Stackelberg Planning and its Application in Security Analysis

Jörg Hoffmann (Saarland University)
Robert Künnemann (CISPA, Saarland University)

Invited talk SPARK 2018

GEFÖRDERT VOM

Joint work

- joint work with Patrick Speicher, Marcel Steinmetz, Milivoj Simeonovski, Giancarlo Pellegrino, Michael Backes

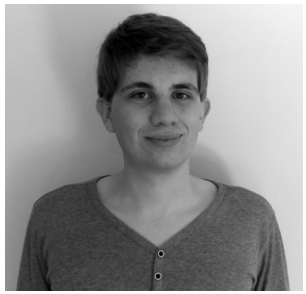


Previously on this Channel

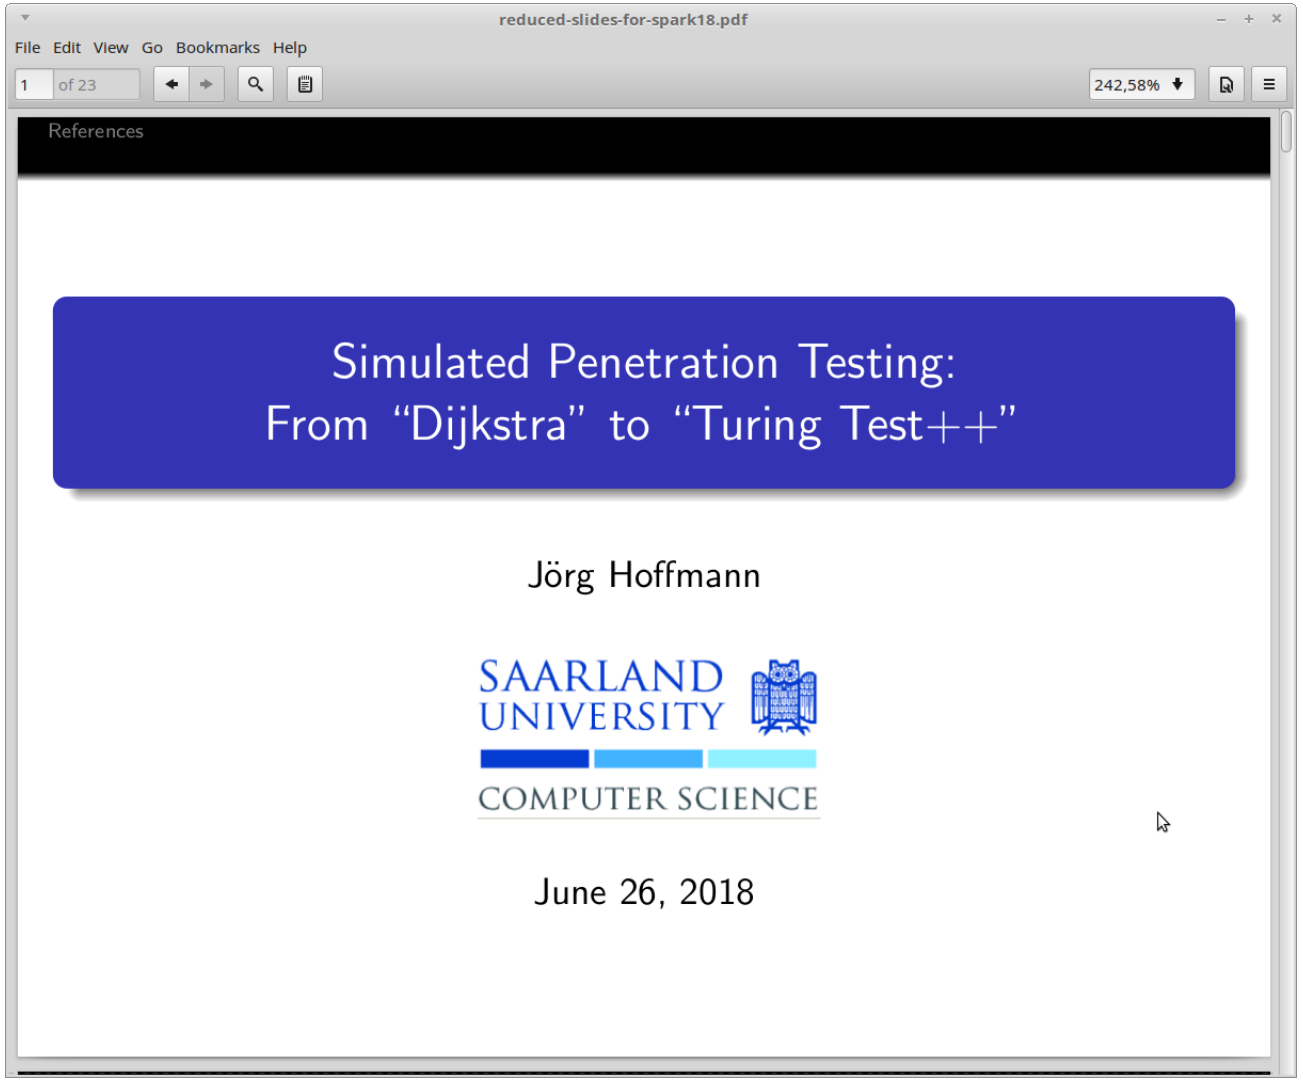

How we got this Idea (outside security testing)

Running Example: Commute-Sabotage

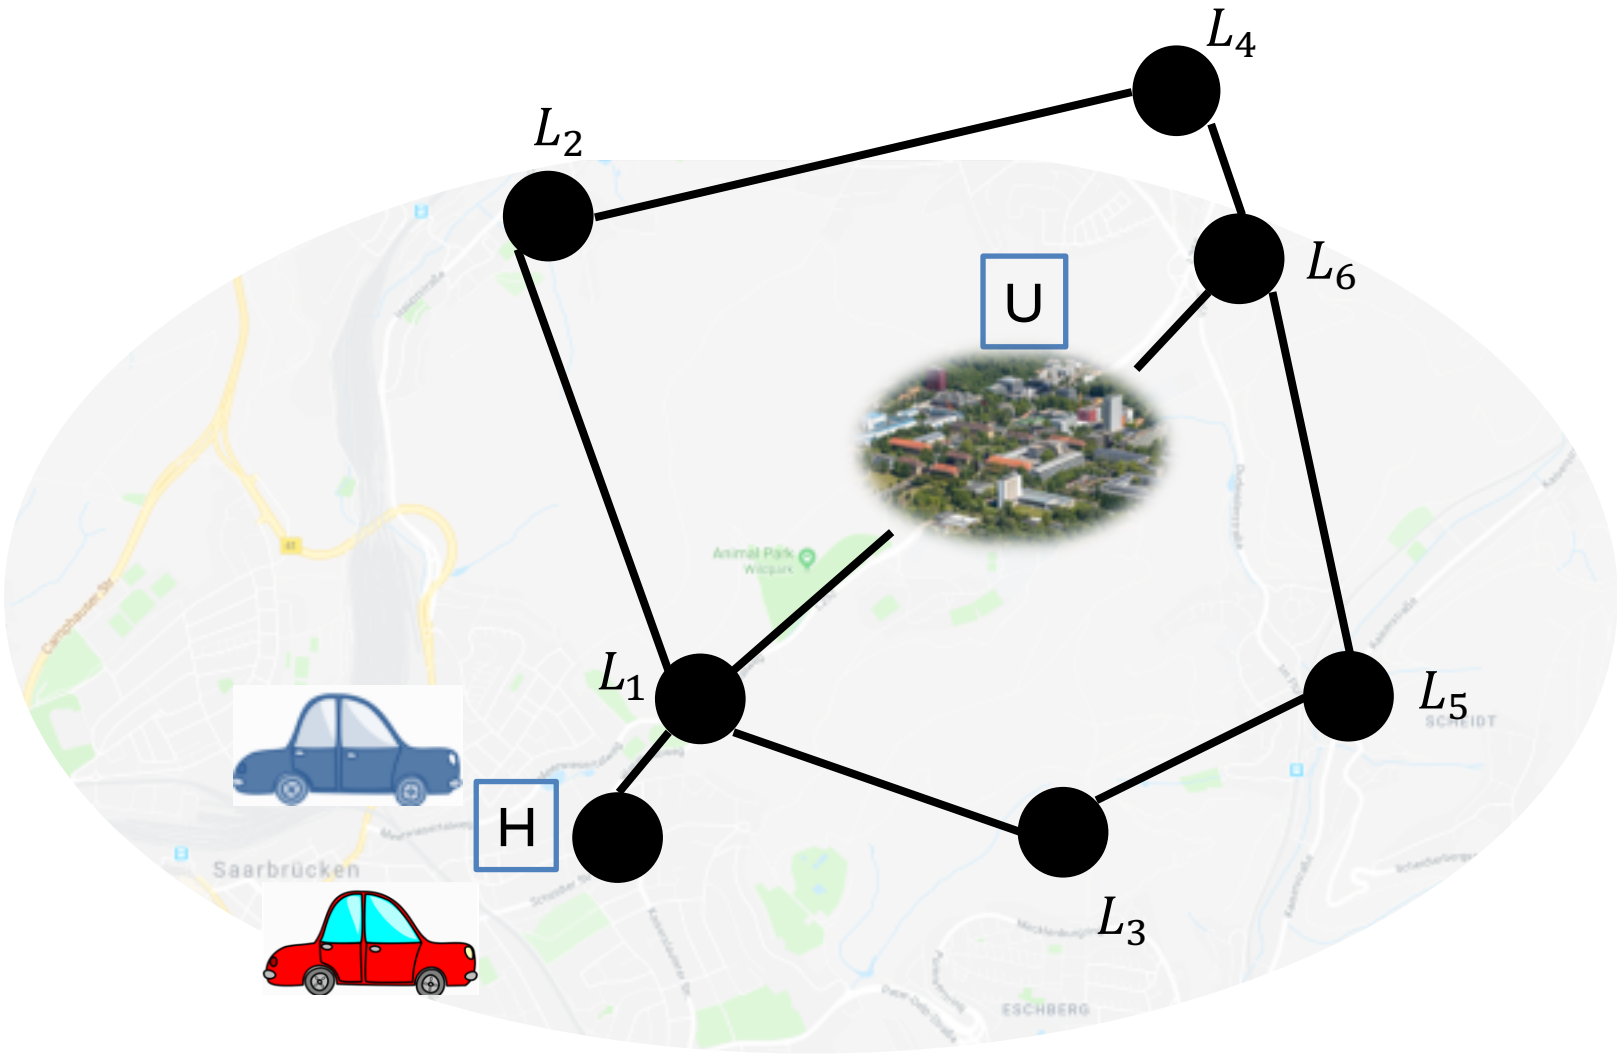

Stackelberg Planning

- **Players:** LEADER (saboteur/network defender) vs. FOLLOWER (me/hacker)
  - Both modelled as full-scale classical planning agents
  - Objectives: Follower has goal; leader maximizes follower's plan cost
- **vs. Planning based games:**
  - Restriction to Stackelberg setting → interesting special case
  - Practically relevant, more feasible algorithmically
- **vs. Security games:** (Tambe et. al., Durkota et. al. in pentesting)
  - Full-scale classical planning models
  - Leader strategy is pure (deterministic) and fully observable to the follower
  - Practically relevant complementary framework

Stackelberg Planning Task

$$\Pi = (\mathcal{P}, \mathcal{A}^L, \mathcal{A}^F, \mathcal{I}, \mathcal{G}^F)$$

- **Syntax:**
  - Predicates, leader actions, follower actions, initial state, follower goal
  - Leader minimal cost to state $L^*(s)$
  - Follower best response $F^*(s)$ (may be ∞)
- **Semantics:**
  - Trade-off between minimizing L^* and maximizing F^*
  - (L, F) dominates (L', F') if $L \leq L'$ and $F \geq F'$, at least one strict
  - S^* set of non-dominated states: the pareto frontier

