

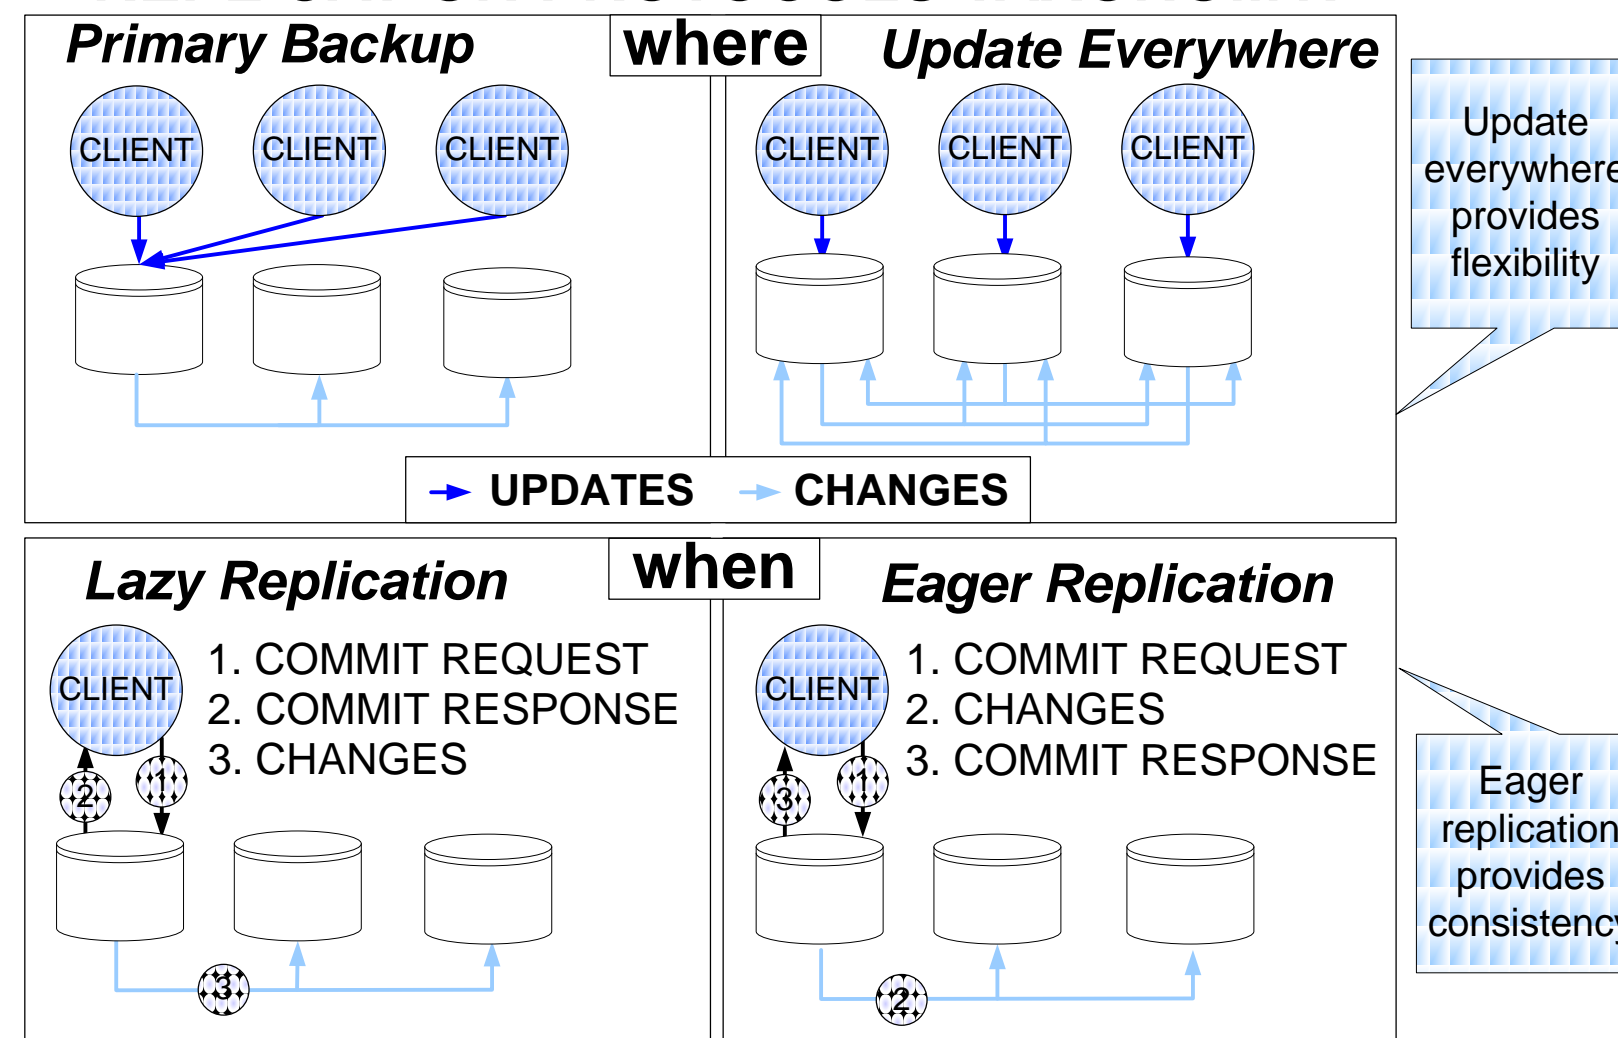
High Performance Database Replication for Storage Area Networks

