Modeling Security Under Partial Observability

Attack defense games

Summary

Attack Trees (ATs) are a widely adopted formalism for modeling security threats. However, their conventional use relies on an unrealistic assumption of perfect knowledge, where the system's entire state and all adversarial actions are fully known. Real-world security interactions are characterized by limited visibility and finite resource constraints for both the attacker and the defender.

To address this gap, we introduce Supervised Attack Trees (SATs), a novel framework that extends ATs to explicitly model the strategic, resource-constrained interaction between an attacker and a defender under conditions of partial observability. In our SAT model, each agent possesses a distinct, limited view of the system's nodes. The defender (supervisor) can dynamically allocate a finite budget to delay ongoing attacks, while the attacker expends a separate budget to compromise nodes.

We formally define the notion of a consistent observation, which represents a partially visible snapshot of the system state, and provide an algorithm for verifying its validity against the underlying SAT structure. Furthermore, we demonstrate that critical security decision problems, such as determining the minimum budget required to guarantee a successful attack and verifying the existence of a purely observation-based defense strategy that perpetually prevents the root compromise, can be systematically reduced to tractable model-checking problems.

An Attckk Tree

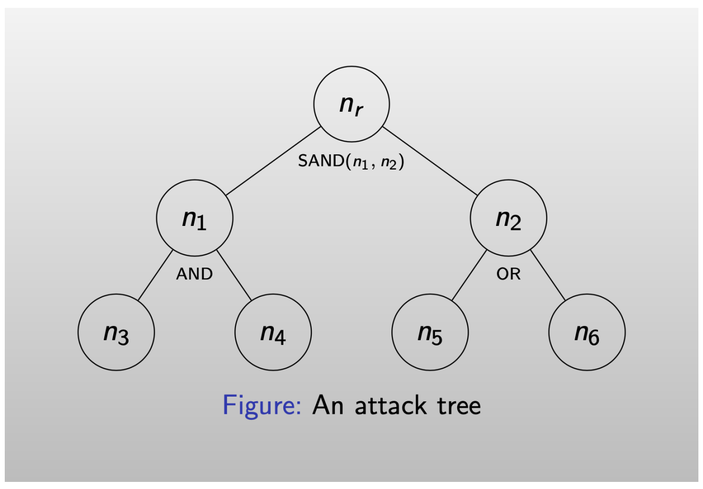

Figure: An attack tree

Attacks:

1. An attack: a set of leaf nodes with time information.
2. A state of an attack at a given time: a set of compromised nodes.
3. The cost of an attack at a given time is calculated from the state.

Extended Attack Trees

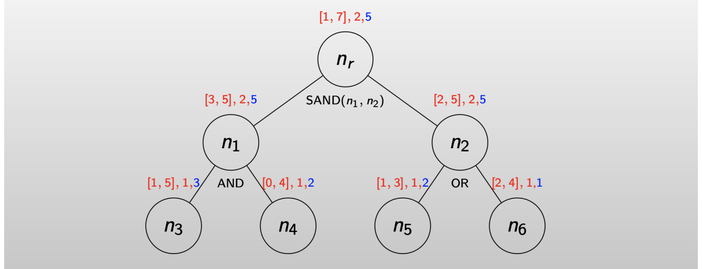

Figure: An example of an attack tree with time and cost constraints on the nodes.

Defense:

1. A defender tries to defend an attack trees.
2. She or he can only influence time when a node is compromised. .
3. A defense has its cost.

Defended Attack Trees

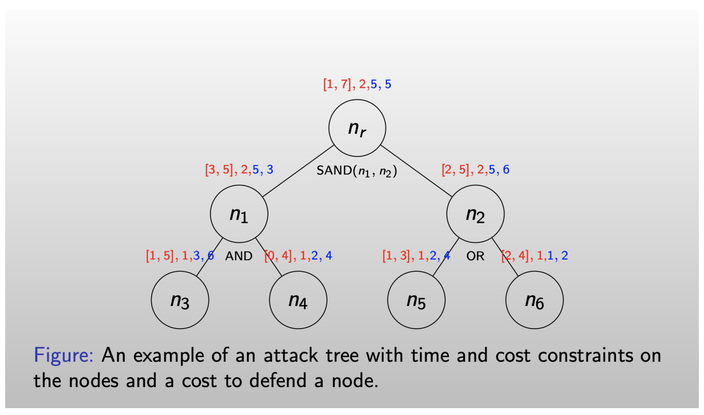

Figure: An example of an attack tree with time and cost constraints on the nodes and a cost to defend a node.

Observations:

1. Both an attacker and defender work with an incomplete information.
2. They can see only a subset of nodes and time.
3. Both have their given budget.

Attack trees and observations

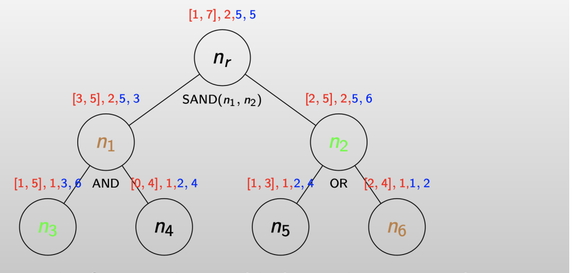

Figure: An example of an attack tree with time and cost constraints on the nodes and a cost to defend a node. Attacker can see green nodes and defender brown ones.

Quetions:

1. Is an observation consistent?
2. For a given observation of an attacker, can it lead to a successful attack? If yes, what to do? Is it within budget?
3. For a given observation of a defender, can a successful attack be avoided? If yes, what to do? Is it within budget?
4. Neural network and explainability by AT?

References

[1] Aliyu Tanko Ali, Damas Gruska, Karam Kharraz, and Martin Leucker.
Analysis of attack time and costs in attack trees via smt resolution.
In *Proceedings of the 8th International Conference on Future Networks & Distributed Systems*, pages 1057–1064, 2024.

[2] Damas Gruska, Aliyu Tanko Ali, and Martin Leucker.
Using attack trees for security education and training: Simplifying threat analysis.
In *IFIP World Conference on Information Security Education*, pages 64–79. Springer, 2025.

Work funded by the EU NextGenerationEU through the Recovery and Resilience Plan for Slovakia under the project No. 09103-03-V04-00095

FAKULTA MATEMATIKY, FYZIKY A INFORMATIKY
Univerzita Komenského v Bratislave

MATFYZ CONNECTIONS

Damas Gruska
Department of Applied Informatics, Faculty of Mathematics, Physics and Informatics, Comenius University

Aliyu Tanko Ali
Institute for Software Engineering and Programming Languages, University of Lübeck, Germany