Intrusion Prevention Through Optimal Stopping

Digital Futures Machine Learning Day

Kim Hammar & Rolf Stadler

kimham@kth.se & *stadler@kth.se*

Division of Network and Systems Engineering
KTH Royal Institute of Technology

Jan 17, 2022

Use Case: Intrusion Prevention

- A **Defender** owns an infrastructure
  - Consists of connected components
  - Components run network services
  - Defender defends the infrastructure by monitoring and active defense
- An **Attacker** seeks to intrude on the infrastructure
  - Has a partial view of the infrastructure
  - Wants to compromise specific components
  - Attacks by reconnaissance, exploitation and pivoting

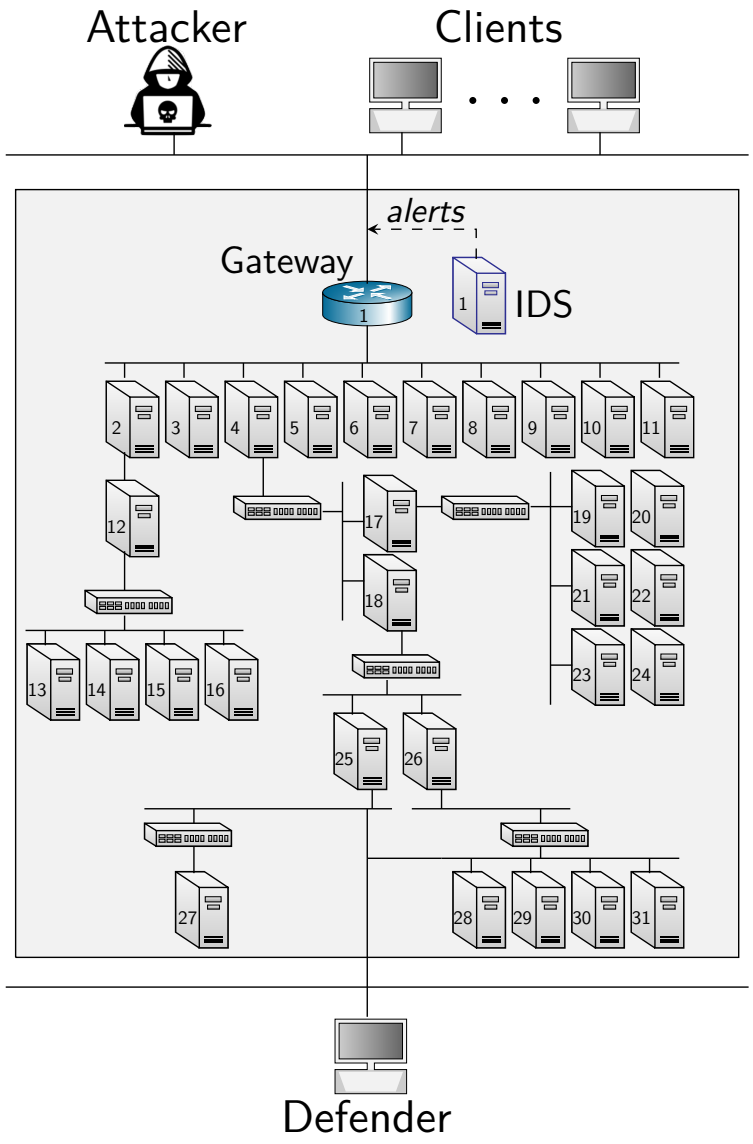

Use Case: Intrusion Prevention

- A **Defender** owns an infrastructure
  - Consists of connected components
  - Components run network services
  - Defender defends the infrastructure

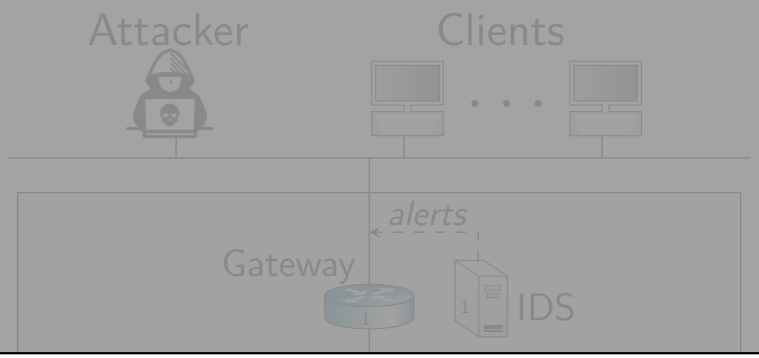

We formulate this use case as an **Optimal Stopping** problem

infrastructure

- Has a partial view of the infrastructure
- Wants to compromise specific components
- Attacks by reconnaissance, exploitation and pivoting

