GRIFIN – Thesis proposal #2 on Cognitive and Programmable Response

Gregory Blanc
SAMOVAR, Télécom SudParis
Institut Polytechnique de Paris

Sébastien Tixeuil
LIP6, CNRS
Sorbonne Université

Thomas Silverston
LORIA, CNRS, Inria
Université de Lorraine

Eric Totel
SAMOVAR, Télécom SudParis
Institut Polytechnique de Paris

Context

Next-generation networks will be characterized by a huge and increasing number of autonomous nodes with computational, communications and storage capabilities, e.g. Internet of Things (IoT). Furthermore, when sensors and actuators are involved, the IoT becomes an instance of the more general class of cyber-physical systems (CPS) [12]. Such systems extend the cyber attack surface into the physical world, with far reaching damages into the society (human casualties included) not only security-wise, but also in terms of safety and privacy. While mission-critical objects such as health devices cannot be halted, attacks against IoT devices usually seek to leverage what little capabilities they have to carry out larger-scale attacks, e.g., botnets. More surreptitious attacks include eavesdropping on the physical environment or the other adjacent devices and leaking the collected information out.

Resilience is the ability of a network to defend against risks (e.g. cyber- attacks, hardware/software faults, human mistakes, natural disasters, etc.) and to maintain an acceptable level of service; security in the sense of self-protection is one important enabling principle for resilience [14]. In GRIFIN, we aim at improving resilience of future network systems by reasoning on their network state in order to select appropriate remediations and actuating the data plane to reliably enforce the selected remediations.

Objectives.

This thesis proposes a cognitive approach to recovering from network incidents or attacks in a future network context. In particular, the approach aims at (i) *modelling the state* of the network, (ii) *automating the selection* of appropriate countermeasures, (iii) *intelligently mapping* high-level measures to security configurations, and (iv) *reliably verifying* that the measures have been deployed and their impact to the network state.

Departing from a static approach to countermeasure selection, the approach will lift locks in adaptive security using artificial intelligence: by modelling the network state continuously, we will be able to observe anomalies and recovery from anomalous situations, while being able to try several possible responses. The approach should converge to a set of optimal policies able to address most common incidents/attacks.

At this level of abstraction, reasoning does not assume anything about the data plane, and there is a need to refine the high level policy into a course of actions, i.e., a(n ordered) combination of remediations that can be directly deployed into network devices,

both physical and virtual. Enabled by software-defined networking, such decoupling between the control plane and the data plane necessitates reconciliation for reliably implementing the countermeasures on the network state. This will be further verified to ensure that the refining process is (a) sovereign, i.e., no third party could tamper the policies and (b) faithful, i.e., the data plane policy implements the high-level intent.

This thesis also proposes to implement this approach for different 5G related use cases, e.g., IoT-based edge computing, SDN-based network slicing or Cloud-RAN.

State of the Art. According to the MITRE corporation [15], cybersecurity situation awareness involves three key areas: 1) network awareness (configuration management and vulnerability auditing), 2) threat awareness (internal malicious behaviour identification and external threat knowledge incorporation), and 3) mission awareness, which still needs to develop a comprehensive picture of critical dependencies and understand them to support forensic analysis, incident response, risk assessment and cyberdefence planning. SDN offers global visibility of the data plane to the network control contributing to network awareness [4], and combined with legacy network flow monitoring capabilities, it has the potential to identify network threats [7], contributing to threat awareness. With regard to mission awareness, there are a number of avenues for contributions (among others): methodologies for risk assessment (RA) or management (RM); dependencies graphs (DG); or intrusion response systems (IRS). A more recent approach balances the costs of deploying security measures and the mitigated loss due to successful breaches on one hand, and the operational impacts on the business assets on the other hand [10]. Finally, the problem of intrusion response given knowledge of the network state can be aptly modelled using a Markov Decision Process (MDP) [3]. The main drawback of most models, although very comprehensive in their ability to capture diverse aspects of information systems, is their static nature. Lately, reinforcement learning has been employed to control and optimize security responses to events, e.g., a DRL-based framework [9] to defend against DDoS attacks by intelligently learning attack traffic patterns in order to throttle it and prioritize legitimate traffic. In the same vein, we proposed a DRL agent [1] to control dynamic prefix load sharing to mitigate path saturation.