Leader-follower search sketch

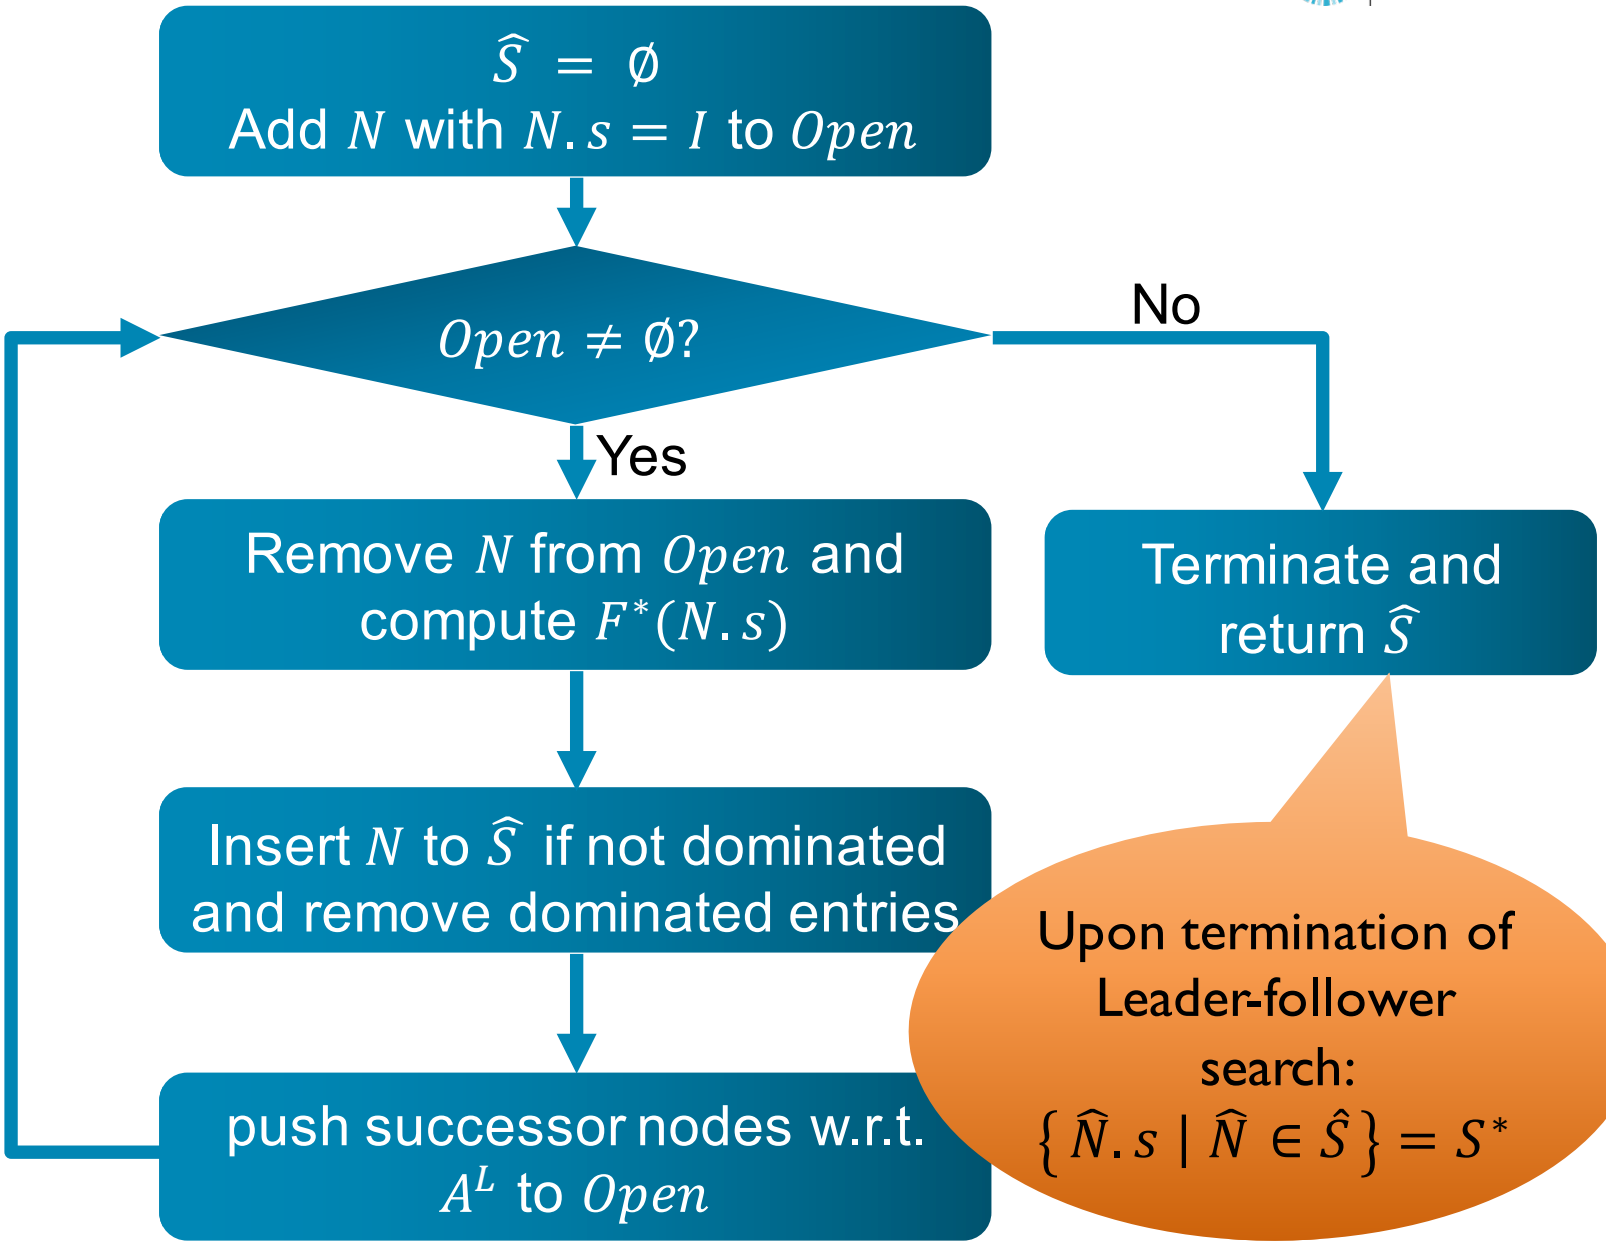

Running Example: Commute-Sabotage

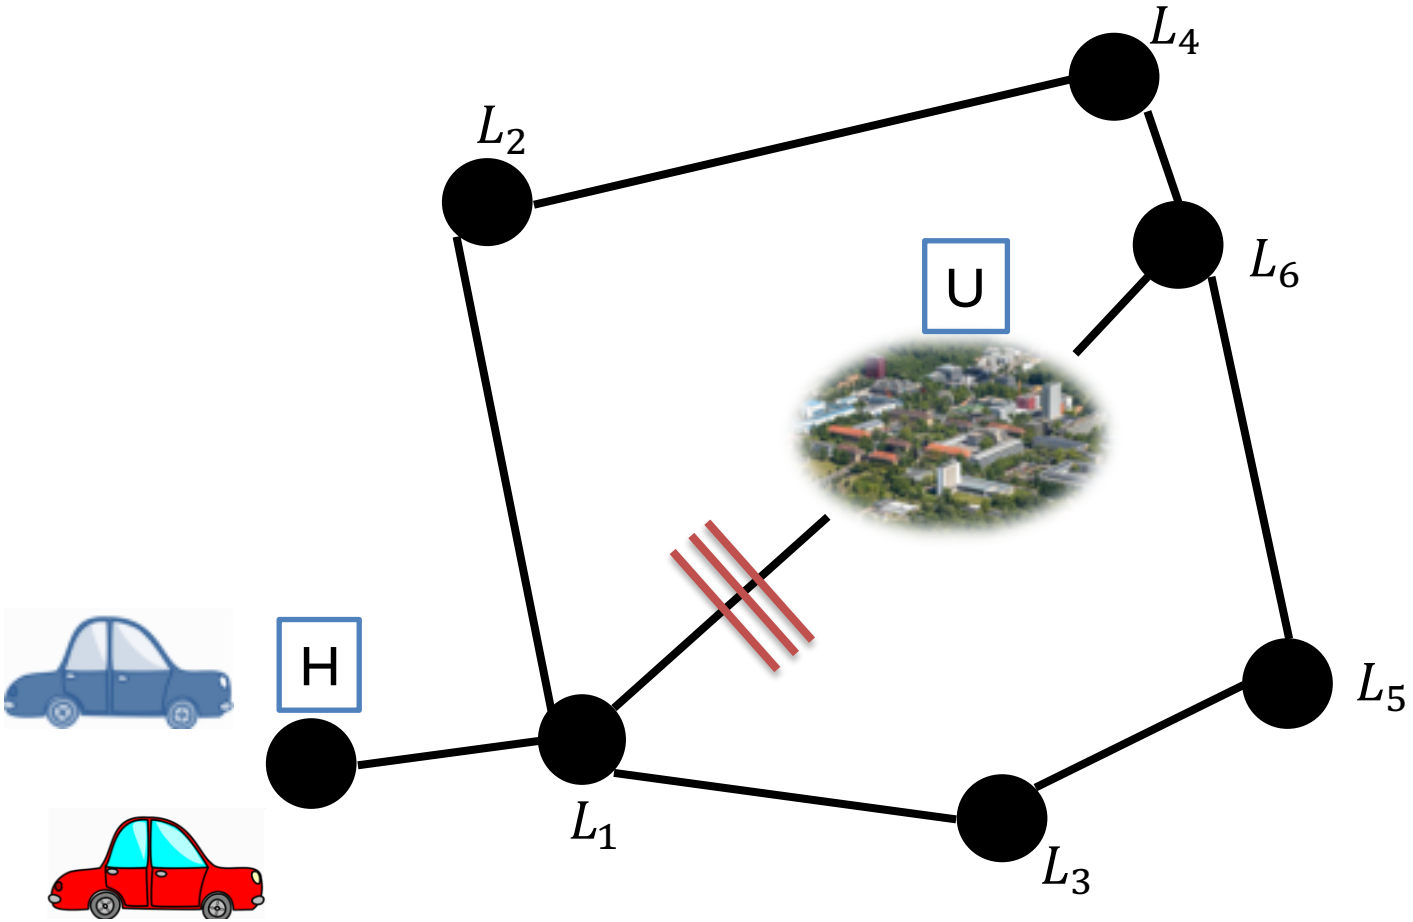

s = {leader-at(L_1), follower-at(L_4), blocked(L_1, U)}

Experiments on IPC Benchmarks

- **Configurations:**
  - With vs. without pruning techniques (omitted here)
  - Optimal vs. satisficing planning for follower
- **Domains:**
  - NoMystery, Logistics
  - Network pentesting
  - Visitall, Sokoban
- **Domain modification:** introduce leader actions!
  - Sabotage road network
  - Block positions in Sokoban
  - Defender (admin) securing parts of network

