Redes de Computadores

Prof. Miguel Elias Mitre Campista

http://www.gta.ufrj.br/~miguel

Part I

Hot topics in networking

The number of smartphones (*mobile devices*) has been increasing at a fast pace in the last few years...

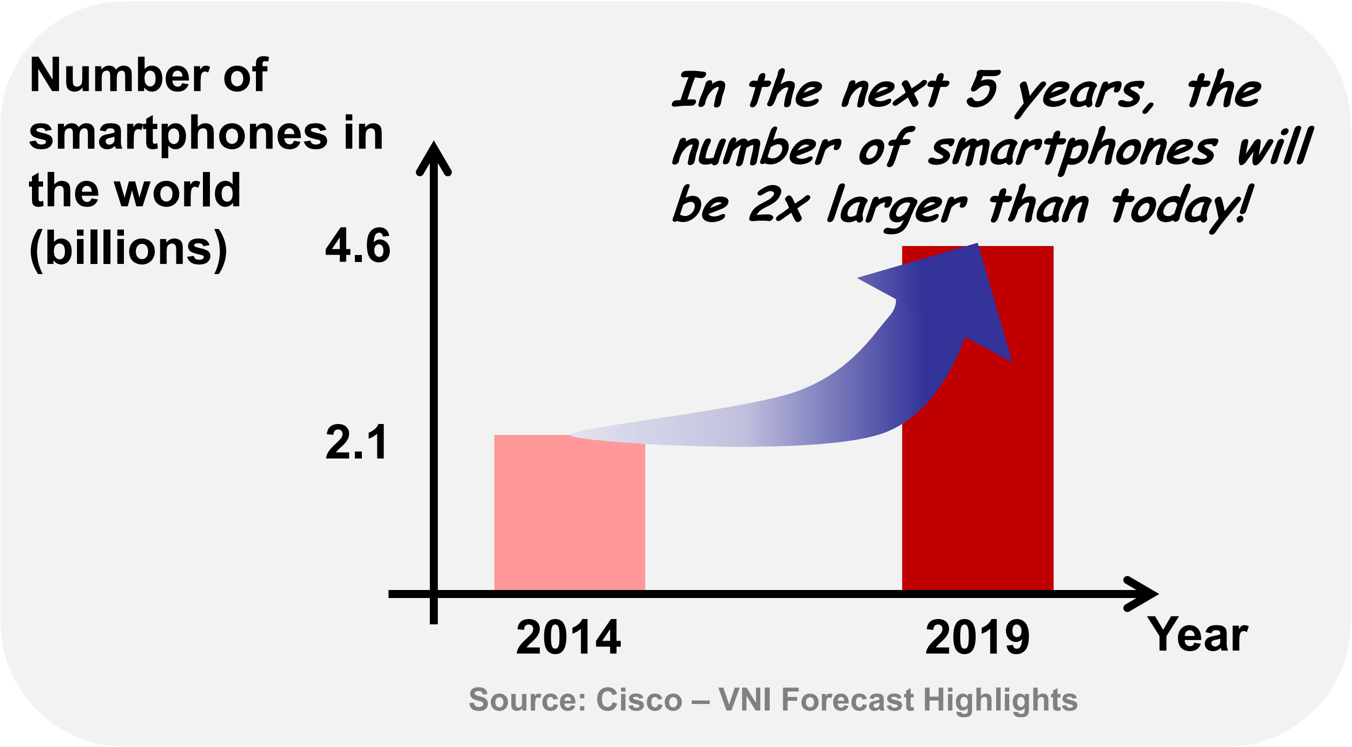

How should we deal with such increasing number of connected devices?

Mobile Cloud Computing (MCC)

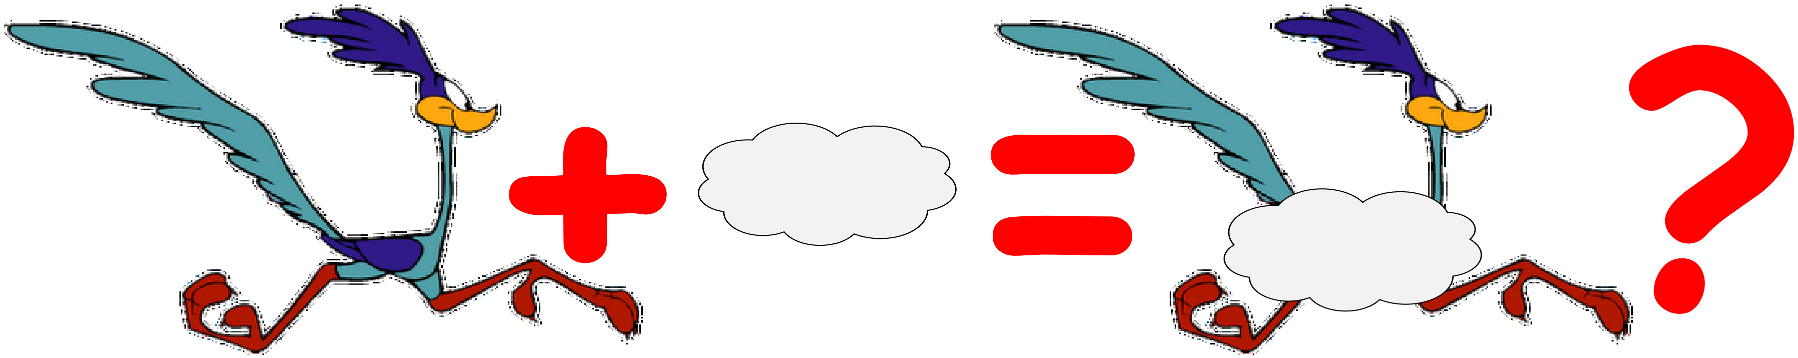

Mobile Cloud Computing (MCC)

- How appealing it is?
  - Mobile users
    - Revenues with richer computing resources and thousands of available applications
  - Network operators
    - Revenues with larger bandwidth plans for users and cloud providers
  - Cloud providers
    - Revenues with an increasing number of users and economies of scale

This can be viewed as a three-layer architecture...

MCC Architecture

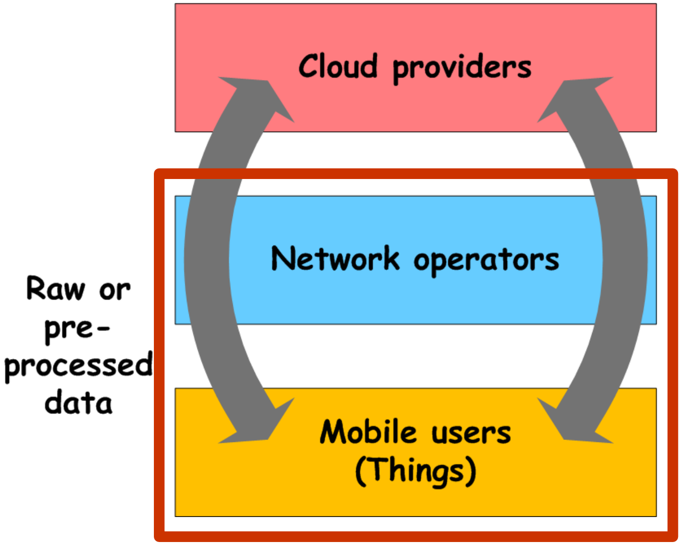

Can you see any challenge?

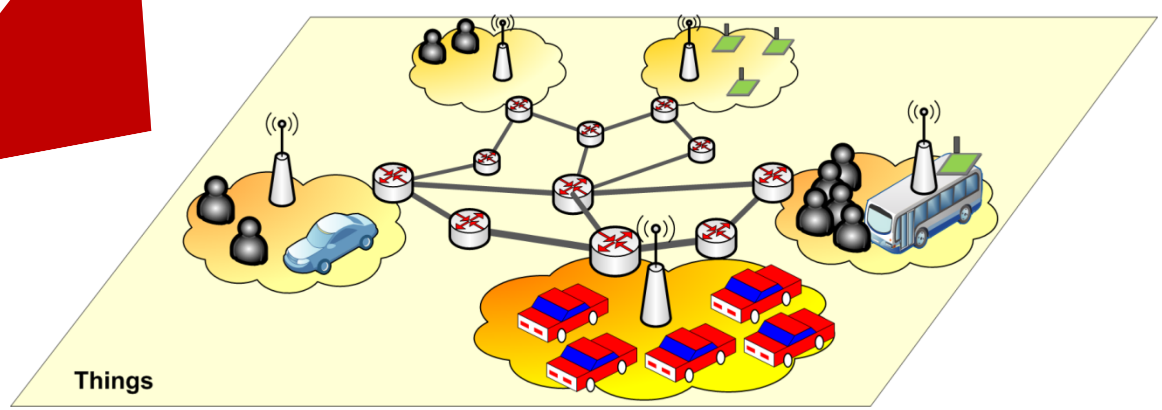

MCC Architecture

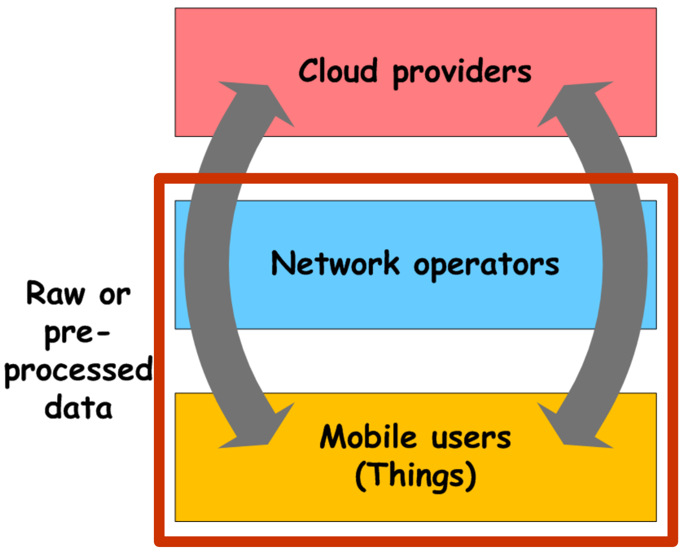

- Well, I can see a few...
  - Mobility
  - Vehicular networking
  - Internet of Things (IoT)
  - Sensor networking/crowdsensing
  - Access networks, ...

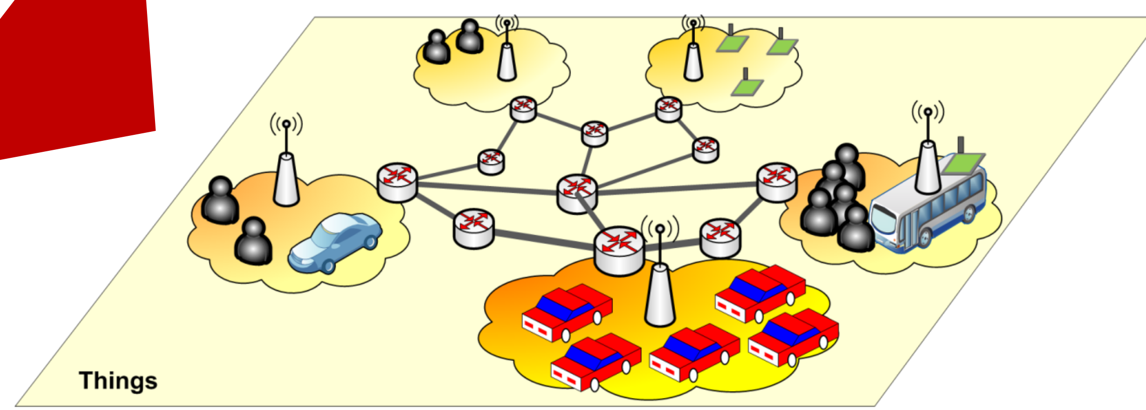

MCC Architecture

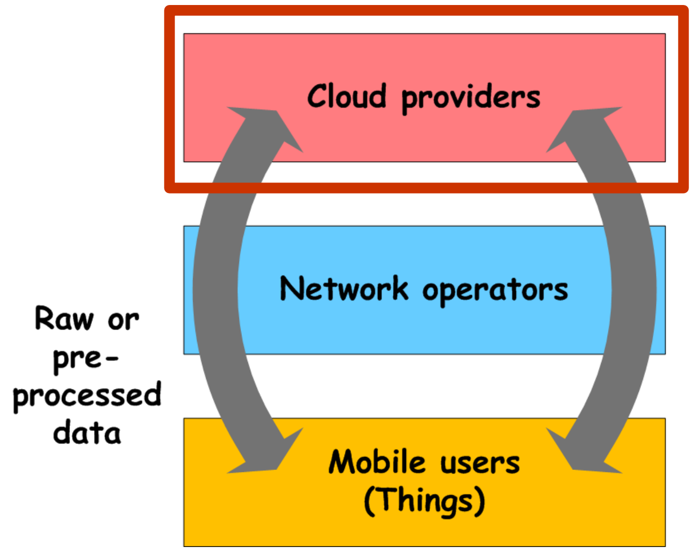

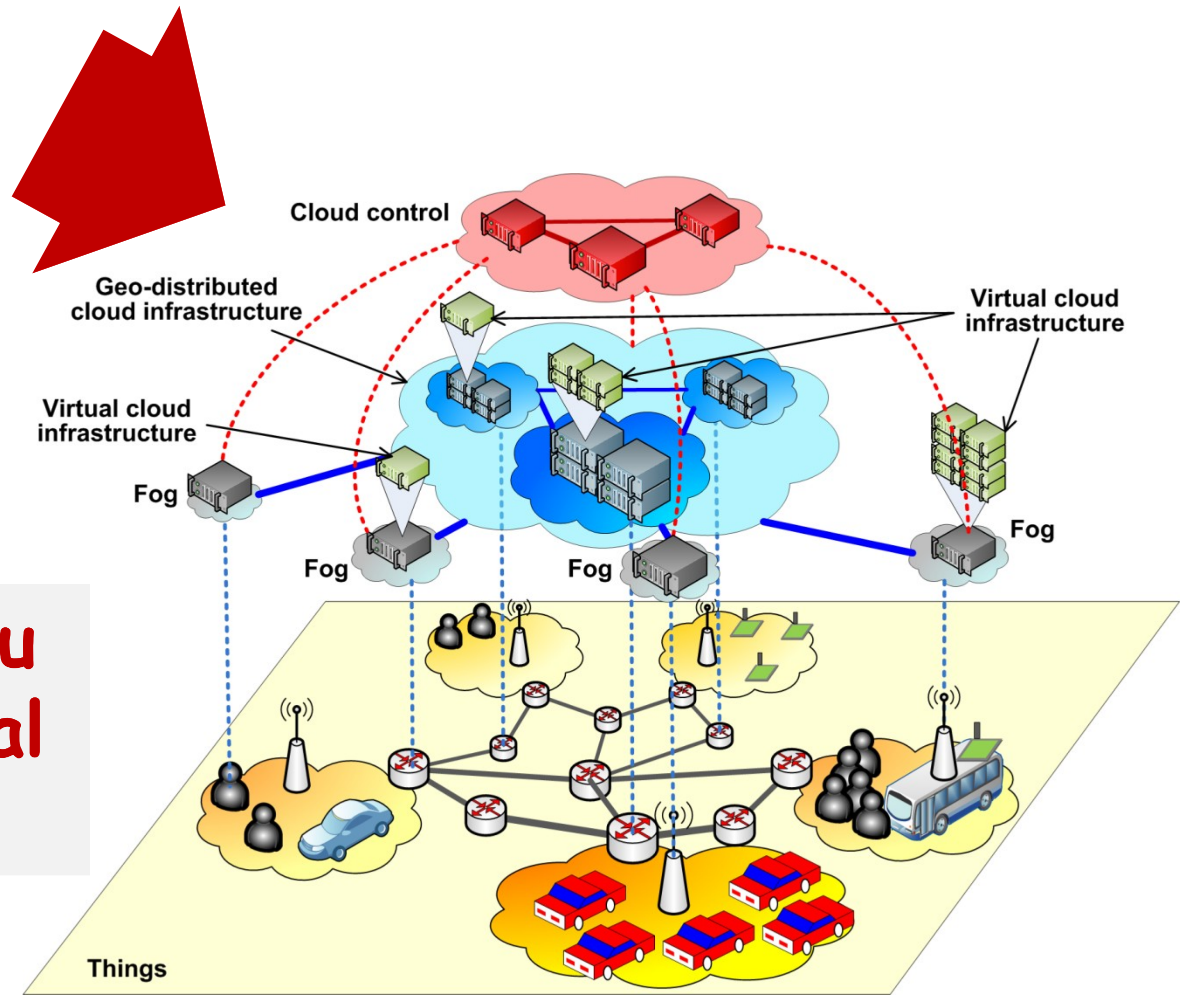

And now, can you see any additional challenge?