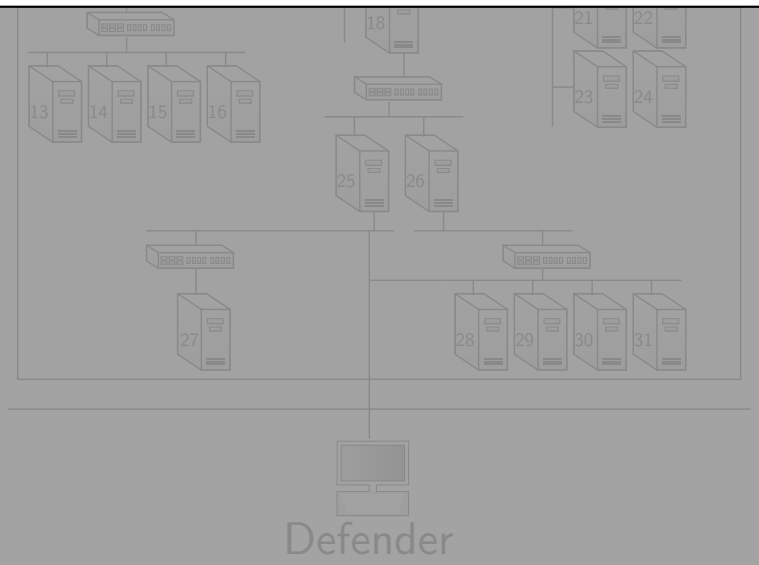

Formulating Intrusion Prevention as a Stopping Problem

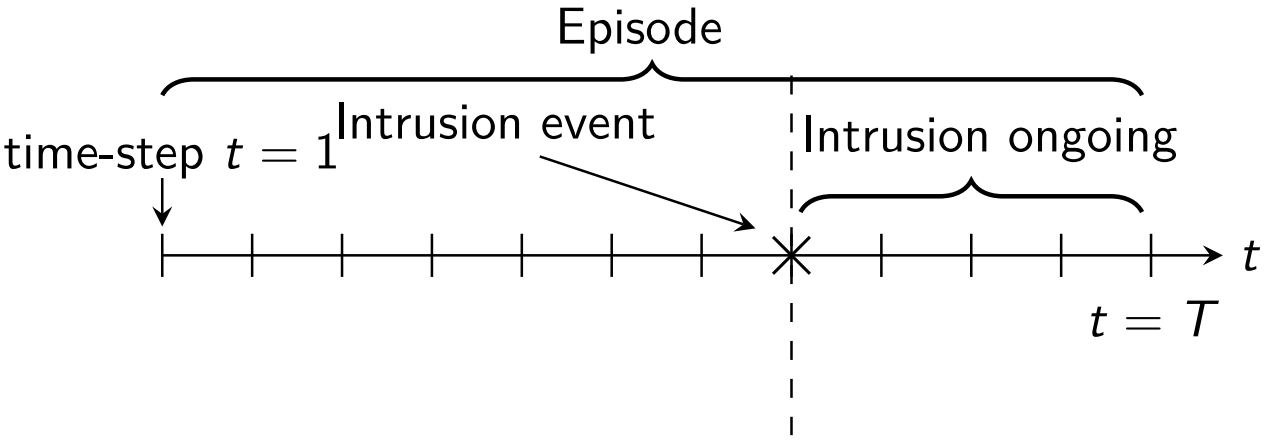

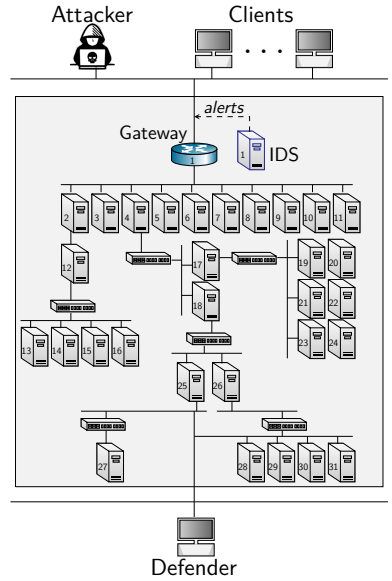

- **Intrusion Prevention as Optimal Stopping Problem**:
  - The system evolves in discrete time-steps.
  - Defender observes the infrastructure (IDS, log files, etc.).
  - An intrusion occurs at an unknown time.
  - The defender can make L stops.
  - Each stop is associated with a defensive action
  - The final stop shuts down the infrastructure.
  - **Based on the observations, when is it optimal to stop?**
  - We formalize this problem with a POMDP

