Monitoring the Cloud-IoT continuum for latency-Aware applications placement

Marco Gaglianese

Service-oriented, Cloud and Fog Computing Research Group

Department of Computer Science

University of Pisa, Italy

Fog computing (Cloud-IoT)

- Everything is a Fog node
  - Routers, Switches
  - Servers
  - IoT devices
- Deploy applications meeting the requirements
  - Computation
  - Storage
  - Quality of Service
    - Latency
    - Bandwidth

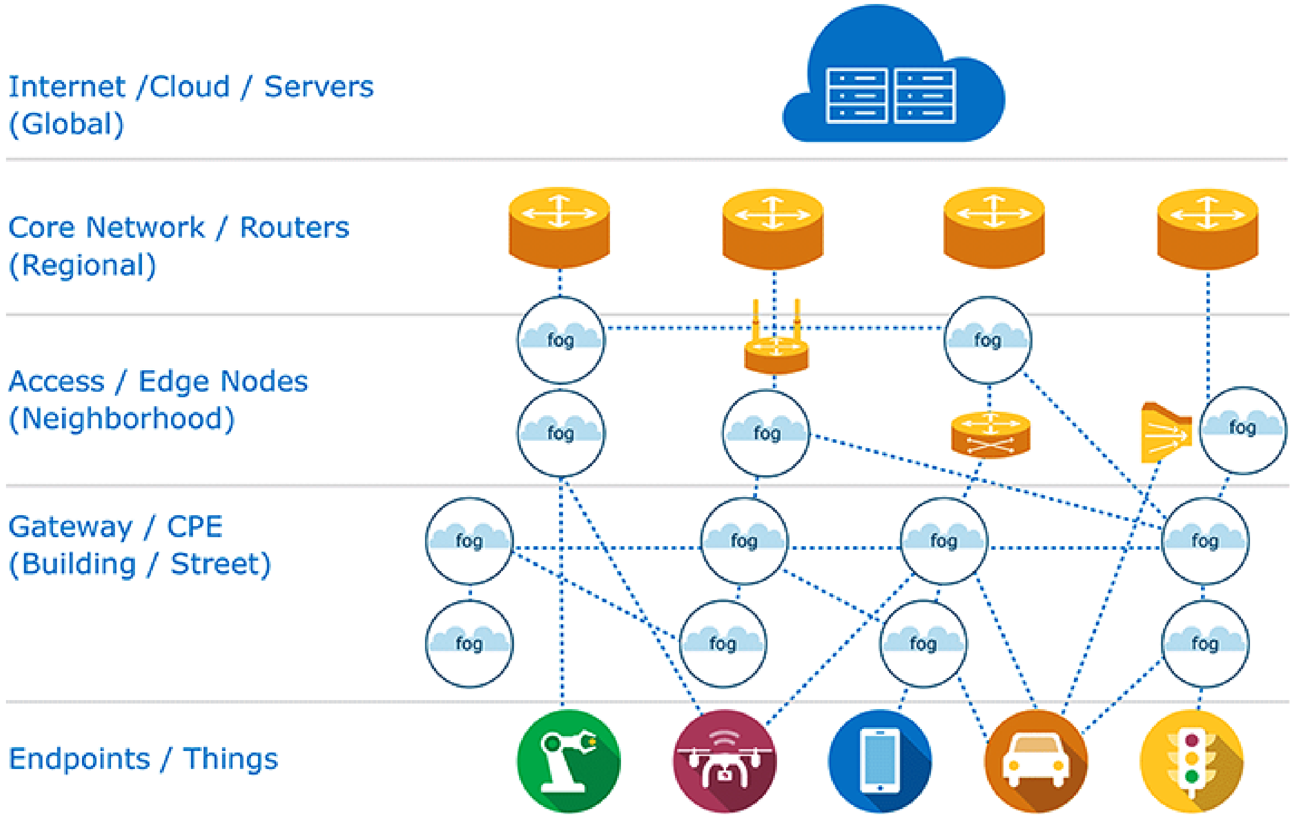

Fog Infrastructure Monitoring

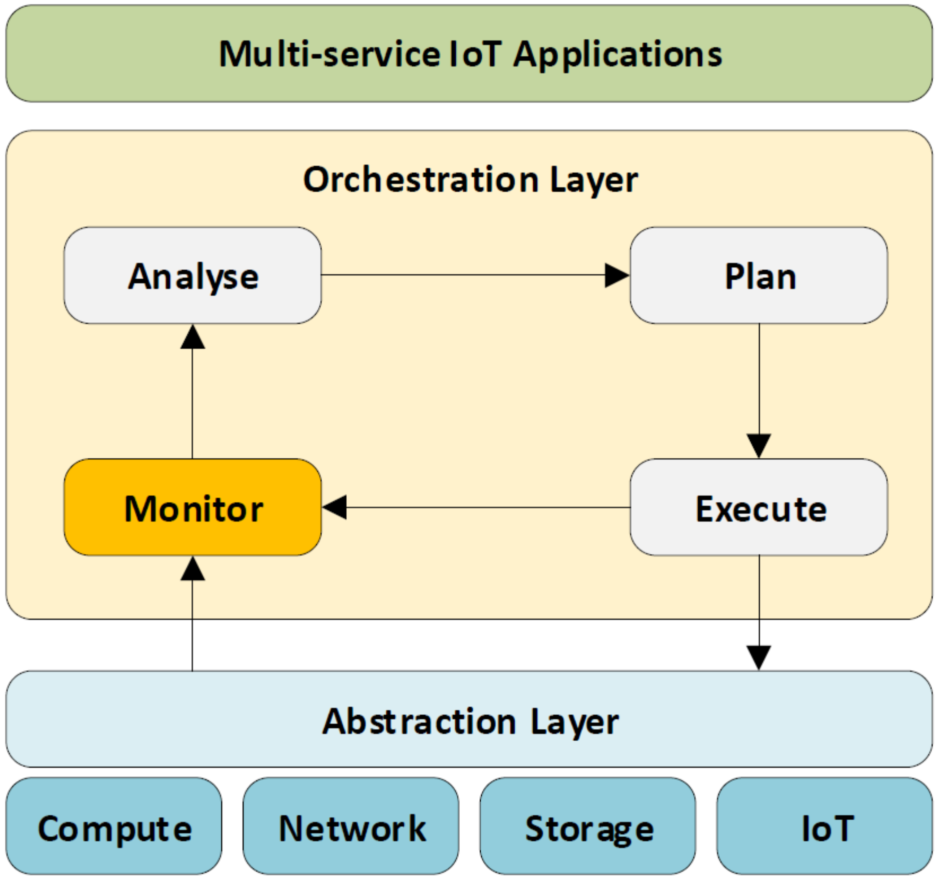

- **Fog orchestration**
  - much work on **Analyse**
  - some work on **Plan** & **Execute**
  - less work on **Monitor**

- **Monitor** pivotal to decide
  1. where to deploy app services at first
