Fabio Palumbo

Tutor: Antonio Pescapé

XXXIII Cycle - III year presentation

Monitoring cloud and mobile infrastructures for latency and bandwidth sensitive applications

Background

- **Graduation:** M. Sc. degree cum laude in Computer Engineering in July 2017
- Currently working in the **COMICS** research group under the supervision of Prof. Antonio Pescapè
- **Fellowship type**: “PON Dottorati Innovativi”
- **Collaborations**:
  Saint Louis University (USA)

  System Management S.p.A. (Italy)

  Ningbo University (China)

Credit summary

	Credits year 1								Credits year 2								Credits year 3												
		1	2	3	4	5	6			1	2	3	4	5	6			1	2	3	4	5	6	7	8	9			
	Estimated	bimonth	bimonth	bimonth	bimonth	bimonth	bimonth	Summary	Estimated	bimonth	bimonth	bimonth	bimonth	bimonth	bimonth	Summary	Estimated	bimonth	bimonth	bimonth	bimonth	bimonth	bimonth	bimonth	bimonth	bimonth	Summary	Total	Check
Modules	20	0	3	12	0	0	7,2	**22**	10	1,2	9	2	0	0	0	**12**	0	0	0	0	0	0	0	0	0	0	**0**	**34**	30-70
Seminars	5	0,3	1,1	0,4	3,2	0	0,4	**5,4**	5	0,8	0,9	2,8	0	0	0	**4,5**	0,5	0,2	0,4	0	0	0	0,2	0	0	0	**0,8**	**11**	10-30
Research	35	9,7	5,9	0	6,8	10	2,4	**35**	45	8	0,1	5,2	10	10	10	**43**	60	9,8	9,6	10	10	10	9,8	10	10	5	**84**	**162**	80-140
	60	10	10	12	10	10	10	**62**	60	10	10	10	10	10	10	**60**	60	10	10	10	10	10	10	10	10	5	**85**	**207**	180

Experience abroad

- 6-months as visiting scholar in collaboration with the **Computer Science Department** and **Pathology Department** (School of Medicine) of Saint Louis University (USA)

- Working on **telepathology** applications

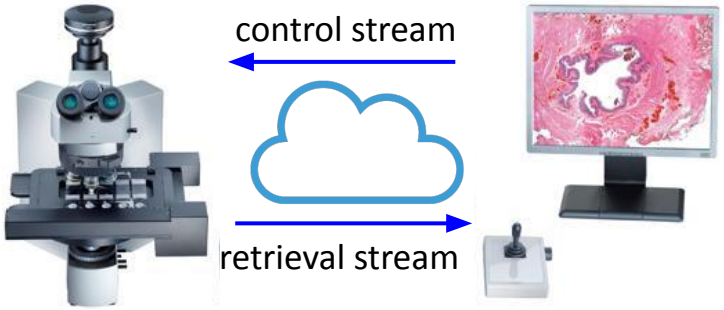

SAINT LOUIS UNIVERSITY
EST. 1818

Multimedia and telemedicine applications

- **Requirements**
  - High **bandwidth**
    - Measured via *Active* and *Passive* approaches
  - Low **latency**
- **Novel scenarios**
  - **Mobile** and **cloud** infrastructures