Coverage: NoMystery Optimal-Follower

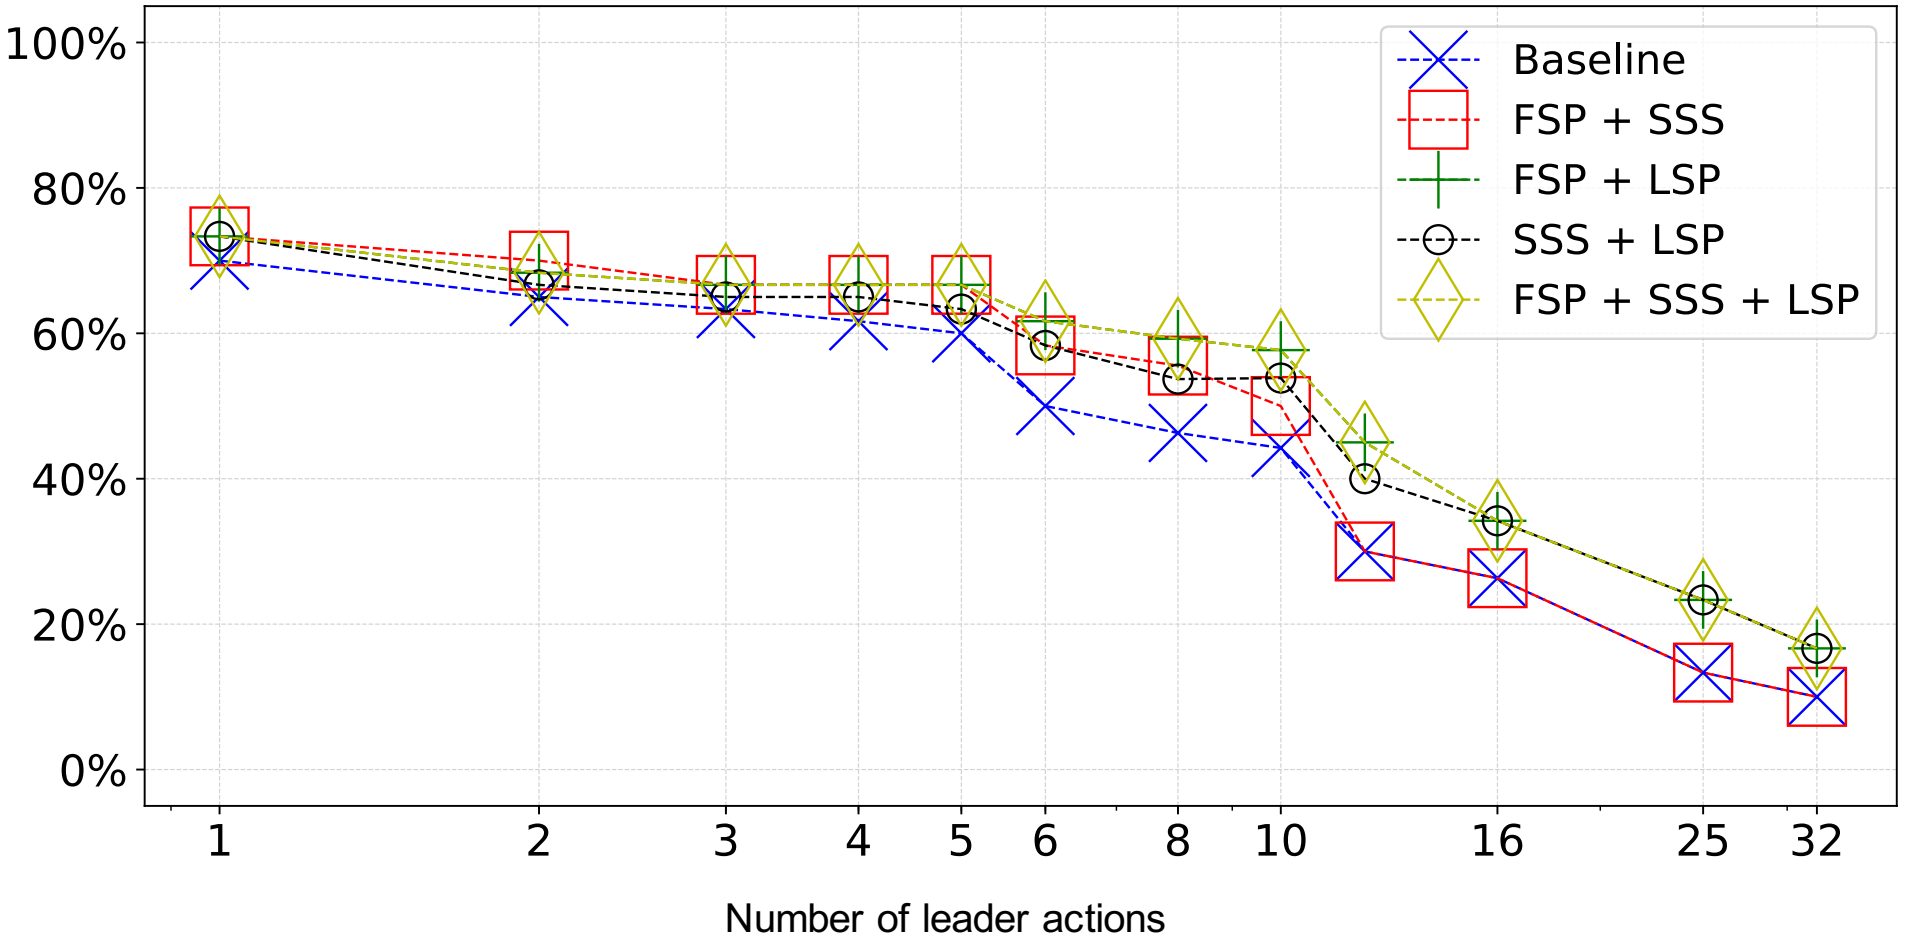

Coverage: Logistics

Optimal follower

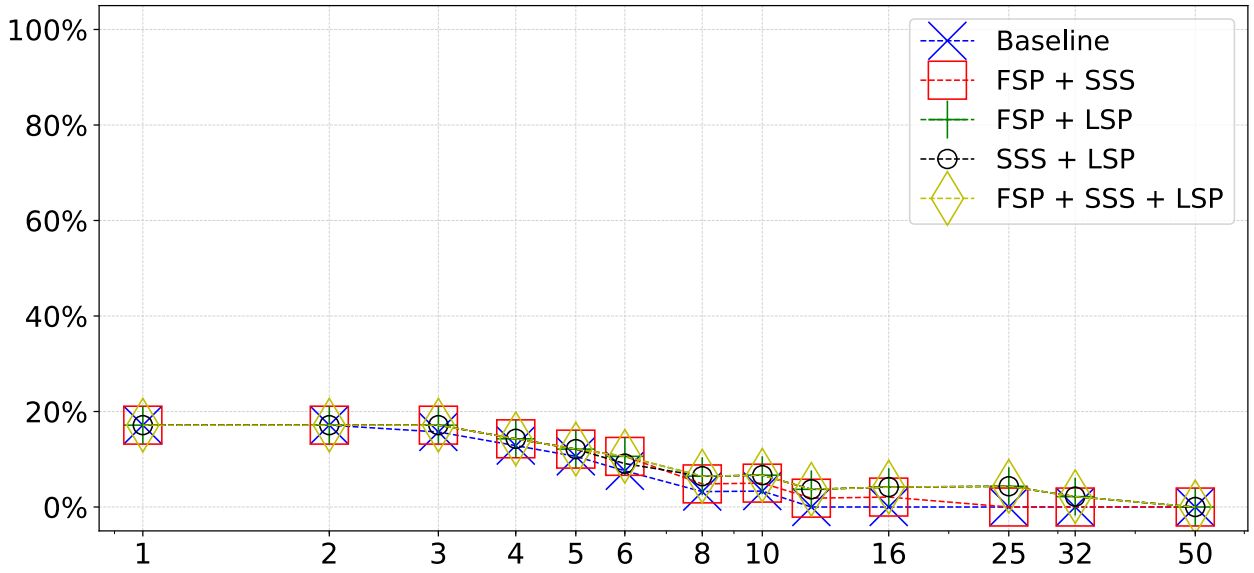

Satisficing follower

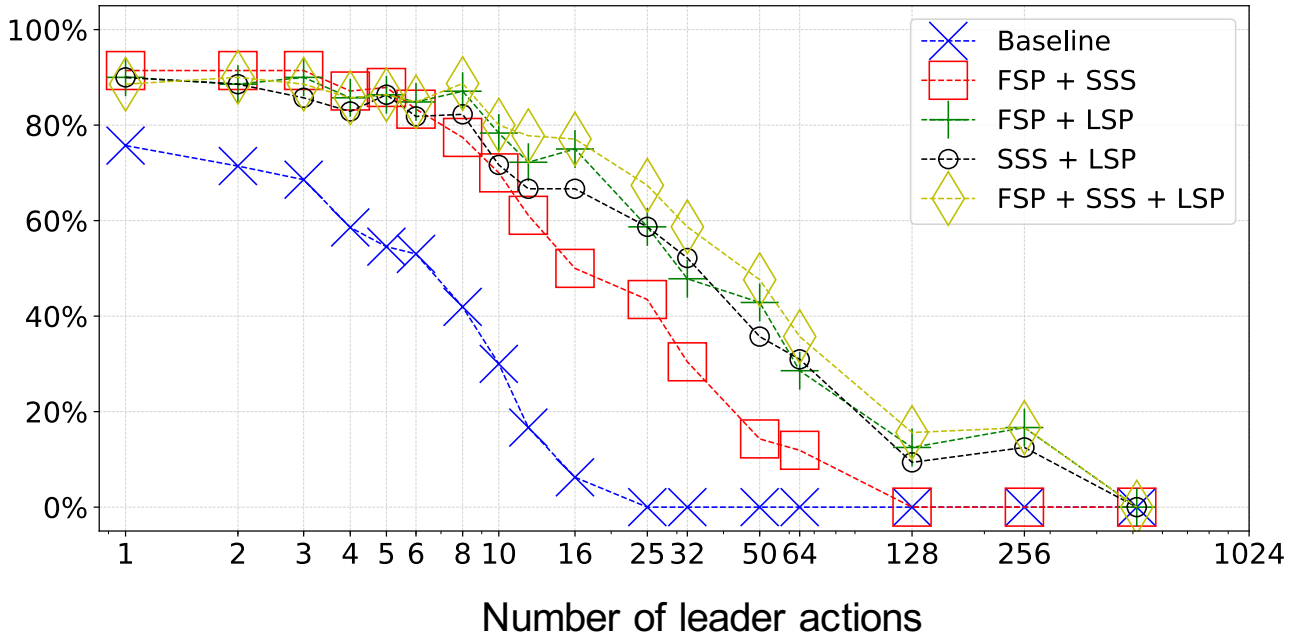

Coverage: Network Pentesting

Optimal follower

Satisficing follower

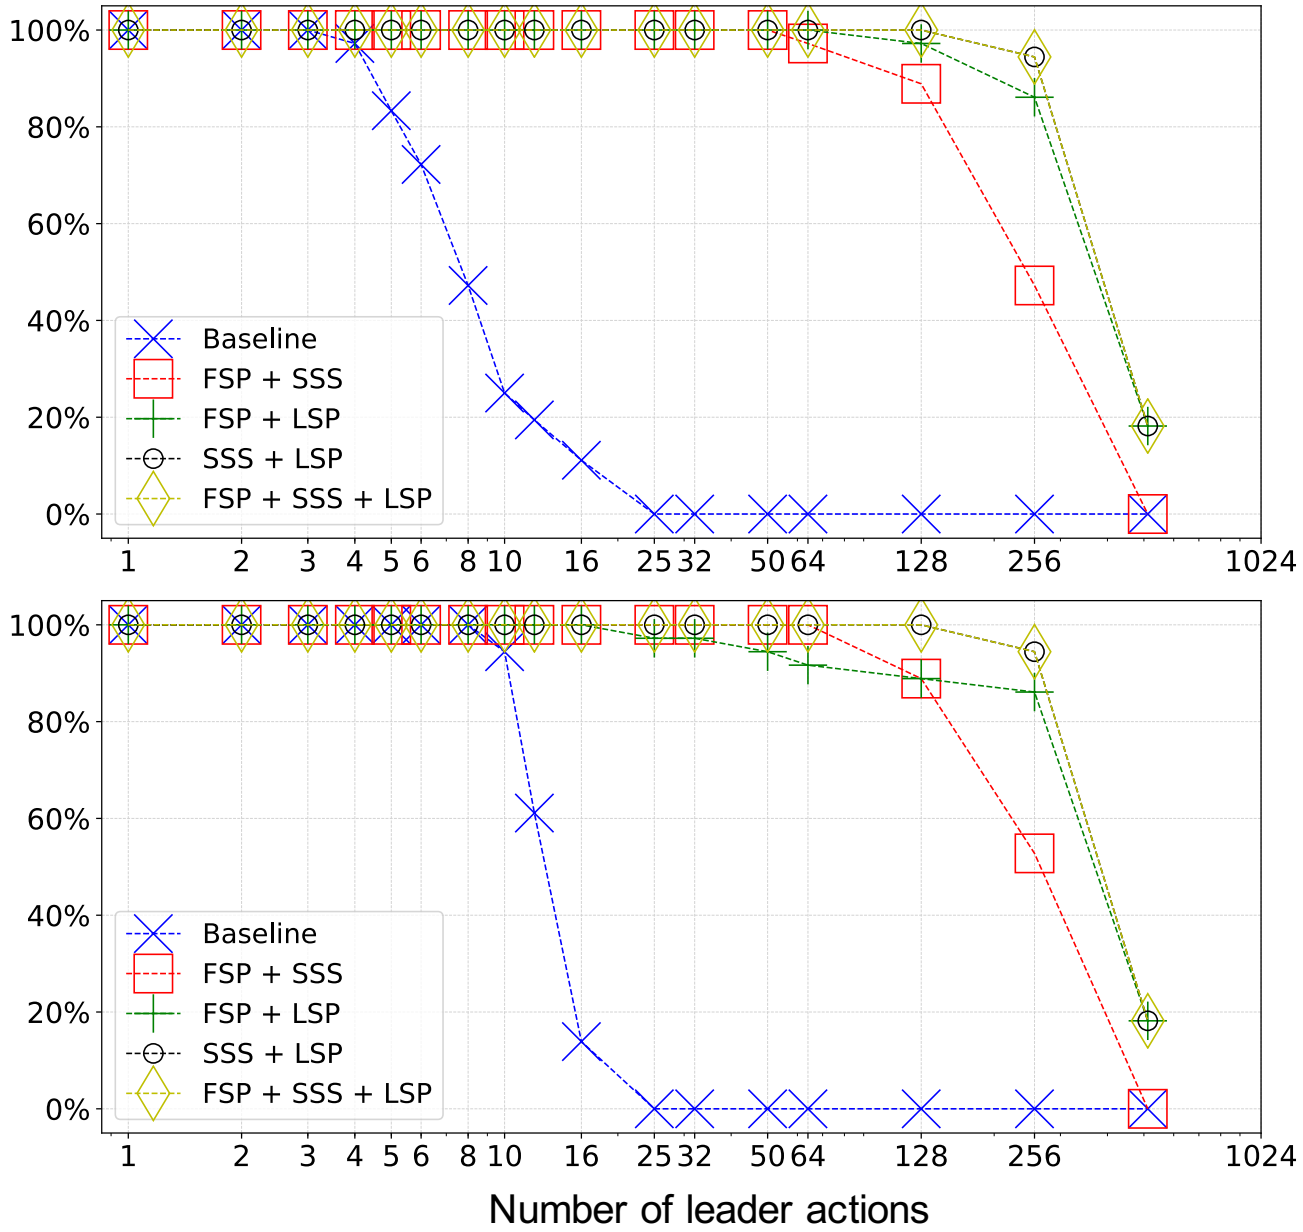

NoMystery Pareto Frontiers

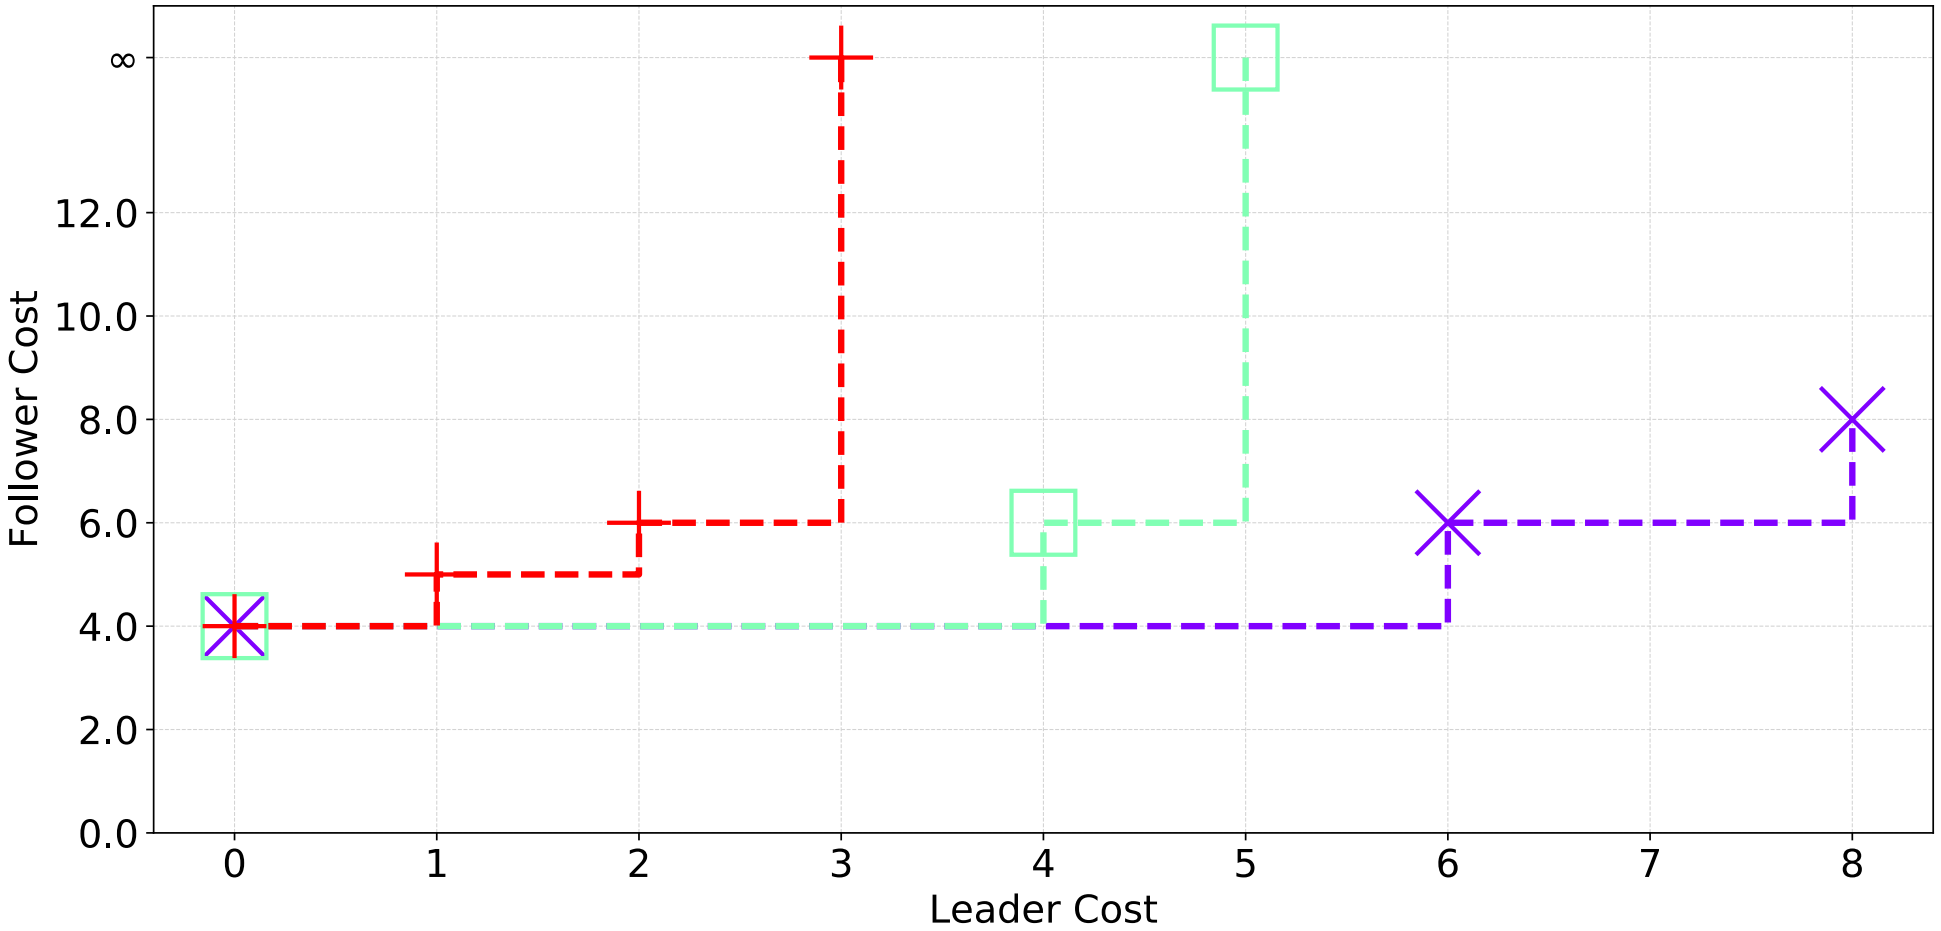

Conclusion of this Part

- **Stackelberg planning:**
  - Middle ground between classical vs game-theoretic planning
  - Applications: robustness measure, defensive planning
- **Algorithms research:**
  - Gap to classing planning techniques not as large → adapt algorithms!