  2. when/where to migrate app services

FogMon 2

https://github.com/di-unipi-socc/FogMon/tree/liscio-2.0

An Open source lightweight fault-resilient monitoring service for Cloud-IoT infrastructures

- The service monitors:
  - hardware resources availability
    - CPU, RAM, STORAGE
  - end-to-end network QoS
- Two types of distributed P2P agents:
  - Followers measure the monitored metrics, and
  - Leaders aggregate the metrics from a **group** of Followers, and gossip them to other Leaders

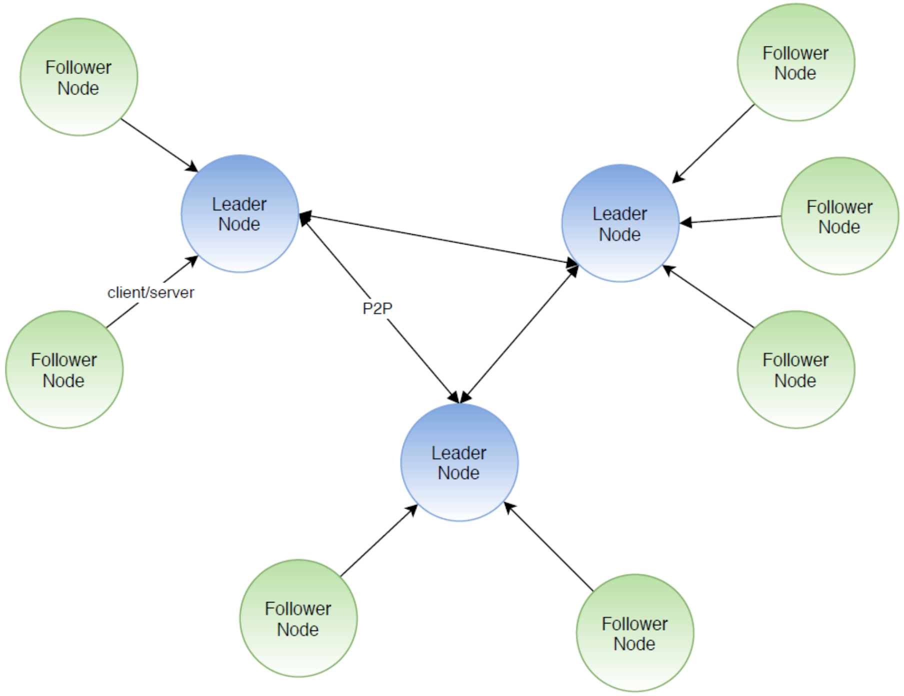

A. Brogi, S. Forti, M. Gaglianese. *Measuring the Fog, Gently*. ICSOC 2019.
S. Forti, M. Gaglianese, A. Brogi. *Lightweight self-organising distributed monitoring of Fog infrastructures*. Future Generation Computer Systems. 2021.