Data plane programming languages have evolved into full-fledge programming languages enabling network operators to design and enforce end-to-end switching, routing, or traffic engineering policies. This obviously comes with a number of challenges as operators at the control plane, or at any higher planes, are not expected to be network engineers. To bridge the high-level policies with the per-packet, per-switch forwarding instructions, many intermediate languages have been proposed [13] to which high-level languages compile to. They are thus expressive and yet low-level enough to be efficiently assembled on various device architectures. Nonetheless, such intermediate representations add yet another layer of abstraction and indirection between the network users and the data plane, which may incur disrupting the network with bugs, unpredictable performance, or security vulnerabilities [5]. Besides symbolic execution [5], other approaches include rule-based verification [16] or set theory [8]. The existence of semantics for a number of data

plane languages, such as P4 [2], encourages researchers to investigate rewriting-based approaches as well.

Proposal

In this thesis, by exposing the network control layer using the SDN paradigm, we aim at improving attack/incident mitigation beyond the state of the art, achieving deep programmability [6], i.e., enabling network owners to specify policies at a high level and select packet processing functions at a low level. To that end, the solution leverages: 1) intent-based networking to enable users in freely designing high-level network policies; 2) software-defined networking to decouple the control plane from the data plane and enables network programmability; 3) ML-based policy optimization to select countermeasures regardless of any prior knowledge; 4) enhanced data plane programming language to facilitate policy pipelining and distribution; 5) formal or auditable verification of policies deployed at enforcement points and network equipments in compliance with the intent of the network owners.

Activity 1 – ML-based countermeasure selection. From a set of known security measures (including redirection, filtering, throttling, blocking), the approach will pair the network states with courses of actions, that is sets of countermeasures to apply to the network environment. Several machine learning approaches will be considered, including model-free and model-based reinforcement learning (RL), as massive IoT data is required for the model-free approach. Thus, resorting to IoT traffic generators, such as the one we proposed [11] will be of high value. Additionally, to cater to a realistic network environment, the actions and state spaces may not be discrete, which actually prompt for exploring Deep RL approaches.

Activity 2 – Intent modeling and pipelining. Once the countermeasures have been selected, the actuating component is left with a policy to implement. Such policy is usually specified using a high- level language. The proposed pipeline should be able to faithfully convey the policy across the data plane, while optimising the load on the data plane devices, in terms of computation and storage, but also in terms of location. This means that the pipelining subcomponent should communicate with the reasoning component to comprehend the situation and output an appropriate pipeline. Additionally, some functions defined in the data plane programs may not be directly amenable by the data plane device, depending on the expressiveness/richness of the data plane language. In such cases, a study of how to approximate data plane functions will be carried out to propose suitable approximations.

Activity 3 – Policy verification. Due to the approximations of the pipelining stage and the functions at the data plane, the deployed policy may actually differ from the high-level policy. Several threats may also affect the resulting policy at the data plane, such as hijacking the control channel between the actuating component and the data plane devices, or

compromising a data plane device. There is a need to ascertain that the running policy is in line with the high-level policy. Leveraging existing semantics, a verifier will be developed to check for discrepancies at the data plane.

Expected Impacts and Perspectives. The research work to be carried out in this thesis will contribute to the ANR PRC GRIFIN project (ANR-20-CE39-0011) which aims at improving the overall resilience of future networks by augmenting the MAPE-K cybersecurity loop with AI systems. In GRIFIN, the Ph.D candidate will have the opportunity to constitute a network of expertise in several communities including 1. *cybersecurity* where security mitigation is a hot topic, as evidenced by C&ESAR 2021 conference theme; 2. *IoT* where the application of DL and SDN to secure them is gaining traction but yet to demonstrate applicable results; 3. *AI* for which novel, or even, incremental results on applying DRL to network environment is emerging. This kind of work, which can be applied to application areas such as healthcare or industry, has potential for short- to medium-term technological transfer. Avenues for collaboration not only include European projects but as well the German-French Academy for the Industry of the Future, or the Japanese Industrial Cybersecurity Centre of Excellence (ICS-CoE), with whom we have lasting collaboration.

