

# Teaching Cloud Computing: Motivations, Challenges and Tools

C. Anglano, **M. Canonico** and M. Guazzone  
massimo.canonico@uniupo.it  
University of Piemonte Orientale

# Outline

- Introduction
- Related works
- Challenges
- Syllabus
- Conclusions and future works

# Introducion

- Teaching Cloud Computing to computer science students (and not only) is important/necessary
  - Cloud computing is everywhere
  - A biological scientist may need the Genomic Analyses Service provided by AWS or the Databiology's platform on the IBM Cloud
- Provide a course with different paths based on students skills and interests
- Re-think how we teach cloud computing
  - Share educational materials in order to easily and exploit any collaborative tools in order to easily update the materials
  - provide tools to easily access to different cloud platforms in order to avoid the vendor lock-in

# Related works

- In our study we considers only free educational materials
  - Coursera, Qwiklabs are not consider
- Open Source Data Science Curriculum (OSDSC)
  - Gitlab and Stack Overflow
  - Last update 6 years ago
- Google Cloud Codelabs
  - More than 100 labs
  - Filter by topic, duration, date of publication (no about student level)
  - Only focused on Google Cloud Platform
- Personal educators websites
  - Resource very specific
  - No open-source nor collaboration mode

# Challenges

1. Find out a real testbed
2. Use different cloud platform to avoid lock-in issues
3. High-quality and wide-variety of educational materials ready to be exploited by educators and students

# Real testbed

- Cloud Computing testbed requires a (possibly big) set of computational resources geographically distributed worldwide.
- Commercial Cloud Platforms (IBM, Google, Amazon)
  - Free resources very poor
    - AWS provides a "free-tier" VM: 1 vCPU with 2.5GHz, 1 GB of system memory, 8 GB of storage memory
  - Good resource expensive for a classroom scenarios
    - Google Cloud Platform: a week with 5 VMs could cost 50 euros
- Solution: OpenStack provided by Chameleon project:
  - 15,000 CPU cores, 5PB of total disk space, hosted across two sites (University of Chicago and Texas Advanced Computing Center) free to use

# Different cloud platforms

```
$ aws ec2 run-instances \  
  --image-id ami-1a2b3c4d \  
  --count 1 \  
  --instance-type c3.large \  
  --key-name MyKeyPair \  
  --security-groups MySecurityGroup
```

(a)

```
$ openstack flavor create --ram 512 --disk 1 --vcpus 1 m1.tiny
```

(b)

```
$ gcloud compute instances create "my-new-instance" \  
  --zone="us-west1-b" \  
  --image-family="tf-latest-cu92" \  
  --image-project=deeplearning-platform-release \  
  --maintenance-policy=TERMINATE \  
  --accelerator="type=nvidia-tesla-v100,count=8" \  
  --machine-type="n1-standard-8" \  
  --boot-disk-size=120GB \  
  --metadata="install-nvidia-driver=True"
```

(c)

Fig. 1. How to run a new virtual machine by using the command-line client software provided by (a) AWS, (b) OpenStack and (c) GCP respectively.

```
***** EasyCloud *****  
  
Select a platform  
-----  
1) Amazon Web Services  
2) Google Cloud Platform  
3) OpenStack  
4) Close application  
  
Please make a choice:  
> █
```

Fig. 2. The EasyCloud main menu.

```
***** EasyCloud *****  
  
What would you like to do?  
-----  
1) Create new instance  
2) Show running instances  
3) Terminare instance  
4) Reboot instance  
5) Manage floating IPs  
6) Manage volumes  
7) Extra functions  
8) Start monitor  
9) Manage rules  
10) Edit configuration file  
11) Back to main manu  
  
Please make a choice:  
> |
```

Fig. 3. The EasyCloud main features.

# Educational materials

- Share all our educational materials
  - Chameleon users mailing list
  - Google cloud platform user mailing list
  - Google cloud faculty experts community → we are thinking about an open access book
  
- Receive comment/suggestion from both educators and students
  
- Edit access available to collaborators by Google docs and Google slides



# Syllabus

- The course is organized in two main paths:
  - the path for beginners
    - Introduction to cloud computing fundamental concepts
    - Students without any programming skills
  - the path for advanced users
    - Computer science students
    - Students with programming skills

# The beginner path

1. Introduction to Cloud Computing
2. Cloud Platforms
3. OpenStack and Chameleon project
  - a. Basic operation (start/stop VM), setup a LAMP example and manage volumes
4. Amazon Web Service
5. Google Cloud Platform
6. Cloud Federation
  - a. Cloud standardization, CloudMesh project and EasyCloud project
7. Beyond Cloud Computing
  - a. Containers, Fog/Edge computing and Femtocloud

# The advanced path

1. OpenStack and Chameleon project
  - a. Load balance, advanced usage of the command line tool, customization scripts
2. Amazon Web Service
  - a. advanced usage of the command line tool
3. Google Cloud Platform
  - a. The GCP's command client, a MapReduce example
4. Cloud federation
  - a. The Boto library, the Apache libcloud
5. Cloud as a platform
  - a. Introduction to Big Data, Data analytics, Stream analytics, Machine learning

# Conclusions and future works

- We provide educational materials and tools to help students and educators to teach cloud computing in a shared and a collaborative manner
- In the future
  - New cloud platforms
    - IBM Cloud
    - Microsoft Azure
  - We are video-recording lessons → MOOC
  - Collaborative tools
    - Moodle
    - MIT Star Platform
    - Piazza