INSTITUT POLYTECHNIQUE DE PARIS

Measurements and analysis in high-speed software networks

Or... the ***Data Uncertainty Principle***

Leonardo LINGUAGLOSSA, Maître de conférences @ Telecom Paris

Measurements are important

- Analyze a system
- Perform predictions
- Evaluate performance
- Detect anomalies
- Optimize resource usage
- …

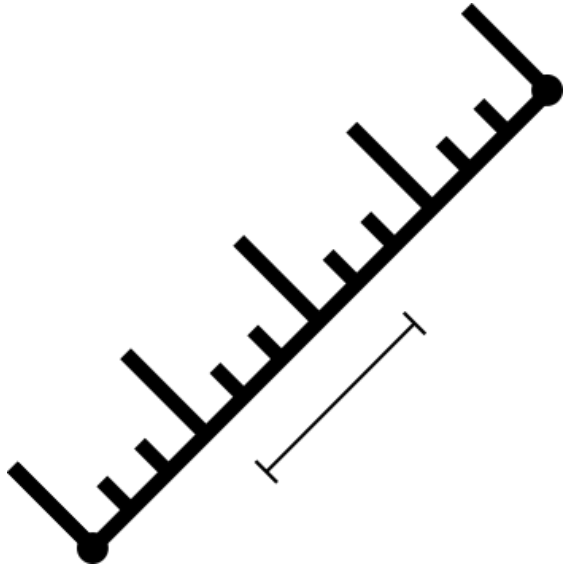

Any analysis is as good as the experimental observations

In physics, at very low level

- Fundamental limit of measurement accuracy of natural systems
- Complementary variables, Heisenberg inequality -> **native property**
- Inherent to all wave-like systems

In physics, the observer effect

- A measured system is altered by the measurement itself
- It can be mitigated by technology or differential measurements
- Inherent to all *macro*scopic systems -> **behavioral alteration**
- Does not set a fundamental limit of the measurement that can be done

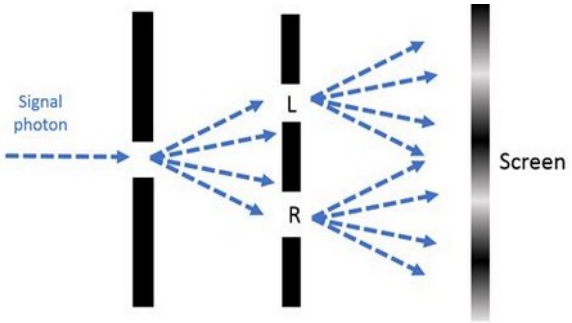

In (data center) networking

- Traffic measurements
- Monitoring performed directly on network devices
- Active or passive
- Direct or indirect
- With or without mirroring
- Deep (per-packet) or sampling (poisson, uniform)
- **Use cases**: anomaly detection, resource allocation, performance enhancement and/or predictions, …

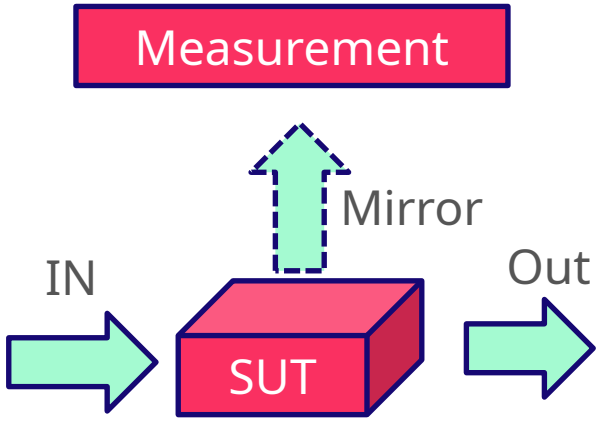

In softwarized networks

- Same as before, but measurements + compute are performed by one (or more) processing elements (CPUs, GPUs, TPUs, DPUs, …)
- All « functions » are implemented as pieces of code

