Coalition-Resistant Peer Rating for Long-Term Confidentiality

Giulia Traverso, *Denis Butin*, Alex Palesandro, Johannes Buchmann

TU Darmstadt, Germany & D2SI, France

Privacy Security Trust 2018

Context & Motivations

Context: Information-theoretic confidentiality

- Sensitive data require **long-term** confidentiality
- Currently used cryptography is unsuitable for long-term confidentiality
- Threats: cryptanalatic progress, quantum computers
- Long-term solution: **information-theoretic** confidentiality
- Can be realised through secret sharing

Context: Proactive secret sharing for long-term confidentiality

- Especially suitable for long-term confidentiality: **proactive** secret sharing
- Periodic renewal of shares
- Resilient to **mobile adversary** (collects shares over time)

Context: Distributed storage with different SSPs

- In an outsourcing scenario, proactive secret sharing can be performed on distributed storage system
- Distributed storage system consists of several **Storage Service Providers** (SSPs)
- Avoids single point of failure (single SSP key management)
- In practice, reliable proactive secret sharing requires high-performing SSPs
- We define *high-performing* in a broad sense; includes reliability

Context: Selecting high-performing SSPs

- How to select high-performing SSPs to build distributed storage system?
- Data owners require reliable guidance for this choice — pointed out by NIST for the special case of cloud infrastructures (NIST-SP 500-291)
- Data owners do not have access to comprehensive performance figures

Context: Obtaining SSP performance figures

- Use a third party to measure and publish performance figures?
- Impractical for a large number of SSPs and frequent measurements
- Alternative: aggregated peer rating — SSPs rate each other's performances. Third party also present, but only as a mediator: only aggregates ratings.

Motivation: Rational SSPs and accuracy

- Problem: SSPs benefit from providing selfish/false performance ratings to undermine competitors
- Naively computed aggregated performance scores unreliable
- For accuracy, need performance scoring mechanism encouraging accurate ratings
- Must model SSPs as rational agents rather than “good”/“bad”
- Natural framework for analysis of rational behaviour: game theory

Contributions

Contributions — Summary

1. Formalisation of computation of aggregate performance scores in game-theoretic framework. Game-theoretic model of SSP peer rating strategies
2. Formalised example of how unincentivised performance scoring mechanisms result in SSPs reporting inaccurate ratings
3. Incentivised performance scoring mechanism with incentive/penalty for accurate/inaccurate ratings, using a TA. Assuming honest majority, accuracy resilient to coalitions (coordinated groups) of SSPs
4. Model of this mechanism as an infinitely repeated game in game-theoretic formalism, and proof of k-resilient equilibrium

Contributions — Remarks

- We do *not* aim at cryptographically improving proactive secret sharing
- Rather, focus is decision support for the selection of high-performing SSPs storing shares
- This supports reliable long-term confidential data storage
- We first show that aggregate performance scores are *not* accurate if participating SSPs are not incentivised to report faithfully
- We then present incentivised scoring mechanism + accuracy proof
- Accuracy margin depends on coalition sizes

Game-theoretic framework (1)

- Idea: game-theoretic formalism models peer rating strategies of the SSPs ("players") $P_1 \ldots P_n$
- $\mathcal{A}_i$ is the set of possible actions of player P_i
- *Action profile*: $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$
- Utility function $u_i : \mathcal{A} \to \mathbb{R}$ of a player defines its preferences
- Strategy $\sigma_i : \mathcal{A}_i \to [0,1]$ for a player P_i: probability distribution
- Game denoted $\Gamma(P_i, \sigma_i, u_i)$, for $i = 1, \ldots, n$

Game-theoretic framework (2)

- Players act rationally: they always play the strategy maximising their utilities
- Non-cooperative game: players choose actions individually
- Cooperative game: players form coordinated coalitions

Game-theoretic framework (3)

Notation: $(\boldsymbol{\sigma}'_C, \boldsymbol{\sigma}_{-C}) = (\sigma'_1, \ldots, \sigma'_{|C|-1}, \sigma'_{|C|}, \sigma_{|C|+1}, \ldots, \sigma_n)$ — players in coalition play strategy σ', outsiders play σ

Strategy σ *dominates* strategy σ' if it always provides its player with a higher utility (pay-off). Denoted $\sigma' \leq \sigma$ (weak) or $\sigma' < \sigma$ (strict)

For cooperative games, we use notion of k-resilient equilibrium:

A joint strategy $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ is a k-resilient equilibrium if $u_i(\boldsymbol{\sigma}'_C, \boldsymbol{\sigma}_{-C}) \leq u_i(\boldsymbol{\sigma}_C, \boldsymbol{\sigma}_{-C})$, for each subset $C \subset \{P_1, \ldots, P_n\}$ of cardinality $n_C \leq k$, where $\sigma'_i \neq \sigma_i$, for $P_i \in C$.

