

Chapter 6: Introduction

Distributed Applications - Motivation

Why do we want to make our applications distributed?

- Applications are inherently distributed.
- A distributed system is more reliable.
- A distributed system performs better.
- A distributed system scales better.

Only paradigm to support use cases like Google, Facebook and Amazon!

Data-Centric Distributed Applications

Union of two technologies:

- Database Systems + Distributed Systems
- Database systems provide
 - data independence (physical & logical)
 - centralized and controlled data access
 - integration
- Distributed System provide scaling

Chapter 6: Introduction

Distributed Applications - Motivation

Why do we want to make our applications distributed?

- Applications are inherently distributed.
- A distributed system is more reliable.
- A distributed system performs better.
- A distributed system scales better.

Only paradigm to support use cases like Google, Facebook and Amazon!

Data-Centric Distributed Applications

Union of two technologies:

- Database Systems + Distributed Systems
- Database systems provide
 - data independence (physical & logical)
 - centralized and controlled data access
 - integration
- Distributed System provide scaling

Goals of Data-Centric Distributed Applications

- 1 Transparent management of distributed and replicated data
- 2 Reliability/availability through distributed transactions
- 3 Improved performance
- 4 Easier and more economical system expansion

Focus of this course: Distributed transactions!

If you are interested in the other aspects, visit my course in the winter term.

Challenges of Distributed/Replicated Data

- Storing copies of data on different nodes provides high availability and reliability
- Data needs be consistent
 - Synchronizing concurrent access
 - Detecting and recovering from failures
 - Deadlock management
- Fundamental conflicts between requirements (see CAP theorem)

Goals of Data-Centric Distributed Applications

- 1 Transparent management of distributed and replicated data
- 2 Reliability/availability through distributed transactions
- 3 Improved performance
- 4 Easier and more economical system expansion

Focus of this course: Distributed transactions!

If you are interested in the other aspects, visit my course in the winter term.

Challenges of Distributed/Replicated Data

- Storing copies of data on different nodes provides high availability and reliability
- Data needs be consistent
 - Synchronizing concurrent access
 - Detecting and recovering from failures
 - Deadlock management
- Fundamental conflicts between requirements (see CAP theorem)

Why *transactions*?

Transactions form a reasonable *abstraction concept* for many classes of real-life data processing problems.

- Transactions cope in a very elegant way with the subtle and often difficult issues of keeping data consistent in the presence of highly concurrent data accesses and despite all sorts of failures.
- This is achieved in a generic way invisible to the application logic so that application developers are freed from the burden of dealing with such system issues.
 - The application program simply has to indicate the boundaries of a transaction by issuing `BEGIN TRANSACTION` and `COMMIT TRANSACTION` calls.
 - The execution environment considers all requests receiving from the application programs execution within this dynamic scope as belonging to the same transaction.
 - For the transaction's requests and effects on the underlying data certain properties are guaranteed: ACID properties.

Challenges inherent to the transaction concept demonstrated by some examples

(Expl.1a) Debit/credit

Consider a debit/credit-program of a bank which transfers a certain amount of money between two accounts. Executing the program will give us the following transaction T:

```

BEGIN
% Withdraw
READ current value VA of account A from disk into T's local main memory;
decrement VA by amount X;
WRITE new value VA' = VA - X of account A from T's local main memory onto disk;
% Deposit
READ current value VB of account B from disk into T's local main memory;
increment VB by amount X;
WRITE new value VB' = VB + X of account B from T's local main memory onto disk;
COMMIT;

```


T
balance

- Assume when executing T the system runs into a failure, e.g. after writing A and before reading B. A customer of the bank has lost X money!
- Assume debit/credit-transaction T1 is running concurrently to a transaction T2, which computes the balance of the accounts A and B. Then the READ and WRITE accesses of both transactions may be interleaved. Assume that T2 is executed after T1 writing A and before T1 writing B, then the balance computed will be incorrect.

(Expl.1b) Distributed debit/credit

Assume that different branches of the bank are involved, where each branch maintains its own server. Assume further, at Branch1 a debit/credit-transaction is started and at Branch2 a balancing transaction, where both involve the same accounts. Transactions shall have access to accounts on remote server via remote procedure calls (RPC), a synchronous communication mechanism transparent to the programmer. We assume procedures *withdraw(account, amount)*, *deposit(account, amount)* and *getBalance(account)*.