Damian Serrano Garcia
 Prof: Marta Patiño Martinez
 {dserrano, mpatino}@fi.upm.es

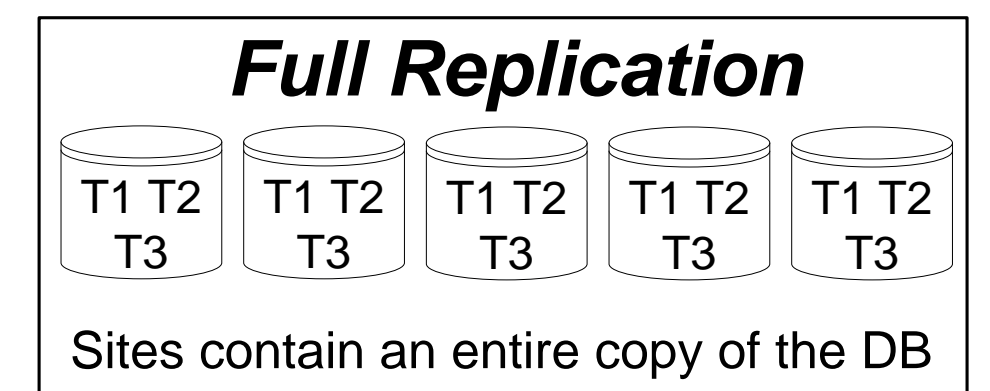
Background

- High availability of data has become crucial for the dependability of many applications and services.
- A replicated database replicates data in several locations to increase data availability and scalability.
- Replication protocols coordinate the execution of transactions on a replicated DB and the propagation of changes.

REPLICATION PROTOCOLS TAXONOMY:




- Database replication protocols attain a medium scalability (20-30 replicas) with eager full replication.



