Algorithms and Software Systems for Resource Management in Clouds

Adel Nadjaran Toosi

Faculty of Information Technology
Monash University

Email: adel.n.toosi@monash.edu
Homepage: http://adelnadjarantoosi.info

Outline

- Brief Biography
- Geographical Load Balancing (GLB)
- Resource Provisioning for Data-intensive Applications on Hybrid Clouds
- A Low-Cost Micro Data Center for Software-Defined Cloud Computing
- Summary

Biography and Research Overview

- PhD, University of Melbourne, 2010-2014
  - Thesis: “On the Economics of Infrastructure as a Service Cloud Providers: Pricing, Markets, and Profit Maximisation”
- Postdoctoral Research Fellow, University of Melbourne, 2014-2018
  - Renewable Energy, Data intensive application Scheduling, Software-defined Clouds
- Lecturer, Faculty of Information Technology, Monash University, May 2018
- Research Interests
  - **Distributed Systems**, **Cloud Computing**, Software-Defined Networking (**SDN**) and Network Function Virtualization (**NFV**), **Energy Efficiency and Green Computing**, **Soft Computing**
- Publications
  - **29** publications, 17 Journal Articles (**11 A/A*** ERA Ranking, ACM CSUR,TCC, JCNA, FGCS, TAAS), **11** Conference papers (CloudCom, UCC, HPCC), **1** Book Chapter,
  - h-index: **16** and **1200+** citations (SRC: Google Scholar)

Geographical Load Balancing for Web Applications

Cloud Computing

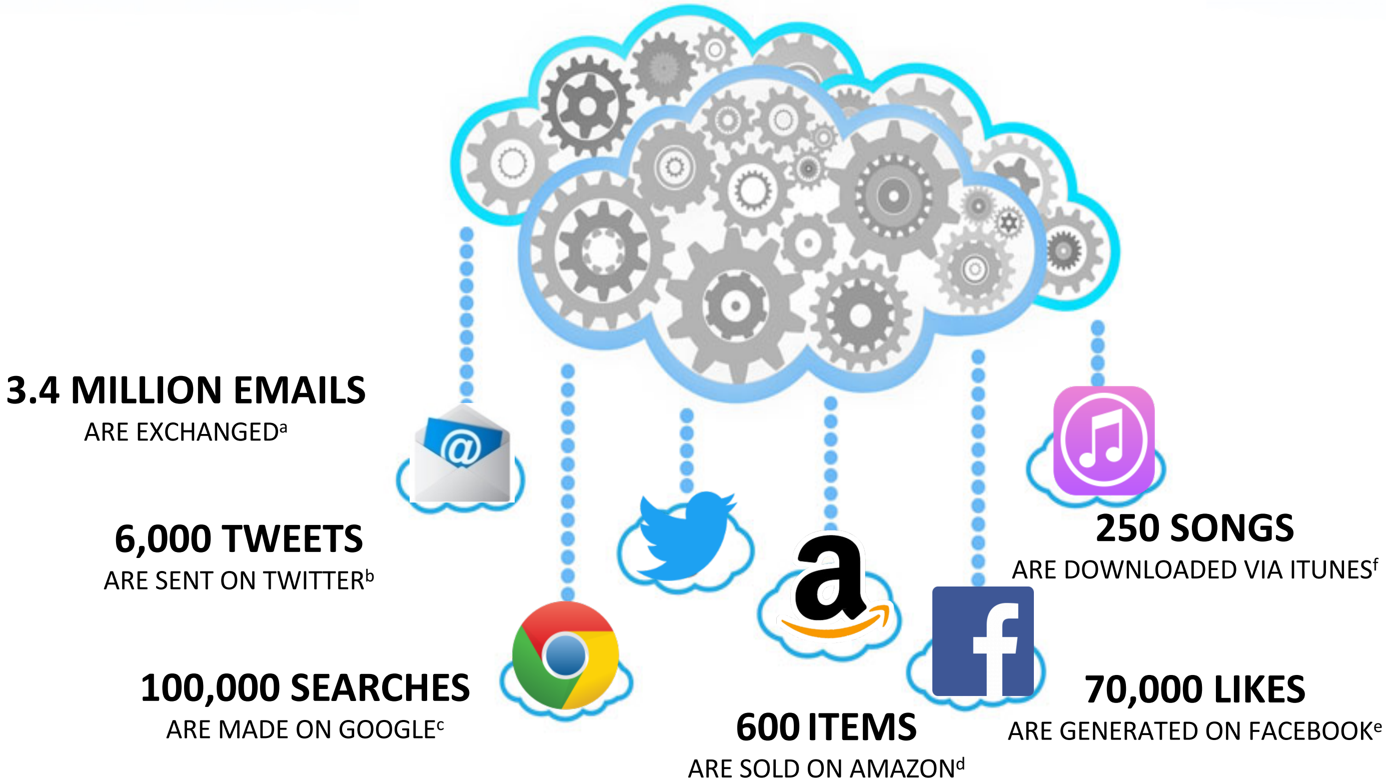

a https://www.lifewire.com/how-many-emails-are-sent-every-day-1171210
b http://www.internetlivestats.com/twitter-statistics/
c http://www.statisticbrain.com/google-searches/
d https://www.inc.com/tom-popomaronis/amazon-just-eclipsed-records-selling-over-600-items-per-second.html
e https://www.brandwatch.com/blog/47-facebook-statistics-2016/
f http://www.billboard.com/biz/articles/news/1538108/itunes-crosses-25-billion-songs-sold-now-sells-21-million-songs-a-day

Power Hungry Clouds

- Cloud data centres consume large amounts of electricity
  - High **operational cost** for the cloud providers
  - High **carbon footprint** on the environment
- US Data Centres
  - 70 billion kilowatt-hours of electricity in 2014
  - = Two-year power consumption of all households in **New York**
  - = The amount consumed by about **6.4 million** average American homes that year
  - Projected nearly **50 million tons of carbon** pollution per annum in 2020.
    - *Source: US Natural Resources Defense Council (NRDC)*

Renewable Energy and Challenges

- Cloud providers aims
  - Reduce energy consumption
  - Abate dependence on brown energy
- Renewable energy
  - Google, Microsoft and Amazon
- Challenges:
  - Non-dispatchable, Intermittent and Unpredictable
  - Powering data centres entirely with renewable energy sources is difficult
- Mixed sources of energy for data centres:
  - Grid power or brown energy
  - Renewable energy sources or green energy
- Challenges:
  - Minimising brown energy usage
  - Maximising renewable energy utilisation

Source: https://aws.amazon.com/about-aws/sustainability/

Geographical Load Balancing (GLB)

- Geographical load balancing (GLB) potentials:
  - **Follow-the-renewables**
- GLB approach benefits cloud providers but it raises an interesting, and challenging question:

*"With limited or even **no a priori knowledge** of the future workload and **Dynamic** and **unpredictable** nature of renewable energy sources, how to optimise the **overall renewable energy use** and **cost**?"*

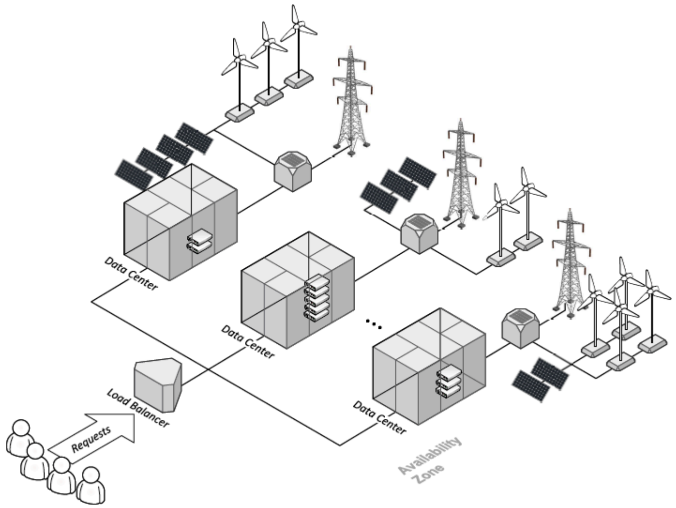

System Architecture for Web Applications

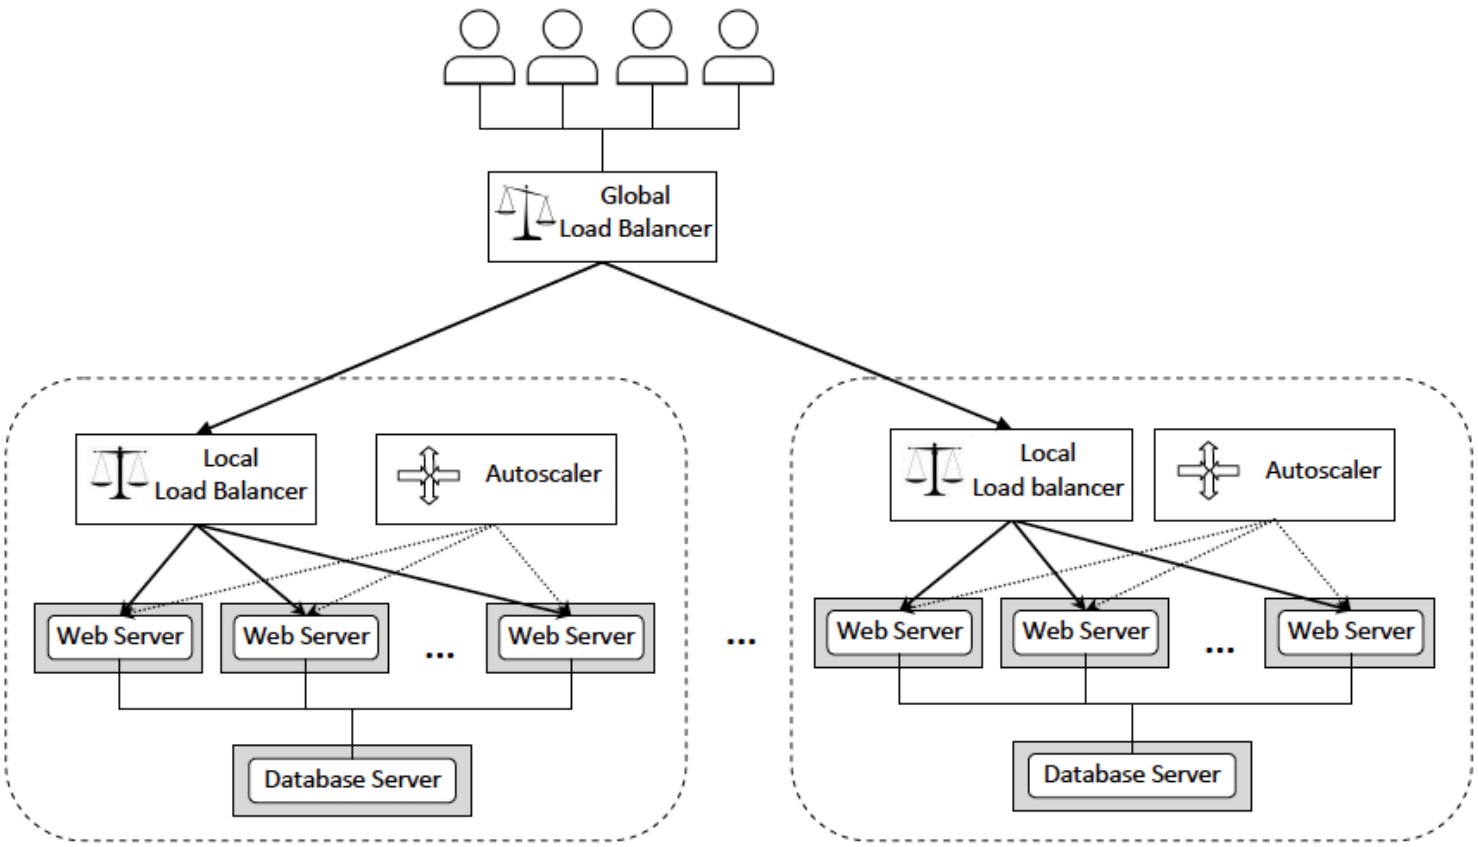

Grid'5000 Testbed

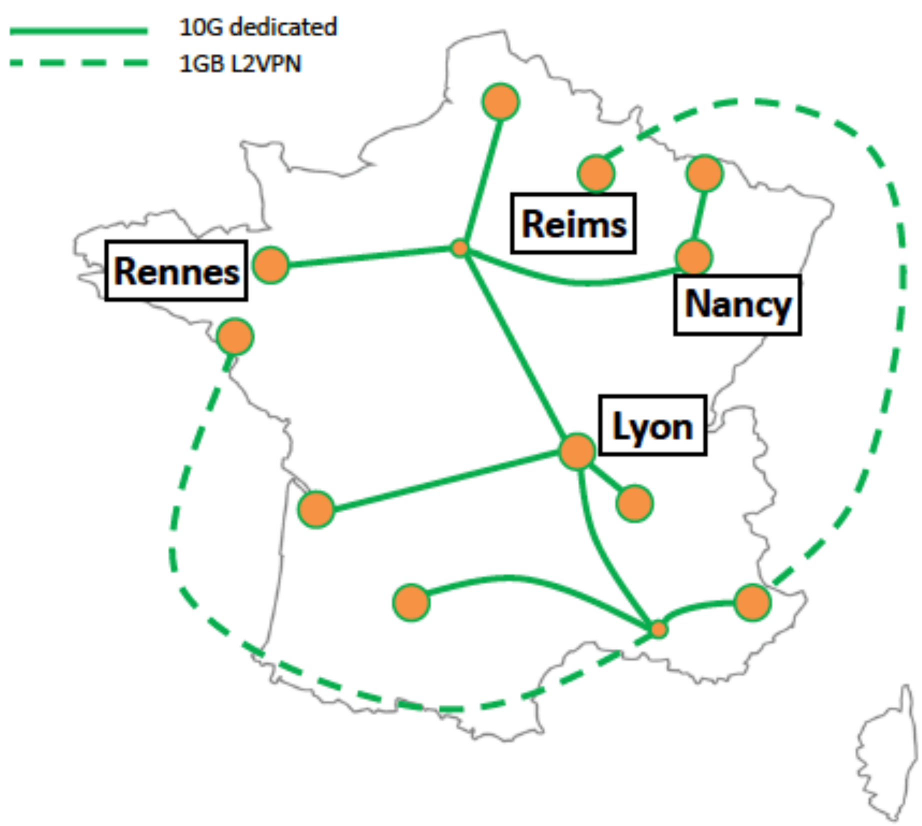

Workload Traces

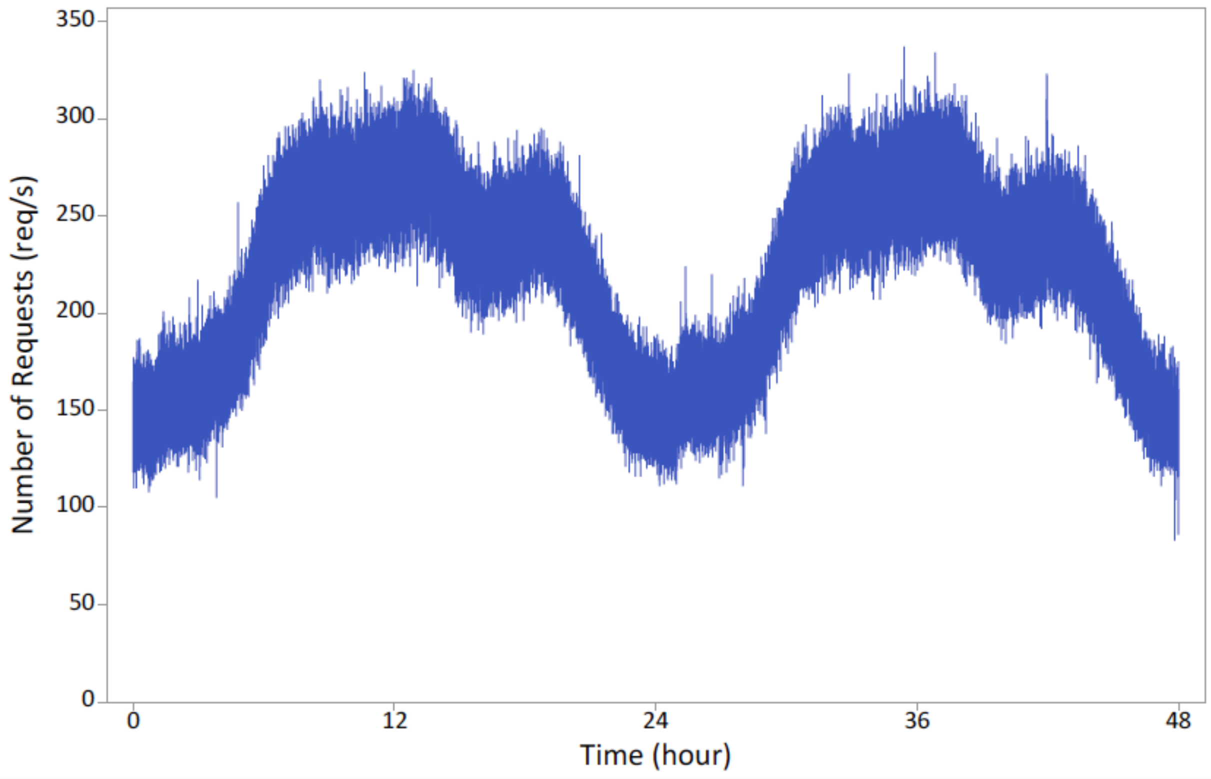

Wikipedia

A Prototype System

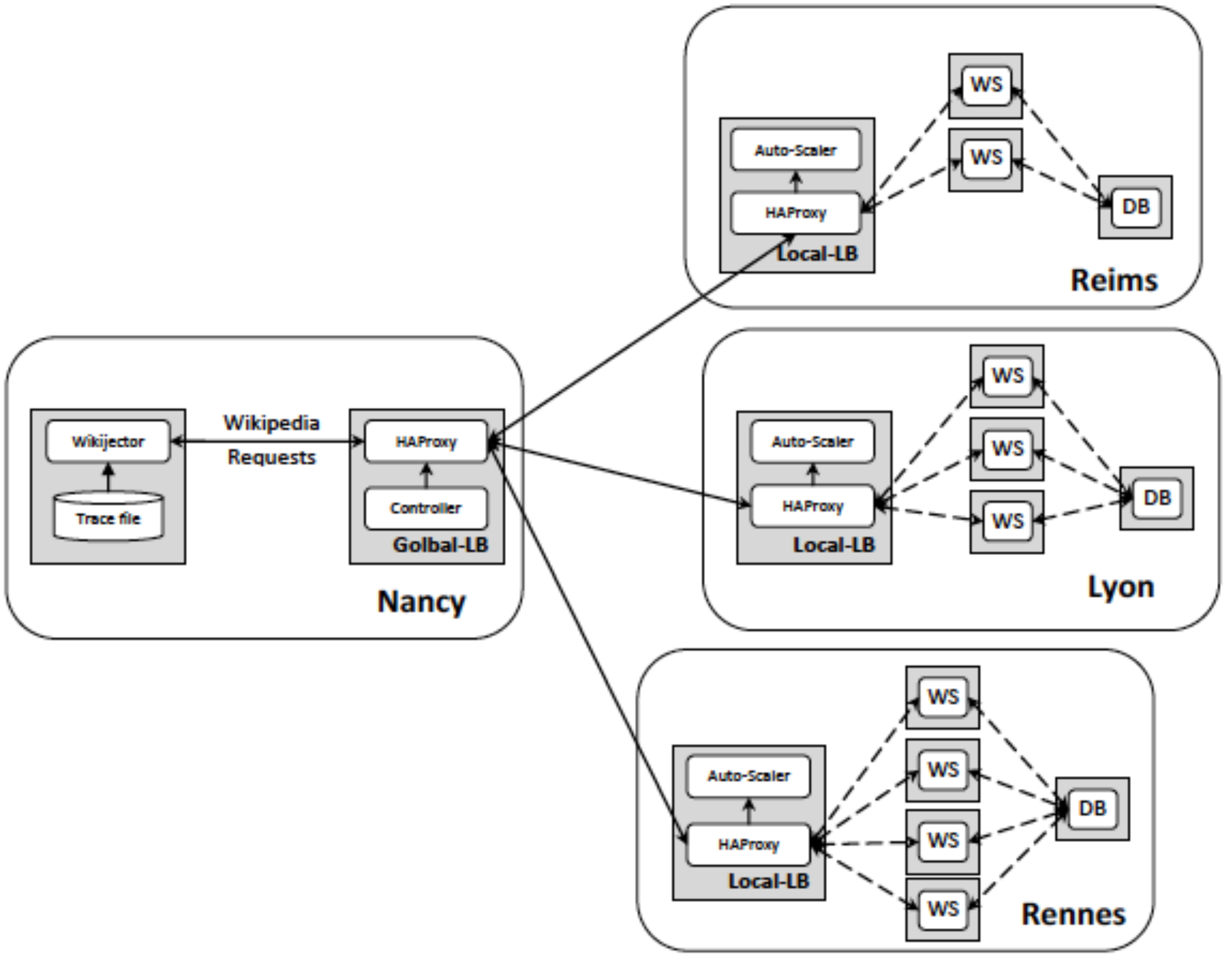

Renewable Power and Electricity Prices

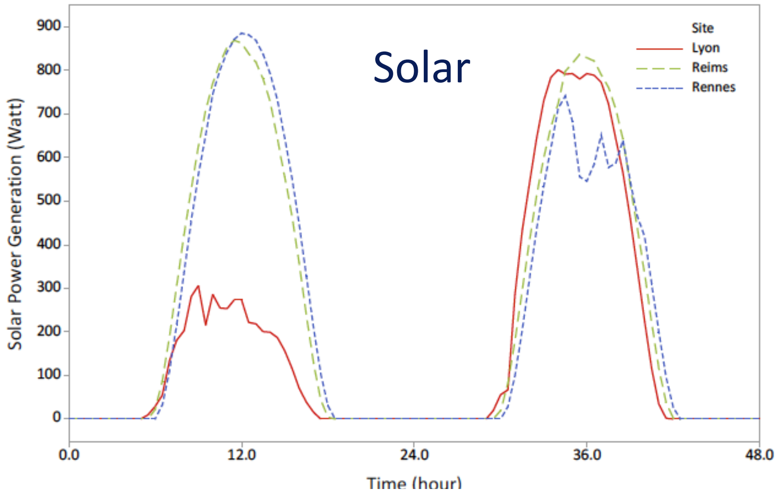

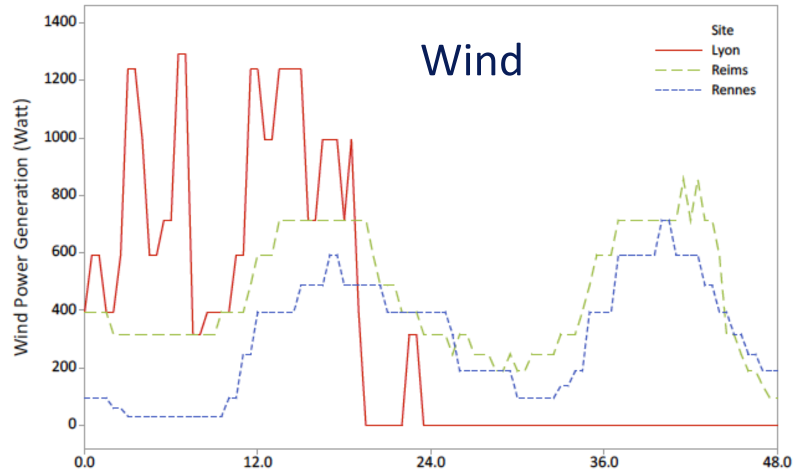

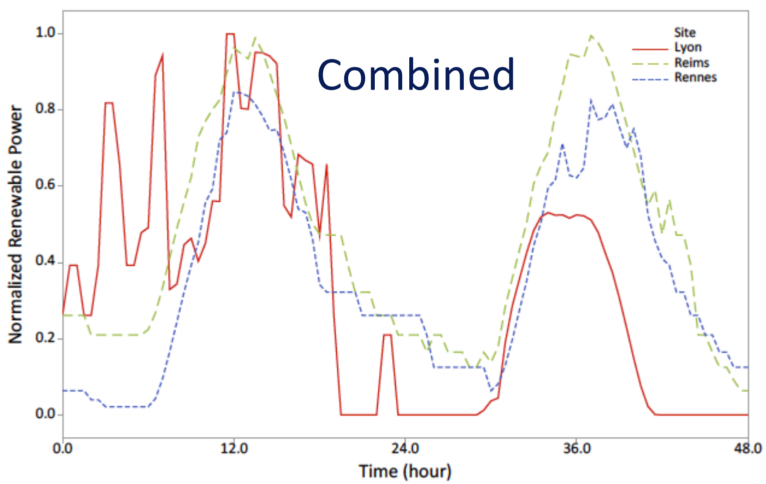

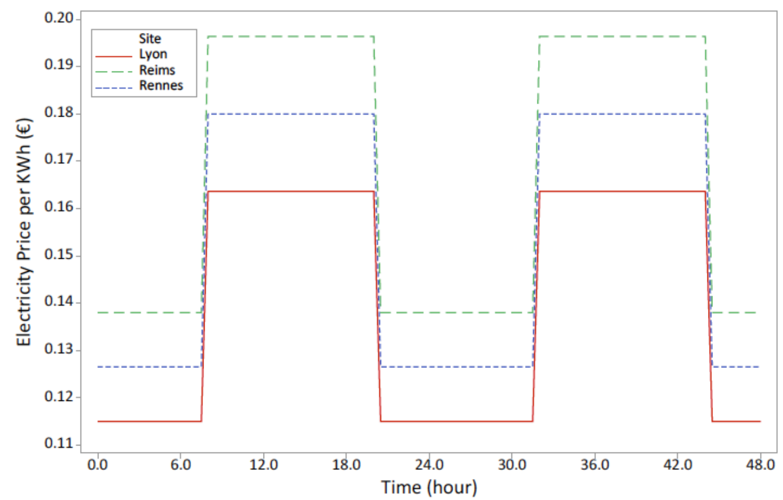

Results

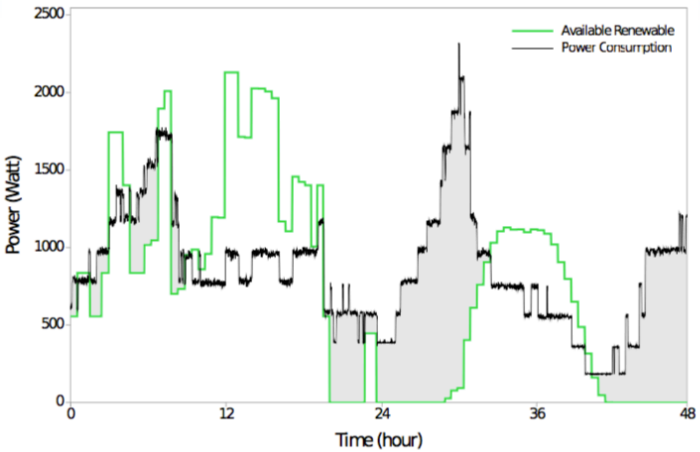

Lyon

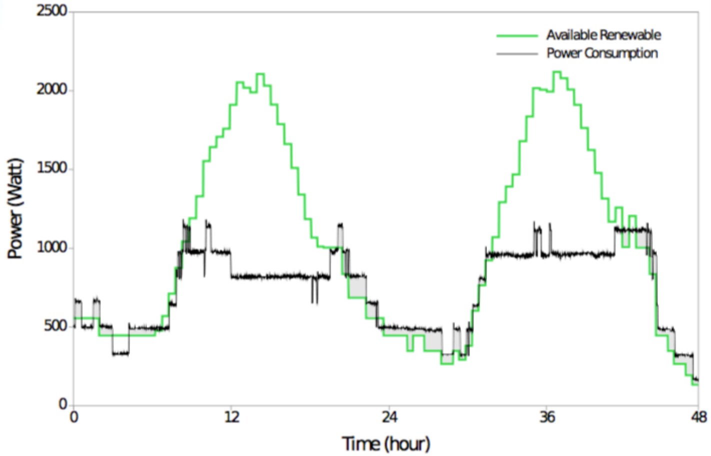

Reims

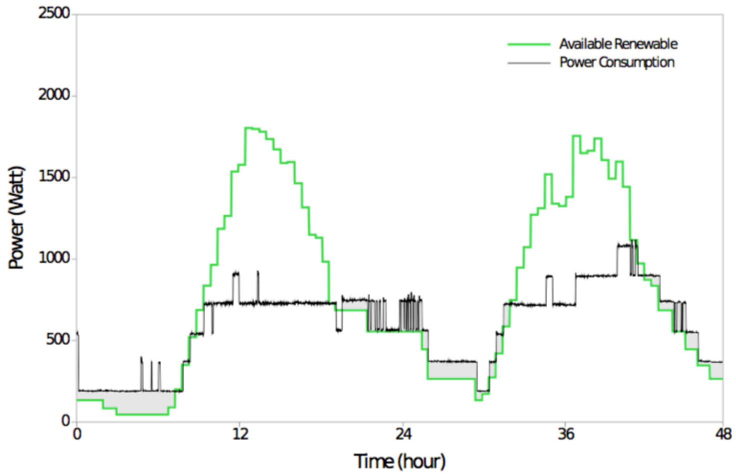

Rennes

Results

Site	Metric	RR	Capping	GreenLB
Lyon	Power Consumption (kWh)	36.3	42.9	**41.2**
	Brown Consumption (kWh)	13.3	19.0	**16.9**
	Cost (€)	1.71	2.31	**2.01**
Reims	Power Consumption (kWh)	32.5	[illegible]	[illegible]
	Brown Consumption (kWh)	2.1	[illegible]	[illegible]
	Cost (€)	0.42	[illegible]	[illegible]
Rennes	Power Consumption (kWh)	36.4	[illegible]	[illegible]
	Brown Consumption (kWh)	9.3	2.9	[illegible]
	Cost (€)	1.23	0.39	**0.35**
Total	**Power Consumption (kWh)**	**105**	**105**	**105**
	Brown Consumption (kWh)	**25.7**	**23.0**	**21.4**
	Cost (€)	**3.36**	**2.85**	**2.63**

Brown Energy:
17% and 7%
Cost Saving:
22% and 8%

Resource Provisioning for Data-intensive Applications on Hybrid Clouds

Background

- Data-intensive applications
  - Analysis of large datasets
  - Explosive growth of data
    - Smart cities, Social networks, Internet of Things (IoT), …
- Cloud computing
  - Preferred platform
- Common Scenario
  - Data is available in local IT infrastructure with limited processing capacity
- Cloud bursting
  - Hybrid Cloud (PaaS, Middleware)

Scheduling Problem

- Locality
  - Location of the data relative to the available computational resources
- Network bandwidth
  - Can become the bottleneck
- Data transfer
  - Not ideal to move the entire data set to the public cloud
- Data-intensive application
  - Data transfer time to the external cloud is often comparable to the computational time

Our Contribution

- Data-aware provisioning and Scheduling algorithm
  - Minimising cost while meeting the deadline requirements of applications
  - Hybrid cloud environments.
  - Data transfer time, available bandwidth, locality
- Plugged into PaaS
  - Aneka platform
  - Support dynamic resource provisioning for Microsoft Azure
- Experiments in actual hybrid cloud environment
  - Local resources and Azure virtual machines
  - Compared with existing approaches
  - A real-world case study
    - A data-intensive application in the smart cities context

Aneka

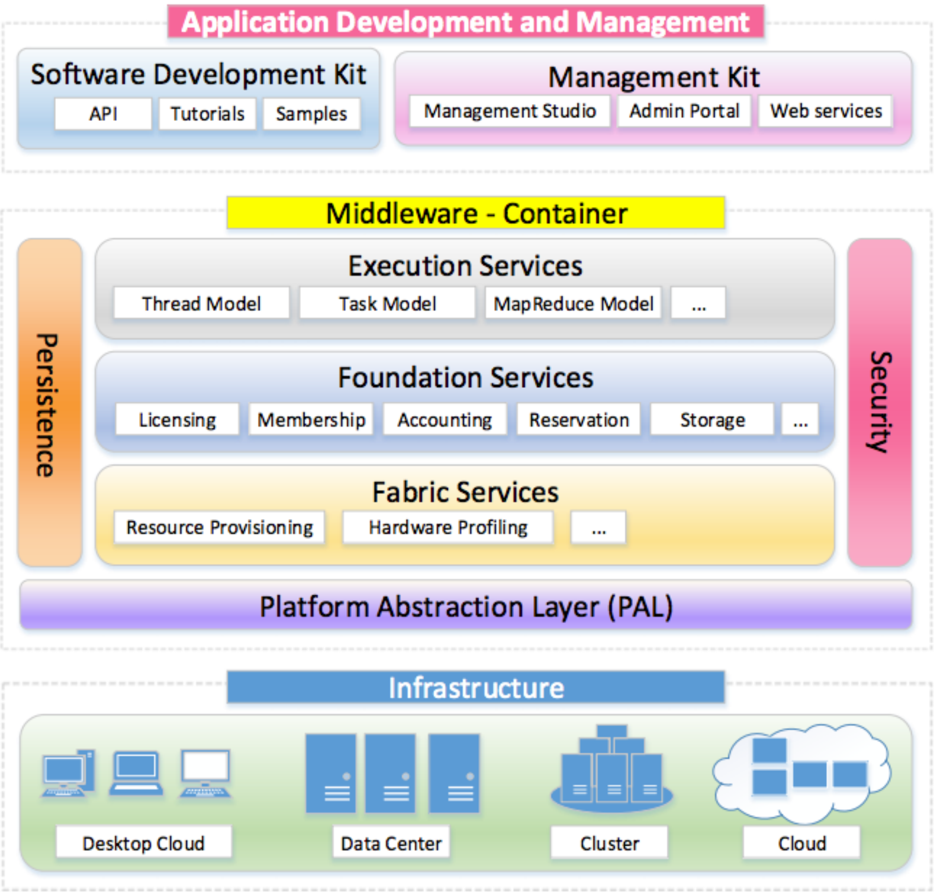

Application Development and Management
Software Development Kit
API
Tutorials
Samples
Management Kit
Management Studio
Admin Portal
Web services
Middleware - Container
Persistence
Execution Services
Thread Model
Task Model
MapReduce Model
...
Foundation Services
Licensing
Membership
Accounting
Reservation
Storage
...
Fabric Services
Resource Provisioning
Hardware Profiling
...
Security
Platform Abstraction Layer (PAL)
Infrastructure
Desktop Cloud
Data Center
Cluster
Cloud

Walkability Index – Melbourne Neighbourhoods

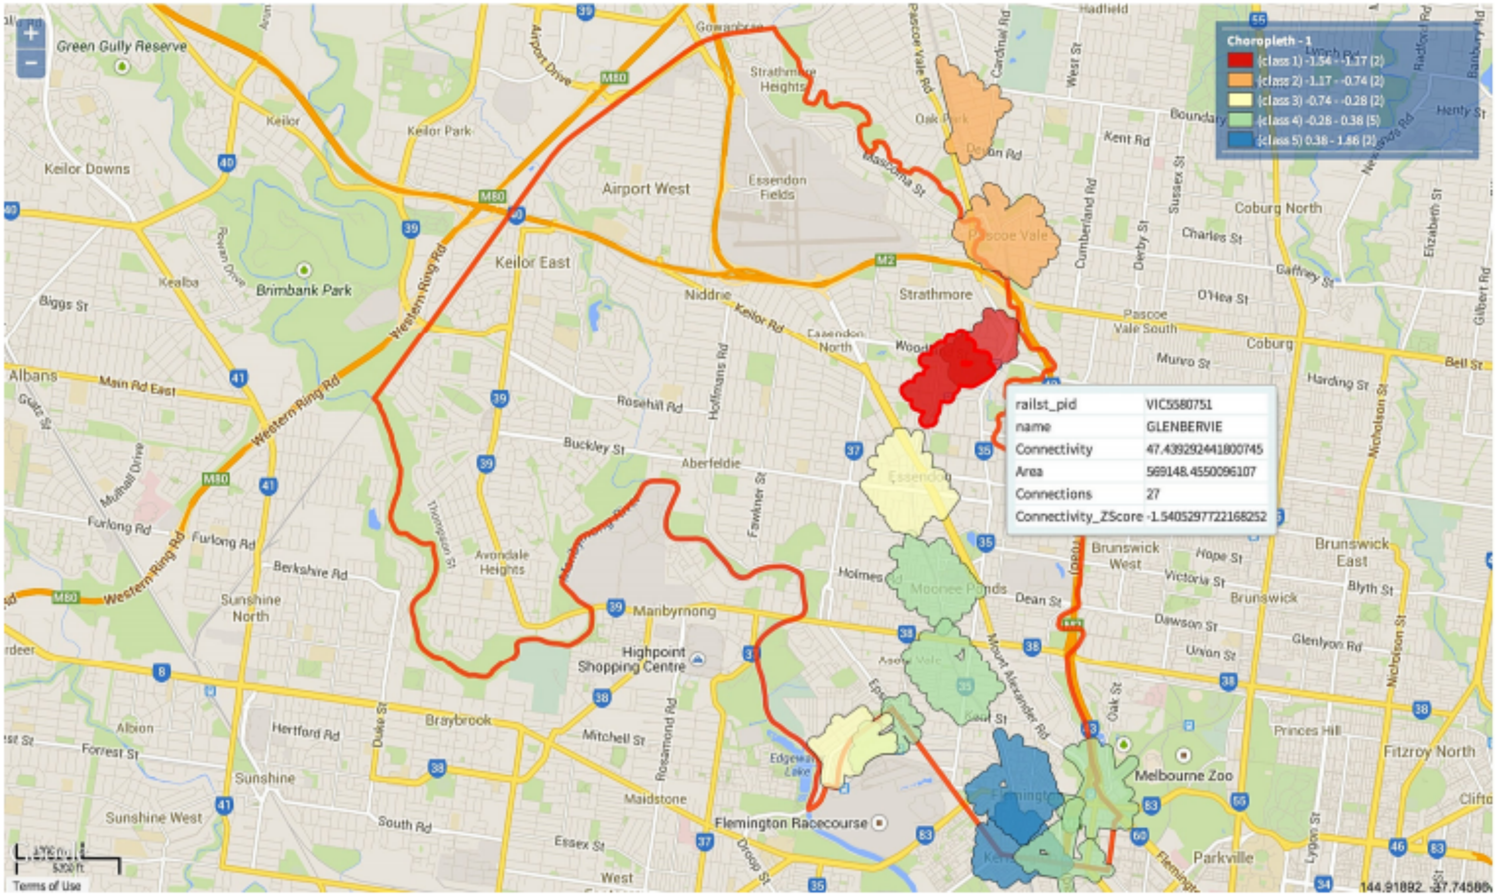

Hybrid Cloud Testbed

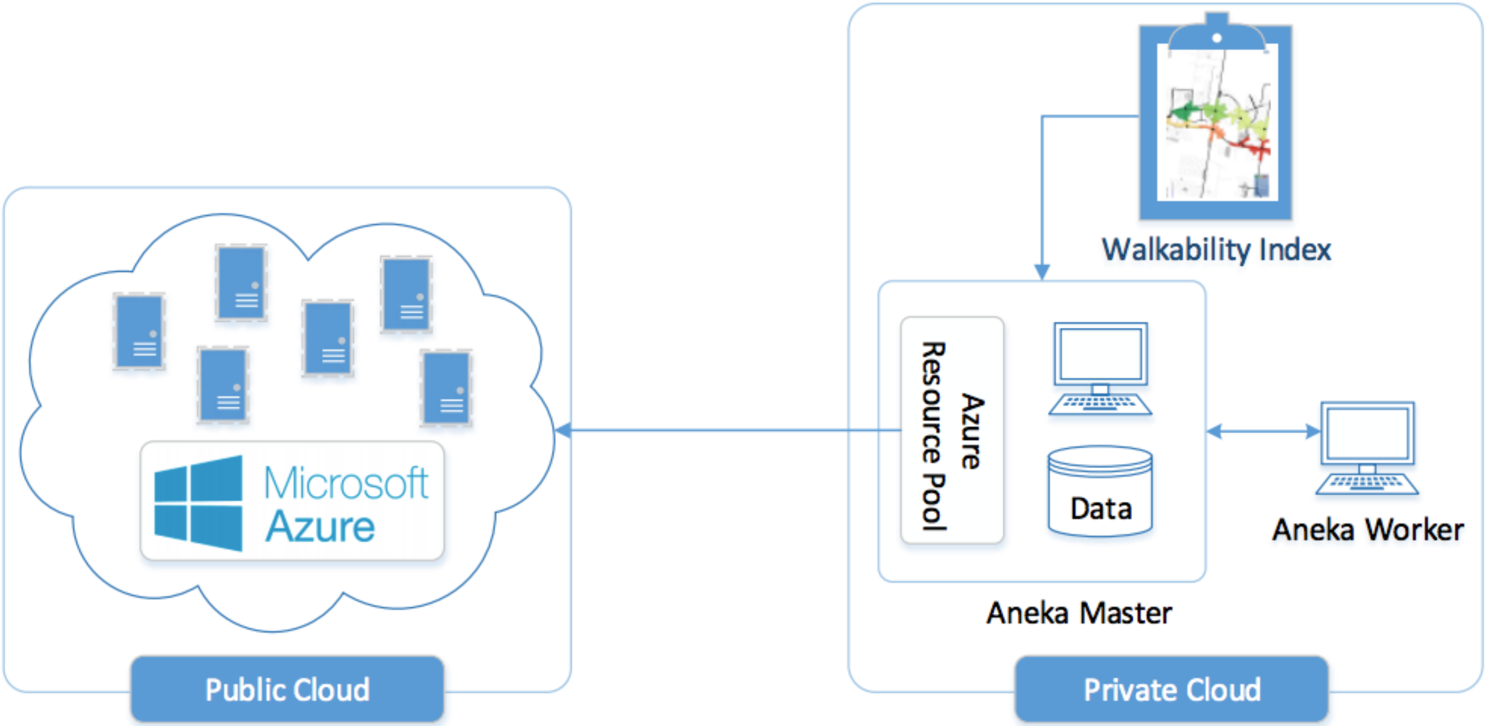

Some Experimental Results

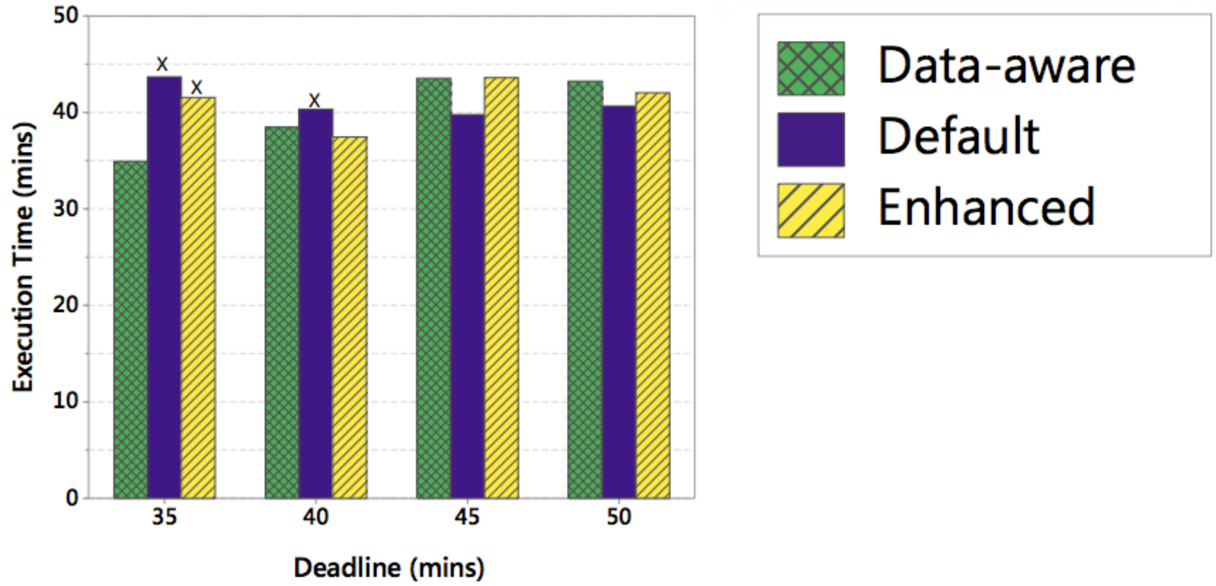

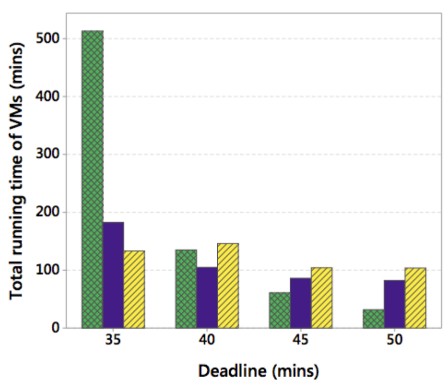

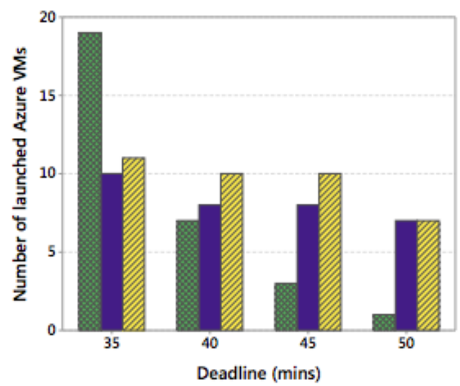

A Low-Cost Micro Data Center for Software-Defined Cloud Computing

Software-Defined Networking

- Separation of control plane from data forwarding plane
- Platform is decoupled from infrastructure
- Centralized controller, network-wide control by controller SW that performs routing and traffic engineering

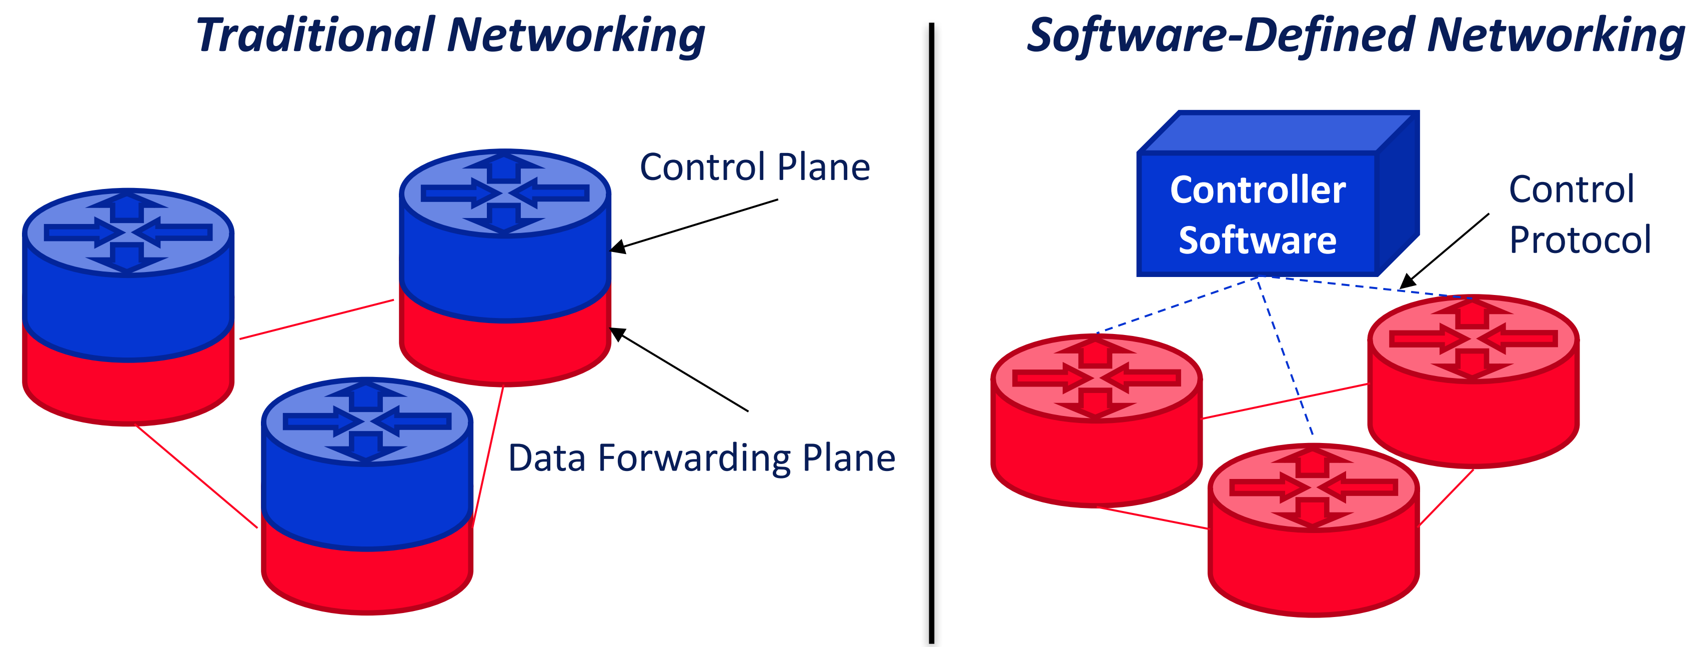

Credit: Jungmin Son

Network Function Virtualization (NFV)

- Migration of network functions to the software layer
  - Firewalls, Network Address Translation (NAT), Intrusion Detection Systems (IDS)
- Virtualized Network Function (VNF)
  - deployable elements of NFV
- Enables better interoperability of equipment and more advanced network functions

Software-defined clouds

- Virtualization in networking
  - Software-defined networking (SDN) and Network Functions Virtualization (NFV).
- Software-defined Cloud Computing (SDC)
  - Extending the concept of virtualization to all resources
    - compute, storage, and network
- Evaluation and Experimentation
  - Complexity, scaling, accuracy, and efficiency.
- A low-cost experimental testbed/infrastructure
  - Conducting practical research in the domain of software defined clouds.

CLOUDS-Pi

- Our recipe for constructing a platform for conducting empirical research in SDCs
  - Easily Repeatable
  - Low-cost (reusing existing servers and Raspberry Pis)
  - Open Source Software
- Hardware
  - Small scale cloud datacenters (9 physical servers, Fat-tree network)
  - Raspberry Pis as SDN Switches
  - Managed enclosure Power Distribution Units (ePDUs)
- Software
  - OpenStack
  - OpenDaylight (ODL)
  - Open vSwitch

Hardware

Machine	CPU	Cores	Memory	Storage
3 x IBM X3500 M4	Intel(R) Xeon(R) E5-2620 @ 2.00GHz	12	64GB (4 x 16GB DDR3 1333MHz)	2.9TB
4 x IBM X3200 M3	Intel(R) Xeon(R) X3460 @ 2.80GHz	4	16GB (4 x 4GB DDR3 1333MHz)	199GB
2 x Dell OptiPlex 990	Intel(R) Core(TM) i7-2600 @ 3.40GHz	4	8GB (2 x 4GB DDR3 1333MHz)	399GB


Eaton Managed Enclosure Power Distribution Units (ePDUs)

USB 2.0 to 100Mbps Ethernet adapters

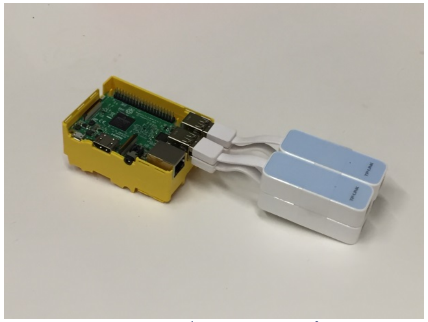

Raspberry Pis
(Pi 3 MODEL B)