Measuring Latency

Intra-groups measurements

ICMP via *PING*

Inter-groups measurements

$$\ell_{A,B} \simeq \ell_{A,L1} + \ell_{L1,L2} + \ell_{L2,B}$$

assuming Leader-Leader latency is higher than Leader-Follower latency

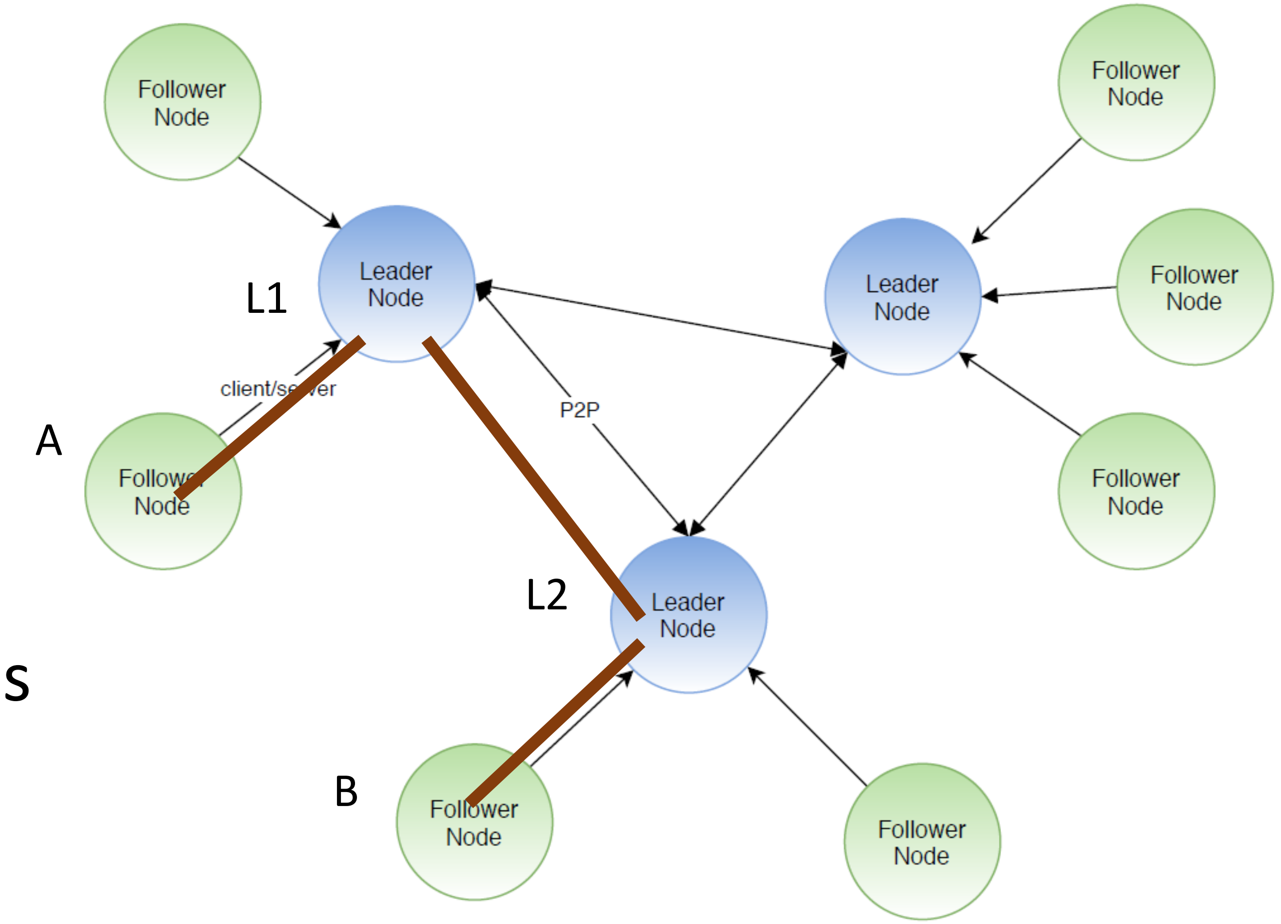

Measuring Bandwidth

Intra-groups:

- Intrusive measurements
  - *Iperf3*

- Passive techniques
  - *Assolo*

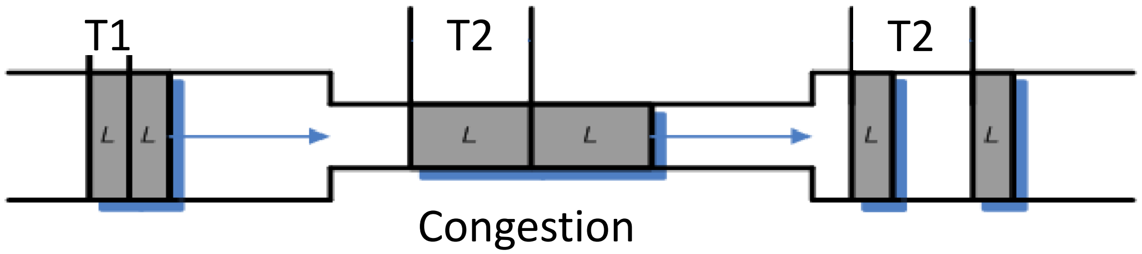

Inter-groups:

$$\beta_{A,B} \simeq \min_{k,h}\{\max\{\beta_{A,k}\}, \max\{\beta_{h,B}\}, \beta_{L1,L2}\}$$

FogMon restructuring mechanism

- Hybrid overlay network, with $\sqrt{N}$ leaders
  - K-medoids with latency.

