

Security and Privacy for Distributed Optimization and Learning

Nitin Vaidya
Georgetown University



5-Hour Tutorial

disc.georgetown.domains → Talks

Secret to happiness is to
lower your expectations to the point
where they're already met

- **Hobbes** (paraphrased)



Goals

- g Background
- g Problem formulation
- g Intuition



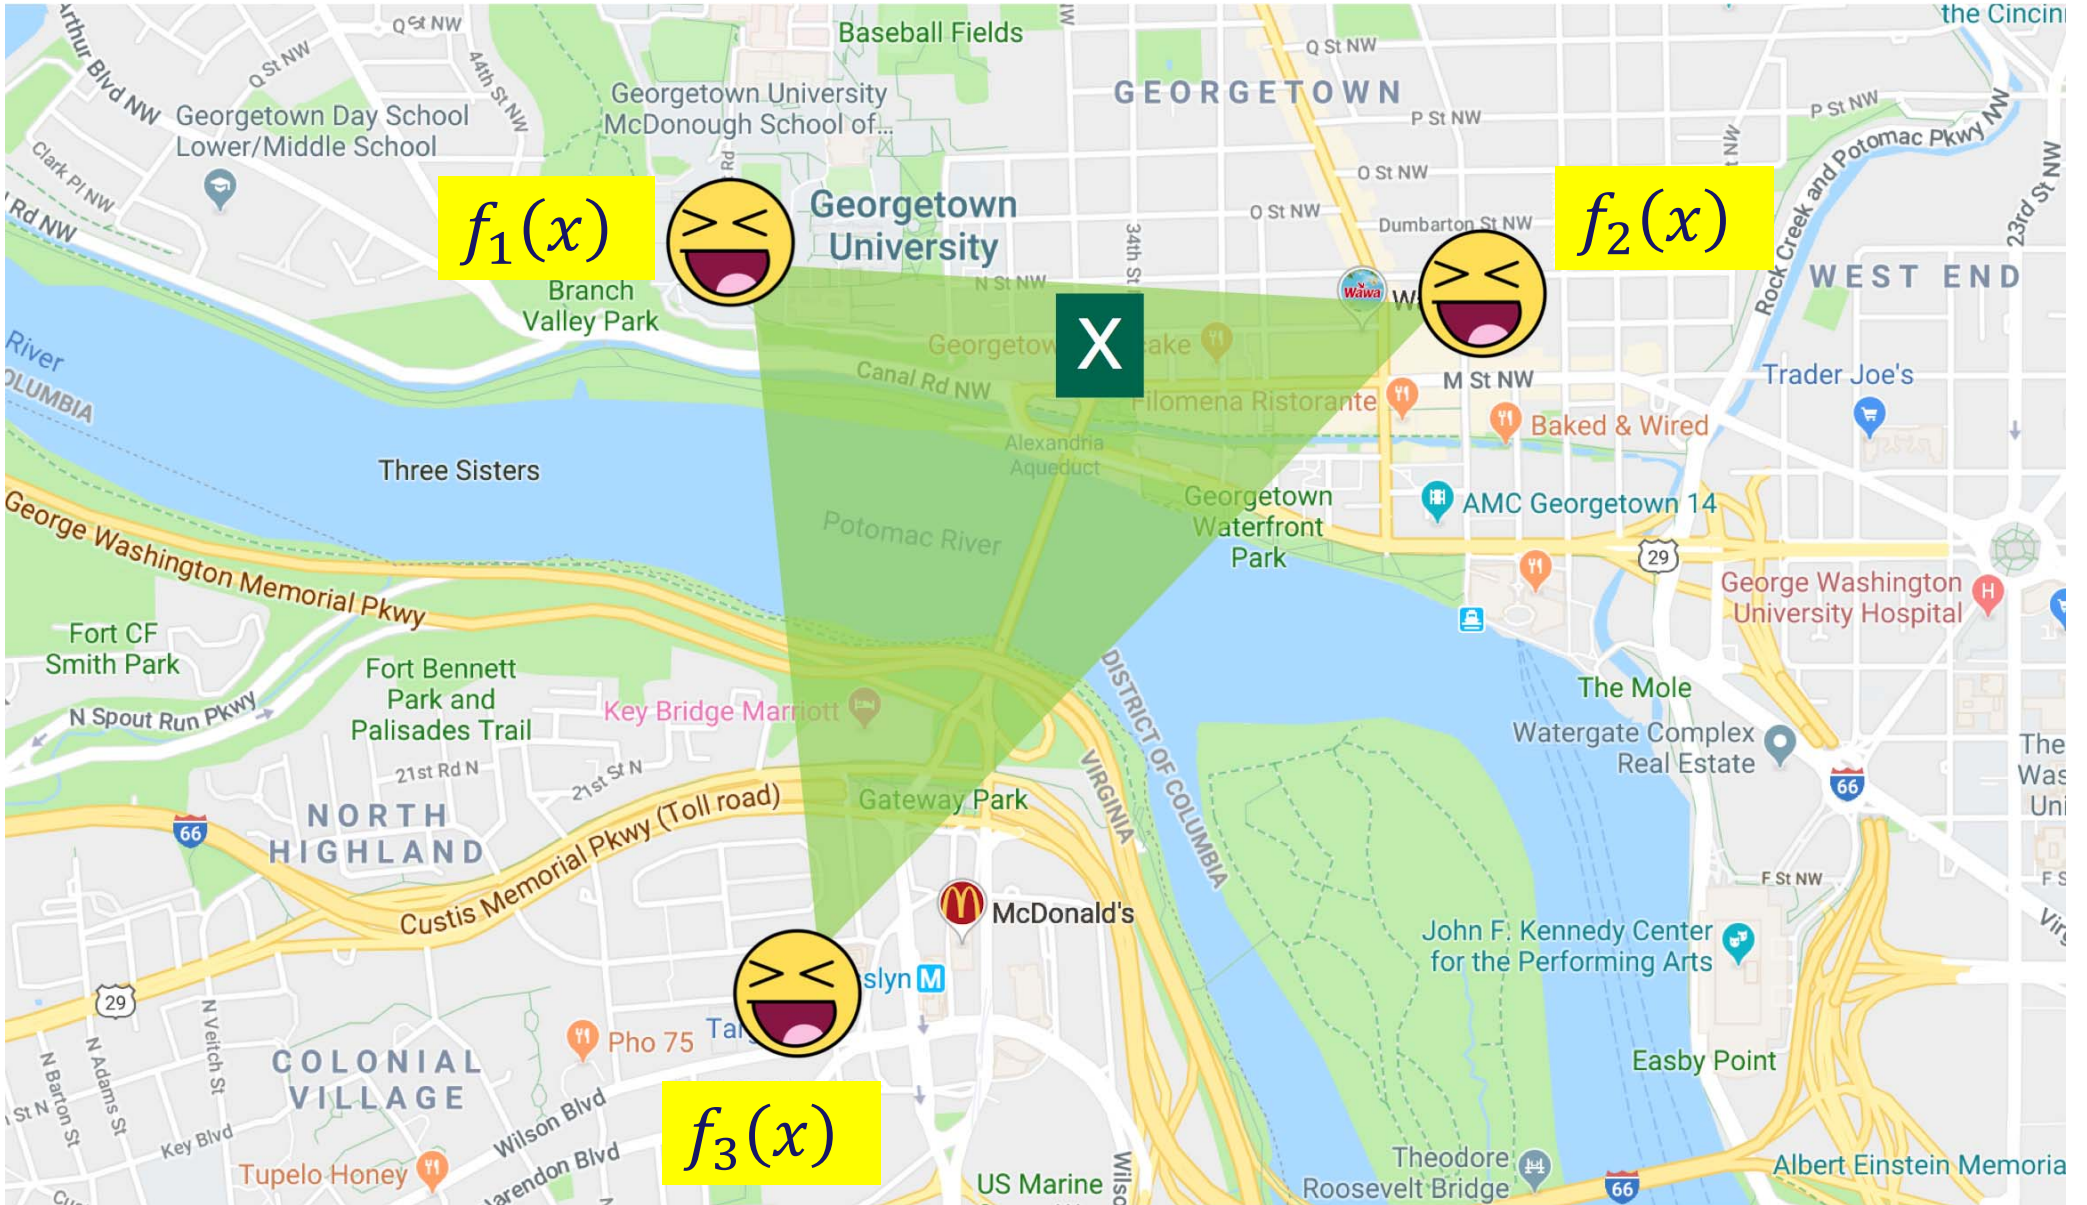
No theorems/proofs

Distributed Optimization

Rendezvous



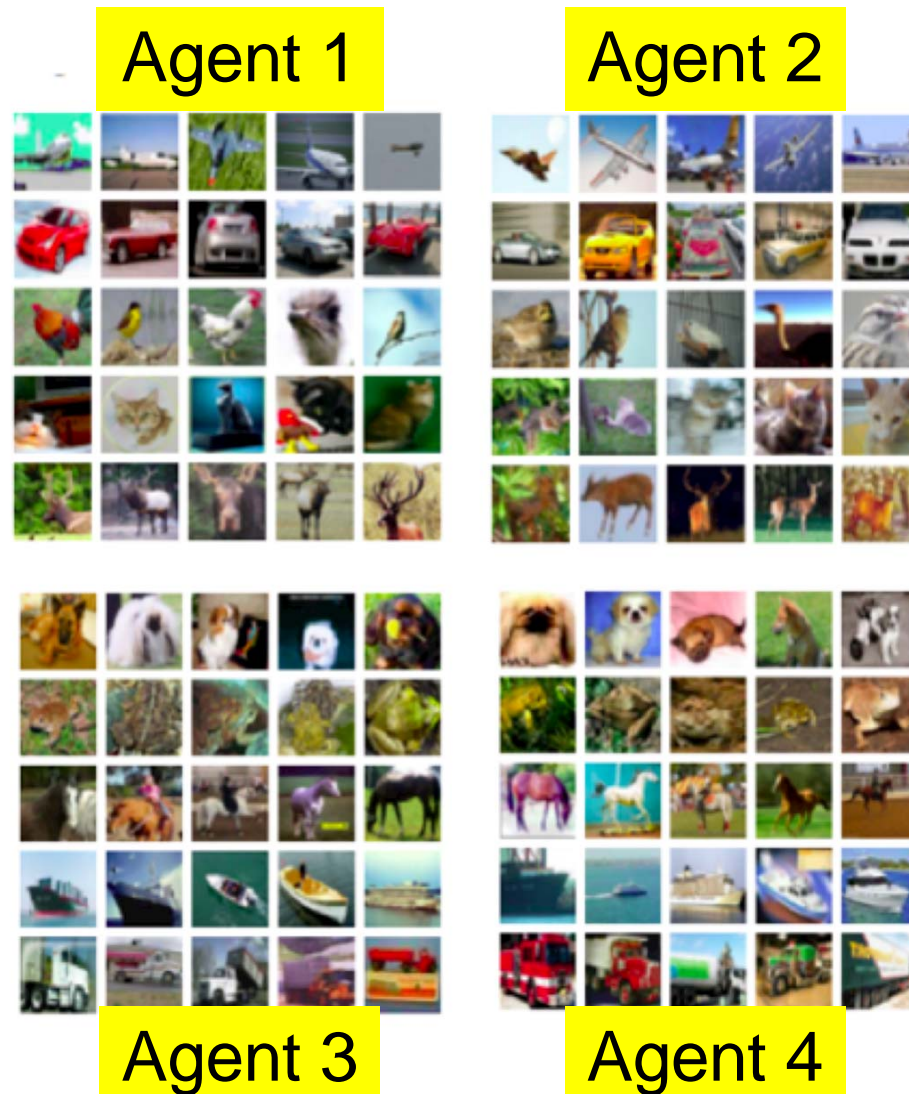
Rendezvous



$$\operatorname{argmin} \sum f_i(x)$$



Machine Learning

