

Energy-Efficient Data Replication in Cloud Computing Datacenters

Dzmitry Kliazovich **University of Luxembourg, Luxembourg**
Pascal Bouvry




Dejene Boru University of Trento, Italy
Fabrizio Granelli



Albert Y. Zomaya University of Sydney, Australia



Data Replication in Data Centers



Data Replication in Data Centers

- Central Database
 - Located in wide area network
 - Hosts complete data for the all cloud applications
- Datacenter-level Database
 - Located at each datacenter
 - Replicates most frequently used data
- Rack-level Database
 - Located within the rack
 - Hosts data frequently accessed by rack server


Data Replication in Data Centers

- Objective 1: Performance of Cloud Applications
- Objective 2: Energy Efficiency

Why Data Replication is Important for Clouds?

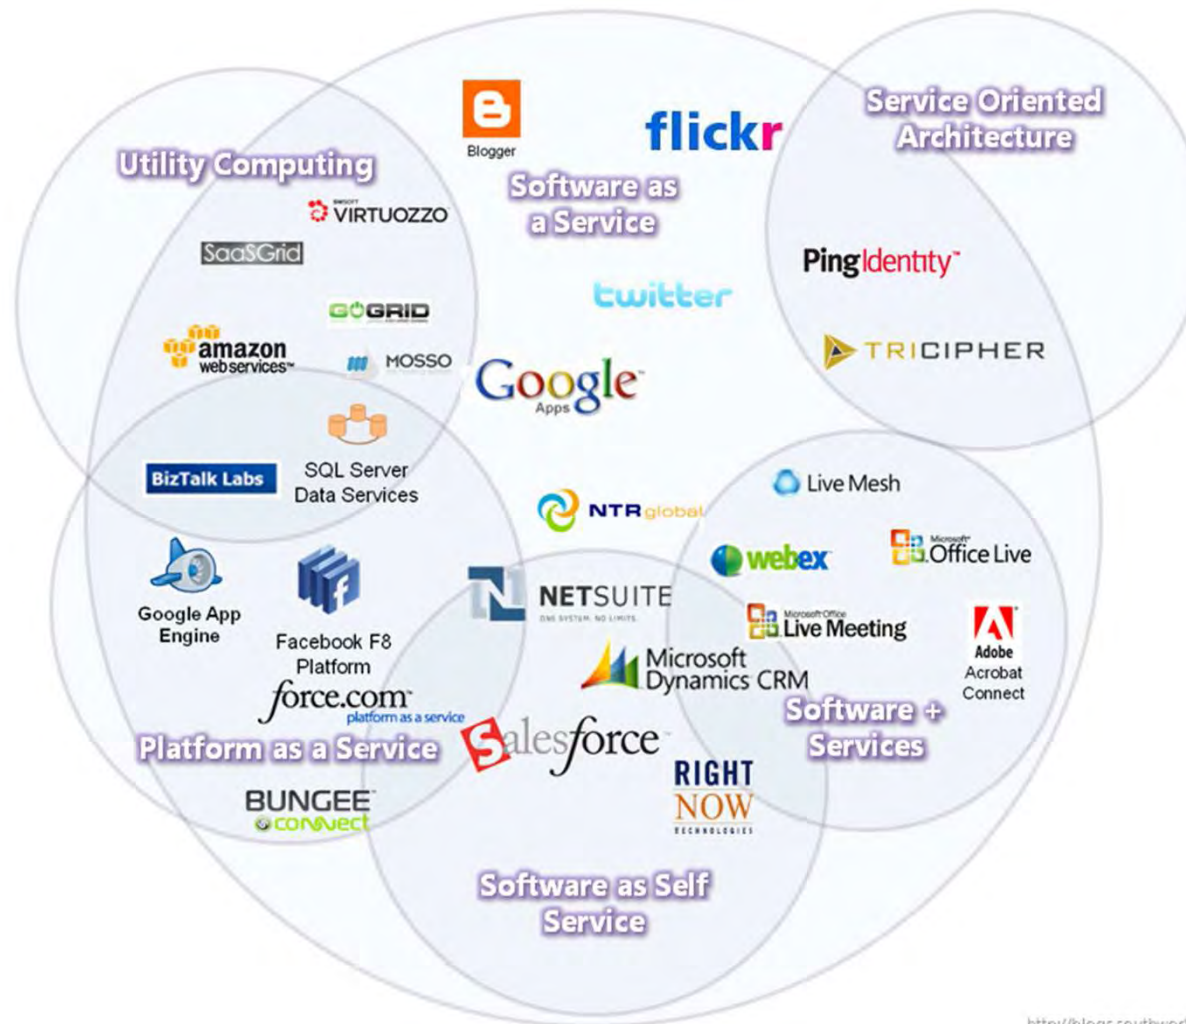
Cloud Computing

- Cloud computing market: \$241 billion in 2020
- Main focus is on Software-as-a-Service (SaaS)



Source: Larry Dignan, "Cloud computing market", ZDNet, 2011.

