Relational Databases &Consistency in Distributed Systems

Homework

(1) The following is a table of a veterinarian's office visit database.

Pet_ID	Pet_Name	Pet_Type	Pet_birthyear	Owner_name	Owner_birthday	Visit_date	Procedure_number	Procedure
246	Rover	Dog	2017	Frederick Douglas	1995	Jan 13/2018	0103	Vaccination Rabies
251	Spot	Dog	2015	Abraham Lincoln	1984	Jan 15/2018	0105	Vaccination Tetanus
341	Morris	Cat	2018	U.S. Grant	1991	Jan 15/2018	0103	Vaccination Rabies
357	Tweedy	Bird	2015	U.S. Grant	1991	Jan 15/2018	2104	Examine and Treat Wound
382	Mr. Ed	Dog	2014	R. Hayes	2000	Jan 15/2018	1005	Eye Wash
…	…	…	…	…	…	…		…
246	Rover	Dog	2017	Frederick Douglas	1995	Feb 19/2018	1003	Heart Worm Test
246	Rover	Dog	2017	Frederick Douglas	1995	Feb 19/2018	2104	Examine and Treat Wound
357	Tweedy	Bird	2015	U.S. Grant	1991	Feb 19/2018	3154	Annual Checkup
212	Mr. Ed	Horse	2011	L. Sherman	1982	Feb 19/2018	3154	Annual Checkup
…	…	…	…	…	…	…		…

Determine whether the following functional dependencies are valid or not. Give examples when you think a functional dependency is not valid and explain why you think if it is valid otherwise. You can assume that owner-name and owner-birthday uniquely identify the human being.

(a) Pet_Name —> Pet_Type

(b) Pet_ID, Visit_date —> Procedure
(c) Pet_ID —> Owner, Owner_birthday
(d) Pet_ID —> Pet_birthyear
(e) Pet_ID —> Pet_Name, Pet_birthyear, Pet_Type, Owner_name, Owner_birthday
(f) Owner_name, Owner_birthday —> Pet_ID, Pet_Name, Pet_Type
(g) Pet_Name —> Pet_Type

(2) Give examples for a write and a delete anomaly in the current veterinarian database table.

(3) Propose a better database scheme for the information in the veterinarians database.

(4) Two transactions access the same data-item and update them. Which of the following schedules suffer from the lost-update problem. The

$$r_1(x)r_2(x)w_1(x)w_2(x)$$

$$r_1(x)w_1(x)r_2(x)w_2(x)$$

$$r_1(x)r_2(x)w_2(x)w_1(x)$$

$$r_2(x)w_2(x)r_1(x)w_1(x)$$

(5) Two transactions access two data items. Transactions 1 updates their values, but Transaction 2 only reads them. Which of the following schedules could result in an inconsistent read?

$$r_1(x)w_1(x)r_2(x)r_1(y)w_1(y)r_2(y)$$

$$r_2(x)r_1(x)w_1(x)r_1(y)w_1(y)r_2(y)$$

$$r_2(x)r_1(x)r_2(y)w_1(x)r_1(y)w_1(y)$$

$$r_2(x)r_1(x)w_1(x)r_2(y)r_1(y)w_1(y)$$

(6) Conflict serializability of a schedule can be checked by permuting operations in the schedule until a serial schedule emerges. For example:

$r_2(x)r_1(x)r_1(y)w_2(x)r_3(x)w_1(z)w_3(x)w_3(y)$	can switch $w_1(z)$ and $r_3(x)$
$r_2(x)r_1(x)r_1(y)w_2(x)w_1(z)r_3(x)w_3(x)w_3(y)$	can switch $w_2(x)$ and $w_1(z)$
$r_2(x)r_1(x)r_1(y)w_1(z)w_2(x)r_3(x)w_3(x)w_3(y)$	can switch $r_2(x)$ and $r_1(x)$
$r_1(x)r_2(x)r_1(y)w_1(z)w_2(x)r_3(x)w_3(x)w_3(y)$	can switch $r_2(x)$ and $r_1(y)$
$r_1(x)r_1(y)r_2(x)w_1(z)w_2(x)r_3(x)w_3(x)w_3(y)$	can switch $r_2(x)$ and $w_1(z)$
$r_1(x)r_1(y)w_1(z)r_2(x)w_2(x)r_3(x)w_3(x)w_3(y)$	which is linear.

Therefore the first schedule is serializable and can be admitted.

Show that the following schedule is conflict serializable by making the right switches. Make one switch at a time and explain why the switch is allowed.

$$r_2(x)w_2(x)r_3(x)w_3(x)r_2(y)r_1(z)w_1(x)w_2(y)$$