

Lian Cheng & Megh: Introduction to Automated Cloud

Debabrota Basu¹, Yang Hong²
Xiayang Wang², Haibo Chen², Stéphane Bressan¹

¹Department of Computer Science, School of Computing, National University of Singapore

²Institute of Parallel and Distributed Systems, Shanghai Jiao Tong University



Overview

Introduction to Clouds



```
graph TD; A[Introduction to Clouds] --> B[Lian Cheng]; B --> C[Megh : Automated Cloud]; C --> D[Discussion];
```

Lian Cheng

Megh : Automated Cloud

Discussion

INTRODUCTION TO CLOUDS

What's the Buzz?


The screenshot displays two news articles. The top article is from Forbes, titled "With Cloud Computing, The Sky's The Limit" by William Craig, dated October 22, 2015. A red box highlights the title. The article text includes: "There are a lot of technical buzzwords floating around the business world right now, but 'cloud' may still be one of the most misunderstood. The National Institute of Standards and Technology (NIST), the official..." and a quote: "Cloud computing is a model for configurable computing resources that can be rapidly provisioned and released. This cloud model promotes several models, and four deployment models..."

The bottom article is from Ars Technica, titled "The direction of computing is only going in one way—to the cloud" by Rupert Goodwins, dated November 14, 2015. A red box highlights the title. The article text includes: "After a rocky start, open-source and hybrid cloud initiatives have righted the ship." Below the article, a red box highlights a headline from another source: "North American Healthcare and Life Sciences Cloud Computing Market Worth 11.43 Billion USD by 2020".

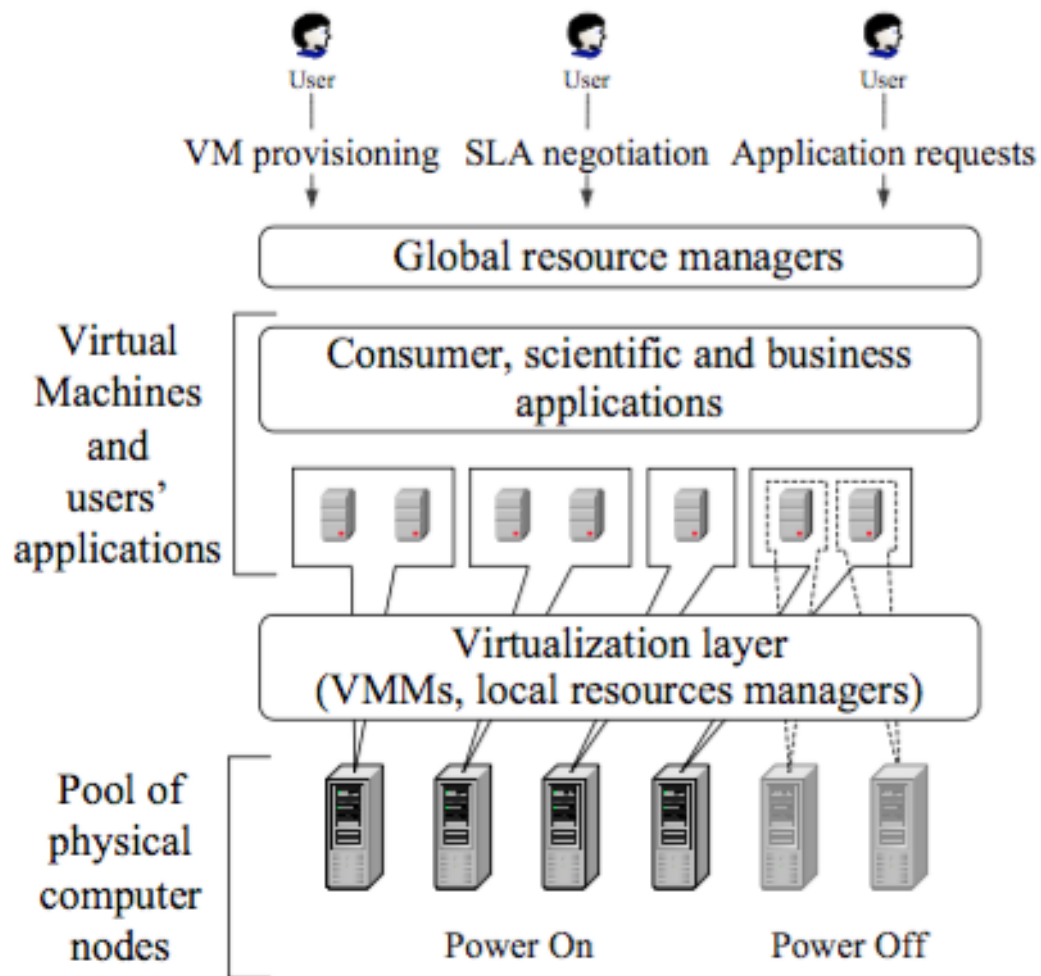
The bottom right section shows a news item from Korea Portal, titled "South Korea To Move 750 e-Government Services To Cloud Computing By 2017 For Improved and Cost-Efficient Public Service" by Therese Agcoppa, dated November 13, 2015. A red box highlights the title. The article text includes: "BY THERESE AGCOPRA / NOV 13, 2015 02:19 AM EST".