  - *E.g. Stackelberg heuristic functions* combining optimistic estimation of leader costs with pessimistic estimation of follower costs
- **Application to security analyses beyond pentesting as we know it:**
  - Robert will tell us all about it now!

Cost-risk analysis in information security: gospel

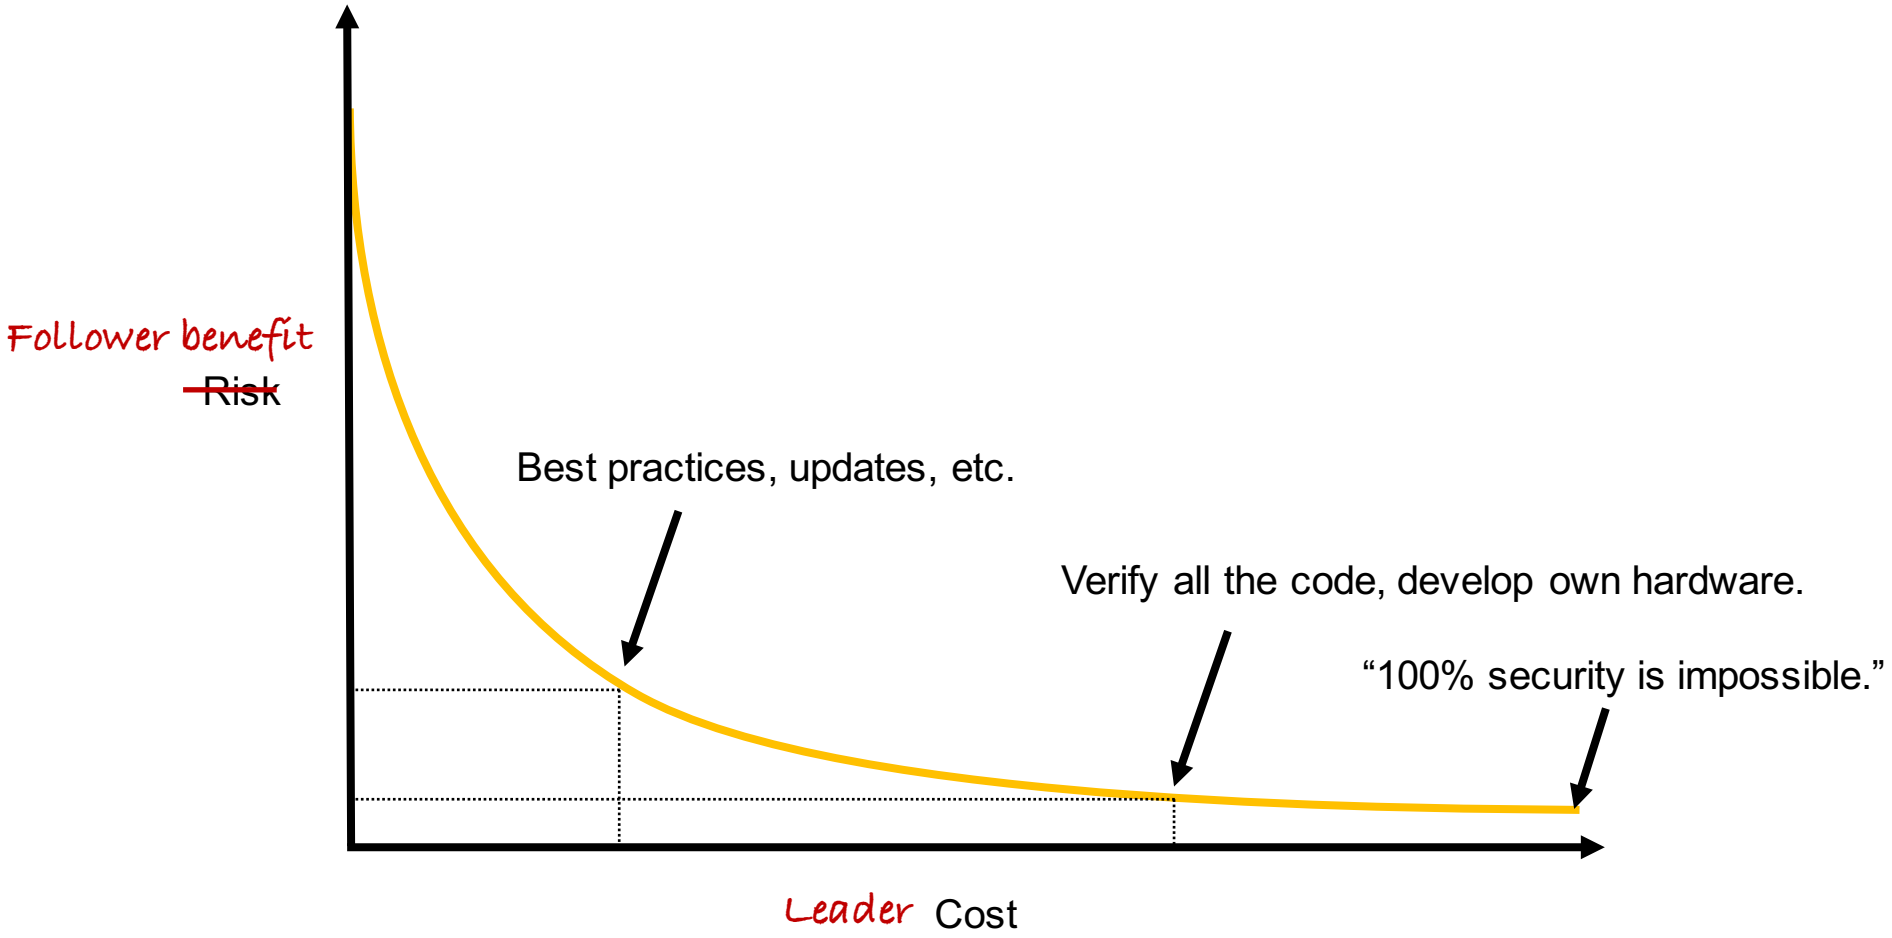

Two main lines of research

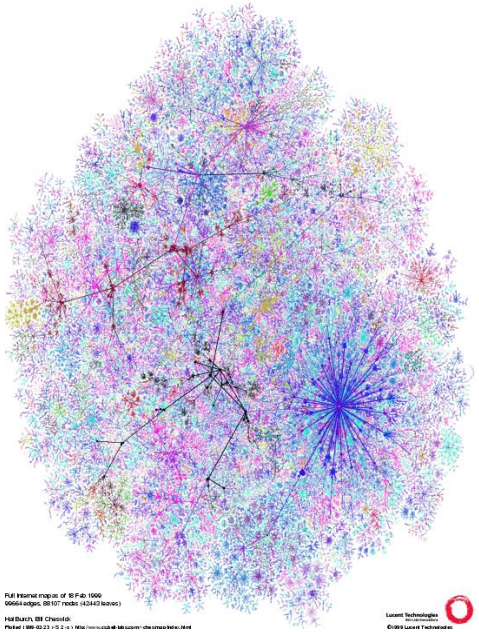

Securing the internet infrastructure:

- Evaluate competing security proposals using “deployment plan”
- Inherently dependent on current infrastructure
- holistic analysis

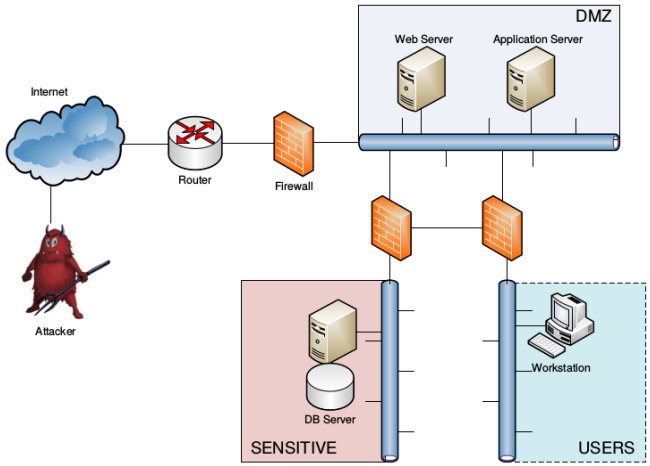

Securing a given network

- Network scanning exposes vulnerabilities
- Develop mitigation plan
- Inherently dependent on software configurations and connectivity

