

School of Engineering and Computer Science
SWEN 435 Database System Engineering
Assignment 1

The objective of this assignment is to test your understanding of database foundations, basic terms, and the relational data model the entity relational model. It is worth 8% of your final grade. The assignment is marked out of 100. The assignment is due on **Friday, 29 March, 23:59 pm**. Please submit your assignment in **pdf** via the submission system.

**QUESTION 1****[20 marks]**

The TOP500 project lists the 500 most powerful non-distributed computer systems in the world (also called supercomputers). Suppose we use a relational database to manage the current and future data of this project. For this purpose, we use a relation schema with the attribute set {Performance, Name, Manufacturer, Country, Year}. The following table shows a portion of the current instance of the SUPERCOMPUTERS relation schema that stores data for some supercomputers. Note that the performance in the table is measured in petaFLOPS.

SUPERCOMPUTERS

Performance	Name	Manufacturer	Country	Year
442010	Fugaku	Fujitsu	Japan	2020
148600	Summit	IBM	United States	2018
94640	Sierra	IBM	United States	2018
93015	Sunway	NRCPC	China	2016
64590	Perlmutter	HPE	United States	2021
63460	Selene	Nvidia	United States	2020
61445	Tianhe-2A	NUDT	China	2013
44120	JUWELS	Atos	Germany	2020
35450	HPC5	Dell EMC	Italy	2020
30050	Voyager-EUS2	Microsoft	United States	2021

- a) [8 marks] For every set of attributes (that is, for every subset of the attribute set) decide whether you can deduce that it is *not* a candidate key, assuming the shown instance is legal. Justify your answer.

Answer:

- b) [4 marks] For every remaining set of attributes (that is, for every set not ruled out as a candidate key in part a)), discuss whether you consider it a suitable candidate key? Justify your answer.

Answer:

- c) [2 marks] Which of the candidate keys identified in part b) would you choose as the primary key? Justify your answer

Answer:

- d) [2 mark] Add a new tuple for a computer into the SUPERCOMPUTERS relation. How would you check that the primary key identified in part c) is still valid?

Answer:

- e) [2 mark] Create a relation that shows for each country in the table above the country and the capital, i.e., use a relation schema with attribute set {Country, Capital}. How many records are in your relation?

Answer:

- f) [2 mark]. Consider a relation schema with attribute set {Manufacturer, City} and assume that both attributes have a domain with ten values each. What would be the maximum number of records in an instance of this relation schema?

Answer:

QUESTION 2

[10 marks]

Suppose your software company has developed a relational database for the grocery store “Fruits and more”. The underlying database schema contains the following relation schemas:

- COMPANY (Cid: STRING, Name: STRING, Location: STRING) with primary key {Cid}
- FRUITS (Fid: STRING, Name: STRING, Cid: STRING, InStock: INTEGER, Price: INTEGER) with primary key {Fid, Cid} and foreign key Cid \subseteq COMPANY[Cid]

Below you find instances of these two relation schemas:

FRUITS

Fid	Name	Cid	InStock	Price
557	Apple	23XY	50	21
85520	Pear	A15F	0	78
63311	Pear	FVT35	211	49
36773	Kiwi	23XY	50	21
36773	Kiwi	FVT35	29	22

COMPANY

Cid	Name	Location
23XY	GreatFruits	Wellington
FVT35	Yummy	Wellington
F15A	GreatFruits	Levin
A15F	BetterFruits	Lower Hutt
5AB32	NiceFruits	<i>null</i>

Your tasks are as follows.

- a) [5 marks] Decide which of the following tuples can be added or removed, respectively. *Justify your answers!*

1. Insert tuple (*null*, ‘Tasty’, ‘Wellington’) into COMPANY

Answer:

2. Insert tuple (‘FVT35’, ‘SweetFruits’, ‘Porirua’) into COMPANY

Answer:

3. Delete tuple (‘A15F’, ‘BetterFruits’, *null*) from COMPANY

Answer:

4. Delete tuple (‘23XY’, ‘GreatFruits’, ‘Wellington’) from COMPANY

Answer:

5. Insert tuple (‘XYZ4’, ‘Wellington’, ‘Yummy’) into COMPANY

Answer:

b) [5 marks] Decide which of the following tuples can be added or removed, respectively.
Justify your answers!

1. Delete tuple ('36773', 'Kiwi', '23XY', 50, 21) from FRUITS

Answer:

2. Insert tuple ('55555', *null*, 'F15A', 2, 99) into FRUITS

Answer:

3. Insert tuple ('54556', 'Lemon', 'FV35', 20, 43) into FRUITS

Answer:

4. Insert tuple ('53557', 'Apple', '5AB32', 500, 1) into FRUITS

Answer:

5. Delete tuple ('46557', 'Apple', '23XY', 1, 21) from FRUITS

Answer:

QUESTION 3

[20 marks]

The Wellington Foreign Trade Office needs to translate hundreds of documents every day. To ensure professional translation in a timely manner the office cooperates with several translation agencies and expert translators in New Zealand. The processing of the data about translations as well as the checking of deadlines and quality requirements is time consuming and error prone if this is done manually on paper.

Therefore, the office wants to build a new database to record all relevant data that is needed for processing and checking translations. Suppose the following relation schemas have been proposed to belong to the database schema for the new database.