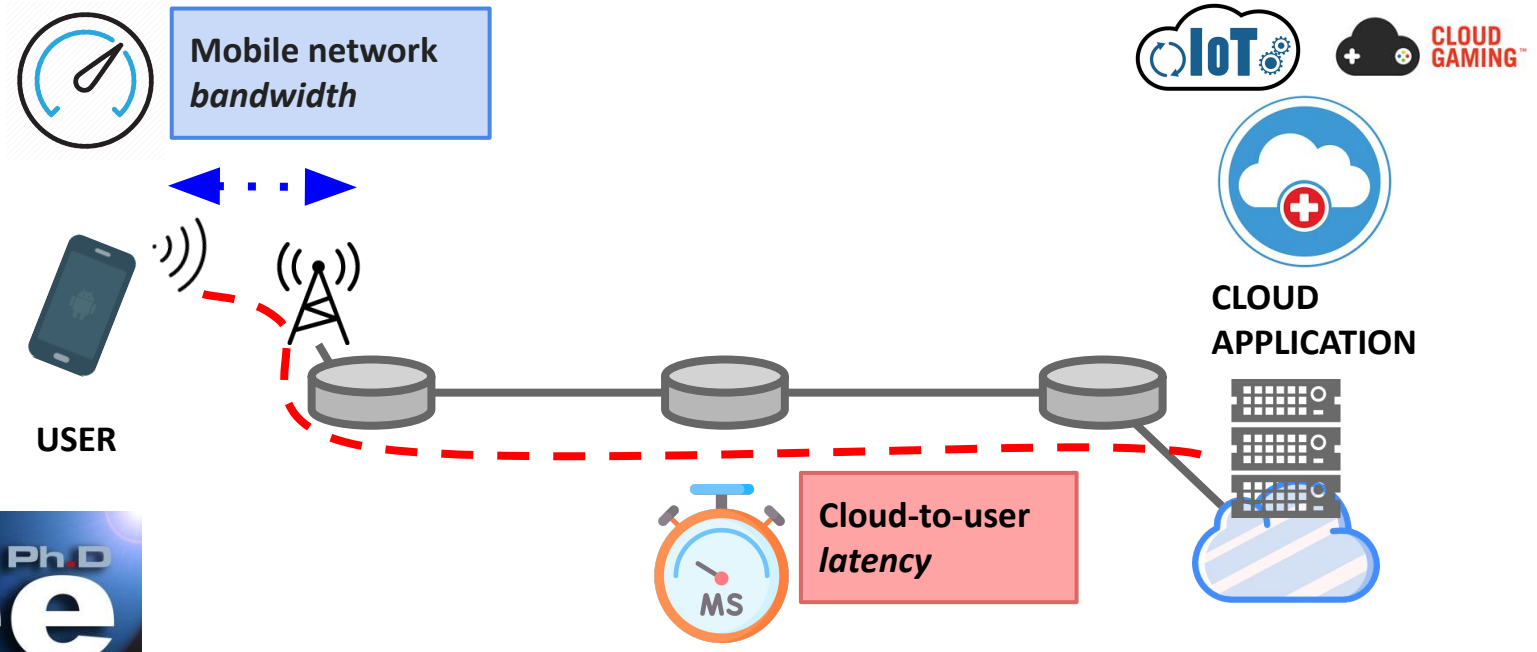

Cloud computing for latency and bandwidth sensitive applications

- **Cloud computing adoption allows to**
  - + reduce costs
  - + dynamically adapt to user demands
  - + provide innovative services
- **but implies limited visibility into network performance**
  - unknown *bandwidth* and packet loss
  - increased *latency*

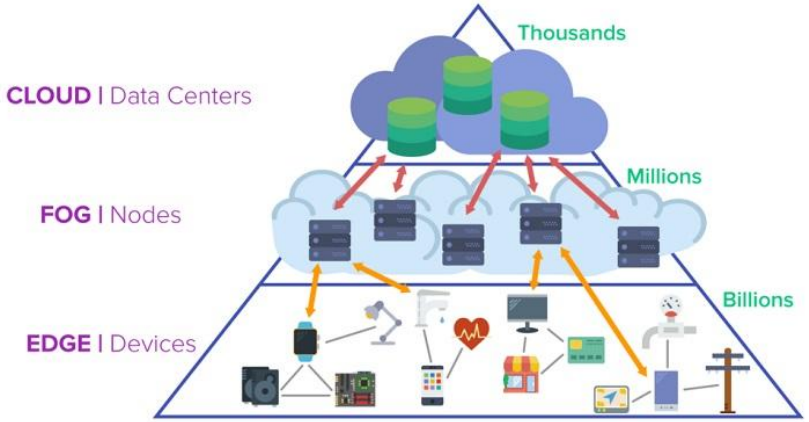

Mobile network infrastructure for latency and bandwidth sensitive applications

- **Mobile cloud computing** can
  - \+ tackle **storage**, **computational** and **battery** constraints of mobile devices and support **innovative applications**
- **but mobile networks are characterized by**
  - constrained **bandwidth**
  - high *variability*, impacted by mobility
  - Only **54%** of mobile traffic expected to be from higher speed 5G infrastructures by 2026

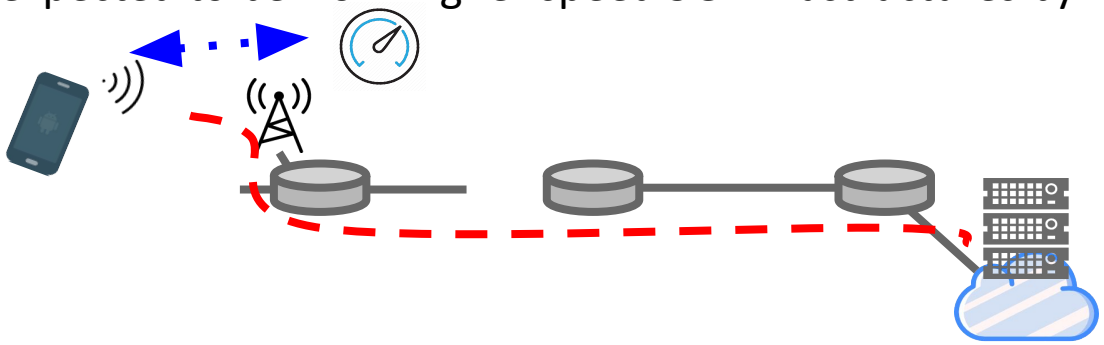

→ ***Latency and bandwidth require proper monitoring methodologies in this scenario***
→ ***Need for open research testbeds, testing real network environments***

Research activities

- ***Design* and *evaluation* of network-level monitoring methodologies for *latency* and *bandwidth* sensitive applications in mobile scenarios**
  - Available *bandwidth* estimation in *mobile broadband networks (MBB)* via *active* [C1] and *passive* (*SDN*-based) [C2] techniques in *real* scenarios
  - Design of an *SDN* controller *migration protocol* [C4]
- **Analysis of network performance of public cloud providers**
  - User-perceived *latency* towards *commercial cloud providers* from distributed nodes [C3, J2]

Active bandwidth estimation in MBB (1/3) [C1]

