INFS7907 Advanced Database Systems

Tutorial 3.2: Distributed Transactions and System Integration (T4)

Semester 1, 2004

Question 1: Implement the Basic Timestamp Ordering Algorithm and determine how these schedules will complete.

The initial timestamp values are

read-TS (X) = 1.0
read-TS (Y) = 1.2
read-TS (Z) = 2.2
write-TS (X) = 1.4
write-TS (Y) = 1.6
write-TS (Z) = 2.4

The timestamps of the transactions are

TS (T1) = 3.0
TS (T2) = 3.2

Schedule 1

T1	T2
read-item (X);	
	read-item (Z);
read-item (Y);	
	Z:= M + N;
	write-item (Z);
Y:= Y + X;	
write-item (Y);	

Schedule 2

T1	T2
read-item (X);	
	read-item (X);
	read-item (Y);
	X:= X + Y;
	write-item (X);
X:= X + N;	
write-item (Y);	
	read-item (Z);
	Z:= Z + Y;
	write-item (Z);

Question 2:
Consider the following Employees and Departments relations:

Employees(eid: integer, did: integer, sal: real)

Departments(did: integer, mgrid: integer, budget: integer)

Suppose that the Employees relation is stored in a database in Adelaide and the tuples with sal ≤ 100,000 are replicated in a database in Sydney. Consider the following three options for lock management: all locks managed at a single site, say, Brisbane; primary copy with Adelaide being the primary for Employees; and fully distributed. For each of the lock management options, explain what locks are set (and at which site) for the following queries. Also state which site the page is read from.

(a) A query submitted in Darwin wants to read a page containing Employees tuples with sal ≤ 50,000.

(b) A query submitted in Adelaide wants to read a page containing Employees tuples with sal ≤50,000.

Question 3:

Suppose that 2PC with Presumed Abort is used as the commit protocol. Explain how the system recovers from failure and deals with particular transaction T in each of the following cases:

(a) A subordinate site for T fails before receiving a *prepare* message.
(b) A subordinate site for T fails after receiving a *prepare* message but before making a decision.
(c) The coordinator site for T fails before sending a *prepare* message.
(d) The coordinator site for T fails after writing an *abort* log record but before sending any further messages to its subordinates.

Question 4:

Suppose there are three sites; A, B and C, all three contain copies of objects O1, O2 and O3. Assume that the read-any write-all technique is used.

There are three transactions, T1, T2 and T3 issued at sites A, B and C accordingly.

T1	**T2**	**T3**
read O1		
	read O2	
		read O3
write O2		
	write O3	
		write O1

(a) Draw the local waits-for graph at each site.

(b) Does deadlock exist in this situation?

(c) Draw the global waits-for graph.

System Integration (T4)

Question 5: A publication database is being maintained at two different sites. The tables maintained at the two sites, db1 and db2, are given below. These two databases now have to be integrated into a global conceptual schema.

Schema for db1:

scientist [name, age, state-birth]
publication [title, keyword, year]
book [num, title, year]
series [name, editor]
belongs-to [book-nr, series-name]
publisher-of [series-name, pub-name]
written-by [publ-title, auth-name]
author [name, address]
of-interest [auth-name, topic]

Schema for db2:

scientist [name, position, degree, city-of-birth]
borrower [pub-title, sci-name]
dealt-with [paper-title, topic-code]
paper [title, author]
journal [title, stack]
proceedings [title, stack]
book [title, author, stack]
topic [code, name]
concerns [pub-title, topic-code]
belongs-to-proc [paper-title, proc-title]

Integrate the above by defining views, where each view definition is given by the following statement:

CREATE VIEW view-name (attr1, attr2, ... attrn) AS SELECT statement

Tips: Try to establish semantic similarity between the attributes that are named and structured differently in the two schemas. For example written-by@db1.publ-title is the same as paper@db2.title

Not all attributes are found in both schemas. For example scientist@db1.age is not found in scientist@db2. Similarly scientist@db2.position is not found in scientist@db1.

SELECT statements used in the view creation would often have unions among result relations retrieved from various database relations belonging to different sites.

For example to create a global view of all topics, we can use

CREATE VIEW topic (*topic-name*) *AS*
(*SELECT keyword FROM publication@db1*)
UNION
(*SELECT name FROM topic@db2*);

Question 6:

Stockbroker A has a database about daily stock prices and trading volumes, where attribute "company" consists of stock exchange code (such as ECP, TLS, etc):

stock(company, date, low, high, close, volume)

Stockbroker B also has a database about daily stock closing prices as well as this stockbroker's trading information (the amount bought, sold and held), where there is one table for each stock (such as ECP, TLS, etc):

ECP(day, month, year, price, bought, sold, balance),
TLS(day, month, year, price, bought, sold, balance)…

Primary keys are underlined.

Identify the following data representation differences in the two databases (using one example from the above databases for each problem).

(a) Format conflicts

(b) Structural differences

(c) Missing Data

Question 7:

Discuss integration issues in the process of global schema design for the following two existing database systems.

System 1

The following information is stored and maintained. (A sample of data is given.)

Product	**(ID#,**	**name,**	**producer,**	**price)**
	32	A	IBM	500
	39	A	MS	475
	141	A	HP	512
	34	B	IBM	32
	017	C	HP	140
	211	D	IBM	750

System 2

For every company (producer) they store information in different tables. Additionally, they store the manager (sales rep) responsible for a given product.

The following is a sample of data:

IBM	**(ID#,**	**name,**	**slsrep,**	**price)**
	32	A	Bob	550
	211	D	Bob	750
	315	E	John	600
HP	**(ID#,**	**name,**	**slsrep,**	**price)**
	141	A	Keith	580
	018	C	Keith	200
	017	C	Keith	160
MS	**(ID#,**	**name,**	**slsrep,**	**price)**
	15	A	John	200
	39	A	John	490
	200	E	Bob	300

Question 8: <u>Discussion question</u>

Why is query optimisation in Mutidatabase Systems difficult? Compare this to parallel computation.