Compositionality in Cybersecurity

Arman Khouzani, Pasquale Malacaria, Chris Hankin, Andrew Fielder and Fabrizio Smeraldi

p.malacaria@qmul.ac.uk

Queen Mary University of London

Imperial College London

Simons Institute, Berkeley

Compositionality

Dec 8, 2016

Decision support for cybersecurity: summary

- we consider the problem of optimal cybersecurity planning
- it is an adversarial problem so natural framework is game theory
- the state space in our real world case study has 10^{15} "pure" strategic behaviours
- we show how we can efficiently (under 1 second) find optimal solutions (equilibria) on this space by reasoning compositionally over security controls

Decision support for cybersecurity: timeline

- started in 2013 considering stochastic games, abstract interpretation and interactions between game theory and game semantics
- moved to Stackelberg games (security games) and affine transformation of zero sum games in 2014
- introduced multi-objective, multiple choice binary knapsack (2015)
- Mixed Integer Linear Programming MILP (2016) = Subgame Perfect Nash Equilibria, so Stackelberg solutions

Cyber-Security Planning

Definition

A **Cyber-Security Plan** is a set of defensive measures (a.k.a., controls) that are applied across an enterprise to improve its overall state of security.

- There are many security controls and each can be implemented at different intensity levels.
  Examples of controls: encryption, access control, firewall, patching, secure OS configuration, pen testing, password policy, etc ....

Cyber-Security Planning: Costs

- Each cyber-security control addresses a specific set of vulnerabilities
  $\Rightarrow$ A cyber-security plan should be composed of a combination of the measures to provide a well-rounded defense.

However, an exhaustive implementation of controls at maximum intensity is neither economically feasible nor managerially desirable for a SME.
Beside the overall **security risk**, must also be wary of:

- Aggregate Direct (Monetary) Costs: e.g. limited cyber-security budget
- Aggregate Indirect (Usability) Costs: e.g.: a low-budget but undesirable plan:
  - `force-install every patch upon release`
  - `min 16-char high-ent pwd, to be changed weekly`
  - `minimal whitelisting`
  - `maximal blacklisting`
  - `minimal priviledges, ...`

Queen Mary
University of London

Challenge I: Multi-Objective Optimization

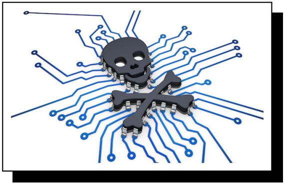

Security Costs (Risk)

Direct Costs

Indirect Costs

- Question: Why not simply minimize a weighted combination?
  - These costs are of hetrogeneous nature.
    - e.g. probabilistic and in-future vs deterministic and at-present
    - hard budgetary limit vs soft tolerance for usability costs
  - Require an a-priori vague determination of the weights:
    - e.g. if a small increase in one cost can improves the others significantly, one may relax her a-priori preference
- Solution Concept: **Pareto-Optimality**

Queen Mary University of London

Challenge II: Cyber-Security Risk depends on Threat Type

Passive

Reactive

APT

- Challenge: **implementation a security plan $\rightarrow$ changes the vulnerability profile $\rightarrow$ attack profile may adapt accordingly.**
- Classical "Risk Management" approaches assume the threat profile is *passive*.
  - e.g. the probability of occurence and intensity of natural disasters do not change based on defensive measures. But security is essentially adversarial (reactive)

Challenge III: Composing Security Controls

- Efficacy of an individual security control: The reduction in the success probability of exploitation attempts per each vulnerability when only that control is implemented (stand-alone).

Question: Often, the same vulnerability can be (partially) mitigated by more than one security measure, then what is the *combined efficacy*?

- **Additive:** assumes complementary defense mechanisms $\Rightarrow$ overestimates, mildly non linear
- **Multiplicative:** assumes "independent" defense mechanisms $\rightarrow$may still be an overestimation, also highly non linear
- **Best-of:** (per each vulnerability) the combined efficacy := (only) the highest efficacy among the implemented controls
  - captures positive "correlations" in defensive mechanisms, but non linear

Main Contributions

- Converting the Non-Linear Multi-Objective Optimizations into MILP: Mixed Integer Linear Programming for all 6 different settings.
  - E.g.: In our case-study, 10^{15} possible plans: state-of-the-art (Genetic Algorithms) will take weeks with no guarantee of optimality, but our MILPs return the exact Pareto-Front in seconds! in seconds.
  - Conducted the largest numerical evaluation to date
    - 37 most common vulnerabilities,
    - 27 distinct controls, each with multiple levels of implementation leading to 10^{15} distinct plans.

Modeling and Notations

- $\mathcal{C}$: set of (cyber-security) *controls*
- $\mathcal{L}_c = \{1, \ldots, L_c\}$ to denote the set of available implementation levels of control c.