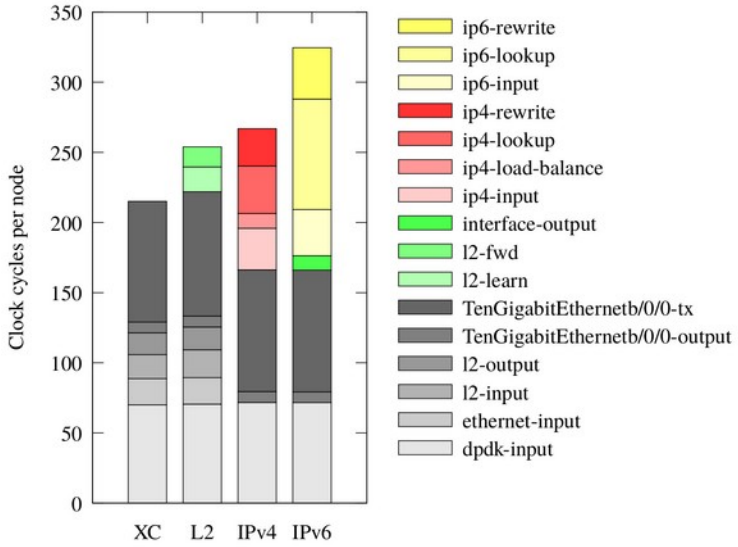

Bonus: **open the white-box!**

Fine-grained data previously not accessible

Data flow in Artificial Intelligence pipelines

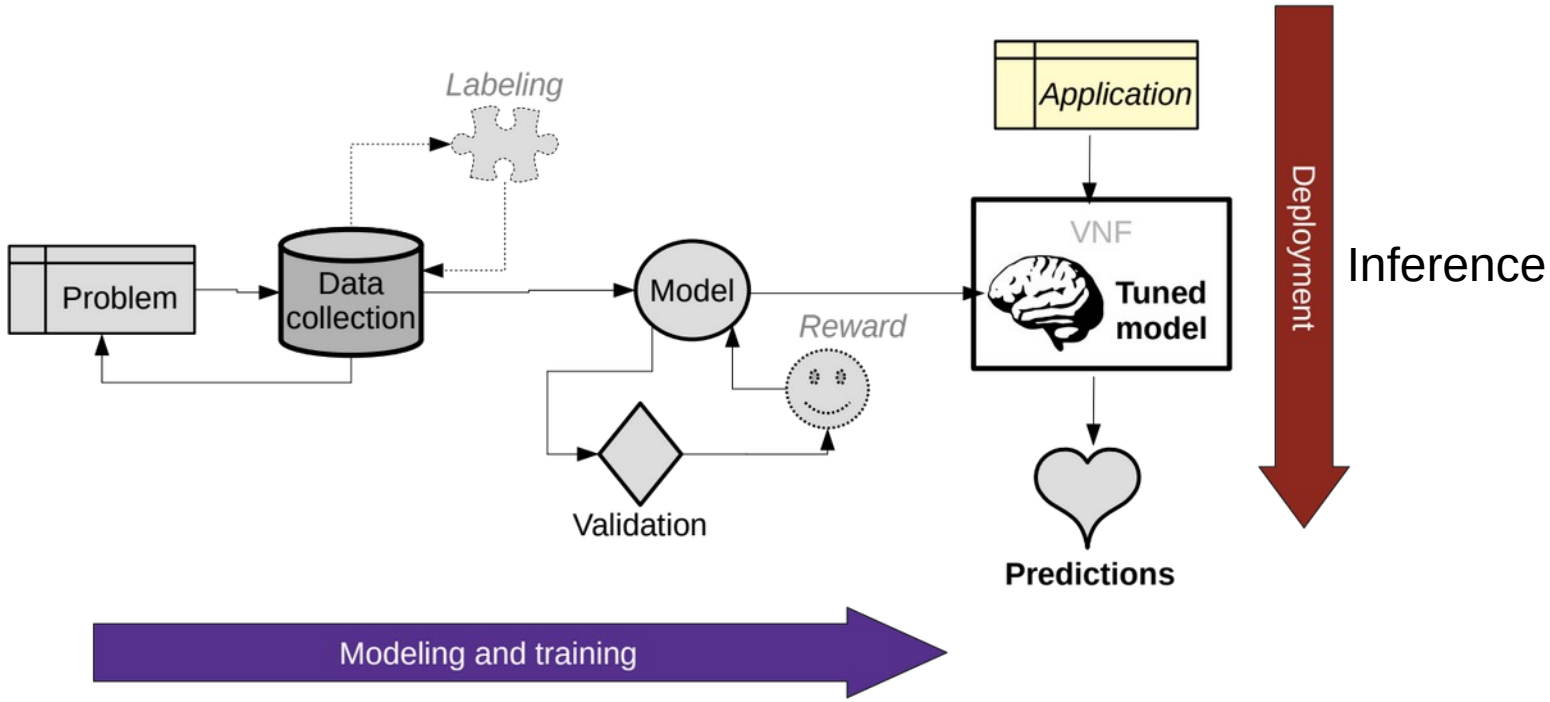

:) Large amount of data for training and use in large-scale data centers

:(Lack of understanding about the coupling of data → models

NOT everything can be measured

(Or do it at your own risk)

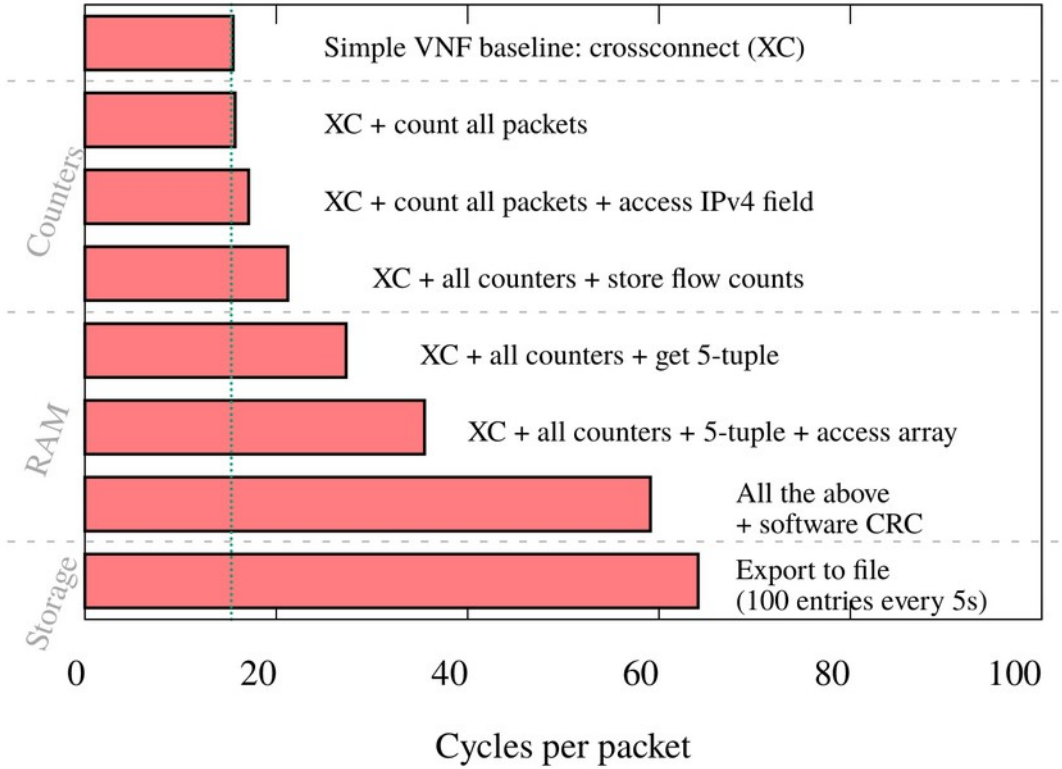

Figure from [TMA2019]

The IONOS^{DX} project

- Measurement problem in high-speed software network: uncertainty/observer effects
- Exploratory project (grant ANR JCJC ANR-22-CE25-0003)
  - Limits of the uncertainty principle
  - Design of non-invasive measurement techniques

In this presentation:

- The measurement problem
  - ~~High-speed software networks, COTS hardware and ML-enhanced functions~~
- Our methodology for non-invasive analysis
- Some use cases and results

