Differentially Private Resource Allocation

Joann Qiongna Chen[1], Tianhao Wang[2], Zhikun Zhang[3,5], Yang Zhang[3], Somesh Jha[4], Zhou Li[1]

[1]UC Irvine [2]University of Virginia [3]CISPA [4]UW-Madison. [5]Stanford University

Motivation: Unintentional Information Leakage

Storage controllers

Network rate limiters

Messengers

Outline

- Threat model
- Possible solutions: AKR
- Our solution by precise modeling
- Evaluation

Outline

- Threat model
- Possible solutions: AKR
- Our solution by precise modeling
- Evaluation

Threat Model: Assumptions

- Available resources (public) are allowed to be less than the number of requests.
- The request sender is aware of whether their requests are being fulfilled.
- A Resource Allocator (RA) is able to work fairly without seeing user identity.

Threat model

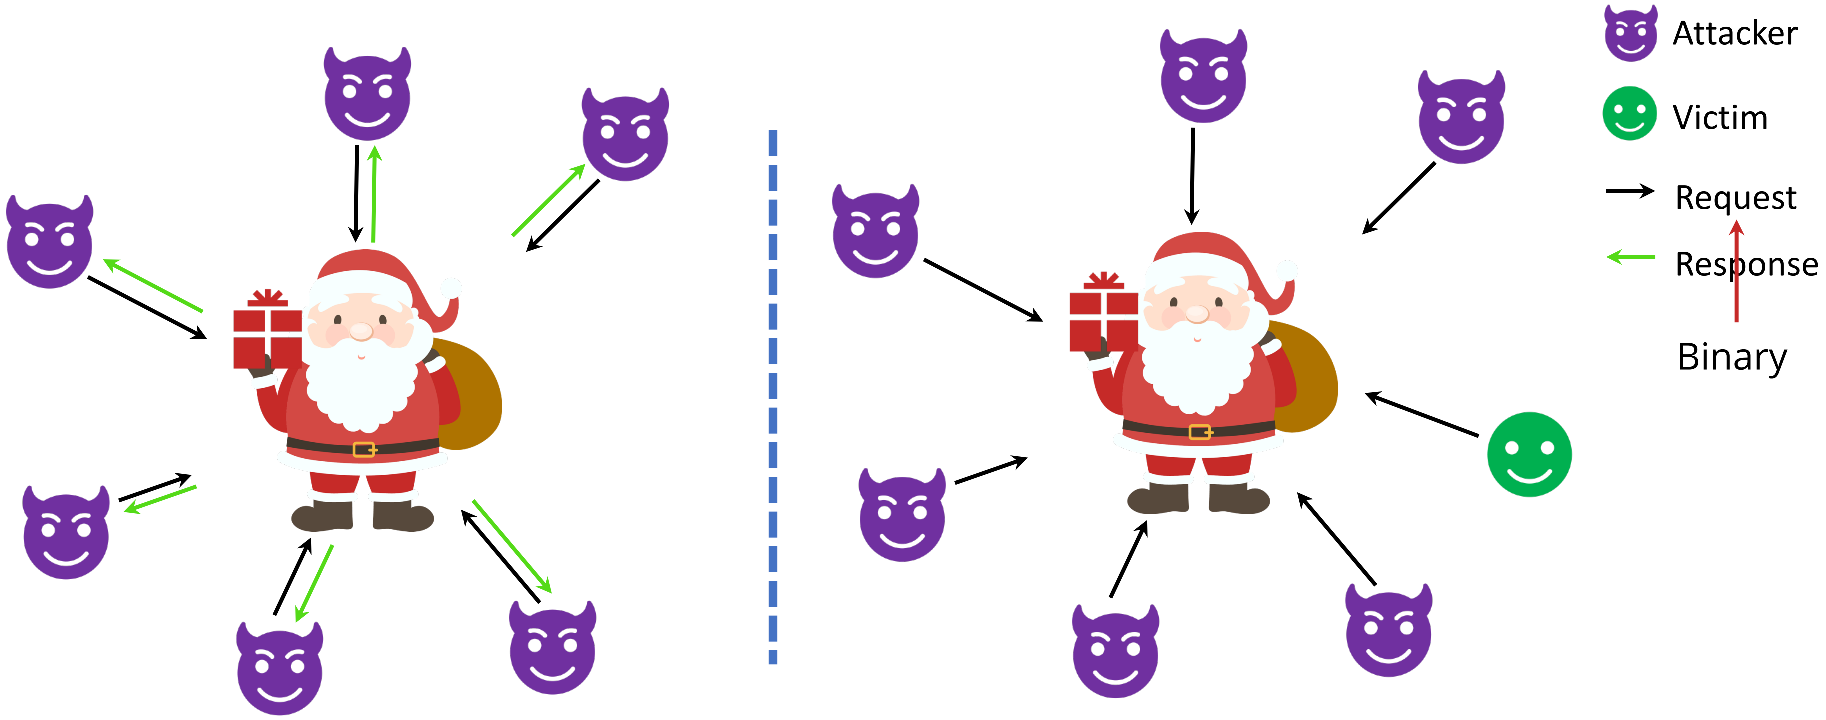

Attacker gets all resources

Threat model

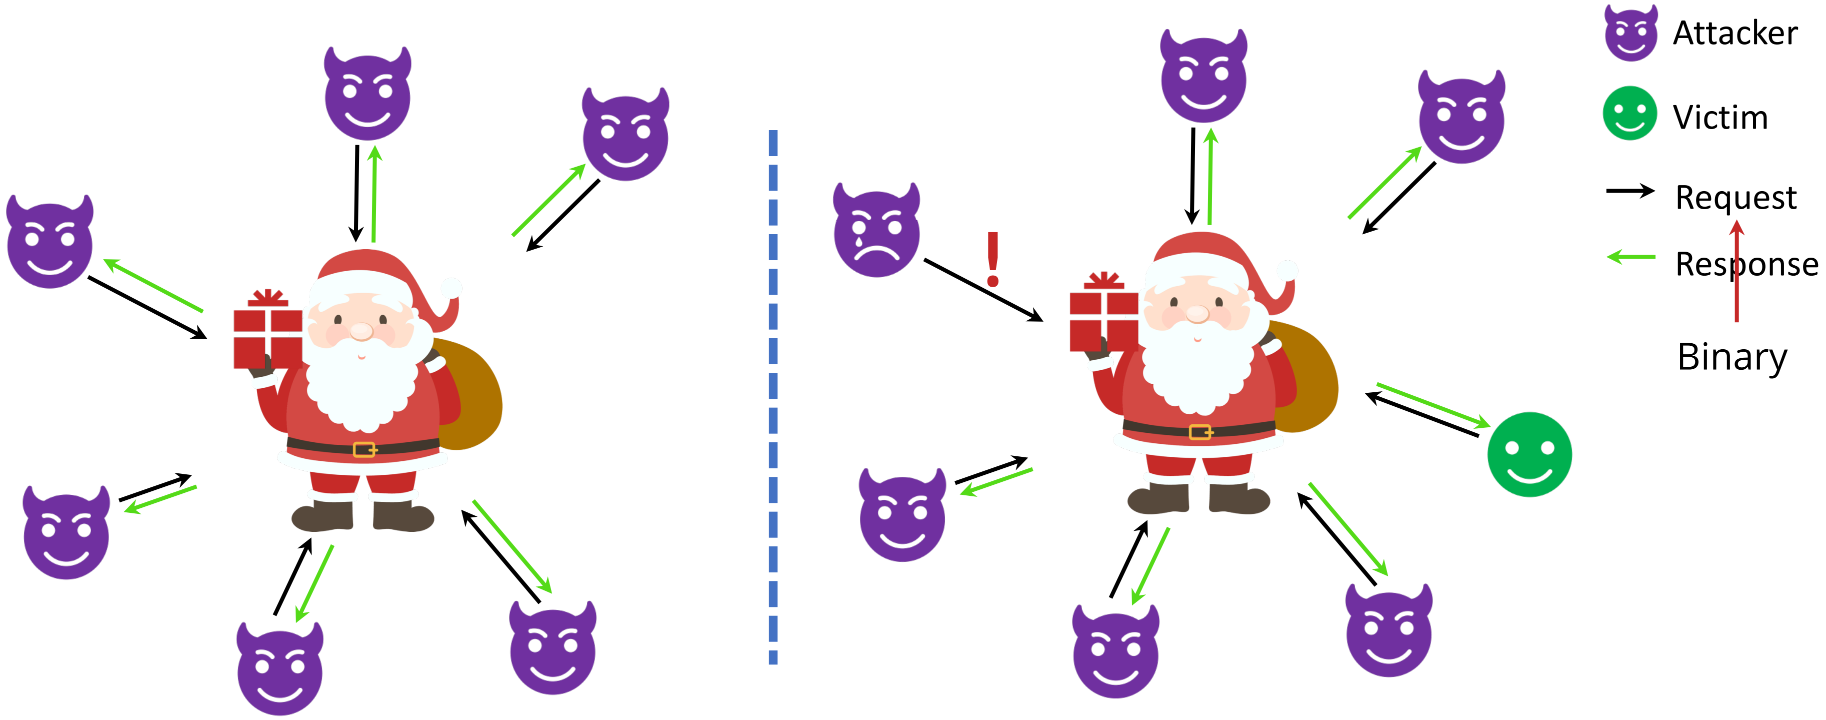

Attacker gets all resources

Attacker gets less resources than requested

Threat model

Attacker's goal:
To learn the existence of the victim

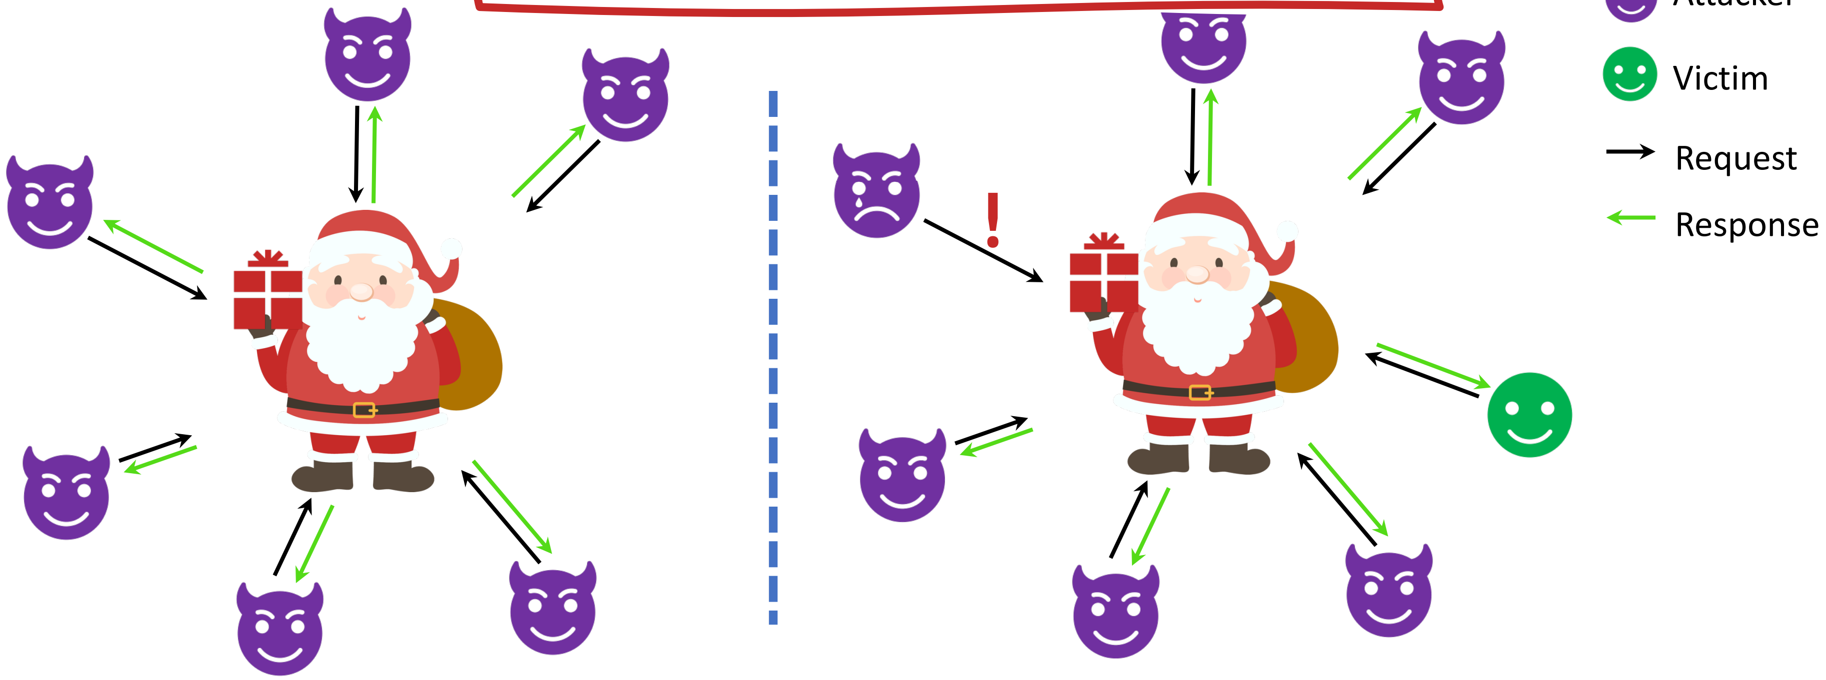

Attacker gets all resources

Attacker gets less resources than requested

Threat model

Attacker's goal:
To learn the existence of the victim

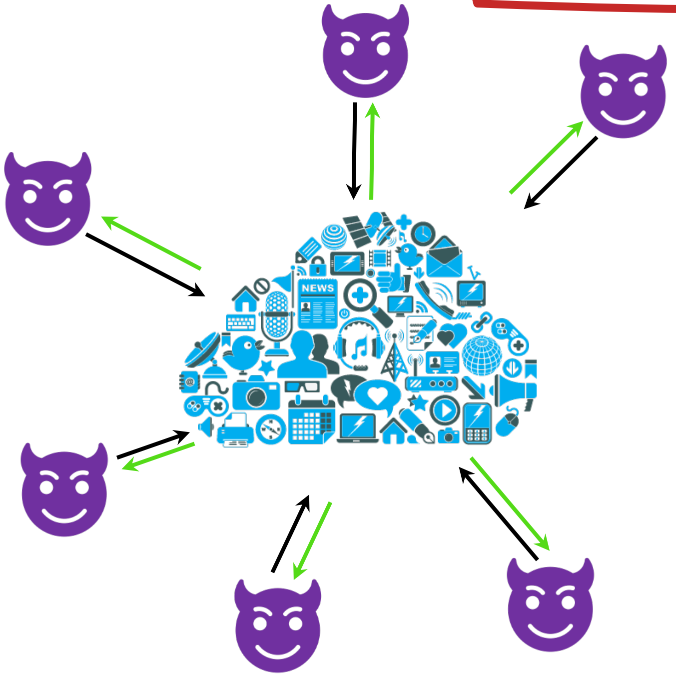

Attacker gets all resources

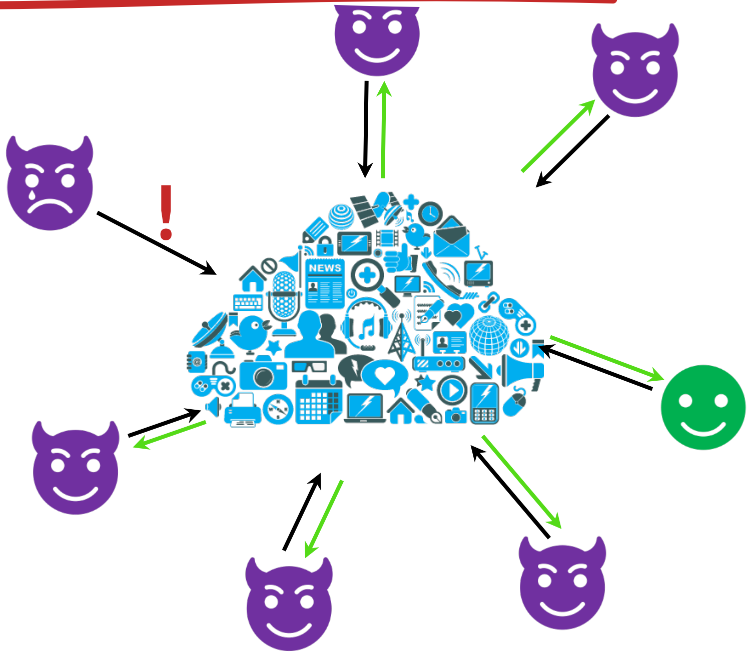

Attacker gets less resources than requested

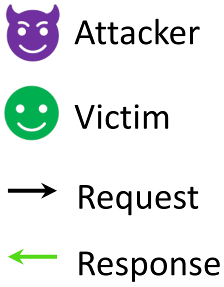

Threat model

Attacker's goal:
To learn the existence of the victim

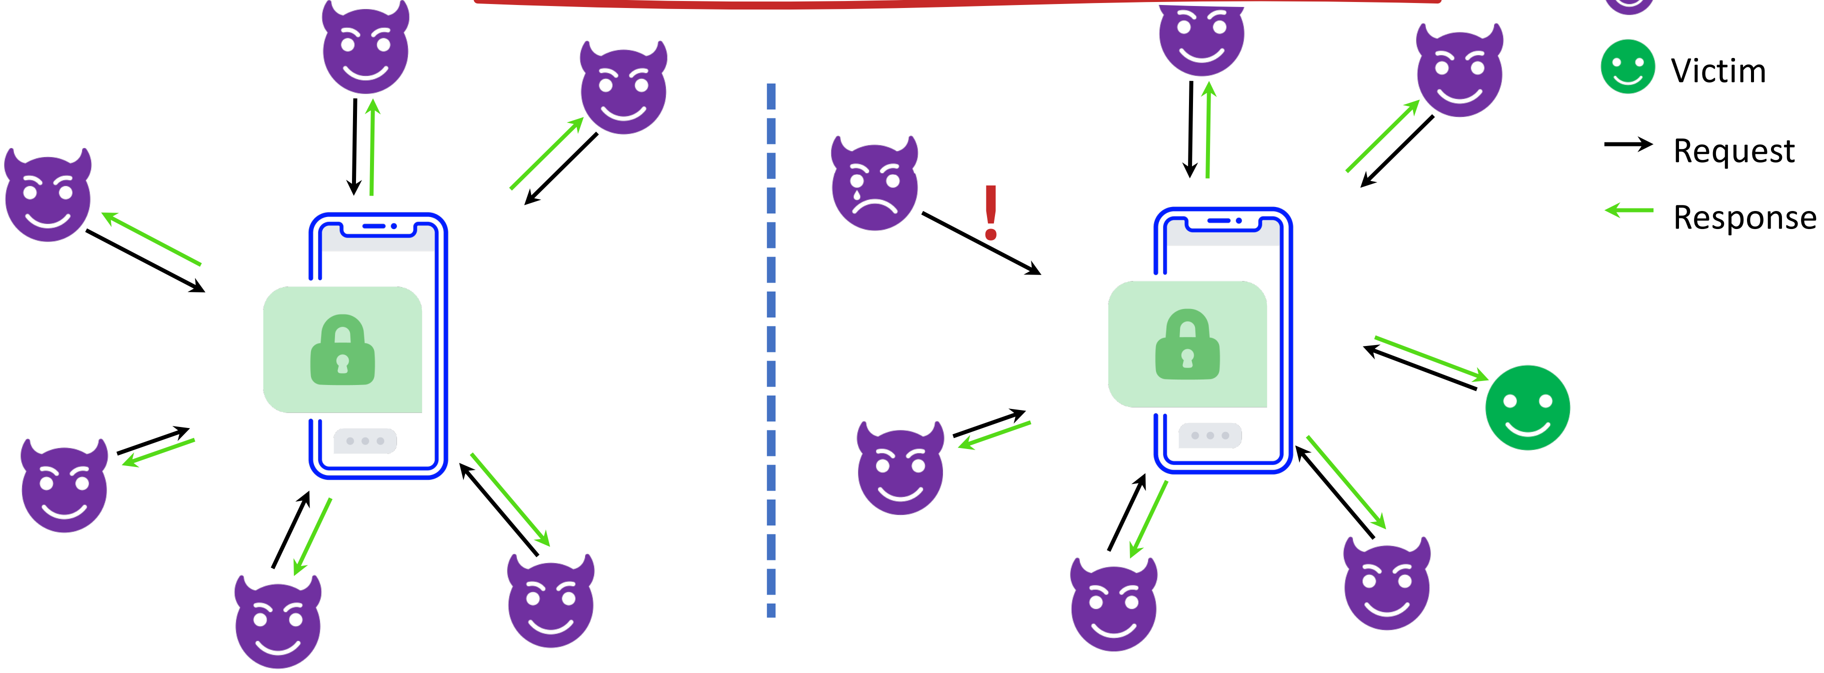

Attacker gets all resources

Attacker gets less resources than requested

Threat model

Attacker's goal:
To learn the existence of the victim

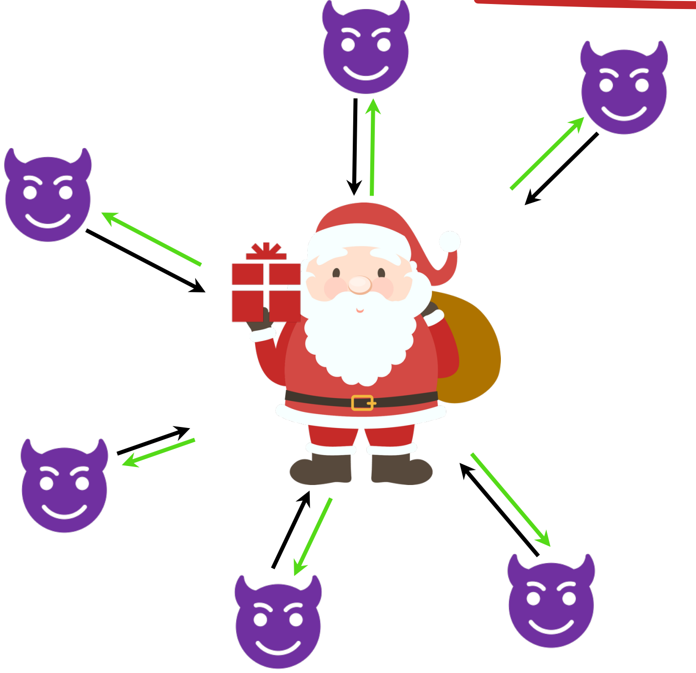

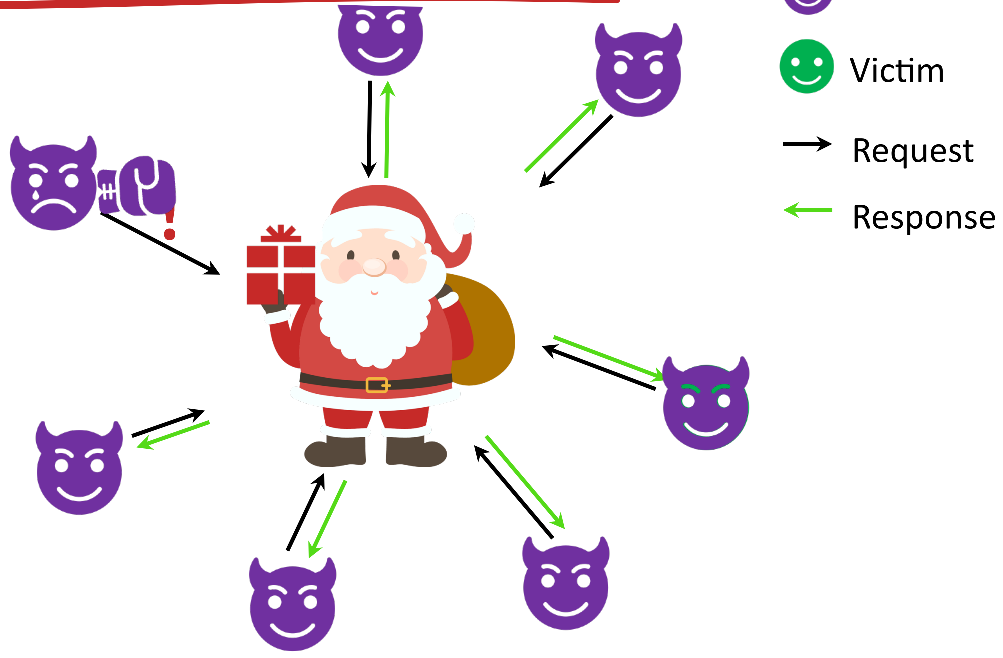

Attacker gets all resources

Attacker gets less resources than requested

Threat model

Attacker's goal:
To learn the existence of the victim

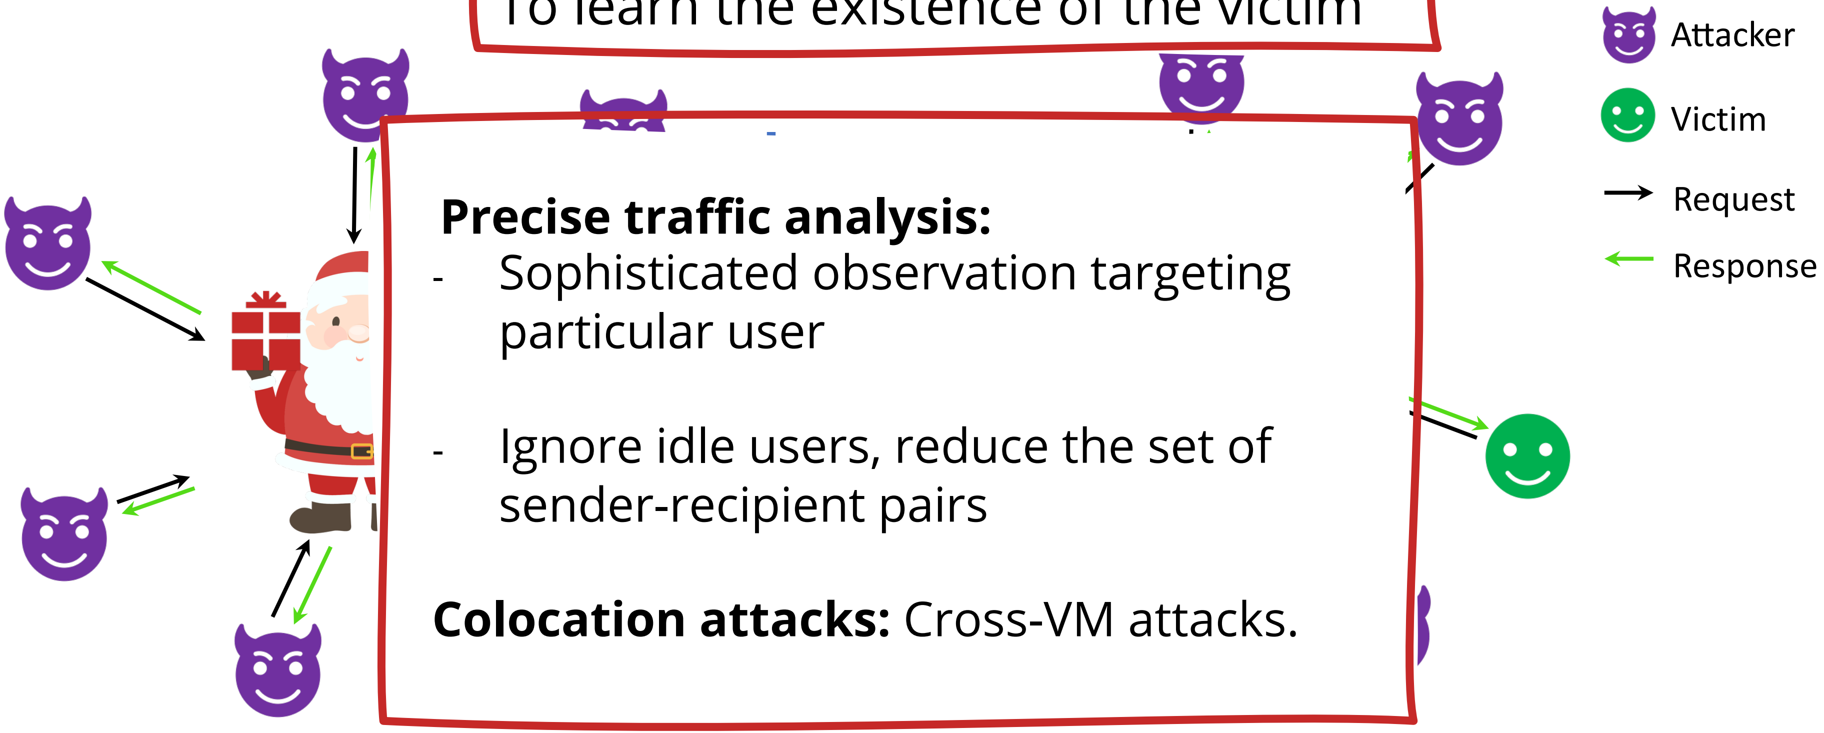

Precise traffic analysis:

- Sophisticated observation targeting particular user
- Ignore idle users, reduce the set of sender-recipient pairs

Colocation attacks: Cross-VM attacks.

Attacker gets all resources

Attacker gets less resources than requested

Overview

- Threat model
- Possible solutions: AKR
- Our solution by precise modeling
- Simulation Results

Possible Solutions (AKR): Private Resource Allocators and Their Applications

1. Slot-based resource allocator
2. Randomized resource allocator
3. Differentially private resource allocator (DPRA)

[AKR20] Sebastian Angel, Sampath Kannan, and Zachary Ratliff. Private resource allocators and their applications. IEEE S&P (Oakland), 2020.

Possible Solution(AKR): Private Resource Allocators and Their Applications

1. Slot-based resource allocator
2. Randomized resource allocator

★ **Differentially private resource allocator (DPRA)**

[AKR20] Sebastian Angel, Sampath Kannan, and Zachary Ratliff. Private resource allocators and their applications. IEEE S&P (Oakland), 2020.