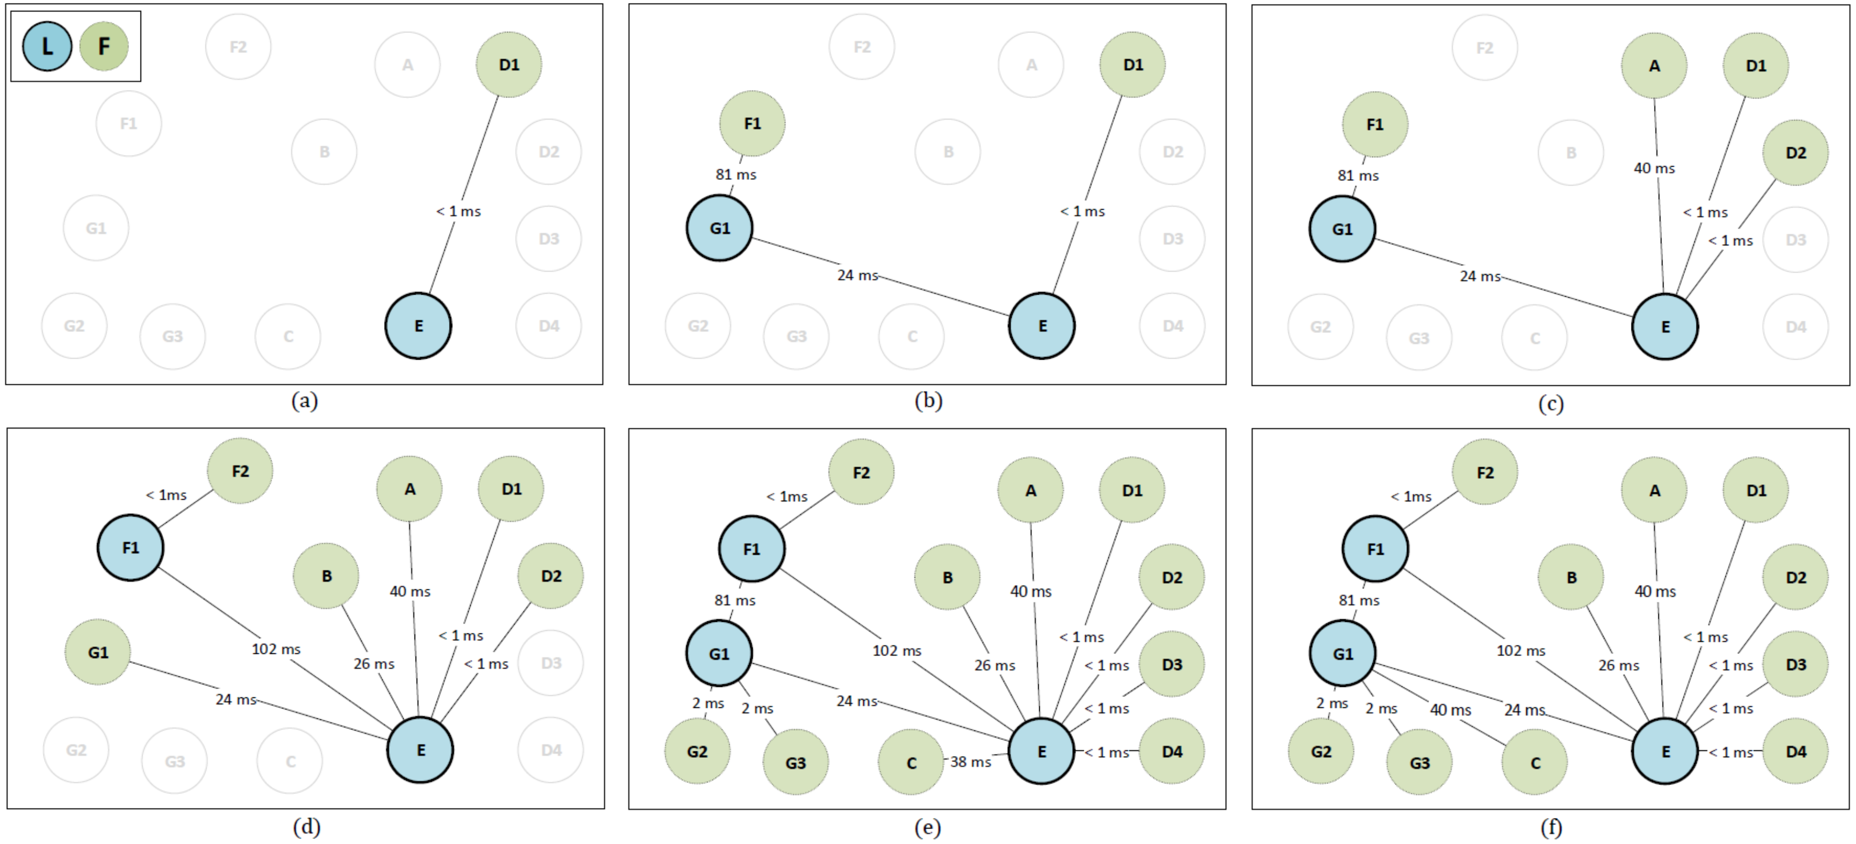

FogMon communications costs

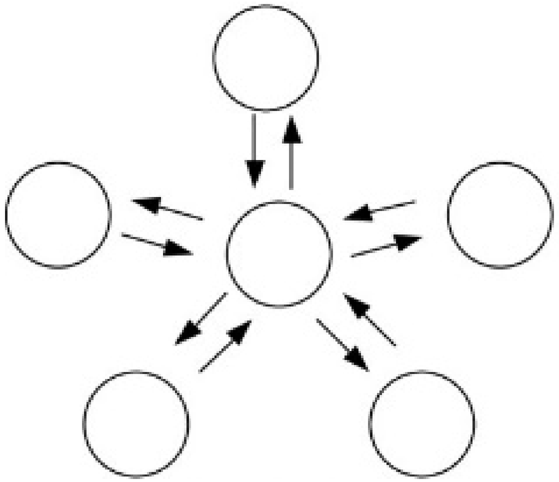

- Leaders collect measurements from Followers in their groups.

- Leaders spread data to other Leaders via Gossiping
  - $O(\log L)$ rounds to spread information on avg
  - $O(L \log L)$ messages exchanged overall

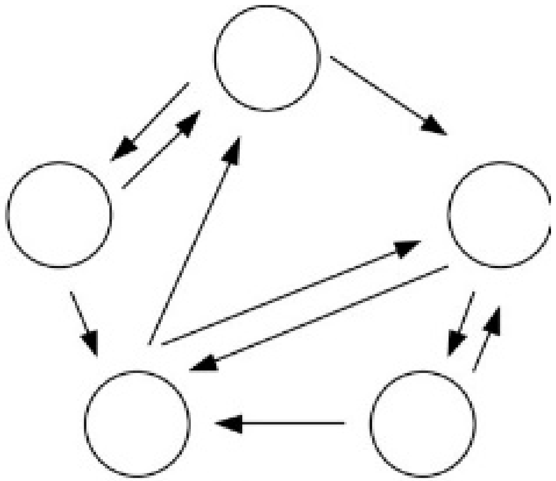

FogMon Fault-tolerance and Scalability

Fault-tolerance

- Data replication at Leaders guarantees tolerance wrt some Leader failures.
- Followers rearrange into other groups when their Leader fails.
- Groups keep working in case of network interruption between Leaders.

