

# GRIFIN Research Internship: Verifiable network resilience policy translation for stateful data planes

Sébastien Tixeul, Thomas Silverston, Gregory Blanc

Internship start: February 2022

**Keywords** — resilience, network policy, intent-based networking, artificial intelligence, stateful data plane

## Context

With the advent of *Software-Defined Networking* (SDN) and stateful data planes [1], network operation has shifted to increased control and programmability. So much, that pioneers in the field believe that Deep Programmability [2] is round the corner, enabling network owners to gain complete control of their network from top (specifying policies) to down (deciding on which functions for packet processing): in particular, this creates the opportunity for network owners to verify that the network works as intended.

Network verification is proving necessary as bugs, misconfigurations, faults or attacks may introduce inconsistencies between the network control and the data plane [3]. Most approaches to network verification are either static or symbolic, either dealing with intrinsic errors at the programming level or faults at the data plane level. In the former case, the analysis is insufficient as programs alone do not determine how packets are processed and forwarded, but also based on the input. More dynamic approaches are proposed for data plane verification: e.g., by comparing an expected configuration derived from the intended policy and an actual report of the forwarding behaviour from telemetry data [3]. Although this approach yields improved consistency, there are still doubts stemming from the dynamic nature of SDNs, where both the control and data planes may evolve over time, requiring near to continuous inconsistency checks. Additionally, relying on passive monitoring may only provide reactive action upon inconsistency detection, while actively probing could enable prevention. Finally, verification could benefit from formal semantics such as the ones proposed in P4K [4], to provide a more complete and sound verification.

Another avenue to ensure the correctness of the data plane policy, with respect to the control plane policy, may lie in the implementation of a stateful data plane. In this paradigm, a controller offloads two main components to the data plane: state management and processing logic [1]. With this, some network functions can be partially or even fully implemented in the data plane. Still remains the issue of the distribution of the policy into a set of network switches, so as to not overwhelm them with additional processing, in addition to packet forwarding. The policy should be partitioned so as to preserve the intent, and distributed so as to reduce the burden in terms of resource overhead, while still implementing the policy correctly (available functions, consistent chaining).

Thus, Deep Programmability seems promising with respect to controlling the whole pipelining from intent specification to packet processing across a set of switches.

This internship aims at surveying network verification for stateful data planes, and proposing a deep programmable approach including aspects of (i) intent-based networking for specifying high-level policies, (ii) self-driving to compute optimal policy translation using for example artificial intelligence (e.g., reinforcement learning [5]), and (iii) in-network computing to enable precise implementation of intents in the form of both telemetry and packet processing.

## Activities

- survey of network verification in stateful data planes
- study of deep programmable networking and related concepts (intent-based networking, self-driving networking, in-network computation)
- design or extension of data plane language
- design of a translation process from high-level policies to data plane processing
- evaluation of the network verification for network threat mitigation use cases

## Practical information

The internship will take place at SAMOVAR laboratory, Télécom SudParis, Institut Polytechnique de Paris, in Palaiseau. It will be 5 to 6 months long.

Applicants are about to complete their Master 2 level degree (or equivalent engineering school degree) and should have the following skills:

- intermediate to strong knowledge in programmable networks (SDN, stateful data plane, in-network telemetry)
- fundamentals of machine/deep learning (ideally, reinforcement learning)
- concepts in cybersecurity, in particular network attack mitigation
- practice in code development

The internship topic is linked to a Ph.D offer in the context of the GRIFIN project (funded by ANR), a research collaboration between Télécom SudParis, Sorbonne Université and LORIA.

Applications (resume, motivation letter, academic transcripts, recommendation letters) must be sent to `gregory.blanc[at]telecom-sudparis.eu`, `sebastien.tixeuil[at]lip6.fr` and `thomas.silverston@loria.fr`.

## References

- [1] Xiaoquan Zhang, Lin Cui, Kaimin Wei, Fung Po Tso, Yangyang Ji, and Weijia Jia. A survey on stateful data plane in software defined networks. *Computer Networks*, 184:107597, 2021.
- [2] Nate Foster, Nick McKeown, Jennifer Rexford, Guru Parulkar, Larry Peterson, and Oguz Sunay. Using deep programmability to put network owners in control. *ACM SIGCOMM Computer Communication Review*, 50(4):82–88, 2020.

- [3] Apoorv Shukla, Seifeddine Fathalli, Thomas Zinner, Arthur Hecker, and Stefan Schmid. P4consist: Towards consistent p4 sdns. *IEEE Journal on Selected Areas in Communications, Series on Network Softwarization & Enablers (JSAC-SI)*, 2020.
- [4] Ali Kheradmand and Grigore Rosu. P4k: A formal semantics of p4 and applications. *arXiv preprint arXiv:1804.01468*, 2018.
- [5] Elkin Aguas, Anthony Lambert, Grégory Blanc, and Hervé Debar. Automated saturation mitigation controlled by deep reinforcement learning. In *2020 IEEE 28th International Conference on Network Protocols (ICNP)*, pages 1–6. IEEE, 2020.