AKR: Differentially Private Resource Allocator

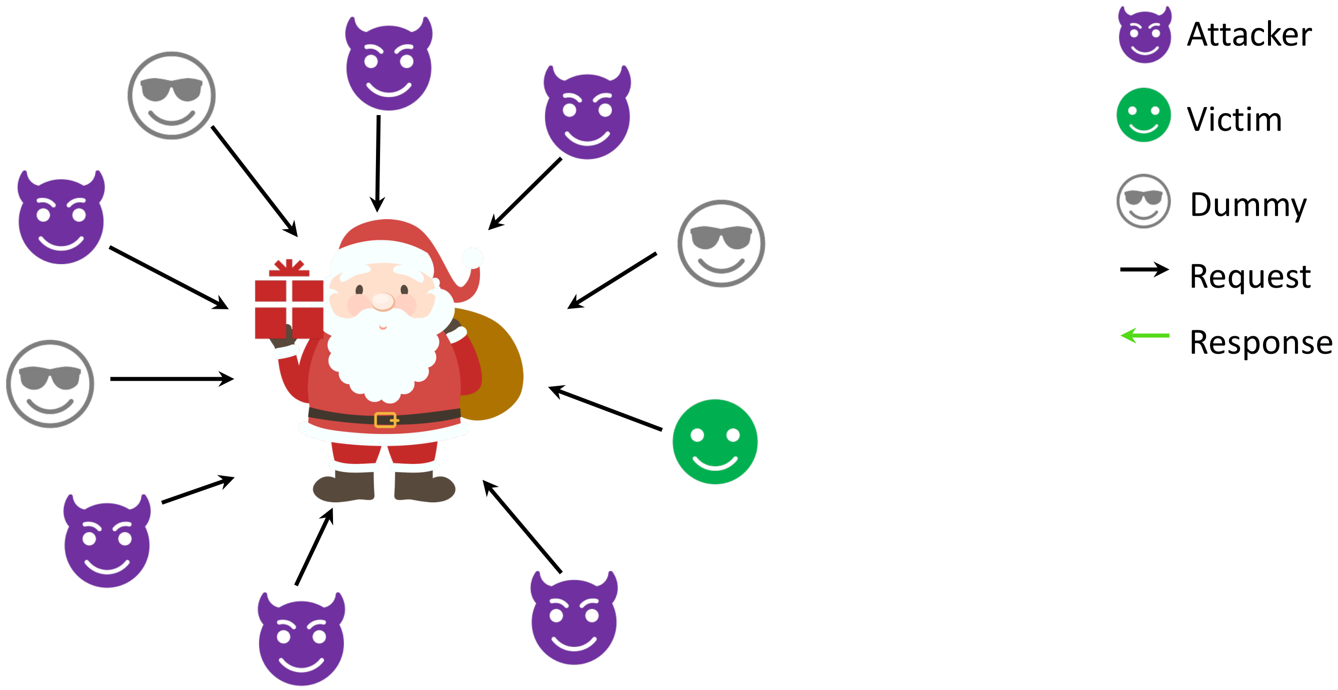

AKR: Differentially Private Resource Allocator

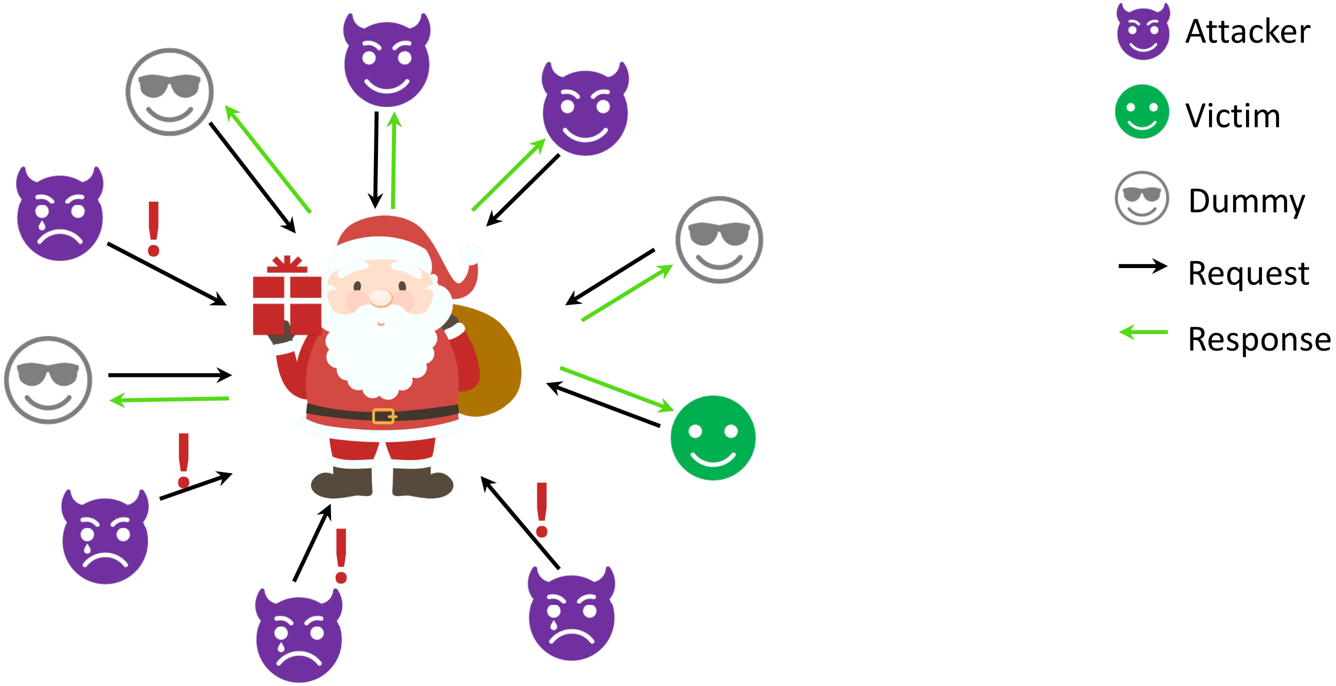

(ϵ, δ)-Differential Privacy

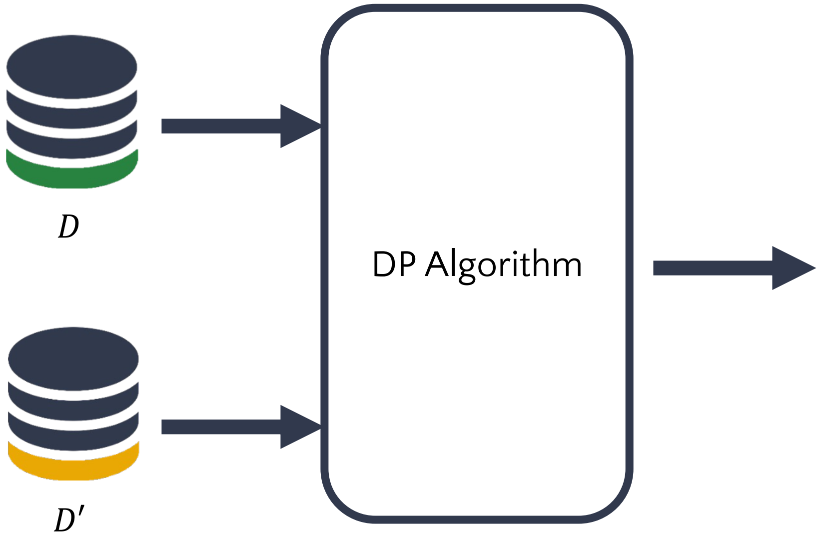

Privacy Budget/ Privacy Loss

Relaxation

$$\Pr[\mathrm{A}(D) \in Y] \leq e^{\epsilon} \Pr\left[\mathrm{A}(D') \in Y\right] + \delta$$

