Using Control Theory to Reduce Disk Congestion Caused by Unpredictable I/O in Cloud Computing

Per3S Workshop - Thomas Collignon

28/05/2024

Qarnot infrastructure

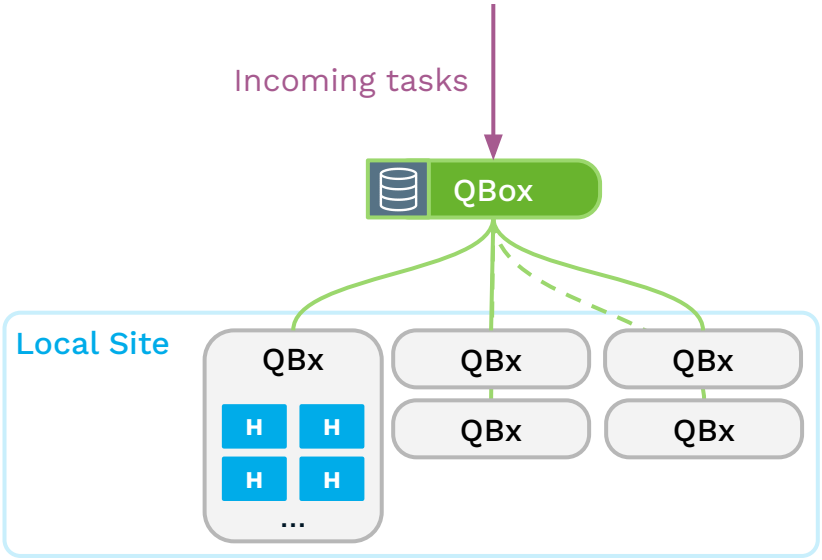

QBox

- Task distribution
- Download of tasks' input data
- Upload of results
- Shared storage for the tasks

I/O congestion on the QBox's storage

Background tasks

- Data traffic
  - Downloads
  - Uploads
- Task Checkpointing
- Cache handling

Computing tasks

- Varied I/O profiles

The performance of Computing Tasks can be degraded by I/Os interferences with other tasks.

How to improve the performances of computing tasks by controlling disk I/Os ?

Control Theory

- Autonomic Computing
- Actions on the system at runtime
- Supports disturbances (new computing tasks ...)

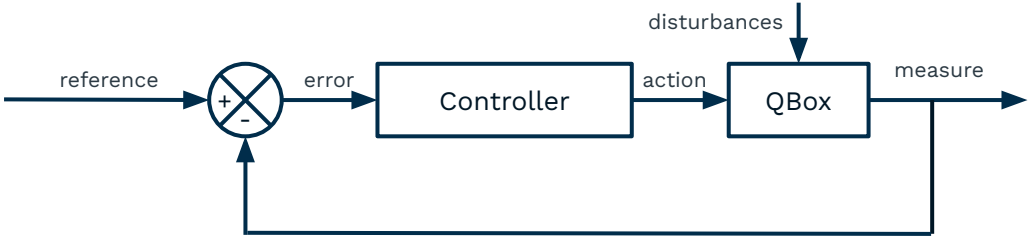

The I/O problem is inherently hard to predict so Control Theory is a good candidate to solve it at runtime.

Control strategies

Actuators for the selected problems :

Data traffic
- Bandwidth
- Delay

Cache Handling
- Garbage collector
- Cache strategy

Checkpointing
- Delay, but user constraints

Actuator for all the tasks:

cgroups
- Memory
- I/O bandwidth
- CPU

Using Control Theory to Reduce Disk Congestion Caused by Unpredictable I/O in Cloud Computing

Thomas Collignon

Thank you for your attention,

Let’s discuss at the poster session !