Securing the email infrastructure

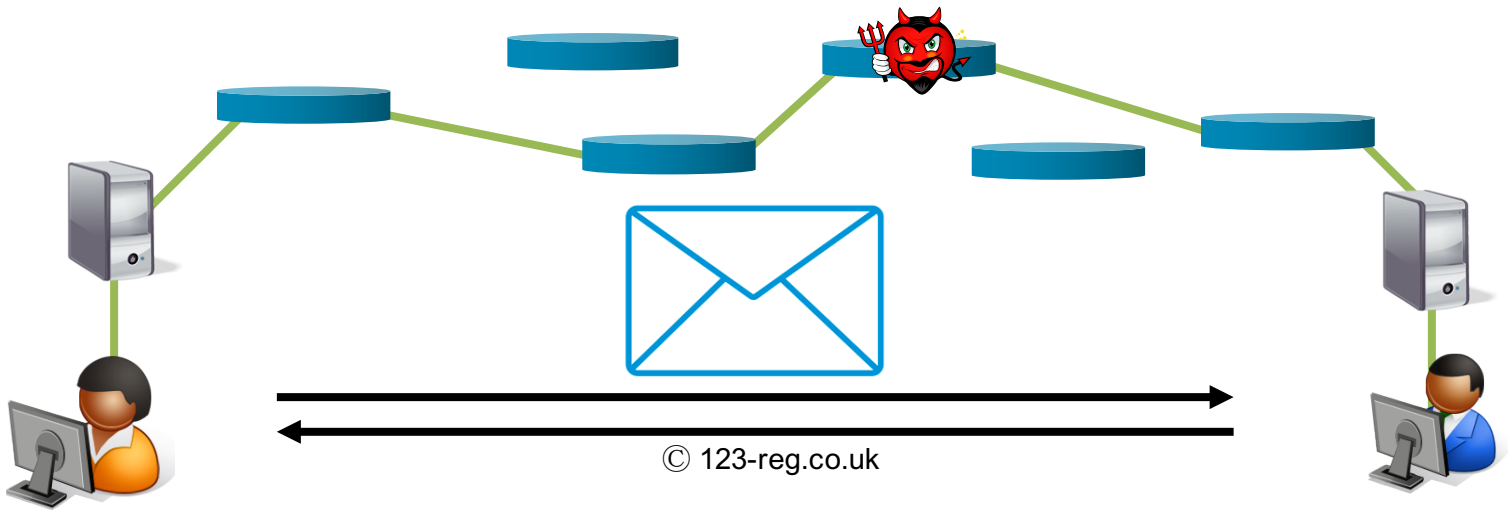

- Plenty of proposals, but which to deploy?
- Cost depends on current infrastructure and potential amortization
- Benefit depends on attack vectors, and thus infrastructure

Securing the email infrastructure

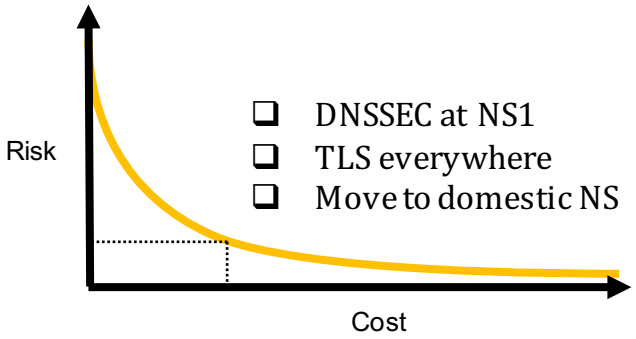

Administrators

Plan investments; secure infrastructure

Policy makers

Evaluate impact of policies; improve self-reliance

Standardization bodies

Focus efforts

Protocol designers

Assess deployment issues with regard to status quo

Running example

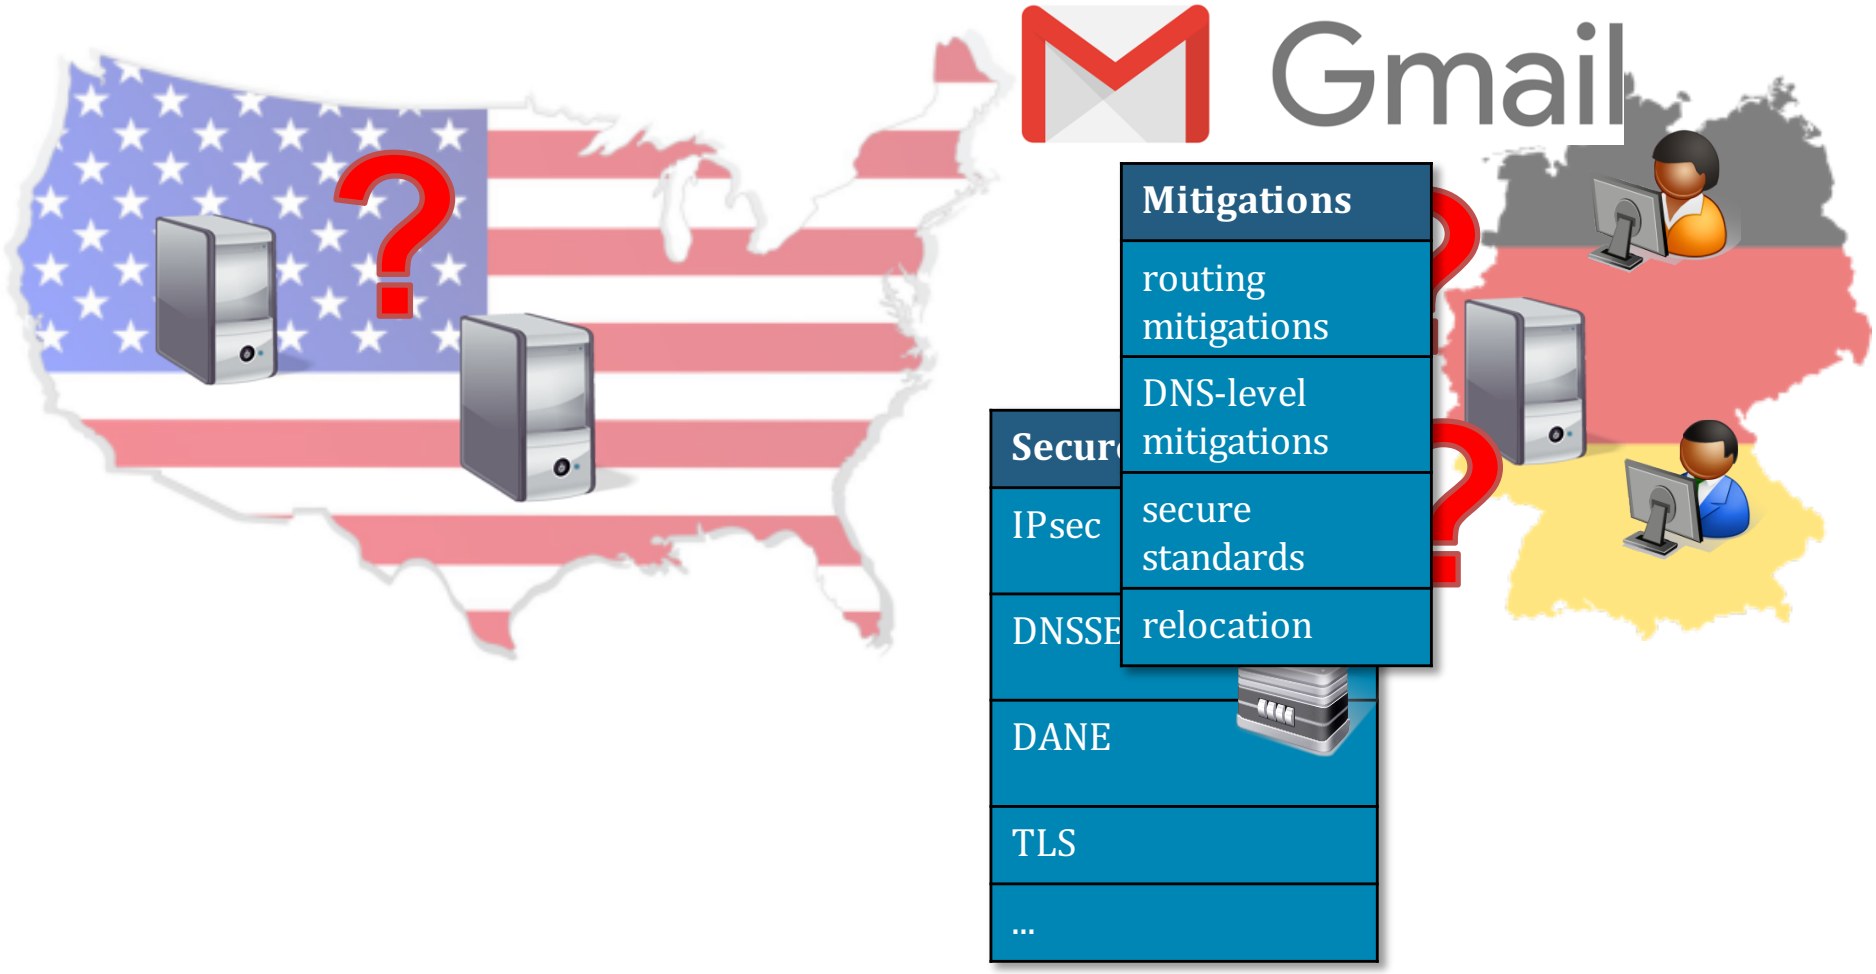

Property Graph

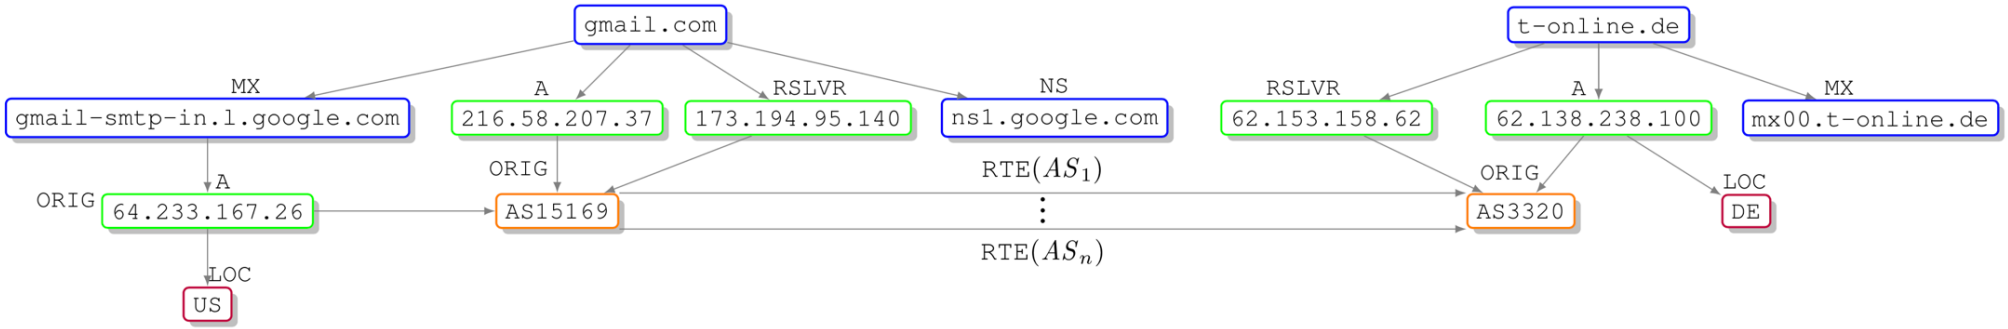

Attacker model: What is initially compromised?

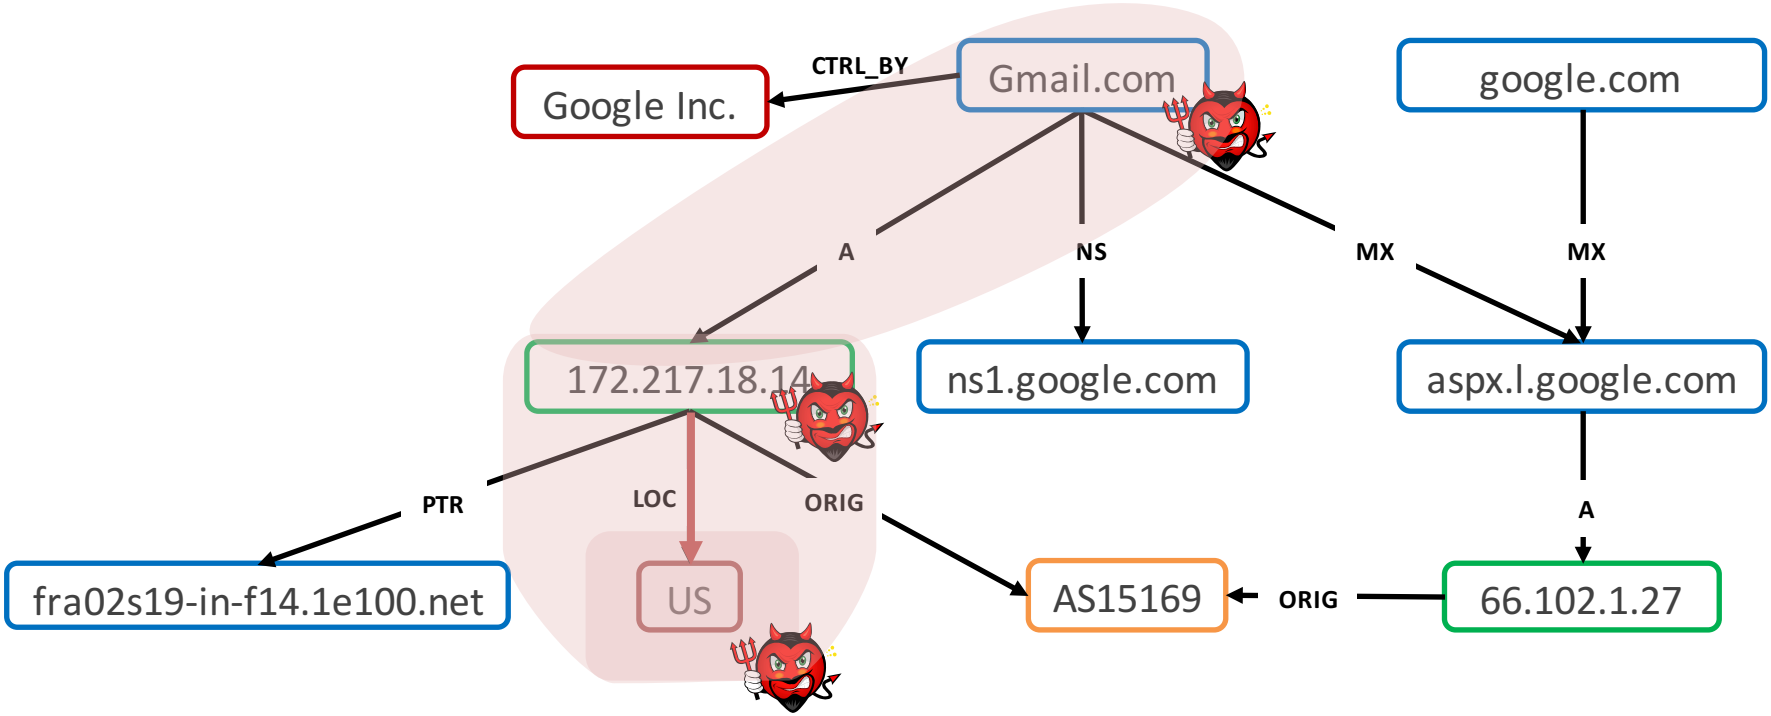

Attacker model & mitigations

- **Attacker (follower):**
  - **direct compromise:** mail servers (MX), name ser. (NS), aut. systems (AS)
  - **routing integrity:** affects email, but also name resolution
  - **name resolution integrity:** resolution chain
  - **communication integrity:** TLS, host validation, DANE etc.
- **Defender (leader):**
  - **VPN/IPSEC:** encrypt connections between two Ases
  - **DNSSEC:** NS signs responses to DNS requests
  - **TLS enforcement:** lose customers, still rely on name resolution
  - **RFC7817:** validate hostname with MX entry and @domain part
  - **DANE:** DNSSEC records indicate that TLS needs to be used
  - **relocate infrastructure:** e.g., use trusted NS