Laplace Mechanism

Sensitivity: Maximum change in output caused by the input dataset

D

Query answer $q(D)$

$$A(D) = q(D) + \mathrm{Lap}(S/\epsilon)$$

$$\Pr[\mathrm{A}(D) \in Y] \leq e^{\epsilon} \Pr[\mathrm{A}(D') \in Y]$$

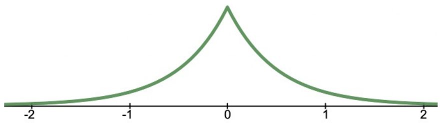

D'

Query answer $q(D')$

$$A(D') = q(D') + \mathrm{Lap}(S/\epsilon)$$

AKR: Laplace Mechanism

⚠️ Number of dummies should be a positive integer. Rounding up the Laplace noise introduces δ

D

Query answer $q(D)$

$A(D) = q(D) + \max(\mathbf{0}, \lceil \mathrm{Lap}(S/\epsilon) \rceil)$

$$\Pr[\mathrm{A}(D) \in Y] \leq e^{\epsilon} \Pr\left[\mathrm{A}(D') \in Y\right] + \delta$$

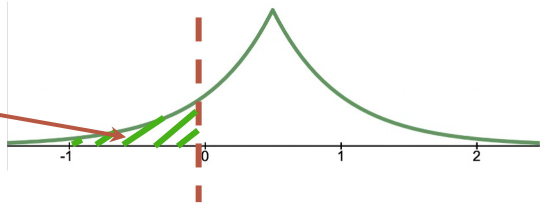

D'

Query answer $q(D')$

$A(D') = q(D') + \max(\mathbf{0}, \lceil \mathrm{Lap}(S/\epsilon) \rceil)$

Privacy Goal: small δ

AKR: Laplace Mechanism

D

Query answer $q(D)$

$A(D) = q(D) + \max(\mathbf{0}, [\mathrm{Lap}(S/\epsilon)])$

$$\Pr[\mathrm{A}(D) \in Y] \leq e^{\epsilon} \Pr\left[\mathrm{A}(D') \in Y\right] + \delta$$

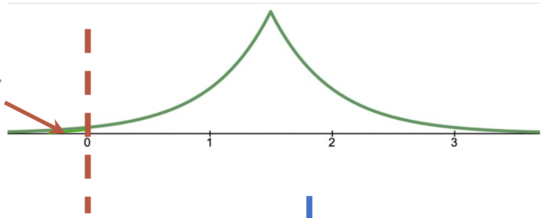

D'

Query answer $q(D')$

$A(D') = q(D') + \max(\mathbf{0}, [\mathrm{Lap}(S/\epsilon)])$

Problem: large bias

AKR: Differentially Private Resource Allocator

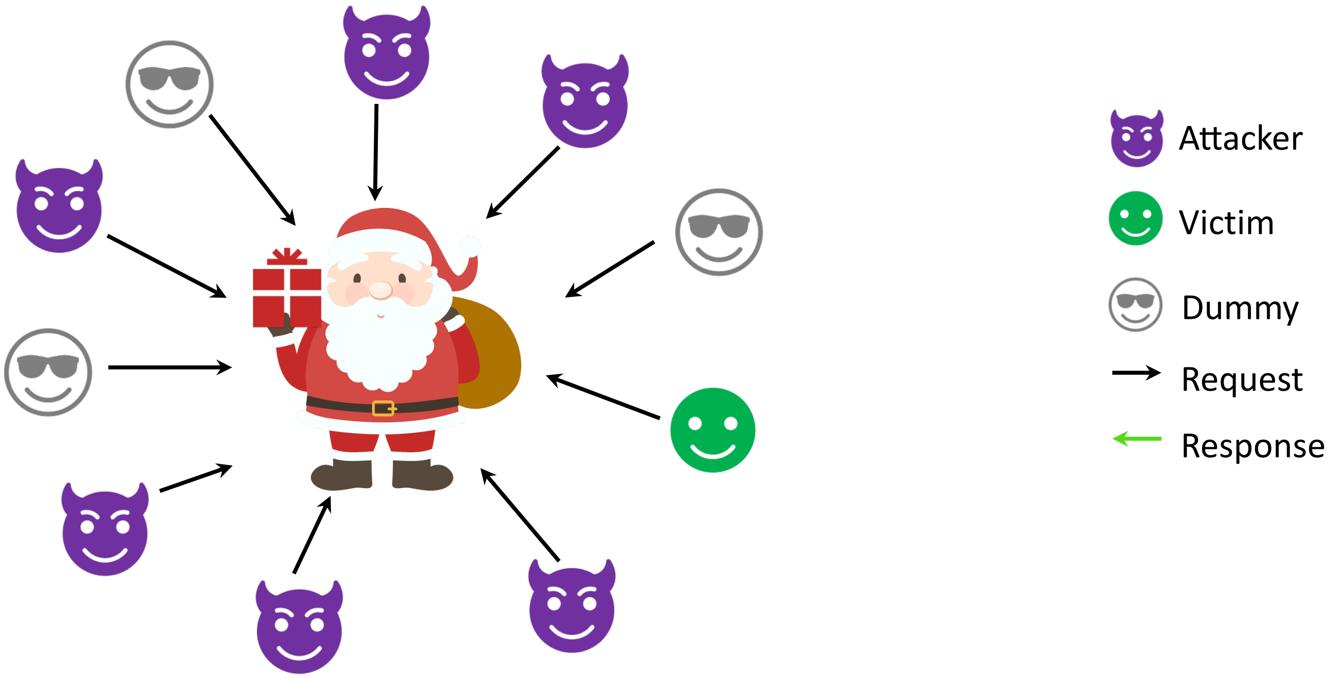

Outline

- General threat model
- Possible solutions: AKR
- Our solution by precise modeling
- Simulation Results