The measurement problem

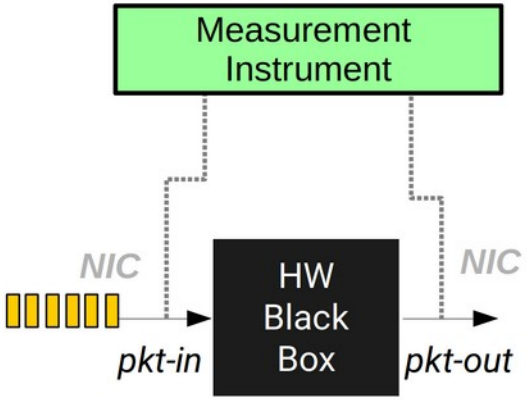

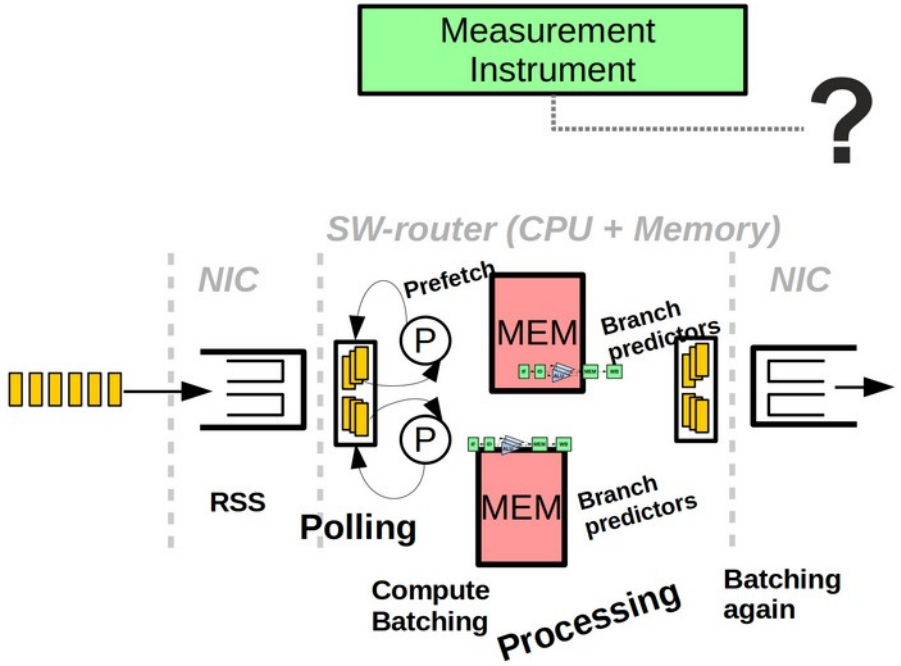

Typical measurements and data collection

- Pre-input (Data collector)
- Input NIC level (mirror, inspection)
- VNF level (pure software)
- CPU level (system level)
- Output NIC level (source demux)
- Post-output (Data collector)

In general: what should we do?

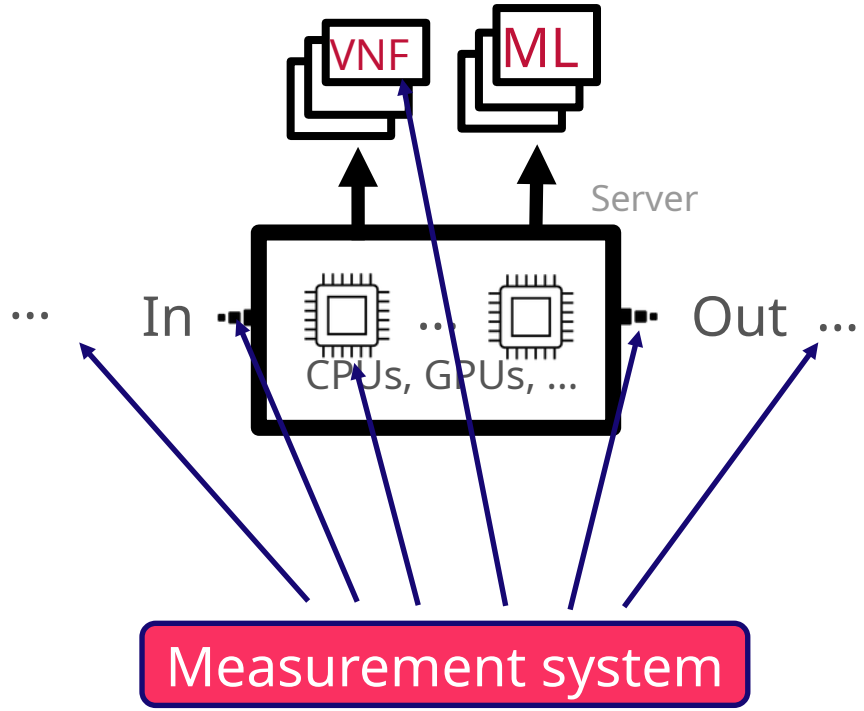

Typical inference scenario

- **Sample application:** detect traffic/state anomalies within a time window
- Precollect several measurements and train very simple ML models
- Deploy the trained model in the data path (for instance, within the orchestrator)
- Infer current state or predict a future state (regression/classification)
  - Predict the evolution of traffic
  - Trigger an alarm in case of anomalies