Supervising Team. Supervision will be shared between the R3S team at SAMOVAR (Télécom SudParis, IP Paris), the NPA team at LIP6 (Sorbonne Université, CNRS) and LORIA (Université de Lorraine, Inria, CNRS). SAMOVAR has a long expertise in attack mitigation in SDN environments and cyber-physical systems, as well as in countermeasure selection and optimization. R3S team has collaborated on such topics in European projects (ITEA3 ADAX, FP7 NECOMA (includign Japanese partners), FP7 PANOPTESEC, H2020 SUPERCLOUD, H2020 SOCCRATES). LIP6 has recognized expertise on fault and attack tolerance in networks and distributed systems (FP7 MOTO, ANR SHAMAN and SAFEOS), as well as mechanical verification and certification of distributed algorithms (ANR SAPPORO). LORIA is highly involved in next-generation network topics, including IoT, related to cybersecurity in general, and traffic analysis in particular (ANR DOCTOR). Additionally, in GRIFIN, the Ph.D candidate will have opportunities to collaborate with esteemed academics in the UK (UCL, MDX) as well as integrate the results with dynamic SMEs (Montimage).

References

[1] E. Aguas, A. Lambert, G. Blanc, and H. Debar. Automated Saturation Mitigation Controlled by Deep Reinforcement Learning. In *Proc. of ICNP'20*.

[2] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, and D. Walker. P4: Programming protocol-independent packet processors. *ACM SIGCOMM Computer Communication Review*, 44(3):87–95, 2014.

[3] F. Charmet, G. Blanc, and C. Kiennert. Optimizing Resource Allocation for Secure SDN-based Virtual Network Migration. In *Proceedings of NCA'2019*.

[4] J. Chen, Y. Ma, H. Kuo, C. Yang, and W. Hung. Software-Defined Network Virtualization Platform for Enterprise Network Resource Management. *IEEE Transactions on Emerging Topics in Computing*, 4:179–186, 2016.

[5] M. Dobrescu and K. Argyraki. Software Dataplane Verification.

[6] N. Foster, N. McKeown, J. Rexford, G. Parulkar, L. Peterson, and O. Sunay. Using deep programmability to put network owners in control. *ACM SIGCOMM Computer Communication Review*, 50(4):82–88, 2020.

[7] K. Giotis, C. Argyropoulos, G. Androulidakis, D. Kalogeras, and V. Maglaris. Combining OpenFlow and sFlow for an effective and scalable anomaly detection and mitigation mechanism for SDN environments. *Computer Networks*, 62:122–136, 2014.

[8] T. Inoue, R. Chen, T. Mano, K. Mizutani, H. Nagata, and O. Akashi. AN Efficient Framework for Data-Plane Verification with Geometric Windowing Queries. *IEE Transactions on Network and Service Management*, 14(4):1113–1127, 2017.

[9] Y. Liu, M. Dong, K. Ota, J. Li, and J. Wu. Deep Reinforcement Learning based Smart Mitigation of DDoS Flooding in Software-Defined Networks. In *Proc. of CAMAD'18*.

[10] A. Motzek, G. Gonzalez-Granadillo, H. Debar, J. Garcia-Alfaro, and R. Möller. Selection of Pareto-efficient response plans based on financial and operational assessments. *EURASIP Journal on Information Security*, 2017.

[11] H. Nguyen-An, T. Silverston, T. Yamazaki, and T. Miyoshi. Generating IoT Traffic: a Case Study on Anomaly Detection. In *Proc. of LANMAN'20*.

[12] F. Restuccia, S. D'Oro, and T. Melodia. Securing the Internet of Things in the Age of Machine Learning and Software-Defined Networking. *IEEE Internet of Things Journal*, 5(6):4829–4842, 2018.

[13] M. Shahbaz and N. Feamster. The Case for an Intermediate Representation for Programmable Data Planes.

[14] J. Sterbenz, D. Hutchison, E. Cetinkaya, A. Jabbar, J. Rohrer, M. Schöller, and P. Smith. Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. *Computer Networks*, 54(8):1245–1265, 2010.

[15] The MITRE Corporation. Cybersecurity Situation Awareness.

[16] Y. Tseng, Z. Zhang, and F. Naït-Abdesselam. SRV: Switch-based rules verification in sotfware defined networking. In *Proc. of NetSoft'16*.