Cloud Computing Applications

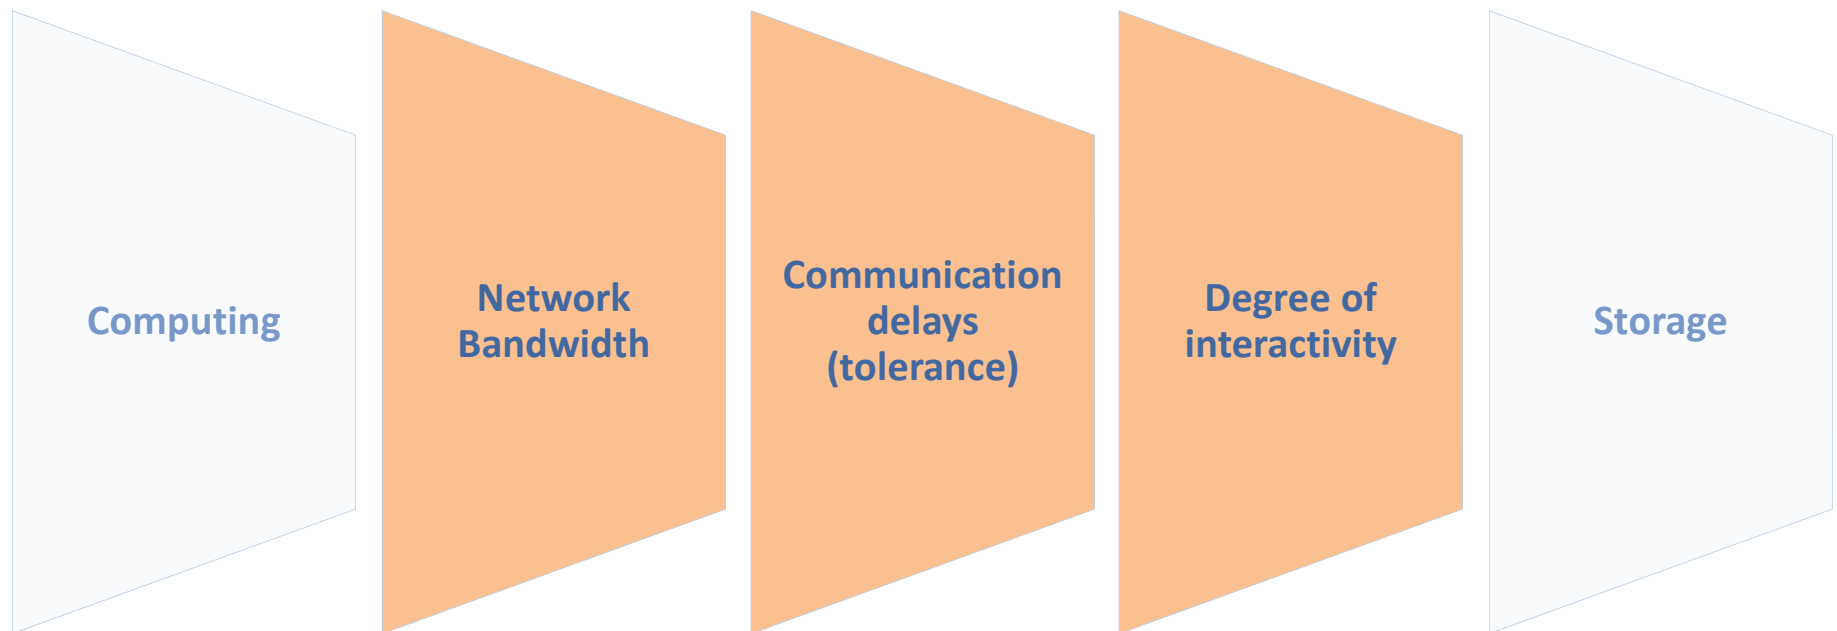


<http://blogs.southworks.net/mwołoski>

Resource Requirements of Cloud Applications



Resource Requirements of Cloud Applications



Cloud Computing Applications

CLASSIFICATION OF CLOUD APPLICATIONS

Cloud application	Resource requirement			Storage
	Computing	Bandwidth	Low communication delay	
Cloud gaming Video conferencing Online office Collaborative editing CRM Remote desktop Cloud Synchronization Video streaming Cloud storage Cloud backup Voice conferencing Social networking HPC				