MCC Architecture

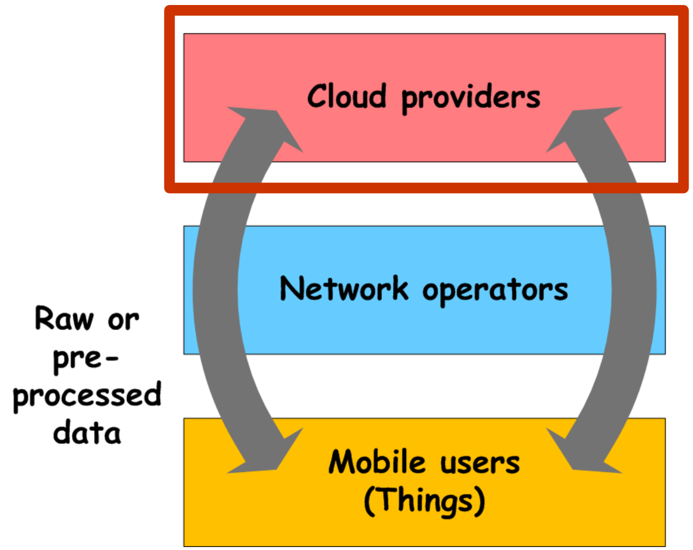

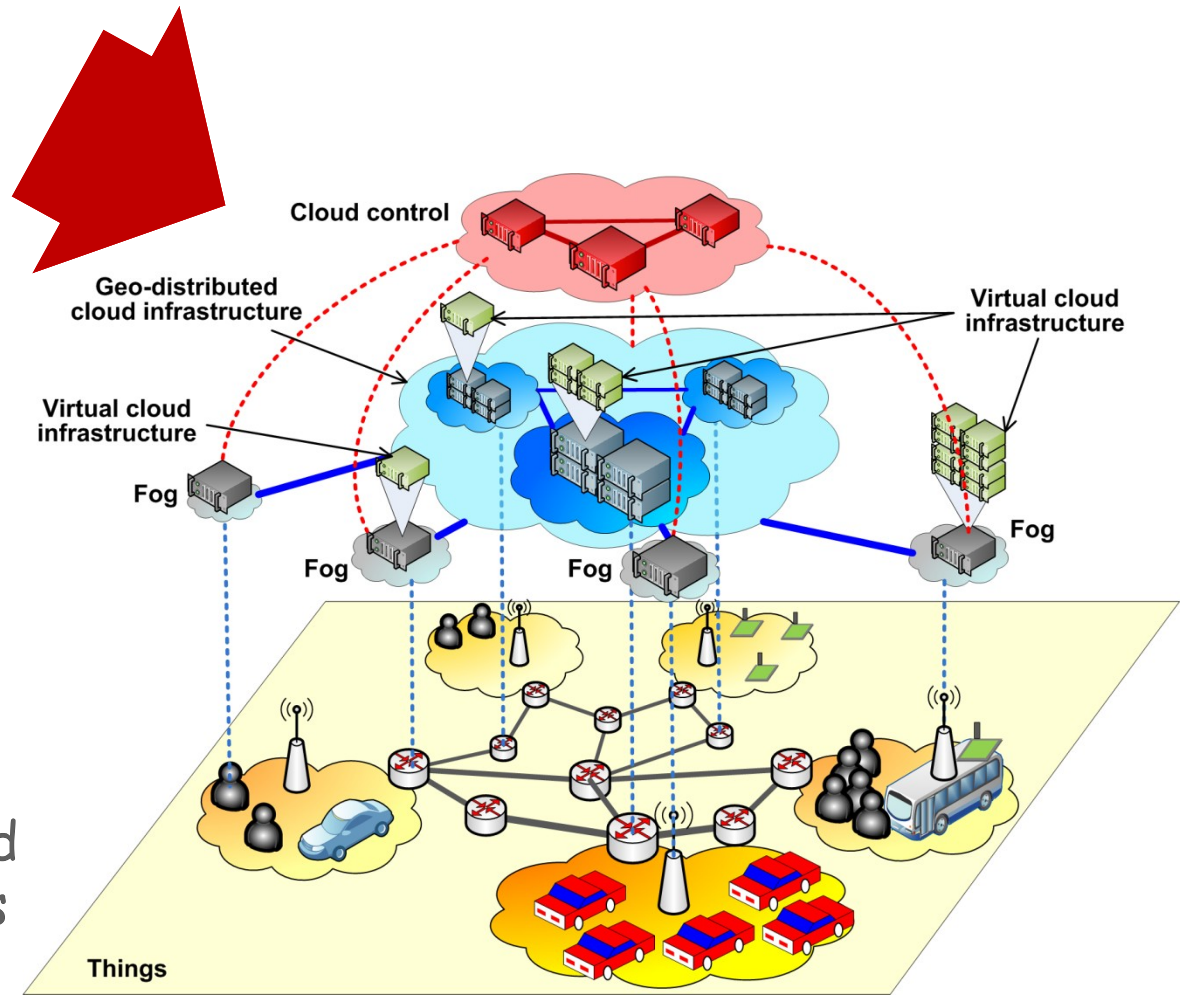

- Here is my list…
  - *Cloud computing*
  - *Cloud control*
  - *Geo-distributed and collaborative clouds*
  - *Fog computing, …*

My Ongoing Work

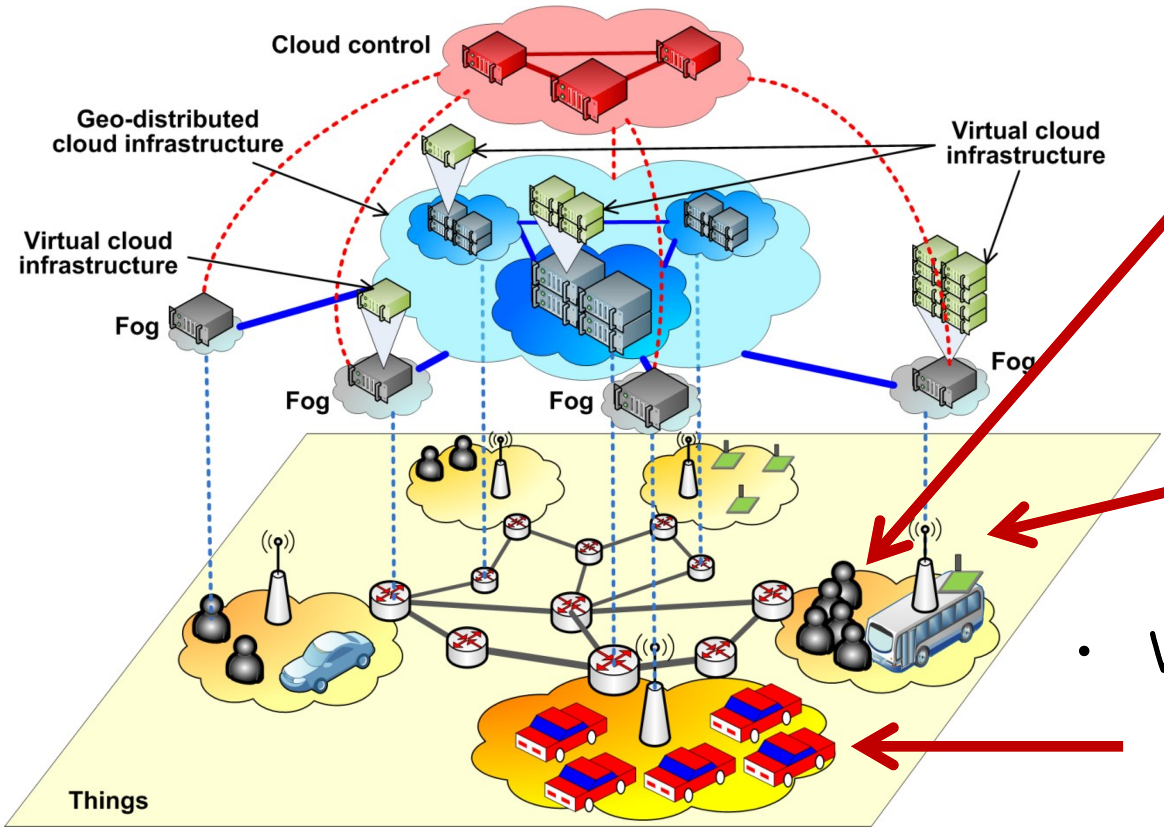

- Vehicular social applications…
  - Do they generate data traffic?
- Vehicular sensing…
  - Can we leverage node mobility for sensing?
- Vehicular mobility…
  - Can we count on multihop communications?

Vehicular Social Applications

- Social data from users to the cloud
  - **Waze**: 10-day dataset from Boston, Massachusetts

Weekdays

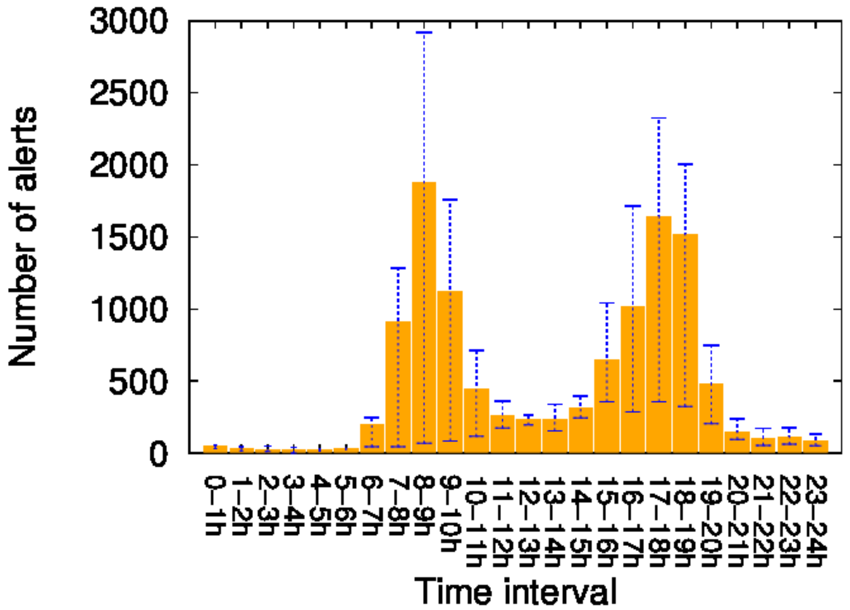

Weekends

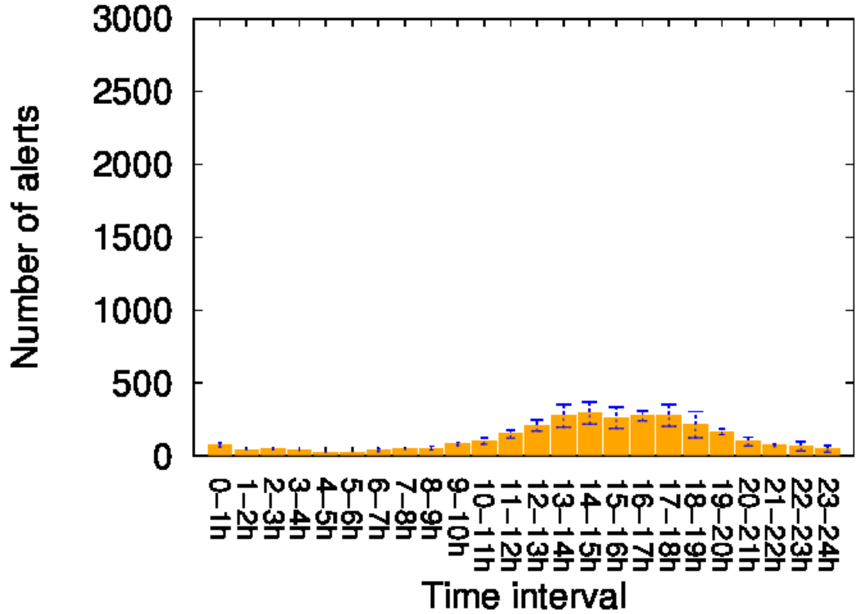

Ribeiro Neto, V., Medeiros, D. S. V., Campista, M. E. M., "*Analysis of Mobile User Behavior in Vehicular Social Networks*", Network of the Future (NoF), Búzio, Rio de Janeiro, November 2016

Vehicular Social Applications

- Social data from users to the cloud
  - **Waze**: 10-day dataset from Boston, Massachusetts

People contribute more at lower speeds!

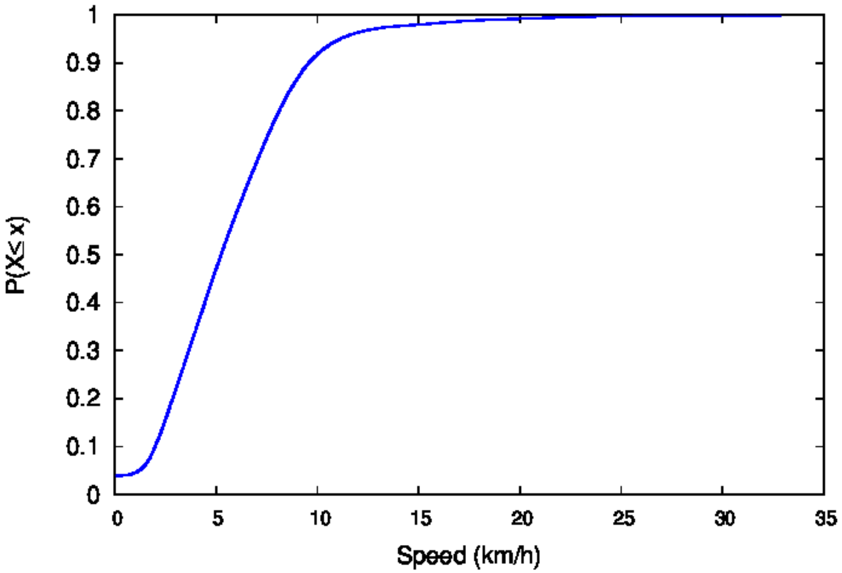

Ribeiro Neto, V., Medeiros, D. S. V., Campista, M. E. M., "*Analysis of Mobile User Behavior in Vehicular Social Networks*", Network of the Future (NoF), Búzio, Rio de Janeiro, November 2016

Vehicular Social Applications

- Social data from users to the cloud
  - **Waze**: 10-day dataset from Boston, Massachusetts

People do not contribute much, but when they do, they send consecutive alerts

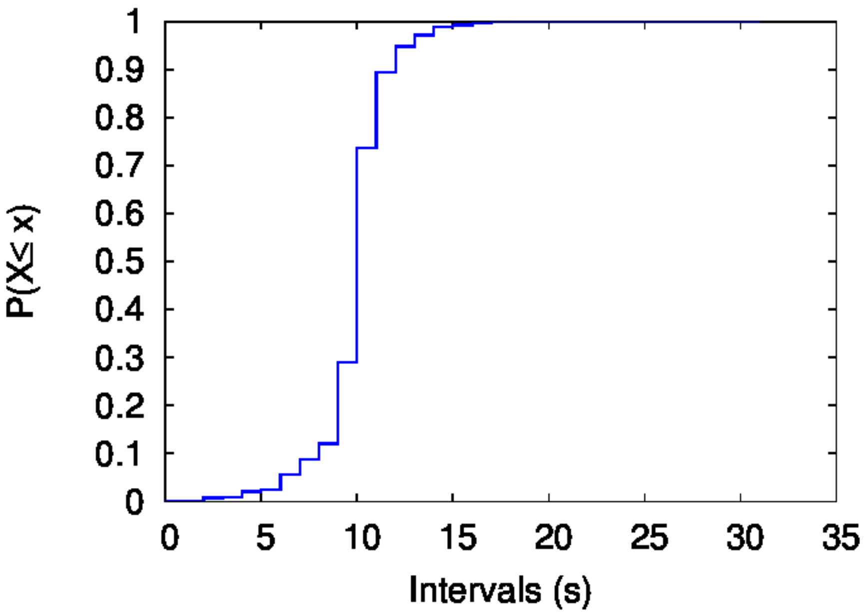

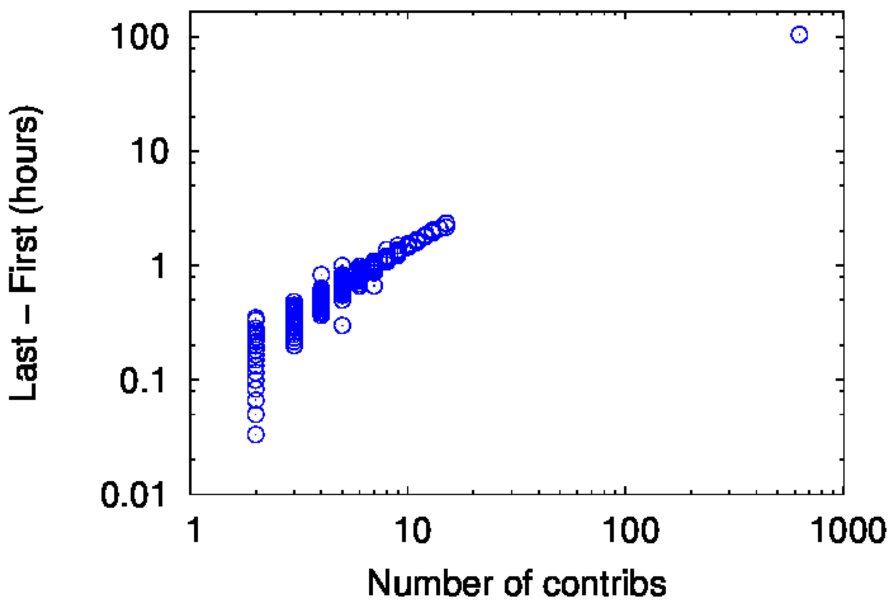

Ribeiro Neto, V., Medeiros, D. S. V., Campista, M. E. M., "*Analysis of Mobile User Behavior in Vehicular Social Networks*", Network of the Future (NoF), Búzio, Rio de Janeiro, November 2016

Vehicular Social Applications

- Social data from users to the cloud
  - **Waze**: 10-day dataset from Boston, Massachusetts

People contribute mostly with information about traffic jams...

