Resilient Observation Selection in Adversarial Settings

Aron Laszka[1], Yevgeniy Vorobeychik[2], and Xenofon Koutsoukos[2]

[1] University of California, Berkeley
[2] Vanderbilt University

Monitoring Spatially-Distributed Systems

- To dynamically control a system, we must have accurate information about its evolving state
- In a number of applications, such as electrical grids or traffic networks,
  - system to be monitored can extend over a vast area
  - there can be a large number of possible points of observation

Where should we place sensors for monitoring a system?

Cyber-Attacks Against Sensors

Example: in 2008, a major Turkish oil pipeline suffered a cyber-attack

- attackers **disabled** the pressure and flow **sensors**, which allowed them to super-pressurize the oil in the pipeline, causing an **explosion**
- control room did not learn of the blast for 40 minutes after it happened

How can we improve the attack-resilience of monitoring?

Resilient Sensor Placement

- Resilience:
  even if some of the sensors are disabled by an attack,
  we are able to reliably estimate the state of the system
- ~~Redundant placement~~
  - ~~simply place more sensors~~
  - ~~requires excessive spending~~
- Limited budget
  - limited number of sensors
- Resilient placement:
  placing sensors so that together they are resilient to attacks

Problem Formulation

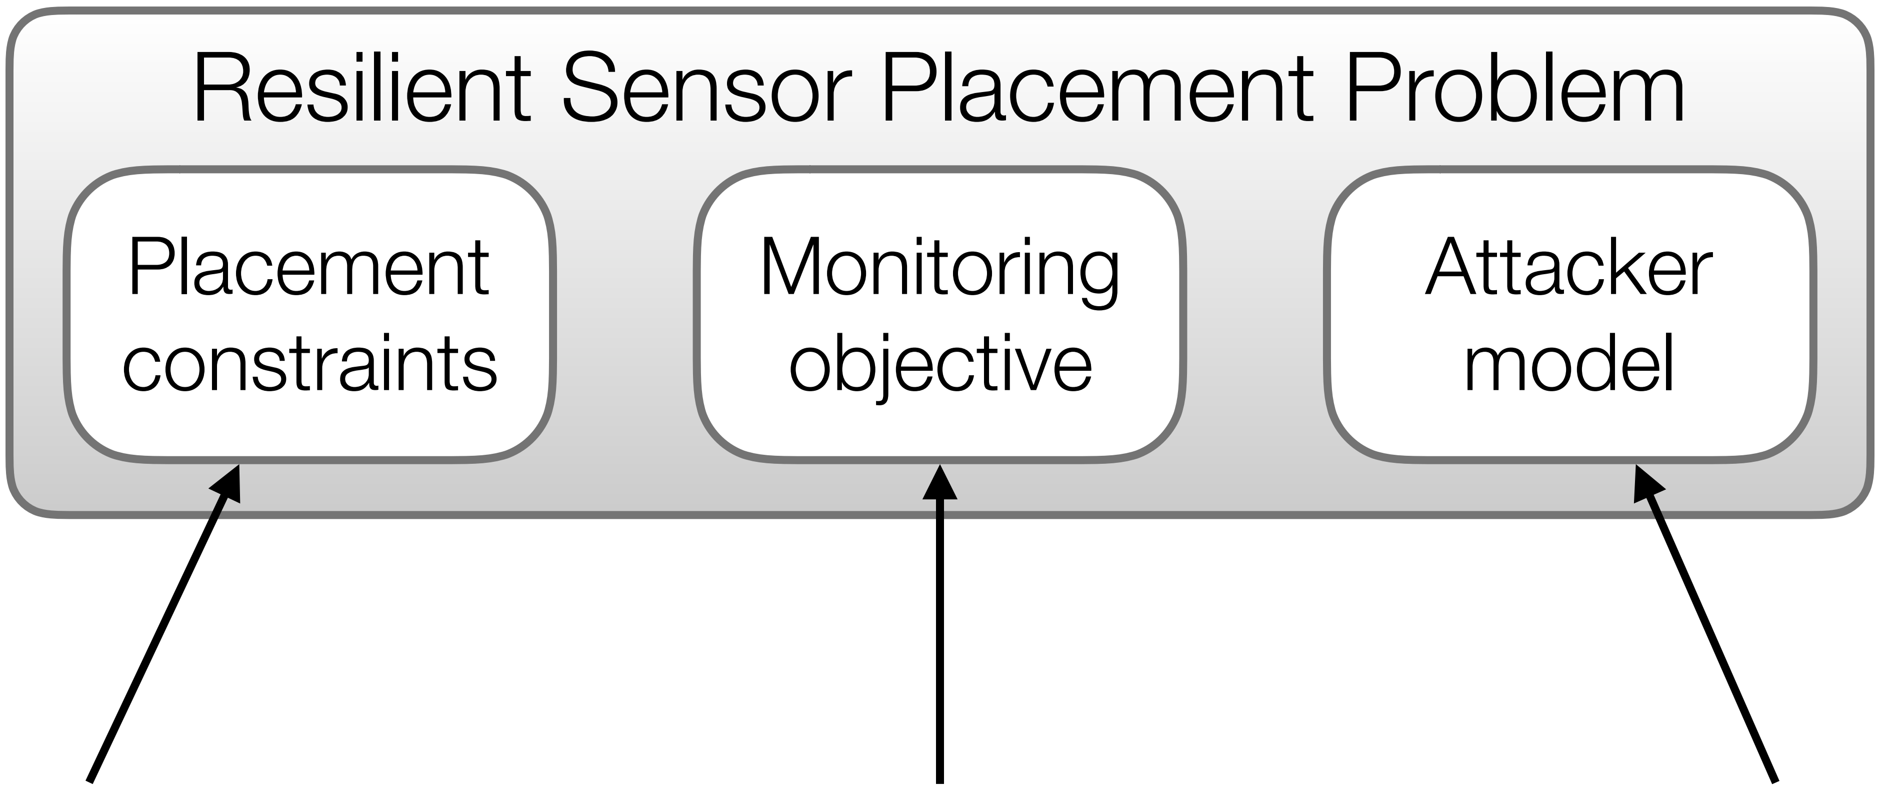

set of locations $\mathcal{V}$,
number of sensors N

$$\mathcal{S} \subset \mathcal{V} : |\mathcal{S}| = N$$

minimizing the variance
of Gaussian-process
based regression

$$\min \sigma^2_{Y|\mathcal{S}}$$

worst-case disabling
attacks against at
most K sensors

$$\mathcal{A} \subset \mathcal{S} : |\mathcal{A}| \leq K$$

Objective

- Gaussian-process based regression
  - kernel-based machine learning method
  - assumes that the joint distribution of the observations and the predictor variable is a Gaussian
  - has been used to predict various physical processes, e.g., road traffic
- Objective: variance of the predictor variable
  - depends only on the prior (i.e., where sensors are placed)
  - proportional to mean squared error or uncertainty
- Resilient sensor placement problem:

$$\min_{\mathcal{S} \subset \mathcal{V} : |\mathcal{S}| = N} \; \max_{\mathcal{A} \subset \mathcal{S} : |\mathcal{A}| \leq K} \; \sigma^2_{Z|(\mathcal{S} \setminus \mathcal{A})}$$

How much can we improve resilience through placement?