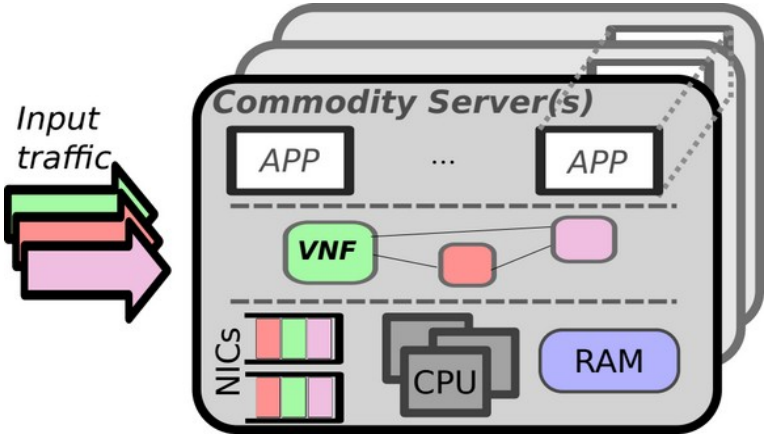

How to collect data on-the-fly

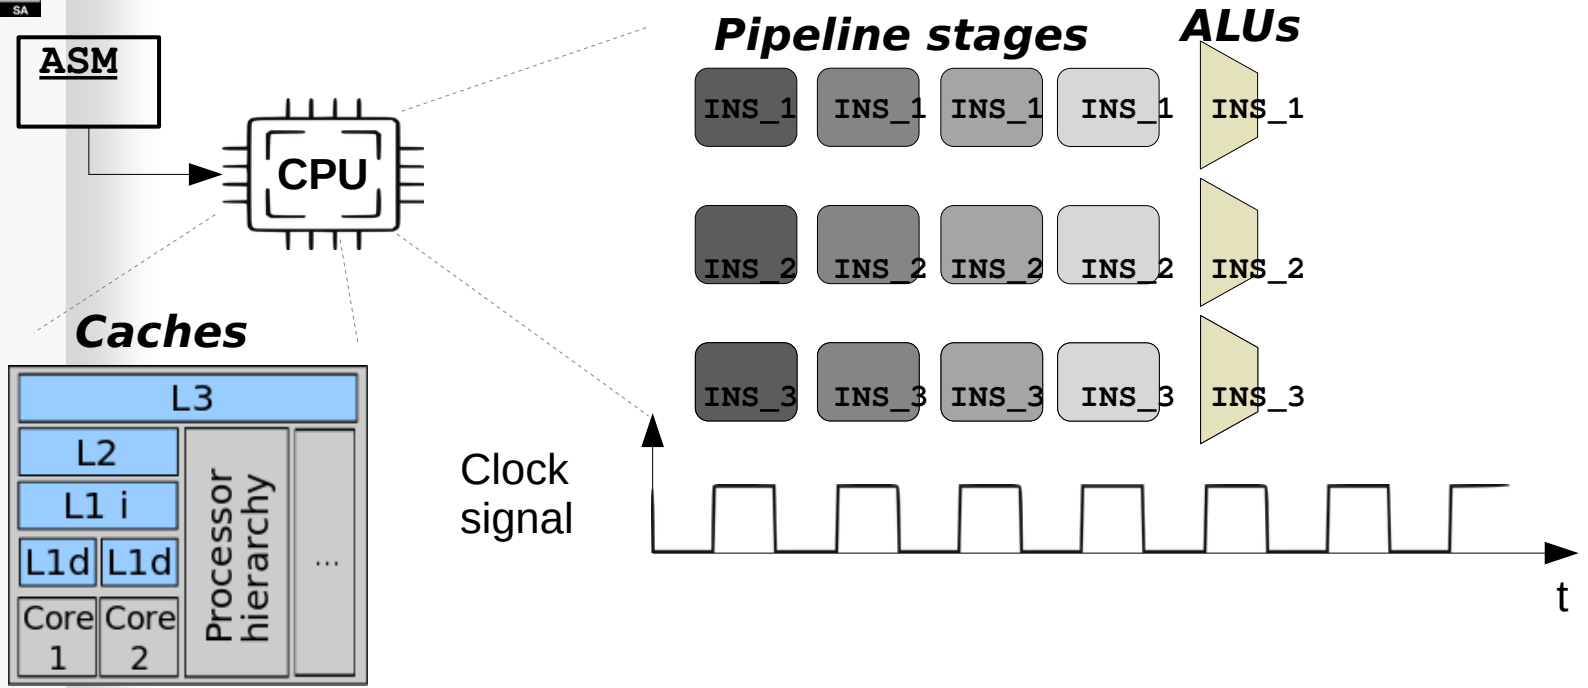

How to collect data on-the-fly

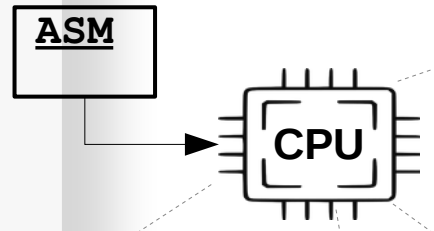

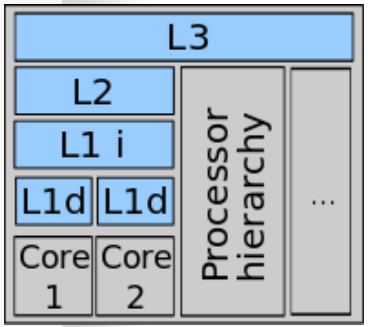

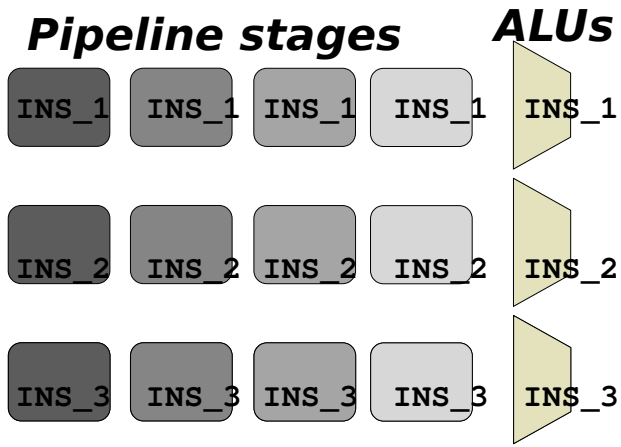

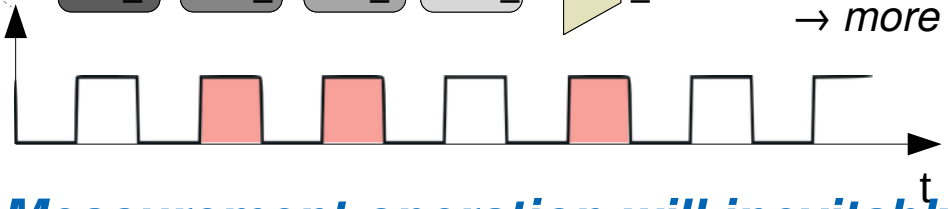

(Clock cycles used for measurement)

More data → more altering

Complex analysis → more altering

Measurement operation will inevitably alter the state of the system.