Weekdays

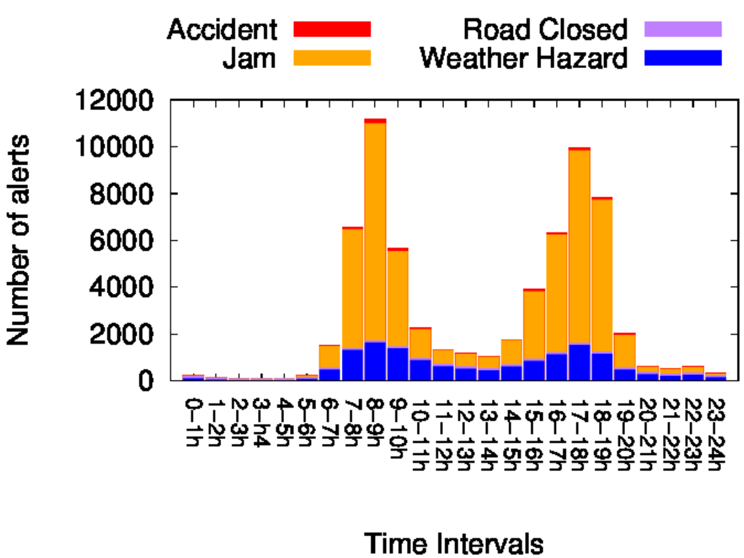

Weekends

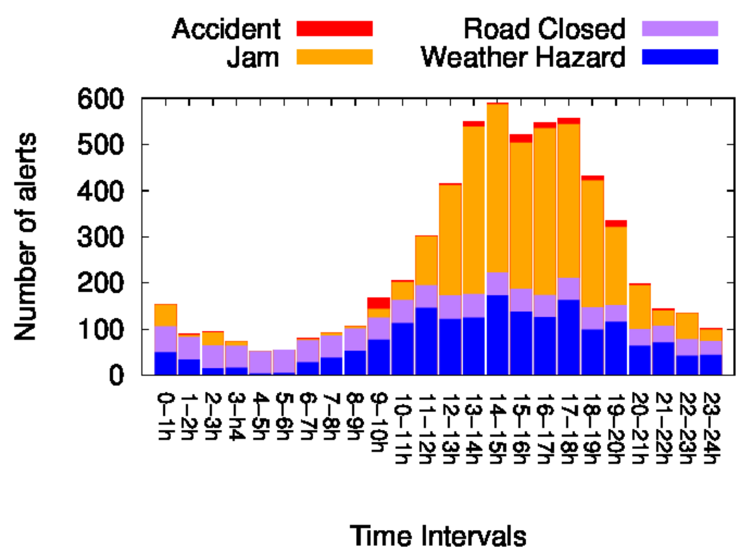

M., "
Network of the Future (NoF), Búzio, Rio de Janeiro, November 2016

Vehicular Sensing

- Environmental readings: Mobile *Vs.* Static sources
  - Mobile nodes can enrich the amount of data collected

Node architecture

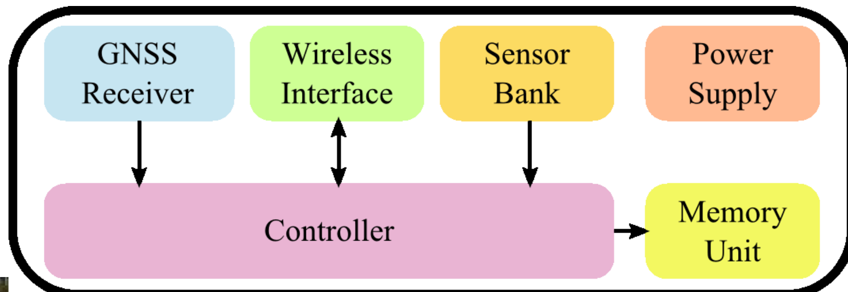

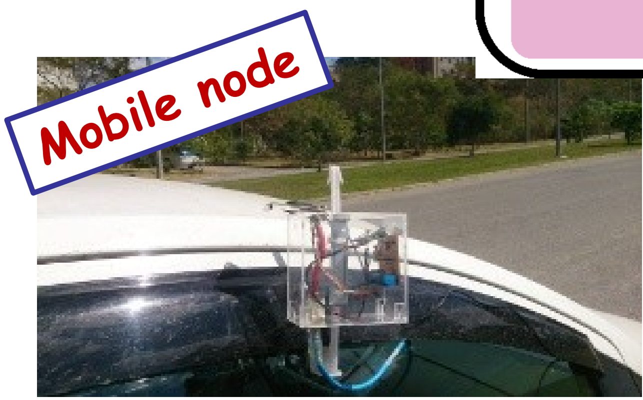

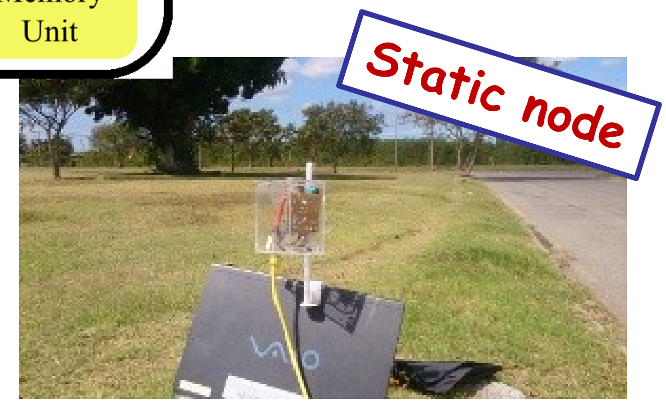

Cruz, P., Pinto Neto, J. B., Campista, M. E. M., Costa, L. H. M. K., "*On the Accuracy of Data Sensing in the Presence of Mobility*", Network of the Future (NoF), Búzio, Rio de Janeiro, November 2016

Vehicular Sensing

- Environmental readings: Mobile *Vs.* Static sources
  - Mobile nodes can enrich the amount of data collected

Vehicle trajectory

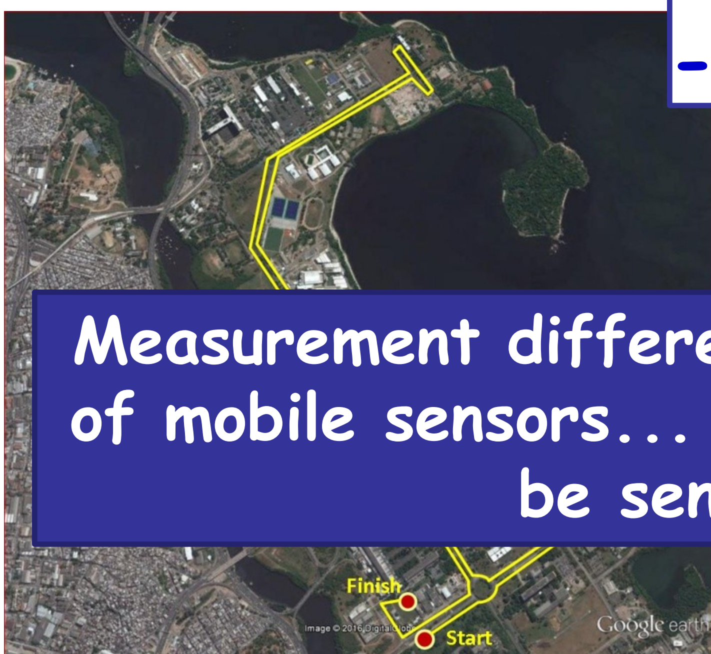

Mobile node

Static node

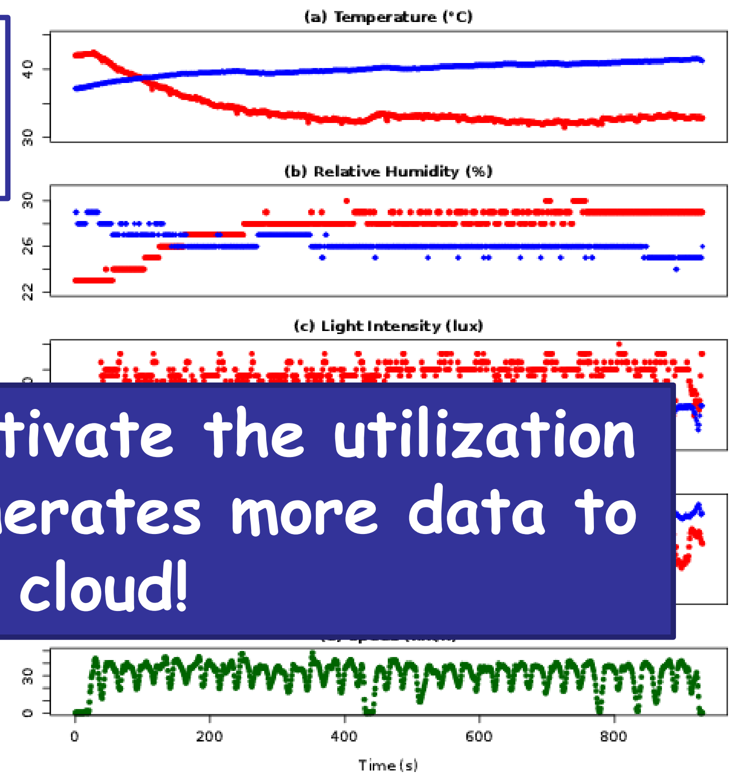

Measurement differences motivate the utilization of mobile sensors... This generates more data to be sent to the cloud!

Cruz, P., Pinto Neto, J. B., Campista, M. E. M., Costa, L. H. M. K., "*On the Accuracy of Data Sensing in the Presence of Mobility*", Network of the Future (NoF), Búzio, Rio de Janeiro, November 2016

Vehicular Mobility

- How should we evaluate the importance of a node?
  - ***The most important node is the one able to keep the network connected...***

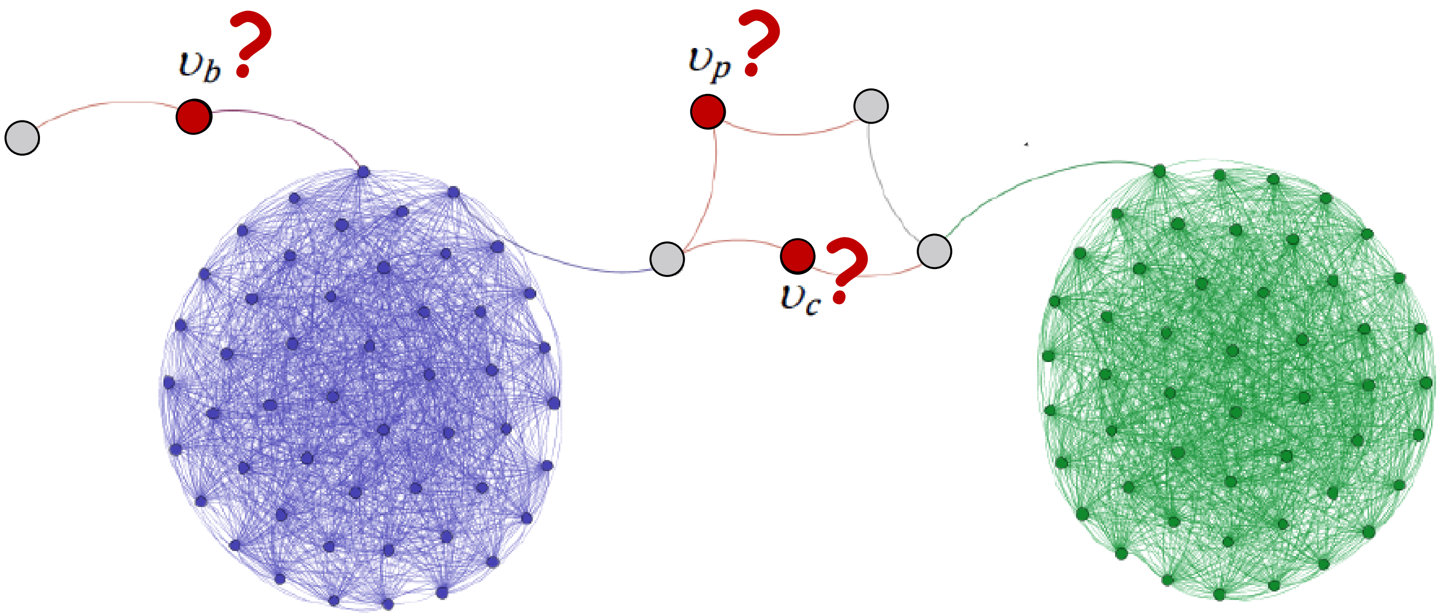

Medeiros, D. S. V., Campista, M. E. M., Mitton, N., Amorim, M. D., Pujolle, G., "*Weighted Betweenness For Multipath Networks*", IEEE Global Information Infrastructure and Networking Symposium (IEEE GIIS 2016), Porto, Portugal, October 2016

Vehicular Mobility

- v_c, v_b, and v_p: which one is the most important?
  - **Betweenness**: *The more shortest paths a node falls in, the more important it is!*

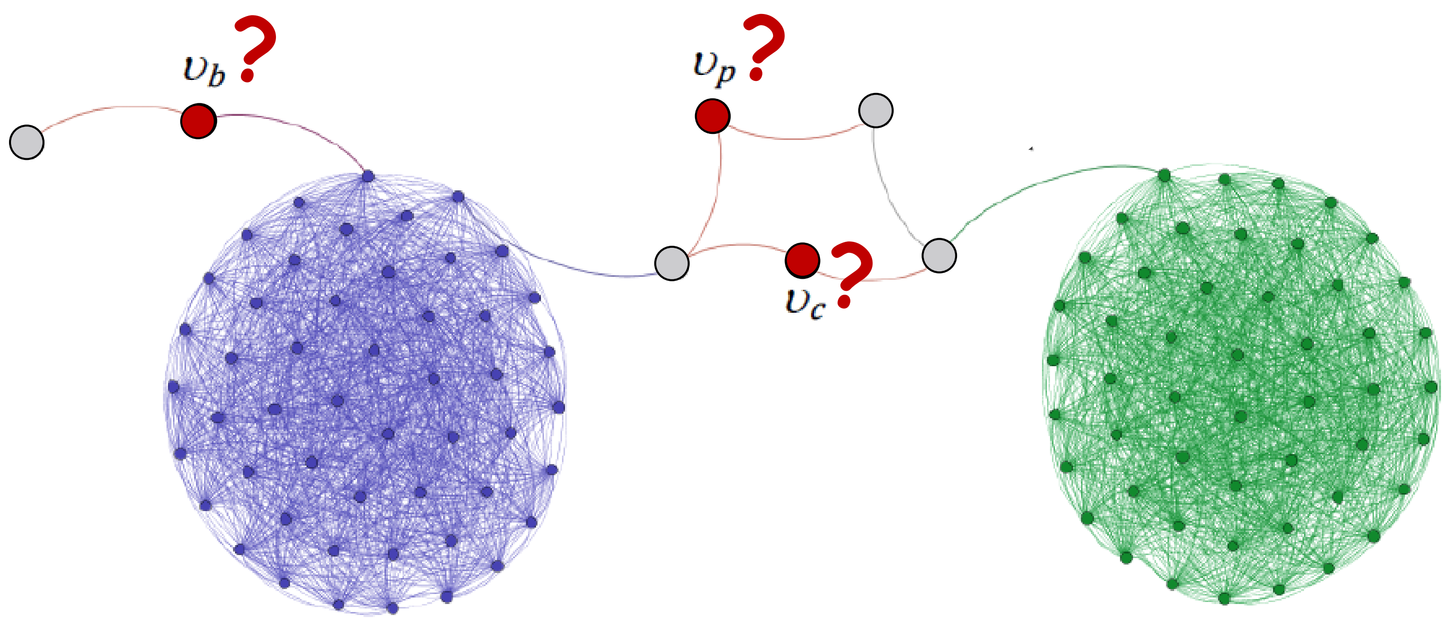

Medeiros, D. S. V., Campista, M. E. M., Mitton, N., Amorim, M. D., Pujolle, G., "*Weighted Betweenness For Multipath Networks*", IEEE Global Information Infrastructure and Networking Symposium (IEEE GIIS 2016), Porto, Portugal, October 2016

Vehicular Mobility

- Nevertheless, for betweenness...
  - v_b is more important than v_p, as it keeps the distant node connected...

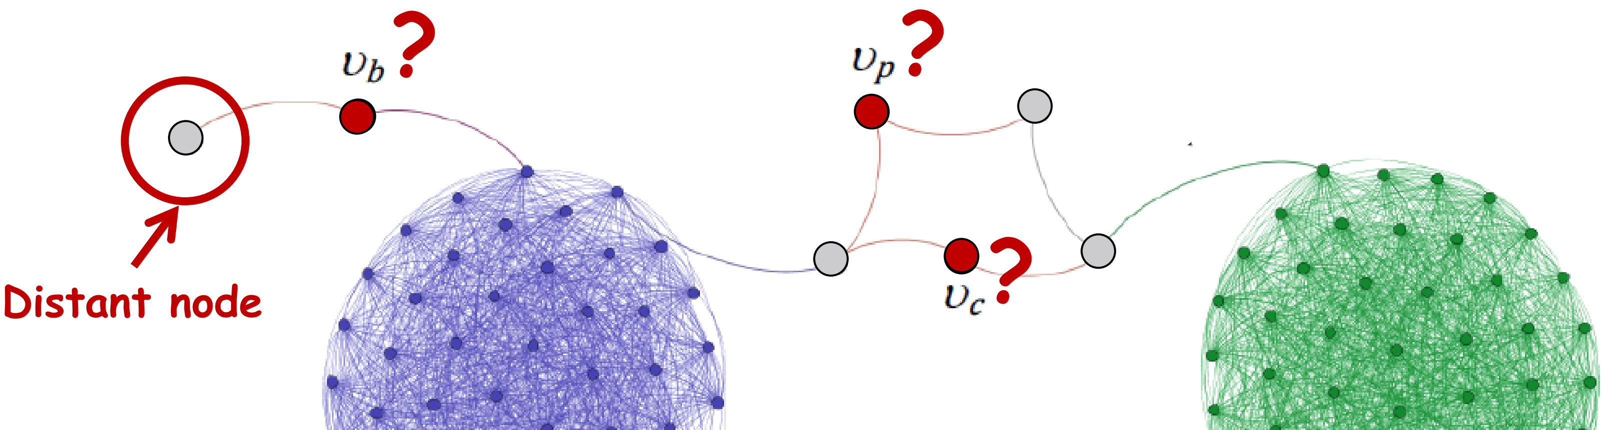

Why v_p is less important than v_b if it is so similar to v_c?