Cost functions/scenarios and attacker reward

- considered several different cost functions (money), e.g.,
  - actual spending (operational cost for 1 year)
  - loss of users when enforcing TLS (optimistic/pessimistic)

- Attacker reward is number of unconfidential connections
  - calculation based on market share

- additional scenario:
  - sneaky adversary (does not dare to forge DNSSEC signatures → remove actions)

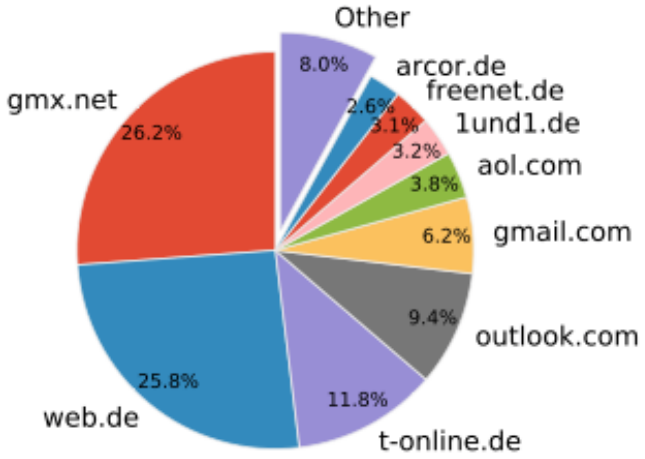

Data acquisition

- popular email providers
- DNS lookup
  - MX, A, RSLVR
- public available resources
  - Ripe Atlas, MaxMind

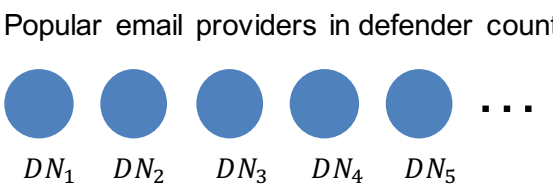

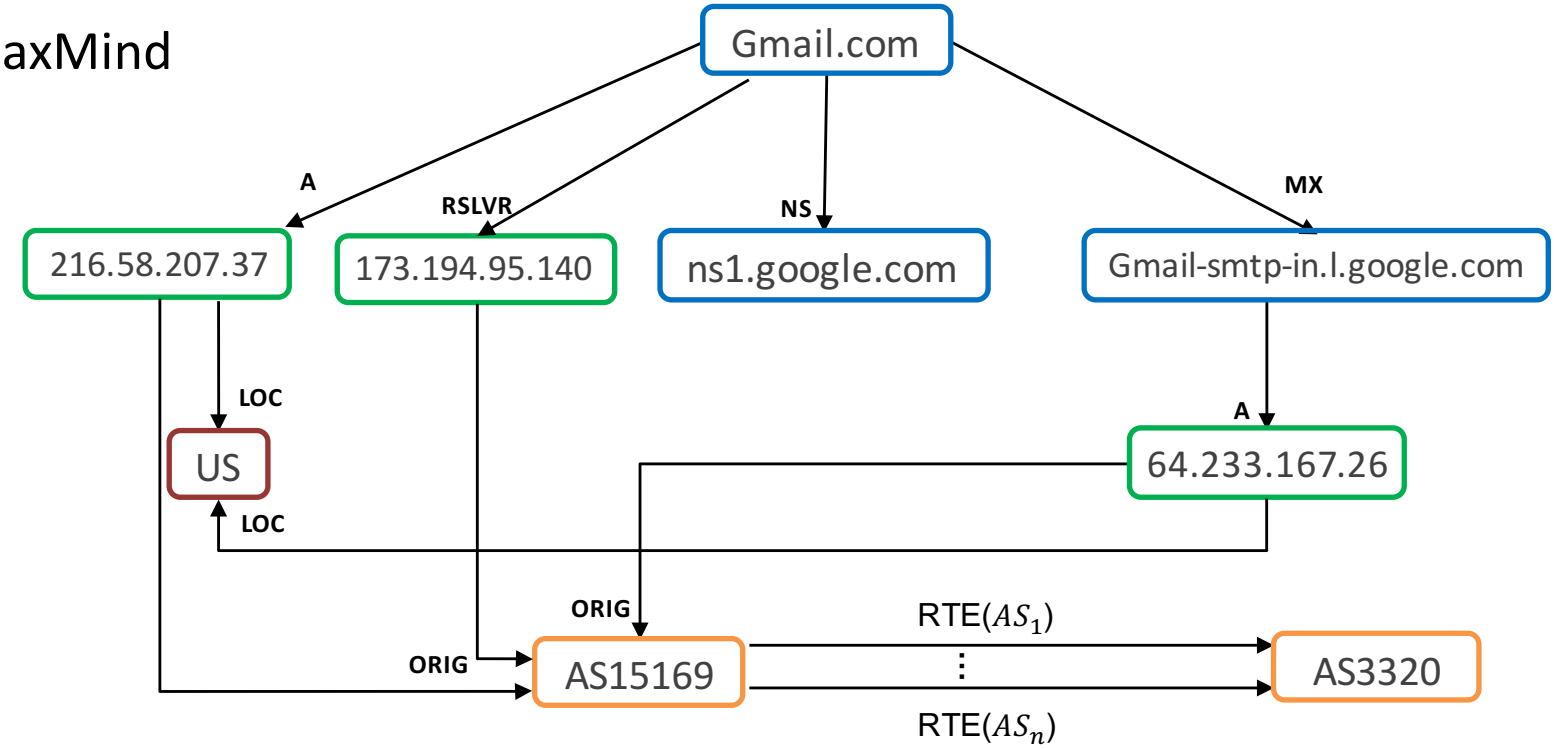

Results: Different attackers against Germany

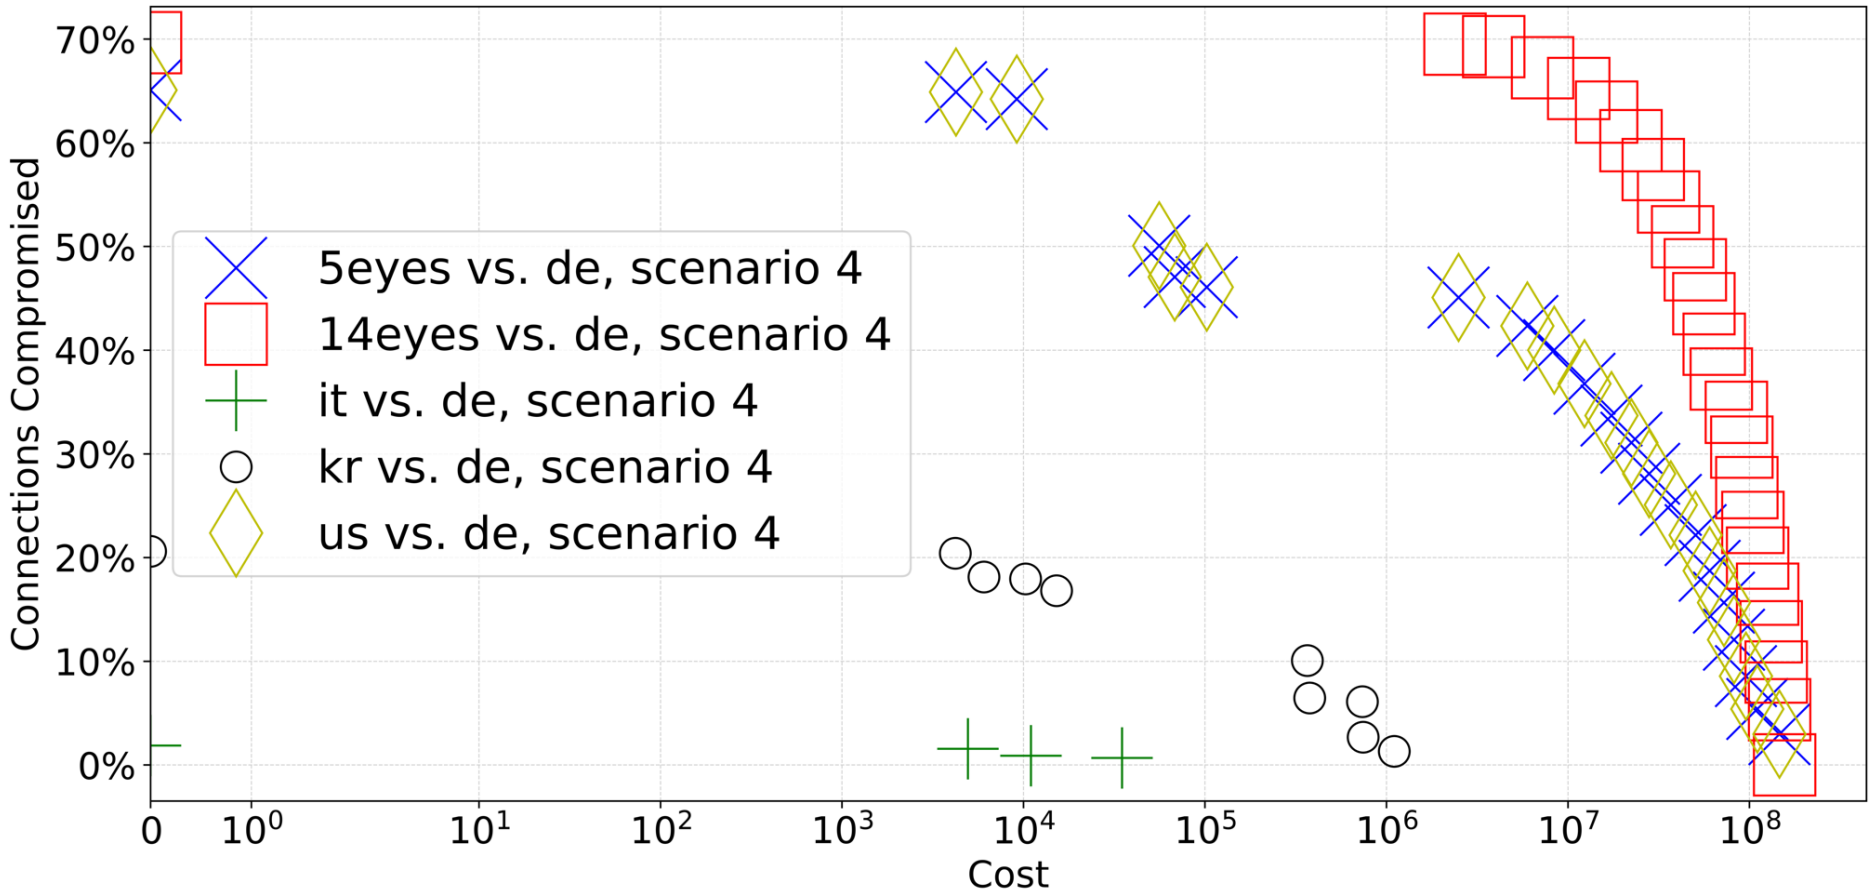

Results: some insights

- Germany: Considering only actual cost, **enforcing TLS with RFC 7817 is** both a **low-cost and high-security solution**. Foreign MXes need to be relocated, however, name server infrastructure is largely domestic.
- TLS would have a great effect in most countries, but **loss of functionality is prohibitive**:
  - In Germany, relocation is the next best option.
  - If we enforce TLS, should **implement RFC 7817** right away
- **DNSSEC** useful only if country wants to avoid leaving forensic evidence
- **Vast difference between countries** (Brasil, Russia)

Results: Check our website!