- MONROE EU testbed for *MBB* (3G/4G) *network*
- ***Active*** bandwidth evaluation: D-ITG and Yaz
- **4 countries** and **4 Mobile Network Operators** tested
- **30+ days** experimental campaign

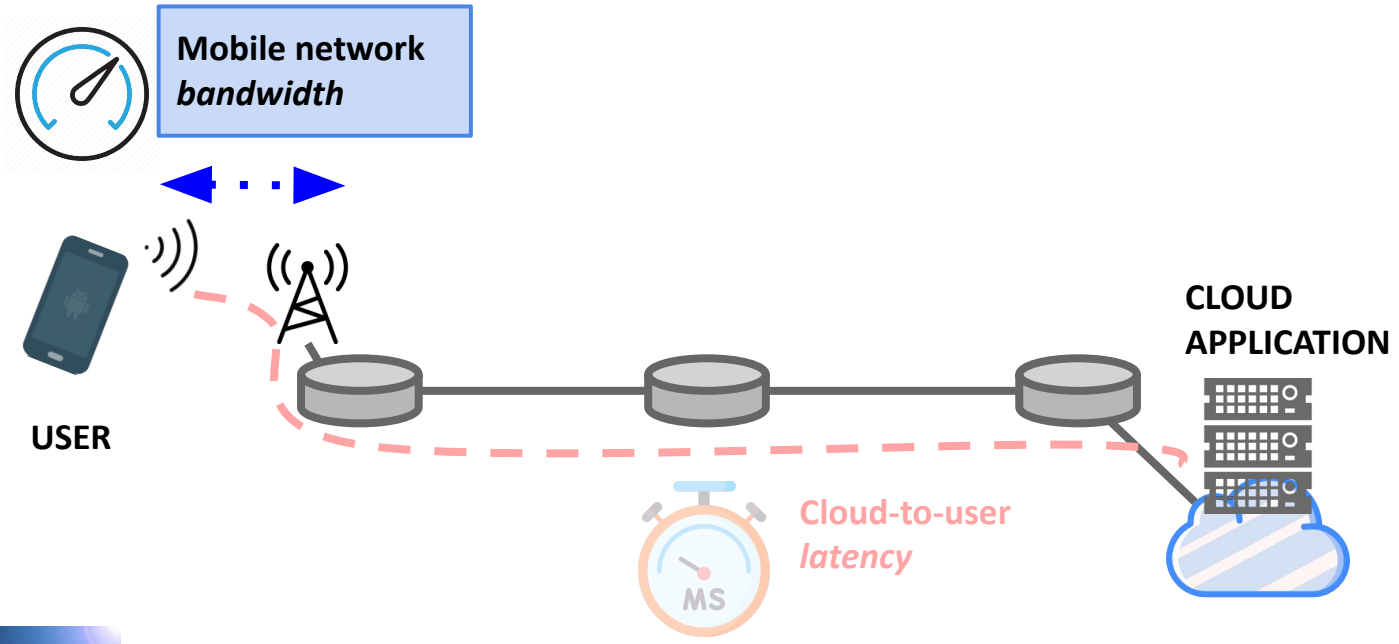

Active bandwidth estimation in MBB (2/3) [C1]

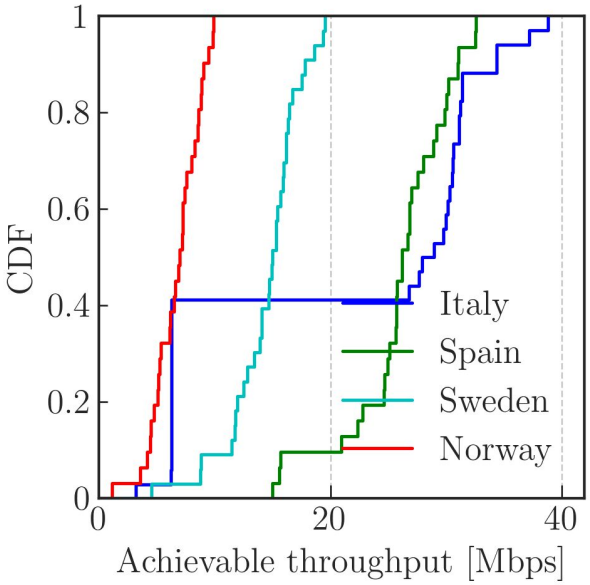

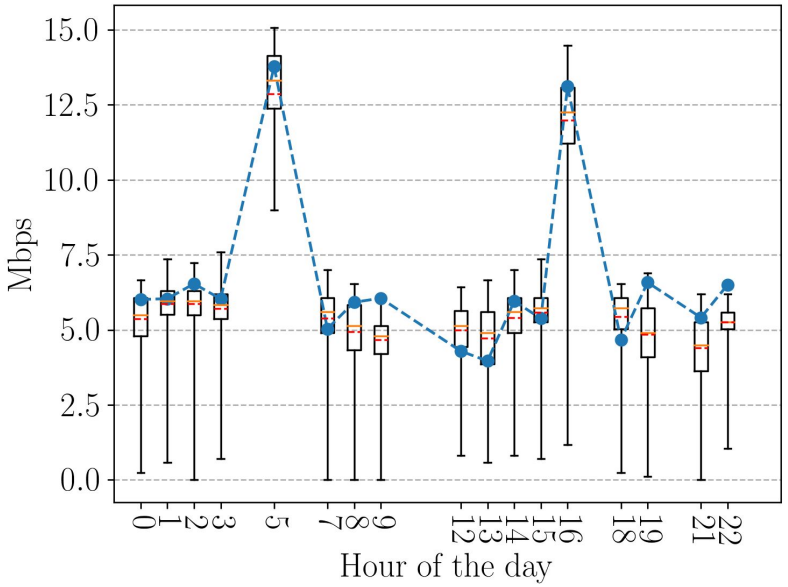

- Assessing **throughput** variability over **space** and **time**

→ *Bandwidth depends on country/operator, time of the day*

Active bandwidth estimation in MBB (3/3) [C1]