Definition

A *cyber-security plan*, or a cyber-security investment portfolio $\boldsymbol{x} = (x_c)$ is a vector in $\mathcal{X} := \times_{c \in \mathcal{C}} (\{0\} \cup \mathcal{L}_c)$

- $B \in \mathbb{R}^+$ total cyber-security *budget*
- $D, I, R : \mathcal{X} \to \mathbb{R}^+$ respectively denote the (total) *direct cost*, (total) *indirect cost*, and the (aggregate) *"security risk"*

Problem Statement:

$$\min_{\boldsymbol{x} \in \mathcal{X}} (D(\boldsymbol{x}), I(\boldsymbol{x}), R(\boldsymbol{x})) \qquad \text{s.t.:} \quad D(\boldsymbol{x}) \leq B$$

Modeling and Notations

Aggregate Direct and Indirect Costs:

$$D(\boldsymbol{x}) = \sum_{c \in \mathcal{C}} d_c(x_c), \qquad I(\boldsymbol{x}) = \sum_{c \in \mathcal{C}} i_c(x_c)$$

Success Rate of Attempts on Vulnerability v:

Additive: $$S_v(\boldsymbol{x}) = \Big(1 - \sum_{c \in \mathcal{C}_v} e_{cv}(x_c)\Big)^+$$

Multiplicative: $$S_v(\boldsymbol{x}) = \prod_{c \in \mathcal{C}_v} s_{cv}(x_c)$$

Best-of: $$S_v(\boldsymbol{x}) = \min_{c \in \mathcal{C}_v} s_{cv}(x_c)$$

Queen Mary University of London

Modeling and Notations

Security Risk:

Passive: $$R(x) = \sum_{v \in \mathcal{V}} \mathbf{P}(v) S_v(\boldsymbol{x}) \lambda_v$$

Reactive: $$R(x) = \max_{v \in \mathcal{V}} \left(S_v(\boldsymbol{x}) \lambda_v \right)$$

Connection to Game Theory:

Proposition

Any strategy of the enterprise (the leader) in a Subgame Perfect Nash Equilibrium (SPNE) of the above non-zero-sum sequential two player game with "perfect information" is a Pareto-optimal solution to the multi-objective problem where the security cost is according to the "reactive threat" model. Conversely each point on the Pareto front is a SPNE in the game defined by that point direct and indirect costs.

Solving the Multi-Objective Optimization (MOOP)

Scalarization - I:

$$\min_{\boldsymbol{x}\in\mathcal{X}} \left[w_d D(\boldsymbol{x}) + w_i I(\boldsymbol{x}) + w_r R(\boldsymbol{x})\right] \qquad \text{s.t.: } D(\boldsymbol{x}) \leq B.$$

Scalarization – II

$$\min_{\boldsymbol{x}\in\mathcal{X}} R(\boldsymbol{x}) \qquad \text{s.t.: } I(\boldsymbol{x}) \leq \epsilon, \qquad D(\boldsymbol{x}) \leq B.$$

Still, highly non-linear optimizations → Tricks to convert them to MILPs (details in the paper).

Example: MILP formulation of best-of reactive optimization

main trick: use "flow variables" y_{vcl} to linearize the problem

$$\min_{(x_{cl}, y_{cvl})} \left[z + \delta_0 \sum_{v \in \mathcal{V}} P_v \lambda_v \sum_{\substack{c \in \mathcal{C}_v \cup \{0\} \\ l \in \mathcal{L}_c}} y_{vcl} s_{cv}(l) \right] \quad \text{s. t.:} \quad \Big(\sum_{l \in \mathcal{L}_c} x_{cl} \leq 1, \forall c \in \mathcal{C} \Big),$$

$$\big(x_{cl} \in \{0,1\},\ \forall l \in \mathcal{L}_c, \forall c \in \mathcal{C} \big), \quad \sum_{c \in \mathcal{C}} \sum_{l \in \mathcal{L}_c} d_c(l) x_{cl} \leq \epsilon_D,$$

$$\sum_{c \in \mathcal{C}} \sum_{l \in \mathcal{L}_c} i_c(l) x_{cl} \leq \epsilon_I, \Big(0 \leq y_{vcl} \leq x_{cl},\ \forall v \in \mathcal{V}, \forall c \in \mathcal{C}_v, \forall l \in \mathcal{L}_c \Big),$$

$$\Big(\sum_{c \in \mathcal{C}_v \cup \{0\}, l \in \mathcal{L}_c} y_{vcl} = 1,\ \forall v \in \mathcal{V} \Big),$$

$$\Big(z \geq \lambda_v \sum_{c \in \mathcal{C}_v \cup \{0\}, l \in \mathcal{L}_c} y_{vcl} s_{cv}(l),\ \forall v \in \mathcal{V} \Big).$$

Queen Mary University of London

Conclusions

- Compositionality and linearization allows us to solve complex strategic problems efficiently.
- We can linearize multiplicative, best-of composition of controls and their custom mixtures.
- The best-of reactive model is "validated" by comparing our tool solutions with official recommendations from GCHQ and SANS.
- Not clear what other compositionality principles are relevant.
- Is there a general theory of flow variables?

Queen Mary University of London