Methodology: focus on inference phase

- Objective: detect *external KPI* given the *internal stats*
- Computing function: **F, G, …** with given cost *C*
- Measurement function: **m**
- Models: log-reg and MLP (pre-trained)

Ideal processing rate

$$R(F) = \frac{\alpha \cdot f}{C}$$

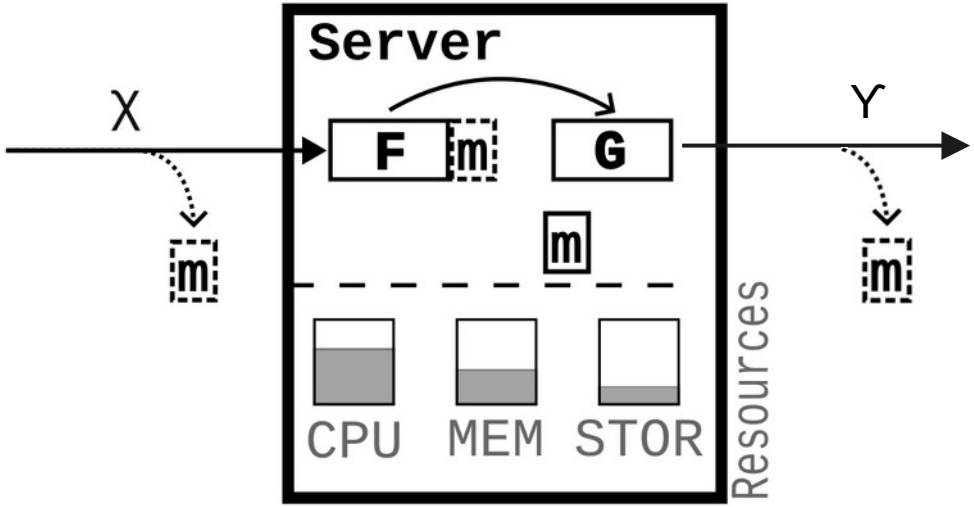

$Y = F \cdot G\,(X)$

Does “m” placement affect the processing rate?
(Yes, *[CNET2024]*)

Idea: use **indirect** and **non-invasive** observations

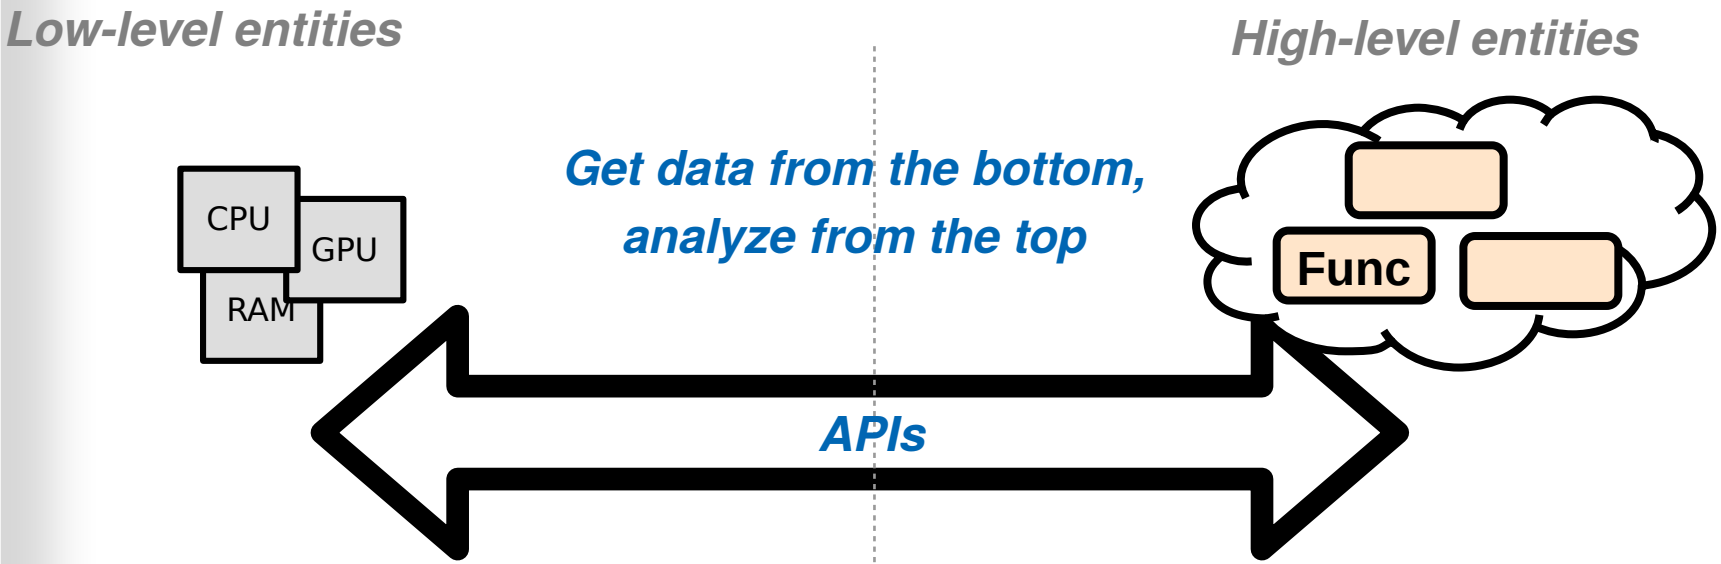

Easy to ***Monitor***
Hard to ***Interpret***
Data availability: ***Huge***

Hard to ***Monitor***
Easy to ***Interpret***
Data availability: ***Low-to-medium***

Case A: detect input traffic (load inference)