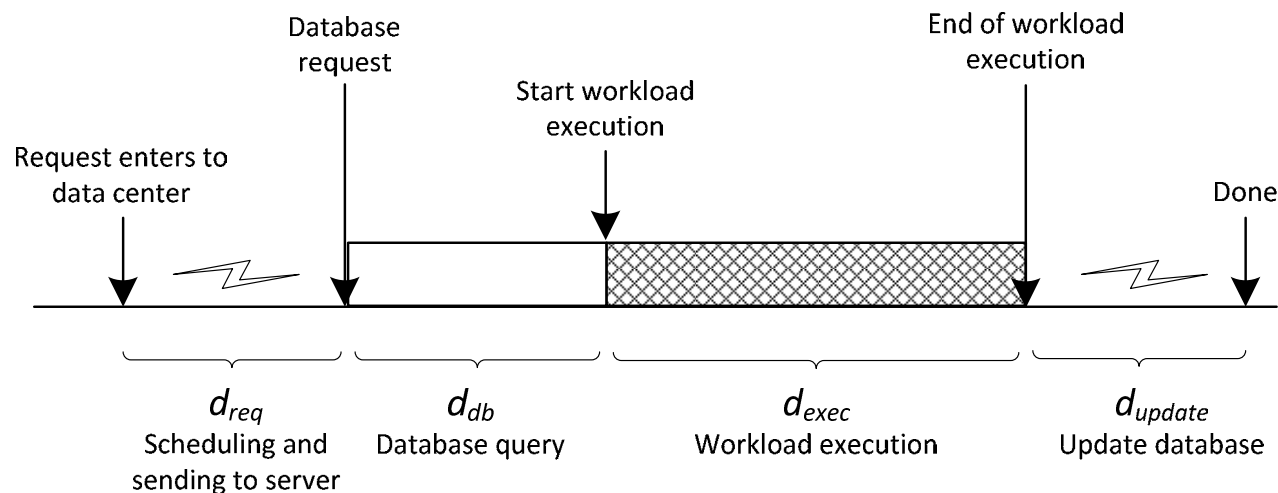
H: High, M: Medium and L: Low

Communication resources

Cloud Computing Applications

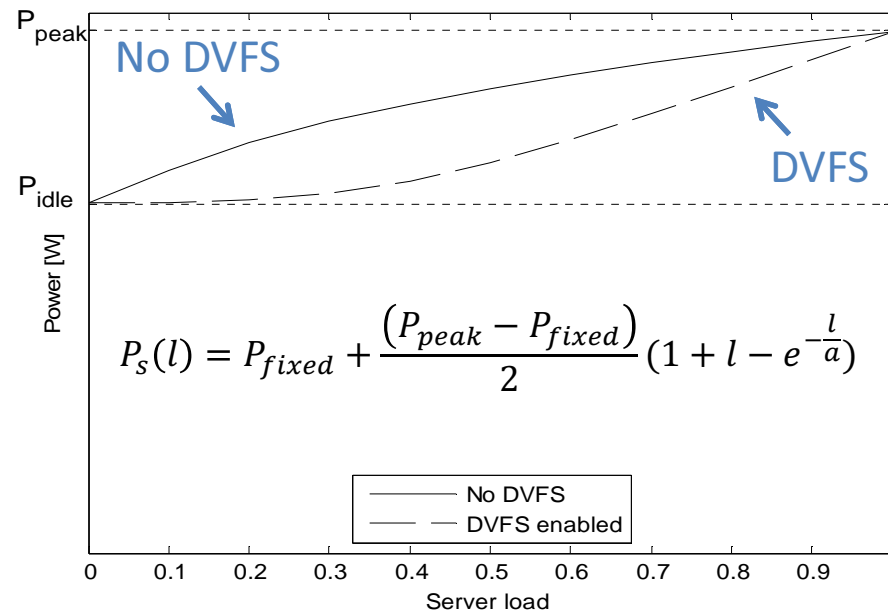
- Workload execution timeline

- $Delay = d_{req} + d_{db} + d_{exec} + d_{update}$



Energy Models

- Computing Servers Power Consumption
 - Non-linear
 - Idle servers consume 50-70% of their peak load



Energy Models

- Network Switch Power Consumption

$$P_{switch} = P_{chassis} + n_{linecards} \cdot P_{linecard} + \sum_{i=0}^R n_{ports,r} \cdot P_r$$



Chassis
~ 36%



Linecards
~ 53%



Port transceivers
~ 11%

Evaluation Results

Cloud Computing Simulator



Simulating Energy-Efficient Clouds


available at
<http://greencloud.gforge.uni.lu>

- Measures cloud performance and energy efficiency
- First to simulate cloud **communications with packet-level precision**
- Implements network-aware scheduling
- Implements complete TCP/IP protocol stack

Evaluation Results


- Downlink Bandwidth

- Bandwidth is proportional to size of data items
- Exceeding available capacity will trigger replication



