

# Replication


INF 5040 autumn 2007

lecturer: Roman Vitenberg

INF5040, Roman Vitenberg

1

## Replication architecture



INF5040, Roman Vitenberg

2

## Why replication I?

- Better performance
  - Multiple servers offer the same service – parallel processing of client requests
  - Geographical distribution
    - creating copies of data/objects closer to the clients leads to smaller network delay and possibly smaller network load

INF5040, Roman Vitenberg

3

## Why replication II?

- Better availability (continuous operation despite failures of individual components)
  - For many services it is important that availability with acceptable response time approaches 100%, despite that ...
    - Server processes may fail
    - Parts of the network may fail
    - Data may get corrupted
- Example: 5% chance of a server failure within a given period - two independent servers give 99.75% availability

INF5040, Roman Vitenberg

4


## Challenges of replication

- Complex mechanisms
  - Placement of replicas (and search for them)
  - Propagation of data (e.g., updates) among the replicas
  - Consistency maintenance
  - Monitoring and failover mechanisms
- These protocols also consume bandwidth
- Some of this complexity is exposed to the clients
  - Impossible to achieve complete replication transparency

INF5040, Roman Vitenberg

5

## Placement of replicas



INF5040, Roman Vitenberg

6

## Placement of replicas

- Permanent replicas
  - Clusters of servers
  - Geographically dispersed web mirrors (Akamai)
- Server-initiated caches
  - Placement of hosting servers
  - Placement of caches
    - Flash crowds in the Web
- Client-initiated caches
  - Enterprise proxies or web browser caches

INF5040, Roman Vitenberg

7

## Propagation of updates among the replicas

- Push-based propagation
  - A replica pushes the update to the others
  - May push the new data or parameters of the update operation
- Pull-based propagation
  - A replica requests another replica to send the newest data it has

Issue	Push-based	Pull-based
State at server	List of client caches	None
Messages sent	Update	Poll and update
Response time at client	Immediate (or update time)	Fetch-update time

INF5040, Roman Vitenberg

8

## Propagation of updates among the replicas

- Pushing data vs pushing updates
  - Pushing updates reduces traffic
  - Requires more processing power on each replica
  - Requires deterministic operations
- Hybrid push-pull approaches
  - Lease-based propagation
  - Pushing invalidations
    - A replica that performs the update notifies other replicas
    - A replica informed that a newer version is available will fetch the new version at a later point

INF5040, Roman Vitenberg

9

## Lack of consistency

### *Client 1*

deposit<sub>B</sub>(x, 100)

deposit<sub>A</sub>(y, 100)

### *Client 2*

balance<sub>A</sub>(y) → 100


balance<sub>A</sub>(x) → 0

INF5040, Roman Vitenberg

11

## Consistency

- A contract between the application developer and replicated service provider
  - The provider guarantees that the data will be updated according to some criteria
  - The application developer will need to devise application with these criteria in mind
- The consistency-efficiency-simplicity triangle



INF5040, Roman Vitenberg

12

## Sequential consistency

- Intuition: for each interleaved execution of the system there should exist a sequential execution that:
  - Performs the same operations as the original execution
  - Fulfills the specification of a single copy
  - The sequence of operations satisfies the “happened before” order

INF5040, Roman Vitenberg

13

## Example revisited

### Client 1

deposit<sub>B</sub>(x, 100)  
deposit<sub>A</sub>(y, 100)

### Client 2

balance<sub>A</sub>(y) → 100  
balance<sub>A</sub>(x) → 0

*This is not sequentially consistent, because there is no corresponding sequential execution of a non-replicated system*

## More examples

C1	C2	C1	
Deposit(a, 50)	Balance(a) = empty Balance(a) = 50	Deposit(a, 50) Balance(a) = empty	
C1	C2	C1	C2
Deposit(a,50)	Balance(a) = 50 Balance(a) = empty	Bal(b)=30 Bal(a)=em	Bal(a)=50 Bal(b)=em
		C3	C4
		Dep(a,50)	Dep(b,30)

## Active replication (replicated state machine)

- The idea:
  - Every replica sees exactly the same set of messages in the same order and will execute them in that order
- Benefits:
  - Every server is able to respond to client queries with updated data
  - Immediate fail-over
- Limitations:
  - Waste of resources, since all replicas are doing the same
  - Update propagation only, which requires determinism
- Different implementation levels
  - Machine instruction level (or VM), e.g., Tandem™
  - Logical state (software-based active replication)

INF5040, Roman Vitenberg

16

## Passive replication (primary-backup replication)

- One server plays a special primary role
  - Performs all the updates
  - May propagate them to backup replicas eagerly or lazily
  - Maintains the most updated state
- Backup servers may take off the load of processing client requests but only if stale results are ok
- Does not require deterministic operations
- Typically easier to implement than active replication
- Less network traffic during the normal operation but longer recovery with possible data loss
- Several sub-schemes (cold backup, warm backup, hot standby)

INF5040, Roman Vitenberg

17



## Primary-backup replication (cold backup)

- Only the primary is active
- Periodically checkpoints its state to backup storage
  - Stable storage or shared storage (SAN)
- When the primary fails, the backup is initiated, it loads the state from storage, and takes over
  - Slow recovery
    - Need to start the backup (run applications, obtain resources, etc.)
    - Either the backup replays the last actions from a log file, or it may miss the last updates since the most recent checkpoint
  - Most resource-efficient
- It is possible to have several backups to survive multiple failures