AKR: Differentially Private Resource Allocator

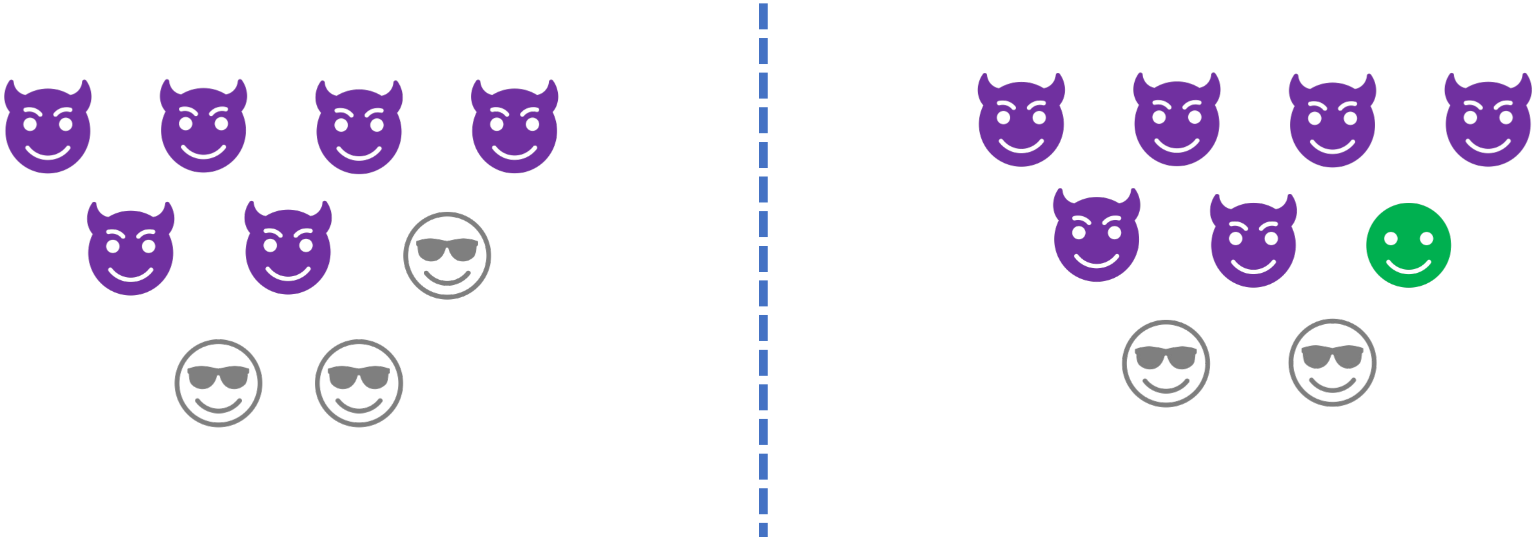

$$\Pr[\mathrm{A}(D) \in Y] \leq e^{\epsilon} \Pr\left[\mathrm{A}(D') \in Y\right] + \delta$$

Ours: Differentially Private Resource Allocator

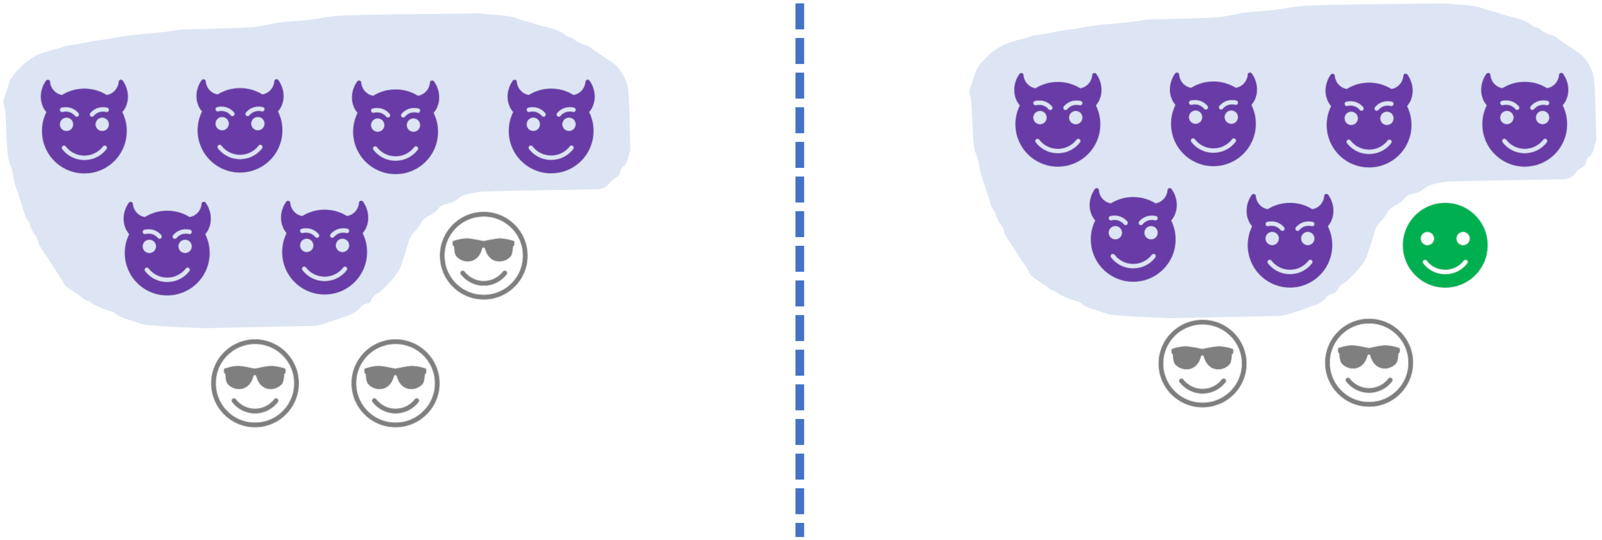

$$\Pr[\mathrm{A}(D) \in Y] \leq e^{\epsilon} \Pr\left[\mathrm{A}(D') \in Y\right]$$

Ours: Differentially Private Resource Allocator

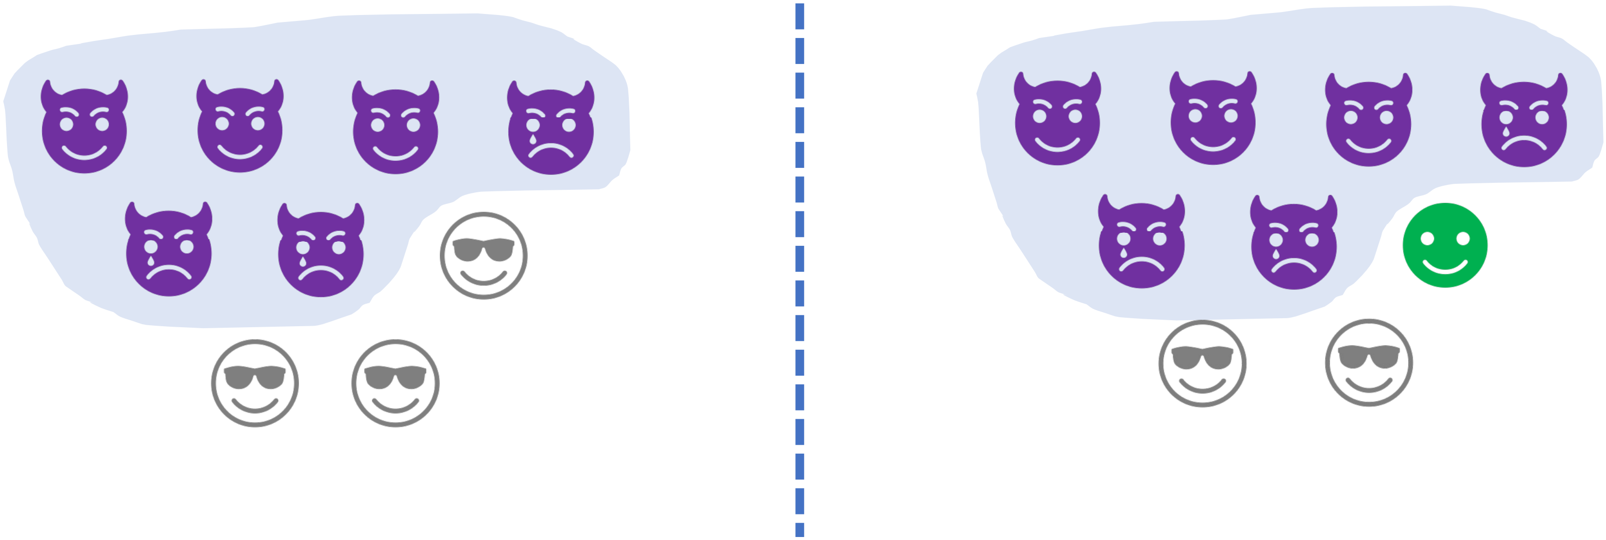

$$\Pr[\mathrm{A}(D) \in Y] \leq e^{\epsilon} \Pr\left[\mathrm{A}(D') \in Y\right]$$