Medeiros, D. S. V., Campista, M. E. M., Mitton, N., Amorim, M. D., Pujolle, G., "*Weighted Betweenness For Multipath Networks*", IEEE Global Information Infrastructure and Networking Symposium (IEEE GIIS 2016), Porto, Portugal, October 2016

Vehicular Mobility

- New metric: **Spread betweenness**

N	Traditional	Distance-scaled	Spread
v_7	48	9,4	53,1
v_9	8	2,2	67,8
v_{16}	15	3,5	15,0

$$w(v_k) = \sum_{\substack{i=1 \\ i \neq k}}^{|\mathcal{V}|} \sum_{\substack{j=1 \\ j \neq i \\ j \neq k}}^{|\mathcal{V}|} \frac{n^+_{i,j}}{n_{i,j}(v_k)} \times \frac{\delta^+_{i,j}}{\delta^+_{i,k} + \delta^+_{k,j}}$$

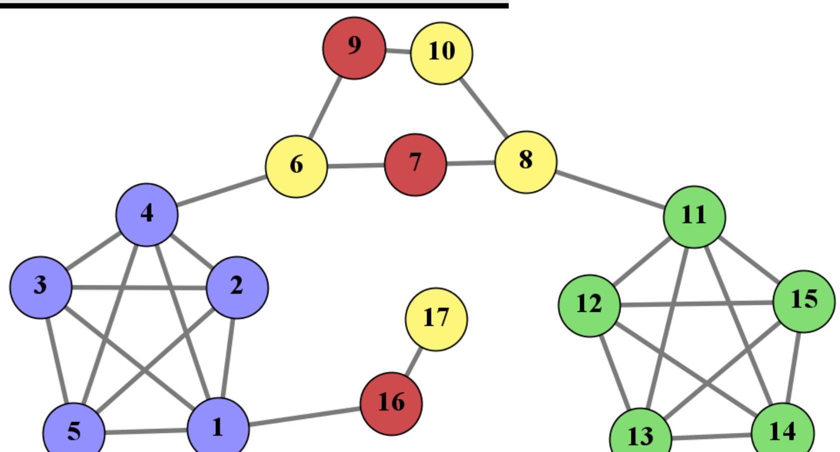

Medeiros, D. S. V., Campista, M. E. M., Mitton, N., Amorim, M. D., Pujolle, G., "*Weighted Betweenness For Multipath Networks*", IEEE Global Information Infrastructure and Networking Symposium (IEEE GIIS 2016), Porto, Portugal, October 2016

My Ongoing Work

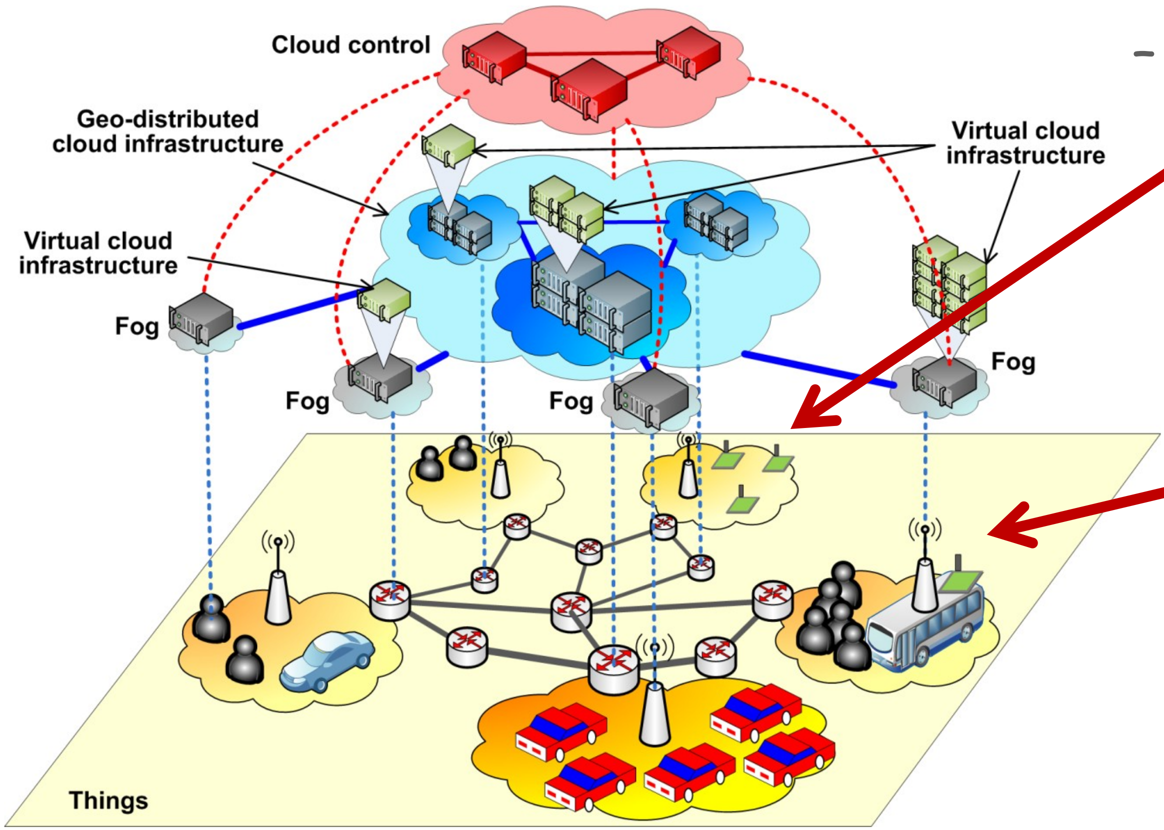

- Trace merging before data transmission …
  - Should we send to the cloud all traces?
- Vehicular sensing…
  - Should we send to the cloud all the data sensed?

Trace Merging Before Data Transmission

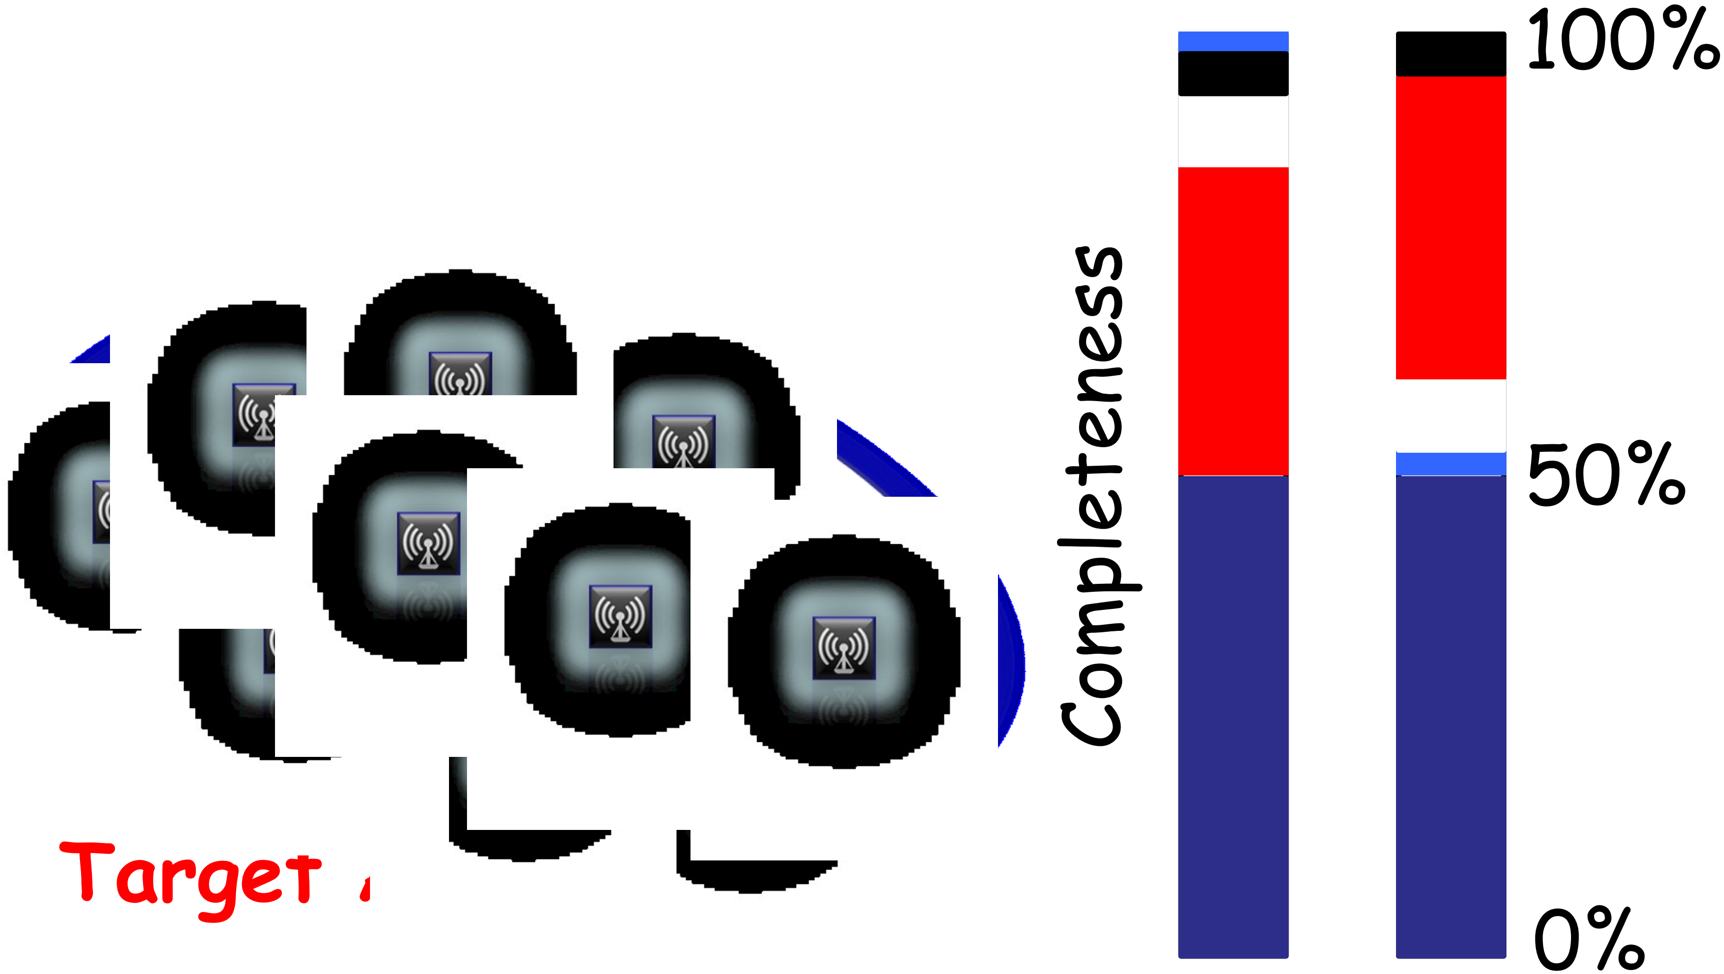

Sammarco, M., Campista, M. E. M., Amorim, M. D., *"Scalable Wireless Traffic Capture Through Community Detection and Trace Similarity"*, IEEE Transactions on Mobile Computing, vol. 15, no. 7, pp. 1757-1769, July 2016

Trace Ranking

- Trace ranking: We consider a fully connected graph
  - v_i: is the trace captured by the i-th sensor
  - e_{ij}: has a weight linearly proportional to the similarity between the i-th and the j-th trace

Ranking all the nodes in according to the minimum Hamiltonian path is a good way to iteratively select traces to merge

Sammarco, M., Campista, M. E. M., Amorim, M. D., "*Scalable Wireless Traffic Capture Through Community Detection and Trace Similarity*", IEEE Transactions on Mobile Computing, vol. 15, no. 7, pp. 1757-1769, July 2016

Experimental Setup

Scenario 1

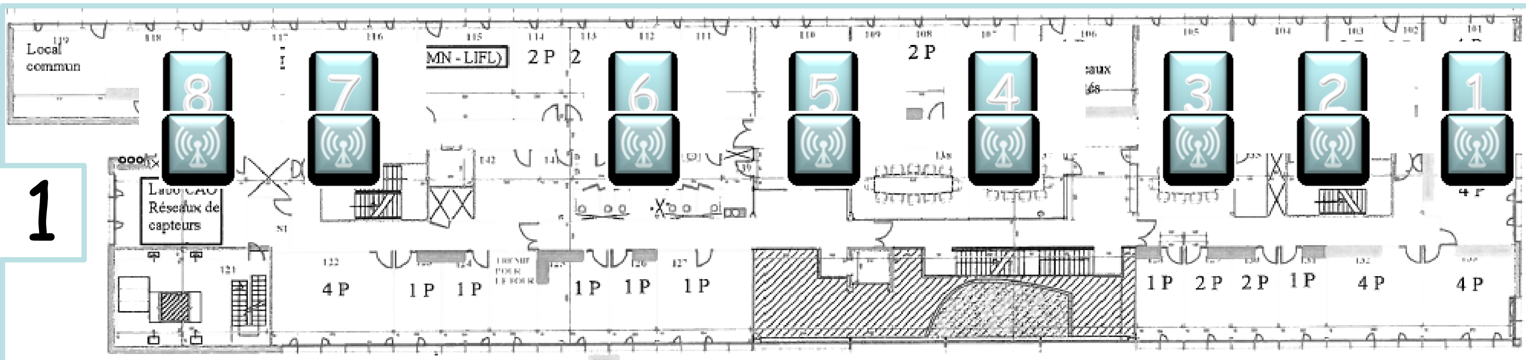

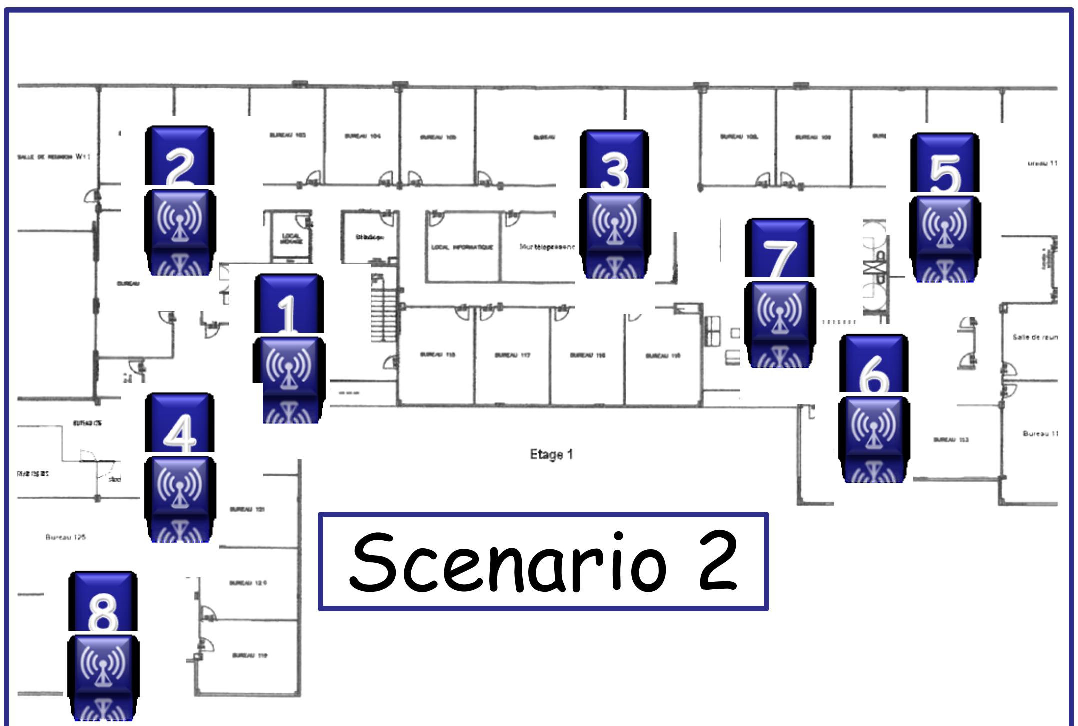

Scenario 2

8 sensors
100 minutes

Similarity: Scenario 1

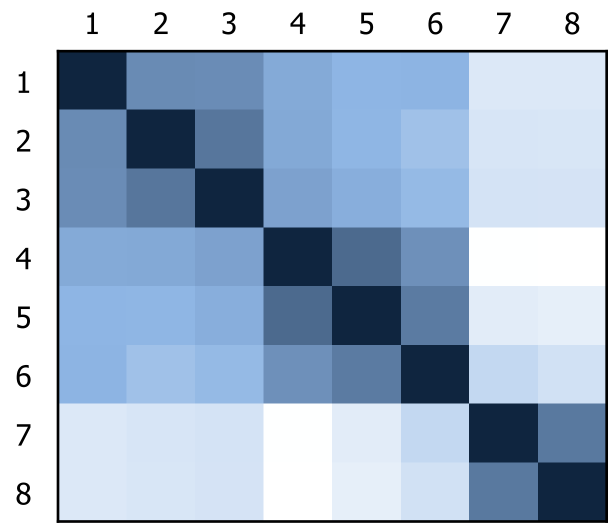

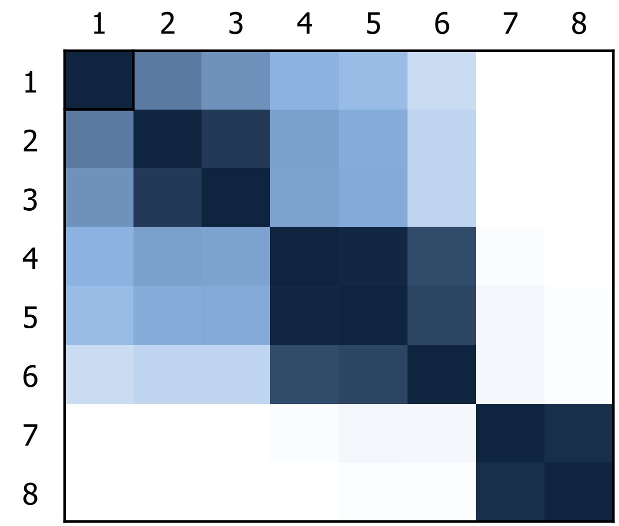

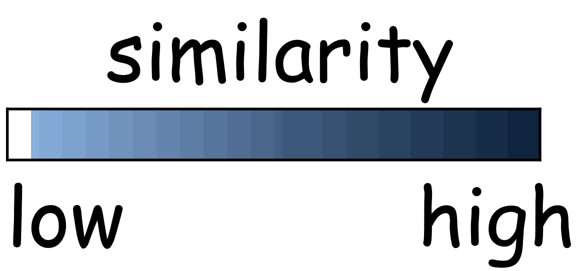

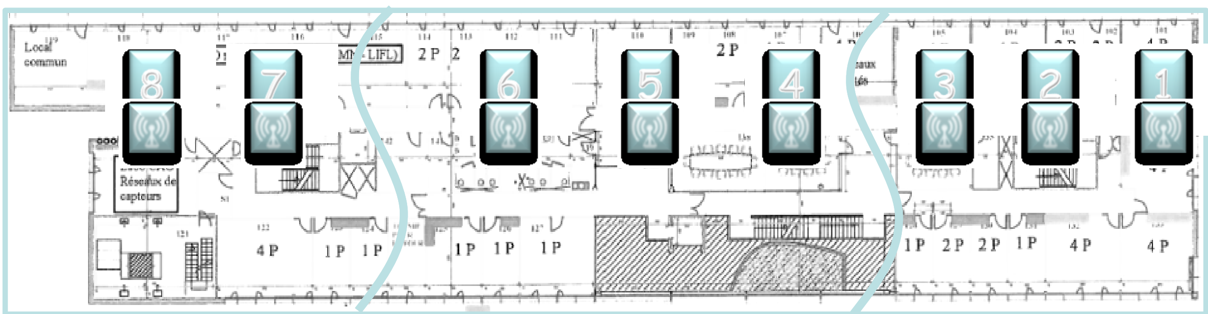

Sammarco, M., Campista, M. E. M., Amorim, M. D., *"Scalable Wireless Traffic Capture Through Community Detection and Trace Similarity"*, IEEE Transactions on Mobile Computing, vol. 15, no. 7, pp. 1757-1769, July 2016