- Assessing **available bandwidth** as a proxy for **achievable throughput**
- Evaluating **traffic volume** required for estimates

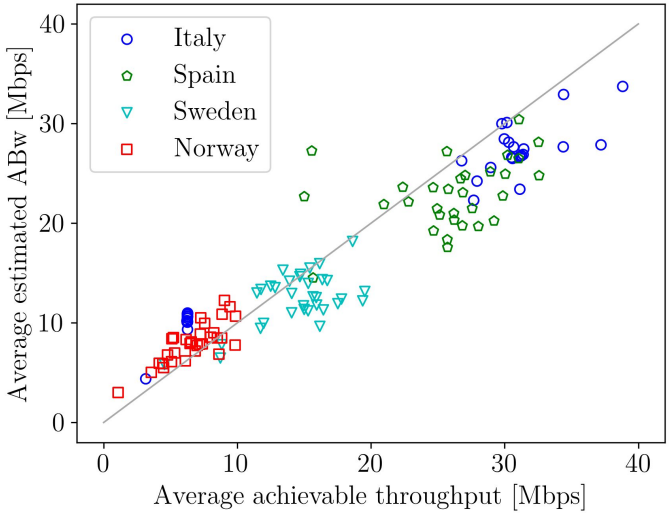

→ *Depending on the country/operator, available bandwidth achieves good accuracy at a fraction of the traffic volume required by TCP throughput*

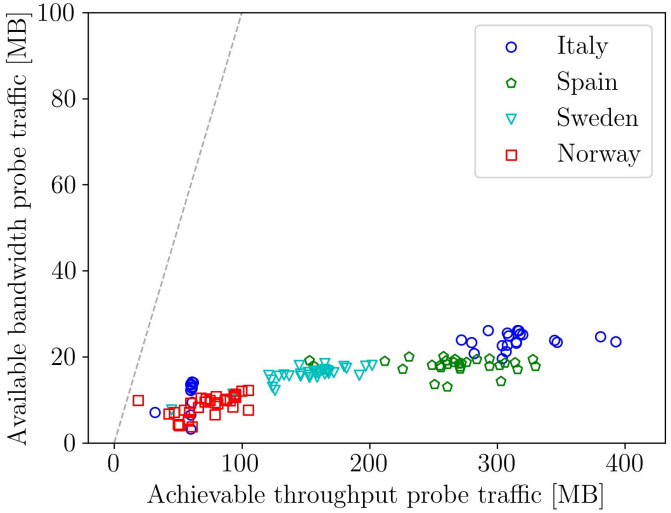

SDN-based bandwidth estimation in MBB (1/2) [C2]

- Evaluation of SDN-based ***passive*** **bandwidth estimation** technique in MBB
- Collecting **traffic volume counters** (*V*) from **switches** over time and estimating bandwidth (*B*)
- **Timestamping** (*T*) impacts estimation accuracy
- **Local** and **remote** controller deployment

$$B_i = \frac{V_i - V_{i-1}}{T_i - T_{i-1}} = \frac{\Delta V_i}{\Delta T_i}$$

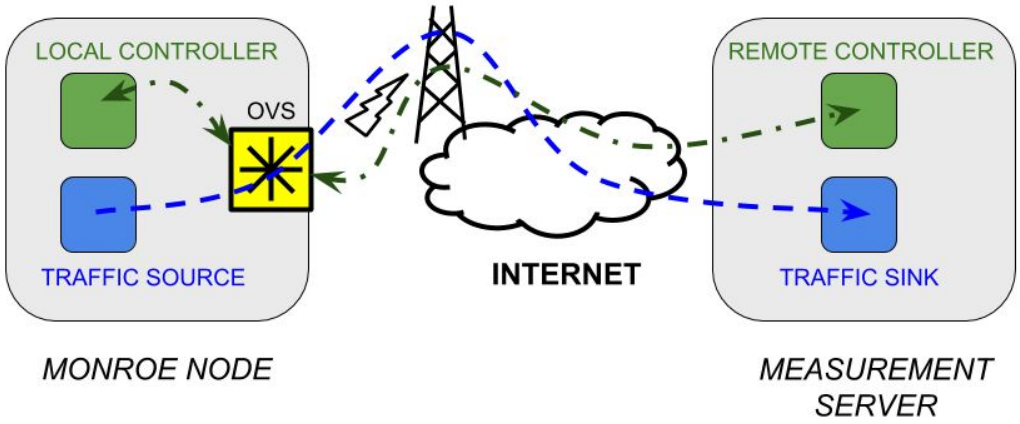

SDN-based bandwidth estimation in MBB (2/2) [C2]