What is Cloud Computing?

- **Cloud computing** is a model that enables *on-demand network access* to a shared *pool of configurable computing resources* (e.g., networks, servers, storage, applications, and services)



What is in there?




Services

SaaS



Allows users to run their applications to a remote physical machine via internet.

PaaS



Provides hosting OS and optional building block services that allow users to run and build their own tailor-made applications.

IaaS



Outsources the computing resources including servers, networking, storage, and data center space on a pay-as-you-go basis.


Location

Public



Allows users to run their applications to a remote physical machine via internet.

Private



Provides hosting OS and optional building block services that allow users to run and build their own tailor-made applications.

Hybrid



Outsources the computing resources including servers, networking, storage, and data center space on a pay-as-you-go basis.

The Contract

- **SLA** (*Service Level Agreement*)

- Performance objectives
 - Availability
 - Response time
 - Capacity
- Security objectives
- Data management objectives



- **SLA violation**

- Power failure
- Hardware fault
- Oversubscription
- The demand of the CPU performance exceeds the available capacity

Pros

- Agile
- Cheap (Pay-as-you-go)
- Independent of Device and Location
- Scalable
- Flexible
- Low Maintenance

Cons

- Reliability
- Security
- Lock-in to service
- Regulations and legal issues
- Complex interfacing


LIAN CHENG

A Glimpse




- Lian Cheng: The cloud deployed by NUS and SJTU
- Features:
 - Open source
 - Hybrid deployment
 - Scientific computing friendly
 - Special workflow structure for queries

Demo



Demo









MEGH: AUTOMATED CLOUD







Challenges of Cloud Provider

- Managing and scheduling resources
- Satisfying increasing user basis
- Maintaining SLA
- Providing better performance

Before Migration



After Migration



Which VM to move?
When to move?
Where to move?



Issues

- But it takes migration time T_m to move and for that the performance is decremented for down time
- To keep SLA intact, T_d has to be as low as possible
- To eliminate disruption in running applications, its better to start it before the overflow or “LIVE migration” T_d
- Inefficient migration schedule causes higher expenditure of energy consumption


What if the system
learns to answer
these questions
from past experiences?

Megh: Automated Cloud Center

- We propose an automatic live migration strategy that learns from its experience
- The process of *live VM migration* is modelled as a *Markov decision process* (MDP)
- This tries to answer **THE *three questions*** by optimizing the cost function
- MEGH learns the stochastic nature of workloads to predict optimal schedule of migration

What is Reinforcement Learning?