A possible interleaving when both transactions are running in parallel.

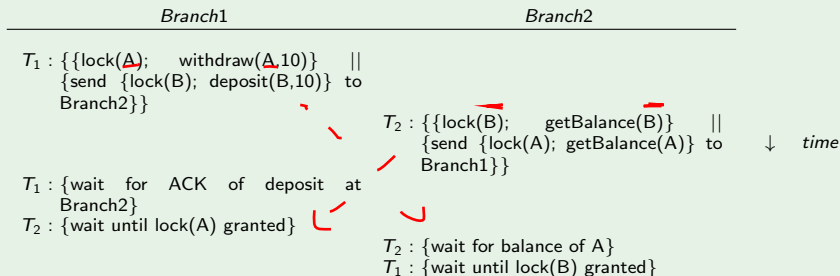


An incorrect balance will be displayed!

(Expl.1c) Distributed debit/credit

Assume that different branches of the bank are involved, where each branch maintains its own servers. Assume further, at Branch1 a debit/credit-transaction is started and at Branch2 a balancing transaction is started, where both involve the same accounts. Finally assume, that each transaction implements exclusive access to both accounts during execution. Communication is explicitly implemented by exchanging messages between the involved servers.

A possible interleaving when both transactions are running in parallel.



A deadlock has occurred which is difficult to detect!

(Expl.2) Electronic commerce *Amazon*

Consider the following purchasing activity, which covers several different servers located at different sites:

- A client connects to a bookstore's server and starts browsing and querying the catalog.
- The client gradually fills a shopping card with items intended to purchase.
- When the client is about to check out she makes final decisions what to purchase.
- The client provides all necessary information for placing a legally binding order, e.g. shipping address and credit card.
- The merchants's server forwards the payment information to the customer's bank or to a clearinghouse. When the payment is accepted, the inventory is updated, shipping is initiated and the client is notified about successful completion of her order.

- The final step of the purchasing is the most critical one. Several servers maintained by different institutions are involved.
- Most importantly it has to be guaranteed, that either all the tasks of the final step are processed correctly, or the whole purchasing activity is undone.

(Expl.3) Mobile computing

Assume that the described purchasing activity is performed via a smartphone. Then the described picture is even more complicated.

- The smartphone might be temporarily disconnected from the mobile net. Thus it is not guaranteed, that the state of the catalog as seen by the client reflects the state of the catalog at the server.
- If the client enters a dead spot during processing of the final step of the purchasing activity, confusion may arise, e.g. the purchasing is started again resulting in double orders.

Transaction Concept

ACID properties

- A tomicity: A transaction is executed completely or not at all.
- C onsistency: Consistency constraints defined on the data are preserved.
- I solation: Each transaction behaves as if it were operating alone on the data.
- D urability: All effects will survive all software and hardware failures.

⇒ *Concurrency Control* (I) and *Recovery* (A, D) provide the mechanisms needed to cope with the problems demonstrated by Expl.1-3.

Concurrency Control Refresh

Basics

- Set of transactions $\mathcal{T} = \{T_1, \dots, T_n\}$.
- A transaction is given as a sequence of read (R) - and write (W)-actions over database objects $\{A, B, C, \dots\}$, e.g.

$$T_1 = R_1A \ W_1A \ R_1B \ W_1B$$

$$T_2 = R_2A \ W_2A \ R_2B \ W_2B$$

$$T_3 = R_3A \ W_3B$$

- Let WX be the j -th action of transaction T and assume that RA_1, \dots, RA_n are the read actions of T being processed in the indicated order before WX . Then the value of X written by T is given by $f_{T,j}(a_1, \dots, a_n)$, where $f_{T,j}$ is an arbitrary, however unknown function and the a 's are the values read in the indicated order by the preceding read actions.
- A concurrent execution of a set of transactions is called schedule and is given as a - possibly interleaved - sequence of the respective actions, e.g.

$$S_1 = R_1A \ W_1A \ R_3A \ R_1B \ W_1B \ R_2A \ W_2A \ W_3B \ R_2B \ W_2B$$

$$S_2 = R_1A \ W_1A \ R_3A \ R_1B \ W_1B \ R_2A \ W_2A \ W_3B \ R_2B \ W_2B$$

$$S_3 = R_3A \ R_1A \ W_1A \ R_1B \ W_1B \ R_2A \ W_2A \ R_2B \ W_2B \ W_3B$$

- A schedule is called serial, if it is not interleaved.