Similarity: Scenario 2

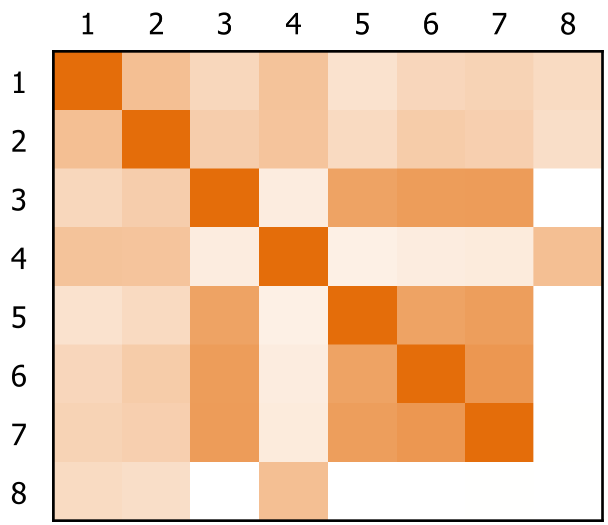

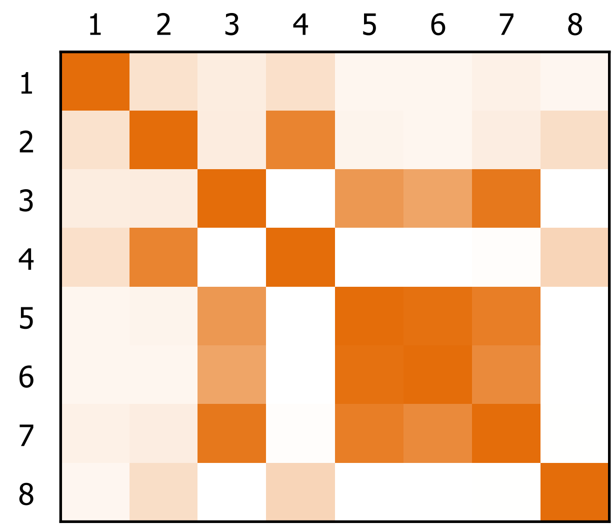

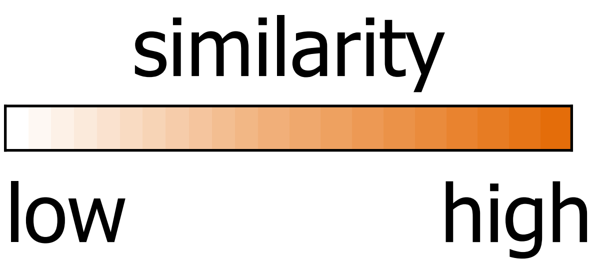

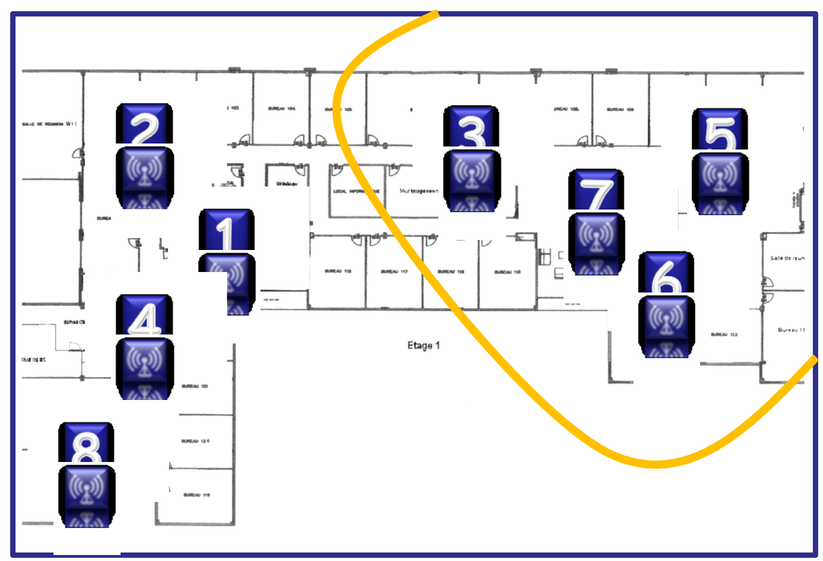

Minimum Hamiltonian Path

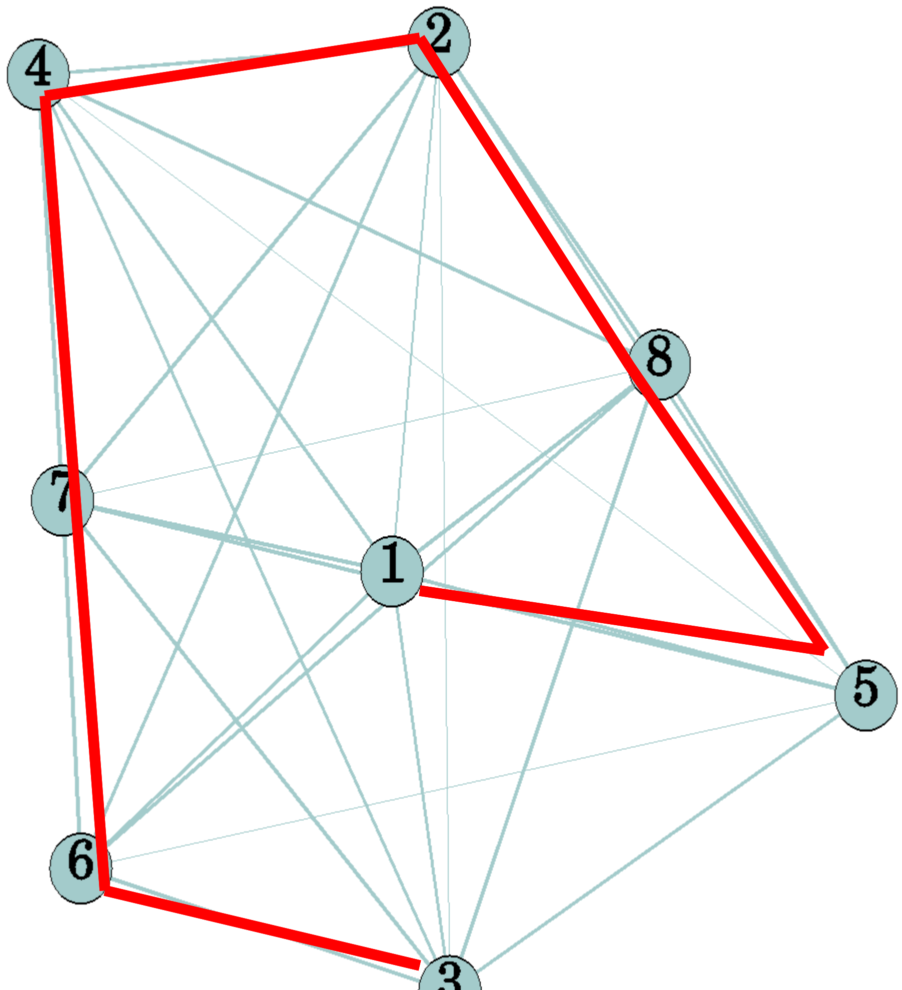

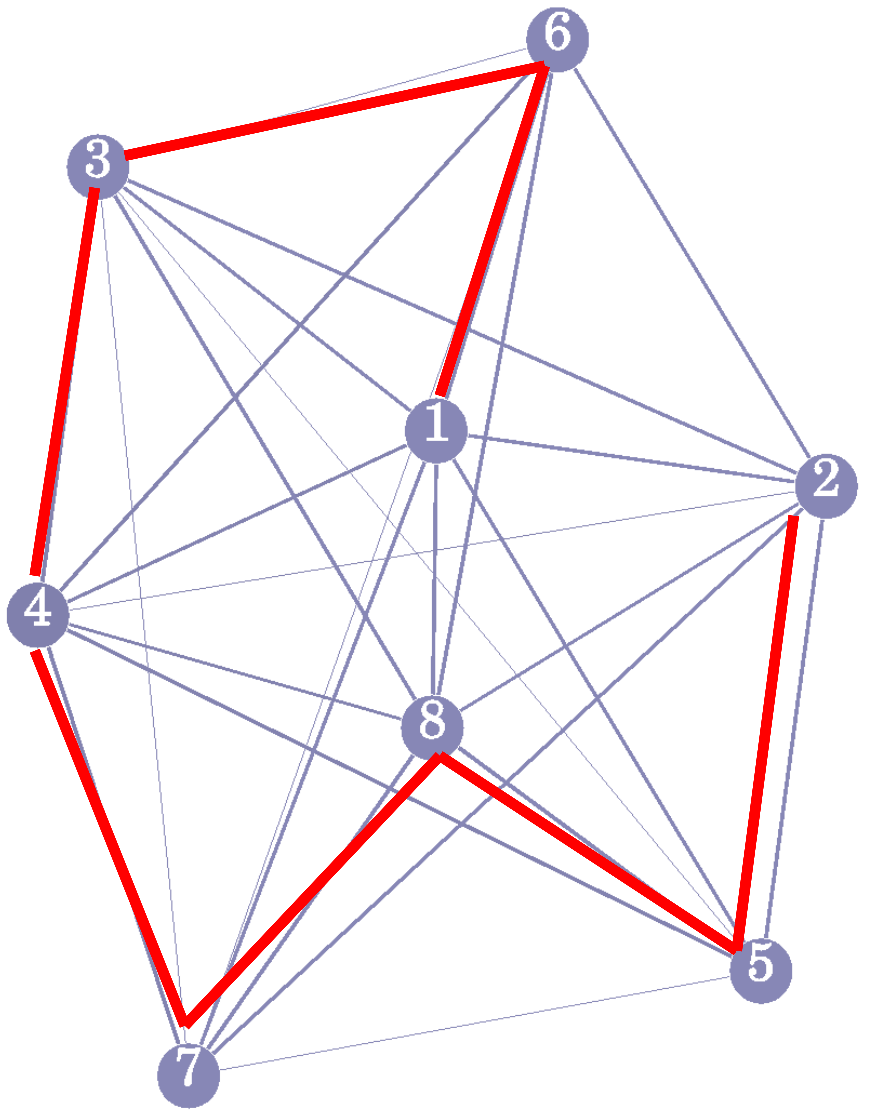

Results

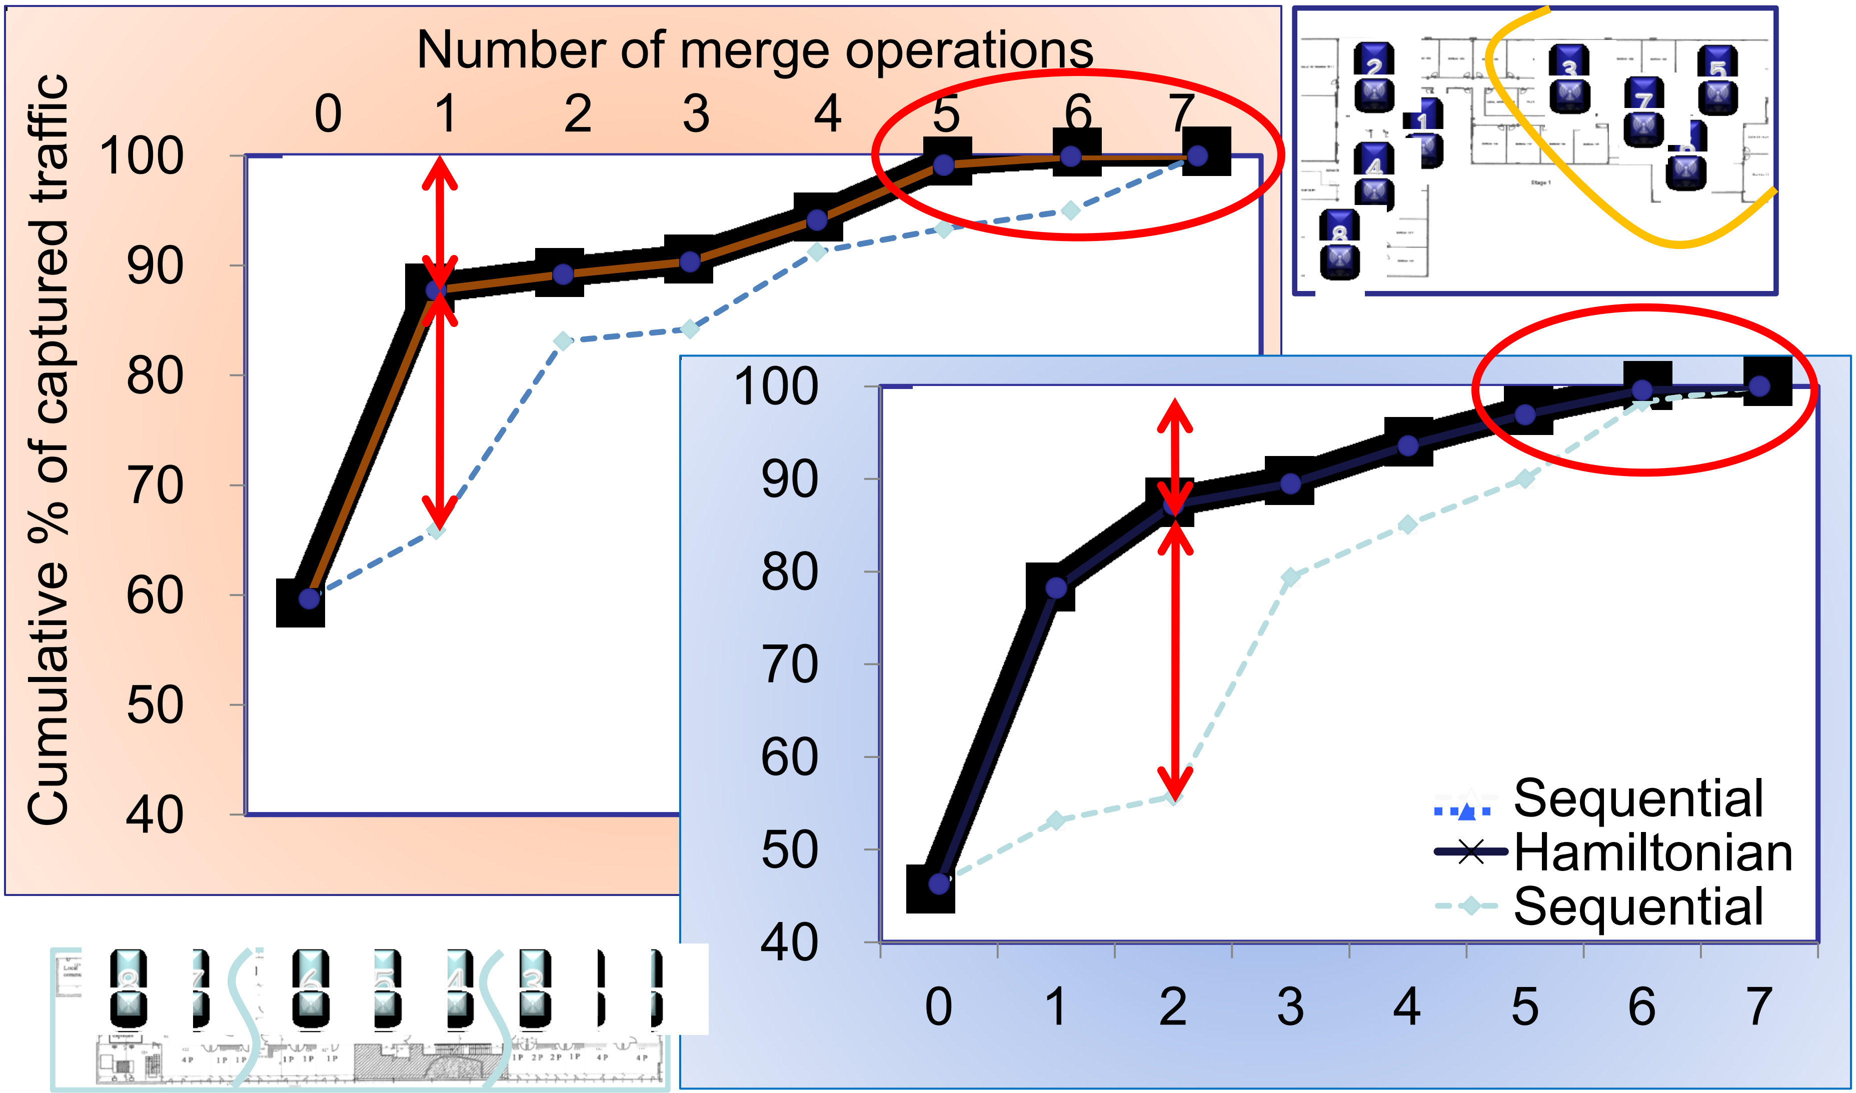

Number of merge operations
Cumulative % of captured traffic
0
1
2
3
4
5
6
7
100
90
80
70
60
50
40
Sequential
Hamiltonian
Sequential

Vehicular Sensing

- PLATT (*Piecewise Linear Automatically Tuned Trending*)
  - Compresses data with automatic parameterization
  - Approximates the signal as sequence of lines
  - Compresses finite-length signals
    - Signals are compresses in batches
    - Each batch is processed as a finite-length signal

Lages, J. E. M., "*PLATT: Um Novo Algoritmo de Compressão de Dados de Processo com Parametrização Automática",* Undergraduate Final Project, Universidade Federal do Rio de Janeiro, Engenharia Eletrônica e de Computação, September 2016 (Supervisor: MEM Campista)

PLATT

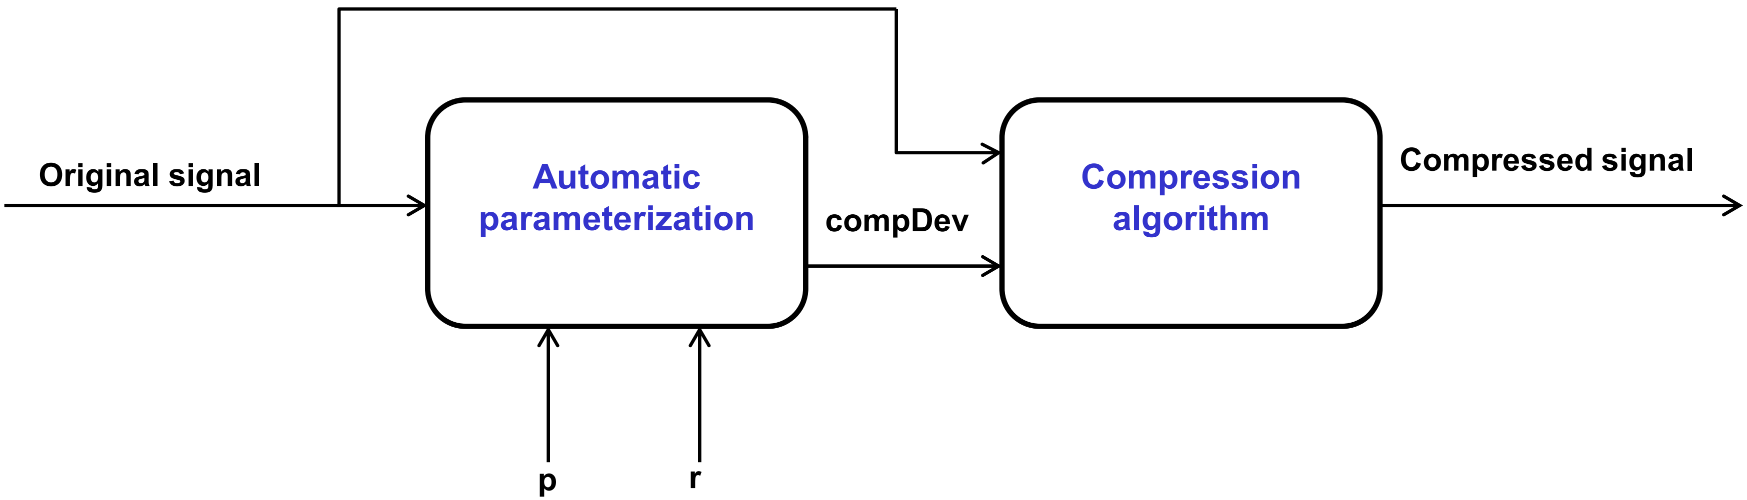

Lages, J. E. M., "*PLATT: Um Novo Algoritmo de Compressão de Dados de Processo com Parametrização Automática",* Undergraduate Final Project, Universidade Federal do Rio de Janeiro, Engenharia Eletrônica e de Computação, September 2016 (Supervisor: MEM Campista)