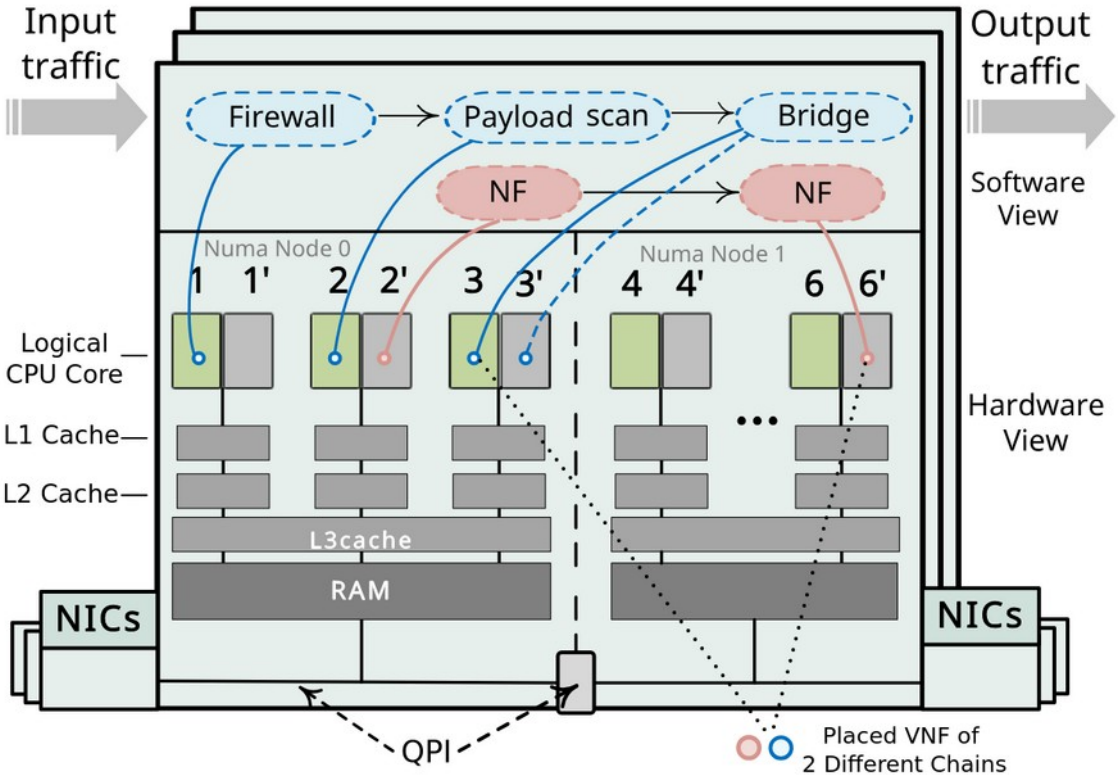

Are low-level features good predictors ?

(Yes, *[ConextStu2019]*)

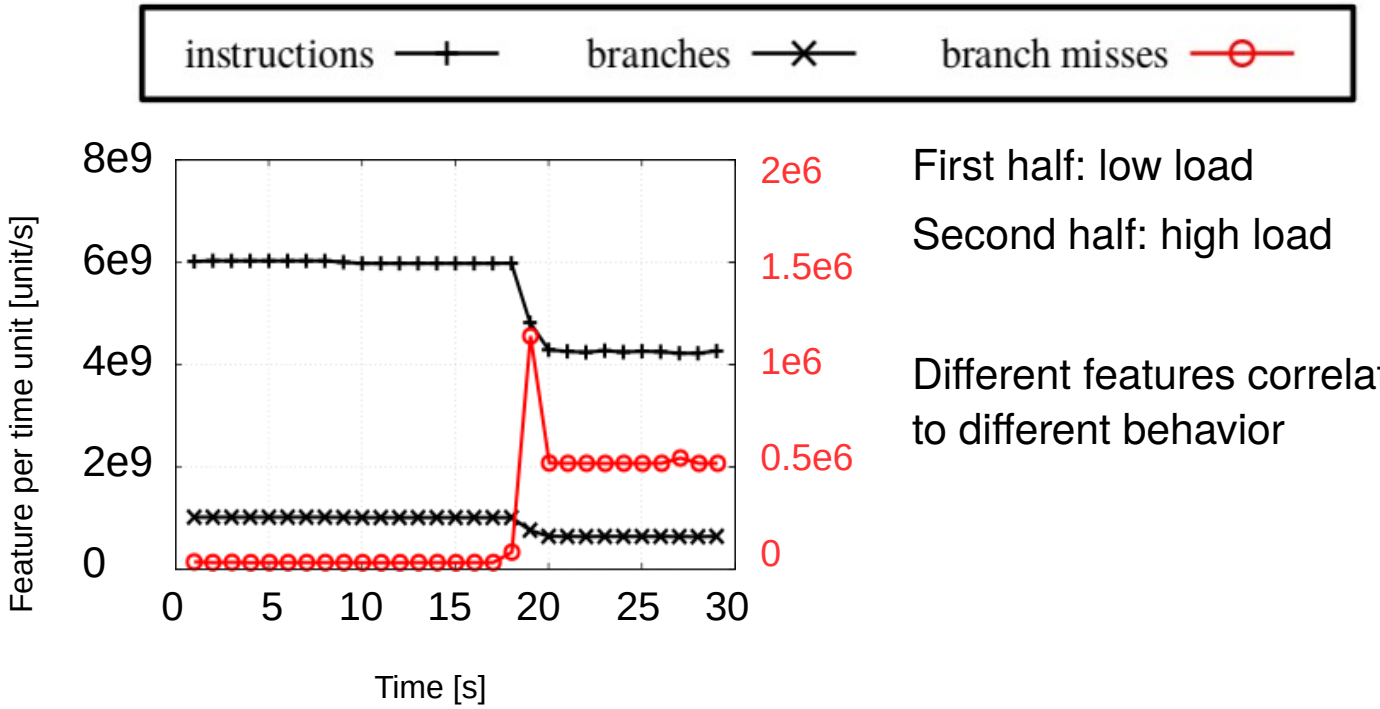

First half: low load

Second half: high load

Different features correlates to different behavior

Are the model performance acceptable? (AI context)