$T_1 \quad T_2 \quad T_3$


Correctness

- an)
 A schedule is called (conflict-)serializable¹ if there exists a (conflict-)equivalent serial schedule over the same set of transactions.
- For a given schedule S over a set of transactions, the conflict graph $G(S)$ is given as $G(S) = (V, E)$, where the node set V is the set of transactions in S and the set of edges E is given by so called conflicts as follows:
 - $S = \dots W_i A \dots R_j A \dots \Rightarrow T_i \rightarrow T_j \in E$, if there is no other write-action to A between $W_i A$ and $R_j A$ in S .
 - $S = \dots W_i A \dots W_j A \dots \Rightarrow T_i \rightarrow T_j \in E$, if there is no other write-action to A between $W_i A$ and $W_j A$ in S .
 - $\hat{S} = \dots R_i A \dots W_j A \dots \Rightarrow T_i \rightarrow T_j \in E$, if there is no other write-action to A between $R_i A$ and $W_j A$ in S .
- A schedule is serializable, iff its conflict graph is acyclic.

¹We consider only conflict-serializability and therefore talk about serializability in the sequel, for short.

Example

Schedule S_1 : R_1A W_1A R_3A R_1B W_1B R_2A W_2A W_3B R_2B W_2B
 Schedule S_2 : R_3A R_1A W_1A R_1B W_1B R_2A W_2A R_2B W_2B W_3B



S_1 is serializable, S_2 is not.

To exclude not serializable schedules, a so called *transaction manager* enforces certain transaction behaviour.

T_1 T_3 T_2

2-Phase Locking (2PL)


- Serializable schedules are guaranteed, if all transactions obey the 2PL-protocol:
 - For each transaction T , each RA and WA has to be surrounded by a lock/unlock pair LA, UA :

$$T = \dots \underline{R/WA} \dots \Rightarrow T = \dots \textcircled{LA} \dots R/WA \dots \textcircled{UA} \dots$$

- For each A read or written in T there exists at most one pair LA and UA . *lock once*
- For each T and any LA_1, UA_2 there holds: $T = \dots LA_1 \dots UA_2 \dots$
 \Rightarrow No more locking after the first unlock!
- In any schedule S , the same object A cannot be locked at the same time by more than one transaction:

$$S = \dots L_i A \dots L_j A \dots \Rightarrow S = \dots L_i A \dots U_i A \dots L_j A \dots$$

- Every schedule according to 2PL is serializable, however
 - Not every serializable schedule can be produced by 2PL.
 - Deadlocks may occur.



Example 1

$$T_1 = L_1A R_1A L_1B U_1A W_1B U_1B,$$

$$T_2 = L_2A R_2A W_2A U_2A,$$

$$T_3 = L_3C R_3C U_3C.$$

$$S = L_1A R_1A L_1B \underline{U_1A L_2A} R_2A L_3C R_3C U_3C W_1B U_1B W_2A U_2A$$

$T_1 \rightarrow | \setminus$

Example 2

$$T_1 = L_1A R_1A L_1B U_1A W_1B U_1B,$$

$$T_2 = L_2A R_2A W_2A U_2A,$$

$$T_3 = L_3C R_3C U_3C.$$

$$S = L_1A R_1A L_1B U_1A L_2A R_2A L_3C R_3C U_3C W_1B U_1B W_2A U_2A$$

The **lock point** of a transaction using 2PL is given by the first unlock of the transaction.

2PL guarantees serializability of schedules.

Let S be a schedule of a set of 2PL-transactions $\mathcal{T} = \{T_1, \dots, T_n\}$.

concrete example

Assume, S is not serializable, i.e. the conflict graph $G(S)$ is cyclic, where w.l.o.g.