PLATT

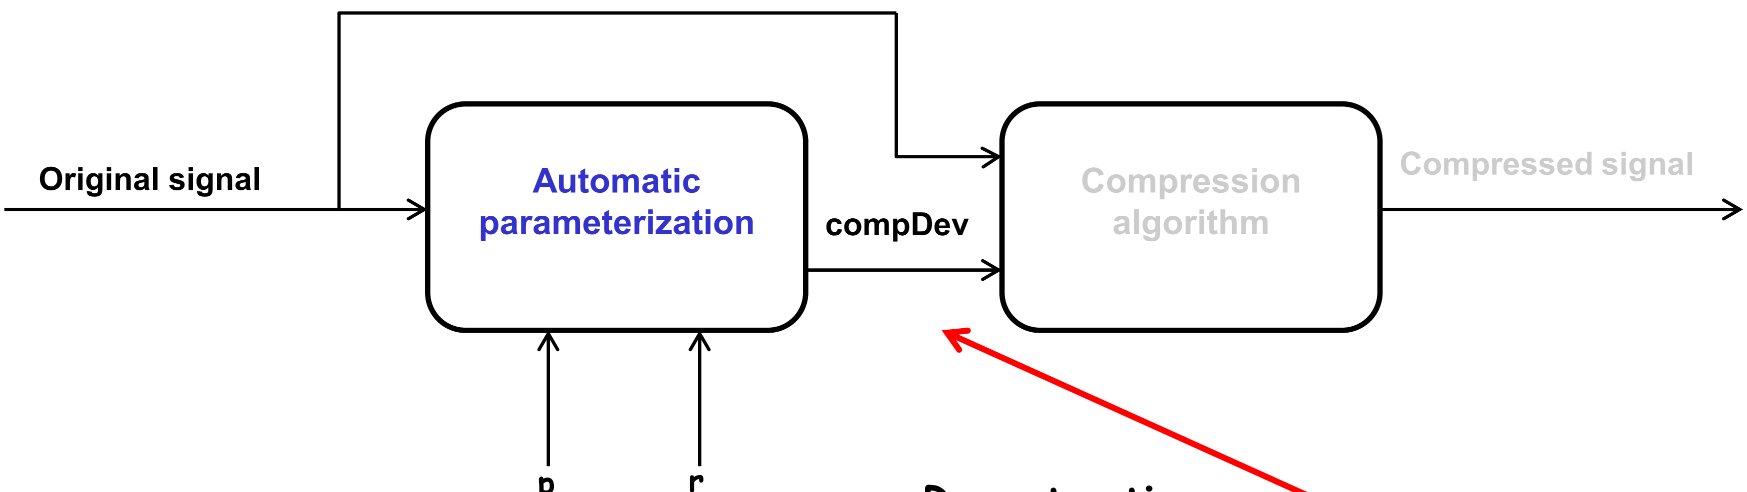

Repeat r times

1. Randomly choose p*N points in both halves
2. Adjust the compDev to each half
3. Compute the average value to each half

Compute the average value to each run r

PLATT

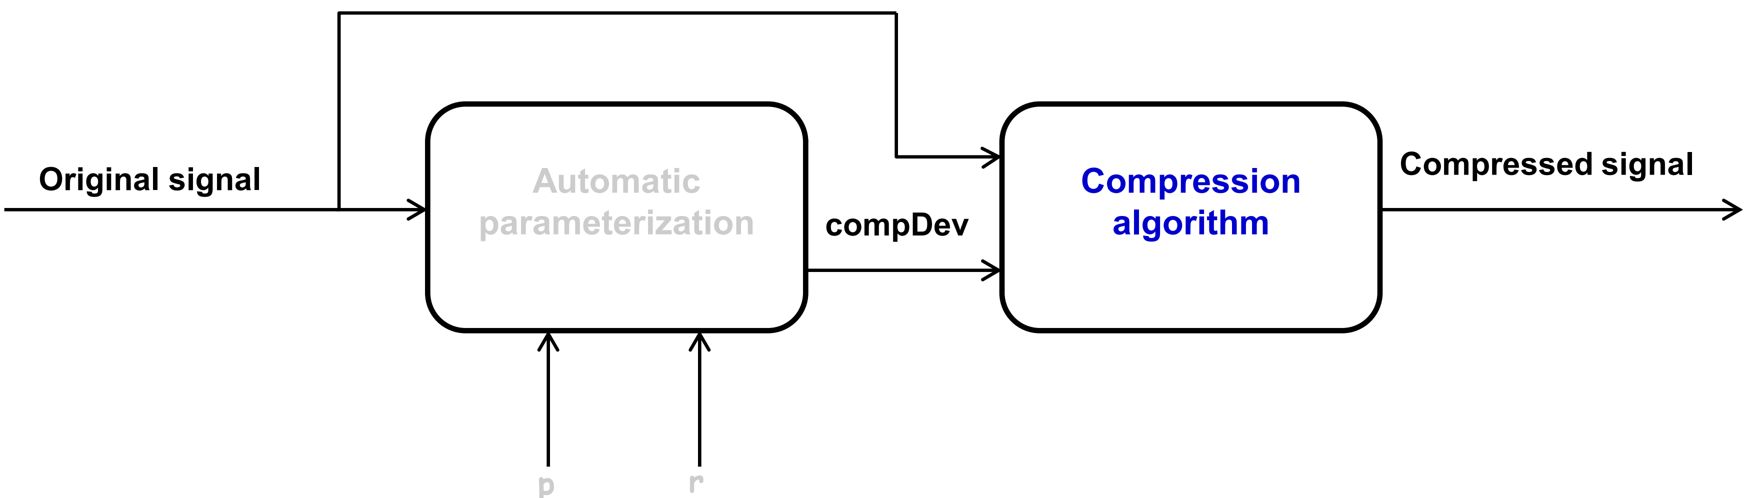

Lages, J. E. M., "*PLATT: Um Novo Algoritmo de Compressão de Dados de Processo com Parametrização Automática",* Undergraduate Final Project, Universidade Federal do Rio de Janeiro, Engenharia Eletrônica e de Computação, September 2016 (Supervisor: MEM Campista)

Compression Strategy

Tolerance bars with length equal to 2*compDev

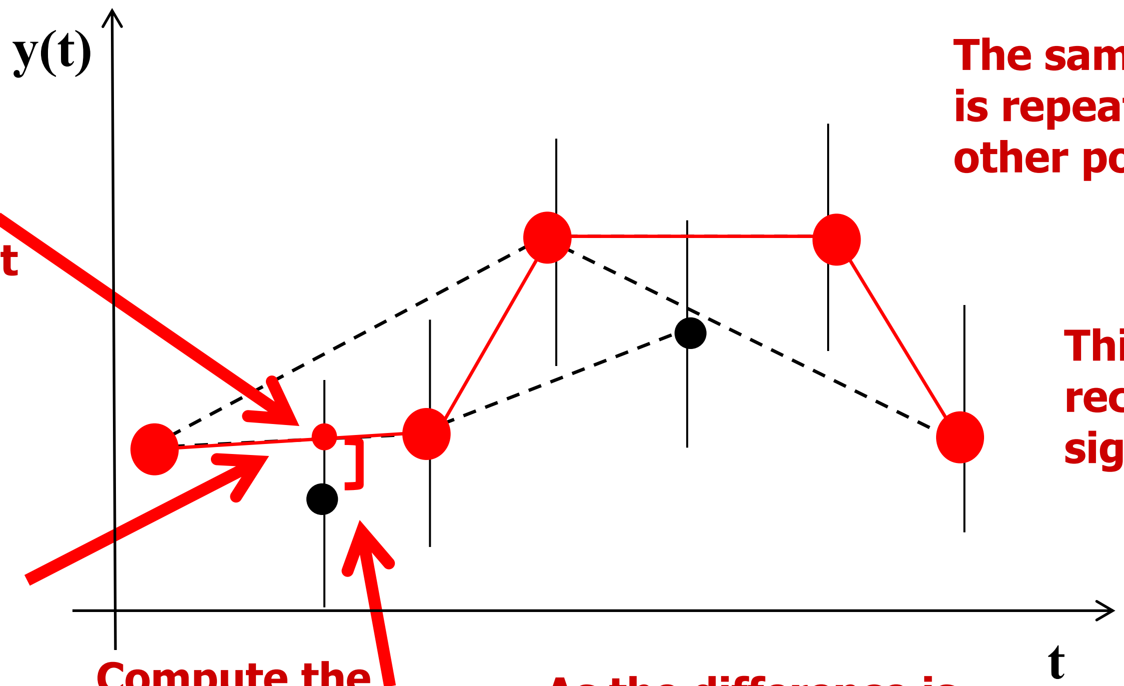

Compute the value in which the test point would have if it were over the line

Compute the line inclination between the last stored point and current point

Compute the difference between the value computed by the line and the real value of the point

As the difference is lower than compDev, the point is discarded

The same strategy is repeated to all other points

This would be the reconstructed signal

Real Signals

- The following signals were used:

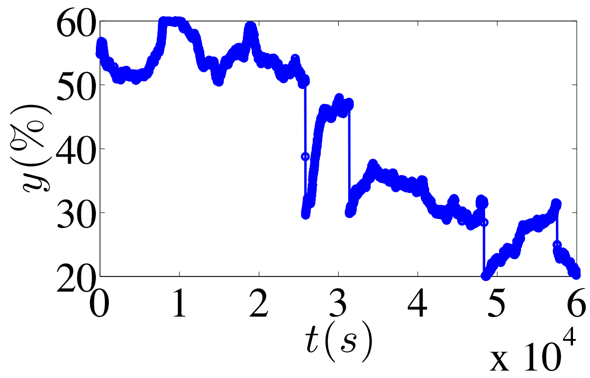

R_1 (posição de uma válvula)

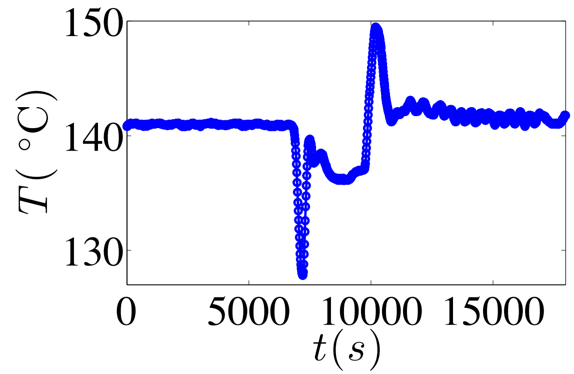

R_2 (temperatura)

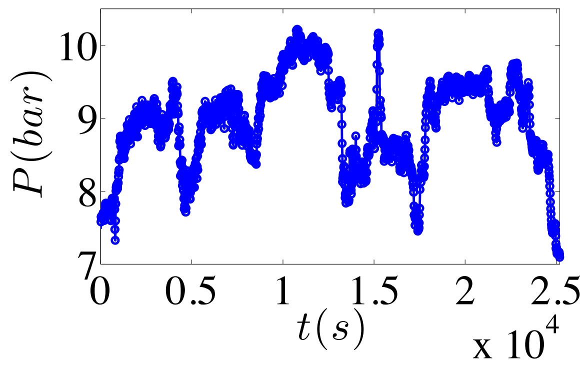

R_3 (pressão)

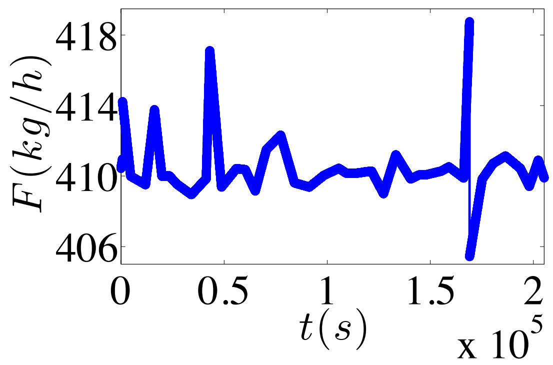

R_4 (vazão)