- RL is more like LIFE!!
 - It is about an agent learning from its environment through interactions

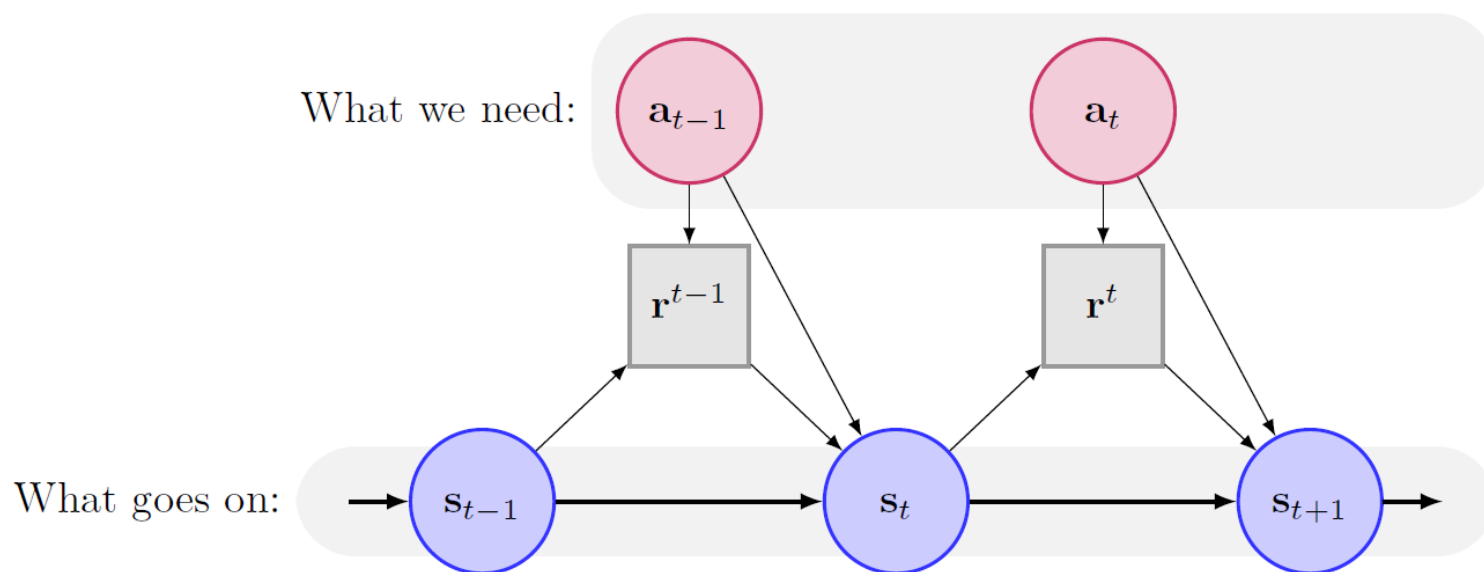


- Components:
 - State
 - Action
 - Reward
 - Utility

Markov Decision Process (MDP)

- **Markov assumption:** The future is independent of the past given the present



$$\mathbf{P}(s_{t+1} | s_t) = \mathbf{P}(s_{t+1} | s_0, s_1, \dots, s_t)$$



- Using Bellman optimality condition, the **best policy** can be given by

$$\pi^*(s) = \operatorname{argmax}_{a \in A} \left(R(s, a) + \gamma \sum_{s' \in S} T(s, a, s') V^*(s') \right)$$

Applications of RL



More Applications



CROC: The Hunter

Application in MEGH

- **Model:** MDP
- **State:** VM configurations on physical machines (PMs) of cloud
- **Action:** Configuration changes due to varying workload

- **Penalty function:** Cost of the action

$$\text{Cost}(t) = \int_{\tau=0}^t C_p(\tau) d\tau + C_v(t)$$

Cost of energy consumption Cost of SLA violation

- **Transition function:** Transition from one state to another on an action. It has to be learned
- **Policy:** A sequence of configuration changes depending on the workloads

A Few Observations

- A basic version of Megh is tested on Cloudfins dataset
- It is converging to give optimal policy before 100 iterations
- The workload variation can be modelled as an exponential family that gives us a sound mathematical structure

DISCUSSIONS

Conclusion

- Features of cloud computing raises it as the “new” paradigm of computing.
- But still there are challenges to solve.
- We are building up Lian Cheng as an open-source cloud system.
- Its unique GUI and workflow makes it user-friendly specially for large query processing.
- The automated system Megh shows an efficient way to provide better performance for this cloud.

Future works

- Providing Lian Cheng as a live cloud service soon.
- Deploying an efficient automated system that learns to handle its workloads and resources by itself.
- Solving more challenges of cloud computing using Lian Cheng.
- Developing reinforcement learning techniques to solve problems from different fields.

Thank you...