- g Data is distributed across different agents

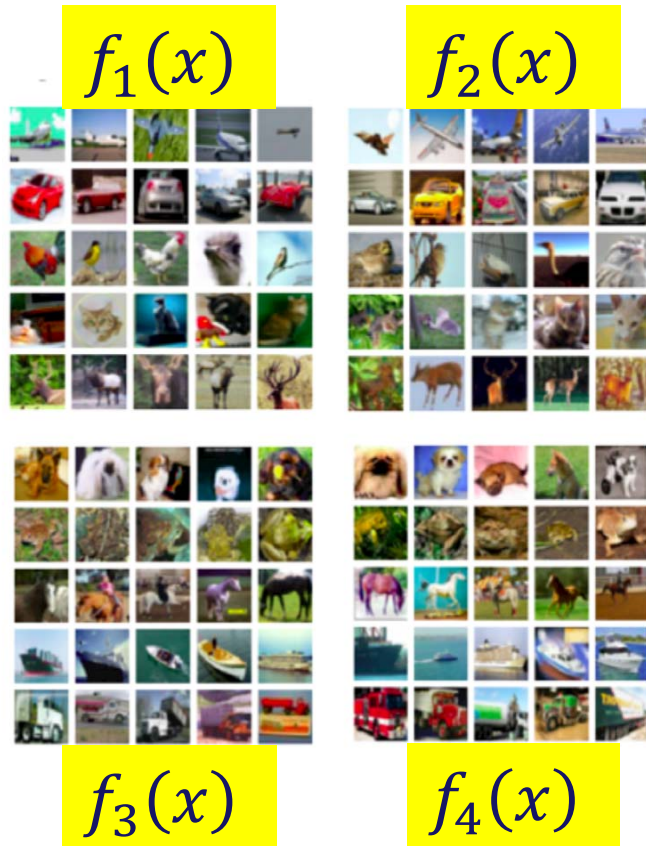


g Data is distributed across different agents

→ Collaborate to learn



Machine Learning




Minimize global loss

$$\operatorname{argmin} \sum f_i(x)$$

$$\operatorname{argmin} \sum f_i(x)$$


Gradient Method

$$f(x) = \sum f_i(x)$$



Gradient Method


$$f(x) = \sum f_i(x)$$



$$x_{k+1} \leftarrow x_k - \lambda \sum_i \nabla f_i(x_k)$$

Gradient Method

$$f(x) = \sum f_i(x)$$





$$x_{k+1} \leftarrow x_k - \lambda \sum_i \nabla f_i(x_k)$$

