Securing Dynamic Routing for Parallel Queues against Reliability and Security Failures

Qian Xie[1,2], Li Jin[1,3]

[1]New York University, [2]Cornell University, [3]Shanghai Jiao Tong University.

Introduction

- Network systems rely on data collection and transmission
  - Intelligent transportation systems (ITSs)
  - Manufacturing systems (production lines)
  - Communication networks
- Cyber components susceptible to data loss and data errors
  - E.g., traffic sensors and traffic signals/lights can be intruded and manipulated
  - Need secure-by-design features

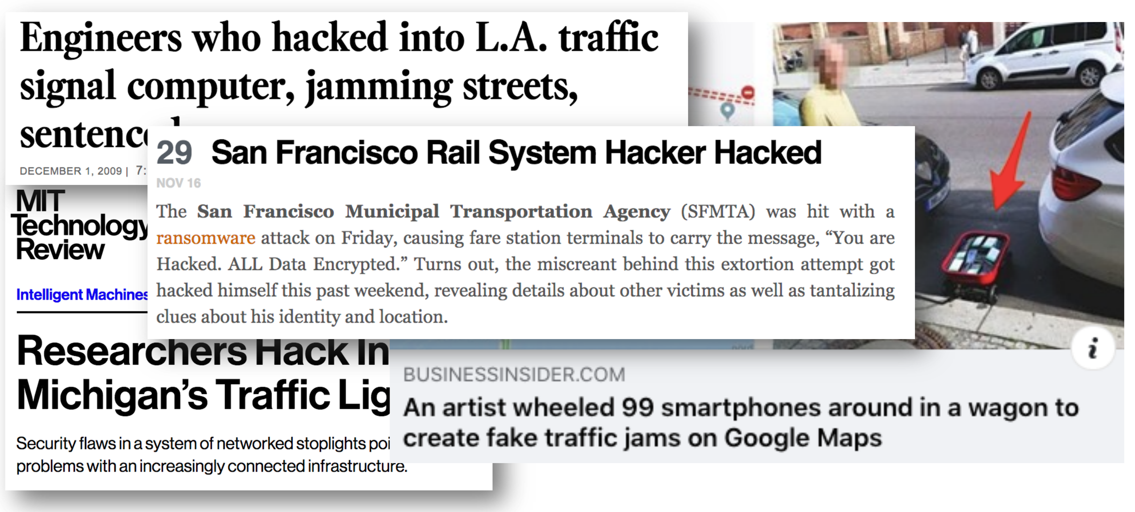

Example: dynamic routing in ITSs

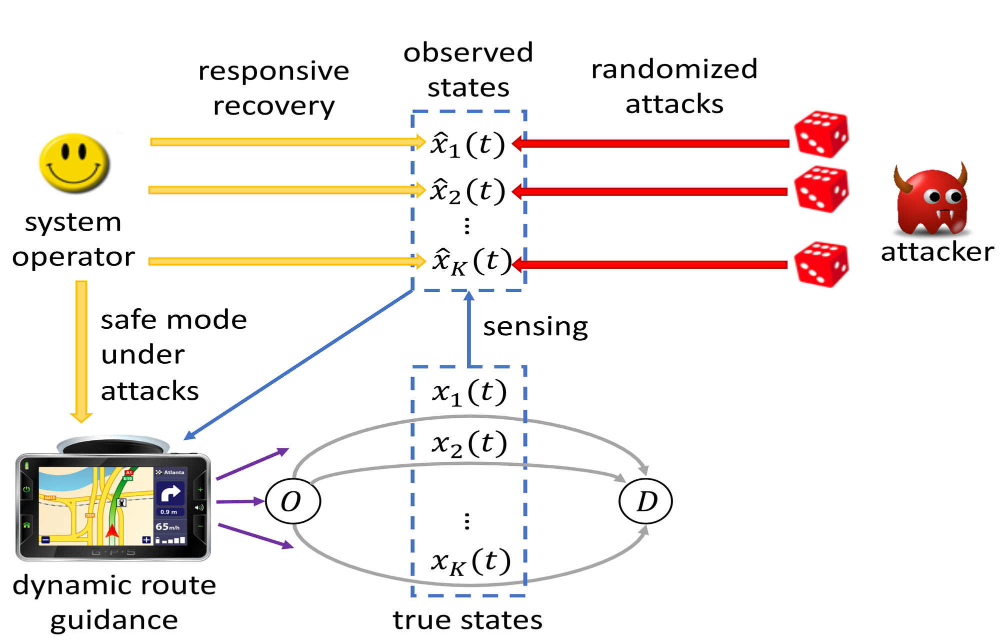

Research questions

Modeling & analysis

- How to model stochastic & recurrent faults/attacks?
- How to quantify attacker's incentive?
- How to quantify the impact due to faults/attacks?
- How to evaluate various security risks?

Resource allocation

- How to allocate limited/costly security resources, including redundant components, diagnosis mechanisms?