$T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T_k \rightarrow T_1$.

- Each edge $T \rightarrow T'$ implies T and T' having conflicting actions, where the action of T precedes the one of T' .
- Because of surrounding actions by lock/unlock and the 2PL-rule, T' can execute its action only after the lock-point of T . This implies the following structure of S , where A_1, \dots, A_k are data items:

$$\begin{aligned}
 S &= \dots U_1 A_1 \dots L_2 A_1 \dots, \\
 &\vdots \\
 S &= \dots U_{k-1} A_{k-1} \dots L_k A_{k-1} \dots, \\
 S &= \dots U_k A_k \dots L_1 A_k \dots
 \end{aligned}$$

- Let l_1, \dots, l_k be the lock points of the involved transactions. Then we have l_1 before l_2 , \dots , l_{k-1} before l_k and l_k before l_1 . However this is a contradiction to the structure of a schedule. Therefore S is serializable.

2PL guarantees serializability of schedules.

Let S be a schedule of a set of 2PL-transactions $\mathcal{T} = \{T_1, \dots, T_n\}$.

Assume, S is not serializable, i.e. the conflict graph $G(S)$ is cyclic, where w.l.o.g.

$T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T_k \rightarrow T_1$.

- Each edge $T \rightarrow T'$ implies T and T' having conflicting actions, where the action of T precedes the one of T' .
- Because of surrounding actions by lock/unlock and the 2PL-rule, T' can execute its action only after the lock-point of T . This implies the following structure of S , where A_1, \dots, A_k are data items:

$$\begin{aligned}
 S &= \dots U_1 A_1 \dots L_2 A_1 \dots, \\
 &\vdots \\
 S &= \dots U_{k-1} A_{k-1} \dots L_k A_{k-1} \dots, \\
 S &= \dots U_k A_k \dots L_1 A_k \dots
 \end{aligned}$$

- Let l_1, \dots, l_k be the lock points of the involved transactions. Then we have l_1 before l_2 , \dots , l_{k-1} before l_k and l_k before l_1 . However this is a contradiction to the structure of a schedule. Therefore S is serializable.

2PL guarantees serializability of schedules.

Let S be a schedule of a set of 2PL-transactions $\mathcal{T} = \{T_1, \dots, T_n\}$.

Assume, S is not serializable, i.e. the conflict graph $G(S)$ is cyclic, where w.l.o.g.

$T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T_k \rightarrow T_1$.

- Each edge $T \rightarrow T'$ implies T and T' having conflicting actions, where the action of T precedes the one of T' .
- Because of surrounding actions by lock/unlock and the 2PL-rule, T' can execute its action only after the lock-point of T . This implies the following structure of S , where A_1, \dots, A_k are data items:

$$\begin{aligned}
 S &= \dots U_1 A_1 \dots L_2 A_1 \dots, \\
 &\vdots \\
 S &= \dots U_{k-1} A_{k-1} \dots L_k A_{k-1} \dots, \\
 S &= \dots U_k A_k \dots L_1 A_k \dots
 \end{aligned}$$

- Let l_1, \dots, l_k be the lock points of the involved transactions. Then we have l_1 before l_2 , \dots , l_{k-1} before l_k and l_k before l_1 . However this is a contradiction to the structure of a schedule. Therefore S is serializable.

2PL guarantees serializability of schedules.

Let S be a schedule of a set of 2PL-transactions $\mathcal{T} = \{T_1, \dots, T_n\}$.

Assume, S is not serializable, i.e. the conflict graph $G(S)$ is cyclic, where w.l.o.g.

$T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T_k \rightarrow T_1$.

- Each edge $T \rightarrow T'$ implies T and T' having conflicting actions, where the action of T precedes the one of T' .
- Because of surrounding actions by lock/unlock and the 2PL-rule, T' can execute its action only after the lock-point of T . This implies the following structure of S , where A_1, \dots, A_k are data items:

$$\begin{aligned}
 S &= \dots U_1 A_1 \dots L_2 A_1 \dots, \\
 &\vdots \\
 S &= \dots U_{k-1} A_{k-1} \dots L_k A_{k-1} \dots, \\
 S &= \dots U_k A_k \dots L_1 A_k \dots
 \end{aligned}$$

- Let l_1, \dots, l_k be the lock points of the involved transactions. Then we have l_1 before l_2 , \dots , l_{k-1} before l_k and l_k before l_1 . However this is a contradiction to the structure of a schedule. Therefore S is serializable.