LR

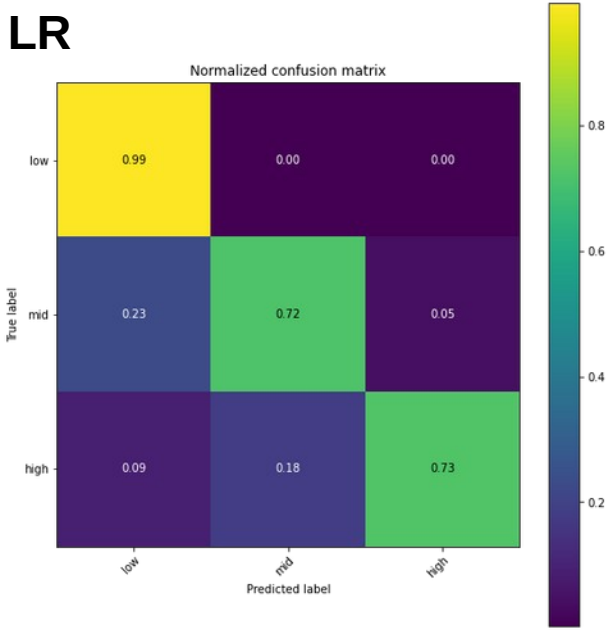

MLP

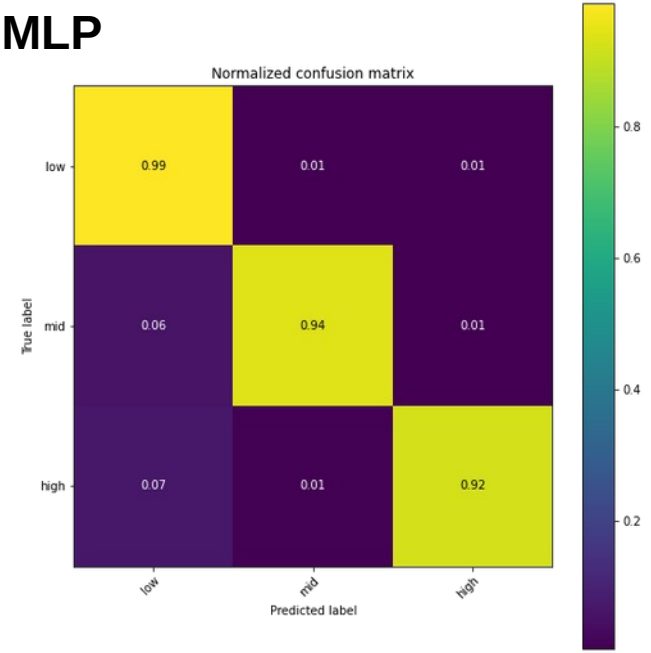

- Objective: detect an anomaly in the processing load
- Left: logistic regression
- Right: multilayer perceptron | 5 layers, 10 neurons per layer

Are the model performance acceptable? (networking context)

It depends [UNDERSUBMISSION]

- The model that performed *the worst*, has 0% packet loss
  - In the data path, there is **not** a **visible alteration**
  - It depends on the initial load

- The model that performed *the best*, has ~45% packet loss
  - For every packet received, a packet is not transmitted
  - The model is altering the state of the system
  - The “anomalies” cannot be detected
  - The system is not the same system to be monitored

- A real-life scenario will have unknown/unquantifiable interactions

Reality is of course even more complex

- Massive data centers, several servers
- Heterogeneous systems
- Multi-tenancy, service function chains

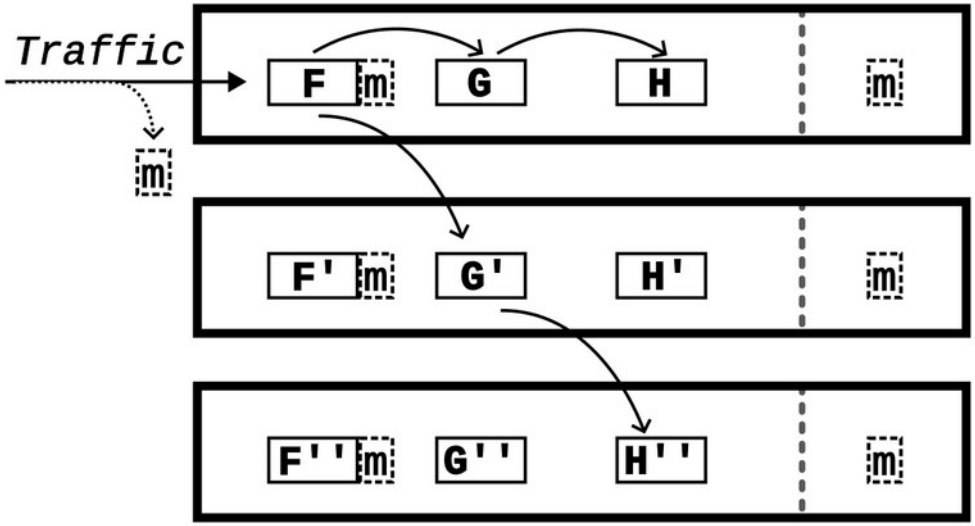

Case B: detect output traffic (prediction)

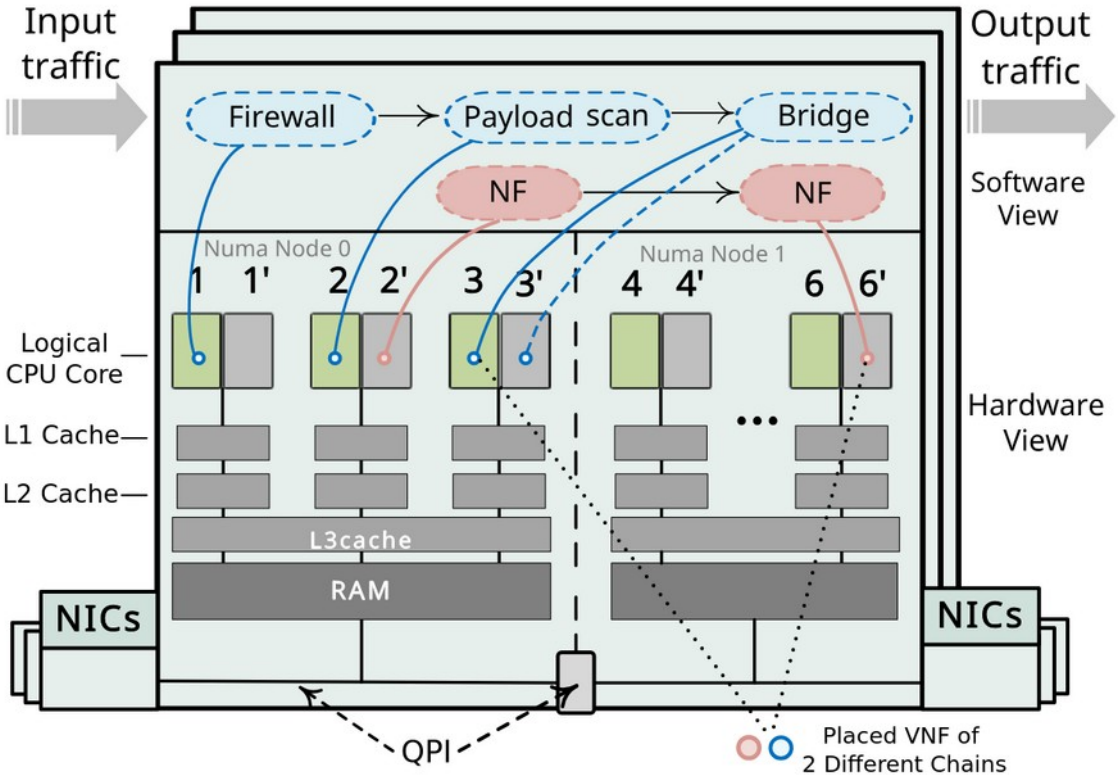

Tradeoffs : speed vs accuracy

- Expected traffic VS predicted traffic
- At high-rates the model struggles -> performance/accuracy tradeoff

