

Atomic Transaction (cont.)	Transaction
	Definition
□ Transaction T atomic \Rightarrow $Q_C = (P \lor Q)$.	A sequence of
 Immediate a set of a commit momentarily transition from P to Q <i>commit</i> <i>abort</i> (<i>roll back</i>) keeps P A crash before <i>commit</i> must leave the system in state P A crash after <i>commit</i> must leave the system in state Q <i>commit</i> is an acknowledgement to the client user that the transaction will hold "forever". T restartable if it has not made commit 	 read write end ≡ commit abort or rollback that is sent by a Client program to a data storage system. all or none of the <i>write</i> commands should be performed no other transaction should see intermediary data (no dirty reads)
5 (16) - DISTRIBUTED SYSTEMS Distributed Transactions - Sven Arne Andreasson - Computer Science and Engineering	s 6 (16) - DISTRIBUTED SYSTEMS Distributed Transactions - Sven Arne Andreasson - Computer Science and Engineering (CHALMER Nested Transactions
 Härder and Reuter, 1983 A transaction must obey the following properties: Atomicity — all or nothing Consistency — take the system from one consistent state to another consistent state, e.g. in a banking system it might be that the sum of all accounts is constant. Isolation — no illegitimate influence among different transactions. Durability — the write operations of a committed transaction must hold for the future. 	 Nested Transactions are transactions composed of other transactions, subtransactions. subtransactions at one level may run concurrently with other subtransactions at the same level. This could be done on different servers. subtransactions can commit or abort independently. Commit rules: A transaction may commit or abort only after its child transactions have completed. When a subtransaction completes, it makes an independent decision either to commit provisionally or to abort. A decision to abort is final. When a parent aborts, all of its subtransactions are aborted, even if they have committed provisionally. If the top level transaction commits, then all of the subtransactions that have provisionally committed can commit too, provided that none of their ancestors has aborted.



11 (16) - DISTRIBUTED SYSTEMS Distributed Transactions - Sven Arne Andreasson - Computer Science and Engineering

CHALMERS 12 (16) - DISTRIBUTED SYSTEMS Distributed Transactions - Sven Ame Andreasson - Computer Science and Engineering

() CHALMERS

two-phase commit	two-phase commit (cont.)
D phase 1:	
• Lock transaction data.	□ If after a crash the intention list
• perform reads and writes in normal order but	• is marked as valid the transaction was committed before the crash occurred.
 write on an intention list instead of changing the corresponding memory addresses. The intention list should be stored on stable storage, i.e. memory that survives a processor crash, e.g. a hard disk. <i>commit</i> - the intention list is made valid phase 2: write to the memory addresses according what has been stored on the intention list. Then the locks can be released. 	 The write operations on the intention list are performed then the locks are released. is not marked as valid the transaction was not committed before the crash occurred. The locks are released. The transaction might be restarted.
13 (16) - DISTRIBUTED SYSTEMS Distributed Transactions - Sven Arne Andreasson - Computer Science and Engineering	14 (16) - DISTRIBUTED SYSTEMS Distributed Transactions - Sven Ame Andreasson - Computer Science and Engineering
Distributed two-phase commit Protocol	three-phase commit
 The coordinator sends a <i>canCommit</i> request to each of the transaction participants. When a participant receives a <i>canCommit</i> request it replies with its vote, <i>Yes</i> or <i>No</i>, to the coordinator. If it answers <i>No</i> it aborts its part of the transaction. If it answers <i>Yes</i> it puts its new data on an intention list on stable storage. 	Takes care of the problem that the coordinator might crash.
□ phase 2 (completion phase):	
 phase 2 (completion phase): The coordinator collects the votes including its own: 	
 phase 2 (completion phase): The coordinator collects the votes including its own: If all votes are <i>Yes</i> the coordinator decides to commit the transaction and sends a <i>doCommit</i> message to all the participants. 	
 phase 2 (completion phase): The coordinator collects the votes including its own: If all votes are <i>Yes</i> the coordinator decides to commit the transaction and sends a <i>doCommit</i> message to all the participants. Otherwise the coordinator sends a <i>doAbort</i> message to all the participants that voted <i>Yes</i>. 	
 phase 2 (completion phase): The coordinator collects the votes including its own: If all votes are <i>Yes</i> the coordinator decides to commit the transaction and sends a <i>doCommit</i> message to all the participants. Otherwise the coordinator sends a <i>doAbort</i> message to all the participants that voted <i>Yes</i>. A participant that voted <i>Yes</i> is waiting for a <i>doCommit</i> or <i>doAbort</i> message. 	
 phase 2 (completion phase): The coordinator collects the votes including its own: If all votes are <i>Yes</i> the coordinator decides to commit the transaction and sends a <i>doCommit</i> message to all the participants. Otherwise the coordinator sends a <i>doAbort</i> message to all the participants that voted <i>Yes</i>. A participant that voted <i>Yes</i> is waiting for a <i>doCommit</i> or <i>doAbort</i> message. If it is a <i>doCommit</i> message it marks the intention list as valid and sends a <i>haveCommited</i> message to the coordinator. Then it copies the data on the intention list to the corresponding storage. 	