Numerical Results Based on Traffic Data

- Dataset
  - from the Caltrans Performance Measurement System (PeMS)
  - real-time data from sensors spanning across all major metropolitan areas of California
  - we used hourly traffic data from January 2015
- Setup
  - we selected 37 locations from the Bay Area as possible sensor locations
  - predictor variable is the average traffic over the area

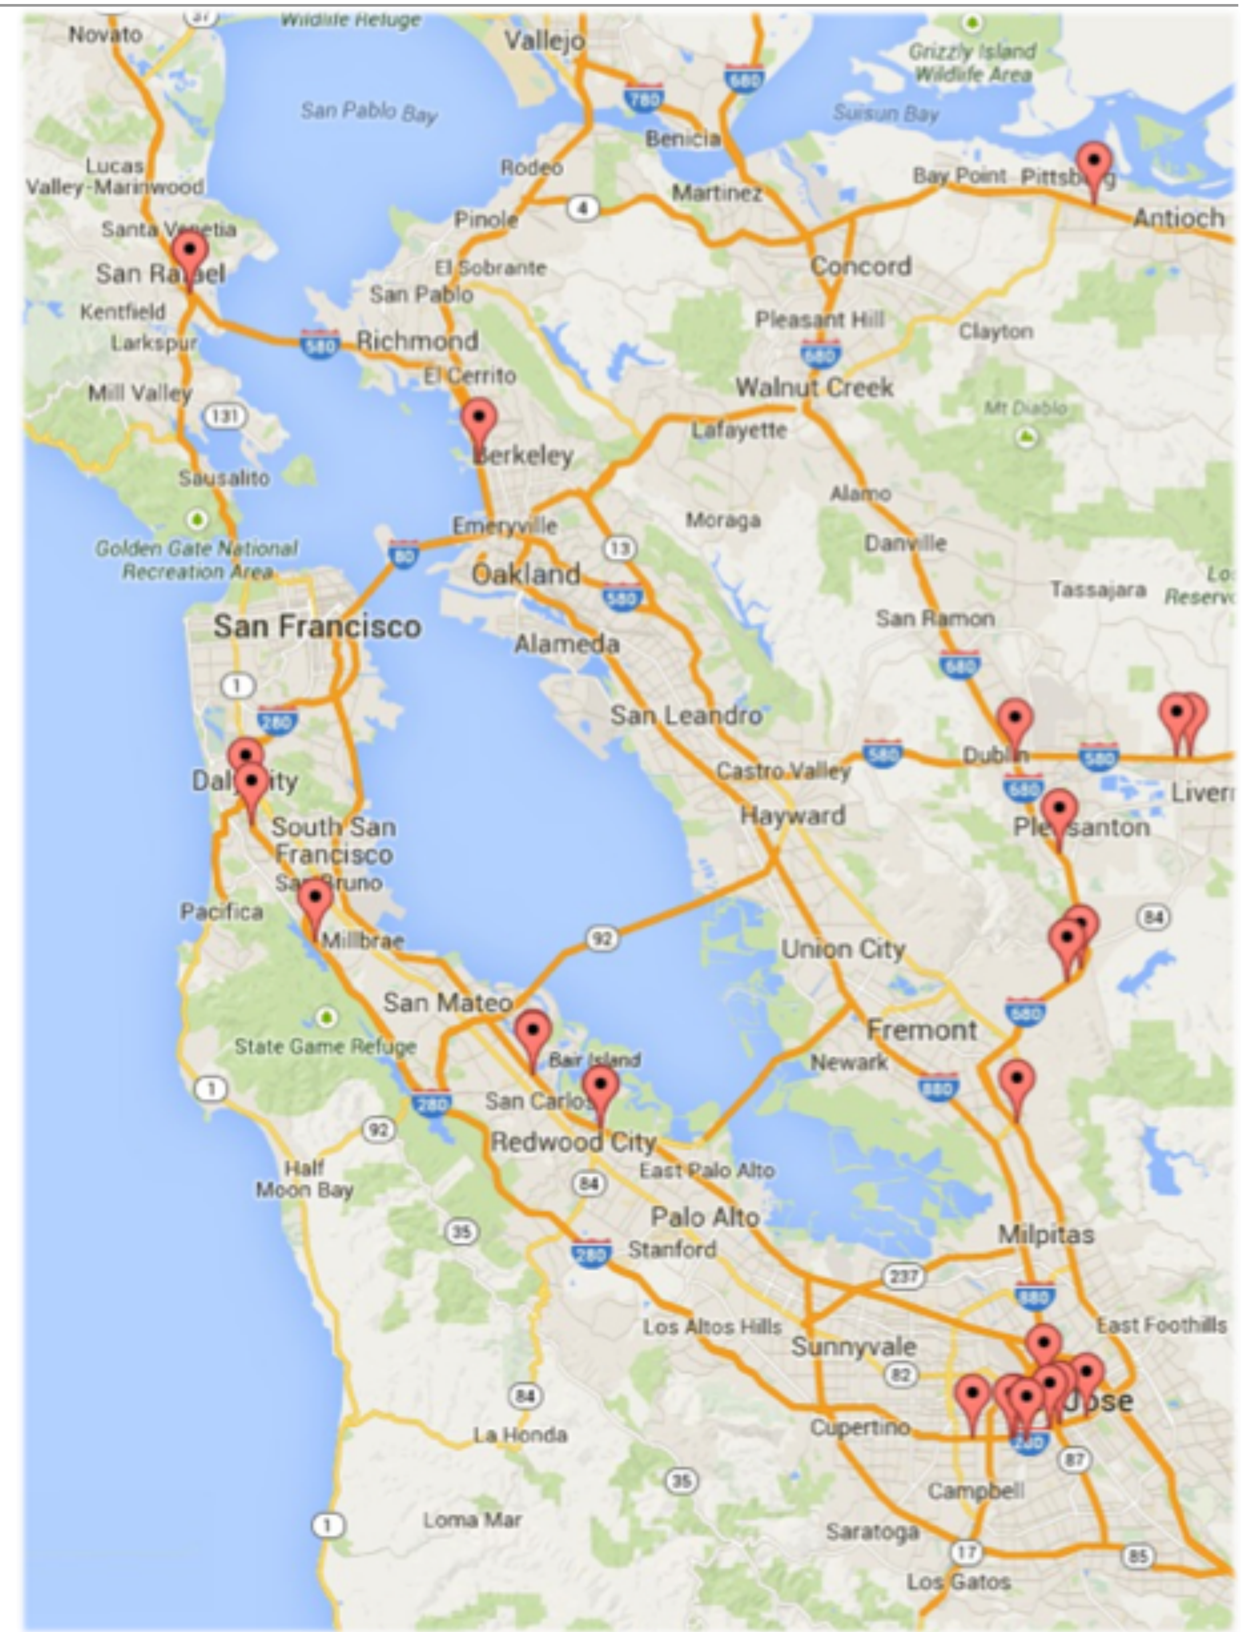

Resilient and Non-Resilient Sensor Placements

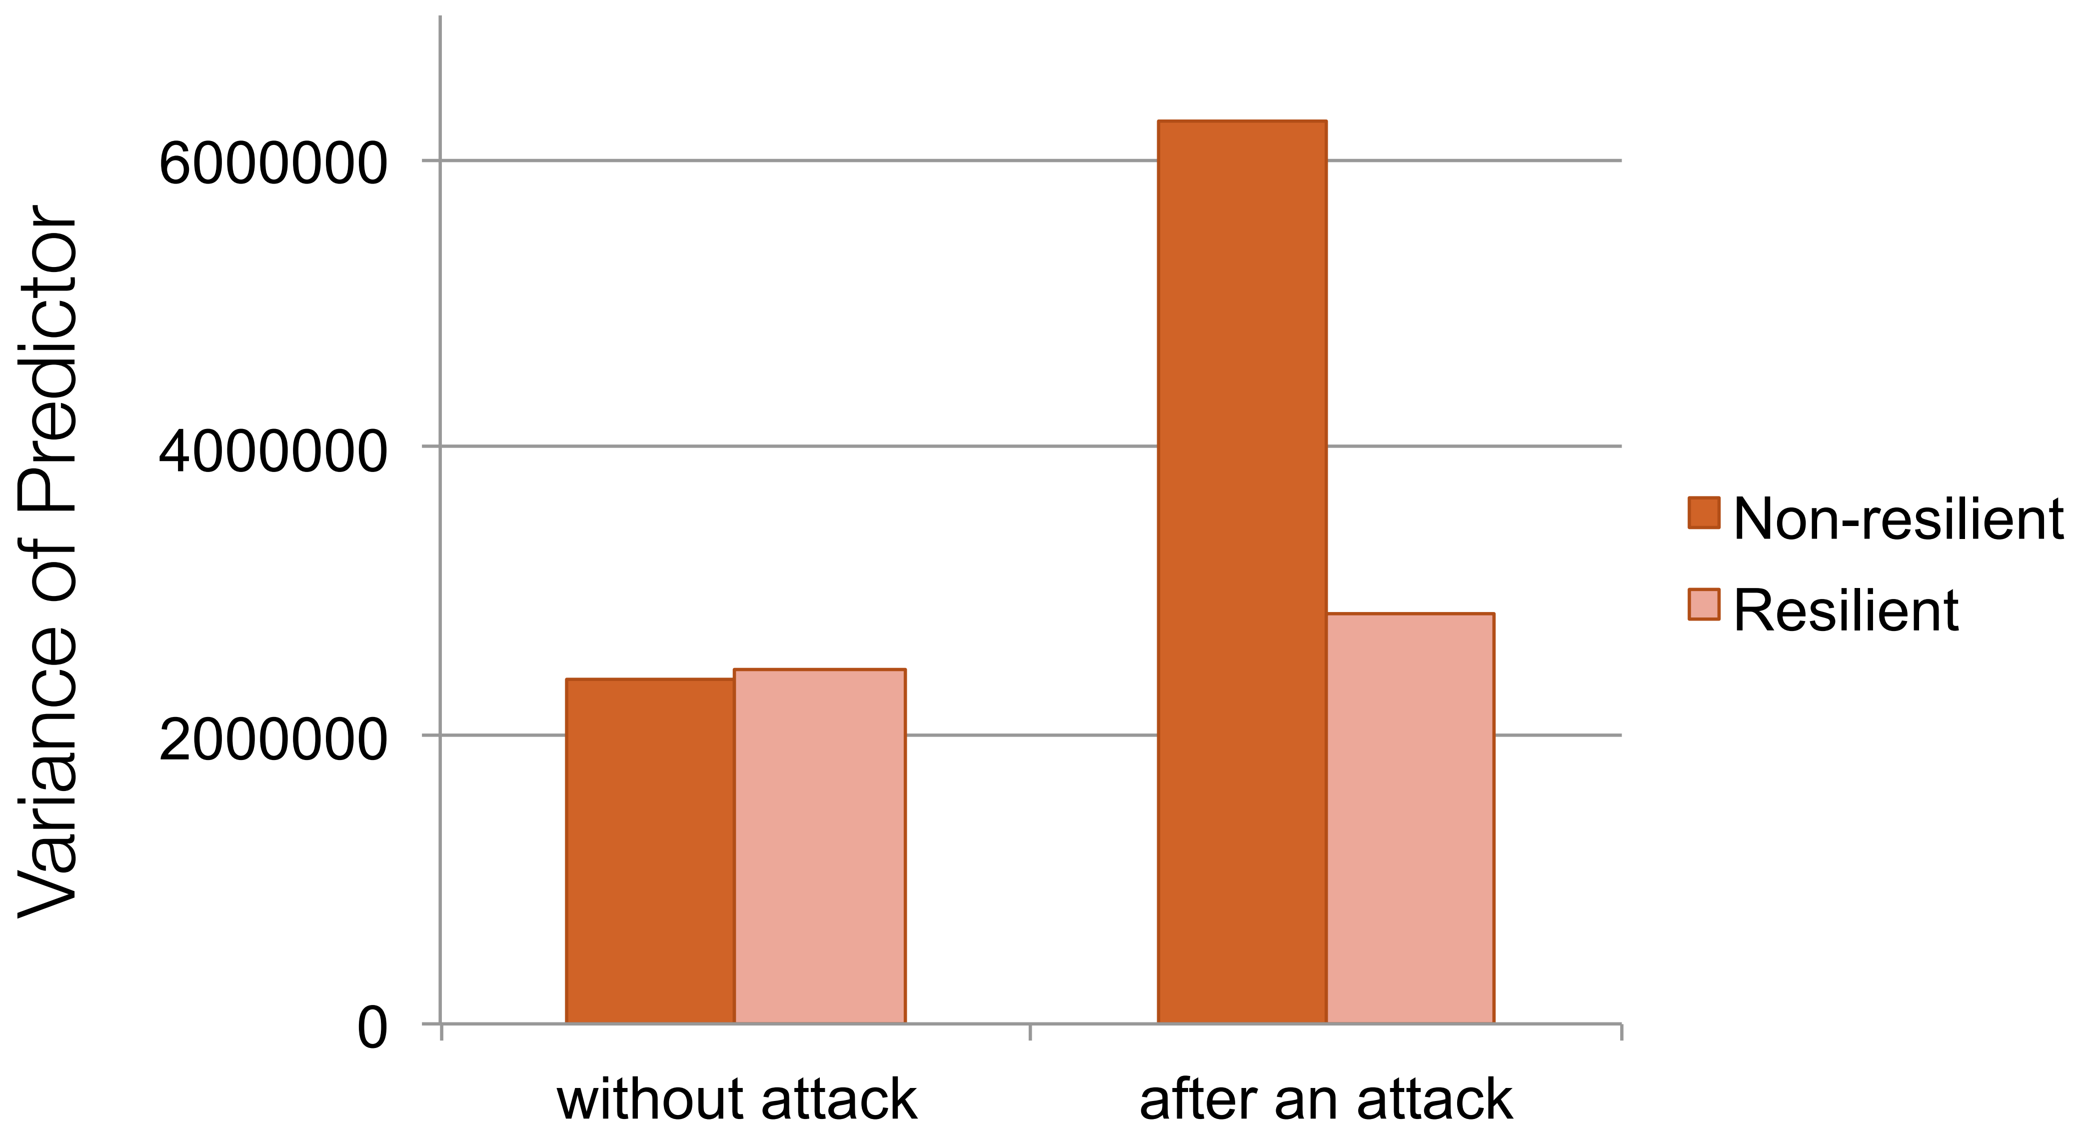

$N = 8$ sensors and $K = 1$ attack size

Computational Complexity

- Exhaustive search over all subsets is not feasible in practice
- *Can we find an optimal placement in polynomial time?*

Theorem: The resilient sensor placement problem is **NP-hard**.

