Intrusion Tolerance for Networked Systems through Two-Level Feedback Control

CDIS Spring Conference, 22 May 2024

Kim Hammar (kimham@kth.se) Rolf Stadler (stadler@kth.se)

Center for Cyber Defense and Information Security (CDIS)
KTH Royal Institute of Technology

Contributions

1. We present TOLERANCE, a **control architecture for intrusion-tolerant systems.**
2. We prove properties of the optimal control strategies and design efficient algorithms for computing them.

Two-level Feedback Control

Node controllers $\pi_1, \ldots, \pi_{N_t}$ compute $b_1, \ldots, b_{N_t}$ and make recovery decisions; a **system controller** π manages the replication factor N_t.

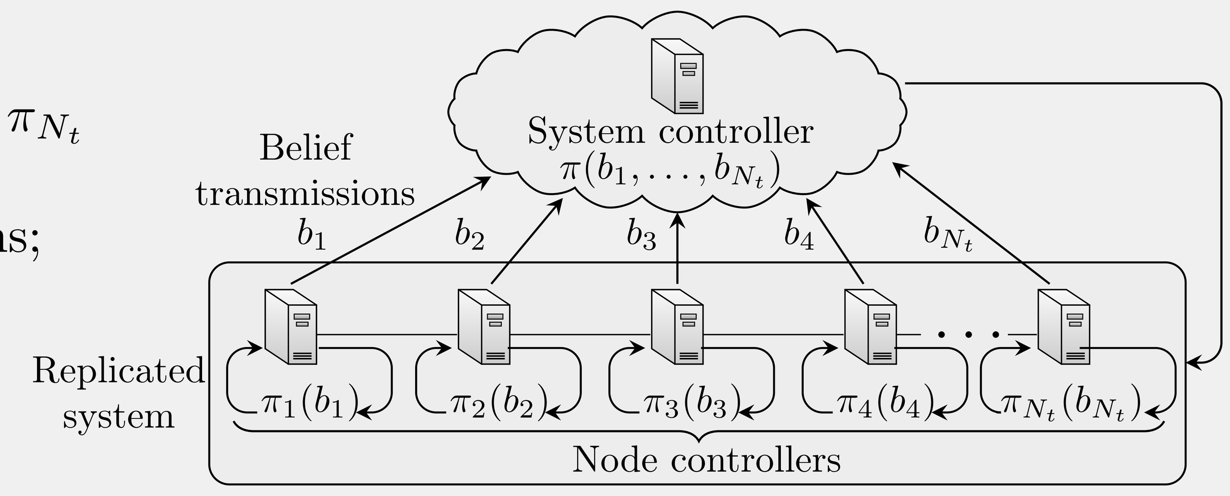

The TOLERANCE Architecture

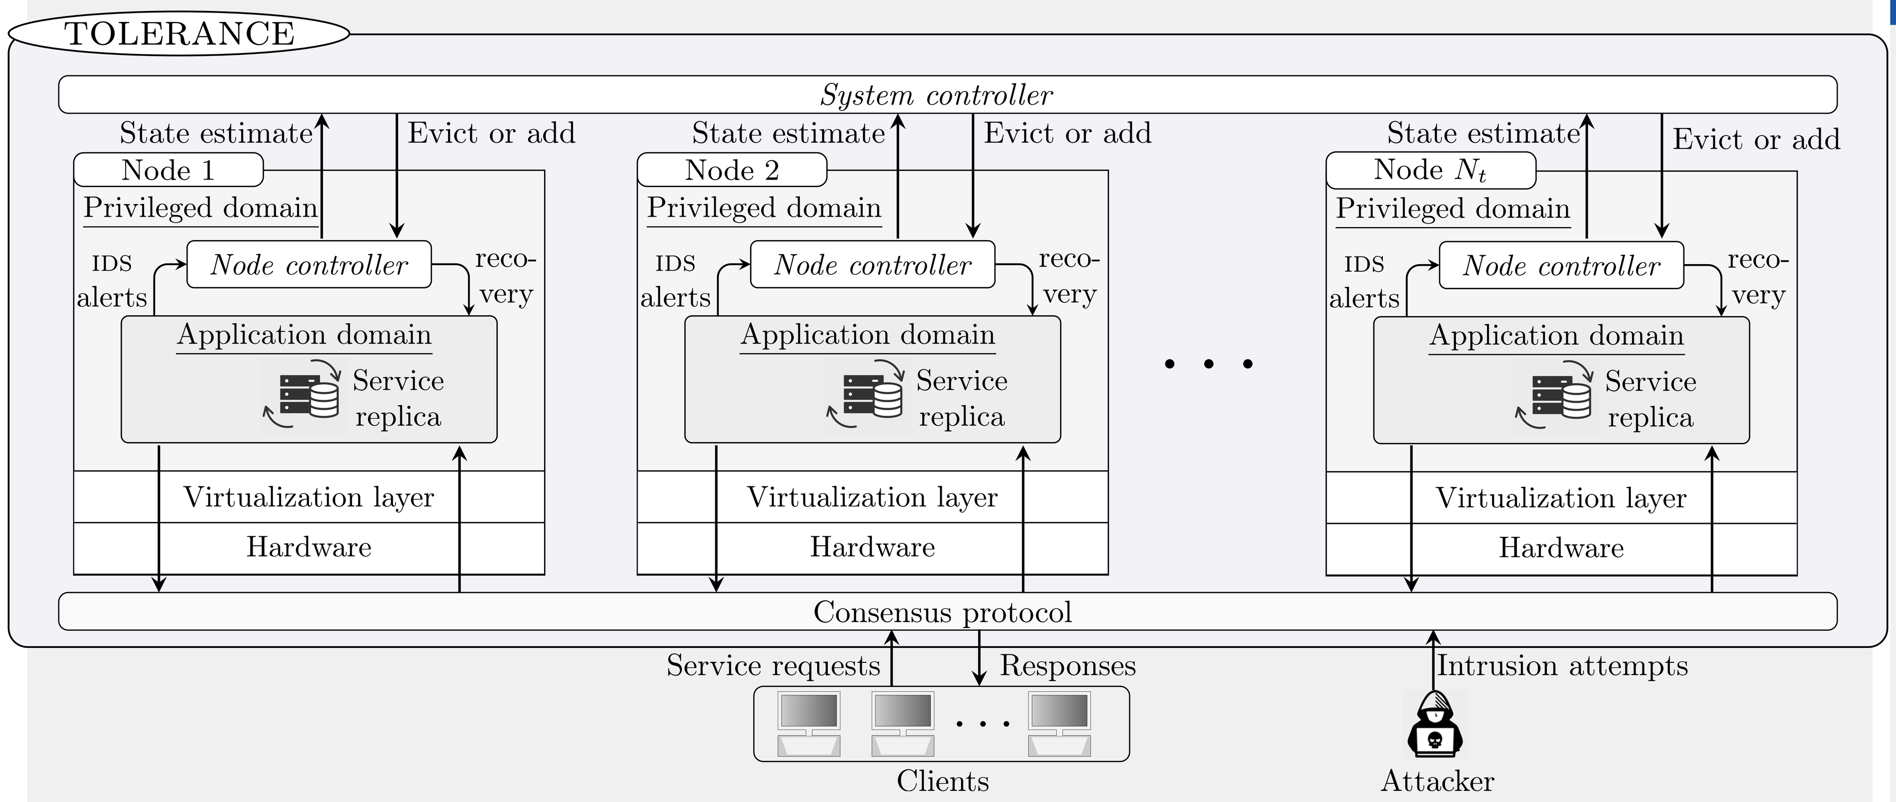

Proposition 1. TOLERANCE provides correct service if the following holds:

(a) An attacker can not forge digital signatures.
(b) Network links are authenticated and reliable.
(c) At most k nodes recover simultaneously and at most f nodes are compromised or crashed simultaneously.
(d) $N_t \geq 2f + 1 + k$ at all times t.
(e) The system is partially synchronous.

