DREAM:Lab

SE252:Lecture 17, Mar 12

IL04: Application Execution Models on Cloud

Yogesh Simmhan

Summary of ILO3

Algorithms and Programming Patterns for Cloud Applications

- Application decomposition into Dataflow Task Graph
  - Analysis for degrees of parallelism
- Task, Data & Pipeline Parallelism
- Sync, Async, Blocking, Non-blocking
- Data Locality

Summary of ILO3

Algorithms and Programming Patterns for Cloud Applications

- MapReduce Dataflow Model
  - Big Data analytics for tuple-based data
  - Algorithm design for MapReduce
- Hadoop MapReduce & HDFS
  - Distributed, Data-local execution
  - Resilience
- Pregel/Giraph Vertex Centric Model
  - Big Data analytics for graph-based data
- Bulk Synchronous Parallel execution
  - Message passing

ILO 3

- Algorithms and Programming Patterns for Cloud Applications
  - *Examine* the design of task and data parallel distributed algorithms for Clouds and ✓
  - *use* them to construct Cloud applications. ➲Project
  - *Demonstrate* the use of task graphs and Map-Reduce programming model. ✓
  - *Apply* Amdahl's law and data locality principles to ✓
  - *analyse* and *characterize* the potential speedup of Cloud applications. ➲Project

DREAM:Lab

Lecture 16: ILO 4

ILO 4

- Application Execution Models on Clouds:
  - *Characterize* resource allocation strategies to leverage elasticity and heterogeneity of Cloud services
  - *Design* and *implement* Cloud applications that can scale up on a VM & out across multiple VMs ➲Project
  - *Illustrate* the use of load balancing techniques for stateful and stateless applications.
  - *Illustrate* the use of NoSQL Cloud storage for information storage and retrieval. ➲Project

Scheduling & Resource Allocation

- Map application to resources
  - Translate abstract application definition to their runtime execution
- Ensure goals are met
  - Reduce makespan of application
  - Increase resource utilization
  - Reduce total resource cost

Scheduling Jobs in a HPC Cluster

Table 2.7 Job Scheduling Issues and Schemes for Cluster Nodes

Issue	Scheme	Key Problems
Job priority	**Non-preemptive**	Delay of high-priority jobs
	Preemptive	Overhead, implementation
Resource required	**Static**	Load imbalance
	Dynamic	Overhead, implementation
Resource sharing	**Dedicated**	Poor utilization → *Not multi-tenant*
	Space sharing	Tiling, large job
Scheduling	**Time sharing**	Process-based job control with context switch overhead
	Independent	Severe slowdown
	Gang scheduling	Implementation difficulty
Competing with foreign (local) jobs	**Stay**	Local job slowdown
	Migrate	Migration threshold, Migration overhead

K. Hwang, G. Fox, and J. Dongarra, *Distributed and Cloud Computing: from Parallel Processing to the Internet of Things, Morgan Kauffmann Publishers*, 2011

But things are different in the Cloud

- Different types of applications
  - Stateless tasks
    - » Many independent tasks
    - » Typically, small resource needs, latency sensitive
  - Dataflows & Directed Acyclic Graphs (DAGs)
    - » Data dependency between tasks forces ordering
    - » Data exchanges
    - » Synchronization
  - Continuous dataflows
    - » Long running, latency sensitive, …

Cloud Resource Constraints: Distributed

- Operating across multiple data centres
  - hybrid private+public, multiple clouds, …
- Operating across Mobile+Cloud
- What to run where?
  - Cost-benefit trade-offs

Cloud Resource Constraints: Exclusivity

- Exclusive access to & control over VM
  - Limited by VMM sandboxing
  - *How about Docker containers?*
- Multi-core scaling on a single machine
  - Threads/tasks
  - Well suited for stateless tasks
- App framework has to manage local resource in a VM
  - May be better than many smaller VMs

Cloud Resource Constraints: Resource Sizing

- Custom resource sizing
  - Small, Medium, Large, … VMs
- Allows best fit for an application
  - But also requires more decisions
- Non-linear performance behaviour
  - See variability
- Resource non-locality for different sizes?
- Resource reuse

Cloud Resource Constraints: ₹eal Co$ts

- Trade-off between application value and resource cost
  - Not an abstract cost problem
- Pricing mechanisms
  - Hourly billing, Resource reuse
  - Asymmetry in bandwidth costs
  - Spot priced markets
- Intelligent use of local & Cloud resources

Cloud Resource Constraints: Elasticity

- Ability to acquire & release resources on-demand
  - Intelligent use of elasticity
  - Maximize hourly allocation
- Elasticity *not instantaneous*
  - VM startup time can be non-trivial
  - Need to plan ahead

Cloud Resource Constraints: Variability

- VMM sandboxing of multi-tenants not perfect
  - Changes over space and over time
  - I/O can be punitive, but CPU also impacted
- *Scale Up* out-performs Scale Out
  - Though price increase of both is similar
  - *t1.** have special characteristics
- VM placement can impact network bandwidth
  - Topology has to be inferred for N/W intensive workloads
  - Moving task to data vs. Data to task