- **Accuracy** evaluation of bandwidth estimates (*relative* error)
- Different polling periods and cross-traffic rates

→ *Mean error is low even in the most demanding cases*

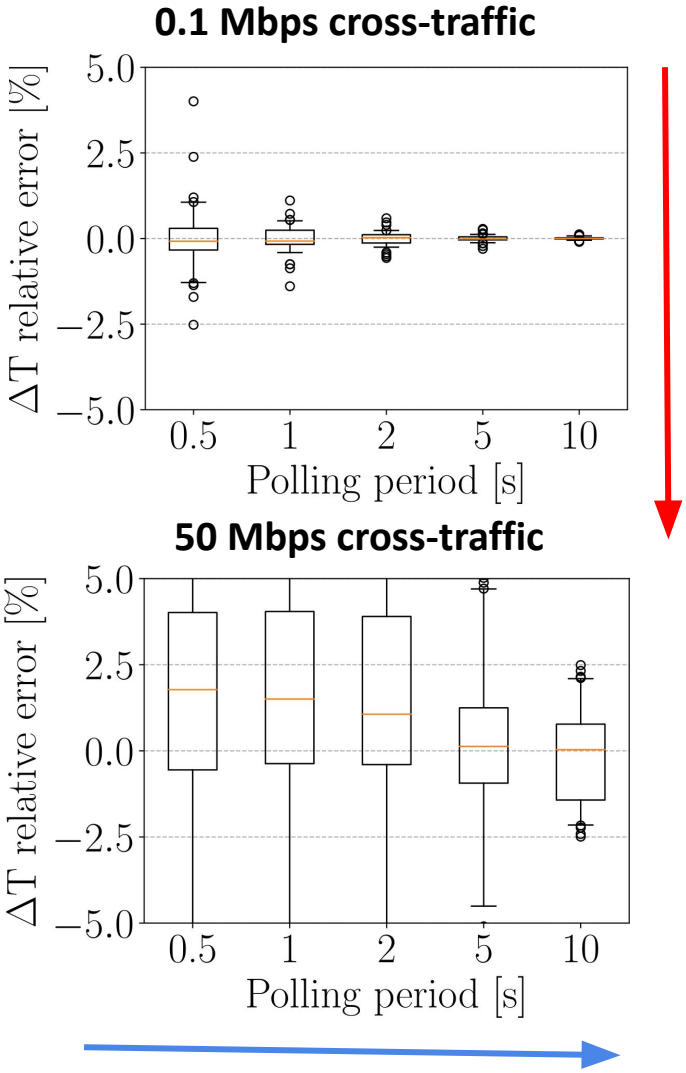

SDN controller migration (1/2) [C4]

- **Increased complexity** of network scenarios, specially with 5G and edge cloud deployments
- **SDN** paradigm eases **manageability**
- **Multiple controllers** can increase **reliability**
- **Design** and **evaluation** of a **migration protocol** to handle multiple SDN controllers

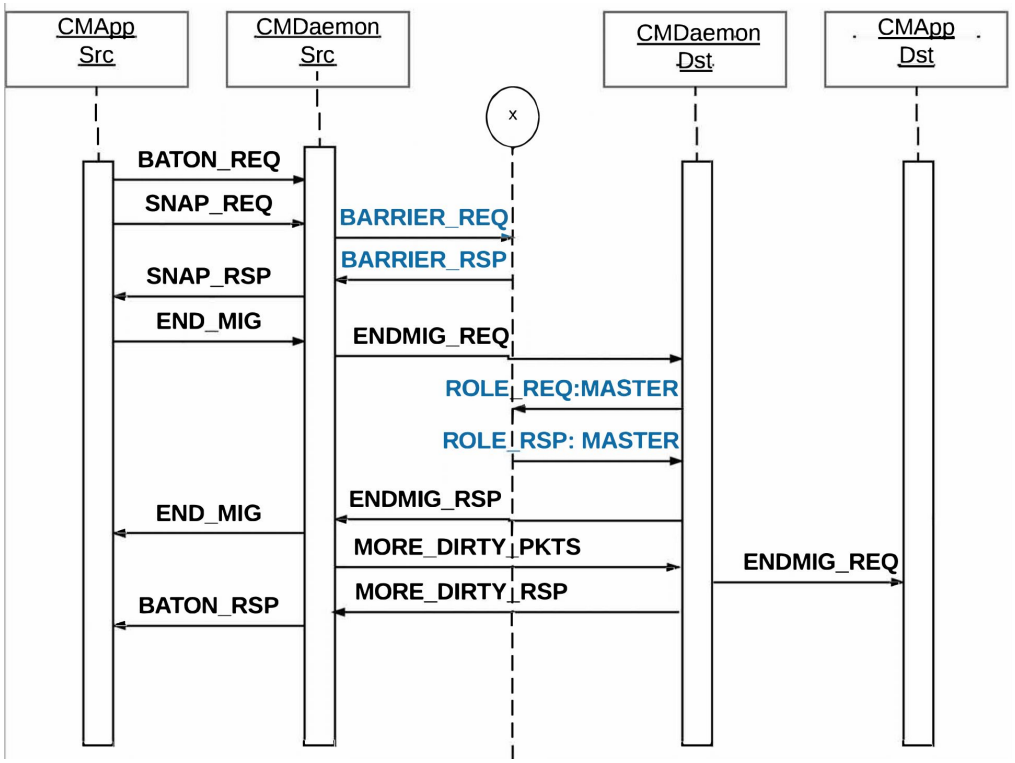

SDN controller migration (2/2) [C4]