Current Status

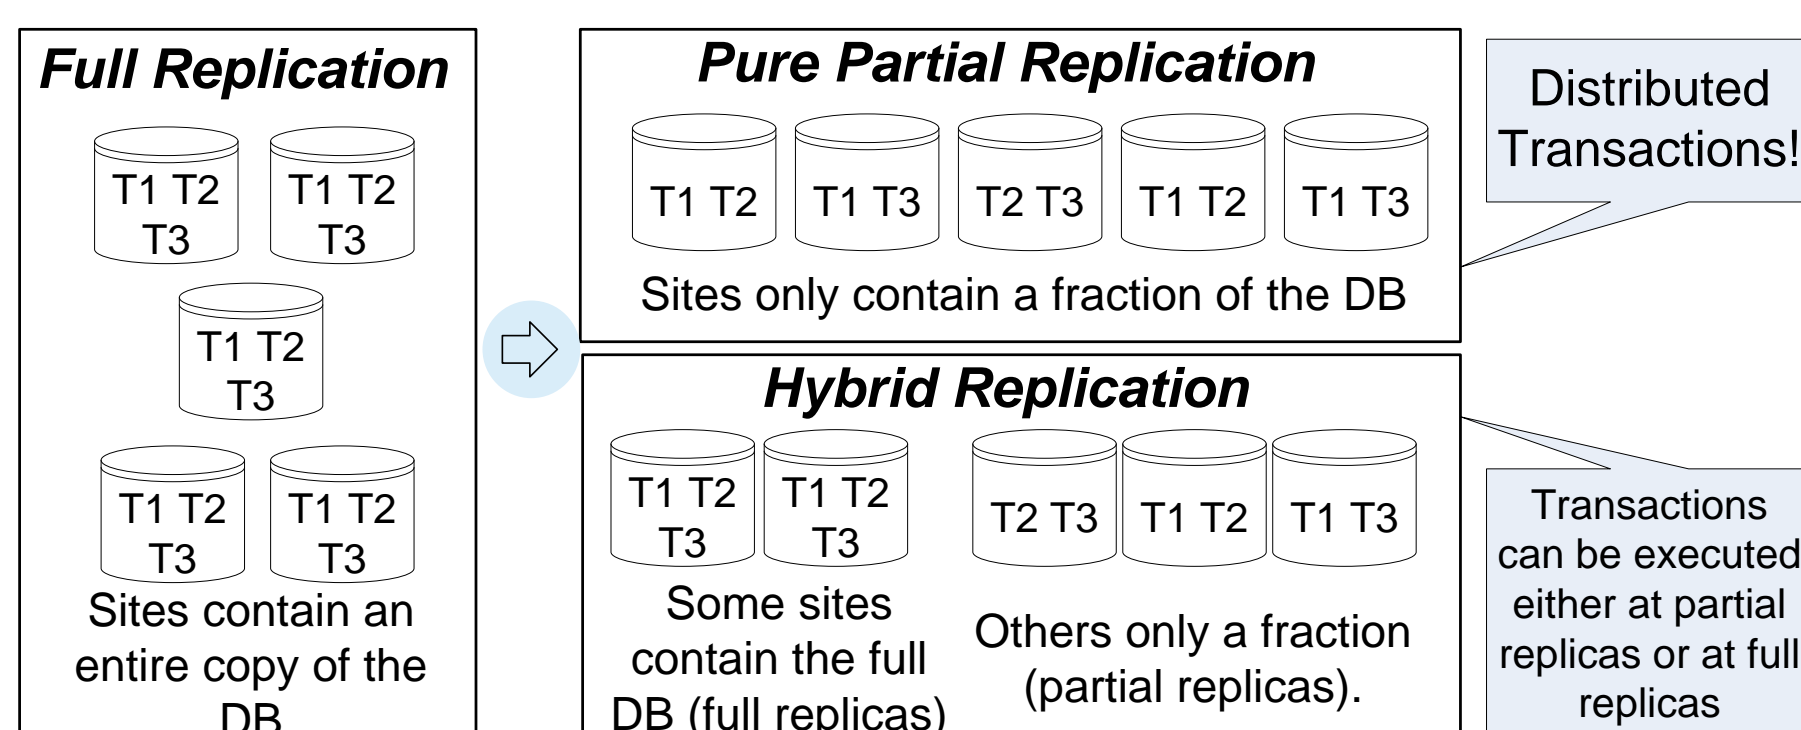
- Full replication does not scale beyond certain limits and is inappropriate for large scale systems:



- Full replication performance is bounded by the amount of remote work to be done.