Scalability

- N nodes, L leaders
- $\frac{N}{L}$ nodes per group (per Leader)
- $\frac{N^2}{L^2}$ e2e measurements for bw and latency
- If $L \simeq \sqrt{N}$ then $O\left(\frac{N^2}{\sqrt{N}^2}\right) = O(N)$ e2e measurements.

Recent experiments of FogMon

LiSClo experiment - "Lightweight Self-adaptive Cloud-IoT Monitoring across Fed4FIRE+ Testbeds"

FED4FIRE
Federation for FIRE Plus

- Fed4FIRE+ federation provided the testbed infrastructure

M. Gaglianese, S. Forti, A. Brogi. Lightweight Self-adaptive Cloud-IoT Monitoring across Fed4FIRE+ Testbeds. At INFOCOM 2022. Workshop CNERT.

Experiment Setup and Plan

Measuring

- **footprint** on hardware and bandwidth
- **relative error** on measurements & estimates **against setup ground-truth** (configured via GRE-tunnels and tc)
- **time to reach stability**

on

- 3 types of **Leader & Follower Failures** (NF)
- 2 types of **Link Failures** (LF)

- **20 (S), 30 (M) and 40 (L) nodes across VIRTUALWALL and CityLab**
- **default** vs **reactive configurations**

	VIRTUALWALL (physical nodes)	CityLab (wireless nodes)
S	10	10
M	15	15
L	30	10

Total **30** experiments

Experiment tooling

FogMonEye

WebGUI for monitoring LiSCIo experiments

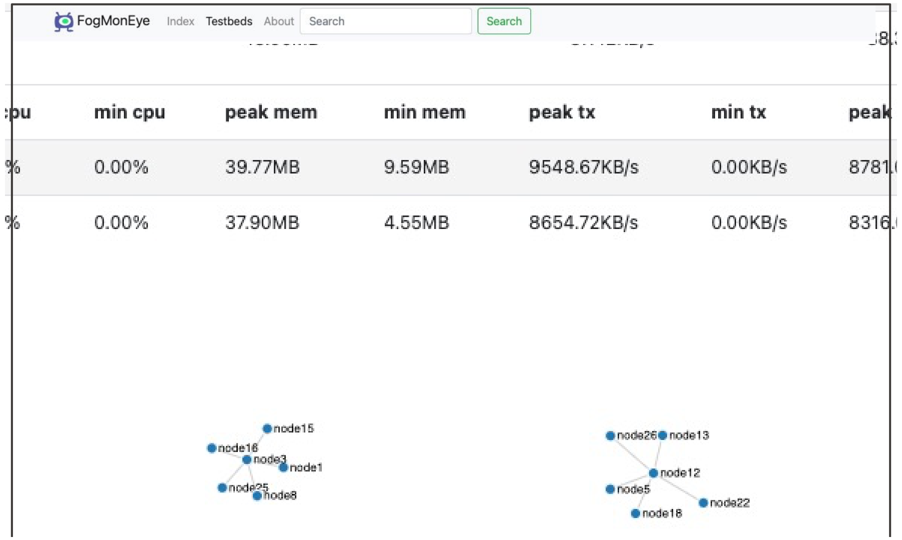

Topology Builder

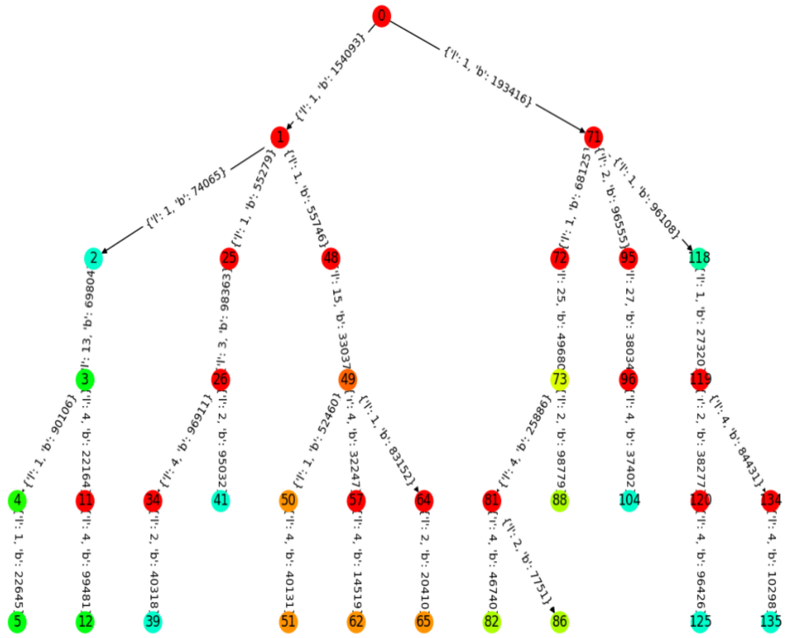

Experiment tooling (cont.)