Decision making

- How to make protecting (resp. defending) decisions in the face of random faults (resp. malicious attacks)?

Model: Parallel-queueing system

- Poisson arrivals of rate λ
- Parallel servers with service rate μ
- State: vector of queue lengths
- Dynamic routing: dynamically allocat vehicles, components, data packets) to servers
- Provably optimal routing policy: join-the-shortest-queue (JSQ)
- Existing works based on perfect observation of system state and perfect implementation of dynamic routing
- Faulty/failed closed-loop can be worse than open-loop (e.g., round robin or Bernoulli routing)
- Research gap: designing fault-tolerant dynamic routing

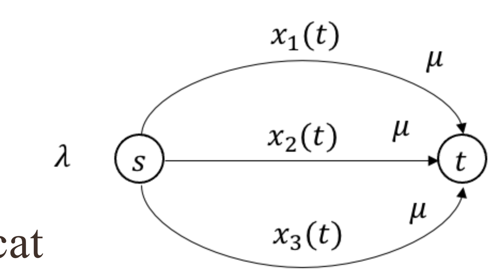

Model: Protection against reliability failures

- Random malfunction: operator fails to send routing instructions
- Denial-of-service: operator loses observation temporarily
- With constant probability α, a job joins a random queue
- Operator protects routing with state-dependent probability $\beta(x)$
- M[illegible] cost