Distributed Optimization

- g Each agent i knows own cost function $f_i(x)$
- g Need to cooperate to minimize $\sum f_i(x)$


→ Distributed algorithms

Architectures



Parameter Server

- g Server maintains estimate x_k




Parameter Server

g Server maintains estimate x_k

In each iteration

g Agent i

i Receives x_k from server




Parameter Server

- g Server maintains estimate x_k

In each iteration

- g Agent i
 - i Receives x_k from server
 - i Uploads gradient $\nabla f_i(x_k)$




Parameter Server

- g Server maintains estimate x_k

In each iteration


- g Agent i
 - i Receives x_k from server
 - i Uploads gradient $\nabla f_i(x_k)$



- g Server updates estimate

$$x_{k+1} \leftarrow x_k - \lambda \sum \nabla f_i(x_k)$$

Many Variations



- ... stochastic optimization
- ... asynchronous
- ... gradient compression
- ... acceleration
- ... shared memory

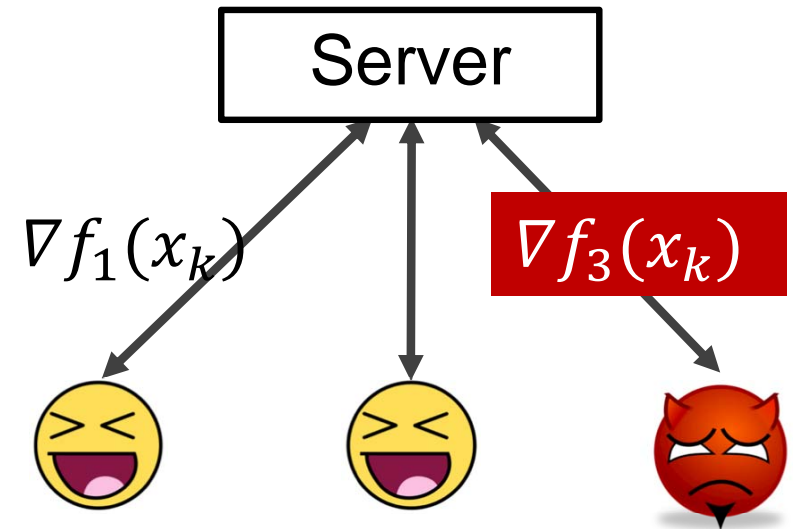
Challenges

Challenges

- g **Fault-tolerant**
distributed optimization

$$f_1(x) + f_2(x) + f_3(x)$$

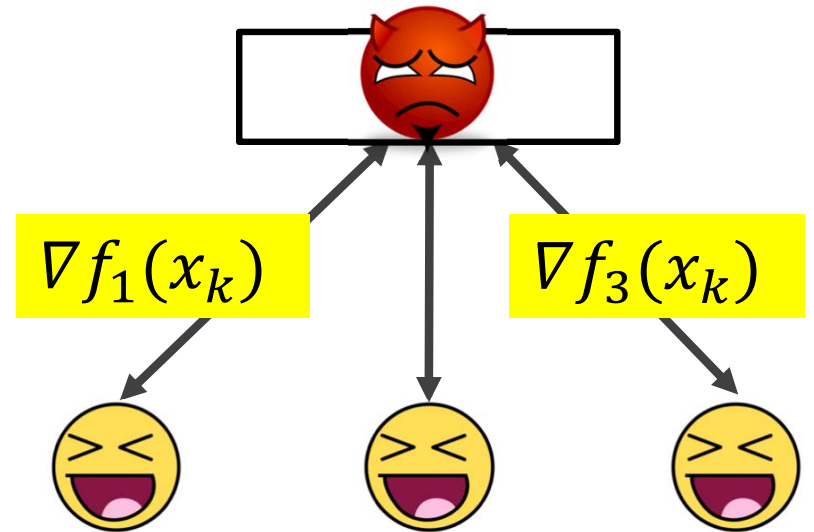
How to optimize
if agents inject
bogus information?



