

Databases in the cloud, ~~they must be different~~ will be

Gustavo Alonso

Systems Group

Department of Computer Science

ETH Zurich, Switzerland



Systems Group@ETH



www.systems.ethz.ch

What is the cloud?

Primarily a change in business model

-

A shift from manufacturing to services

What is big data?

Primarily a change in business model

-

A shift toward customization and personalization of “services”

What does it all mean?

- General purpose no longer the overall trend
- Many applications with enough value to justify customization
- Customization can be achieved by combining standard elements
- Customization supported by abundant surplus computing power

Specialization is the
name of the game
- also in the cloud


The role of hardware

- Oracle RAPID: High scale parallel data processor
- Microsoft Cypherbase: FPGA for encrypted database processing in the cloud
- Microsoft Catapult: FPGA acceleration for search tasks (page rank like algorithms)
- Intelligent storage systems: NetApp, Oracle, ...
- Intel HARP, IBM's CAPI, ...
- HPC has been doing this for years ...

How will databases be
different?

1. The Network
2. The Processing
3. The Processors

Multicore vs cluster



Scalability through more, smaller
machines rather than through large,
multi-socket machines

=

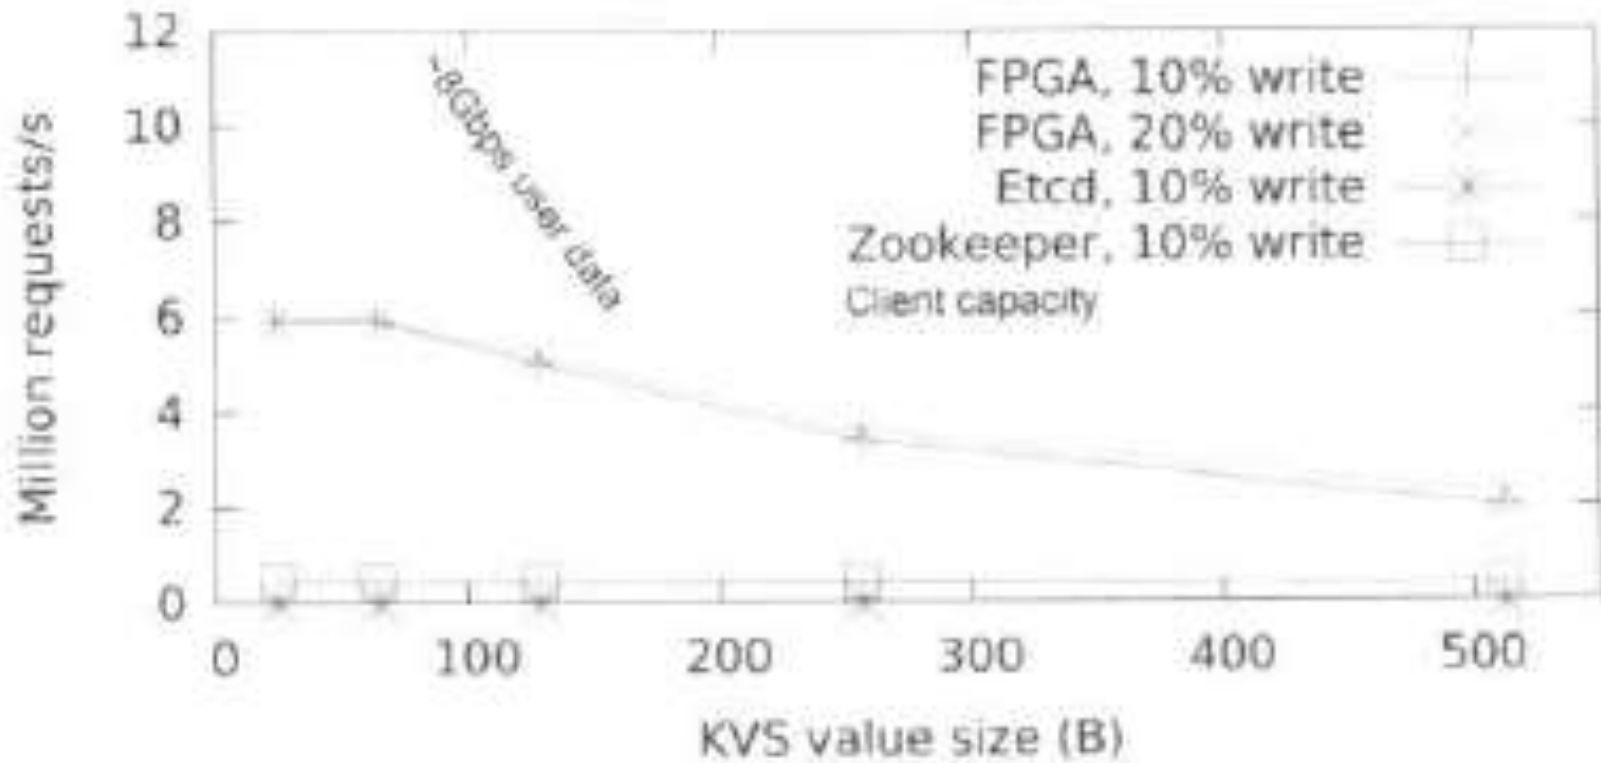
fast networks, RDMA

Short term (forget TCP)

Software Defined Networks

Optimized application level protocols for data centers

Proof?



Long term (forget the computer)

If the network latency and the bandwidth significantly increase: disaggregation

HP “The machine” based on optical networks

1. The Network
2. The Processing
3. The Processors

Data movement is bad


It costs energy

It takes time

Bandwidth bottlenecks (I/O, network, cache hierarchies, etc.)

Solution:

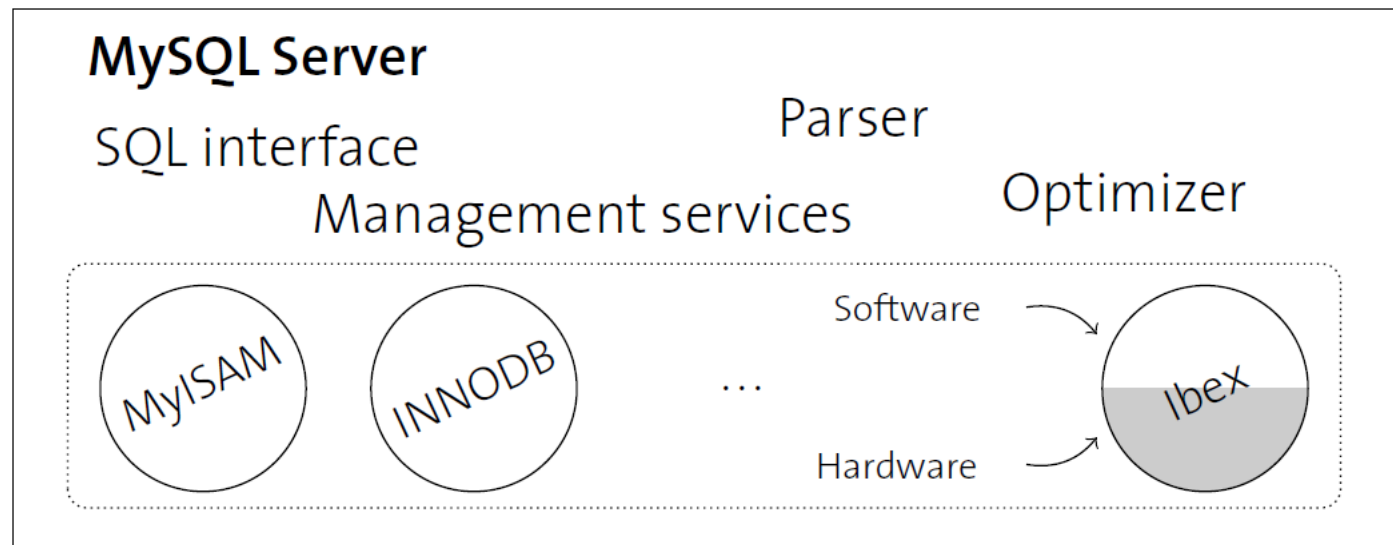

use hardware to process data
in place or as it flows



IBEX

(Woods, VLDB'14; Istvan, SIGMOD'14)

A processor on the data path



A variety of operators

CPU bound operators on accelerators

- Skyline (FCCM'13)
- Complex Event Detection (PVLDB'11)
- Histogram calculation (SIGMOD'14)
- Aggregation (PVLDB'14)
- Simple statistics (PVLDB'09)

Sounds good?