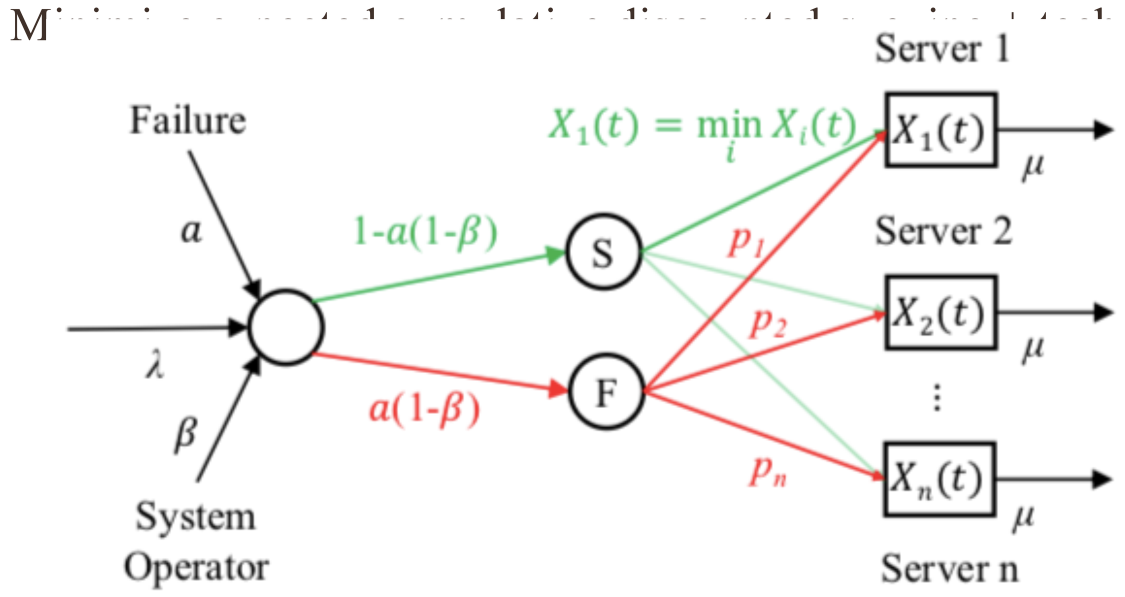

Model: Defense against security failures

- Spoofing: attacker compromises sensing
- Attacker manipulates routing with state-dependent probability $\alpha(x)$ and sends the job to the longest queue
- Operator defends routing with state-dependent probability $\beta(x)$
- Max/minimize expected cumulative discounted reward/loss

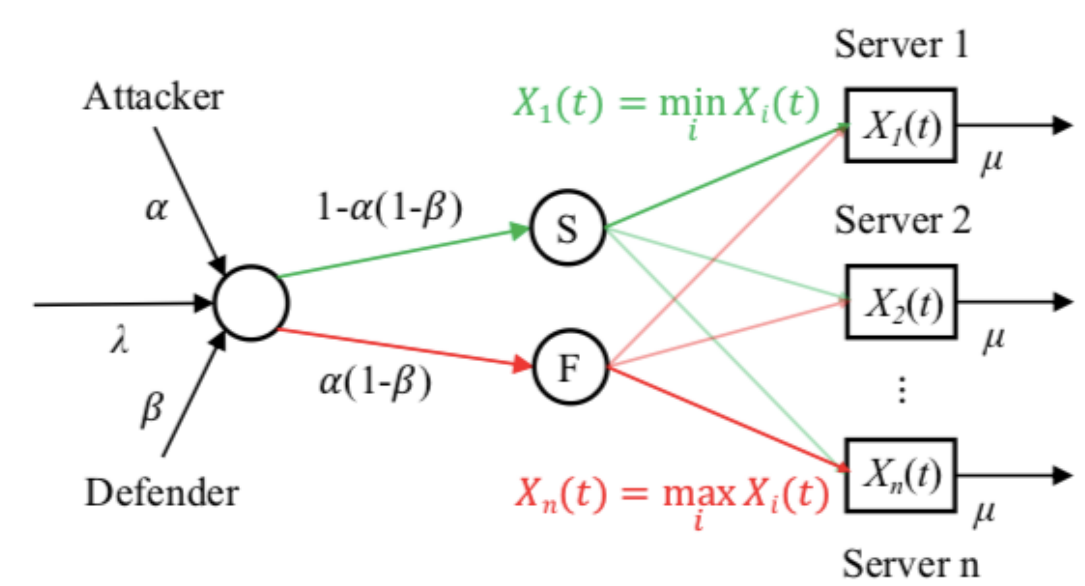

Main results

Theorem 1. The parallel n-queue system with reliability failures is stable if for any non-diagonal vector x,

$$\beta(x) > 1 - \frac{\mu|x| - \lambda x_{min}}{a\lambda(\sum_{i=1}^{n} p_i x_i - x_{min})}.$$

Theorem 2. The parallel n-queue system with security failures is stable if for any non-diagonal vector x,