- TranslationAgency (AgencyNumber, Name) with primary key {AgencyNumber}
- Translator (Name, Phone, Field, IRDNumber) with unknown primary key
- IsExpert (Name, Language) with primary key {Name, Language}
- TranslationOrder (OrderNumber, OrderDate, PageNumber, Budget, FromLanguage, ToLanguage, Deadline) with primary key {OrderNumber}
- Assignment (Agent, OrderNumber, Part, Language, Name) with primary key {OrderNumber, Part}

The following additional constraints are known:

1. Each translator has a unique IRD number, a unique phone and a unique name.
2. For each translator, the IRDNumber must be specified, while Field may be left blank (if not known).
3. Translators can be experts in up to four languages.
4. An agent can assign a translation order to multiple translators who can be distinguished by the assigned part of the order.

Your tasks are as follows:

- a) [3 marks] For the relation schema Translator, identify all suitable candidate keys. Explain your answer. Which candidate key would you choose as the primary key? Justify your answer.

Answer:

- b) [5 marks] For each of the relation schemas, identify all suitable foreign keys (if there are any). Explain your answer.

Answer:

- c) [2 marks] For each of the relation schemas, decide which attributes must be declared as not null. Explain your answer.

Answer:

- d) **[5 marks]** Assume, the translator with name ‘Peter Pan’ in the Translator relation retires. When deleting the record of this translator from the Translator relation, all the assignments made to him should not be lost. How would you ensure this requirement? Explain your answer.

Answer:

- e) **[5 marks]** Assume, a translation order with order number ‘42’ in the TranslationOrder relation is cancelled. Suppose, however, that already some assignments have been made to translate parts of this translation order. When deleting the record of the translation order from the TranslationOrder relation, then all the assignments should be deleted, too. How can this requirement be ensured? Explain your answer.

Answer:

QUESTION 4

[30 marks]

You are asked to design a new database for your grandma’s collection of books. A book has a title, a release year, a unique international standard book number, a number of pages and was published by a certain publishing house.

A book can have one or more authors. The authors of a book are writers. A writer has a first name, a last name, a birthday, a nationality, and a home country. A book can have one or more editors. An editor has a first name, a last name. Editors oversee the emergence of a book from the first manuscript to the print-ready form. A book without authors has at least one editor.

Furthermore, your grandma buys books at certain bookstores which are either physical ones or online ones. Physical bookstores have a name, an address, a contact phone number and opening hours, while online bookstores have a name, a website and a contact email address. There are different genres such as adventure, comedy, crime, mystery, fantasy or science fiction. For every book at most two genres should be recorded in the database.

- a) **[24 marks]** Draw an extended ER diagram for the database above. Write down the corresponding extended ER schema, including declarations of all the entity types (showing attributes and keys) and relationship types (showing components, attributes and keys).
- b) **[6 marks]** There may be information, requirements or integrity constraints that you are not able to represent in your diagram. Give three examples of integrity constraints that have not been represented in your diagram.

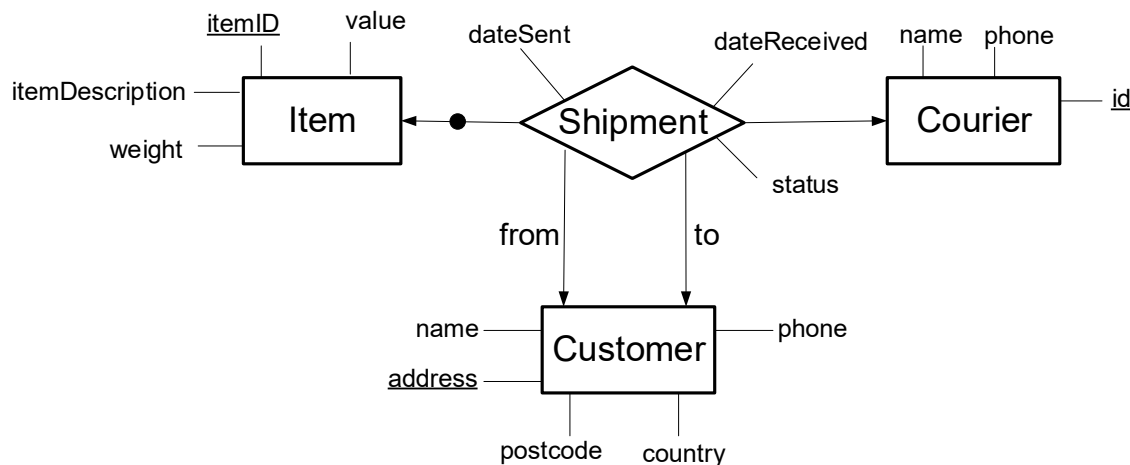
Remark: Whenever you feel that information is missing in the problem description above, add an assumption and make your assumption explicit. In practice you would consult the domain experts or potential users for clarification.

Answer:

QUESTION 5

[20 marks]

Consider the following extended ER diagram:



- a) [5 marks] Present the extended ER schema of the extended ER diagram above.
- b) [10 marks] Transform your extended ER schema into a relational database schema. In particular, list all the relation schemas in your relational database schema. For each relation schema, list all attributes, the primary key, the NOT NULL constraints, and the foreign keys.
- c) [5 marks] We also want to record information about expenses related to shipments. Each related expense has a date, a cost, and a description. The related expenses of a given shipment can be uniquely identified by the description.

Enhance the given extended ER diagram to reflect this additional information. Present the extended ER diagram with your proposed enhancements. Justify your answer.

Answer:
