Geographical Load Balancing for Sustainable Cloud Data Centres

Adel Nadjaran Toosi

Cloud Computing and Distributed Systems (CLOUDS) Laboratory,
School of Computing and Information Systems
The University of Melbourne, Australia

Email: anadjaran@unimelb.edu.au
Homepage: http://www.cloudbus.org/~adel/

Outline

- Brief Biography
- Backgrounds
- Geographical Load Balancing (GLB)
- Optimal offline algorithm and its intractability
- A GLB framework for web applications
- Results and Performance Evaluation
- Summary and future directions

Biography and Research Overview

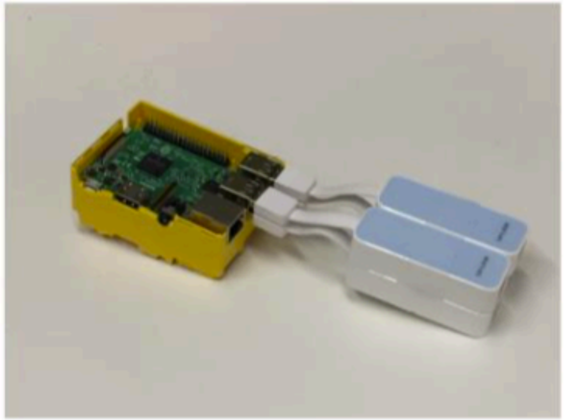

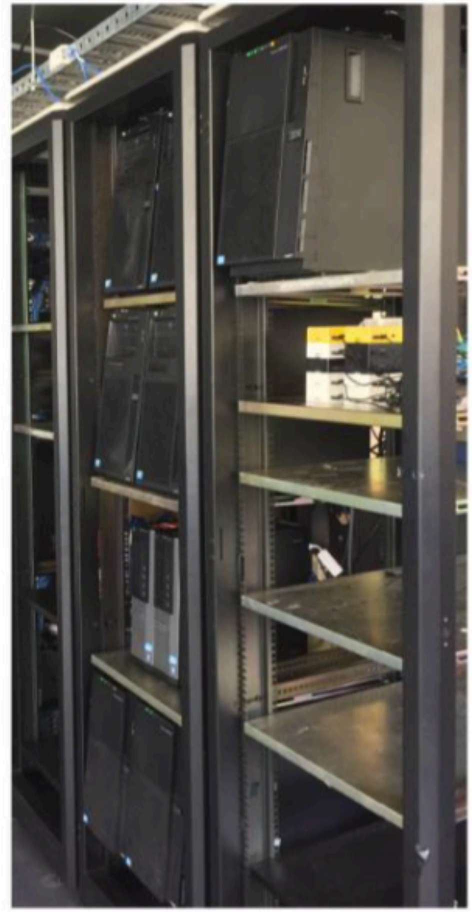

- PhD, U
  - T … cing, Markets,
    a
- Postd … esent
  - A
  1. … npleted]
  2.
  3.
- Resea
  - **D** …) and Network
    F … **Computing**
  - S … uting and
    D
- Publi
  - **27** publications, 15 Journal Articles (**11 A/A*** ERA Ranking, ACM CSUR,TCC, JCNA, FGCS, TAAS), **11** Conference papers (CloudCom, UCC, HPCC), **1** Book Chapter,
  - h-index: **15** and **1200+** citations (SRC: Google Scholar)

Cloud Computing

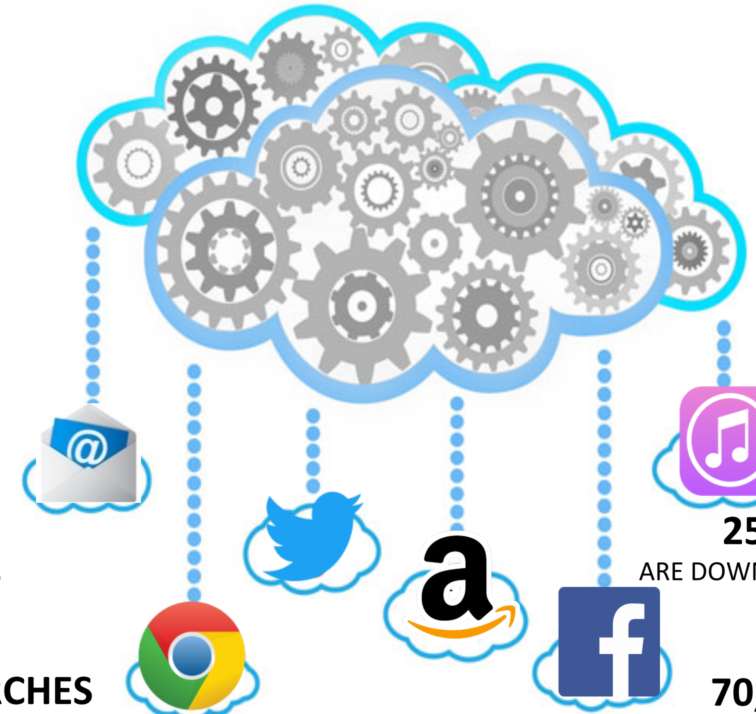

3.4 MILLION EMAILS
ARE EXCHANGED[a]

6,000 TWEETS
ARE SENT ON TWITTER[b]

100,000 SEARCHES
ARE MADE ON GOOGLE[c]

600 ITEMS
ARE SOLD ON AMAZON[d]

250 SONGS
ARE DOWNLOADED VIA ITUNES[f]

70,000 LIKES
ARE GENERATED ON FACEBOOK[e]

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

- Cloud providers aims
  - Reducing energy consumption
  - Dependence on brown energy
- Renewable energy
  - On-site green power generation
  - Google, Microsoft and Amazon

"Amazon Web Services (AWS) has built a wind farm in 2017 and exceeded the goal of 50% electrical usage from renewable energy sources"

- Not only Data Centres
  - Monash Net Zero Project, The University of Melbourne's Sustainability Plan
  - Bitcoin Mining

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

Challenges

- ***Non-dispatchable, Intermittent*** and ***Unpredictable***
  - Renewable energy sources (Wind and Solar)
  - 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

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**?"*

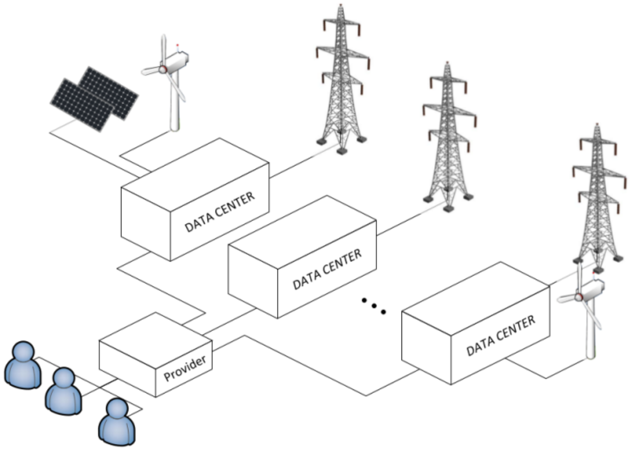

Example: Offline GLB Problem

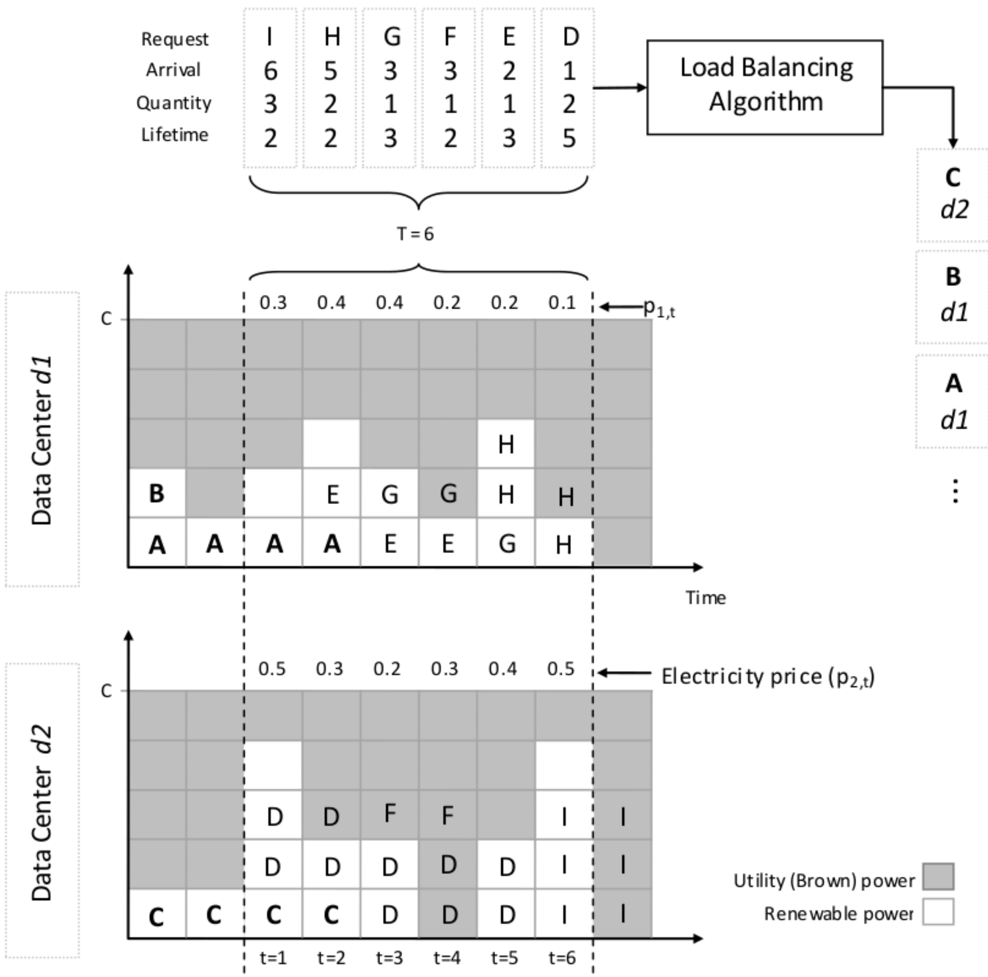

Optimal Offline Algorithm

- Assuming the following information is known for a time window:
  - Future knowledge of renewable energy availability
  - Workload (i.e., number of requests, arrival time, and duration of requests)
- We showed that the optimal strategy is computationally intractable
  - Exponential time complexity
  - Formal proof

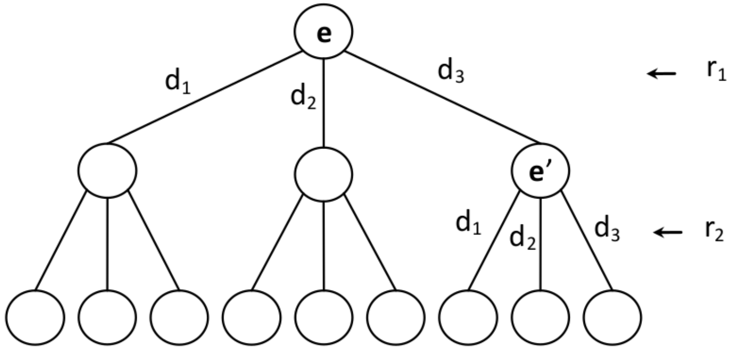

GLB for Web Applications

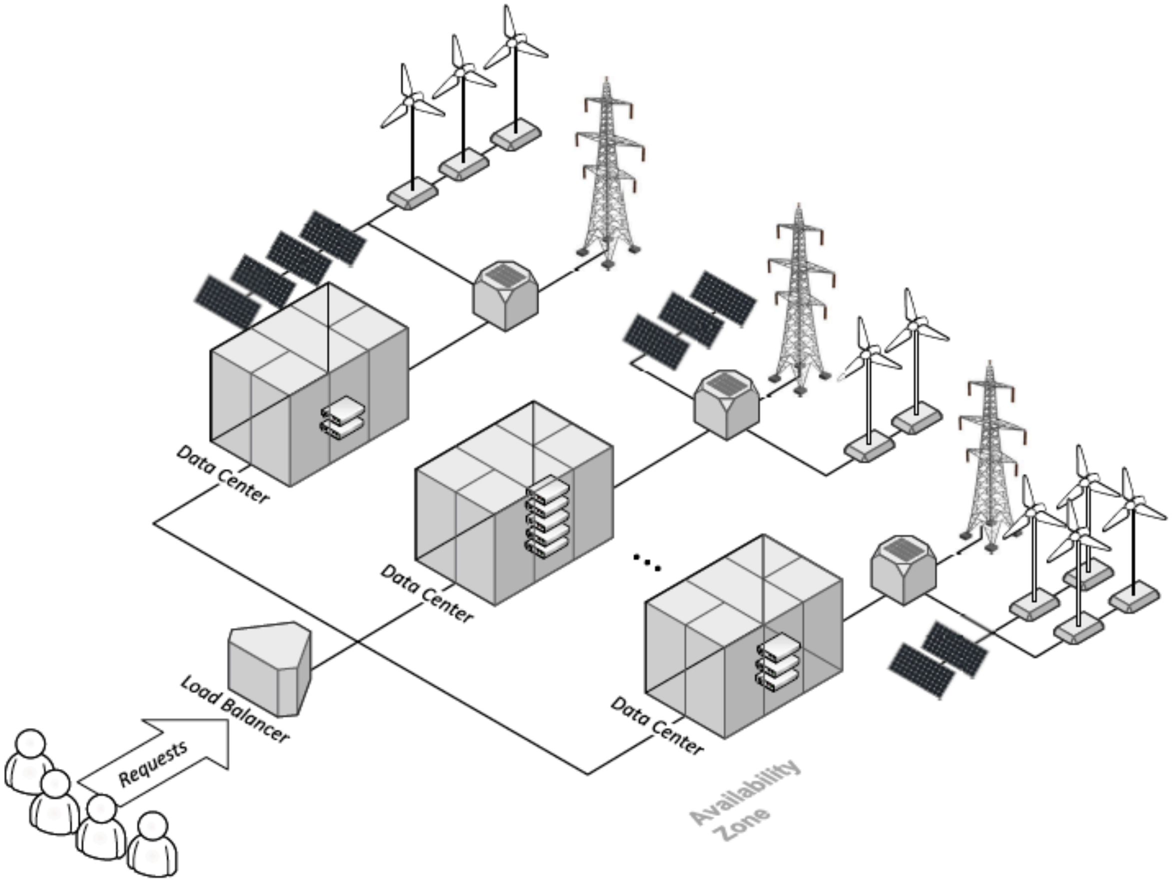

Overall System Architecture

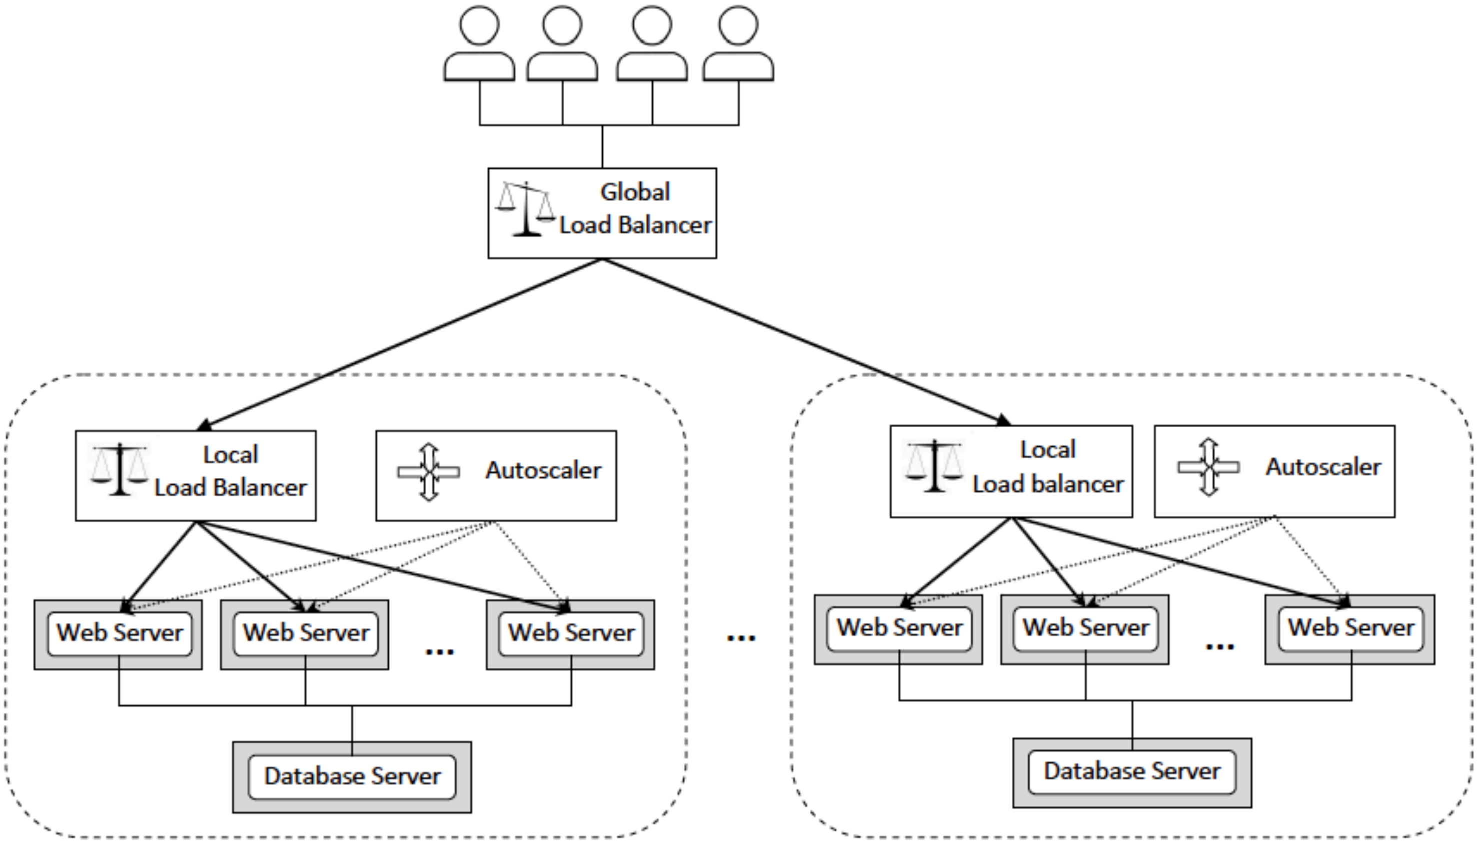

Global Load Balancer (GreenLB)

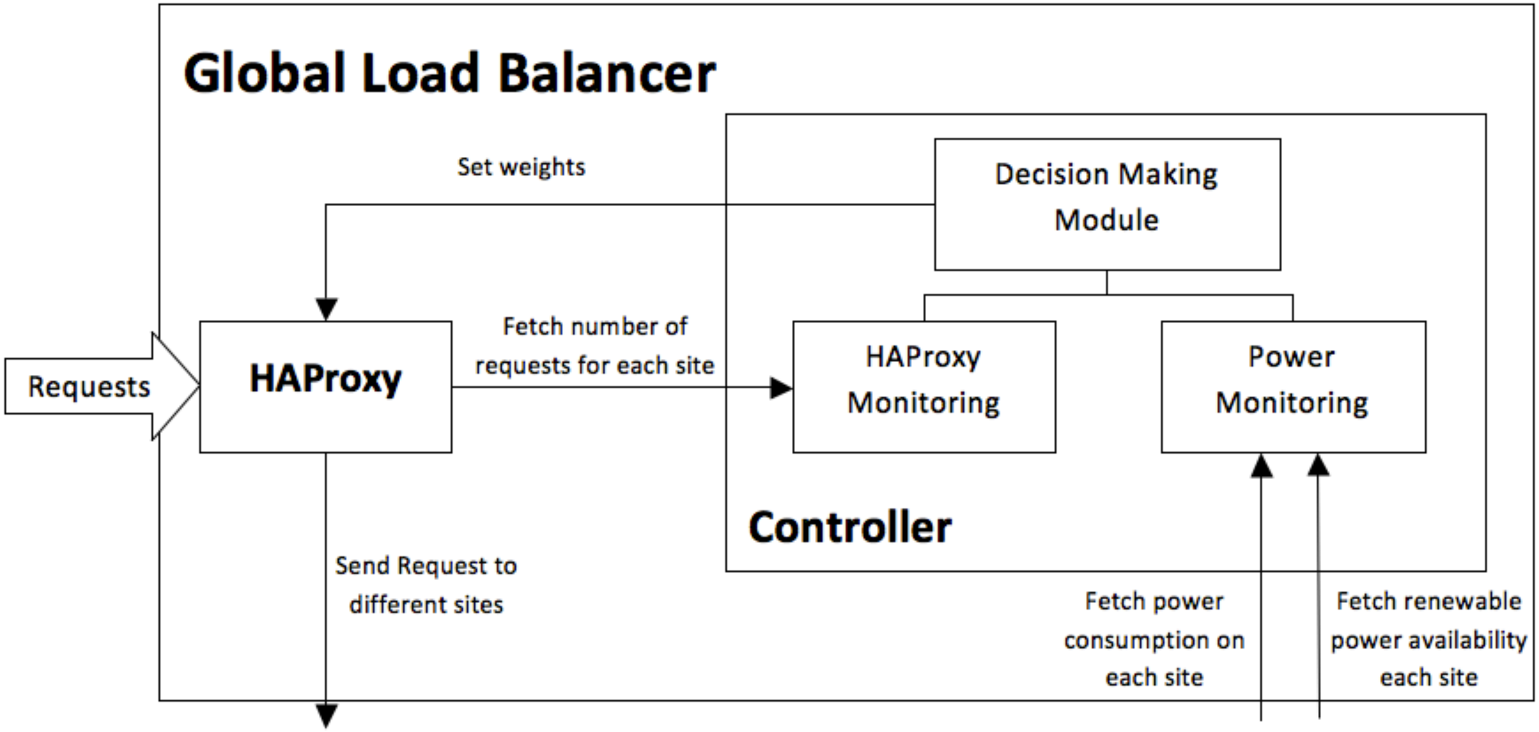

Grid'5000 Testbed

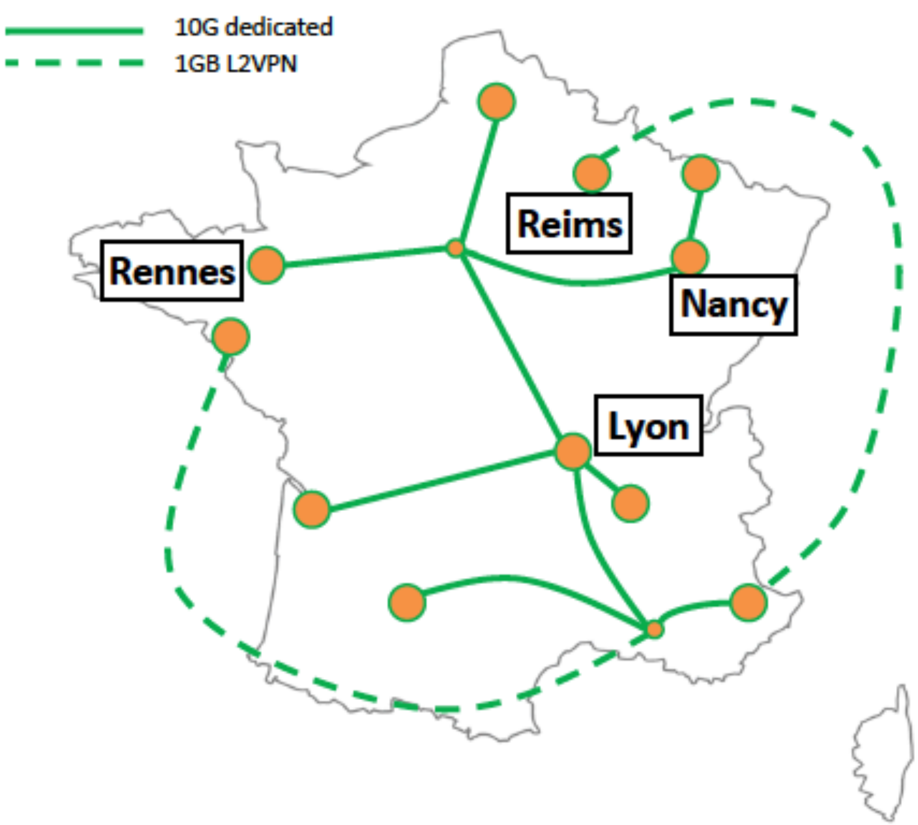

Workload Traces

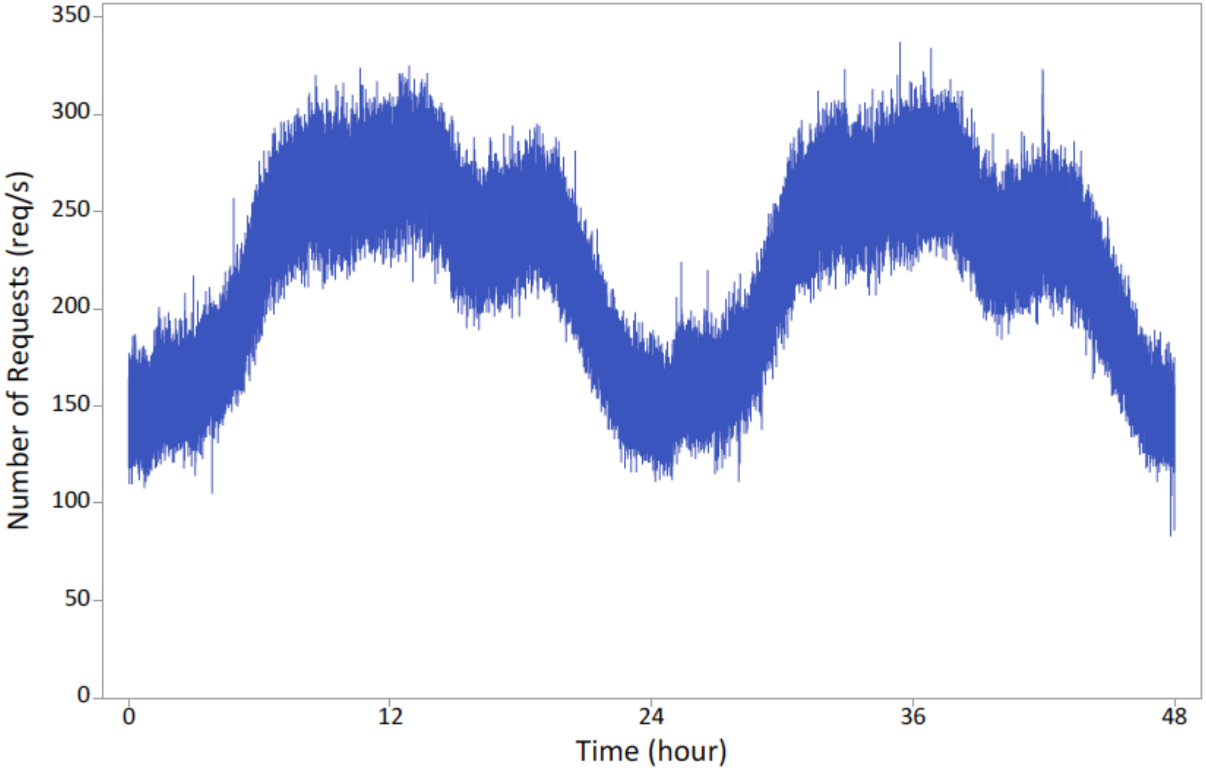

Wikipedia

A Prototype System

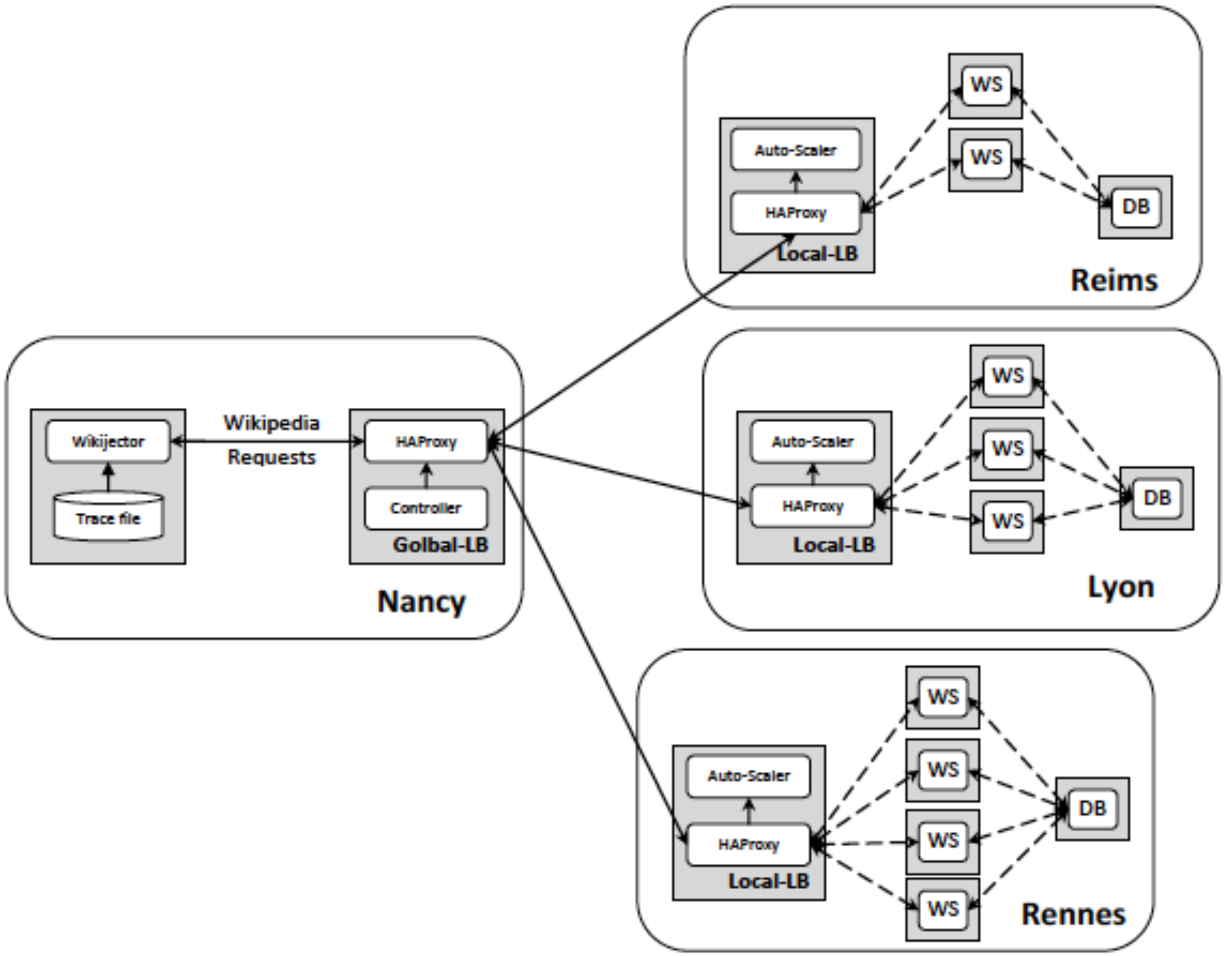

Renewable Power and Electricity Prices

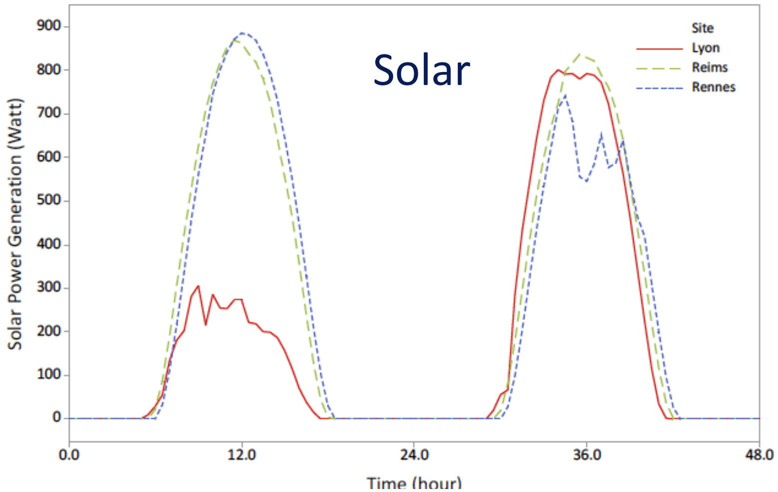

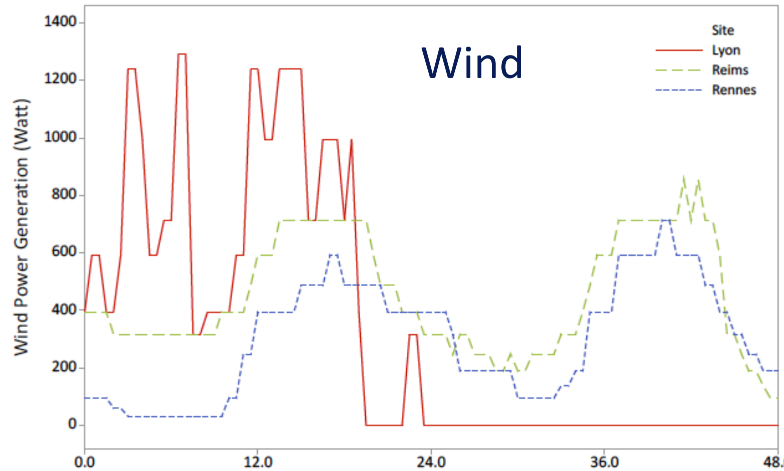

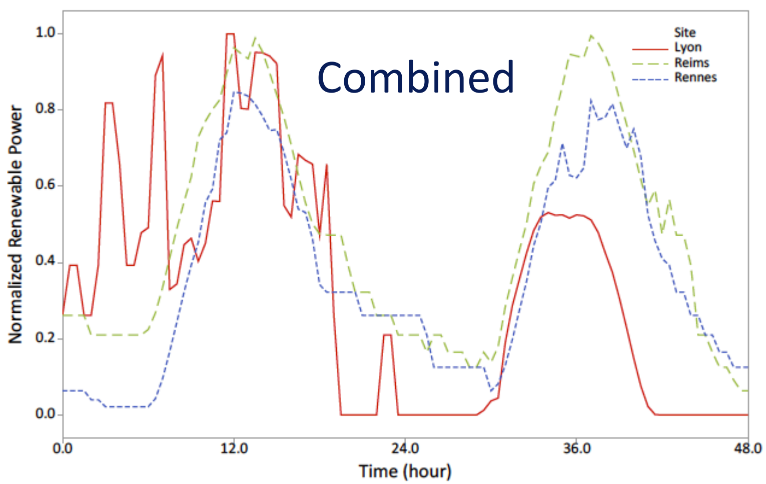

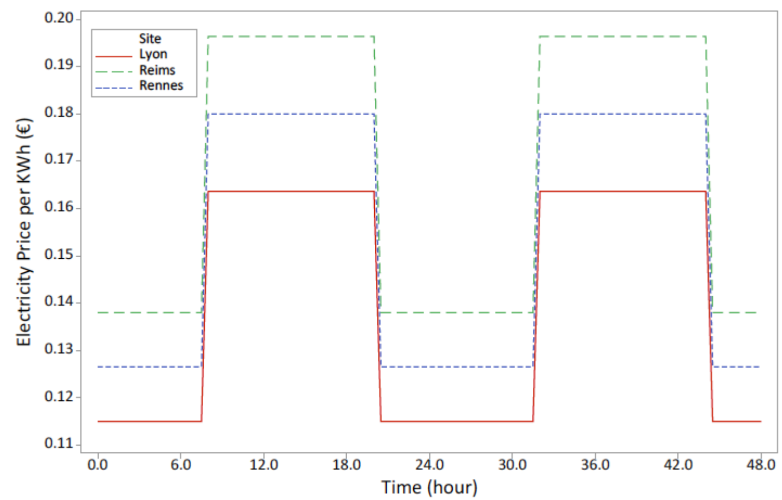

Results

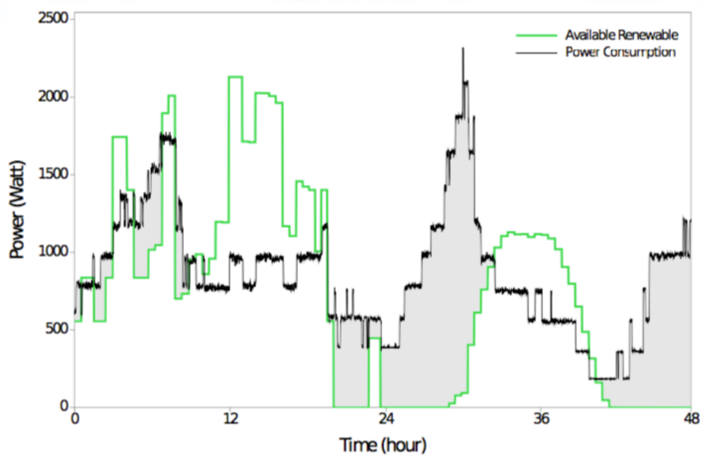

Lyon

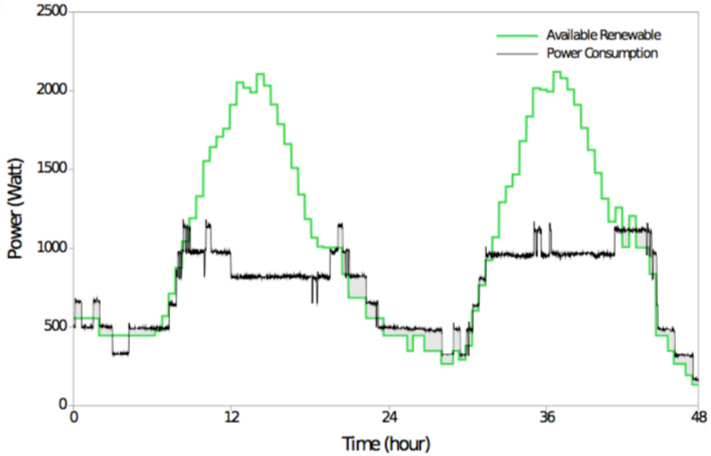

Reims

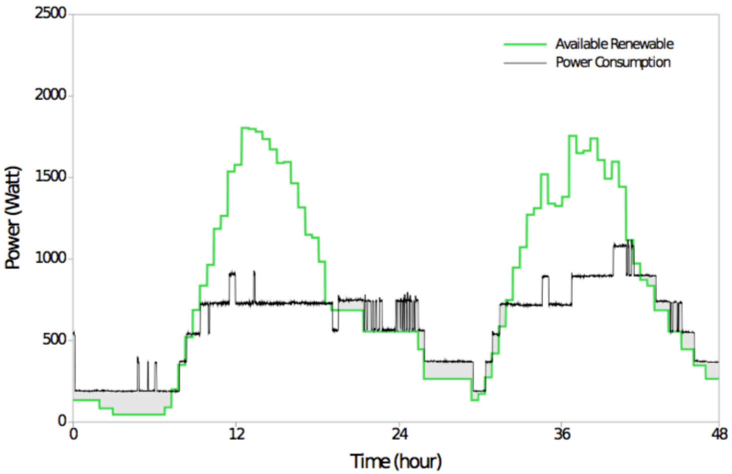

Rennes

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	[illegible]5
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%

Summary and Conclusion

- A framework for **cost** and **energy efficient** load balancing
  - Distributes web application requests among multiple cloud data centres
- A prototype and experimental studies in a real testbed
  - 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.
- Reduces **cost** by 8%, **Brown energy** by 7% in comparison a method by a group of researchers from **Rutgers** and **Princeton** universities.
  - Linear Optimisation
  - Workload and Renewable Energy Prediction

Future Works

- GLB for other types of workloads/applications
  - Scientific workflows, Map-Reduce
  - Web Sticky Sessions
  - Demand response
- Internet of Things (IoT)
  - Healthcare, Smart Vehicles
  - Edge and Fog Computing
  - Sustainability and Reliability
- Challenge
  - Shaping workload to match renewable power supply
    - Offloading IoT tasks to the core clouds
    - Trimming approximation analytics
    - Scheduling deferral tasks
    - Selective microservices power-off
    - Orchestration of network slices

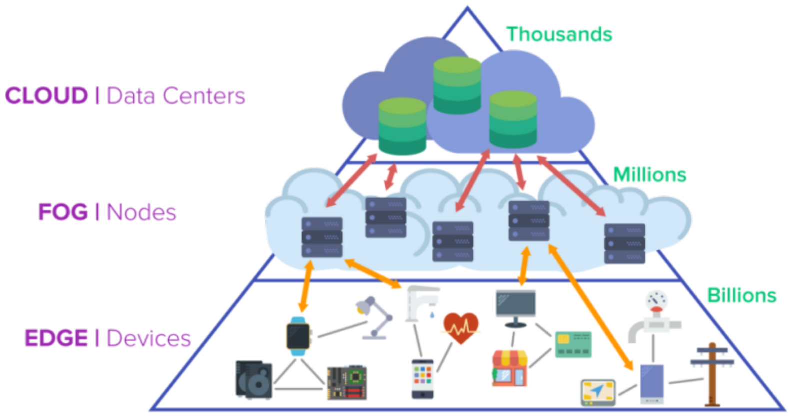

https://erpinnews.com/fog-computing-vs-edge-computing

THANK YOU

Questions?

Results

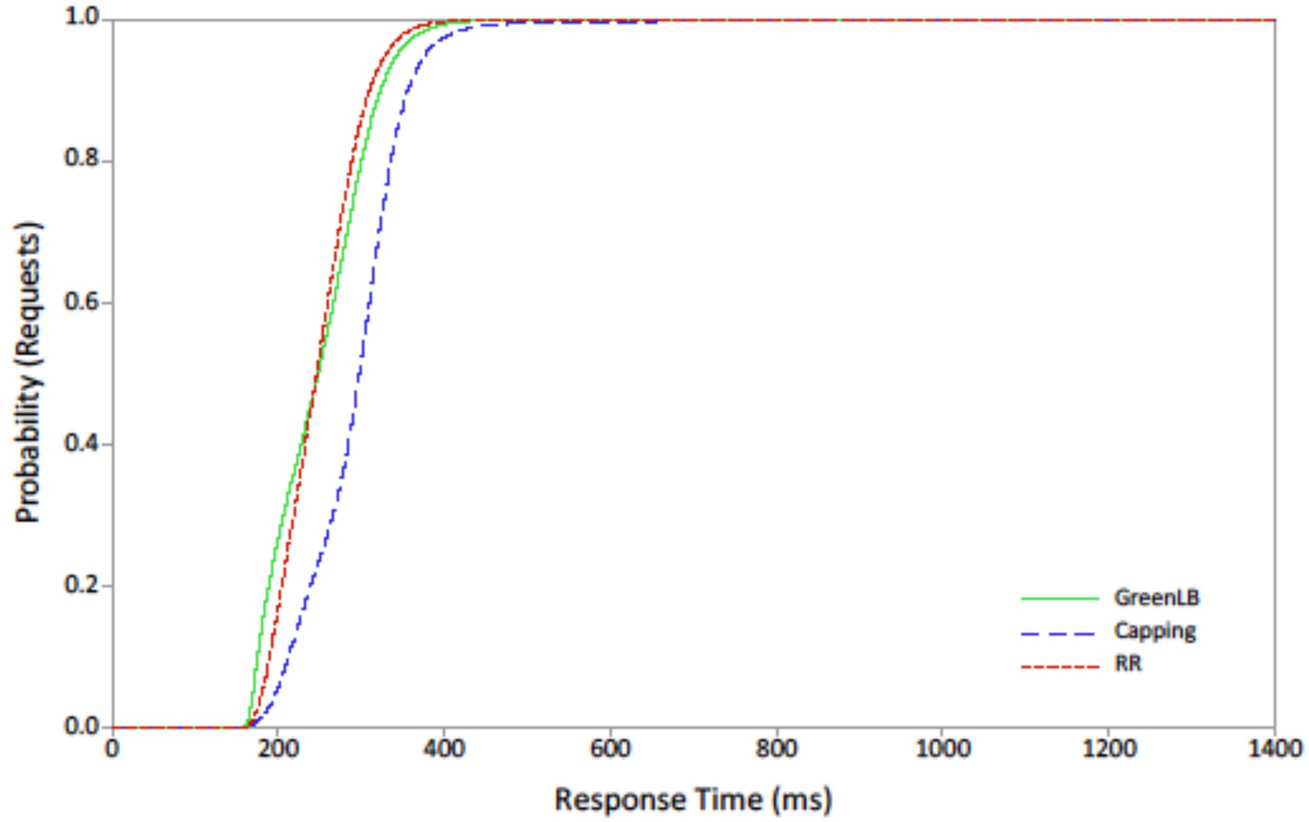

CDF of average response time