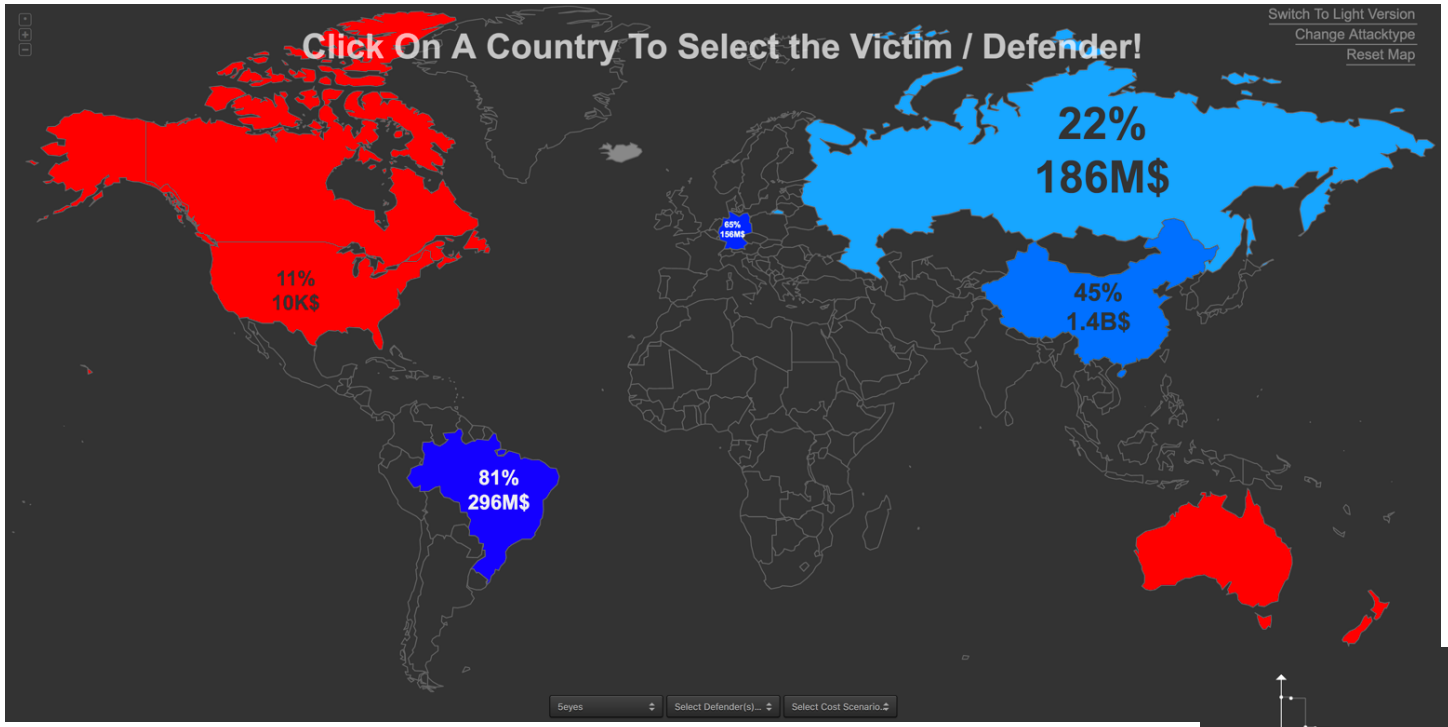

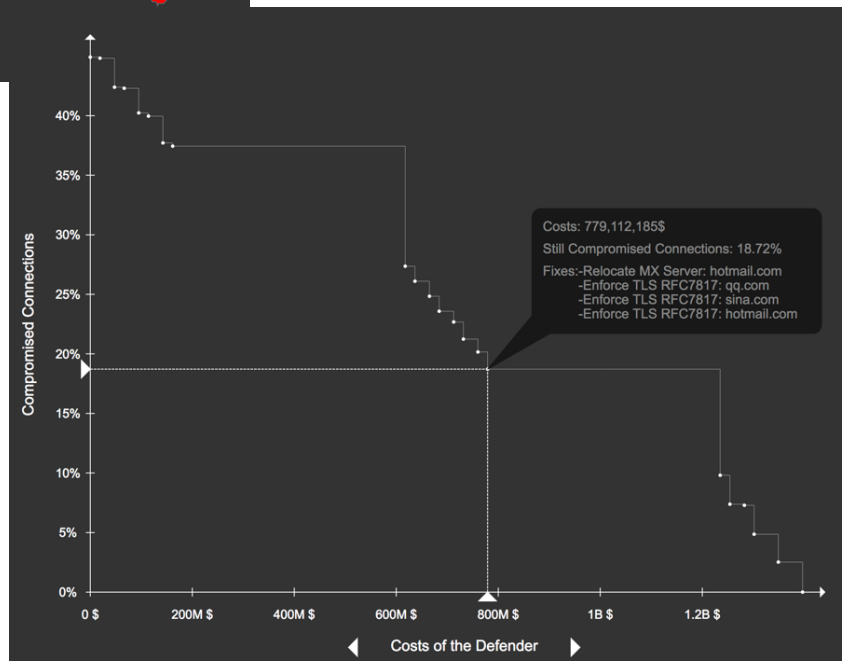

www.whocontrolstheinternet.com

Check our website!

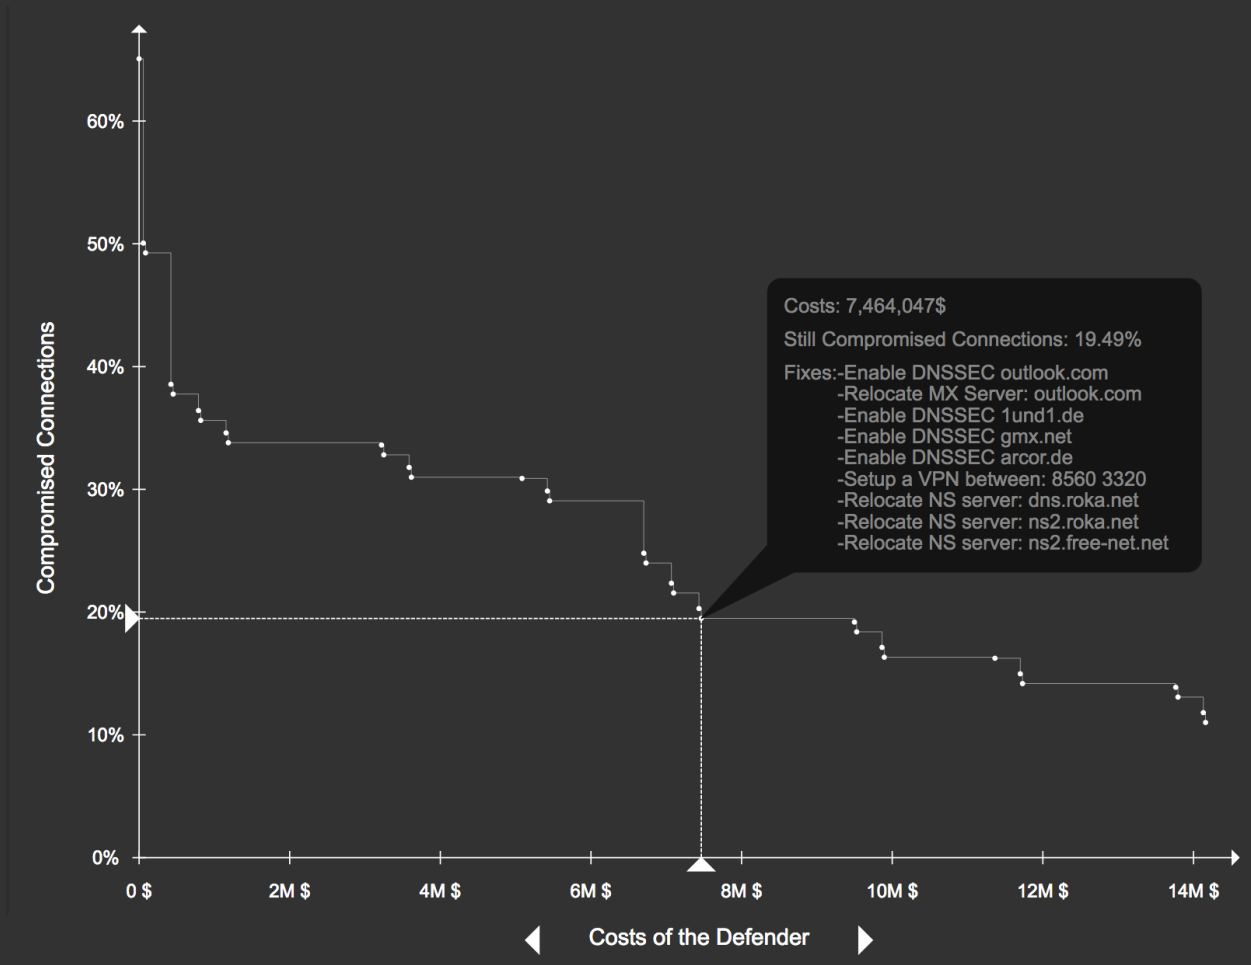

Second line of research: Securing corporate networks

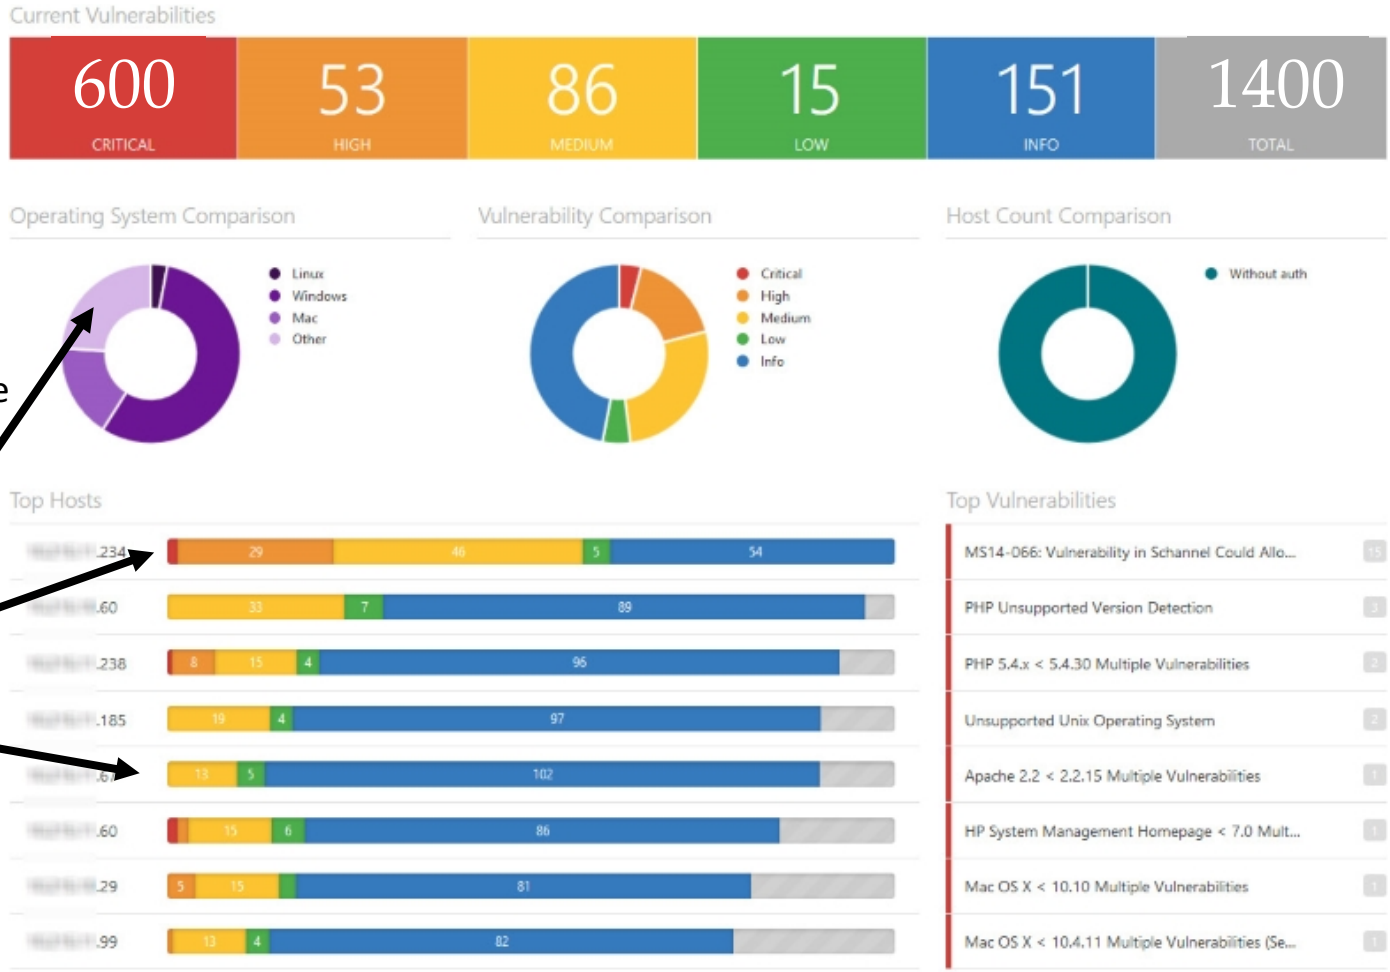

Too many;
Need to prioritize

Reachable from the outside?

Securing corporate networks: approach

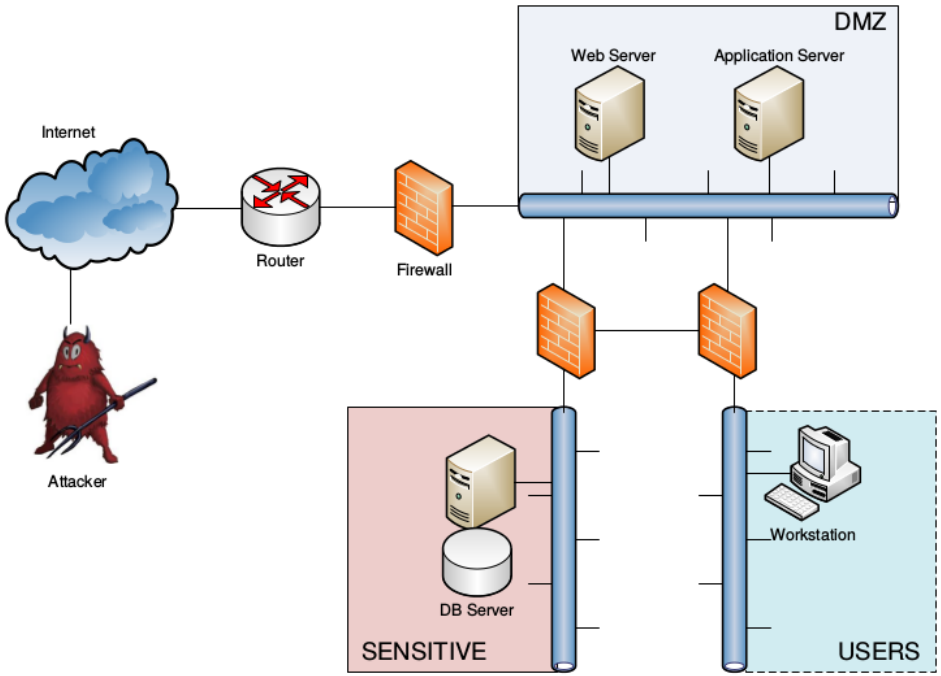

- attacker can hop if (severe, remote-exploitable) vulnerability exists
- wins if reaches critical asset
- defender: apply patch, filter packets, develop work around
- perform ok up to 800 hosts with ~5 vulnerabilities per host

Defender planning versus verification

- "100% security can never be reached"
- Still, we strive for 100%:
  - massive success of provable security in cryptography
  - success of verification at protocol level (will see with TLS 1.3)
  - progress for implementations
- 'Attacker success < x %' is a verification result (in our model!)
  - completeness of the model is the limit (well, and the budget)
  - opposed to simulated pentesting alone ("where to look")

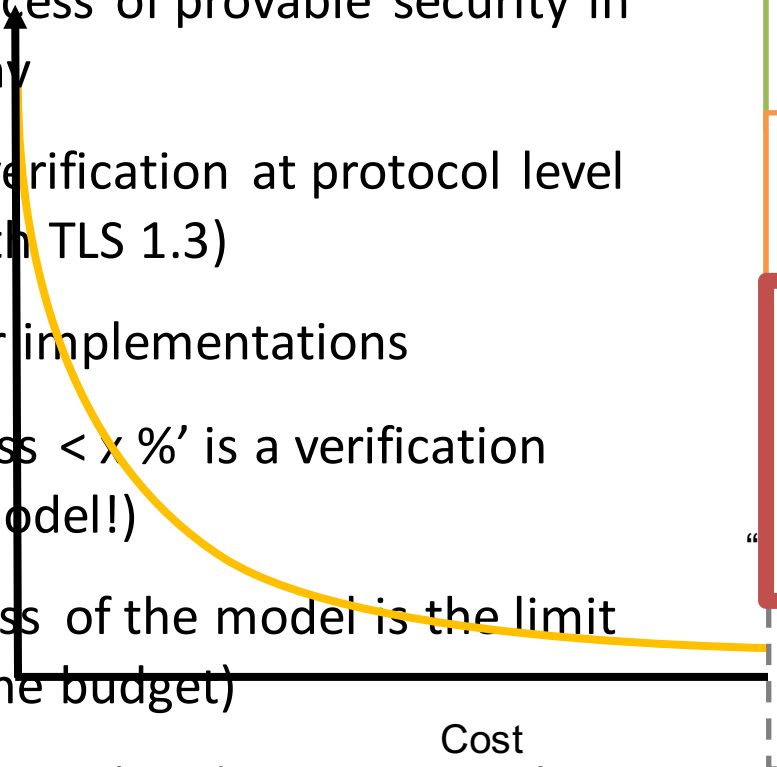

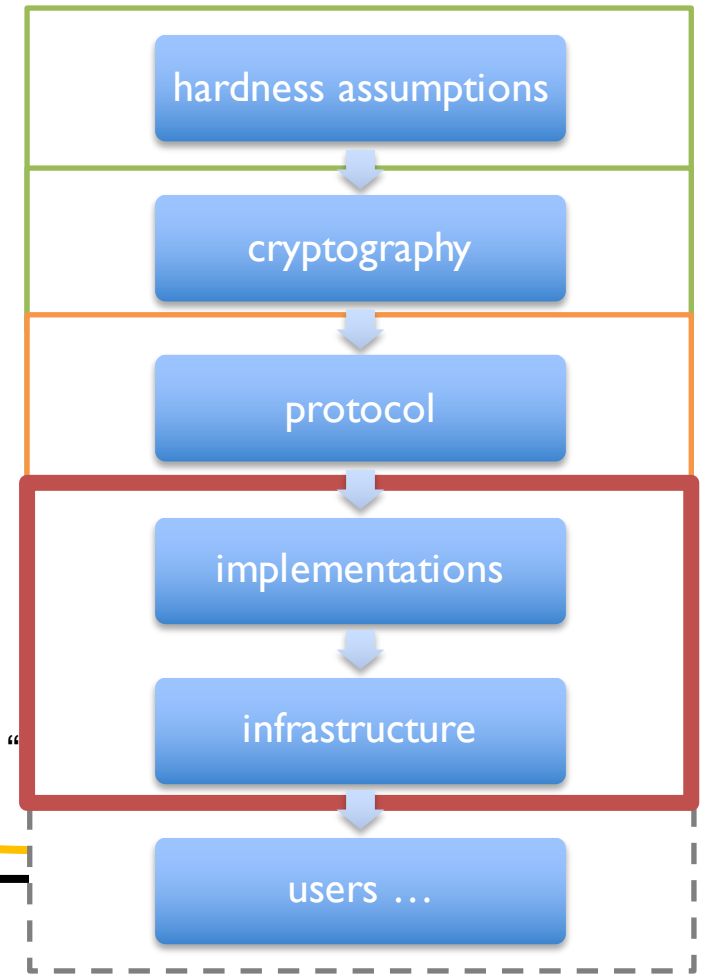

Defender planning versus verification

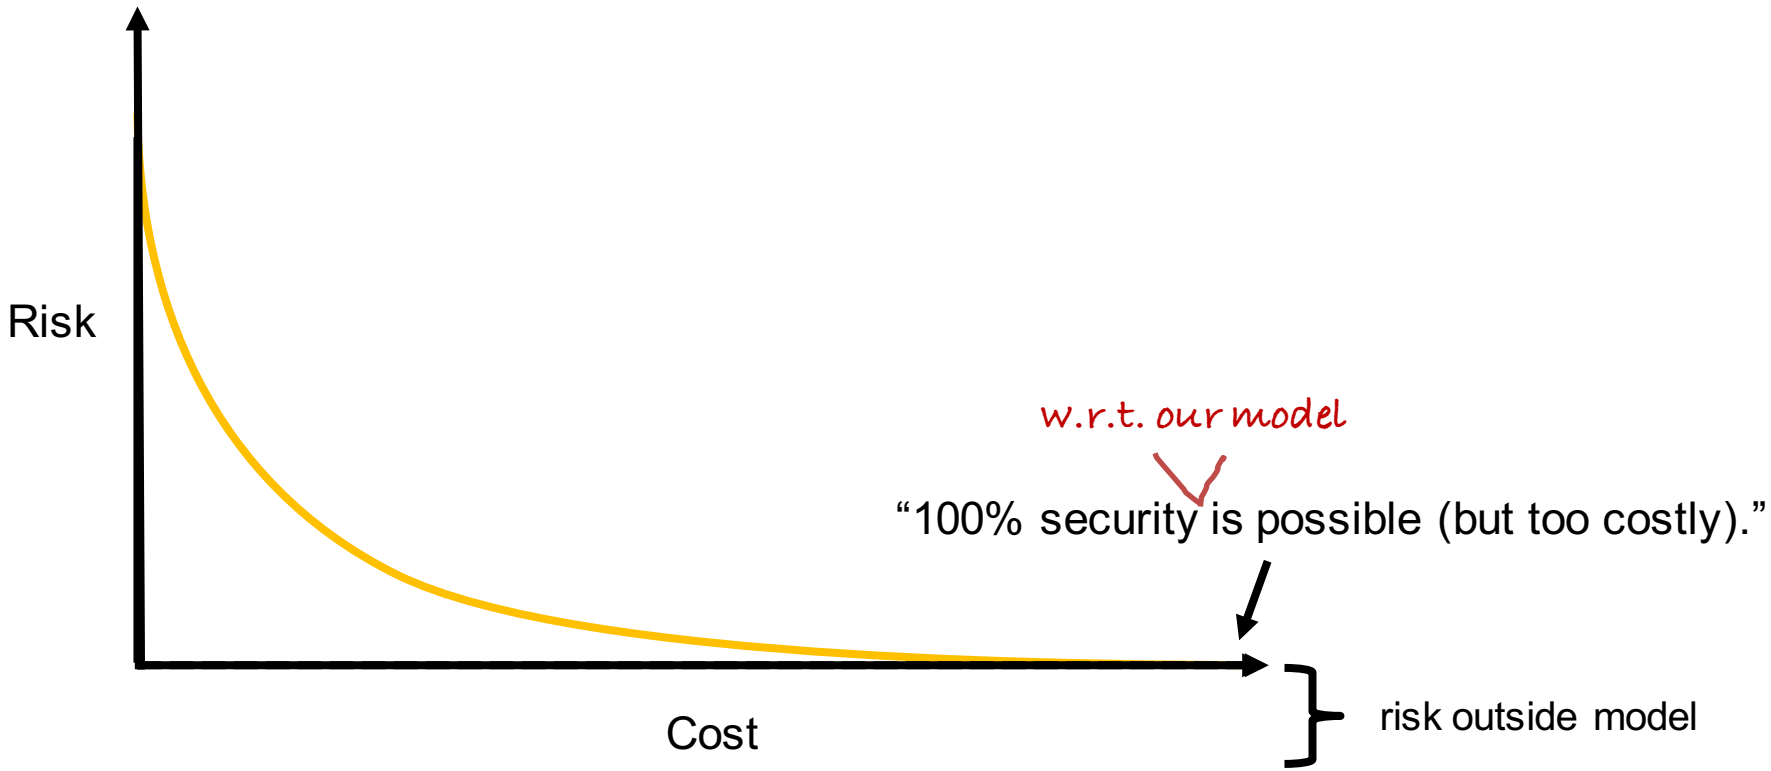

Future challenges (within AI)

- Almost universal disagreement over cost of mitigation
  - Compute region around Pareto frontier, showing deviation as function of action-cost uncertainty (intervals).
- Scalability of attacker task:
  - (Mostly) delete-relaxed, basically set of reachable nodes from initial set
  - Email: ca. 10 providers (representing 92% of users), 2mins
  - Web: ca. 5000 domains, 60k ASes., ??mins
  - Attacker planning problem can be solved/approximated using other techniques (graph algorithms, approximate counting)
- Malware spreading via zero-day exploits:
  - Probabilistic occurrence of vulnerability (MDP)
  - Emergency response (attacker goal: expected infections at time t)

Future challenges (outside AI)

- Infrastructure:
  - Link planning results to verification results for protocols.
  - Better data (scans, route prediction, cost estimates) .. cost estimation field-research required
  - More scenarios:
    - phishing, malware spreading/-response
    - Logical dependencies (shared DH groups in TLS, accumulated cost)
- Corporate networks
  - Better metadata through virtualisation (metasploit)
  - Validation (how to study effectiveness?)
  - Exploit chaining: theory of privilege escalation (work in progress)

CISPA HELMHOLTZ-ZENTRUM i. G.

Thank you for your attention.

Questions, comments?

Conclusion

- **Stackelberg planning:** practically relevant middle ground between full-scale game-theoretic planning and classical planning
- Many applications, despite lack of probabilistic defender actions (randomization is a protocol task!), e.g.
  - Cost-benefit analysis w.r.t. internet infrastructure:
    - next up: the web!
  - Mitigation planning in companies
- Algorithmic challenges in scalability and robustness of results

Excerpt of the Pareto frontier from USA vs. Russia

1) Reward = 0.217100, Cost = $0

$\emptyset$

2) Reward = 0.161200, Cost = $366 342

$\neg\text{nDNSSEC}(mail.ru)$

.
.
.

16) Reward = 0.000000, Cost = $186 404 888

$\neg\text{nDNSSEC}(rambler.ru)$, $\neg\text{nDNSSEC}(mail.ru)$, $\neg\text{nDNSSEC}(yandex.ru)$, $\neg\text{nReloc}(gmail.com)$, $\neg\text{nRFC7817}(yandex.ru)$, $\neg\text{nRFC7817}(rambler.ru)$, $\neg\text{nRFC7817}(mail.ru)$

Future challenges

- Infrastructure research line: link planning results to verification at lower architecture layers.
- Almost universal disagreement over cost of mitigation
  - Cost estimation field-research required
  - Compute region around Pareto frontier, showing deviation as function of action-cost uncertainty (intervals).
- Performance of attacker planning:
  - (Mostly) delete-relaxed, basically set of reachable nodes from initial set
  - Email: ca. 10 providers (representing 92% of users), 2mins
  - Web: ca. 5000 domains, 60k ASes., ??mins
  - Attacker planning problem can be solved/approximated using other techniques (graph algorithms, approximate counting)