Univ.-Prof. Dr.-Ing. Matthias Boehm
Graz University of Technology
Computer Science and Biomedical Engineering
Institute of Interactive Systems and Data Science
BMK endowed chair for Data Management

March 25, 2022

Exam INF.01017UF Data Management (Winter 2021/22, V2a)
Exam 706.010 Databases (Winter 2021/22, V2b)

Important notes: The working time is 90min, and lecture materials or any kind of mobile devices are not allowed. Please, make sure to put your *name* and *matriculation number* on the top right of the first page of the task description, and each additional piece of paper. You may give the answers in English or German, written directly into the task description.

Task 1 Data Modeling (25 points)

(a) Given the Entity-Relationship (ER) diagram above, specify the cardinalities in Modified Chen notation based on the following information. (**9 points**)

- An actor may play roles in an arbitrary number of movies (including none), and every movie has a cast of at least one but potentially many actors.

- A movie is directed by exactly one director, and a single director might produce (i.e., direct) an arbitrary number of movies.

- A movie review (with score, text comment, and date) refers to exactly one movie, but there can be 0, 1, or many reviews per movie.

- A single actor may receive multiple—specifically up to 8—awards (e.g., best actress, best supporting actress) for a specific, single movie. An actor receives such a single award for exactly one movie though. A single award (e.g., best ensemble) for a single movie can be awarded to one or multiple actors.

(b) Map the given ER diagram into a relational schema in third normal form, including data types, primary keys, and foreign keys. Your schema should also ensure that each movie has an associated director, and each review refers to a movie. Note that you only need to provide the final schema and there is no need to explain the normal forms. (**12 points**)

1

Actor: AID, FName, LName, DoB
plays in — Role
Movie: MID, Title, Year
directed by
Director: DID, FName, LName, DoB

DoB ..
day of birth

received
has

Award: PID, Name, Type, Year
Review: RID, Score, Comment, Date

(c) Assume an independent relation Movies(MID, Title, Year) with MID and Title each being unique and defined (not null). List below all super keys, all candidate keys, and select an appropriate primary key. Use (a,b,c) to indicate compound keys. (**4 points**)

- Super Keys:

- Candidate Keys:

- Primary Key:

Task 2 Structured Query Language (30 points)

Movies

MID	Title	Year	Length [min]	Budget [Mio $]	Revenue [Mio $]	GID
1	The Matrix	1999	136	63	455	2
3	Hangover	2009	100	35	470	1
2	Fast and Furious	2001	106	40	210	3
7	Passengers	2016	116	130	300	2
4	Horrible Bosses	2011	98	35	210	1
5	The Hunger Games	2012	142	80	700	2
6	Draft Day	2014	110	25	30	4
8	The Post	2017	116	50	180	5

Genres

GID	Name
1	Comedy
6	Romance
2	Science Fiction
3	Action
4	Sports Drama
5	Historical Drama
7	Documentary

(a) Given the Movies and Genres tables above, compute the results for the following three queries: (**15 points**)

```
Q1: SELECT M.Title, M.Year, G.Name
       FROM Movies M, Genres G
       WHERE M.GID = G.GID
         AND (G.Name LIKE '% Drama'
         OR M.Length BETWEEN 130 AND 140)
```

```
Q2: SELECT Title, Year
       FROM Movies WHERE Year > 2010
    INTERSECT
    SELECT Title, Year
       FROM Movies WHERE Revenue > 250
```

```
Q3: SELECT Name, round(avg(Revenue)) --avg=sum/count
       FROM Movies M JOIN Genres G ON (M.GID=G.GID)
       GROUP BY Name
       ORDER BY avg(Revenue) DESC
```

(b) Given the Movies and Genres tables above, write SQL queries to answer the following questions (in a way that is independent of the shown data): (**15 points**)

- Q4: Which movies belong to the genre "Science Fiction" (return the Title and Year, sorted in ascending order of Title)?

- Q5: Which movie from the years 2005-2015 (both inclusive) yielded the maximum Revenue (return the Title of this movie and its Revenue)?