INF5040, Roman Vitenberg

18


## Primary-backup replication (other than cold backup)

- Warm backup
  - In this case, the backup is (at least) partially alive, so the recovery phase is faster
    - But typically still requires some replaying of last transactions, or losing the last few updates
- Hot standby (leader/follower)
  - The backup is also up, and is constantly updated about the state of the primary
- Local-write scheme
  - The primary migrates between the servers
  - Commonly used in mobile systems

INF5040, Roman Vitenberg

19

## From multicast to reliable group communication




INF5040, Roman Vitenberg

21

## Message ordering problem

➤ Ordering:



INF5040, Roman Vitenberg

22

## Group communication


- Group membership service
  - Dynamic maintenance of groups
  - Failure detection
  - Distributes information about changes in the membership
  - Address expansion – an address for multicast to the entire group

INF5040, Roman Vitenberg

23

## Group communication

- **View**: epoch of system evolution between two consecutive changes of membership
- The evolution of the system can be seen as a global sequence of views
- Illusion of a static system in each view



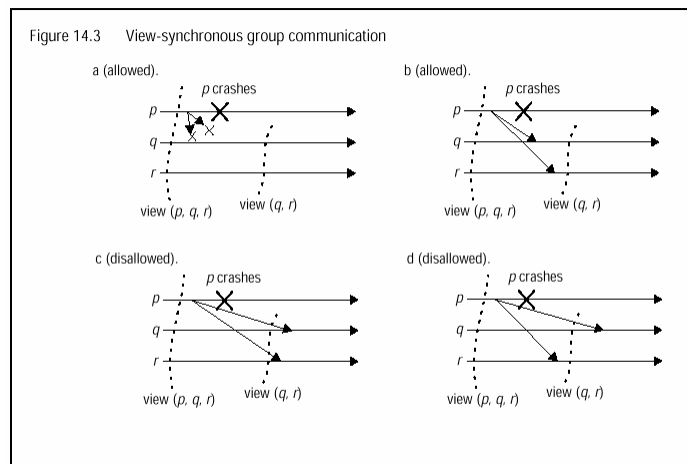
INF5040, Roman Vitenberg

24

# View synchrony

- Synchronization:
  - Processes deliver views and messages in the same sequence of events
  - If two different processes deliver  $m$ , they do it in the same view
- Delivering the same set of messages:
  - If the process  $p$  delivers  $m$  in  $v(g)$  and later delivers  $v(g')$ , then every process  $q$  that delivers both  $v(g)$  and  $v(g')$  delivers  $m$  in  $v(g)$
  - This implies retransmitting missing messages
  - If  $p$  delivers  $m$  in  $v(g)$ , and a process  $q$  does not deliver  $m$  in  $v(g)$ , the next view  $p$  delivers will not include  $q$

# Illustration (from the book)



## Ordered message delivery

- Some rule (binary relation) that establishes that two messages  $m_1$  &  $m_2$  sent in the system are ordered:  $m_1 < m_2$ 
  - Standard relation properties
- Two variants of ordered message delivery
  - Unreliable ordered delivery: if a process delivers  $m_1$  and  $m_2$ , it should deliver  $m_2$  after  $m_1$
  - Reliable ordered delivery: if a process delivers  $m_2$ , it should have already delivered  $m_1$ 
    - To implement, one needs unlimited message buffers
    - Delays while waiting for messages to arrive

INF5040, Roman Vitenberg

27

## Commonly used orderings

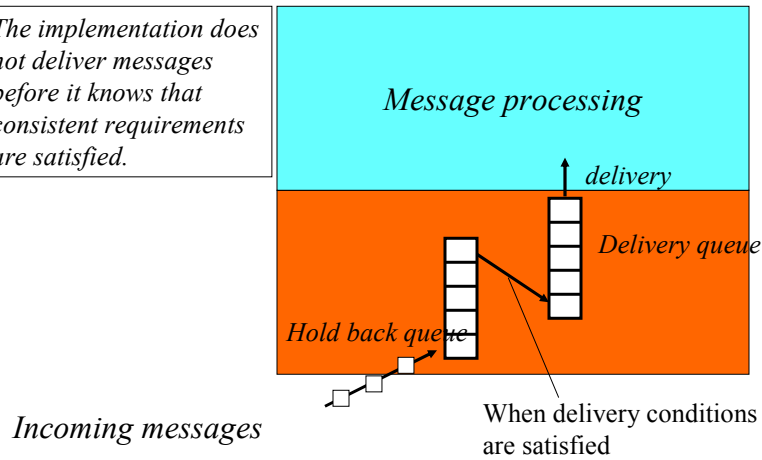
- FIFO
- Causal: two messages are ordered if related by the happen-before relation
  - Many applications require message delivery in an order that preserves cause and effect
    - Publish/subscribe (netnews), email, control systems, root cause determination
- Total: all messages will be received in the same order by all the processes in the group
  - Useful towards implementing the state machine abstraction

INF5040, Roman Vitenberg

28

# Implementing ordered delivery

*The implementation does not deliver messages before it knows that consistent requirements are satisfied.*



INF5040, Roman Vitenberg

29