$\boldsymbol{\sigma}$ yields best pay-off for coalition members, for coalitions of size up to k

Score inaccuracy for unincentivised ratings

- Assume that TA computes aggregate score of each P_i by simply taking into account raw ratings from P_j, $J \neq i$, $1 \leq i \leq n$, i.e. by averaging them
- Aggregate scores output to data owner upon request
- We show formally that such unincentivised score computation does not lead to accurate aggregate scores
- Shown both for the case of non-cooperating and cooperating players
- Proof strategy: consider a number of rating strategies for players, including the one where accurate ratings are given. Show that giving faulty low ratings to all other players is the dominant strategy when goal is to maximize own aggregate score

Outline of new performance scoring mechanism

- At round r, each aggregate score τ_i^r for P_i is computed as convex combination of components τ_i', τ_i'' and τ_i^{r-1} (for $1 \leq i \leq n$)
- τ_i' is aggregate score of all ratings submitted by all players for player P_i being rated for current round r
- τ_i'' is aggregate score of incentives and penalties given to P_i by TA for accurate/inaccurate ratings for current round r — see next slide
- τ_i^{r-1} is aggregate score of P_i at previous round $(r-1)$, previously computed

Notation: round r implicit for τ_i' and τ_i''

Performance scoring mechanism: computing τ_i''

Notation: $\rho^r_{i,j}$ = rating submitted by P_i about P_j for round r

Computing the second component τ_i''

1. t_ε arbitrarily selected before the mechanism starts
2. Select incentive α and penalty β with $0 < \alpha \leq t_\varepsilon$ and $-t_\varepsilon \leq \beta < 0$
3. For all j such that $1 \leq j \neq i \leq n$, compute
   $$o_{i,j} = \begin{cases} \alpha, & \text{if } |\tau_j' - \rho^r_{i,j}| \leq t_\varepsilon \\ 0, & \text{if } t_\varepsilon < |\tau_j' - \rho^r_{i,j}| \leq 2t_\varepsilon \\ \beta & \text{if } |\tau_j' - \rho^r_{i,j}| > 2t_\varepsilon \end{cases}$$
   (penalize outlier ratings by P_i about P_j)
4. $\tau_i'' = \frac{1}{n-1} \sum_{j=1}^{n-1} o_{i,j}$ (averaging over j)

Some more formalism

Weight of player P_j w.r.t. evaluation of P_i: $w^r_{j,i} = \frac{\tau_j^{r-1}}{\sum_{l \neq i} \tau_l^{r-1}}$, with $\sum_{j \neq i} w^r_{j,i} = 1$, i.e. empower highly-rated players

In performance scoring mechanism, parameter ε (weight threshold) computed. Limits the weight of a coalition against all others. Depends on coalition size C_k

Utility functions

Let $M = \{m \mid P_m \in C_k\}$ and $L = \{l \mid P_l \notin C_k\}$.

Utility function for player P_i in coalition C_k with respect to aggregate scores, with $|C_k| = n_k$, $1 \leq i \leq n$:

$$u_i(r) := \begin{cases} \frac{1}{n_k}\sum_{m\in M}\tau_m^r > \frac{1}{n-n_k}\sum_{l\in L}\tau_l^r \implies u_i(r) = 1 \\ \frac{1}{n_k}\sum_{m\in M}\tau_m^r = \frac{1}{n-n_k}\sum_{l\in L}\tau_l^r \implies u_i(r) = 0 \\ \frac{1}{n_k}\sum_{m\in M}\tau_m^r < \frac{1}{n-n_k}\sum_{l\in L}\tau_l^r \implies u_i(r) = -1 \end{cases}$$

Coalition members want to be better rated on average than outsiders

Score accuracy for incentivised ratings

Main result (Theorem 2)

Let $\varepsilon > 0$ be a weight threshold and let C_k be the biggest coalition for ε, with $|C_k| = K$. The infinitely repeated cooperative game $\Gamma(P_i, \sigma_i, u_i)$, for $i = 1, \ldots, n$, with utility $u_i(r)$ and the above mechanism run at every round, reaches a K-resilient equilibrium for the computations of aggregate scores $\tau_1^r, \ldots, \tau_n^r$ if

$$\frac{\sum_{i \in C_k} w_{i,m}^r}{\sum_{j \notin C_k, j \neq m} w_{j,m}^r} \leq \varepsilon.$$

Score accuracy for incentivised ratings — Interpretation

Infinitely repeated: players do not know when last round happens

Honest majority assumed

Weight of biggest coalition C_k bounded depending on accuracy threshold t_ε

Unfair (selfish) ratings deviating from accurate mainstream are detected

Unfair raters penalized w.r.t. own rating

It pays to rate accurately!

Conclusions

Conclusions

- Performance scoring mechanisms can be resilient to coalitions of rational SSPs if majority of SSPs is honest
- Guiding data owners in their SSP selection supports long-term confidentiality
- Experimental validation needed to estimate average number of rounds for aggregate scores to converge

Thank You!

Questions?