Formal models for threat analysis in next-generation networks

Formal models for threat analysis in next-generation networks

Gianmarco Bachiorrini, Daniele Bringhenti, Fulvio Valenza
Dip. Automatica e Informatica
Politecnico di Torino
Torino, Italy
Emails: {first.last}@polito.it

***Abstract*—Modern computing and networking infrastructures are becoming increasingly complex, heterogeneous, and dynamic, integrating paradigms such as cloud computing, cyber-physical systems, and next-generation mobile networks. At the same time, cyber attacks are evolving in sophistication, often unfolding across multiple stages and spanning cyber, physical, and human domains. Threat modeling and analysis are fundamental proactive cybersecurity activities aimed at identifying such attacks. However, current approaches often struggle to model complex attack scenarios and frequently overlook the human and physical dimensions of attacks. This PhD research aims to design a novel threat modeling and analysis methodology combining high modeling expressivity, formal rigor and automation. The approach relies on a threat modeling taxonomy to formally represent system components, vulnerabilities, threats, and relationships, enabling the construction of a structured knowledge base. An automated analysis engine applies formal derivation rules to infer security properties and reconstruct multi-step attack scenarios. An interactive analysis phase further supports explainability and decision-making. The expected outcome is a formally grounded methodology for analyzing complex threats across heterogeneous, multi-domain infrastructures.**

***Index Terms*—formal models, threat modeling, threat analysis**

I. Introduction

Cybersecurity is a scientific field that evolves under constant pressure, as defensive technologies and attack techniques continuously co-evolve. This challenge is further amplified by the increasing complexity and dynamism of modern networked infrastructures, in which both defenders and attackers must adapt their strategies. New networking paradigms are emerging at a rapid pace to address increasing demands in resource allocation and communication efficiency, often coexisting and overlapping with one another. Examples include software-defined networking, cloud and liquid computing paradigms, the Internet of Things (IoT), and 5G/6G technologies. Across these paradigms, a common element that persists is the presence of valuable assets, either critical to service delivery or responsible for storing sensitive data, which therefore require appropriate protection mechanisms. The identification of such protection mechanisms is inherently a proactive cybersecurity activity, for which threat modeling and analysis represent fundamental processes.

Threat modeling, according to NIST [1], is "a form of risk assessment that models aspects of the attack and defense sides of a logical entity, such as a piece of data, an application, a host, a system, or an environment". The primary objective of this process is to construct a threat model, i.e., a structured representation of the assets, vulnerabilities, threats, and relationships relevant to the security of a system or infrastructure. Threat analysis is the next step in the proactive cybersecurity pipeline: by analyzing the threat model, cybersecurity designers can assess the infrastructure's security posture, identify potential attack scenarios, and evaluate the effectiveness of the defensive measures represented in the model.

Despite their importance, many threat modeling and analysis approaches defined in literature for cybersecurity management in computer networks still rely on manual knowledge elicitation and limited automation. As a consequence, they often produce static models that do not easily capture the evolution of attack scenarios over time. Moreover, the heterogeneous nature of modern networking environments and paradigms requires threat modeling ontologies to be sufficiently flexible to adapt to the different contexts in which cybersecurity designers operate. However, this requirement for flexibility often conflicts with the intrinsic nature of threat modeling methodologies, which are typically grounded in formal methods. An additional shortcoming of existing research in this field, which is often overlooked, is the lack of approaches capable of modeling not only the cyber aspects of infrastructures and involved actors, but also their human and physical dimensions. Recent offensive security studies have demonstrated that the interdependencies between these domains can be exploited by attackers performing hybrid attacks, as shown in [2].

In this context, this PhD research aims to design and develop a novel threat modeling and analysis methodology capable of addressing the limitations of current approaches. In particular, the research focuses on combining high modeling expressivity, strong formal foundations, and advanced automation capabilities, in order to support the systematic analysis of complex and evolving threat scenarios. The envisioned methodology is designed to support the modeling and analysis of multi-step attack scenarios, while remaining sufficiently flexible to be applied across heterogeneous technological environments and multi-domain infrastructures, including cyber, physical, and human components.

The remainder of this paper is structured as follows. Section II reviews the literature related to the PhD research field. Sec-

tion III describes the proposed research methodology. Section IV presents the conclusions and future research directions.

II. Related Work

The literature related to this research area is broad and heterogeneous. However, many recent studies focus on applying well-established threat modeling methodologies to specific scenarios and application domains, rather than proposing new threat modeling and analysis methodologies. For example, in [3], the STRIDE methodology is applied to cloud infrastructure, particularly within container-based ecosystems, to model threats and vulnerabilities and subsequently evaluate existing mitigation strategies for the identified threats. Similarly, in [4], the PASTA threat modeling process is employed to analyze threats targeting systems in vehicular networks, with particular focus on AI-based attacks.