- *Can we at least compute how resilient a given placement is (i.e., find an optimal attack)?*

Theorem: Finding an optimal attack for a given placement is **NP-hard**.

- Polynomial-time algorithms for resilient sensor placement
  1. heuristic and approximation algorithms
  2. algorithms for special cases

Heuristic and Approximation Algorithms

$$\underbrace{\min_{\mathcal{S}\subset\mathcal{V}:\,|\mathcal{S}|=N} \underbrace{\max_{\mathcal{A}\subset\mathcal{S}:\,|\mathcal{A}|\le K} \sigma^2_{Z|(\mathcal{S}\setminus\mathcal{A})}}_{\text{attacker's problem}}}_{\text{sensor placement problem}}$$

- First, we propose a greedy heuristic for finding an attack:

1. Let $\mathcal{A} \leftarrow \emptyset$
2. While $|\mathcal{A}| < K$:
    1. Let X be a variable maximizing $\sigma^2_{Z|(\mathcal{S}\setminus(\mathcal{A}\cup\{X\}))}$
    2. Let $\mathcal{A} \leftarrow \mathcal{A} \cup \{X\}$
3. Return $\mathcal{A}$

Greedy Attack: Numerical Results

- Less than 4% difference throughout numerous experiments

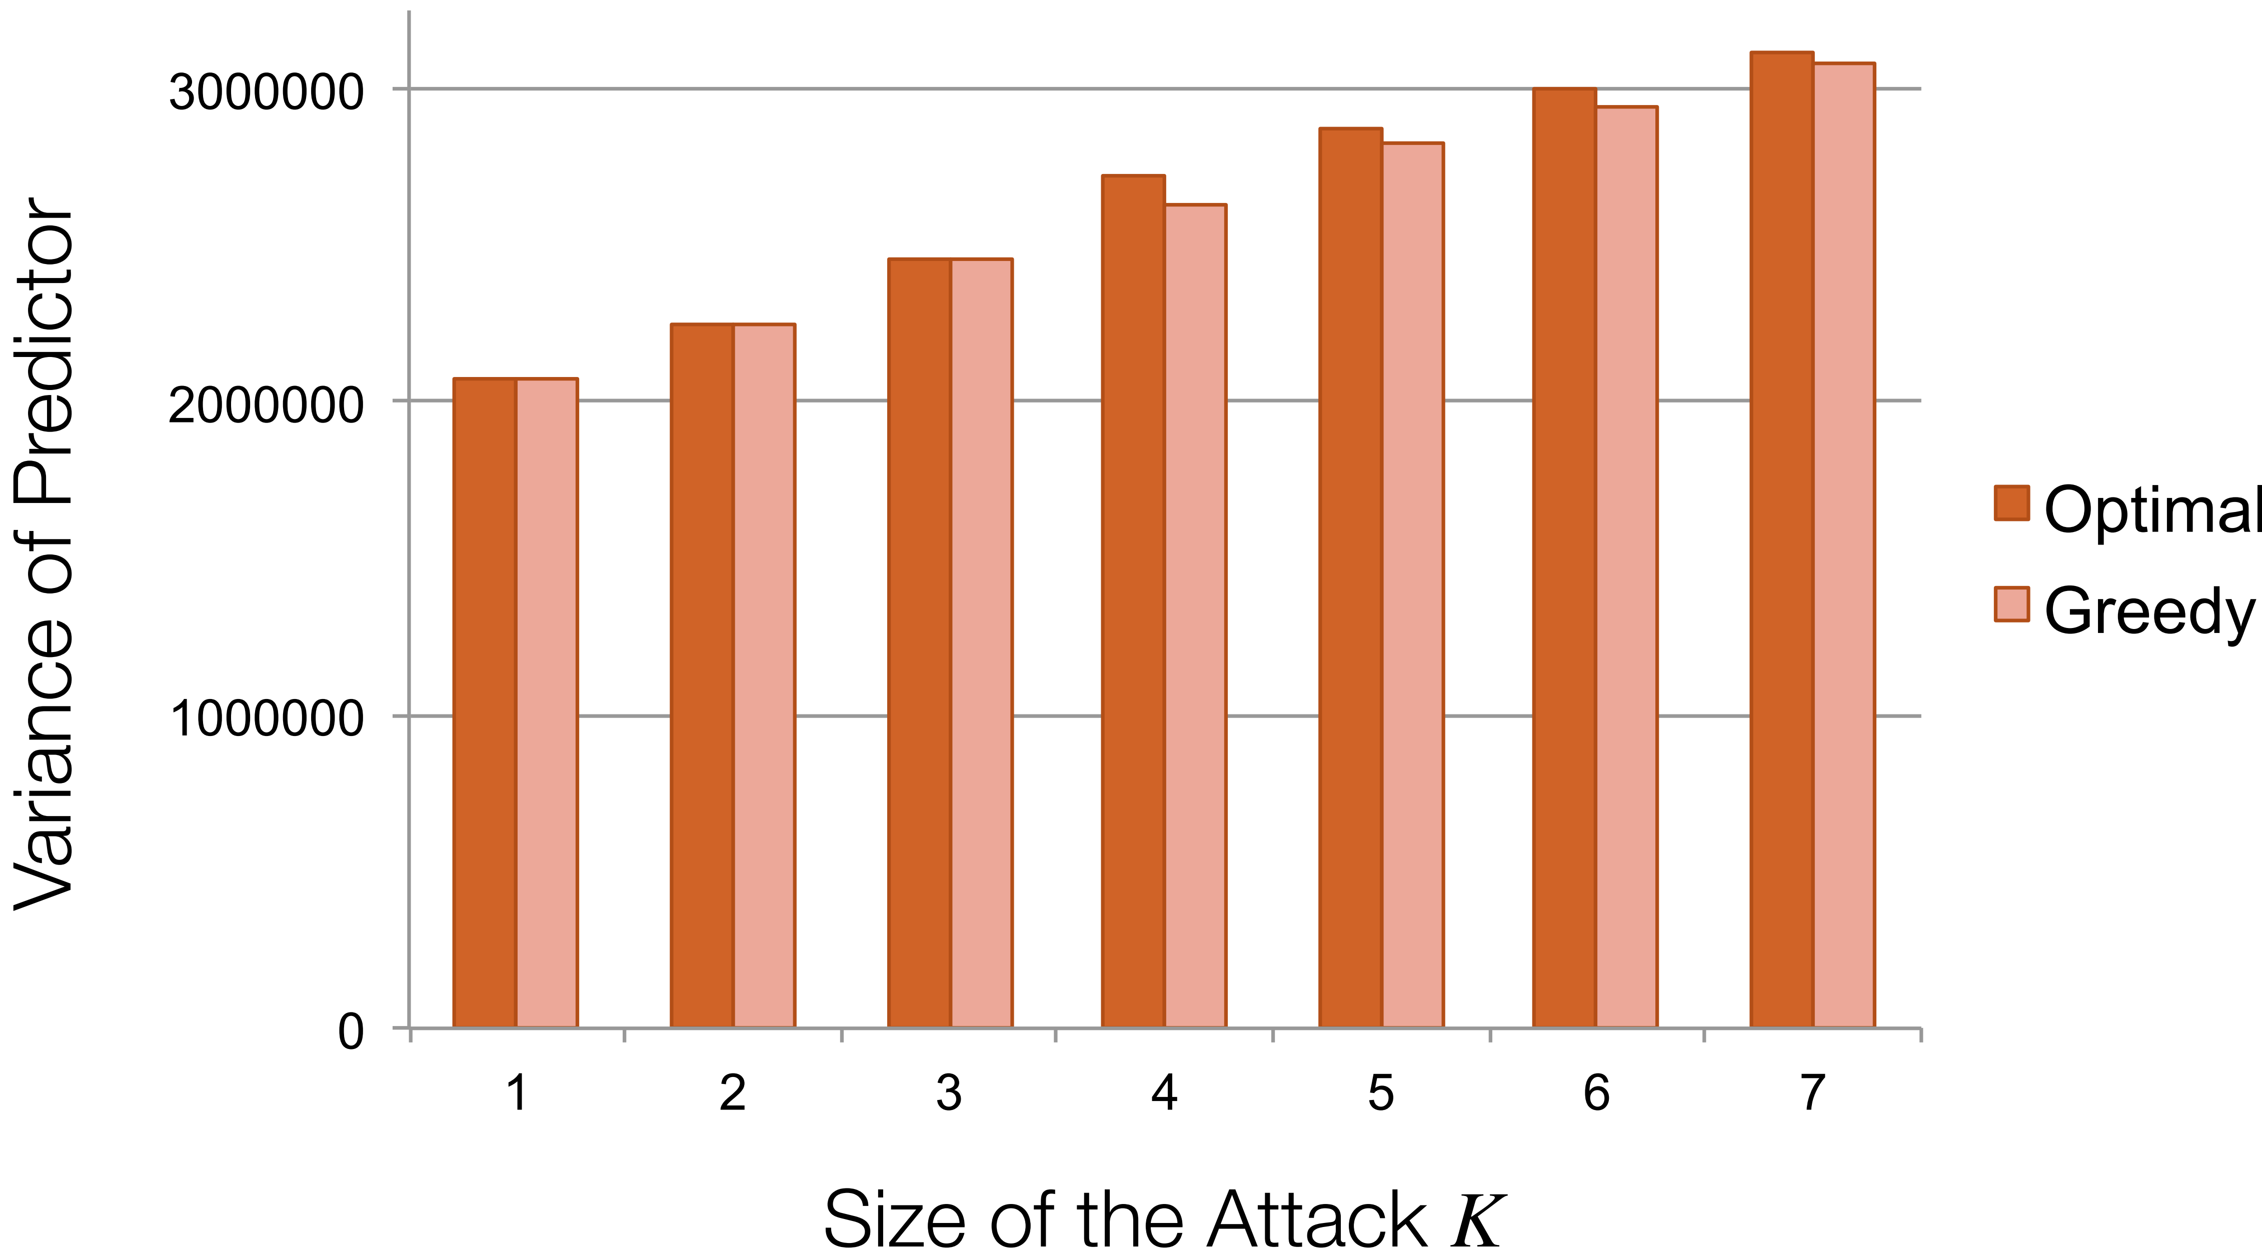

Greedy Algorithm for Resilient Sensor Placement

1. Let $\mathcal{S} \leftarrow \varnothing$
2. While $|\mathcal{S}| < K+1$:
   1. Let $X \notin \mathcal{S}$ be a variable minimizing $\sigma^2_{Z|\{X\}}$
   2. Let $\mathcal{S} \leftarrow \mathcal{S} \cup \{X\}$
3. While $|\mathcal{S}| < N$:
   1. Let $X \notin \mathcal{S}$ be a variable minimizing $Obj(\mathcal{S} \cup \{X\})$
   2. Let $\mathcal{S} \leftarrow \mathcal{S} \cup \{X\}$
4. Return $\mathcal{S}$

- Approximation guarantee:

Theorem: Let OPT be the difference between the prior variance and the optimal posterior. Then, the difference for the greedy selection is at least

