>eid</i>	<i>num</i>
	P.	_CNT
		_ALLCNT

<i>Conditions</i>
CNT = MAX.ALLCNT

Ex 4.5.10

- We require a SUM operation to do this query, but it is not provided in the relational algebra, so this query is not expressible in relational algebra.
- Same as part a), we don't have the required SUM operation in tuple relational calculus, so this query is not expressible in tuple relational calculus.
-

Employees	eid	ename	salary	
			_S	P.SUM._S

2. Ex 4.4.1

- This query will produce an empty result or will be declared wrong by the compiler. This is because the projection of Parts on 'sid' does not make sense.

One possibility is to try a direct translation, but insert an impossible condition on the tuple variable ranging over Parts, as here:

$$\{ s[sname] \mid s \in Suppliers \wedge \exists p (p \in Parts \wedge p[color] = 'red' \wedge p \notin Parts \wedge \exists c (c \in Catalog \wedge c[cost] < 100 \wedge c[sid] = s[sid])) \}$$

If we interpret the question as having a typo, and that the projection was meant to be on 'pid', we get:

$$\{ s[sname] \mid s \in Suppliers \wedge \exists p (p \in Parts \wedge p[color] = 'red' \wedge \exists c (c \in Catalog \wedge c[cost] < 100 \wedge c[sid] = s[sid] \wedge c[sid] = p[pid])) \}$$

- This cannot be expressed in QBE, as there is no 'sid' attribute in the Parts table. However, if we again treated it as a typo we would have gotten:

Suppliers	sid	sname	address
	_SID1	P._S	

Catalog	sid	pid	cost
	_SID1	_PID1	<100

Parts	pid	pname	colour
	_PID1		red

Ex 4.4.3

- a) $\{ s[sname] \mid s \in \text{Suppliers} \wedge \exists p (p \in \text{Parts} \wedge p[\text{color}] = \text{'red'} \wedge \exists c (c \in \text{Catalog} \wedge c[\text{cost}] < 100 \wedge c[\text{pid}] = p[\text{pid}] \wedge s[\text{sid}] = c[\text{sid}])) \wedge \exists p2 (p2 \in \text{Parts} \wedge p2[\text{color}] = \text{'green'} \wedge \exists c2 (c2 \in \text{Catalog} \wedge c2[\text{cost}] < 100 \wedge c2[\text{pid}] = p2[\text{pid}] \wedge \exists s2 (s2 \in \text{Supplies} \wedge s2[\text{sid}] = c2[\text{sid}] \wedge s[sname] = s2[sname]))) \}$

b)

Suppliers	sid	sname	address
	<u>_SID1</u>	P._S	
	<u>_SID2</u>	_S	

Catalog	sid	pid	cost
	<u>_SID1</u>	<u>_PID1</u>	<100
	<u>_SID2</u>	<u>_PID2</u>	<100

Parts	pid	pname	colour
	<u>_PID1</u>		red
	<u>_PID2</u>		green

Ex 4.4.5

- a) $\{ s[sname] \mid s \in \text{Suppliers} \wedge \exists p (p \in \text{Parts} \wedge p[\text{color}] = \text{'red'} \wedge \exists c (c \in \text{Catalog} \wedge c[\text{cost}] < 100 \wedge c[\text{pid}] = p[\text{pid}] \wedge s[\text{sid}] = c[\text{sid}])) \wedge \exists p2 (p2 \in \text{Parts} \wedge p2[\text{color}] = \text{'green'} \wedge \exists c2 (c2 \in \text{Catalog} \wedge c2[\text{cost}] < 100 \wedge c2[\text{pid}] = p2[\text{pid}] \wedge \exists s2 (s2 \in \text{Supplies} \wedge s2[\text{sid}] = c2[\text{sid}] \wedge s[sname] = s2[sname] \wedge s[\text{sid}] = s2[\text{sid}])))) \}$

b)

Suppliers	sid	sname	address
	<u>_SID</u>	P.	