- Evaluating **migration time** according to background traffic rate (affecting number of requests)

→ *Migration time is limited, affected by background traffic rate*

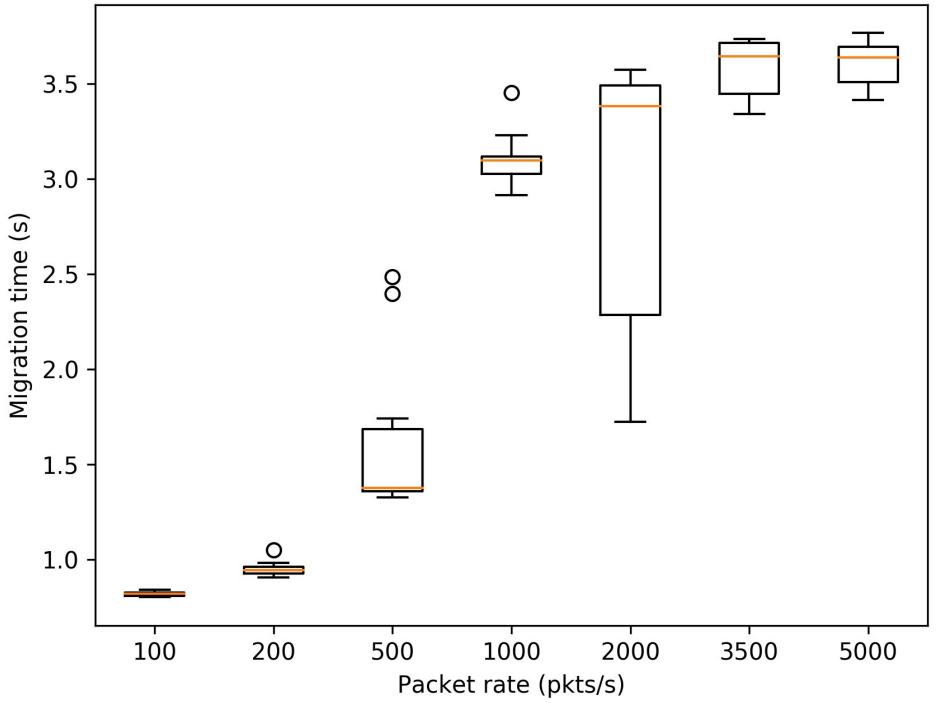

Cloud-to-user latency monitoring (1/3) [C3, J2]