$$\alpha(x)(1-\beta(x)) < \frac{\mu|x| - \lambda x_{min}}{\lambda \ldots}.$$

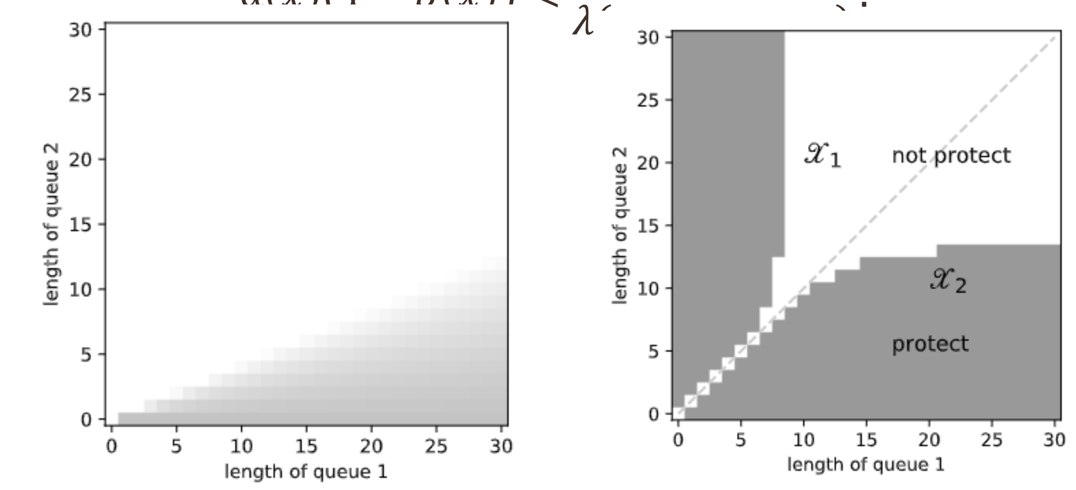

Characterization of the threshold | Characterization of the optimal policy

Markov decision process

Theorem 3. Consider a parallel n-queue system with reliability failures. The optimal protecting policy $\beta^*(x)$ is threshold-based.

- Operator either protects or does not protect (no probabilistic protection), i.e. $\beta^*(x) \in \{0,1\}$;
- Operator is more likely to protect when the queues are 1) less "balanced"; (2) close to empty.

Proof: HJB equation and induction on value iteration.

Attacker-defender stochastic game

Theorem 4. The Markovian perfect equilibrium has the following regimes depending on c_a, c_b and $\delta^*(x) = \lambda(\max_j V^*(x+e_j) - \min_j V^*(x+e_j))$

- $\delta^* < c_a \Rightarrow (0,0)$ (low risk)
- $c_a \le \delta^* < c_b \Rightarrow (1,0)$ (medium risk)
- $\delta^* > \max(c_a, c_b) \Rightarrow (\frac{c_b}{\delta^*}, 1 - \frac{c_a}{\delta^*})$ (high risk)

Equilibrium strategies α^*, β^* are both threshold-based.

Proof: Adapted Shapley's algorithm and induction.

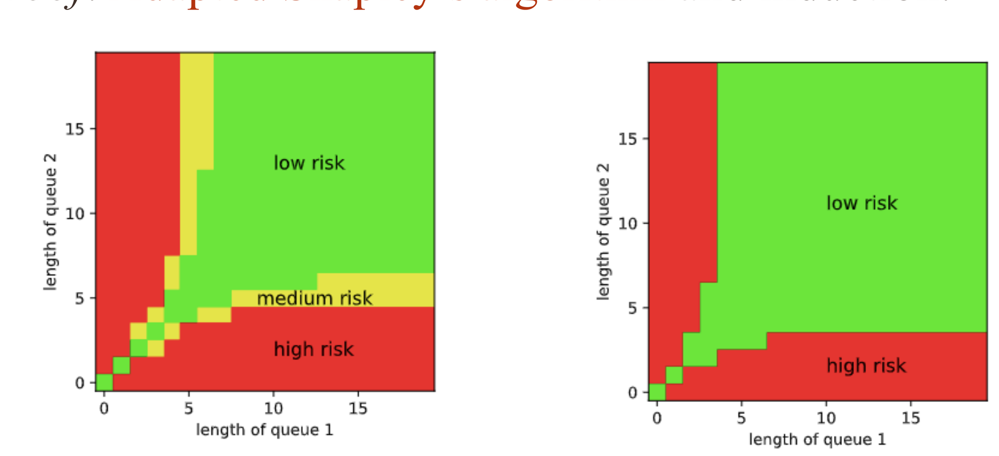

(a) $c_a = 0.1$, $c_b = 0.2$ (b) $c_a = 0.2$, $c_b = 0.1$

Numerical Studies

The incentive to protect is non-decreasing in the failure probability α, non-increasing in the tech cost c_b, and non-decreasing in the throughput λ (estimation of the optimal protecting policy is based on the truncated policy iteration).

