McGill

# COMP 598: Cloud, Fogs, and Virtual Networks

January 2019


**Instructor:** Muthucumaru Maheswaran
**Email:** muthucumaru.maheswaran@mcgill.ca
**Tel:** 398 1465
**Office:** Room 754, McConnell Engineering Building

**Office hours:** MF 1:00-2:00. Appointments can be made for meetings at other times.
**Class:** MW 10:00-11:30 ENGTR 2110
**Credits:** 3
**Class web page:** MyCourses will be used for assignments
**TAs:** None

**Description:** This course covers cloud and fog computing. Fog computing is a variation of cloud, where the resources are pushed towards the edge so that they are accessible with smaller latencies and less prone to Internet-wide disruptions. Both cloud and fog computing rely heavily on different forms of networking because the client devices have to access the resources over the network and cloud resources themselves are interconnected by networks. Therefore, we quickly cover the bare essentials so that this course will be accessible for students without prior networking background. Cloud computing is deployed in variety of contexts from scientific computing that has high-performance batch tasks to telecommunications. In this case, we would focus on cloud computing in the context of telecommunication-centric applications. Also, the focus is more from a system-development point-of-view and less from actual application deployment. This means we will examine how clouds interact with networking and how some of the popular network paradigms such as software-defined networking is impacting the relationship between cloud computing and networking.

## Evaluation:

| | |
|---|---|
| Cloud and networking lab experiments | 20% |
| Midterm 1 | 20% |
| Midterm 2 | 20% |
| Project | 40% |
| **TOTAL** | **100%** |

1. **Lab experiments:** This component will be carried out using the GINI toolkit. We will have about 4 to 5 lab experiments covering legacy networking, software defined networking, and even network function virtualization (tentative). These experiments will be well defined in scope, where you are expected to carry out certain experiments using the GINI toolkit and submit a report and/or show a demo of the experiment.
2. **Midterms:** This component will be based on the lectures and the reading assignments which can include papers, technical report, or selected chapters from reference texts.
3. **Project:** In this component, you will develop a significant cloud framework. It could be something like Kubernetes or service function chaining infrastructure with cloud support. This will be semester long project with different milestones. A skeletal design will be provided, which is more like design requirements. Your grade will depend on how your solution addresses the requirements. Solutions that go beyond the basic requirements will receive bonus marks.

McGill

**Textbook:** *None*

**Reference texts:** **FG:** I. Foster and D. Gannon, *Cloud Computing for Science and Engineering,* MIT Press, 2017.
**M**: D. Marinescu, *Cloud Computing Theory and Practice*, Morgan Kaufmann, 2018.
**SNBP**: Sehgal, Naresh, Bhatt, Pramod Chandra P., *Cloud Computing Concepts and Practices*, Springer, 2018.
**FLRS:** C. Fehling, F. Leymann, R. Retter, W. Schupeck, *Cloud Computing Patterns,* Springer, 2014.
**CGR:** C. Cachin, R. Guerraoui, and L. Rodrigues, *Introduction to Reliable and Secure Distributed Programming*, Springer, 2011.
**BH:** L. A. Barroso and U. Holzle, *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines,* Morgan&Claypool Publishers, 2009.

**All of the above references are available for downloading through the McGill library.**

## Course Schedule (Topics coverage order can change)

| Week | Topics | Remarks |
|---|---|---|
| 1 | A whirlwind tour of networking and related concepts. | Chapter draft from *Applied Networking* |
| 2 | **Cloud concepts:** definitions, Berkeley view of clouds, high-level patterns of clouds, concept of elasticity, cloud ontology, cost of cloud, pricing, spot pricing. | Sections 3.5, 3.6, 4.2, 4.3-4.5 of FLRS |
| 3 | **Encapsulating computation for clouds:** multi-tenancy versus multi-processing, unique challenges of encapsulating computation, virtual machines, containers, comparison of the different encapsulation schemes, serverless computing. | |
| 4 | **Kubernetes – a closer look:** we study Kubernetes as an example of container-based cloud orchestration engine. Previous systems that motivated Kubernetes like Borg and Omega will be studied. If time permits, Kubernetes will be contrasted with Openstack. | |
| 5 | **Containers and virtual machines – in depth:** we will examine the anatomy of containers. We will walk through the design and implementation of a simple container system. We will also examine virtual machines in some depth and compare the two approaches. Some of the emerging systems like Kata containers will be discussed. | |
| 6 | **Virtual networks for cloud computing:** cloud computing needs networking at a massive scale and they are also quite dynamic. That is, virtual machines or containers can be created and destroyed at a very rapid rate. Also, cloud installations can be in different sites and pass through many middleboxes (e.g., firewalls). Many network virtualization schemes have been | |

| | | |
|---|---|---|
| | developed and we will discuss few of them. | |
| 7 | **Software defined networks:** we will describe SDN in detail and examine the connection between SDN and cloud computing. In particular, how SDN can improve virtual network at cloud scale will be examined. Wide-area network optimization using SDN for cloud computing will also be discussed. | |
| 8 | **Network function virtualization:** we will examine the virtual network function concept, service function chains will be introduced, the problems that can map into service function chains and the challenges associated with them will be described. In particular, how NFV interacts with cloud computing will be examined. | |
| 9 | **Programming models and frameworks for clouds:** we will describe various programming models starting from the SPMD (single program multiple data) model. The SPMD model will be examined in the context of high-performance style parallelism. Many task parallelism, map-reduce, bulk synchronous parallel, and graph flow execution will be examined. Finally, agents and micro-services will also be examined. | |
| 10 | **Resource scheduling in clouds:** Many cloud-centric resource management algorithms with the associated fairness and performance issues will be examined. | |
| 11 | **Data analytics in clouds:** Performing data analytics in clouds will be examined. In particular, Apache Spark will be examined in this context. | |
| 12 | **Streaming data in clouds:** Streaming data challenges in cloud will be discussed. Example deployments will be discussed. | |
| 13 | **Mobile, fog, vehicular computing and clouds;** latency, availability issues; fast hand-off related issues; programming challenges of complicated cloud infrastructures | |

## Additional Reading List

Will be posted as needed.