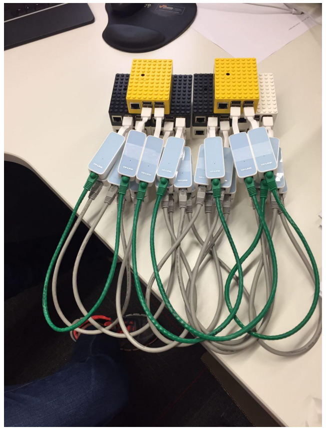

System Architecture

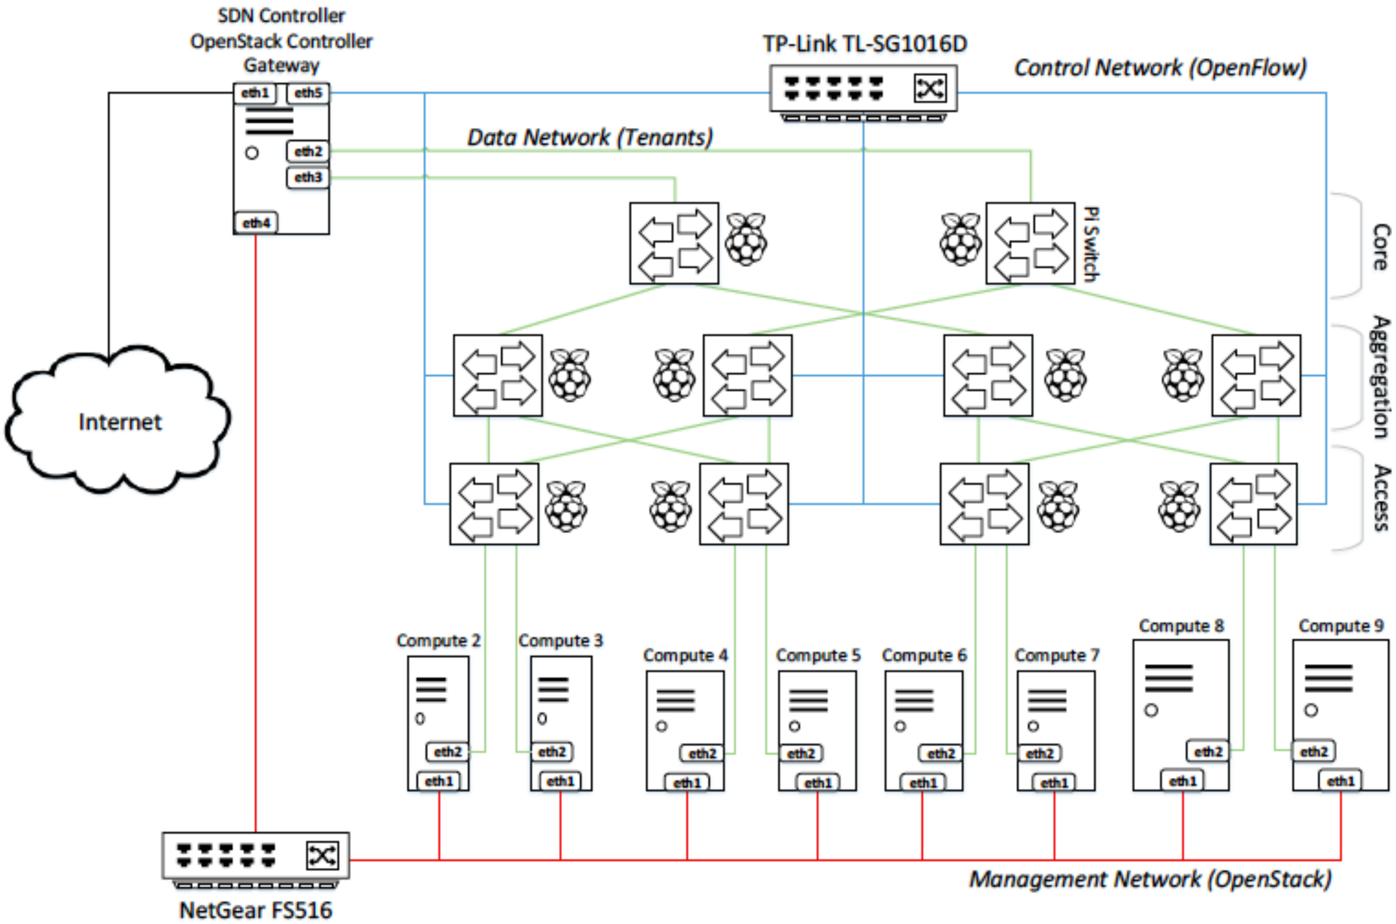

More Photos

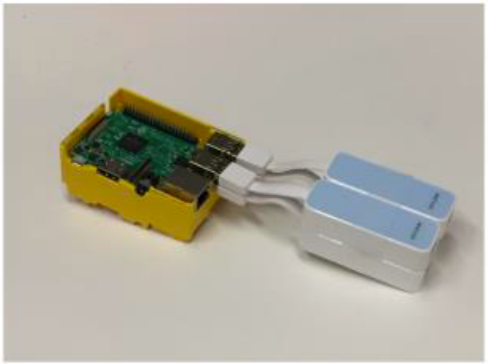

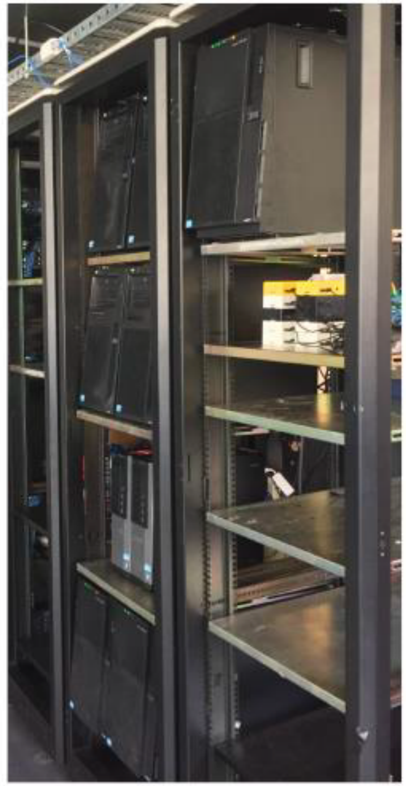

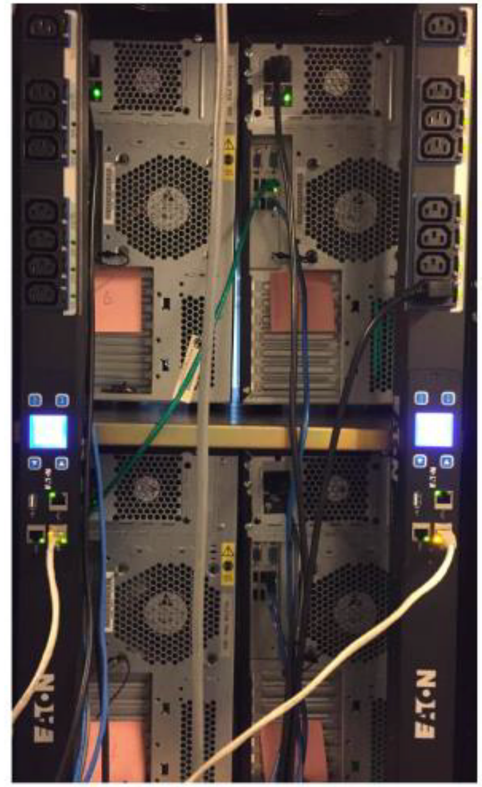

Dynamic Flow Scheduling for Virtual Machine Migration

- "Is it possible to reduce live VM migration time and overhead by dynamically scheduling flows in a cloud data center with multiple paths available between a given pair of physical hosts?"

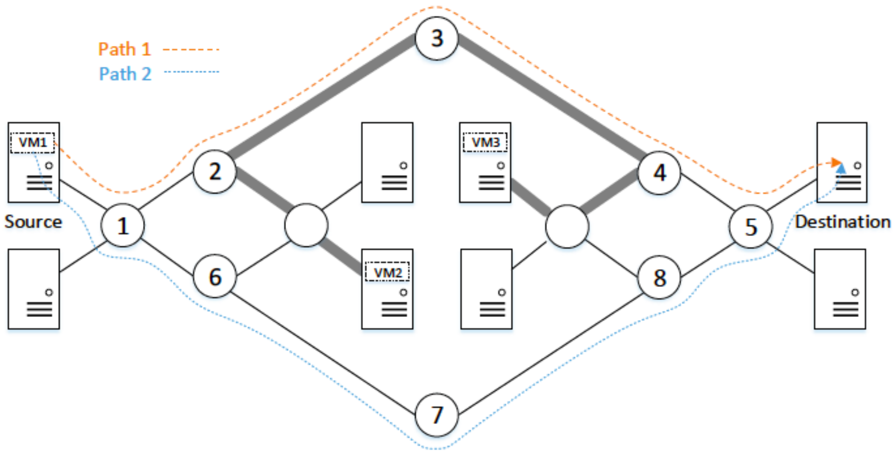

Acinonyx: Proposed Algorithm

- When multiple shortest paths are available between the source and destination
  - As long as the VM migration is in progress it exploit residual bandwidth on multiple available paths
  - Redirect the live VM migration traffic on a path with the lowest load
  - Find a path that has the highest residual bandwidth on its most utilized link
  - Push appropriate flow rules into the switches to redirect traffic

Some Results

Metric	Static Routing	Acinonyx
Migration Time (s)	287	**256**
Average Throughput (Mbs)	32.0	**34.4**

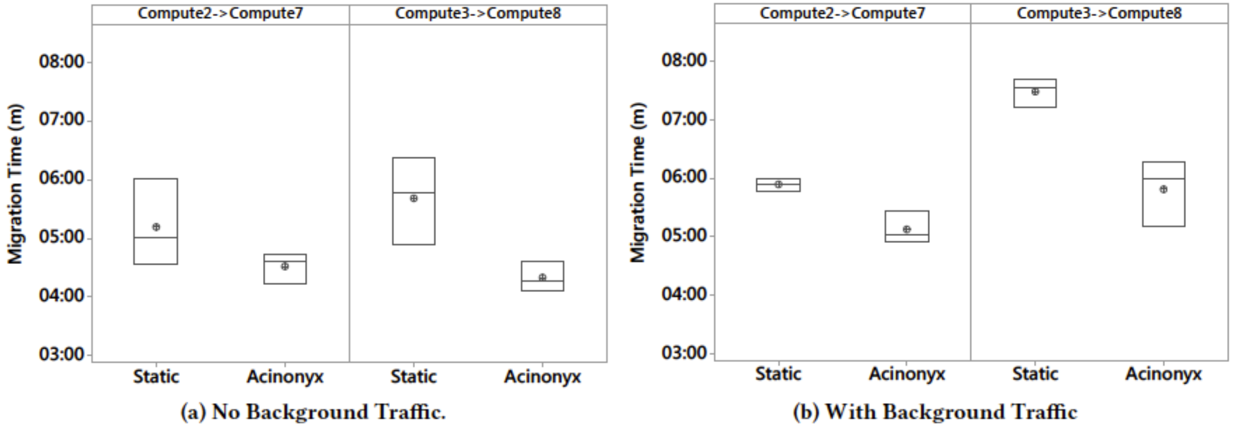

Migration time for two simultaneous migrations when Static and Acinonyx flow scheduling are used.

Summary

- **Cloud computing** is a critical building block of many ICT applications.
- **Geographical load balancing** for maximization of renewable energy usage.
  - Real traces of web requests for English Wikipedia
  - Meteorological data in the location of each data centre to model solar and wind power generation
  - Uses **17%** less **brown energy** and saves **cost** by almost 22% in comparison to round robin policy.

Summary

- Deadline-aware Scheduling and Resource Provisioning Method for Data-intensive Applications on Hybrid Clouds
  - The proposed method is able to meet strict deadlines for a sample data-intensive application to measure the walkability index
  - It minimizes cost and the total number launched instances compared to other existing algorithms.
- Recipe for constructing an economical testbed for **Software Defined Clouds** and conducting practical experiment
  - Dynamic flow scheduling algorithm for live VM migration
  - Migration time is reduced by 12% and network throughput is increased by 7%.

THANK YOU

Questions?