Amazon Micro Instance

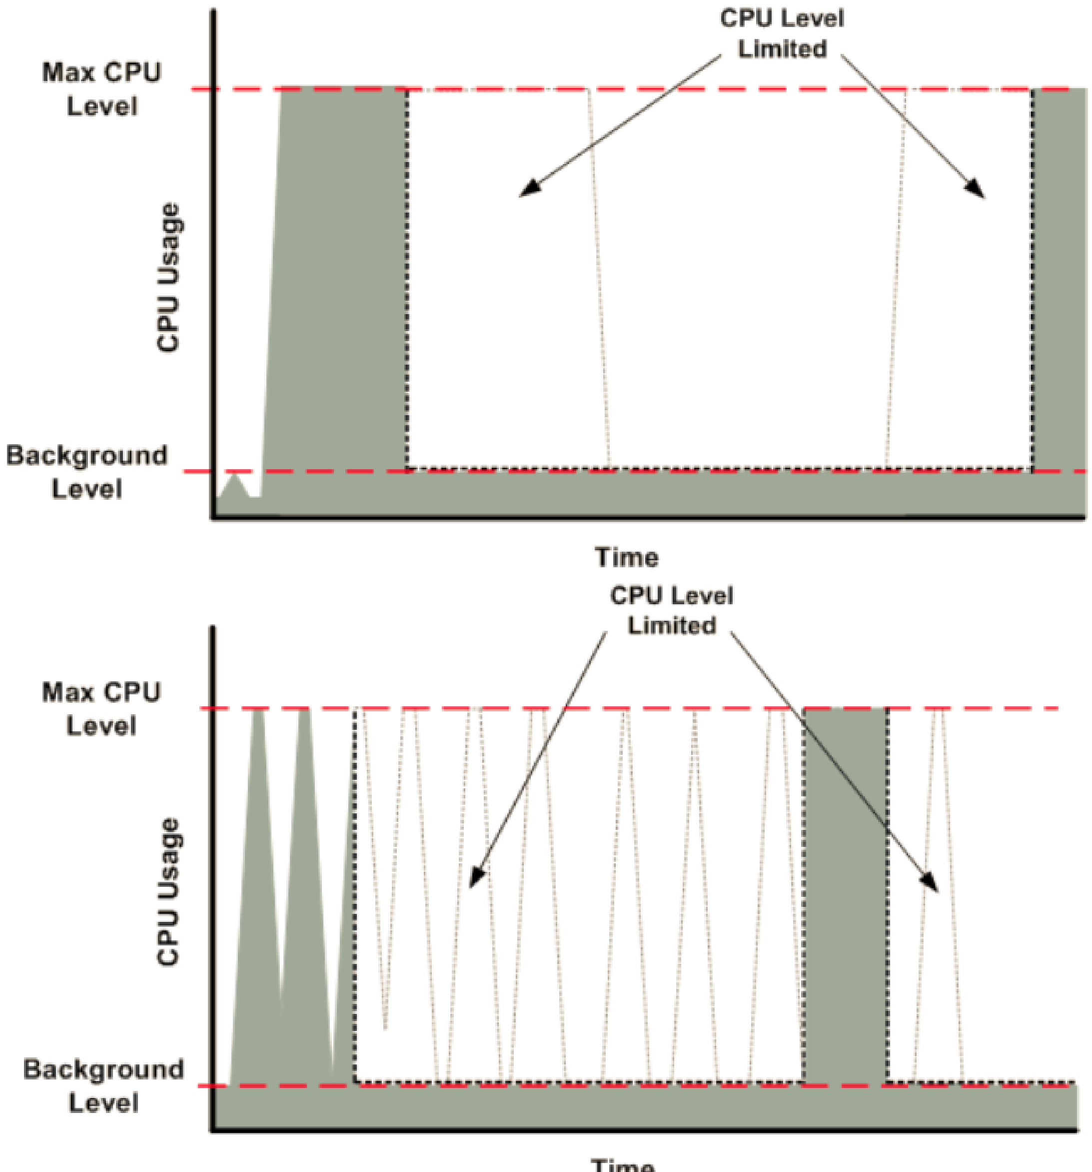

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts_micro_instances.html

VM Performance Variability

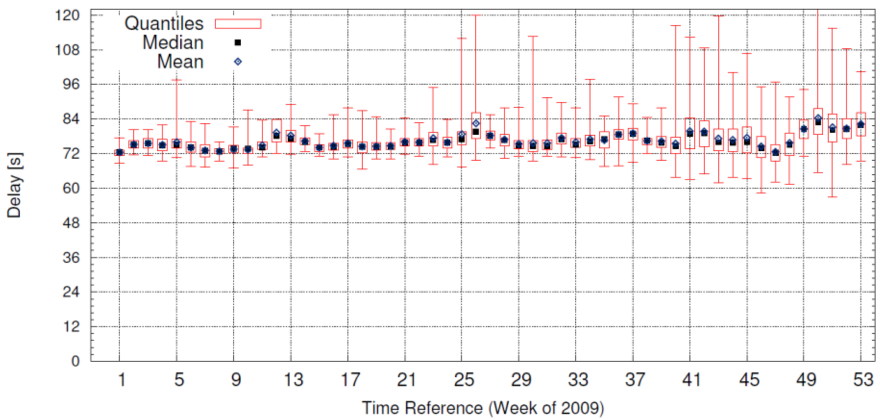

Amazon EC2 weekly statistics for resource acquisition.