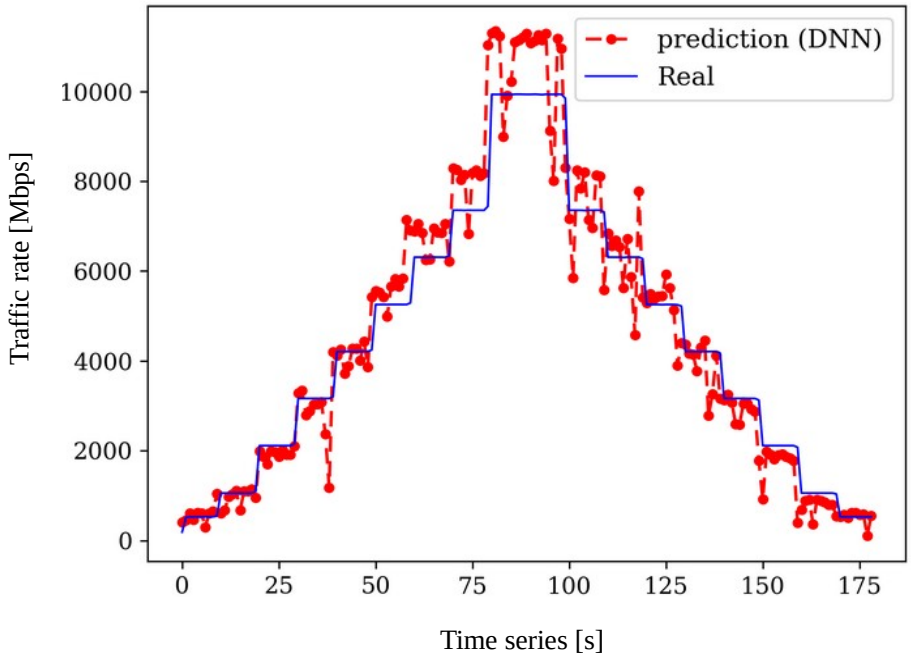

Figure from [INFOCOM2024]

HW architecture may again! completely alter results

- *(Trick: concept of equivalence classes → a tale for another day)*

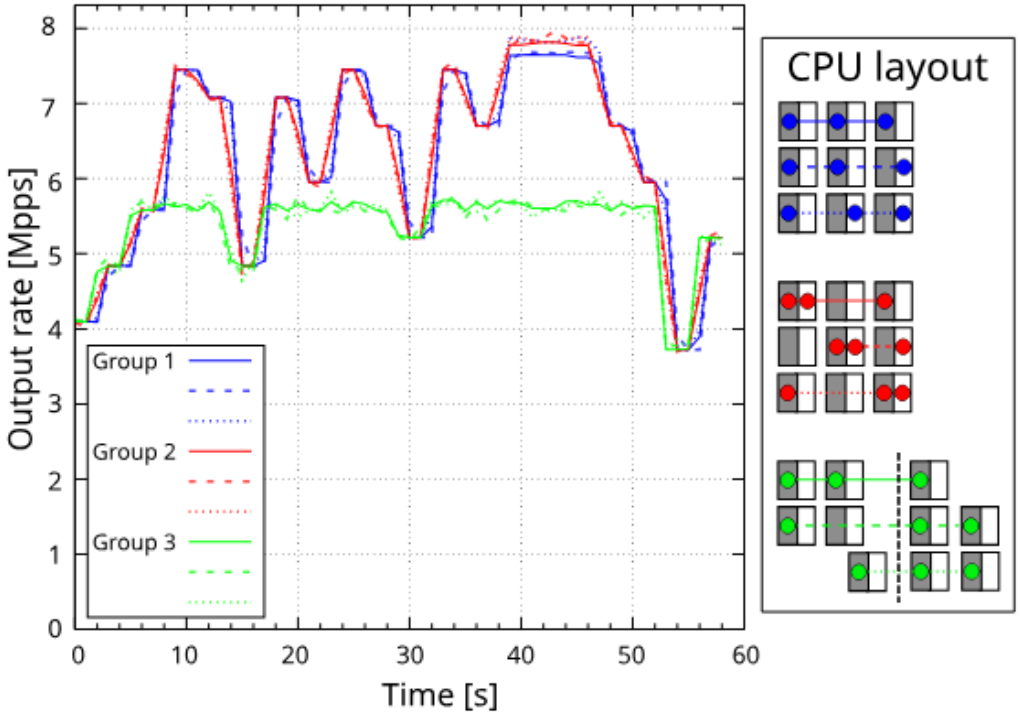

Figure from [UNDERSUBMISSION]

Takeaways and future plan

- Study the **fundamental limits** of SW/HW measurements
- Propose **new methodologies** for network managers, operators and users
- **Key ideas:**
  - (i) indirect measurements
  - (ii) the microarchitecture has an impact
  - (iii) distribute the knowledge sources to improve quality
- **ANR-funded for 4 years**
  - Notion of tradeoff index for a prediction/inference model
    - KPI for multiple metrics to be included
    - Causality vs Data
    - Energy / compute impact of online AI

References and acknowledgements

[IEEEProc2019] Linguaglossa L., Lange S., Pontarelli S., Retvari G., Rossi D., Zinner T., Bifulco R., Jarschel M., Bianchi G., "Survey of Performance Acceleration Techniques for Network Function Virtualization", **Proceedings of IEEE**, 2019

[ConextStu2019] Shelbourne C., Linguaglossa L., Lipani A., Zhang T., Geyer F., "On the Learnability of Software Router Performance via CPU Measurements", **ACM CoNEXT Student Workshop**, 2019

[ITC33] Shelbourne C., Linguaglossa L., Zhang T., Lipani A., "Inference of virtual network functions' state via analysis of the CPU behavior", **ITC 33**, 2021

[CNET2024] Tao W., Linguaglossa L., "A First Look at the Impact of Measurement on Orchestrating Digital Twin Network", **IEEE CloudNet, 2024**

[INFOCOM2024] Liu Q., Zhang T., Linguaglossa L., "Non-invasive performance prediction of high-speed softwarized network services with limited knowledge", **IEEE INFOCOM**, 2024

*(and good luck to my Ph.D. student **Weichen Tao** for his submission 🤞🤞)*

https://ionosdx.wp.imt.fr/

https://anr.fr/Projet-ANR-22-CE25-0003

Thank you for your attention!

Thanks a lot to all my collaborators!

Interested? linguaglossa [at] telecom-paris [dot] fr

Not interested? Feel free to tell me why ^