$$OPT - OPT \cdot \left(1 - \frac{\gamma_{N,K}}{N}\right)^{N-K}$$

Greedy Sensor Placement: Numerical Results #1

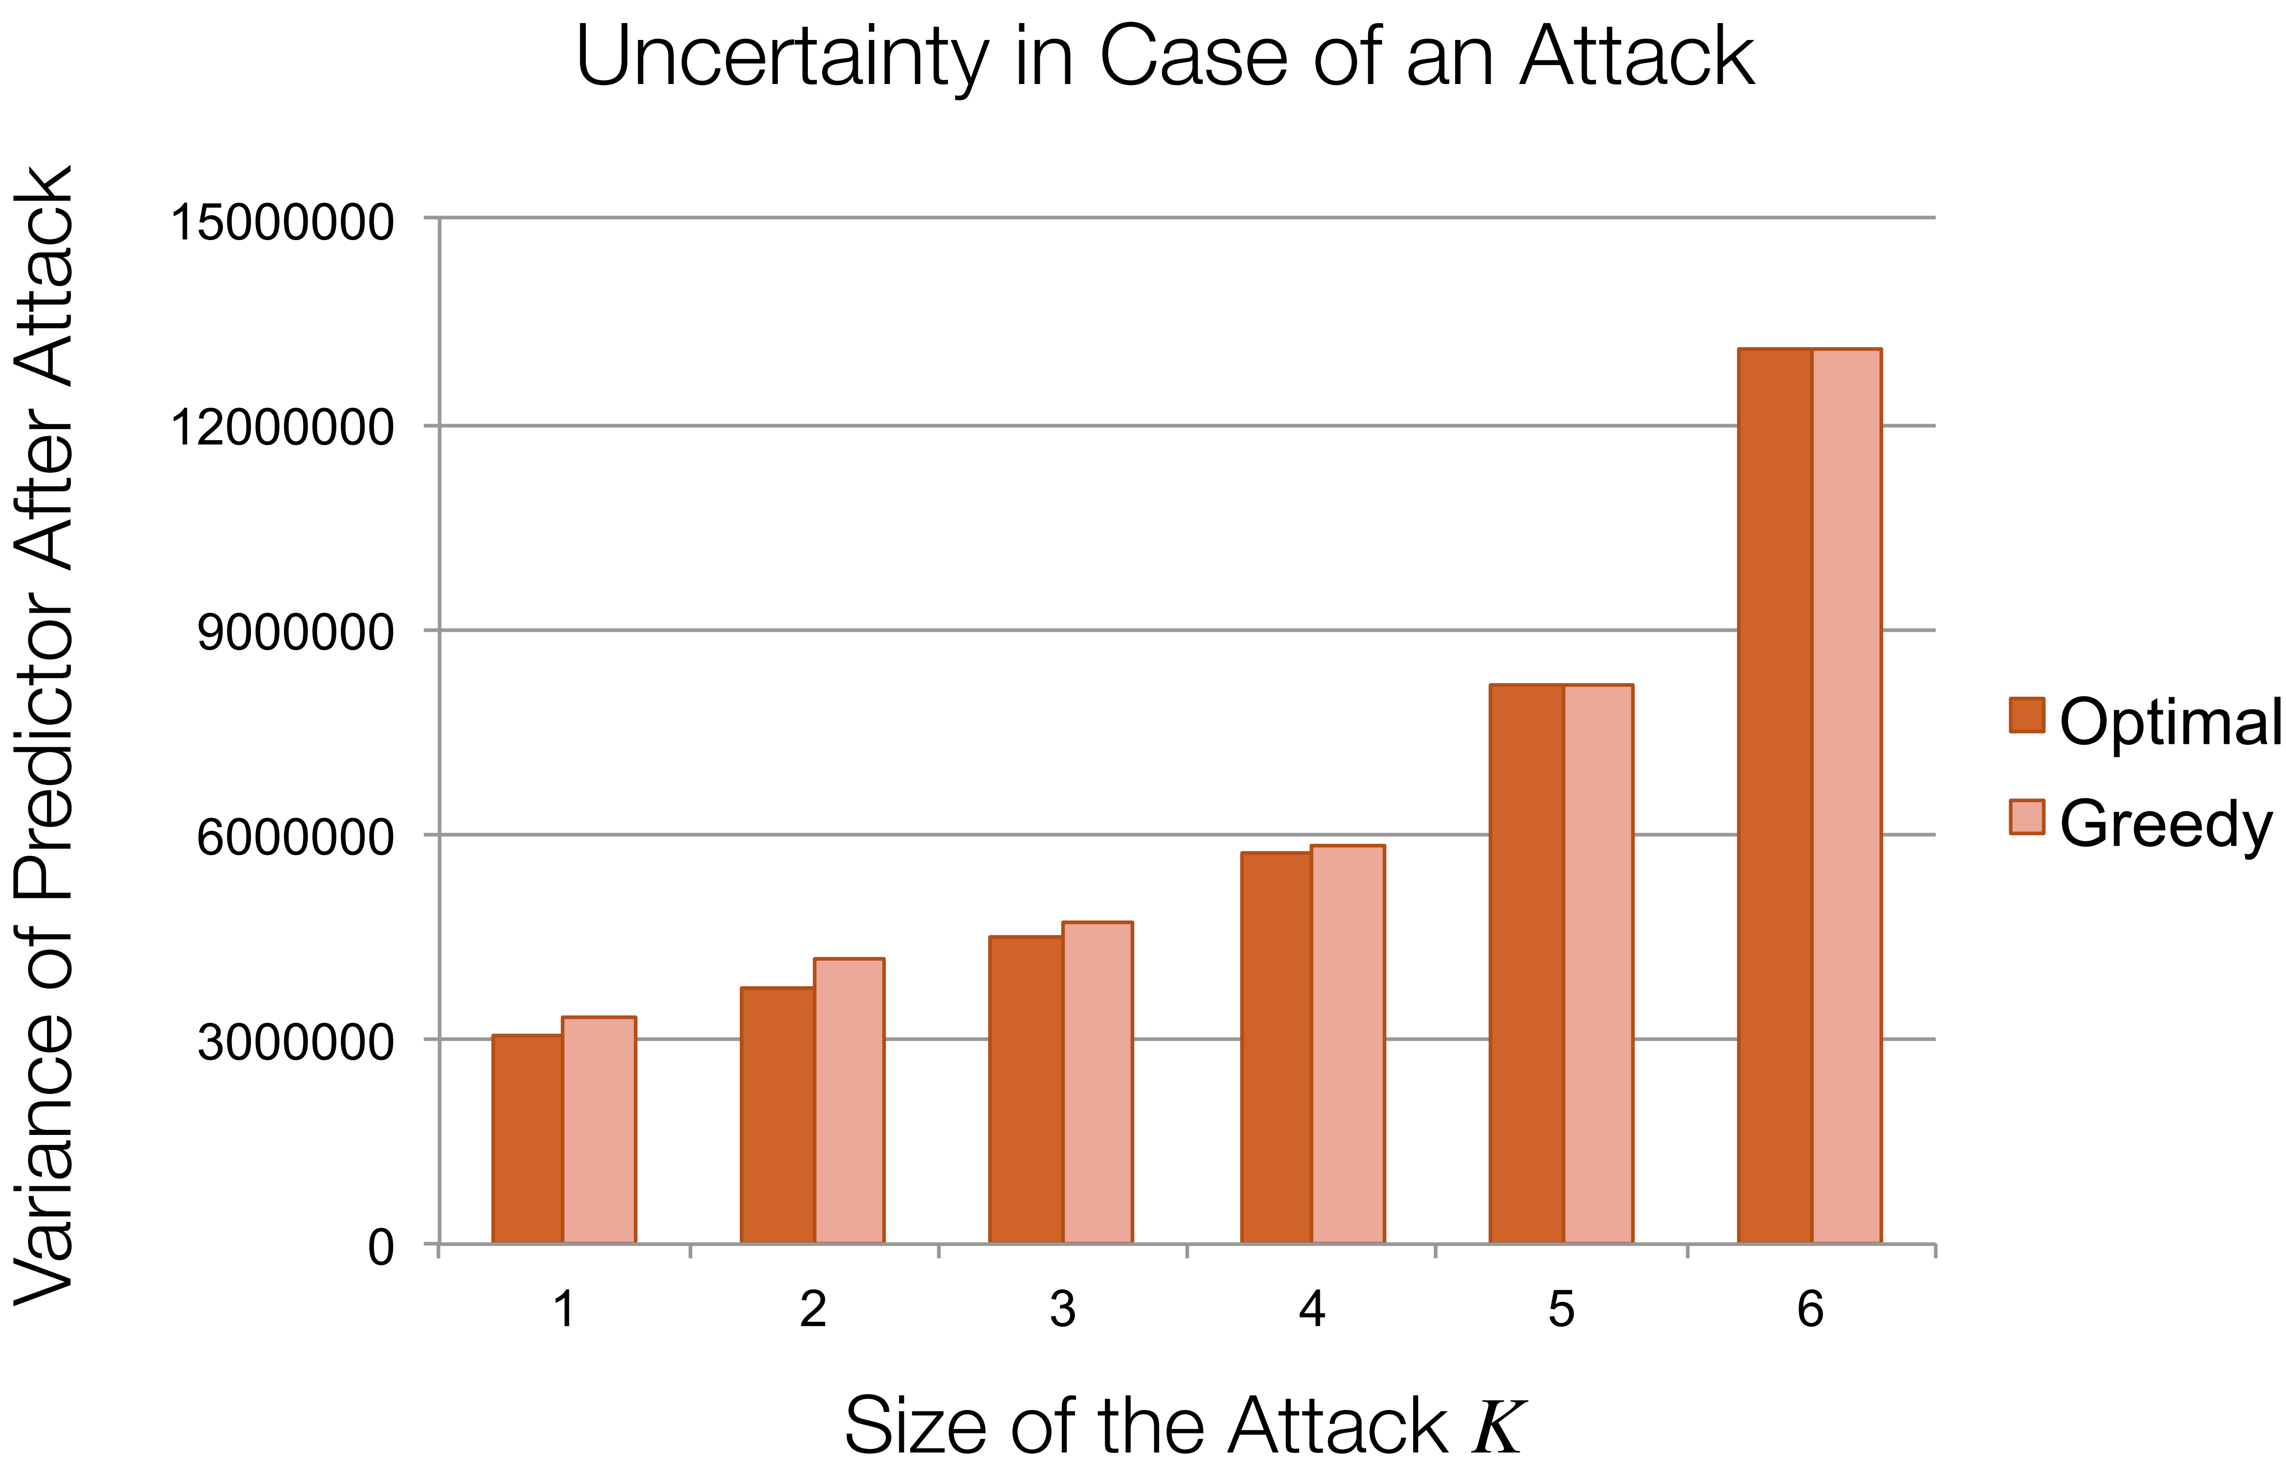

Greedy Sensor Placement: Numerical Results #2

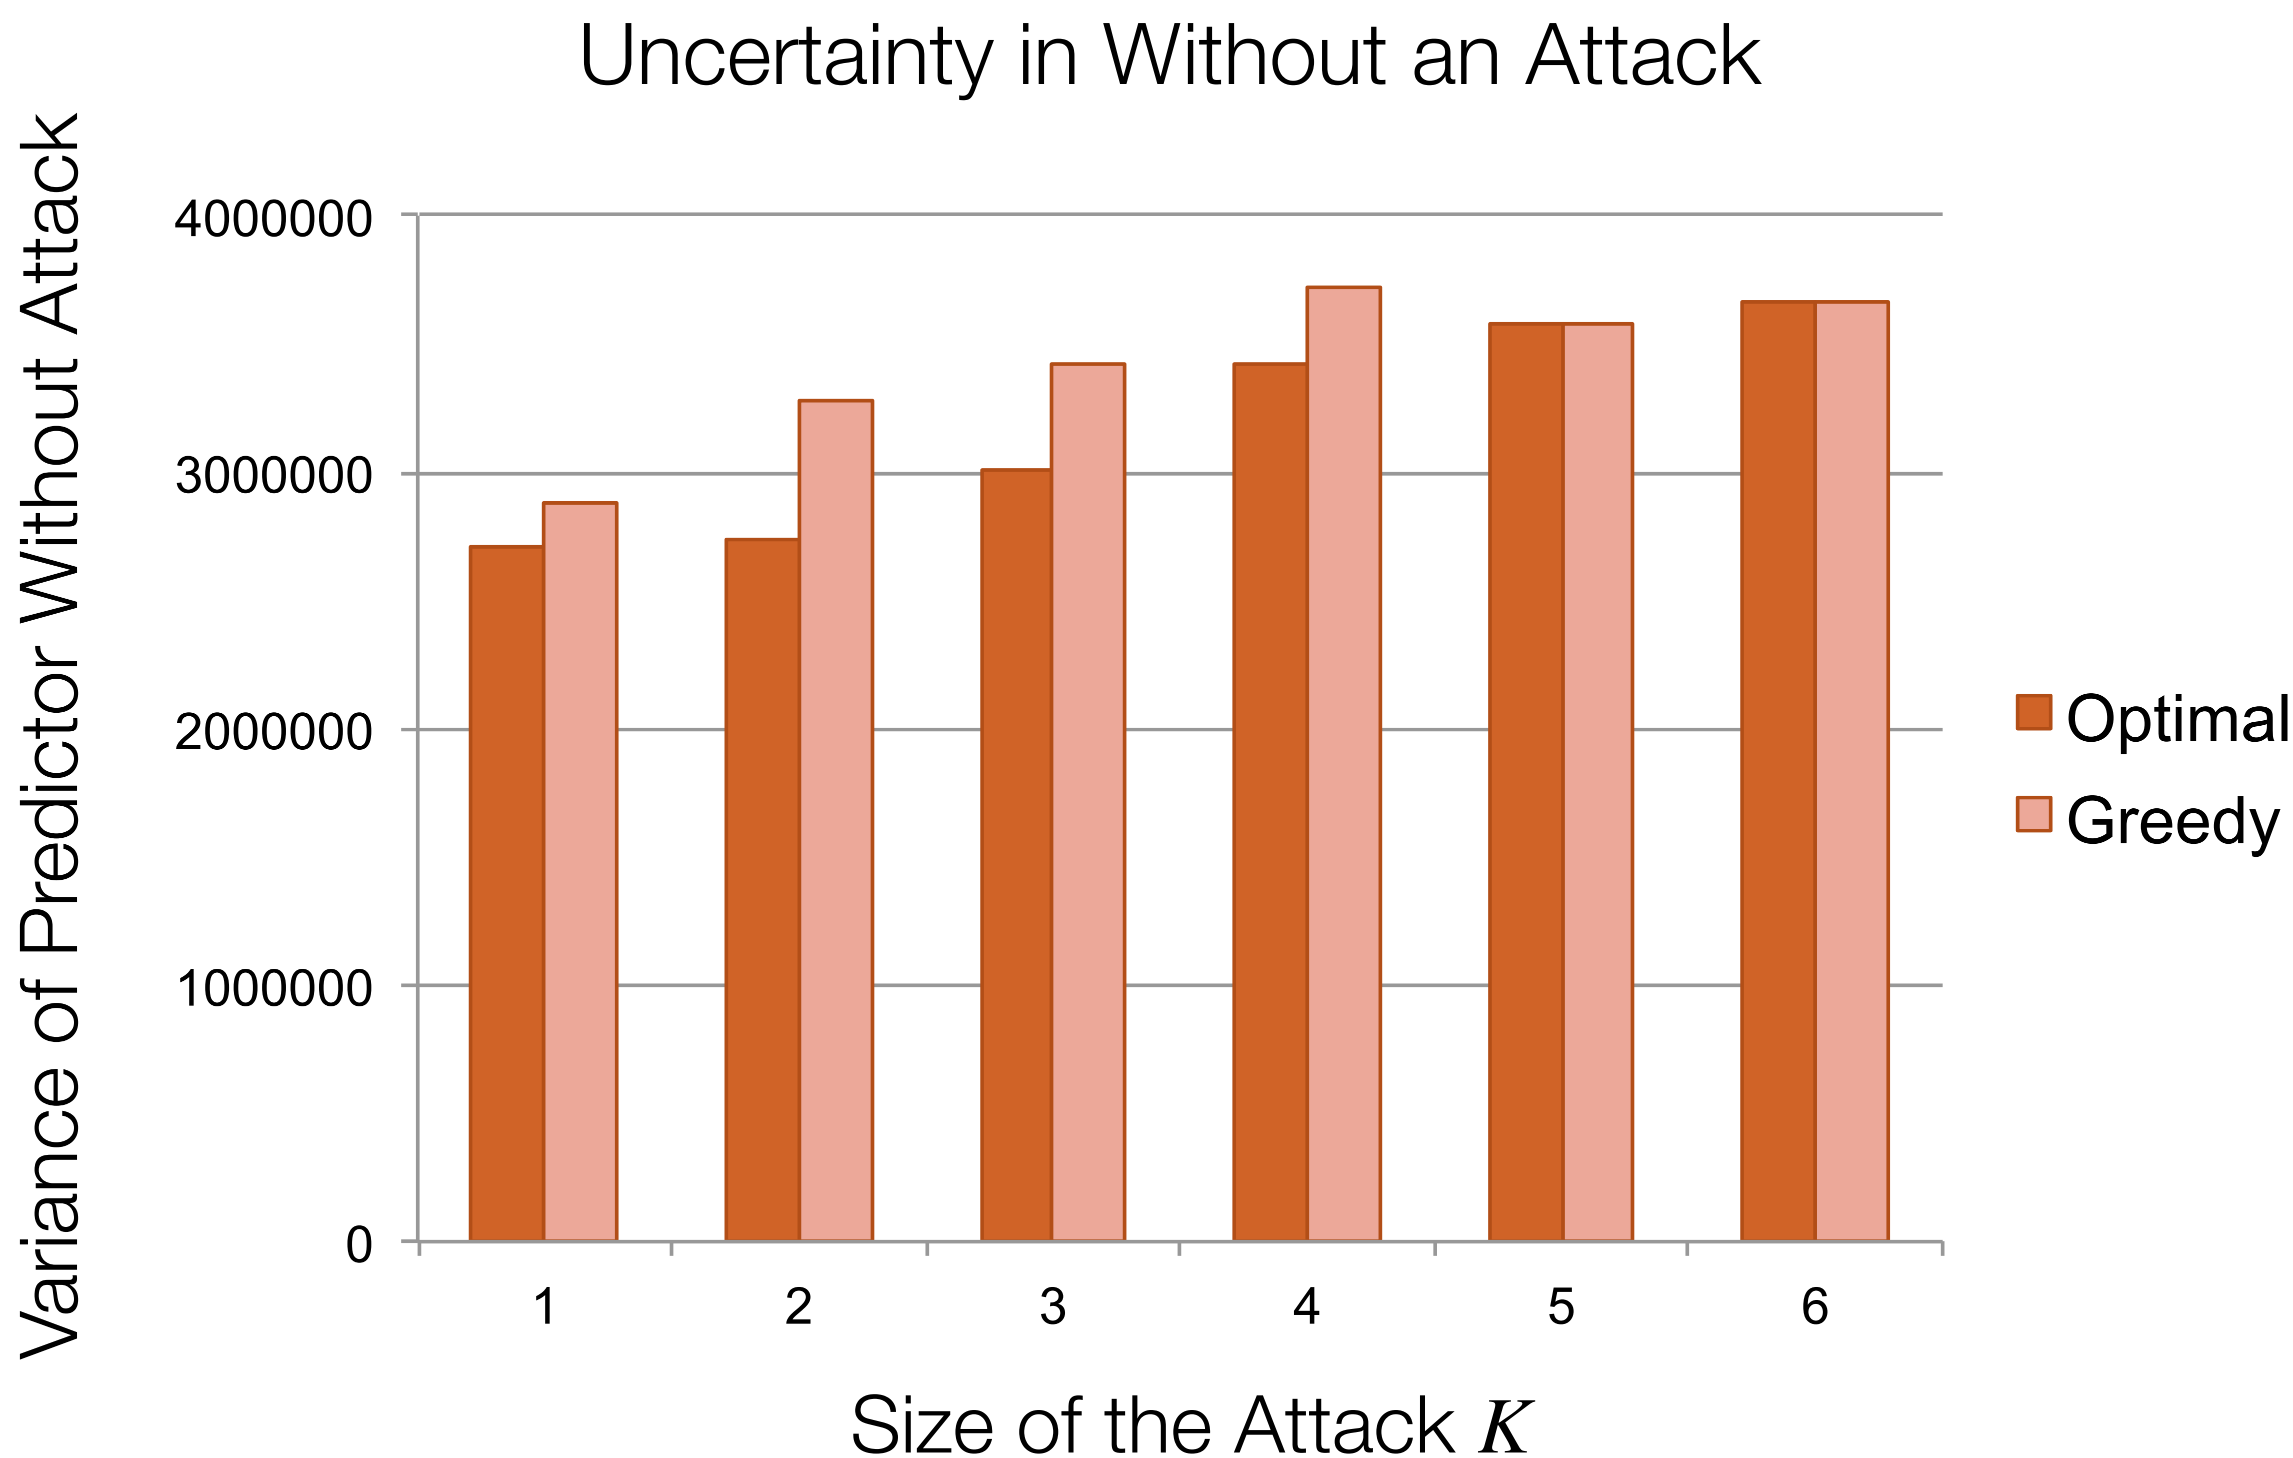

Special Case: Tree Covariance Graphs

- Covariance graph
  - vertices: variables representing observations at potential sensor locations