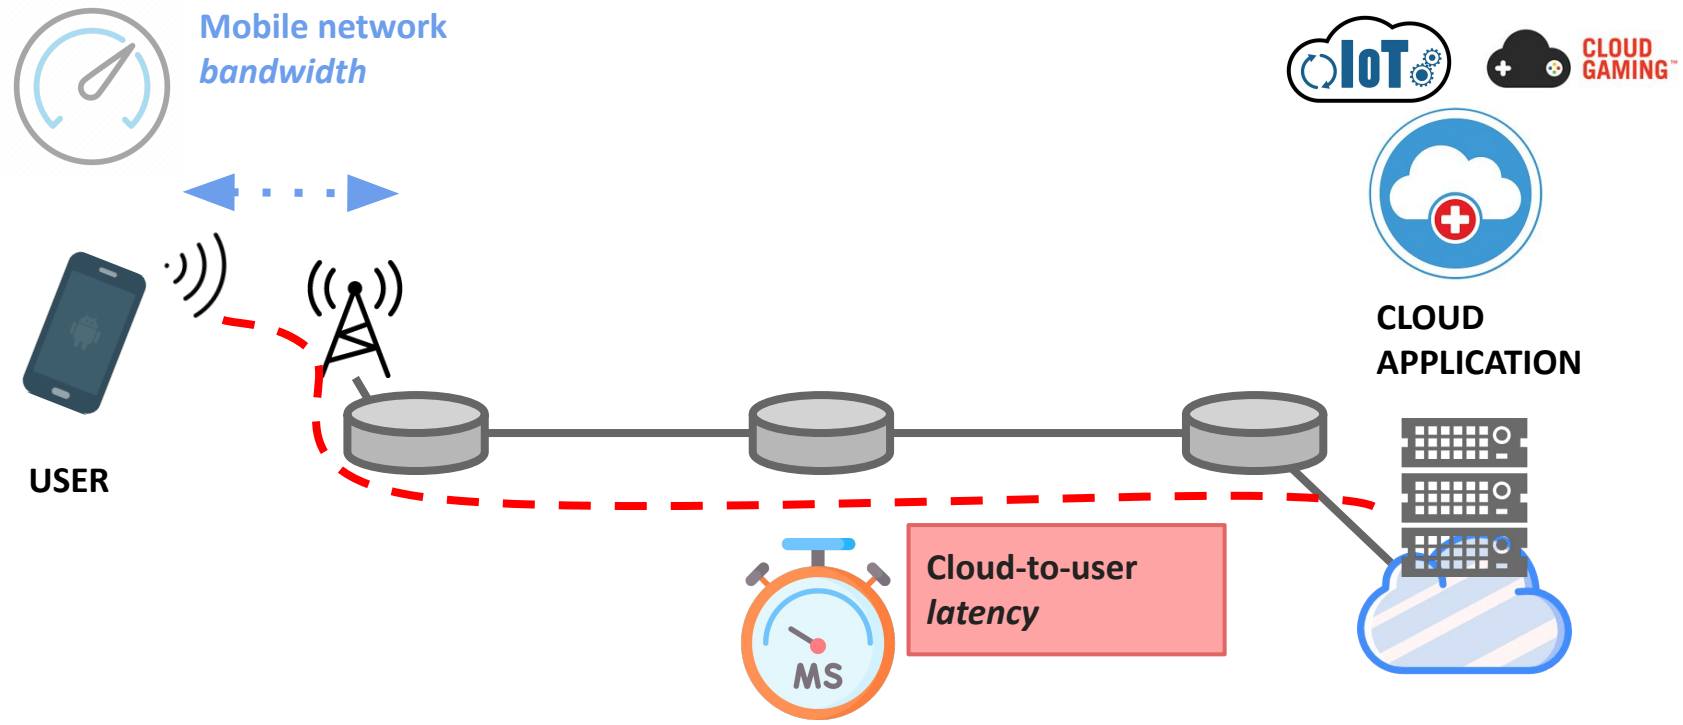

Cloud-to-user latency monitoring (2/3) [C3, J2]

PlanetLab

- Worldwide testbed for network measurements
- **Cloud-to-user latency** (C2U) evaluation
- 2 **providers** in 4 **regions**

amazon web services™ Microsoft Azure

- 25 **source nodes** across the world
- Different **protocols** at different **layers** of the stack tested

= Cloud Region = Source Node

Cloud-to-user latency monitoring (3/3) [C3, J2]

- Assessing latency **variability** over *space* and *time*
- Comparing results from the *two providers*

→ *Cloud-to-user latency depends on the destination region tested, may vary over time*