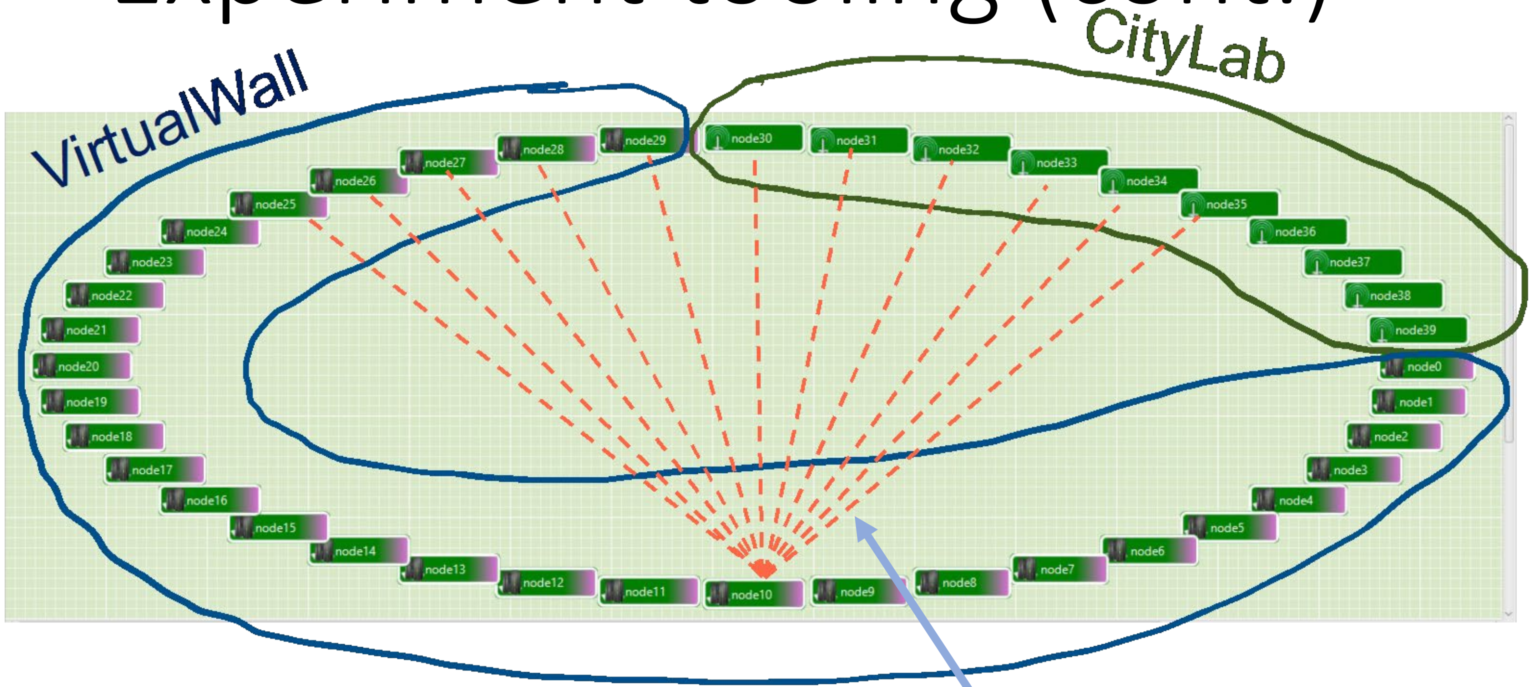

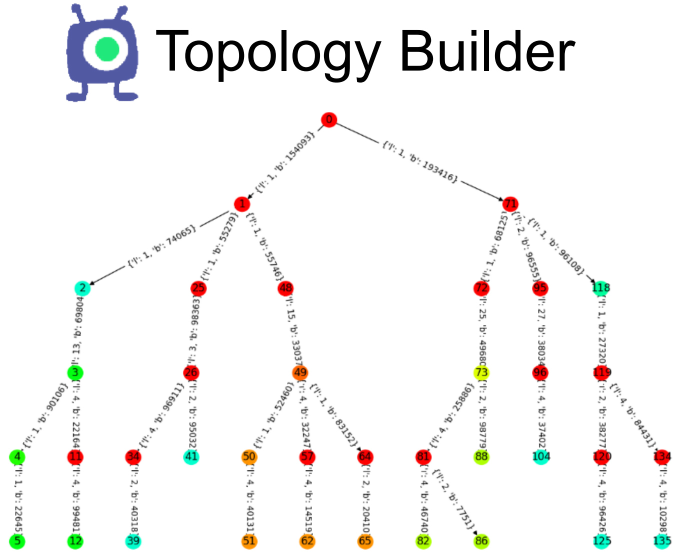

- Topology builder creates all N^2 **end-to-end connections** between all nodes across testbeds **via tc and GRE-tunnels**
- Mimicks a **hierarchical Cloud-IoT network** with Edge (at CityLab), transport (at both testbeds) and Cloud (at VIRTUALWALL) nodes, with lifelike latencies and bandwidths

Results

- Reactive vs Default:
  - Reactive faster in identifying changes but more resource intensive
- Link Failures (**LF**) vs Node Failures (**NF**):
  - Faster to react to Link Failures
  - **Relative errors** are comparable
- Size:
  - slight increase in **footprint** with the size
  - **Relative errors** are almost comparable