Privacy Amplification

The attacker has only a limited view of the resource allocator

Overall View

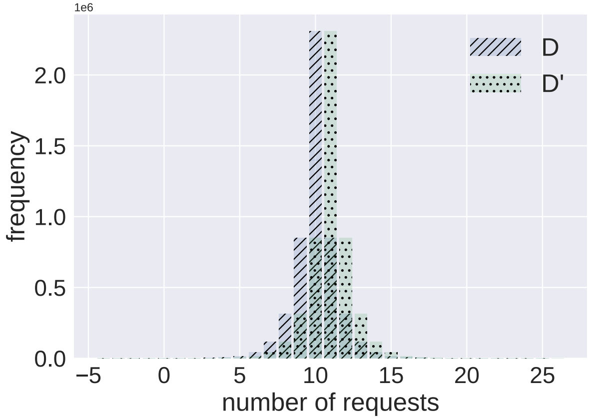

Number of total requests during a round of allocation

Attacker’s View

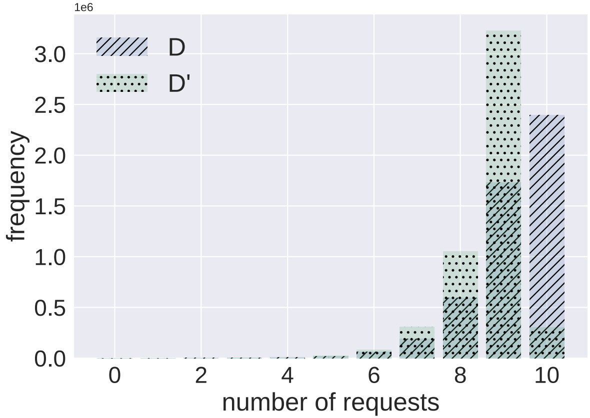

Number of attacker’s fulfilled requests

Privacy Modeling

Traditional DP
$$\frac{\Pr[A(D) = y]}{\Pr[A(D') = y]} = \frac{\Pr[q(D) + \ Lap\left(\frac{S}{\epsilon}\right)]}{\Pr[q(D') + \ Lap\left(\frac{S}{\epsilon}\right)]} \leq e^{\epsilon}$$

Conditional probability of output y

PMF of noise distribution

Precise Modeling
$$\frac{\Pr[\mathrm{View}^{\mathcal{A}}_{\mathcal{M}}(D) = y]}{\Pr[\mathrm{View}^{\mathcal{A}}_{\mathcal{M}}(D) = y]} = \frac{\sum_{i=x_l}^{x_r} \Pr[d = i] \Pr[y|\,|D| + d]}{\sum_{i=x_l}^{x_r} \Pr[d = i] \Pr[y|\,|D'| + d]} \leq e^{\epsilon}$$

Kearns M, Pai M, Roth A, Ullman J. Mechanism design in large games: Incentives and privacy. In Proceedings of the 5th conference on Innovations in theoretical computer science 2014

Our Mechanisms

- Constant Mechanism (CST)
- Uniform Mechanism (UNI)
- Geometric Mechanism (GEO)
- Double Geometric Mechanism (DGEO)

$$\frac{\Pr[\mathrm{View}_{\mathcal{M}}^{\mathcal{A}}(D) = y]}{\Pr[\mathrm{View}_{\mathcal{M}}^{\mathcal{A}}(D) = y]} = \frac{\sum_{i=x_l}^{x_r} \boxed{\Pr[d = i]} \Pr[y \mid |D| + d]}{\sum_{i=x_l}^{x_r} \boxed{\Pr[d = i]} \Pr[y \mid |D'| + d]}$$

Findings

- Precise modeling of resource allocation yields better utility-privacy tradeoff
- Constant noise can already satisfy DP when noise is greater than *k*
- In general, GEO has the best performance

	Privacy	Noise	Noise Sign	DP Condition	Utility (ϵ=0.65)	Utility (ϵ=1.7)	Utility (ϵ=2.3)
CST	ϵ-ADP	Constant	$+$	Noise $c \geq k$	0.50	-	-
UNI	ϵ-ADP	Discrete uniform	$+/-$	Right bound $x_n \geq k$	0.46	0.65	0.70
GEO	ϵ-ADP	One-sided geometric	$+/-$	-	0.47	0.82	0.90
DGEO	ϵ-ADP	Double geometric	$+/-$	-	0.44	0.77	0.98
AKR [2]	(ϵ, δ)-DP	Laplace	$+$	Bias $\mu = 1 - \ln(2\delta)/\epsilon$	0.32	0.53	0.59

A summary of different mechanisms and their utility under some representative ϵ values.

Outline

- Threat model
- Possible solutions: AKR
- Our solution by precise modeling
- Evaluation

Evaluation

Setup

- Following AKR's setting, we set resource capacity $k = 10$ for most of our simulations
- Metrics
  - Privacy (ϵ) is measured by the DP guarantee
  - Utility: percentage of resources allocated to legitimate requests
- Each simulation consists of millions of rounds

Evaluation

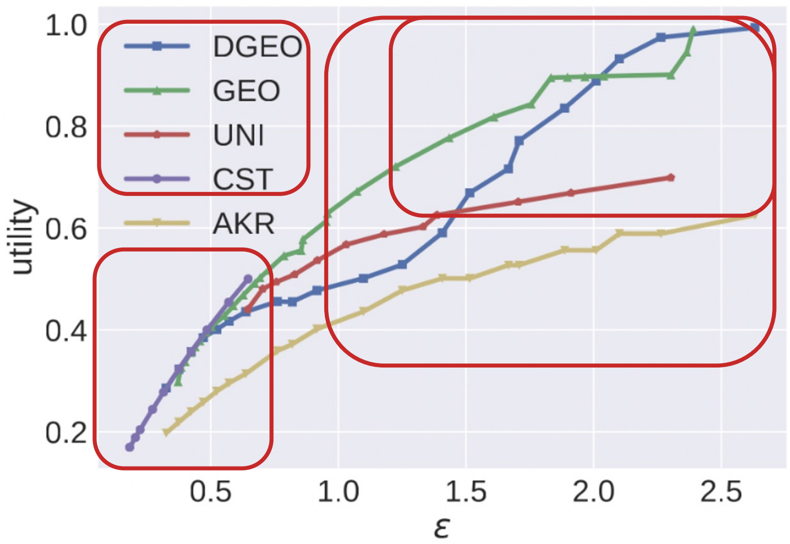

- Utility of of constant mechanism cannot exceed 50%
- DGEO, GEO lead in privacy-utility trade-off, especially when ϵ is large
- Precise privacy modeling improves the privacy-utility trade-off

Conclusion

1. We conduct a rigorous privacy analysis of differentially private resource allocators.
   - Tighter privacy bounds
   - The attacker's view
   - Four noisy mechanisms
2. We theoretically and empirically evaluate our proposed mechanisms.
   - Our mechanism GEO leads to the best privacy-utility tradeoff and outperforms AKR by a large margin
   - Constant noise can already satisfy DP when noise is greater than *k, though the* utility cannot exceed 50%
3. Our code is available at https://github.com/dpra-dp/dpra

Question

Joann Qiongna Chen
(**on the academic job market**)

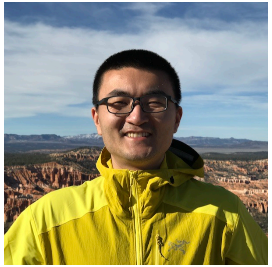

Tianhao Wang

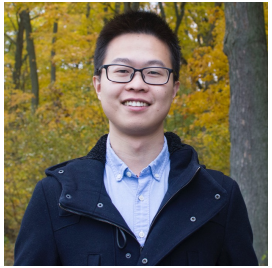

Zhikun Zhang

Yang Zhang

Somesh Jha

Zhou Li