Catalog	sid	pid	cost
	<u>_SID</u>	<u>_PID1</u>	<100
	<u>_SID</u>	<u>_PID2</u>	<100

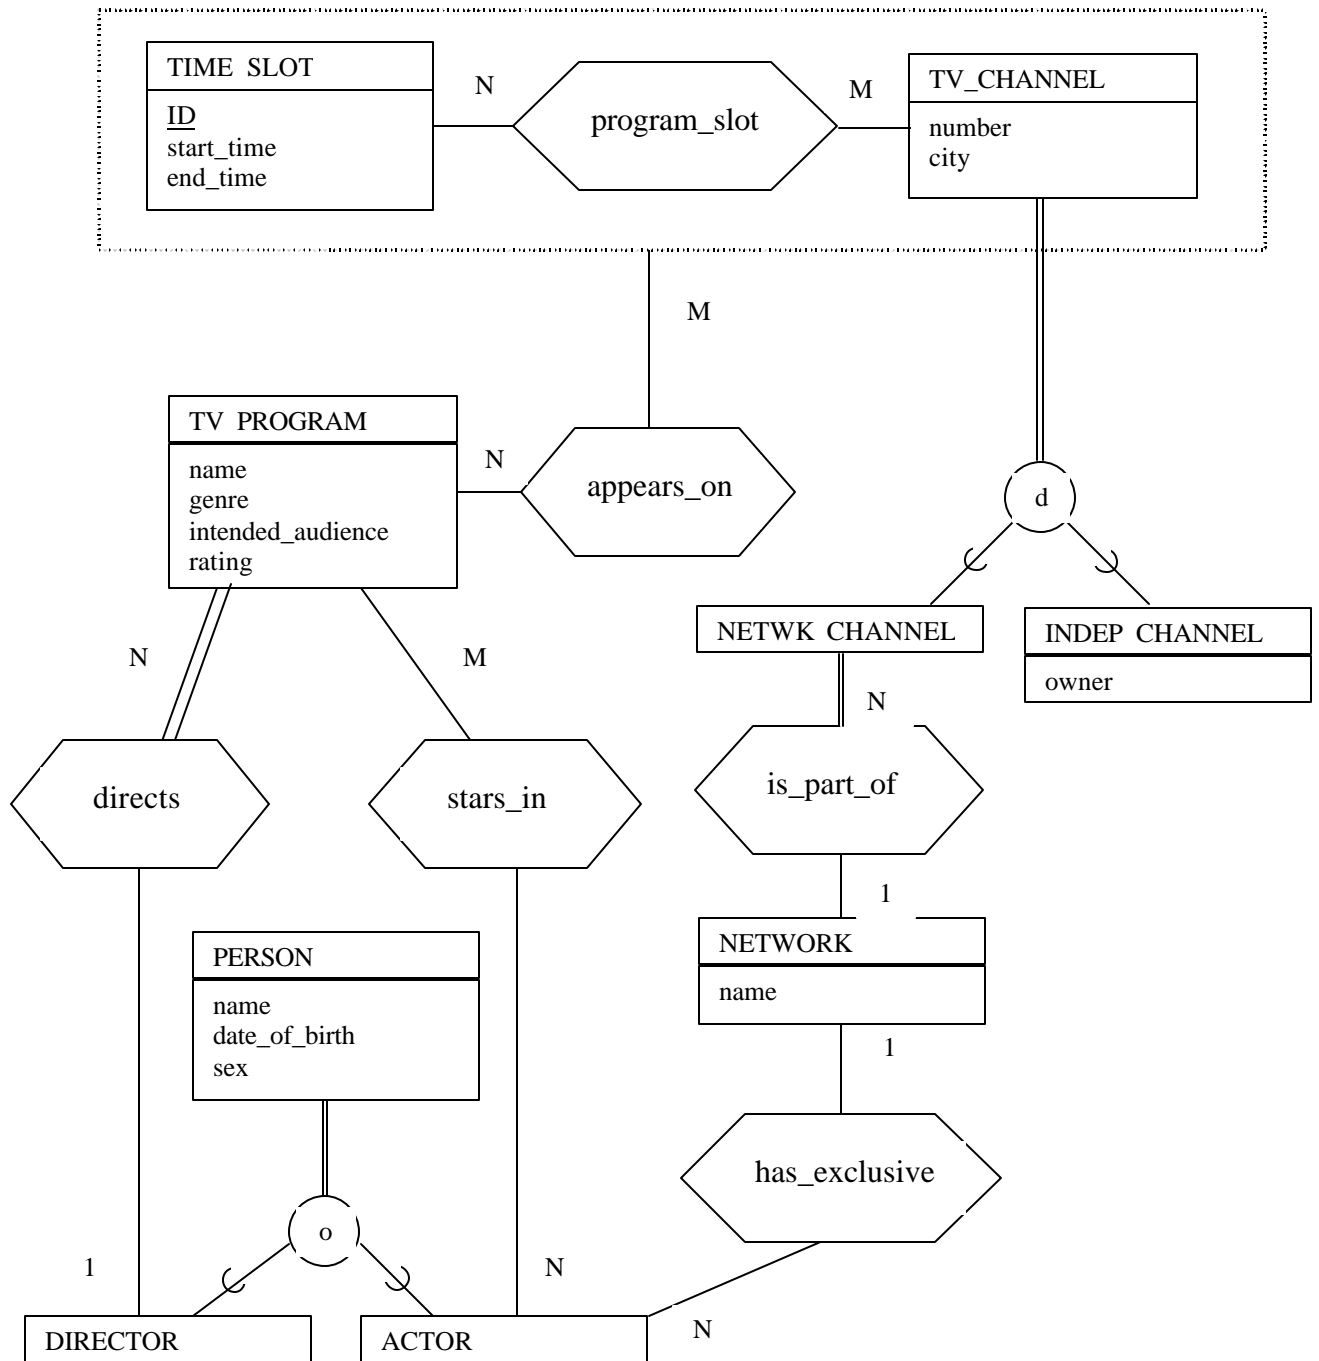
Parts	pid	pname	colour
	<u>_PID1</u>		red
	<u>_PID2</u>		green