At the same time, the literature includes several studies more closely aligned with the objectives of this PhD research, e.g., [5], [6]. For instance, in [5], a novel formal threat model is introduced for the security management of cyber-physical systems. Rooted in the π-calculus, that approach seeks to model in detail both cyber-physical systems and adversary behavior, and uses attack-defence trees (ADT) to reason about the necessary conditions for attacks to succeed. However, this level of modeling detail comes at the cost of reduced generality, as the considered cyber threats are primarily limited to denial-of-service and man-in-the-middle attacks affecting sensor measurements or control communications. Instead, Yang and Zhao [6] focus on providing a general framework for building visualization models based on Knowledge Graphs to support threat modeling and analysis activities, with the objective of integrating them with the MITRE ATT&CK framework. While their approach maintains a lightweight and adaptable structure suitable for different cybersecurity ontologies, it does not provide a formal representation of relationships between entities, nor does it define formal inference rules for updating the knowledge base, instead relying on external tools to derive such information.

Finally, the studies [7] and [8] represent the foundational pillars upon which this PhD research is built. Initially, [7] laid the groundwork for a threat modeling and analysis approach, named FATHoM, based on formal methods and automated reasoning. By formally defining both the relationships among entities and the derivation rules governing the analysis process, FATHoM enables security designers to automatically infer security properties and verify the consistency of the produced model with respect to the designer's assumptions. The main limitation of FATHoM is that it is limited to virtualized environments. Subsequently, TAMELESS, proposed in [8], went beyond this limitation of FATHoM, by extending, modifying, and enriching its original model through the introduction of additional relationship types, security properties, and derivation rules, that enables the modeling and detection of threats across cyber, physical, and human domains, effectively resulting in a hybrid threat model. Still, neither [7] nor [8] provide a threat analysis process suitable for modeling multi-step attack scenarios occurring in heterogeneous and big computer networks. In fact, both approaches produce static threat models, requiring the designer to manually update the initial assumptions to simulate different stages of a multi-step attack. Addressing this limitation is one of the key challenges this research targets.

III. Research Methodology

In light of the gaps identified in the existing literature, this PhD research aims to propose a novel threat modeling and analysis methodology. To advance the research on the topic, the proposed methodology should be designed with the following characteristics in mind:

- **Expressivity**: the methodology should enable cybersecurity designers to accurately represent the infrastructure under analysis through well-defined and unambiguous sets of entities and relationships, suitable for human, cyber and physical domains. This expressivity should support the modeling of complex interactions among system components and facilitate a comprehensive and systematic threat analysis process.
- **Formal rigor**: the methodology should be grounded in a solid formal foundation, ensuring internal consistency, soundness, and reliability. Such rigor is essential to support trustworthy threat modeling and analysis activities in highly dynamic and evolving technological environments.
- **Automation**: the methodology should provide a high degree of automation, minimizing manual effort for cybersecurity administrators and thus improving scalability and operational efficiency when analyzing large-scale, complex infrastructures.

Moreover, as highlighted in Section II, it is fundamental in the current cybersecurity landscape that modern threat modeling frameworks can model and detect multi-step attack scenarios, while adopting appropriate visualization strategies to support their analysis.

At the current stage of the PhD research, a preliminary architecture for the proposed threat modeling and analysis methodology has been defined, along with a fully operational tool implementation that demonstrates the feasibility of the approach and its ability to achieve the aforementioned characteristics. The architecture, illustrated in Fig. 1, also provides a high-level overview of the operational workflow of the proposed approach, which can be decomposed into two fundamental phases.

Threat Modeling

The first phase of the envisioned methodology focuses on threat modeling. In this phase, the cybersecurity designer leverages the proposed Threat Modeling Taxonomy to formally represent the system under analysis by means of entities and relationships. These constructs are used to capture both the constituent elements of the system and the interactions among them. Specifically, the entities defined in the taxonomy enable the formal representation of system components, vulnerabilities, and threats, while relationships are employed