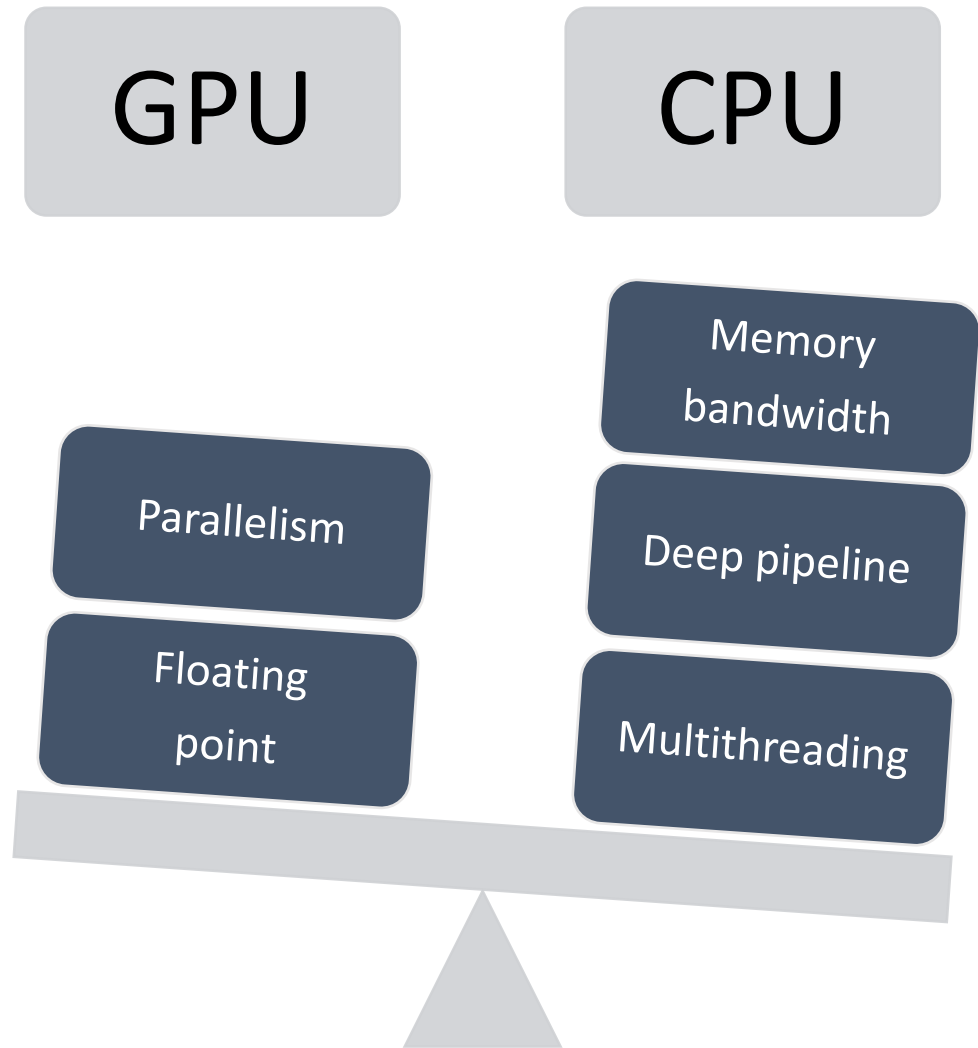
Imagine the same at all levels:

- Smart storage
- On the network switch (SDN like)
- On the network card (smart NIC)
- On the PCI express bus
- On the memory bus (active memory)

