Politecnico di Torino
Database Management Systems

February 28^{th} 2012

1. (7 Points) The following relations are given (primary keys are underlined):

   E-SERVICE(E-ServiceID, Name, URL, Company)
   TASK(TaskID, E-ServiceID, Type, Priority)
   SERVER(ServerID, #Cores, Memory, Hard_Disk, IP_address)
   SCHEDULING(Date, Timestamp, TaskID, ServerID, CPU_time, Memory_usage)

   Assume the following cardinalities:

   - card(E-SERVICE)= 10^4 tuples,
     distinct values of Company $\simeq 100$,
   - card(TASK)= 10^6 tuples,
     distinct values of Type $\simeq 10$,
     distinct values of Priority $\simeq 5$
   - card(SERVER)= 10^3 tuples,
     MIN(#Cores) = 1, MAX(#Cores) = 16,
     MIN(Memory) = 2 GB, MAX(Memory) = 32 GB,
     MIN(Hard_Disk) = 100 GB, MAX(Hard_Disk) = 1 TB
   - card(SCHEDULING)= 10^{10} tuples,
     MIN(Date) = 1/1/2011, MAX(Date) = 31/12/2011.

   Furthermore, assume the following reduction factor for the group by condition:

   - having avg(Memory_usage)<100 MB $\simeq \frac{1}{10}$.

   Consider the following SQL query:

```
select Name, URL
from TASK T, E-SERVICE E
where E.E-ServiceID=T.E-ServiceID
      Priority <> 'Low' and Company = 'ItalBank'
      and TaskID in (select TaskID from SERVER S, SCHEDULING SC
                     where S.ServerID=SC.ServerID
                           and Date ≥ 1/8/2011 and Date ≤ 30/8/2011
                           and Memory = 32 GB and #Cores ≤ 4
                     group by TaskID
                     having avg(Memory_usage)<100 MB)
```

   For the SQL query:

   (a) Report the corresponding algebraic expression and specify the cardinality of each node (representing an intermediate result or a leaf). If necessary, assume a data distribution. Also analyze the group by anticipation.

   (b) Select one or more secondary physical structures to increase query performance. Justify your choice and report the corresponding execution plan (join orders, access methods, etc.).

2. (8 Points) The following relations are given (primary keys are underlined, optional attributes are denoted with *):

```
ROOM (RoomCode, Name, Location, Capacity)
MEETING(MCode, Topic, NumberOfInvitedPeople, MeetingDate*, RoomCode*, NumberOfParticipants*)
DATES_PROPOSED_FOR_MEETING(MCode, Date, NumberOfAvailablePeople)
PEOPLE_AVAILABILITY(EmployeeCode, MCode, Date)
```

A company wants to manage the scheduling of its meetings. The `MEETING` table contains the information about the meetings organized in the company. The `ROOM` table contains the information about the rooms available in the company.

Selection of the date and the room for a meeting. Write the trigger to define the date of a meeting, the room for the meeting, and the number of people who will participate in the meeting on the selected date (attributes `MeetingDate, RoomCode, NumberOfParticipants` in the `MEETING` table).

For each meeting to be scheduled, some possible dates are proposed (table `DATES_PROPOSED_FOR_MEETING`). Each person invited to a meeting expresses her availability for *only one* of the dates proposed for the meeting. When a person expresses her availability for a given meeting date (insert in the `PEOPLE_AVAILABILITY` table), the following actions should be performed. (a) The number of people available on that meeting date (attribute `NumberOfAvailablePeople`) should be updated. (b) It should be checked whether all the people invited to the meeting (attribute `NumberOfInvitedPeople`) have expressed their availability for one of the dates proposed for the meeting. If so, you must select the date for the meeting and the room where the meeting will be held. In particular, among the dates proposed for the meeting, you should select the date for which at least 80% of the people invited to the meeting have expressed their availability. Then, you should select a room whose capacity can accommodate all the people who have expressed their availability for the selected date, and is not occupied by other meetings already scheduled on the selected date. Assume that there is at most one date and at most one room satisfying all these conditions. Finally, the `MEETING` table should be updated. If it is not possible to select a date with a sufficient number of available people, or a free room with adequate capacity, the trigger ends without performing any further activity.

3. Data Warehouse design

Problem specifications

An international company distributes automotive replacement parts (e.g., lamps, suspensions, brakes, and other spare parts). The company would like to evaluate the distribution efficiency of its agencies around the world by analyzing the turnover of each agency and their distribution policies. A spare part can be used for a single car model, while the same car model can use different spare parts. In addition, a spare part can be distributed to several agencies and the same agency can deploy multiple parts.

Currently, each agency has an independent database to manage his business, but the company is interested in designing and implementing a data warehouse allowing to analyze the average daily turnover and average revenue for spare part of its agencies The analysis has to be performed according to:

- date, month, three-month period, four-month period, semester, year of the turnover,
- holiday, day of the week, day of the month,
- agency
- payment method, which means both the payment type (e.g., credit card, bank transfer, etc.) and the payment time (within 30 days, 60 days, etc.),
- spare part, and its type,
- car model (e.g., Punto, Golf, Focus, etc.) where the spare part can be used,
- car manufacturer (e.g., FIAT, Volkswagen, Ford, etc.) where the spare part can be used,
- spare model and manufacturer,
- city, region, state and continent where the agency is located.

The following are some of the frequent analyses the company is interested in:

(a) Considering only Italian agencies, for each agency and each car model, select the monthly turnover, the average daily turnover for each month, and the cumulative monthly turnover since the beginning of the year, separately for each payment type.

(b) For each spare part, select the total yearly turnover for each Cinese agency.

(c) For each spare part of cars manufactured by FIAT and considering only the turnover of the European agencies in 2011, for each city where an agency is located, select the monthly turnover, the monthly number of distributed spare parts, and the total monthly turnover for the region where the agency is located.

(d) For each three-month period in 2006 and 2007, select the total turnover and the average turnover per spare part, separately for each payment type.

The data warehouse contains data related to the years 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011.

Design

(a) (6 Points) Design the data warehouse to address the described issues. In particular, the designed data warehouse must allow efficient execution of all the queries described in the specifications.

(b) (8 Points) Write frequent queries (a) and (c) of the "problem specifications" using the extended SQL language.