- Q6: Compute the number of movies associated with each genre, including genres without any movies (return the genre Name, and count)?

Task 3 Query Processing (15 points)

(a) Assume relations $R(a, b, c)$ and $S(d, e)$, and indicate in the table below whether or not the two relational algebra expressions per row are equivalent in bag semantics ($\checkmark$ for equivalent, $\times$ for non-equivalent). For non-equivalent expressions, briefly explain why. (**5 points**)

Expression 1	Expression 2	Equivalent? Why Not?
$\sigma_{c=3}(\sigma_{b=7}(R))$	$\sigma_{c=3 \wedge b=7}(R)$	
$(\sigma_{c=3}(R)) \cap (\sigma_{b=7}(R))$	$\sigma_{c=3 \vee b=7}(R)$	
$R \bowtie_{a=e} S$	$\sigma_{a=e}(R \times S)$	
$\pi_{b,d}(R \bowtie_{a=e} S)$	$(\pi_{a,b}(R)) \bowtie_{a=e} (\pi_{d,e}(S))$	
$R - (\sigma_{a<b \wedge b<c \wedge a=c}(R))$	R	

(b) Draw a logical query tree for the following query. (**4 points**)

```
Q7: SELECT G.Name, count(*)
      FROM Movies M, Genres G
      WHERE M.GID = G.GID
        AND M.Revenue > 200
      GROUP BY G.Name
```

(c) Describe the conceptual ideas of a nested-loop join, and a hash join. Furthermore, assume $R \bowtie S$ with cardinalities $N = |R|$ and $M = |S|$, and enter the space and time complexity of these operators (in the open-next-close iterator model) in the table below. (**6 points**)

| Operator	Time Complexity	Space Complexity
Nested Loop Join		
Hash Join		

Task 4 Transaction Processing (10 points)

(a) Explain the concept of a database transaction log, and how it helps to ensure Atomicity and Durability of changes made by uncommitted and committed transactions on failures. (**6 points**)

(b) Explain the concept and use of database locks. In this context, also briefly describe lock type compatibility and multi-granularity locking. (**4 points**)

Task 5 Physical Design (20 points)

(a) Given the B-tree with k=2 below, insert the keys 7 and 37, and draw the resulting final B-tree. (**5 points**)

k = 2

```
                    | 14 | 25 |    |    |
         +-------------+------+-------------+
         |             |                    |
    | 3 | 11 |  |  | | 19 | 21 |  |  | | 26 | 27 | 29 | 31 |
```

(b) Given the B-tree with k=2 below, delete the keys 21 and 43, and draw the resulting final B-tree. (**5 points**)

k = 2

```
                    | 25 | 36 |    |    |
         +-------------+------+-------------+
         |             |                    |
| 19 | 21 |  |  | | 26 | 29 |  |  | | 37 | 42 | 43 | 46 |
```

(c) Which of the following trees are valid—i.e., satisfy the constraints of—B-trees with k=1. Mark each tree as valid (✓) or invalid (×) and name the violations. (**4 points**)

(b1)
```
      | 4 |   |
     +-----+
     |     |
  | 1 |  | | 9 | 12 |
```

(b2)
```
      | 4 | 16 |
    +----+----+----+
    |    |         |
| 1 |  | | 9 | 12 | |   |   |
```

(b3)
```
      | 4 | 9 |
     +----+----+
     |         |
  | 1 |  | | 12 | 16 |
```

(b4)
```
      | 5 | 16 |
    +----+----+----+
    |    |         |
| 1 | 6 | | 9 | 12 | | 17 |   |
```

(d) Recall the table `Movies` from Task 2 and perform a horizontal and vertical partitioning, respectively. Specifically, provide—for both scenarios—relational algebra expressions for partitioning `Movies` into two fragments `Movies1` and `Movies2`, as well as its subsequent reconstruction from the two fragments. (**6 points**)

- Horizontal Partitioning (row partitions):
    - Movies1 :=
    - Movies2 :=
    - Movies :=

- Vertical Partitioning (column partitions):
    - Movies1 :=
    - Movies2 :=
    - Movies :=