Every element in the system
(a computer rack)
will be a processing component

1. The Network
2. The Processing
3. The Processors

The architecture of co-processing



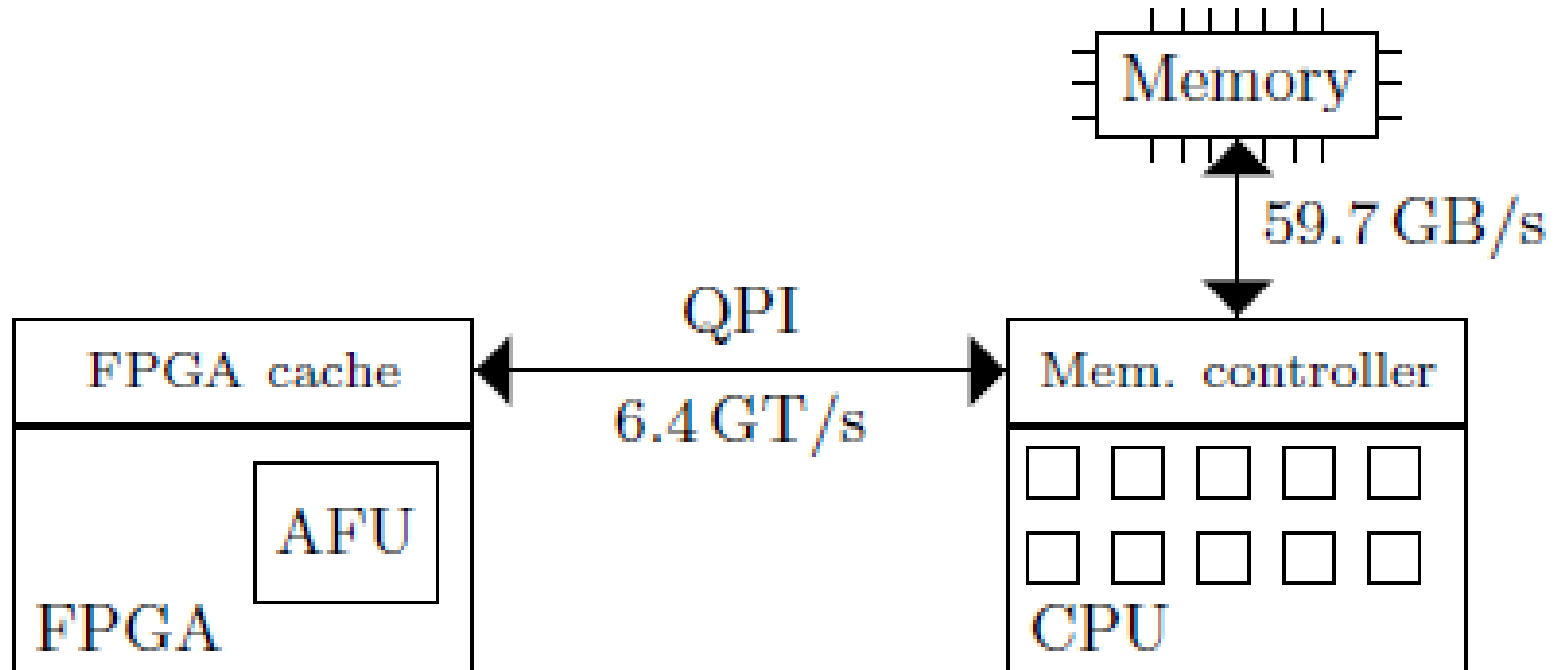
A processor far, far away ...

A co-processor (GPU, Xeon Phi, FPGA) on a PCI bus works only when:

- Load is computationally bound
- Load remains computationally bound on the co-processor
- The data movement cost is less than the performance gain in the co-processor vs the CPU

Relational operators are often memory bound ...

Heterogeneous multicore



INTEL HARP Program

DISCLAIMER: this is pre-production hardware and software, and may not reflect the performance of production or future systems.

User Defined ... Hardware

User Defined Functions extend the functionality of a database

Imagine the same but in hardware (extending performance or functionality)

Perform operations on relational data that no database has ever done before: Skylines, Monte-Carlo, pattern matching, clustering, complex text search, advanced statistics, learning, ...

The End

The agenda ahead of us

- Very interesting times
 - Many opportunities driven by hardware
 - Plenty of use cases justifying specialization
- Many challenges
 - Hardware changes affect the whole stack
 - How to program heterogeneous architectures
- Think beyond one machine
 - Racks, clusters, data centers
 - Pay attention to developments in computer architecture and networking