  - edges: non-negligible covariance values between variables
- Special case: tree covariance graphs

Lemma: Greedy attack is always optimal.

Theorem: Optimal resilient sensor placement can be found in polynomial time using a bottom-up dynamic programming approach.

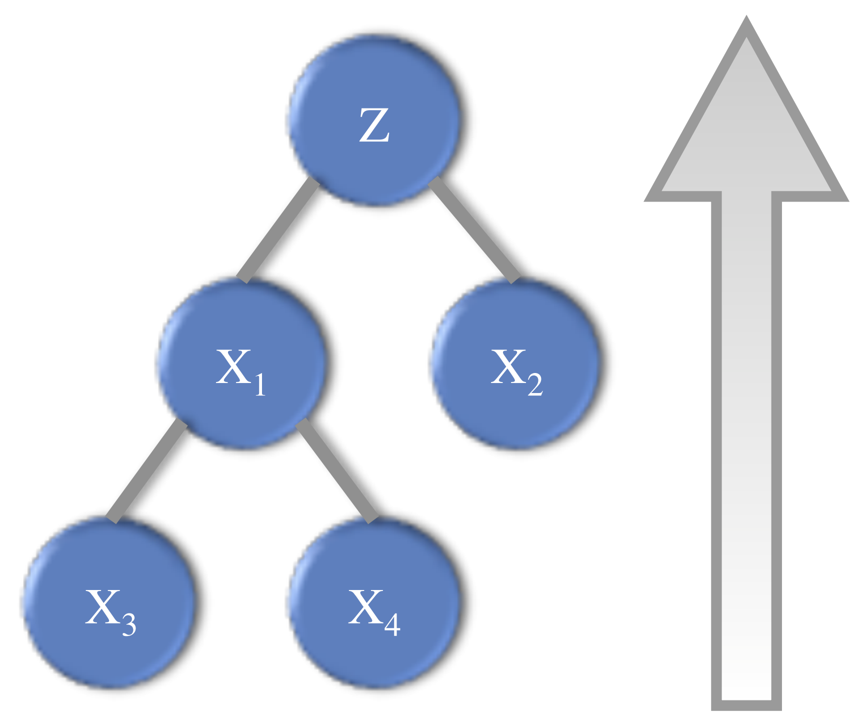

Conclusion

- We formulated a resilient sensor placement problem based on Gaussian-process based regression
- Using numerical results based on real-world data, we demonstrated that we can increase resilience significantly
- We proposed heuristic and approximation algorithms, and an optimal algorithm for a special case
- Open problems and future work
  - resilience to tampering attacks
  - arbitrary trade-off between resilience and accuracy without attack

Thank you for your attention!

Questions?