Contrat doctoral – ED Galilée

Titre du sujet : AI-Driven Diagnosis and Adaptive Observability for Softwarized Network Services

- Unité de recherche : L2TI
- Discipline : Informatique
- Direction de thèse : Thi Mai Trang Nguyen
- Contact : thimaitrang.nguyen@univ-paris13.fr
- Domaine de recherche : Réseaux
- Mots clés : AIOps, 5G, network telemetry, uncertainty quantification

---

Context and motivation. Modern 5G core networks are increasingly softwarized and virtualized, which improves flexibility but also makes runtime diagnosis more difficult. The Network Data Analytics Function (NWDAF) in the 5G network architecture provides a standardized component for analytics-driven service management, supporting the collection of runtime evidence, the production of analytics outputs to be used by other network functions [1]. Connecting such analytics to the internal behavior of high-speed User Plane Functions (UPFs) remains non-trivial because these functions operate under tight per-packet execution budgets. In such settings, observability has a cost and may perturb the forwarding path being observed. This PhD thesis is therefore concerned with diagnosis and adaptive observability in high-speed softwarized 5G networks.

The challenge becomes especially acute in dynamic sliced 5G environments, where observability decisions must be made under shared-resource constraints and heterogeneous service requirements. Nonintrusive observability relies on lightweight signals such as host-side counters, data plane runtime statistics, and coarse service indicators. These signals are attractive because they can be collected continuously with limited overhead, but they often remain indirect and diagnostically ambiguous [2]. By contrast, intrusive observability, including richer in-band telemetry, packet- or flow-level tracing, and fine-grained measurement inside the forwarding path, can provide stronger evidence for diagnosis, but directly competes with forwarding resources and may degrade latency, throughput, or loss performance [3, 4].

Beyond measurement cost, two further difficulties make this problem fundamentally open. First, telemetry from UPFs may change over time because of load variations, co-location effects, or measurement changes, making diagnosis models trained on earlier data unreliable [5]. Runtime drift detection is therefore needed to decide when models should be recalibrated or retrained. Also, diagnosis outputs are also uncertain. Some uncertainty comes from inherent traffic variability and measurement noise, while some comes from incomplete knowledge that richer telemetry could reduce [6]. This raises the question of when adaptive measurement can improve diagnosis, rather than simply adding measurement cost.

Scope of the PhD thesis. The PhD thesis is structured around three successive research directions:

- Establish how lightweight host-side and dataplane signals relate to service-level conditions in highspeed softwarized UPF settings. This includes developing a method on what can and cannot be reliably diagnosed.
- Determine when non-intrusive signals become insufficient for reliable diagnosis, and develop formal criteria for selectively activating richer measurement in a cost- and perturbation-aware manner.
- Investigate how Artificial Intelligence (AI)/ Machine Learning (ML) methods can support the design of a closed-loop interaction between telemetry adaptation and diagnosis within NWDAF.

Methodology. The PhD thesis combines runtime telemetry collection, and service-level diagnosis in highspeed softwarized 5G user-plane environments. Experimental validation will proceed in two stages: first on bare-metal platforms under realistic high-speed UPF conditions, where controlled experiments will establish the diagnostic limits and overhead of different observability regimes [7]; then in more realistic 5G environments, for example SLICES-RI post-5G testbed settings [8], where adaptive observability strategies will be validated under operational constraints.

Collaboration. The PhD thesis is in collaboration with Laboratoire d'Informatique de Paris 6 (LIP6), under the co-supervision of Qiong Liu, Associate Professor at Sorbonne University, and with Anastasios Giovanidis (researcher at Ericsson) as industrial collaborator.

References

[1] 3rd Generation Partnership Project (3GPP). Technical Specification Group Services and System Aspects; Procedures for the 5G System (5GS); Stage 2 (Release 18). Technical Specification TS 23.502, 3GPP, 2024.
[2] Q. Liu et al. Non-invasive performance prediction of high-speed softwarized network services with limited knowledge. In IEEE Conference on Computer Communications (INFOCOM), pages 1–10, 2024.
[3] H.F. Sun et al. Distributed network telemetry with resource efficiency and full accuracy. IEEE/ACM Transactions on Networking, 32(3):1857–1872, 2024.
[4] N. Saha et al. Monarch: Monitoring architecture for 5g and beyond network slices. IEEE Transactions on Network and Service Management, 22(1):777–790, 2024.
[5] Joao Gama et al. A survey on concept drift adaptation. ACM computing surveys (CSUR), 46(4):1–37, 2014.
[6] Eyke Hullermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. Machine learning, 110(3):457–506, 2021.
[7] OpenAirInterface. oai-cn5g-nwdaf. GitLab repository, OpenAir CN 5G for NWDAF, 2026. Available: https://gitlab.eurecom.fr/oai/cn5g/oai-cn5g-nwdaf. Accessed: 2026-03-26.
[8] S. Fdida et al. Slices, a scientific instrument for the networking community. Computer Communications, 193:189–203, 2022.