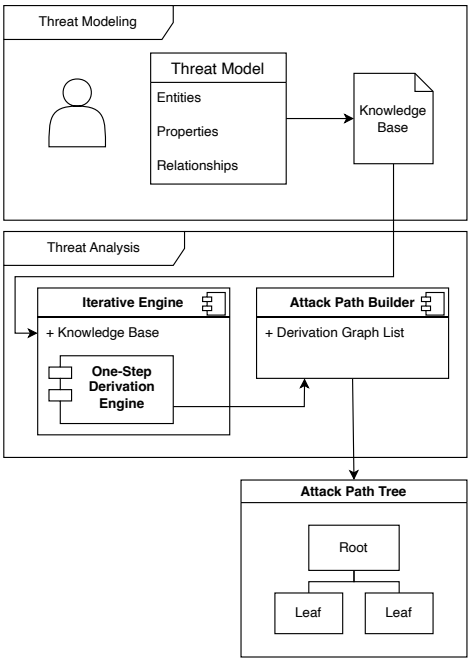

Fig. 1. Logical Architecture

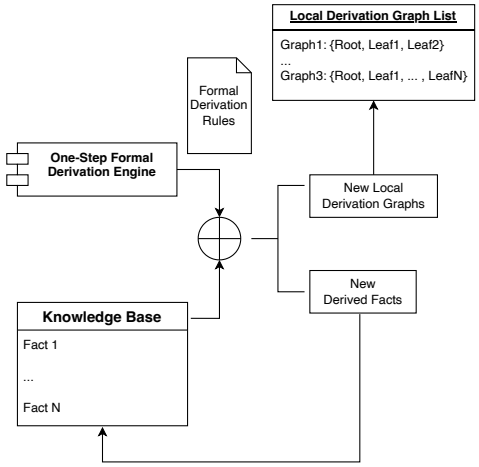

Fig. 2. Iterative Phase of Threat Analysis

to express semantic associations among them, such as the presence of known vulnerabilities in architectural components or the exploitation of specific vulnerabilities by threats to carry out an attack. Completion of this phase results in the construction of a Knowledge Base that encapsulates all the information required to support the subsequent phase of the proposed methodology.

Threat Analysis

While the expressivity of the proposed methodology is ensured by the Threat Modeling Taxonomy, formal rigor and automation are inherent to the Threat Analysis stage. This stage of the methodology can be further divided into two sub-phases: an iterative phase and an interactive phase.

The iterative phase (illustrated in Fig 2) begins when the cybersecurity designer provides the Knowledge Base to the Threat Analysis Engine, which is subsequently processed by the One-Step Derivation Engine. This engine is equipped with a set of formal derivation rules that are automatically applied to infer security-related properties of the entities described in the Knowledge Base. These derived properties include, for example, whether entities are exposed to specific threats, under which conditions they become compromised, and whether they can be successfully defended. Whenever the engine infers a new fact, a corresponding Derivation Graph is generated. This graph represents the reasoning steps leading to the inference and is stored for subsequent use in the interactive phase. After the Knowledge Base has been fully processed, it is enriched with the newly derived information, and the One-Step Derivation Engine iteratively (and automatically) reapplies its derivation rules to the updated Knowledge Base in order to infer additional security properties. This process continues until a fixpoint is reached, meaning that no further facts can be derived.

At this final stage, the cybersecurity designer interacts with the methodology by submitting queries concerning the derivation of specific facts. In response, the methodology exploits the Attack Path Builder module, whose workflow operates as follows. First, the Derivation Graph associated with the queried fact, previously computed during the iterative phase, is retrieved. Subsequently, for each leaf node of the graph, the module searches its internal repository to determine whether a corresponding Derivation Graph has been computed during the iterative phase, effectively assessing whether the leaf represents an intermediate step in the attack chain or an initial condition of the attack scenario. When such a graph is found, the original Derivation Graph is expanded and merged with the retrieved one, as shown in a simple example in Fig. 3.

This recursive process enables the construction of deep Attack Path Trees representing complex, multi-step attack scenarios, where each level of the tree corresponds to an individual attack step.

IV. Conclusion and Future Work

This paper presented an overview of an ongoing PhD research aimed at addressing key limitations of existing threat modeling and analysis approaches in modern networked infrastructures. By combining a highly expressive threat modeling taxonomy with formal reasoning mechanisms and automated analysis, the proposed methodology supports the systematic reconstruction and interpretation of complex, multi-step attack scenarios across cyber, physical, and human domains. The resulting representation enables cybersecurity designers to reason about security properties and attack paths that would

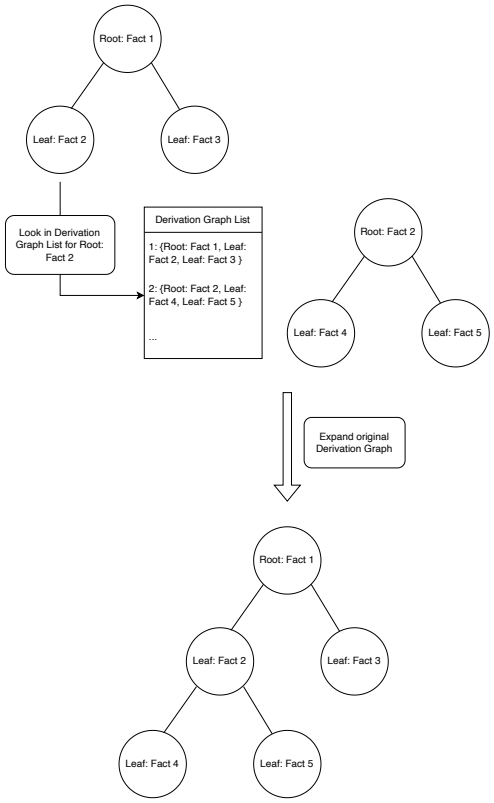

Fig. 3. Example of Attack Path Tree expansion

be difficult to identify manually, especially in heterogeneous and dynamic next-generation networks.