Time to stability

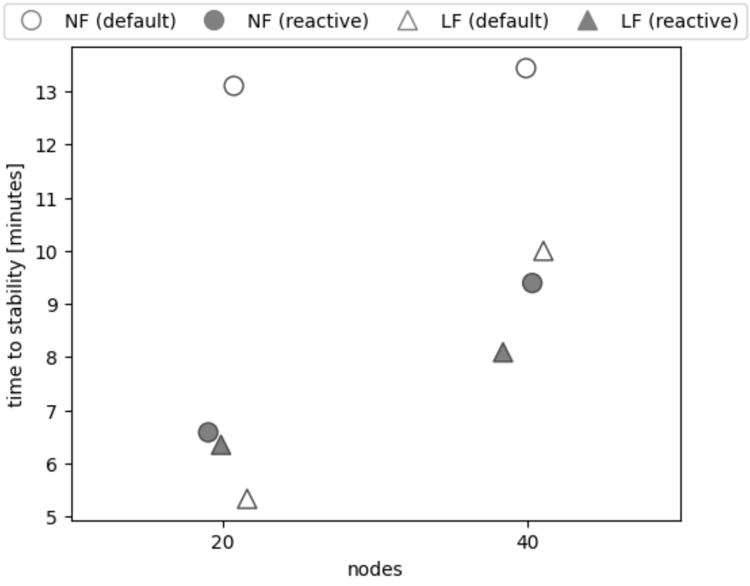

FogMon Footprint

CPU usage

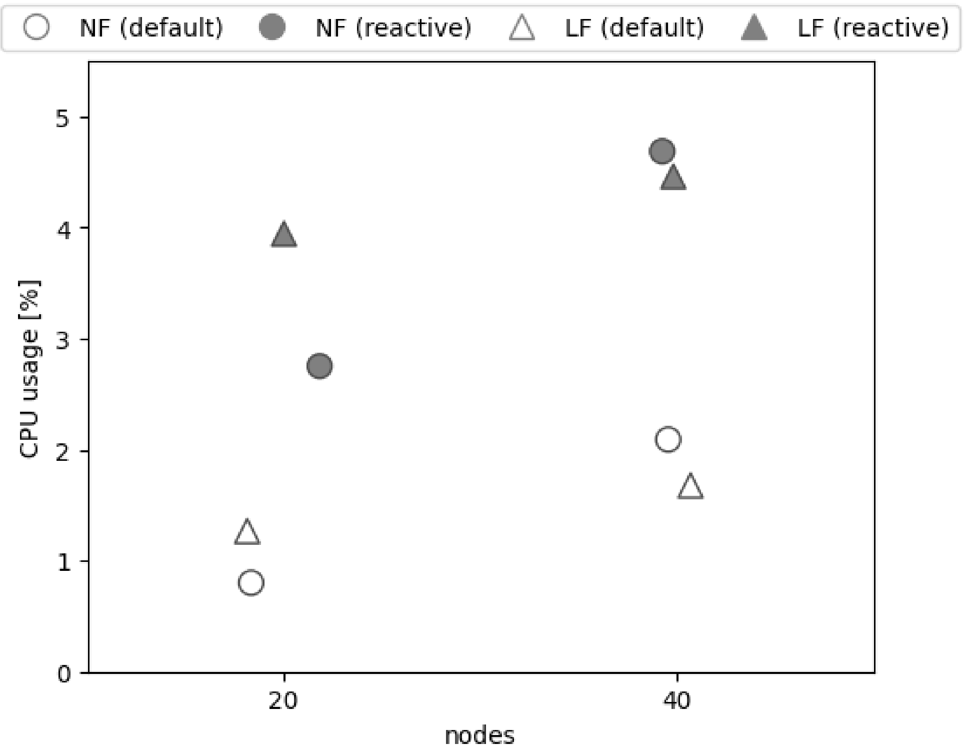

Bandwidth usage

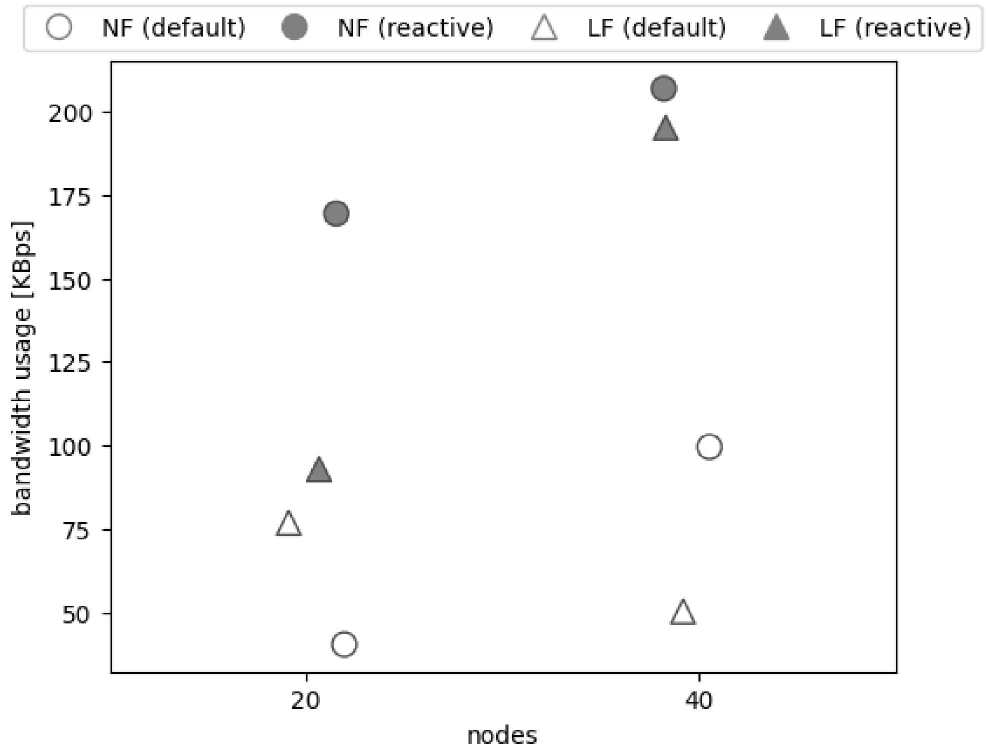

Relative errors

Latency

Bandwidth

Intra group
(Measured)