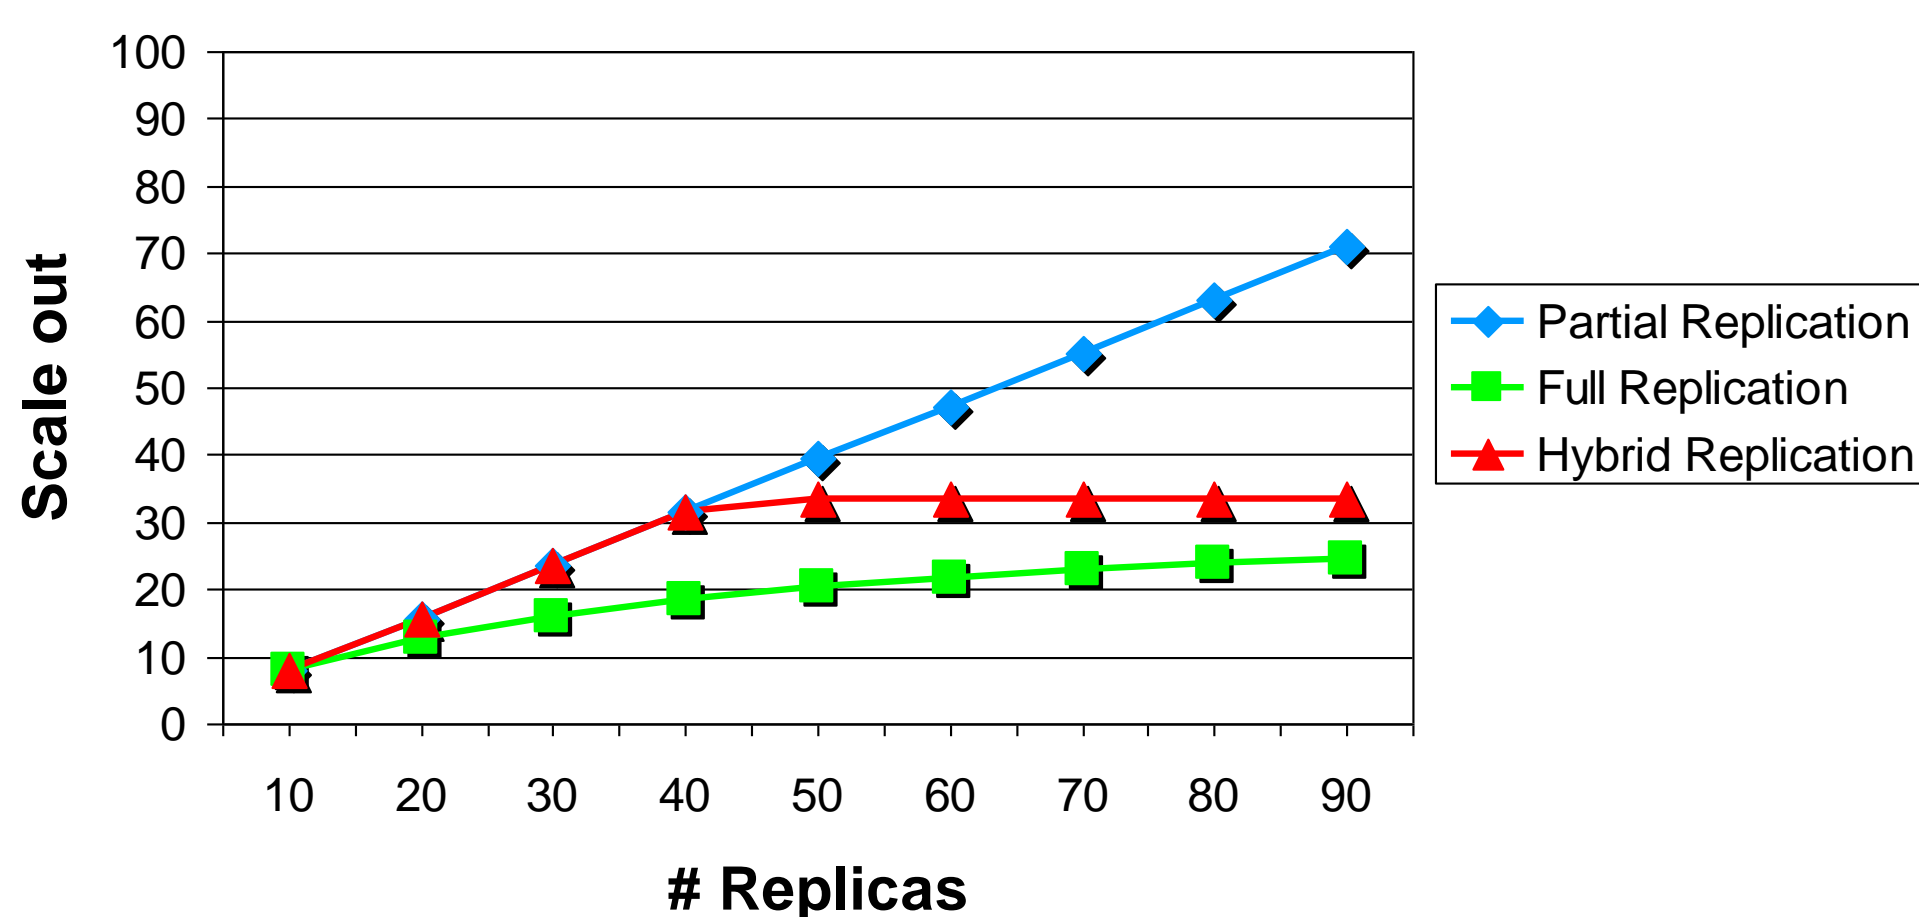
First Stage

- From full replication to partial replication:



Analytical Results

Scale Out



This PhD will

- **explore** partial replication as one additional mechanism to overcome the scalability ceiling of current approaches.
- **exploit** new correctness criteria such as 1-copy snapshot isolation to enhance scalability.
- **build** a highly scalable replicated database upon SANs (low latency and cards that execute firmware to offload protocol processing from the host)
- **develop** autonomic reconfiguration and optimisation techniques to maximise performance.

Analytical Model

- An analytical model for the scale out of hybrid and pure partial replication has been developed.
- The scale out measures how many times the throughput of a non-replicated site is increased by the replicated system.
- The scale out is modelled as:
 - An equation system.
 - An linear program that maximizes the scale out.

Preliminary Conclusions

- Partial replication increases scalability.
- With pure partial data replication the scale out only depends on the number of replicas.
- Hybrid partial replication scales slightly better than full replication.

PhD Timeline