Formal Model of Intrusion Tolerance

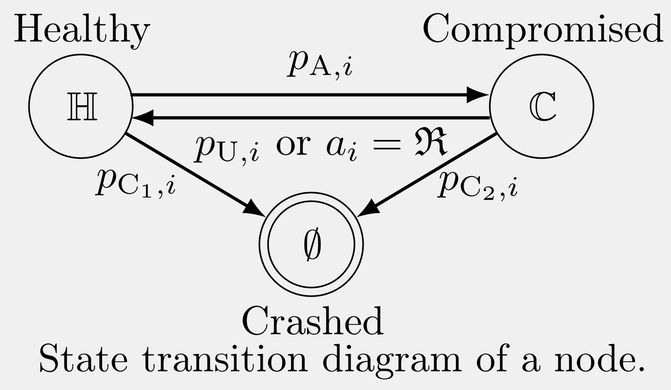

State transition diagram of a node.

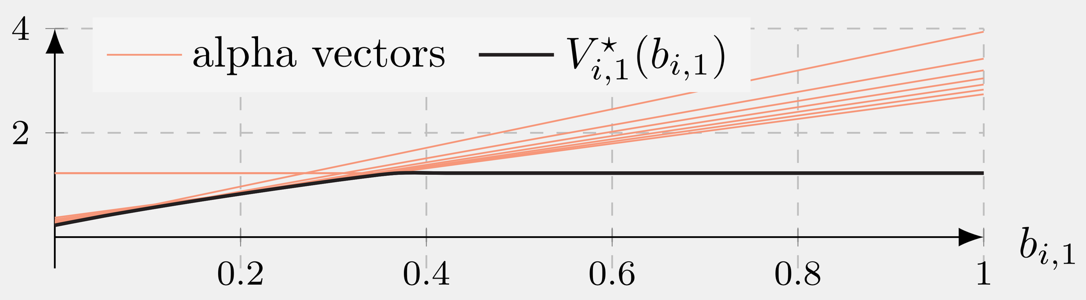

Optimal value function for a node

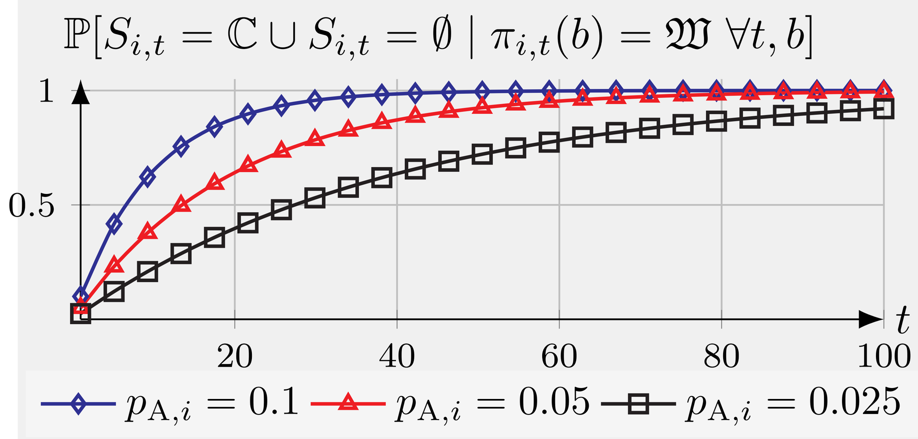

Probability of node compromise or crash

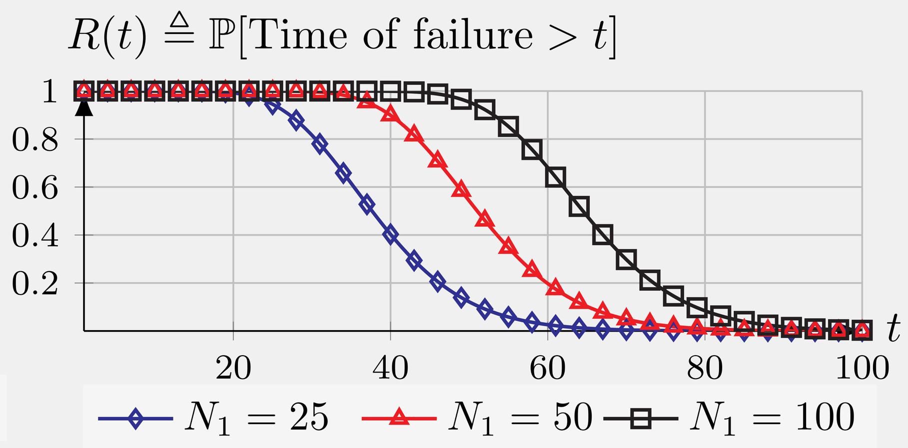

Reliability curves; N_1 = initial # nodes

Structural Results of Optimal Control Strategies

Theorem 1. There exists an optimal **recovery strategy** $\pi^\star_{i,t}$ for each node i that satisfies