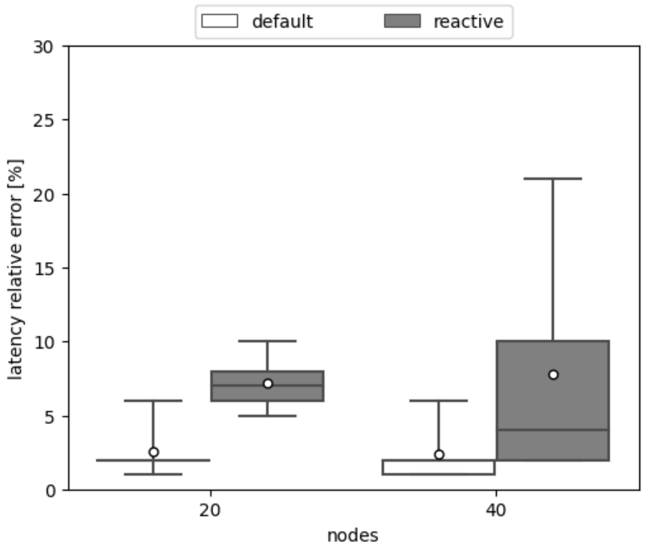

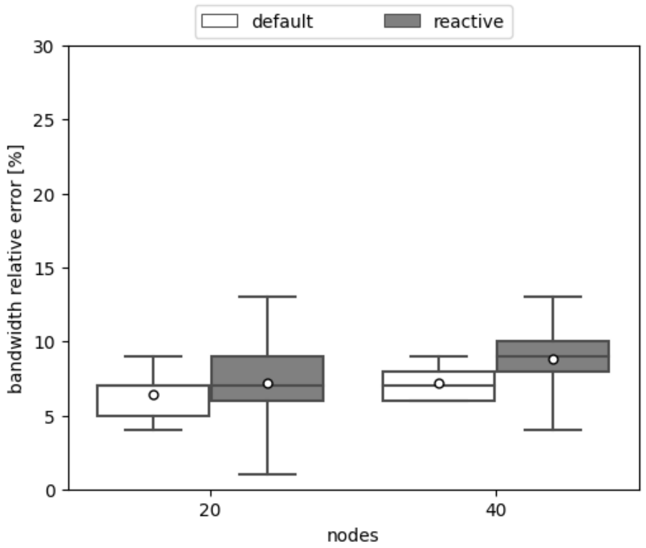

Inter group
(Approximated)

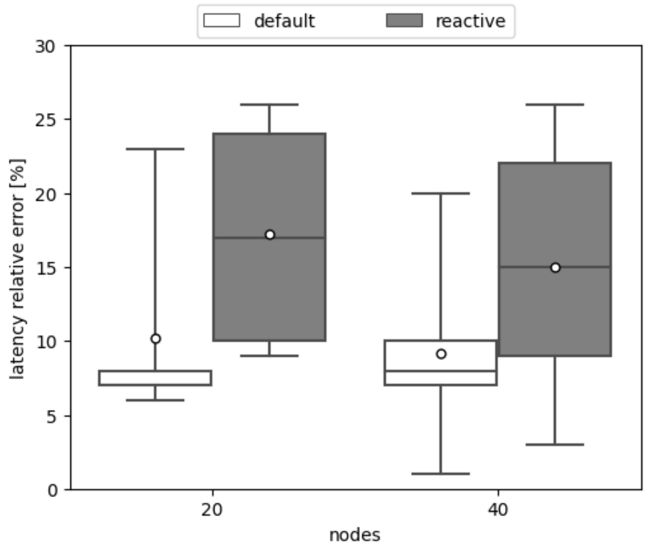

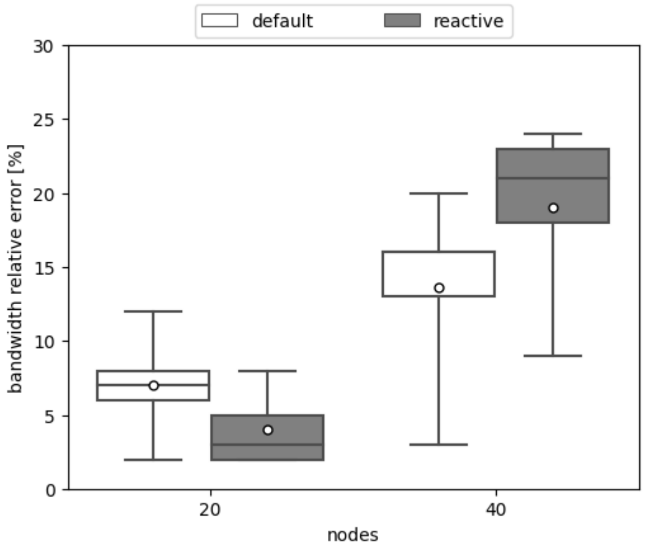

Concluding remarks

Leveraging Fed4Fire+ facilities, we have shown that FogMon:

- can be deployed across network boundaries and heterogenous computing capabilities
- detects and adapts to nodes and link failures
- exhibits low acceptable footprint on nodes and links at increasing infrastructure sizes (from 20 to 40 nodes)

Evolved FogMon from TRL 4 (lab) to TRL 5 (relevant environment, 40 nodes)

M. Gaglianese, S. Forti, F. Buti, F. Paganelli, A. Brogi. *Lightweight Self-adaptive Cloud-IoT Monitoring across Fed4FIRE+ Testbeds (LiSCIo) [Dataset]. Available on Zenodo:* [*http://doi.org/10.5281/zenodo.4682987*](http://doi.org/10.5281/zenodo.4682987) (2021)

Future work

- Hybrid Cloud-Edge architecture
- Improve network QoS estimates
- Improve Leader election mechanism
- further experimental assessment

Thank you
for your attention!

FogMon2

Marco Gaglianese
marco.gaglianese@phd.unipi.it

Service-oriented, Cloud and Fog Computing Research Group
Department of Computer Science
University of Pisa, Italy

FogMon2

https://github.com/di-unipi-socc/FogMon/tree/liscio-2.0