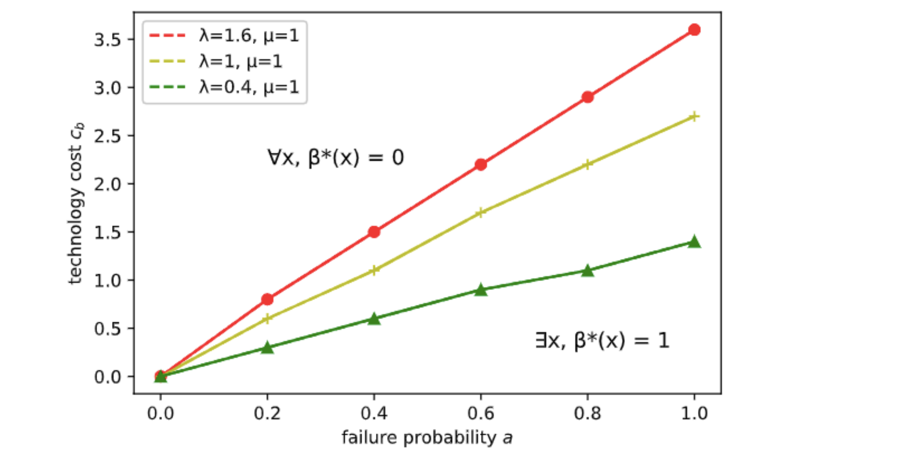

Tipping points of the operator starting to protect

The optimal closed-loop protecting policy β^* performs better in terms of the simulated cumulative discounted cost, compared to the open-loop policies (benchmark) never protect and always protect.

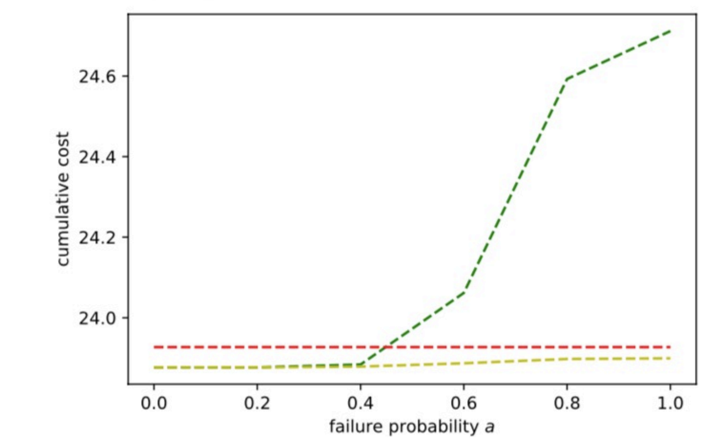

Conclusions

- Without secure dynamic routing, random faults and malicious attacks can destabilize the queueing system
- The optimal protecting strategy and the equilibrium of attacker-defender game have threshold-properties
- System operator has higher incentive to protect when
  - the failure probability is higher
  - the tech cost is lower
  - the throughput is higher
  - the queue lengths are less "balanced"
  - the queues are close to empty
- Our proposed optimal protecting policy (closed-loop) performs better than the benchmark (open-loop)
- Optimal protecting strategy (resp. equilibrium) can be estimated by truncated policy iteration (resp. adapted Shapley's algorithm)

Acknowledgements

The authors appreciate the discussions with Sid Banerjee, Zhengyuan Zhou and Yu Tang. Undergraduate students Jiayi Wang and Dorothy Ng also contributed to this project.