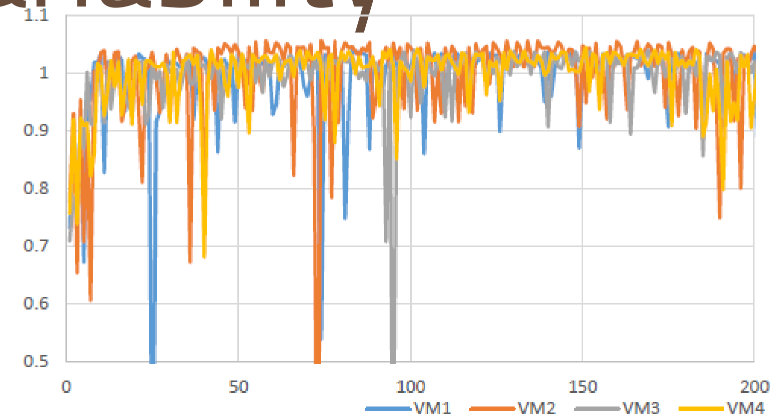

Figure 2: Variations in CPU performance

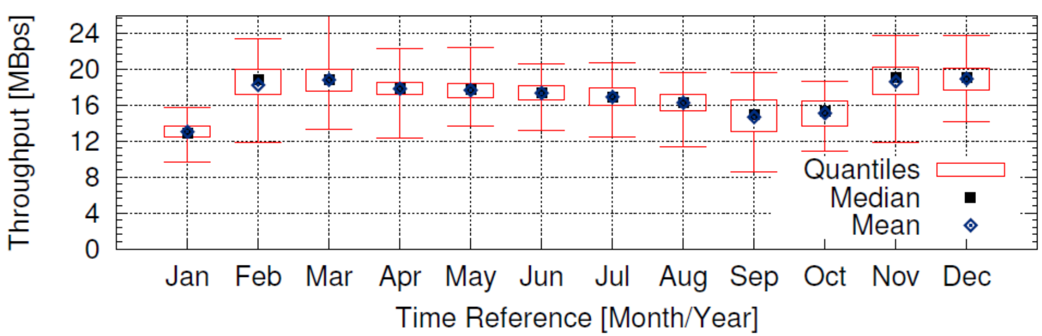

Amazon S3 monthly statistics of GET op.

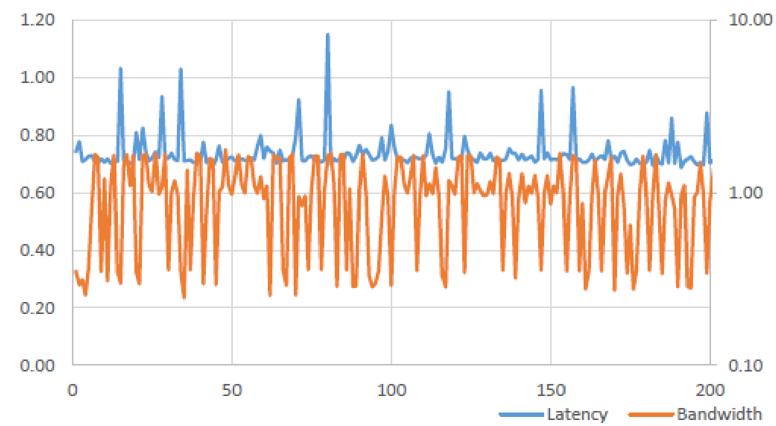

Figure 3: Variations in network performance

On the Performance Variability of Production Cloud Services, Iosup, A., *CCGrid*, 2011
Exploiting Application Dynamism and Cloud Elasticity for Continuous Dataflows, Kumbhare, *SC*, 2013

Cloud Resource Constraints: Reliability/Availability

- Reliable uptime of VM instances *not* guaranteed
  - Local data storage is transient (unless using EBS)
  - State has to be persisted for long running tasks
- VM availability is guaranteed
  - Ask and ye shall receive…*eventually*
  - Docker containers reduce this overhead
- Spot instances have further uptime constraints

Ongoing Assignments

- Project 2, due date extended to Mar 11
- Research report draft due Mar 11

Reading Assignment

- Text book Chs: 2.4.1, 2.4.2, 4.5.2, 6.2.6
- On the Performance Variability of Production Cloud Services, Iosup, A., Yigitbasi, N., and Epema, D., *IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, 2011