$$\pi^\star_{i,t}(b_{i,t}) = \mathfrak{R} \iff b_{i,t} \geq \alpha^\star_{i,t} \qquad \forall t, \tag{1}$$

where $\alpha^\star_{i,t} \in [0,1]$ is a threshold.

Corollary 1. The thresholds satisfy $\alpha^\star_{i,t+1} \geq \alpha^\star_{i,t}$ for $t \in [k\Delta_R, (k+1)\Delta_R]$ and $i \in \mathcal{N}$. As $\Delta_R \to \infty$, all thresholds converge to $\alpha^\star_i$, which is time-independent. (Δ_R is the bounded-time-to-recovery (BTR) constraint.)

Theorem 2.
There exist an optimal **replication strategy** $\pi^\star$ that satisfies

$$\pi^\star(s_t) = \kappa\pi_{\lambda_1}(s_t) + (1-\kappa)\pi_{\lambda_2}(s_t) \qquad \forall t, s_t \in \mathcal{S}_S \tag{2}$$

for some probability $\kappa \in [0,1]$, where λ_1, λ_2 are Lagrange multipliers and $\pi_{\lambda_1}, \pi_{\lambda_2}$ are threshold strategies.

Consequence of the structural results: the optimal control strategies can be computed efficiently.

Comparison to State-of-the-art Intrusion-Tolerant Systems

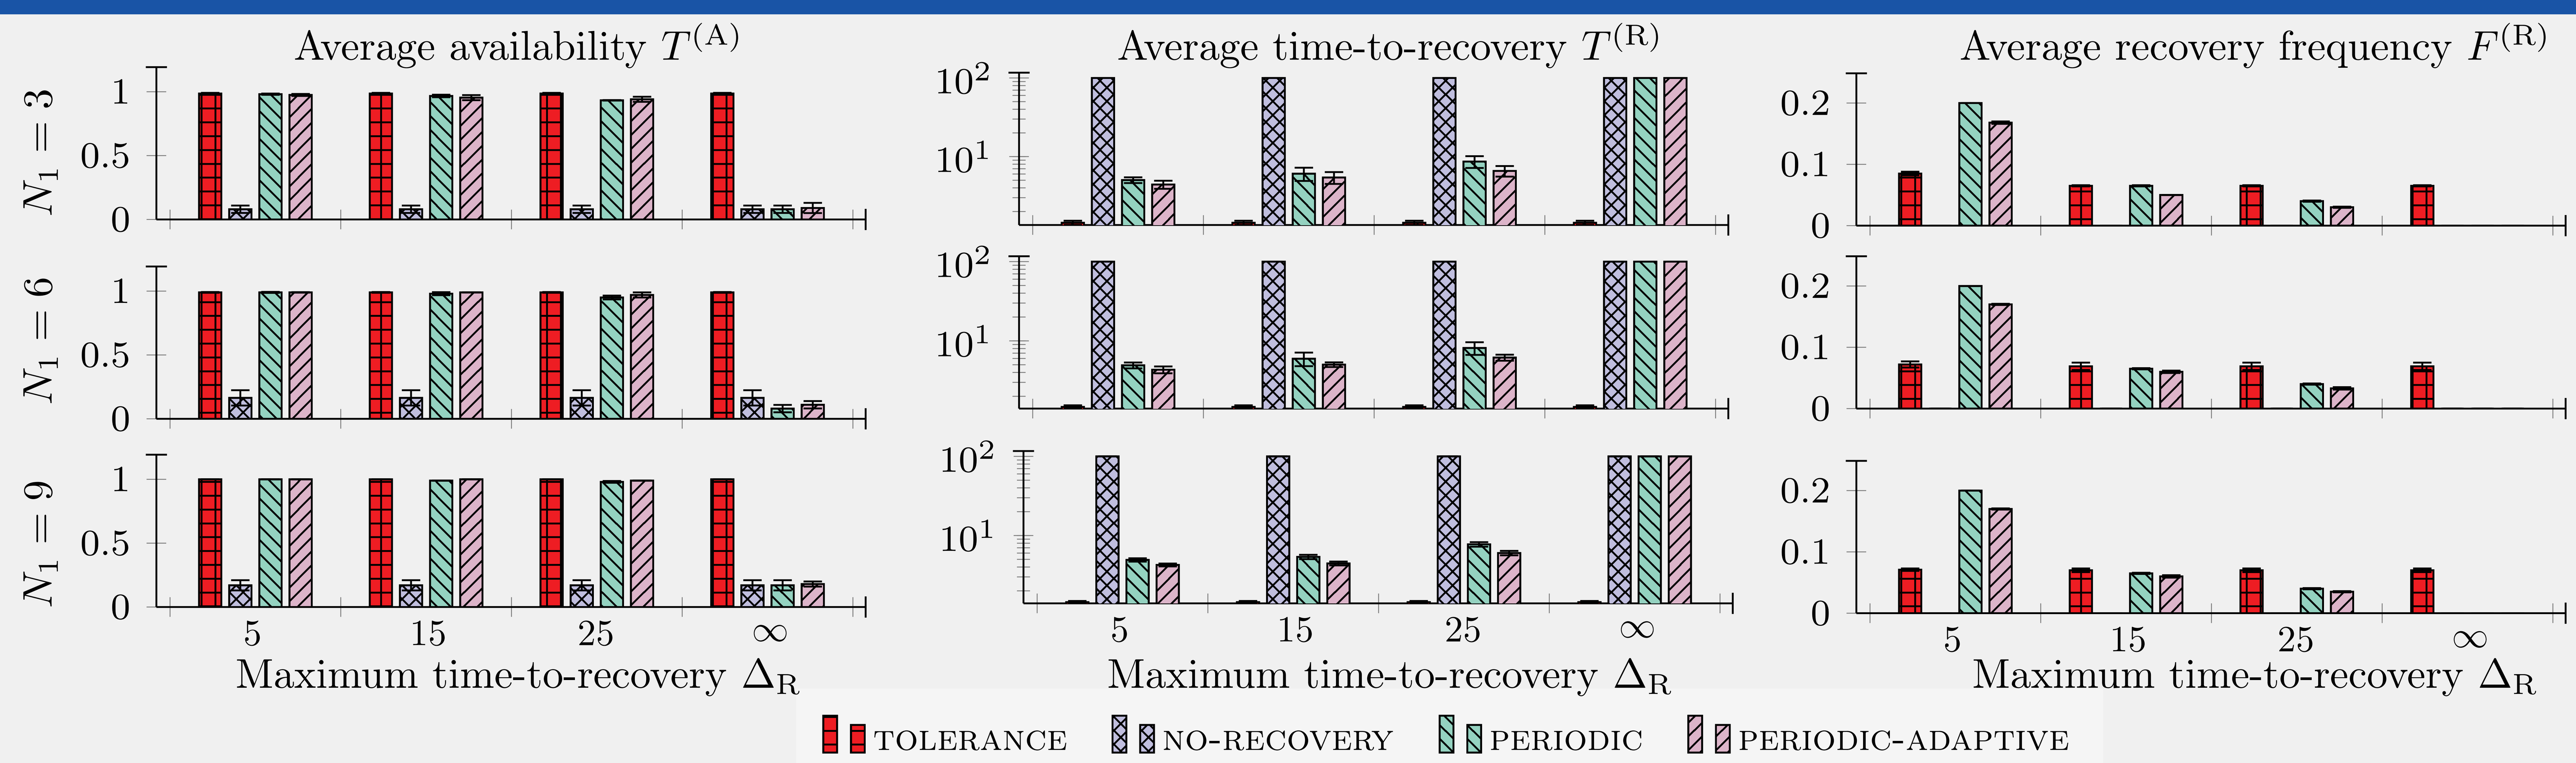

Statistical Intrusion Detection