Impact of *p* variation

- Compression Ration (RC) and Reconstruction Error (ER) for R_1

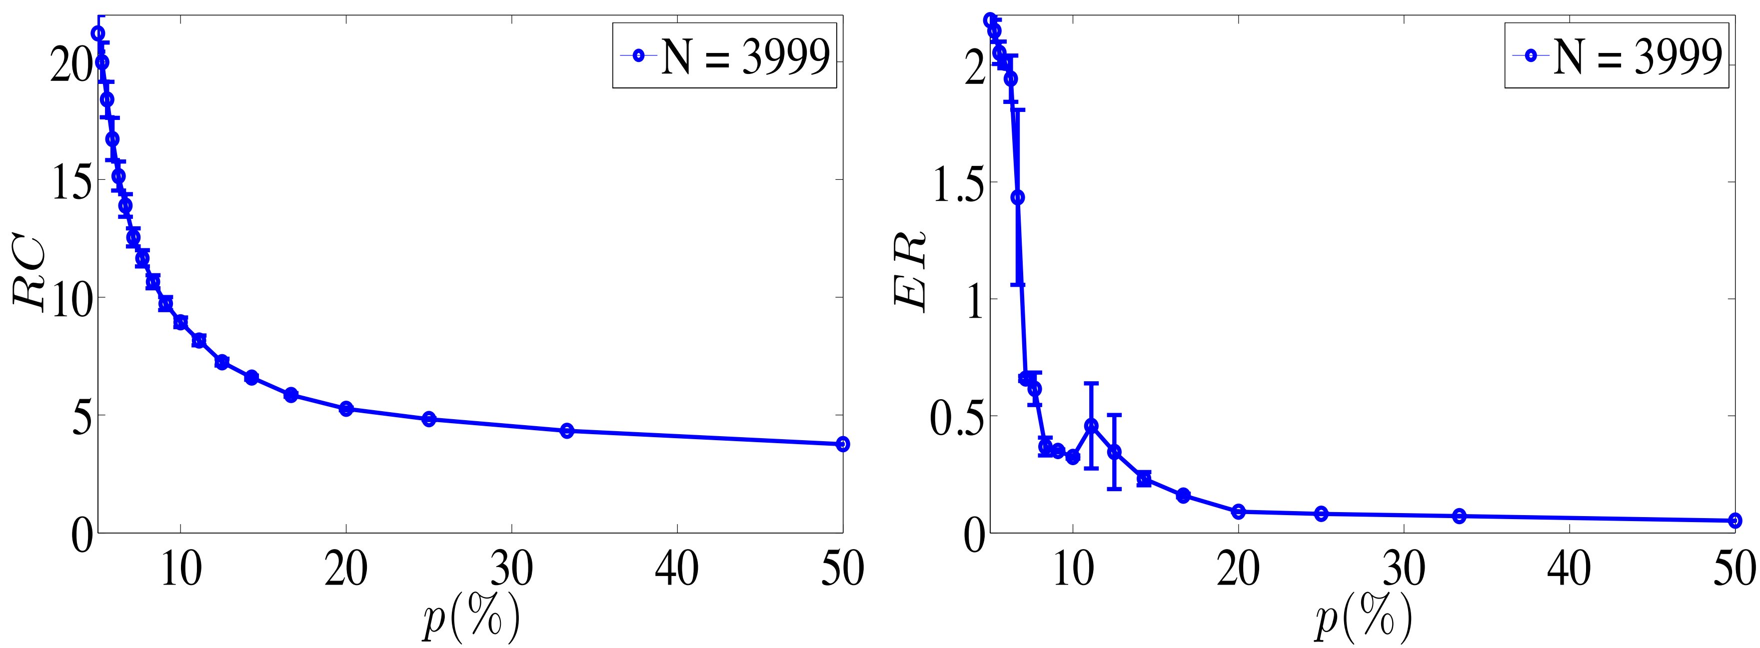

Impact of *p* variation

- Compression of R_1 with p = 10%

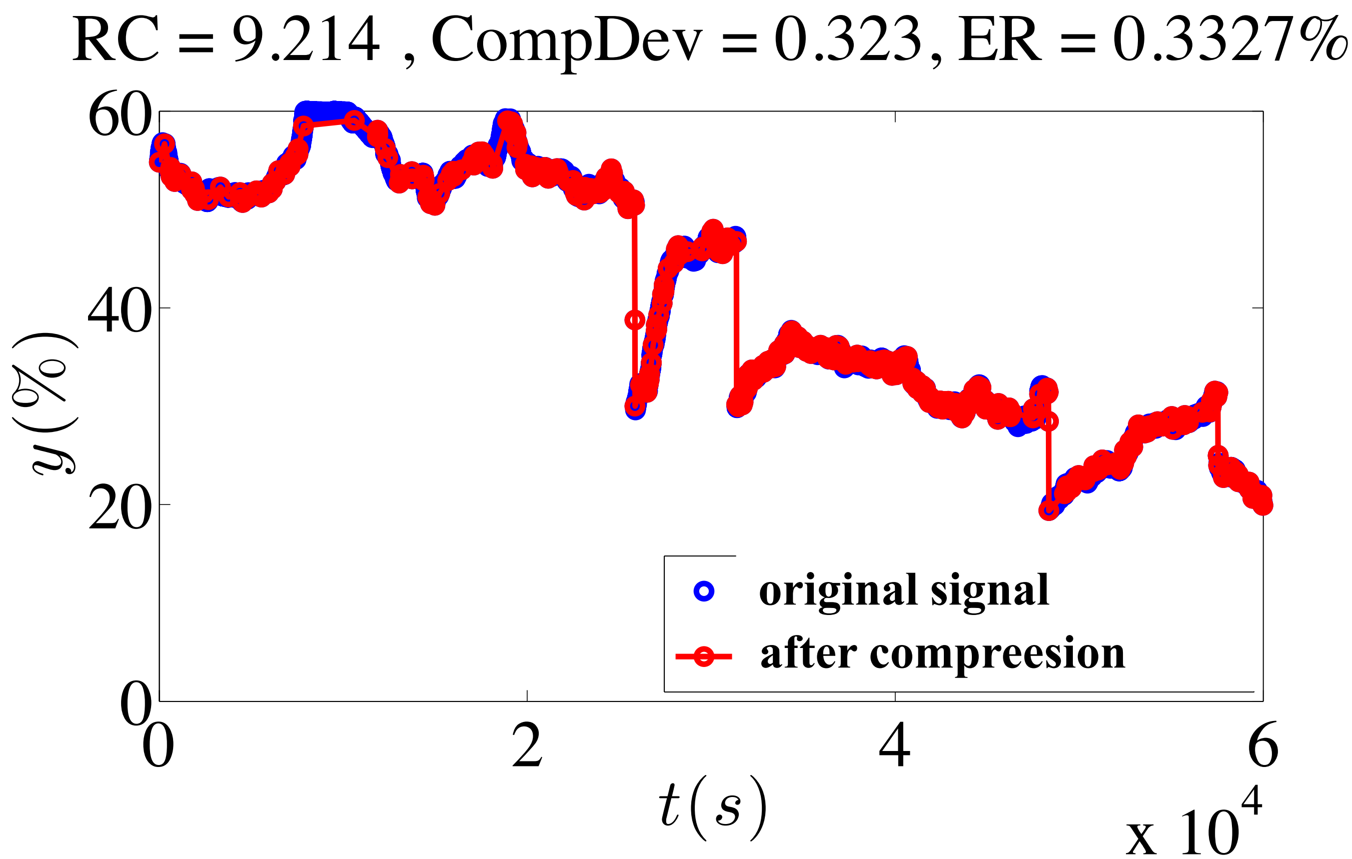

My Ongoing Work

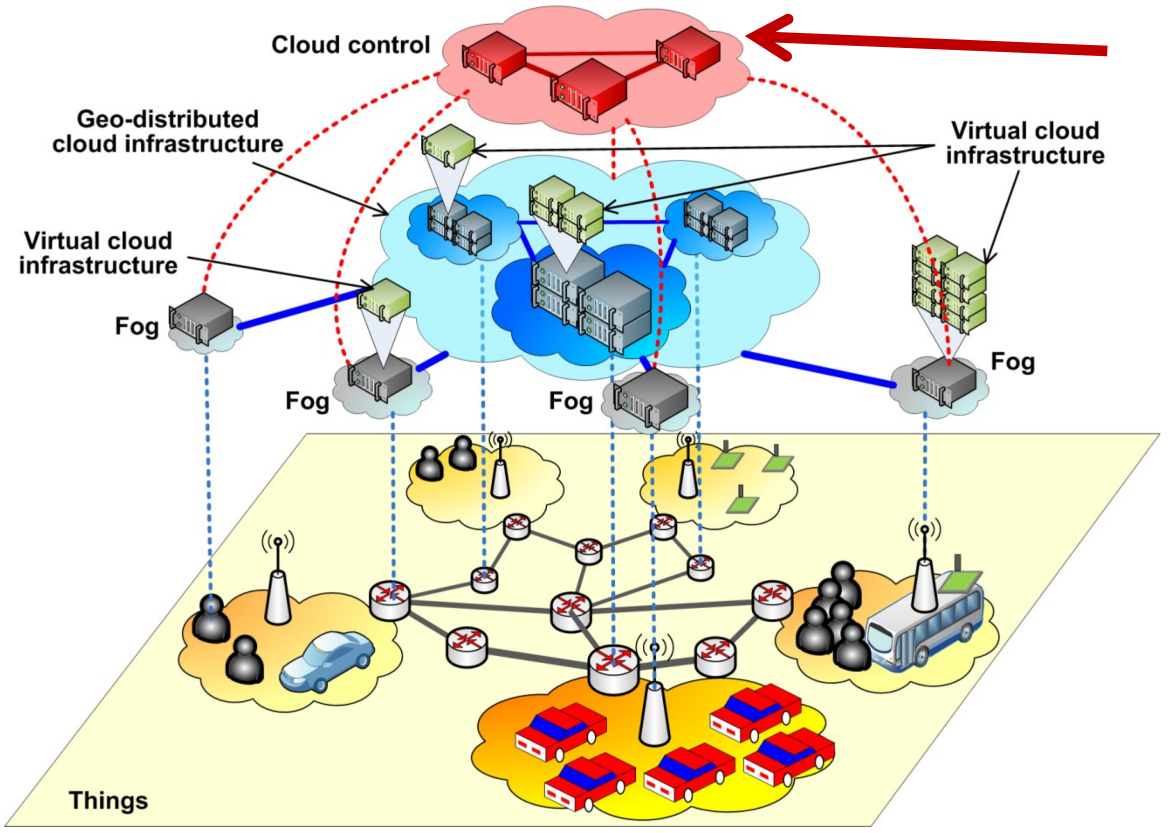

- Collaborative clouds…
  - Should we care about the control traffic on cloud controllers?

Collaborative Clouds

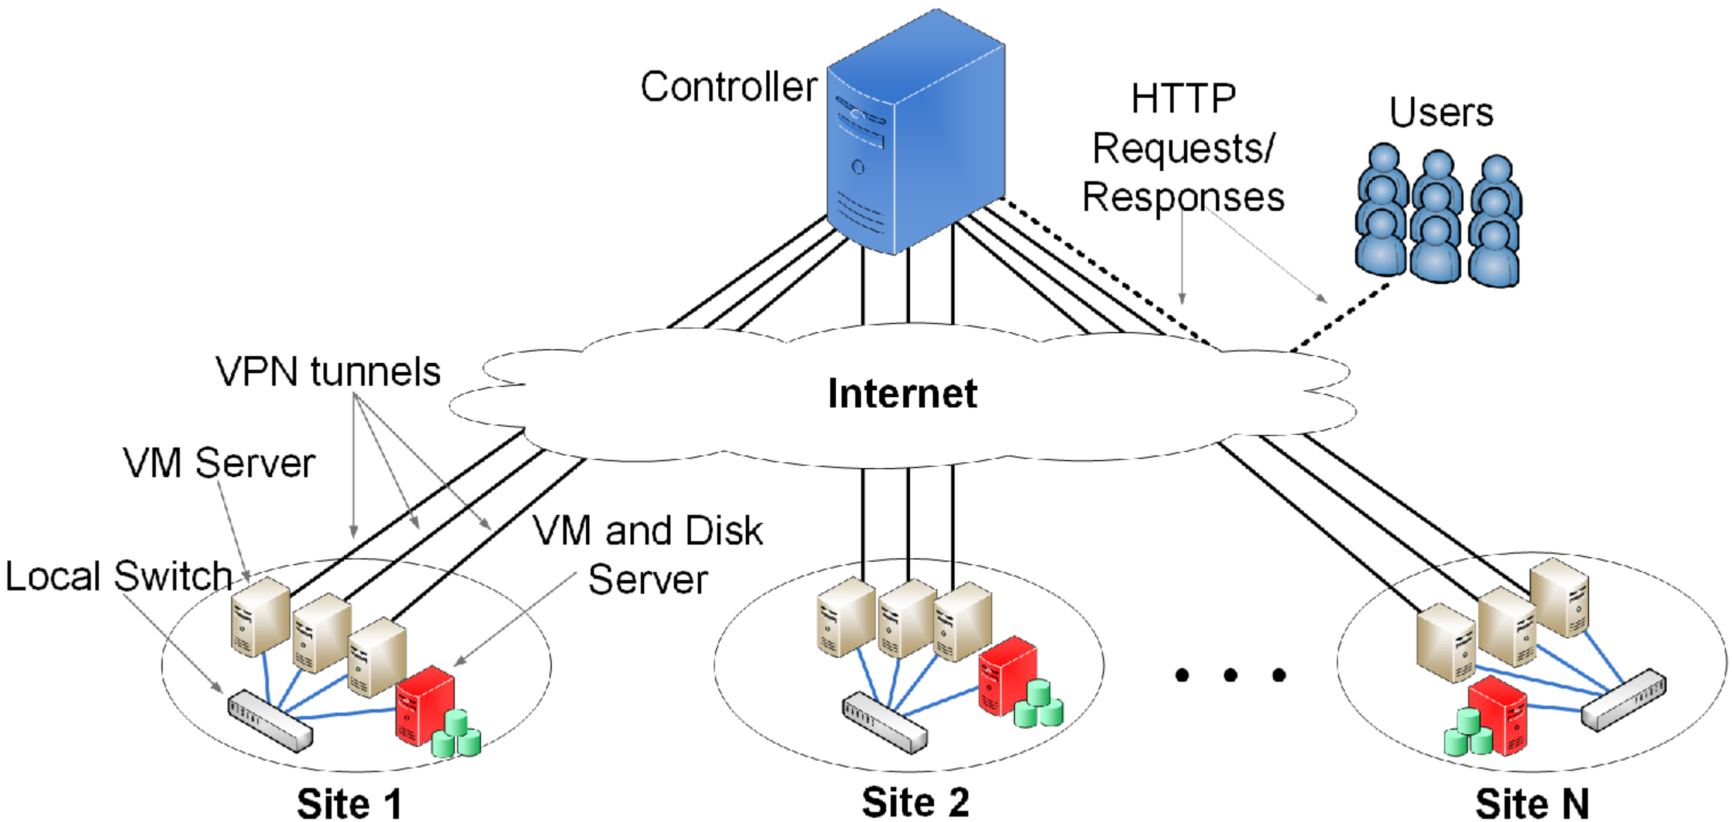

Sciammarella, T., Couto, R. S., Rubinstein, M. G., Campista, M. E. M., Costa, L. H. M. K., *"Analysis of Control Traffic in a Geo-distributed Collaborative Cloud",* IEEE CloudNet, Pisa, Italy, October 2016

OpenStack Projects

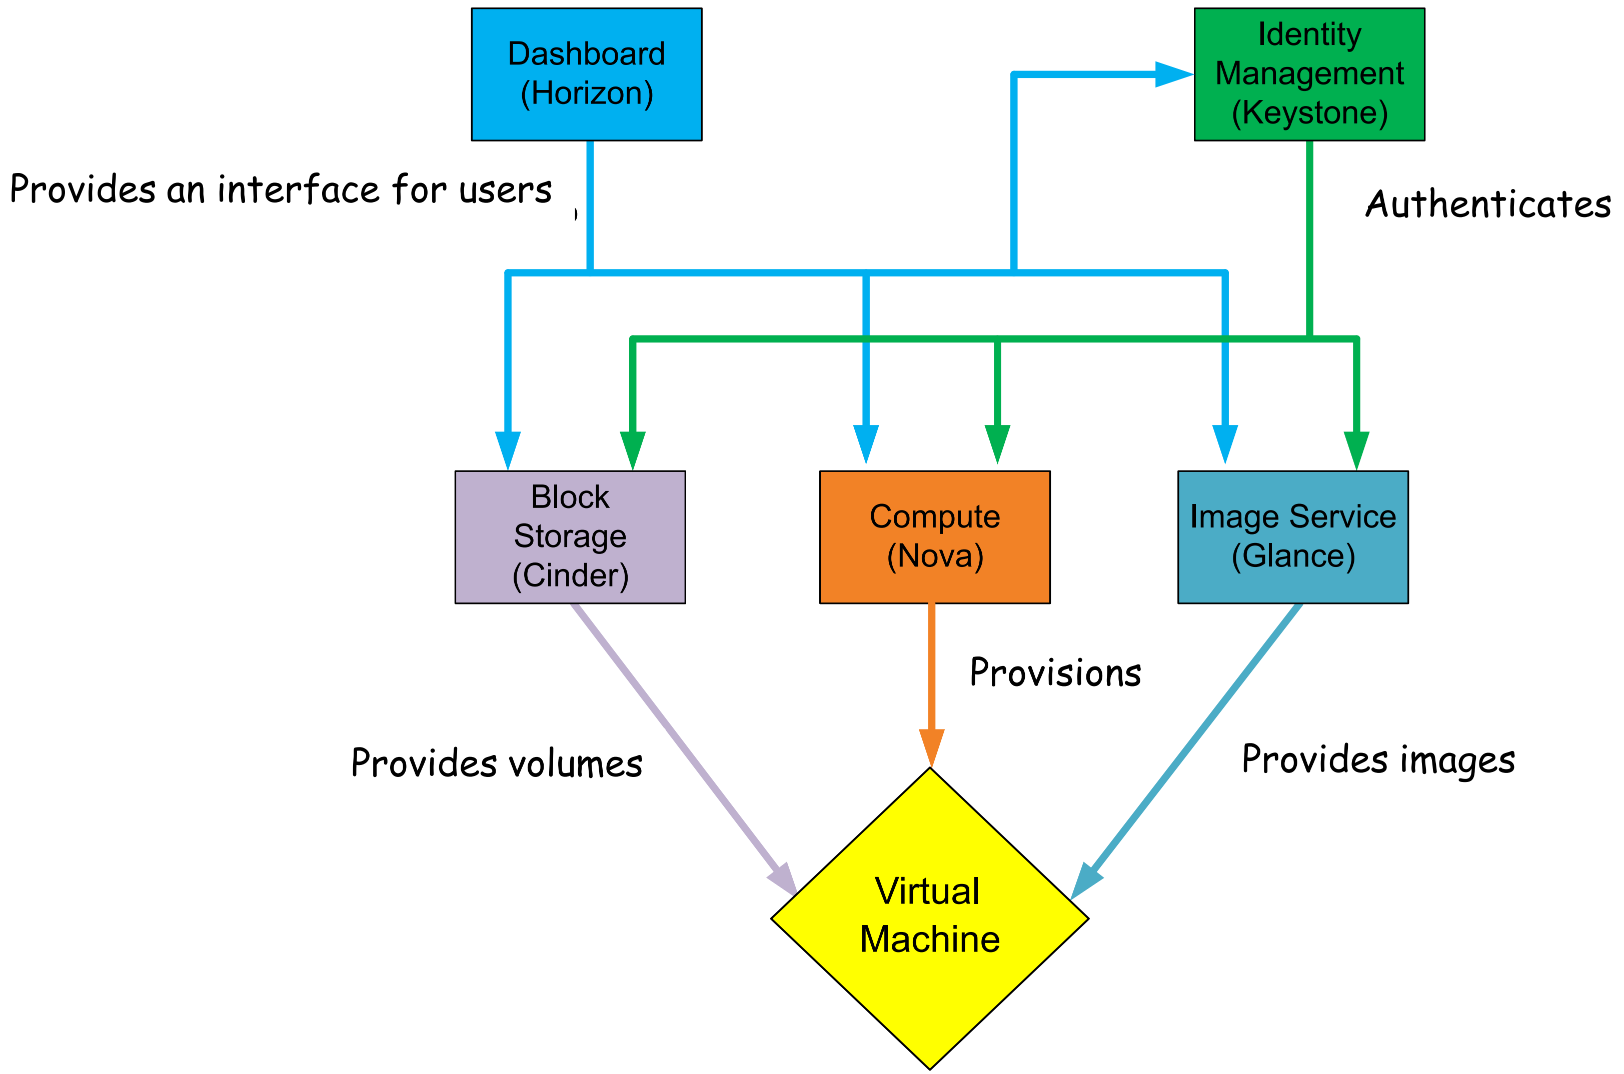

Sciammarella, T., Couto, R. S., Rubinstein, M. G., Campista, M. E. M., Costa, L. H. M. K., *"Analysis of Control Traffic in a Geo-distributed Collaborative Cloud",* IEEE CloudNet, Pisa, Italy, October 2016

OpenStack Modules

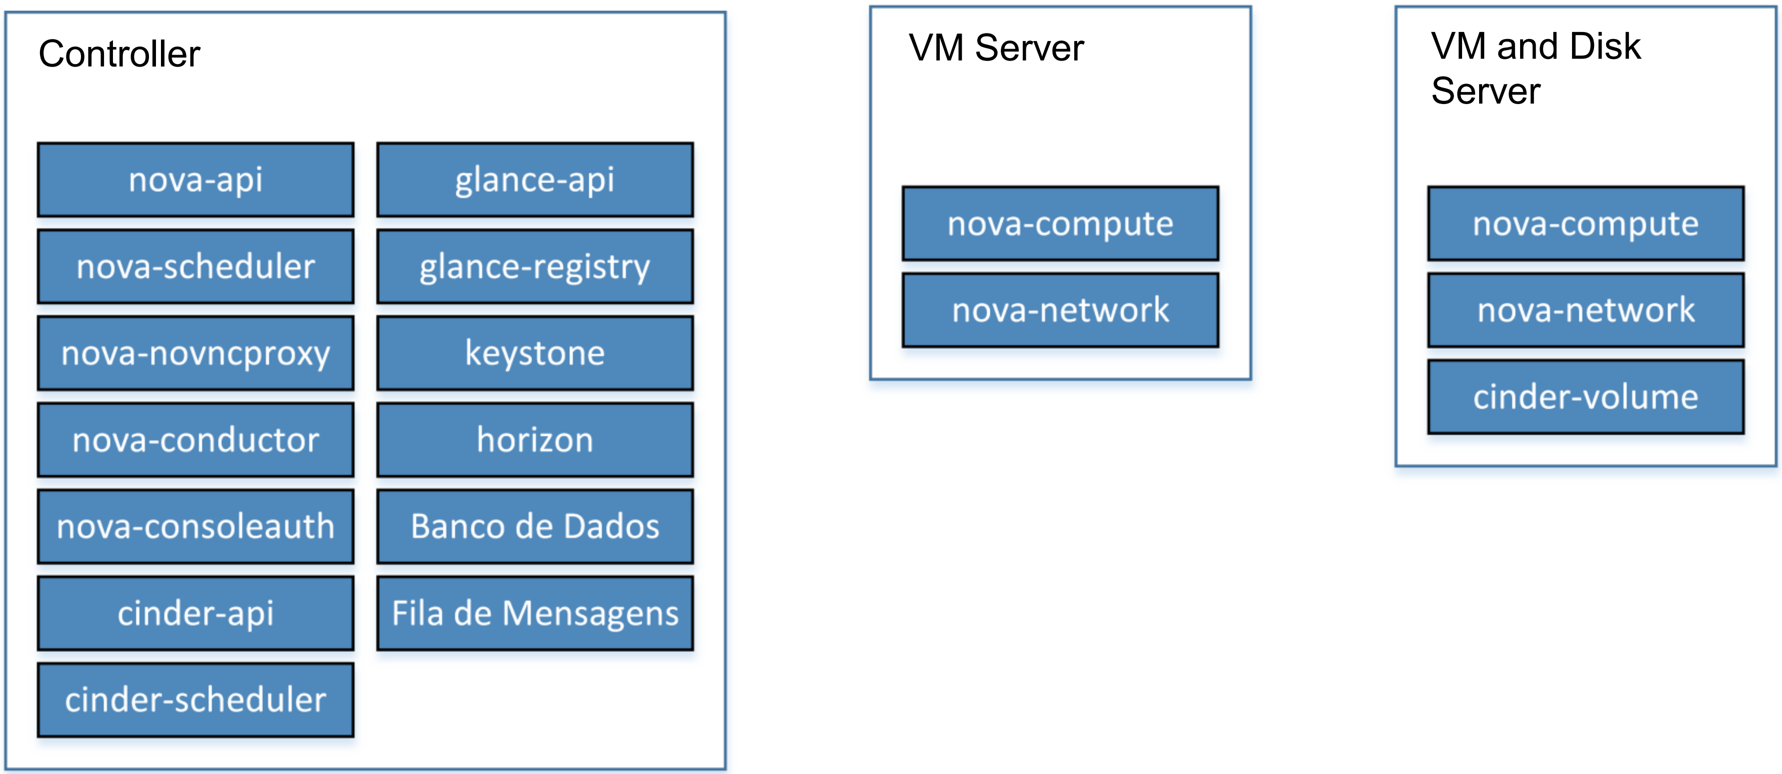

Sciammarella, T., Couto, R. S., Rubinstein, M. G., Campista, M. E. M., Costa, L. H. M. K., *"Analysis of Control Traffic in a Geo-distributed Collaborative Cloud",* IEEE CloudNet, Pisa, Italy, October 2016