Evaluation Results

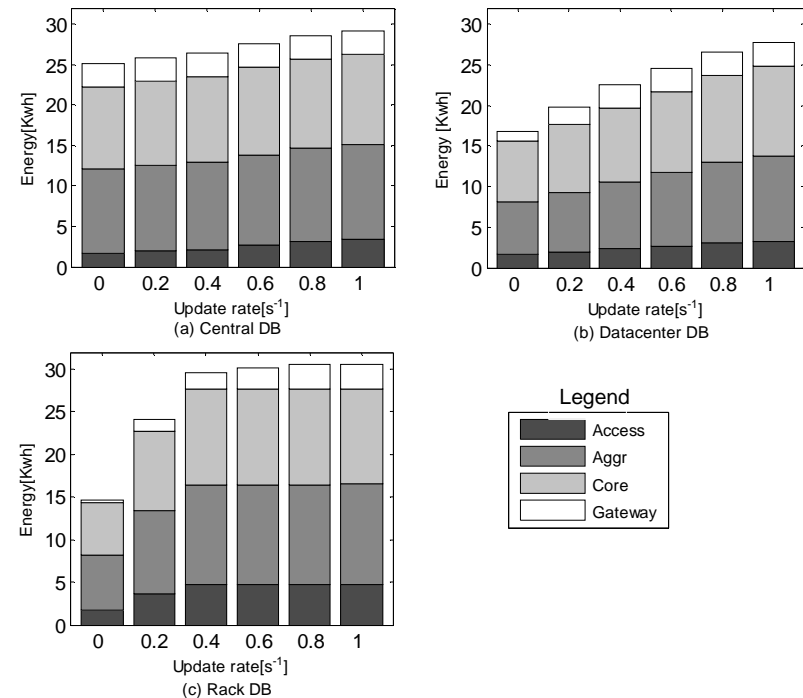
- Energy Consumption of Computing Servers
 - Increases with size of data items as communication delay is included into execution time preventing servers to enter sleep mode



Evaluation Results

- Energy Consumption of Network Switches


- Energy consumption increases with update rate
- Switches at all layers are involved in propagating database updates



Evaluation Results

- Access Delay

- Smaller delays when accessing replicated data
- Increase in data size leads to larger delays



Conclusions

- Data replication is important in geographically distributed cloud computing data centers
- Replicating data closer to data consumers, i.e., cloud applications, **reduces energy consumption, bandwidth usage, and communication delays significantly**
- Future work will be focused on developing a formal mathematical model and testbed implementation of the proposed solution

Thank you!

Contact information:

Dzmitry Kliazovich

University of Luxembourg

dzmitry.kliazovich@uni.lu