3.

- a) $\pi_{\text{ENAME,RESP}} (\sigma_{\text{DUR}>12} (\text{EMP} \bowtie \text{WORKS}))$
- b) $\pi_{\text{PNAME,RESP}} (\sigma_{\text{DUR}>12 \vee \text{BUDGET}>20000} (\text{WORKS} \bowtie \text{PROJ}))$
- c) $\pi_{\text{PNAME,ENAME}} (\text{EMP} \bowtie \text{PROJ} \bowtie (\pi_{\text{ENO,PNO}} (\text{WORKS}) - \pi_{\text{ENO,PNO}} (\rho_{\text{ENO,PNO}} (\rho_{\text{A,WORKS}}) \bowtie_{\text{A.RESP=B.RESP} \wedge \text{A.DUR}<\text{B.DUR}} \rho_{\text{B,WORKS}))))$

4.

a) Note that we used a different notation to fit everything on one page.



Note: There is an additional unstated constraint that the time slots related to each TV channel must be pairwise non-overlapping.

b) Note: Primary keys are underlined and foreign keys are italic

Step 1 - Handling Entities:

TIME_SLOT(ID,start_time,end_time)
TV_CHANNEL(number, city)
TV_PROGRAM(name, genre, intended_audience, rating)
PERSON(name, date_of_birth,sex)
NETWORK(name)

Step 4 - 1:N Relationships

Modify relations to include foreign keys:

TV_PROGRAM(name, genre, intended_audience, rating, *director_name*: f.k. to DIRECTOR.name)
ACTOR(name, *exclusive_network*: f.k. to NETWORK.name)
TV_CHANNEL(number, city, type, owner, *network_name*: f.k. to NETWORK.name)

Add the constraint that network_name is not null iff type = "network"

Step 5: M:N Relationships

Add further tables:

STARS_IN(*program_name*: f.k. to TV_PROGRAM.name, *actor_name*: f.k. to ACTOR.name)
PROGRAM_SLOT(*time_slot*: f.k. to TIME_SLOT.ID, *channel*: f.k. to TV_CHANNEL.number)

Step 8: Specialization

When we use general specialization for Person we get:

ACTOR(name)
DIRECTOR(name)

Add the constraint that Actor.name and Director.name should be in Person.name

When we use disjoint specialization for the TV_Channel we get:

TV_CHANNEL(number, city, type, owner)

Add the constraint that type is either "network" or "independent"

Step 9: Aggregation

APPEARS_ON(time_slot: f.k. to TIME_SLOT.ID, channel: f.k. to TV_CHANNEL.number, program_name: f.k. to TV_PROGRAM.name)

Final set of relations:

TIME_SLOT (ID, start_time, end_time)

TV_CHANNEL(number, city, type, owner, *network_name*)

PROGRAM_SLOT (time_slot, channel)

TV_PROGRAM (name, genre, intended_audience, rating, *director_name*)

APPEARS_ON (time_slot, channel, program_name)

PERSON (name, date_of_birth, sex)

DIRECTOR (name)

ACTOR (name, *exclusive_network*)

STARS_IN (program_name, actor_name)

NETWORK (name)

As well as key constraints and foreign key constraints, we also have various domain constraints (i.e., whatever datatypes we would wish to associate with each attribute) plus inclusion dependencies between DIRECTOR.name and PERSON.name and between ACTOR.name and PERSON.name. We also have further constraint that TV_CHANNEL.network_name is not null iff TV_CHANNEL.type = “network”. (Note that because of this constraint we could eliminate the attribute TV_CHANNEL.type and use a test for NULL in TV_CHANNEL.network_name as the subtype discriminator.)