Communications between OpenStack Modules

- Message queue: RabbitMQ
- Database: MySQL

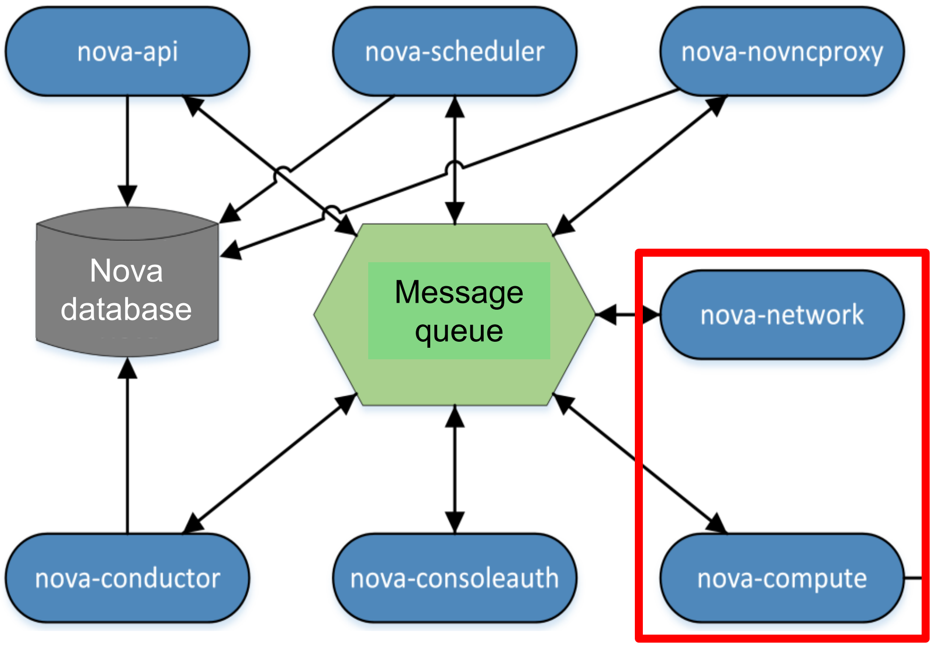

Sciammarella, T., Couto, R. S., Rubinstein, M. G., Campista, M. E. M., Costa, L. H. M. K., "*Analysis of Control Traffic in a Geo-distributed Collaborative Cloud*", IEEE CloudNet, Pisa, Italy, October 2016

Testbed

- Physical setting

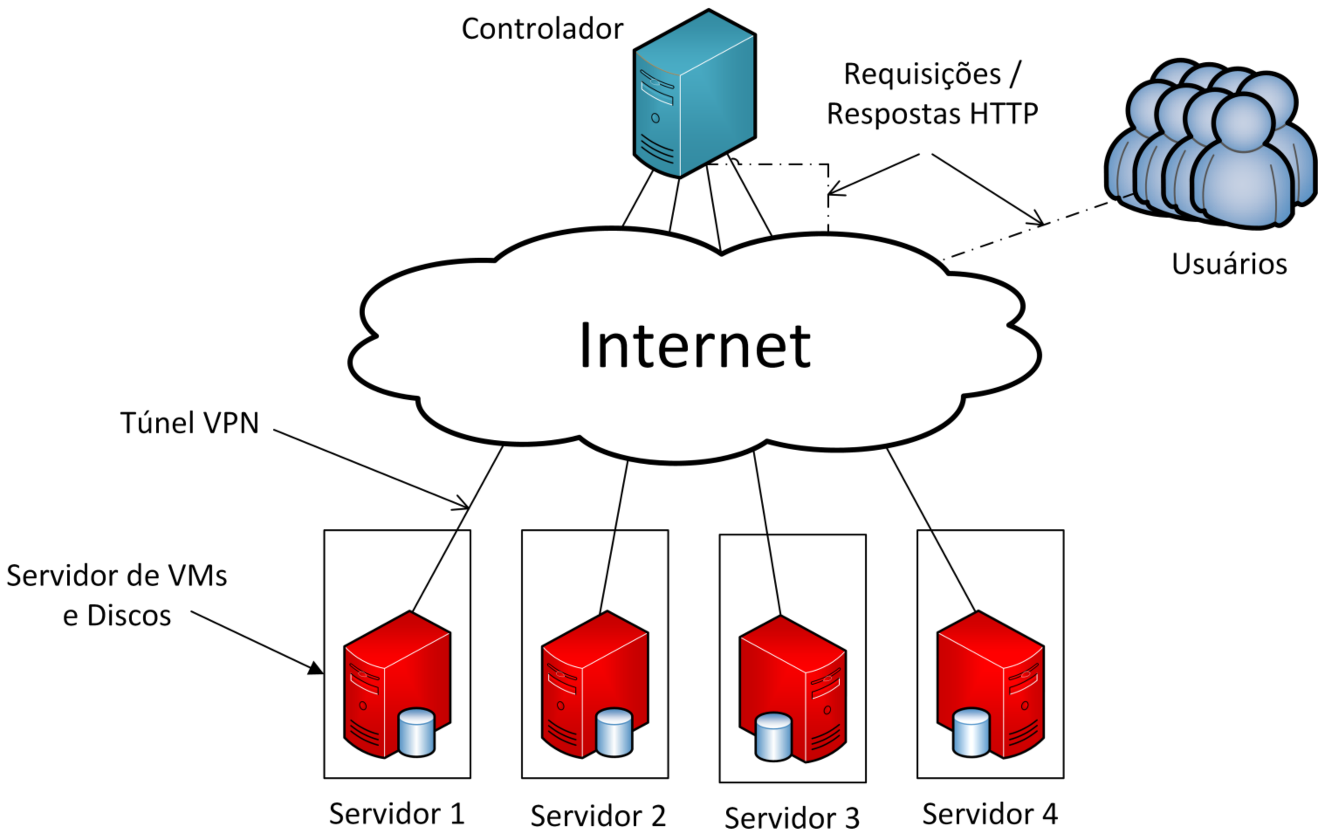

Sciammarella, T., Couto, R. S., Rubinstein, M. G., Campista, M. E. M., Costa, L. H. M. K., *"Analysis of Control Traffic in a Geo-distributed Collaborative Cloud",* IEEE CloudNet, Pisa, Italy, October 2016

Server Periodical Traffic

- Each 10s: Services state update
- Each 60s: VMs state update

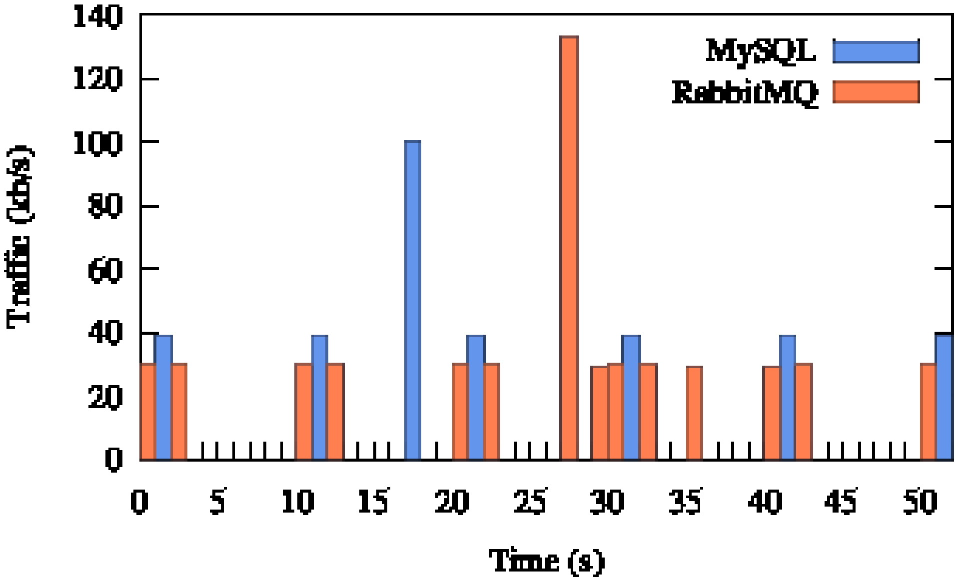

Sciammarella, T., Couto, R. S., Rubinstein, M. G., Campista, M. E. M., Costa, L. H. M. K., “*Analysis of Control Traffic in a Geo-distributed Collaborative Cloud”,* IEEE CloudNet, Pisa, Italy, October 2016

Impact of the Number of VM and Disk Servers

- Servers without instantiated VMs
- Traffic measured during 60s

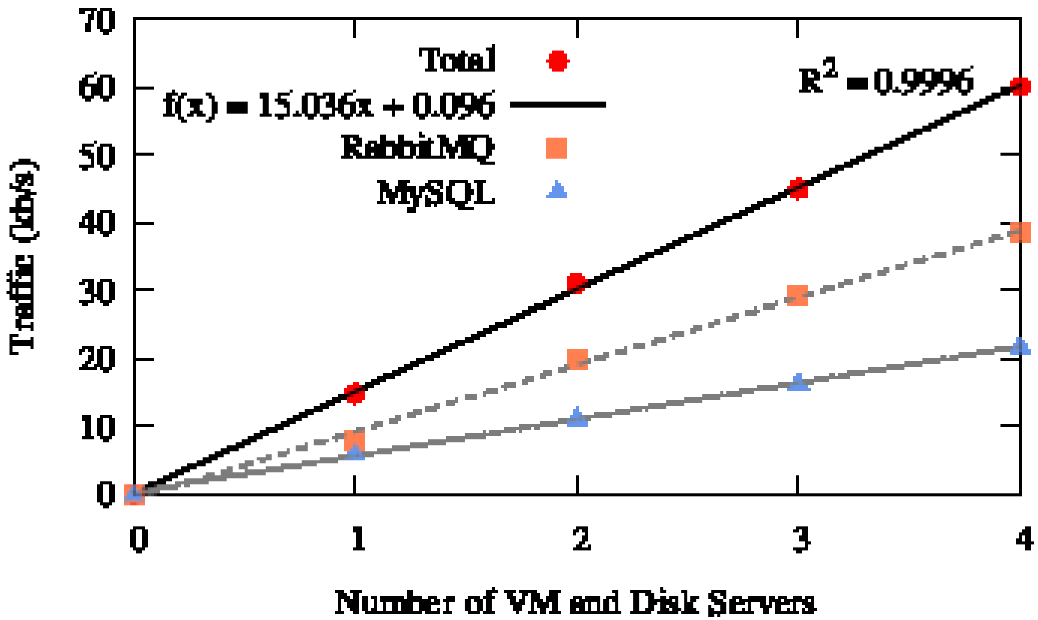

Sciammarella, T., Couto, R. S., Rubinstein, M. G., Campista, M. E. M., Costa, L. H. M. K., "*Analysis of Control Traffic in a Geo-distributed Collaborative Cloud",* IEEE CloudNet, Pisa, Italy, October 2016

Impact of the VM Number per Server

- Linear growth
  - In one server, each VM adds approximately 0.77 kb/s

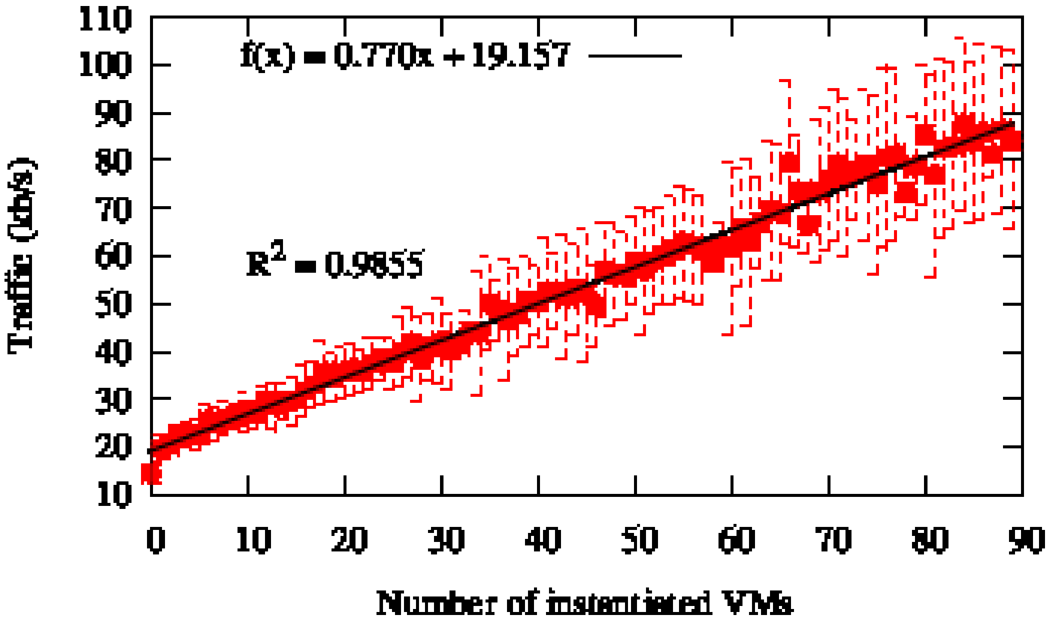

Sciammarella, T., Couto, R. S., Rubinstein, M. G., Campista, M. E. M., Costa, L. H. M. K., "*Analysis of Control Traffic in a Geo-distributed Collaborative Cloud*", IEEE CloudNet, Pisa, Italy, October 2016

Traffic Projection

- In a cloud with 100 servers, 15 VMs each
  - Periodical traffic from servers
    - 100 servers: 100 x 15kb/s = **1,5 Mb/s**
  - Traffic with increasing number of VMs
    - 1500 VMs: 1500 x 0,77 kb/s = **1,2 Mb/s**
  - Total traffic
    - **2,7 Mb/s**

Sciammarella, T., Couto, R. S., Rubinstein, M. G., Campista, M. E. M., Costa, L. H. M. K., "*Analysis of Control Traffic in a Geo-distributed Collaborative Cloud*", IEEE CloudNet, Pisa, Italy, October 2016

Traffic Projection

- In a cloud with 100 servers, 15 VMs each
  - Periodical traffic from s
    - 100 servers: 100 x 15kb
  - Traffic with increasing r
    - 1500 VMs: 1500 x 0,77
  - Total traffic
    - **2,7 Mb/s**

Where should one locate the controller?

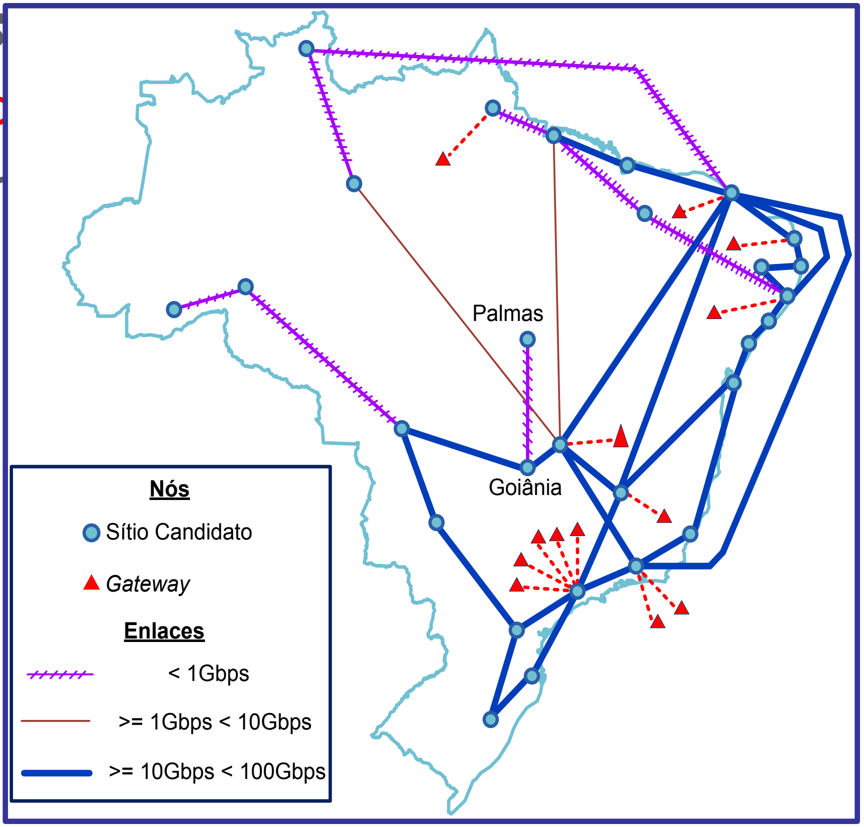

Sciammarella, T., Couto, R. S., Rubinstein, M. G., Campista, M. E. M., Costa, L. H. M. K., "*Analysis of Control Traffic in a Geo-distributed Collaborative Cloud*", IEEE CloudNet, Pisa, Italy, October 2016

Addicional Hot Topics?

- Information-centric Networks
  - Silva, V. B. C., Campista, M. E. M., Costa, L. H. M. K., "*TraC: A Trajectory-aware Content distribution strategy for Vehicular Networks*", Elsevier Vehicular Communications, vol. 5, pp. 18-34, July 2016
- Software-Defined networking (SDN)
- Network Function Virtualization (NFV)
- ...