While the proposed methodology already demonstrates the feasibility and benefits of a formally grounded and automated threat analysis process, several research challenges remain open in order to further improve scalability, adaptability, and practical applicability in real-world deployments. In particular, two future research directions identified as central to this PhD research program are the following:

1) **Automatic extraction of threat intelligence**: Currently, a limitation of the envisioned approach is that it still relies on the cybersecurity designer to manually gather information related to vulnerabilities and threats. A relevant research direction is the integration of automated threat intelligence extraction mechanisms. In particular, LLM-assisted techniques could be explored to automatically extract relevant threat intelligence from trusted sources, such as CVE databases or the MITRE ATT&CK knowledge base, and to contextualize it with respect to the target infrastructure.
2) **Automatic generation of threat hypotheses**: Threat modeling approaches, by design, focus on modeling and analyzing known threats and vulnerabilities, either manually specified by the cybersecurity designer or automatically extracted, as proposed in the first research direction. A complementary research direction is the design of approaches capable of supporting the autonomous generation of plausible threat hypotheses based on the structural and semantic properties of the modeled infrastructure. In particular, AI-based techniques could be explored to suggest potential multi-step attack scenarios involving unknown or emerging threats, which can then be formally analyzed and validated through the reasoning mechanisms of the proposed methodology.

ACKNOWLEDGMENT

This project has received funding from the European Union's Horizon Europe Research and Innovation Programme under grant agreement No. 101168144 (MIRANDA). The views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union or the European Cybersecurity Competence Centre. Neither the European Union nor the granting authority can be held responsible.

REFERENCES

[1] NIST. (2020) Security and privacy controls for information systems and organizations. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf

[2] J. Staggs, D. F. Ferraiolo, and S. Shenoi, "Wind farm security: attack surface, targets, scenarios and mitigation," *Int. J. Crit. Infrastructure Prot.*, vol. 17, pp. 3–14, 2017. [Online]. Available: https://doi.org/10.1016/j.ijcip.2017.03.001

[3] A. Y. Wong, E. G. Chekole, M. Ochoa, and J. Zhou, "On the security of containers: Threat modeling, attack analysis, and mitigation strategies," *Comput. Secur.*, vol. 128, p. 103140, 2023. [Online]. Available: https://doi.org/10.1016/j.cose.2023.103140

[4] N. Pape and C. Mansour, "PASTA threat modeling for vehicular networks security," in *7th International Conference on Information and Computer Technologies, ICICT 2024, Honolulu, HI, USA, March 15-17, 2024*. IEEE, 2024, pp. 474–478. [Online]. Available: https://doi.org/10.1109/ICICT62343.2024.00083

[5] L. O. Nweke, G. K. Weldehawaryat, and S. D. Wolthusen, "Threat modelling of cyber-physical systems using an applied π-calculus," *Int. J. Crit. Infrastructure Prot.*, vol. 35, p. 100466, 2021. [Online]. Available: https://doi.org/10.1016/j.ijcip.2021.100466

[6] Y. Yang and Y. Zhao, "Network security threat intelligence modeling based on knowledge graph," in *2024 5th International Conference on Computer, Big Data and Artificial Intelligence (ICCBD+AI)*, 2024, pp. 154–158. [Online]. Available: 10.1109/ICCBD-AI65562.2024.00034

[7] D. Sgandurra, E. Karafili, and E. Lupu, "Formalizing threat models for virtualized systems," in *Data and Applications Security and Privacy XXX - 30th Annual IFIP WG 11.3 Conference, DBSec 2016, Trento, Italy, July 18-20, 2016. Proceedings*, ser. Lecture Notes in Computer Science, S. Ranise and V. Swarup, Eds., vol. 9766. Springer, 2016, pp. 251–267. [Online]. Available: https://doi.org/10.1007/978-3-319-41483-6_18

[8] F. Valenza, E. Karafili, R. V. Steiner, and E. C. Lupu, "A hybrid threat model for smart systems," *IEEE Trans. Dependable Secur. Comput.*, vol. 20, no. 5, pp. 4403–4417, 2023. [Online]. Available: https://doi.org/10.1109/TDSC.2022.3213577