→ *Best-performing provider varies according to the destination region*

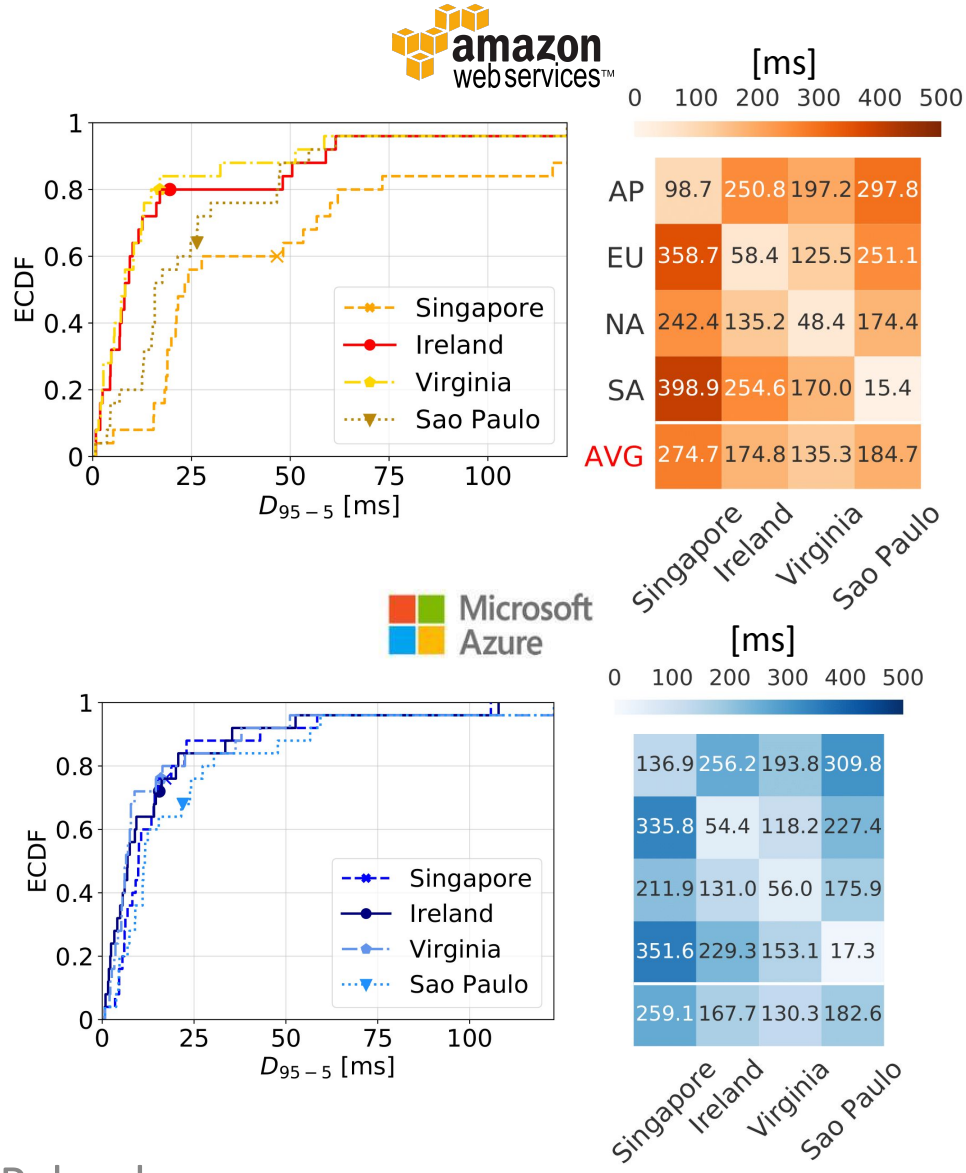

C2U latency application scenarios [C3, J2]

- Leveraging C2U latency measurements to detect **anomalous** events
- Using previously collected samples as **baseline** to detect anomalies

→ *Detecting anomalies and helping troubleshooting*

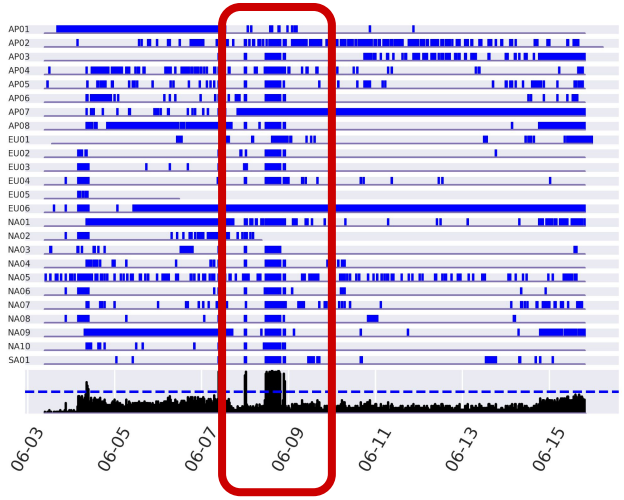

- Assessing the **advantages** of **multi-cloud deployments**
- Evaluating the **improvement** in the *ideal* case, considering the best provider at each instant

→ *Non-negligible improvement if switching provider over-time*

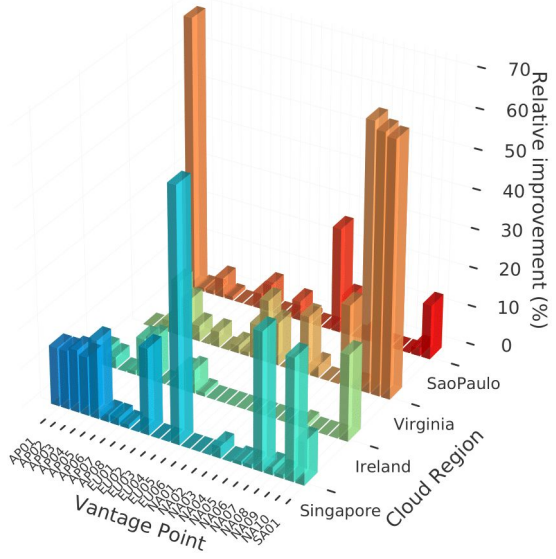

Final remarks

- **Latency** and **bandwidth** are crucial for modern applications, including **multimedia** and **healthcare** ones
- Focus on **Cloud** and **Mobile** scenarios
- Proper **monitoring methodologies** are needed to assess these metrics
- Current infrastructures *can support* demanding applications with careful **monitoring** and **management**

Products

Publications:

C1. G. Aceto, **F. Palumbo**, V. Persico and A. Pescapé, "*Available Bandwidth vs. Achievable Throughput Measurements in 4G Mobile Networks*", 2018 14th International Conference on Network and Service Management (CNSM), Rome, 2018.

C2. G. Aceto, **F. Palumbo**, V. Persico, A. Pescapé and H. Chen, "*Evaluation of SDN-based bandwidth estimation in Mobile Broad Band networks*", 2018 24th Asia-Pacific Conference on Communications (APCC), Ningbo, 2018.

C3. **F. Palumbo**, G. Aceto, A. Botta, D. Ciuonzo, V. Persico and A. Pescapé, "*Characterizing Cloud-to-user Latency as perceived by AWS and Azure Users spread over the Globe*", 2019 IEEE Global Communications Conference (GLOBECOM), Waikoloa, 2019.

C4. C. Contoli, **F. Palumbo**, F. Esposito, F. Callegati and A. Pescapé, "*Flock: A Live Migration Protocol For SDN Controllers*", 2019 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN), Dallas, 2019.

J1. S. Shamshirband, M. Fathi, A. T Chronopoulos, A. Montieri, **F. Palumbo** and A. Pescapè , "*Computational intelligence intrusion detection techniques in mobile cloud computing environments: Review, taxonomy, and open research issues*", Journal of Information Security and Applications (JISA), 2020

J2. **F. Palumbo**, G. Aceto, A. Botta, D. Ciuonzo, V. Persico and A. Pescapé, "*Characterization and Analysis of Cloud-to-User Latency: the case of Azure and AWS*", Computer Networks (ComNet), 2021.

Thank you for your attention