2PL guarantees serializability of schedules.

Let S be a schedule of a set of 2PL-transactions $\mathcal{T} = \{T_1, \dots, T_n\}$.

Assume, S is not serializable, i.e. the conflict graph $G(S)$ is cyclic, where w.l.o.g. $T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T_k \rightarrow T_1$.

- Each edge $T \rightarrow T'$ implies T and T' having conflicting actions, where the action of T precedes the one of T' .
- Because of surrounding actions by lock/unlock and the 2PL-rule, T' can execute its action only after the lock-point of T . This implies the following structure of S , where A_1, \dots, A_k are data items:

$$\begin{array}{l}
 S = \dots U_1 A_1 \dots L_2 A_1 \dots, \\
 \vdots \\
 S = \dots U_{k-1} A_{k-1} \dots L_k A_{k-1} \dots, \\
 S = \dots U_k A_k \dots L_1 A_k \dots
 \end{array}$$

not possible
2PL

- Let l_1, \dots, l_k be the lock points of the involved transactions. Then we have l_1 before l_2 , \dots , l_{k-1} before l_k and l_k before l_1 . However this is a contradiction to the structure of a schedule. Therefore S is serializable.

Recovery Refresh

Basics

- Reliability has to be achieved even though system components are unreliable, in general.
- Log files are a prerequisite for recovery from system failures.
- Log files are maintained as follows:
 - When a transaction T starts executing, (T, \textit{Begin}) is written into the log.
 - For each WA of a transaction T , (T, A, A_{old}, A_{new}) is written into the log, where A_{new} is the new value (After-image) and A_{old} the previous value (before-image) of A .
 - When a transaction commits its execution, (T, \textit{Commit}) is written into the log and otherwise (T, \textit{Abort}) .
- The write-ahead-log-rule (WAL) has to be observed: writing into the log must precede writing into the database.

Consequence of atomicity

- Whenever a transaction has processed a commit action, all its effects are permanent and will survive all failures.
- Whenever a transaction has processed a abort action - respectively is aborted -, all its effects are removed.

Recovery from system failures: Backwards Restart-Algorithm, logging has to be done on page-level

- $Redone := \emptyset$; $Undone := \emptyset$.
- The log is processed backwards. Let (T, A, A_{old}, A_{new}) the next log-entry to be considered. If $A \notin Redone \cup Undone$:
 - Redo:** If $(T, Commit)$ has already been found, then process WA with value A_{new} and perform $Redone := Redone \cup \{A\}$.
 - Undo:** Otherwise perform WA with value A_{old} and perform $Undone := Undone \cup \{A\}$.

Example

	System-failure	State after Restart
<p>T₁ LA RA WA CO UA</p> <p>T₂ LB RB LA RA WB CO UA,B</p> <p>T₃ LC RC WC</p> <p style="text-align: center;"> WA not yet materialized in the database, e.g. read accesses RA are expected </p>		
<p>DB :</p> <p>A₀</p> <p>B₀</p> <p>C₀</p>	<p>$f_2(f_1(A_0), B_0)$</p> <p>$f_3(C_0)$</p>	<p>$f_1(A_0)$</p> <p>$f_2(f_1(A_0), B_0)$</p> <p>C₀</p>
<p>localmemory/ systembuffer :</p> <p>T₁</p> <p>T₂</p> <p>T₃</p>	<p>$f_1(A_0)$</p> <p>$f_2(f_1(A_0), B_0)$</p> <p>$f_3(C_0)$</p>	
<p>Log (reduced):</p> <p>(T₁, A, A₀, f₁(A₀)), (T₁, CO), (T₂, B, B₀, f₂(f₁(A₀), B₀)), (T₂, CO), (T₃, C, C₀, f₃(C₀))</p>		