Threshold Properties of the Optimal Defender Policy

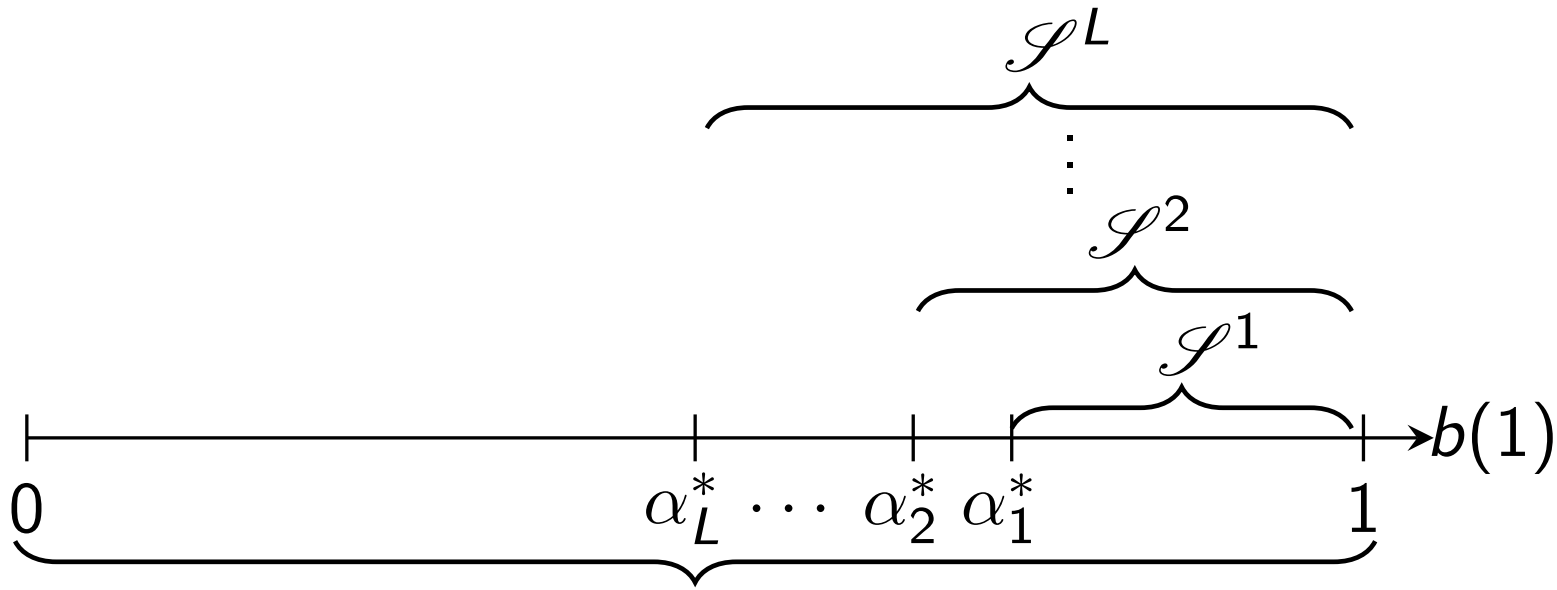

Our Method for Finding Effective Security Strategies

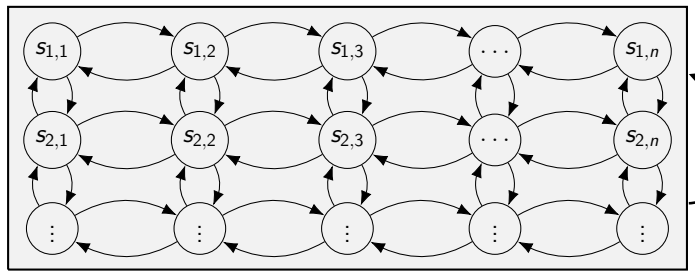

Simulation System

Reinforcement Learning & Generalization

Policy Mapping π

Model Creation & System Identification

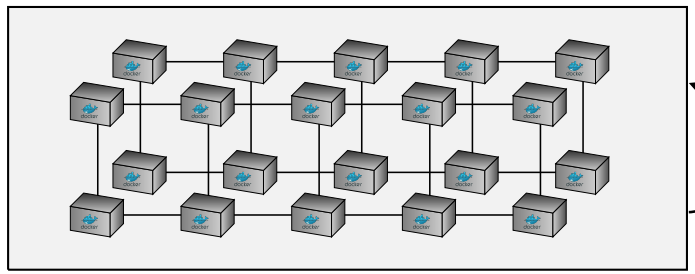

Emulation System

Policy evaluation & Model estimation

Policy Implementation π

Selective Replication

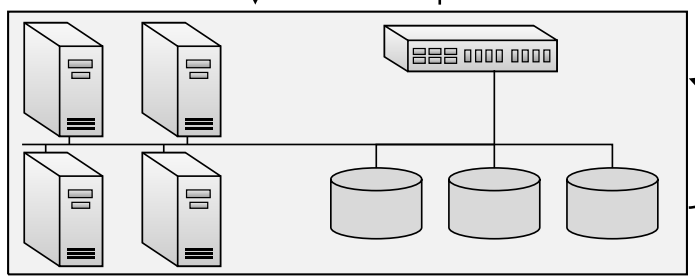

Target Infrastructure

Automation & Self-learning systems

Conclusions

- We develop a *method* to learn automated security prevention policies
  1. emulation system;
  2. system identification;
  3. simulation system;
  4. reinforcement learning
  5. domain randomization and generalization.

- We apply the method to an **intrusion prevention** use case.
  - We formulate intrusion prevention as a multiple stopping problem
  - We model it as a POMDP
  - We apply the stopping theory to establish structural results of the optimal policy