Challenges

- g Privacy-preserving distributed optimization

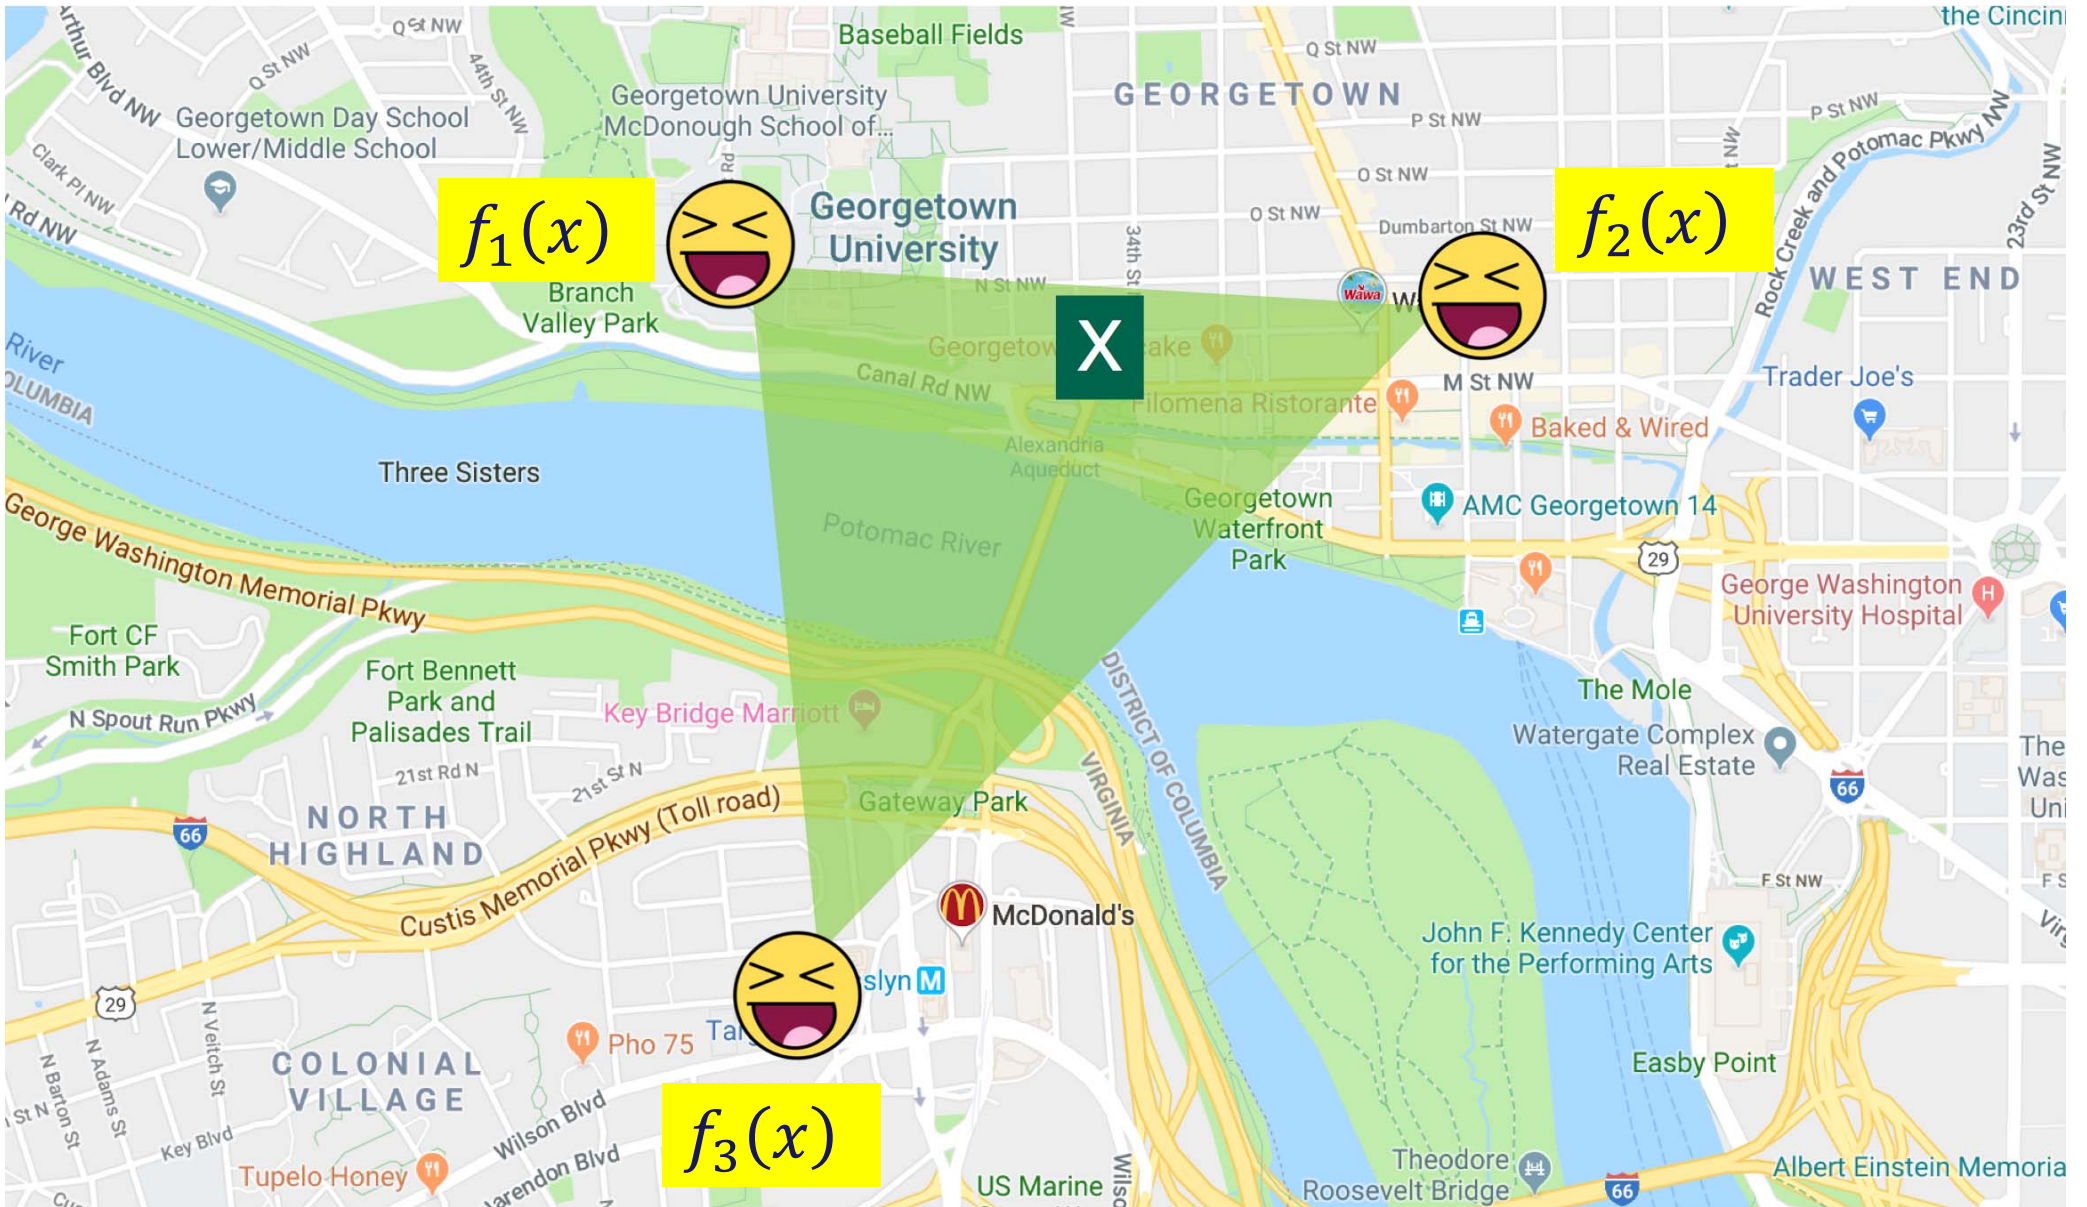
How to collaborate without revealing own cost function?



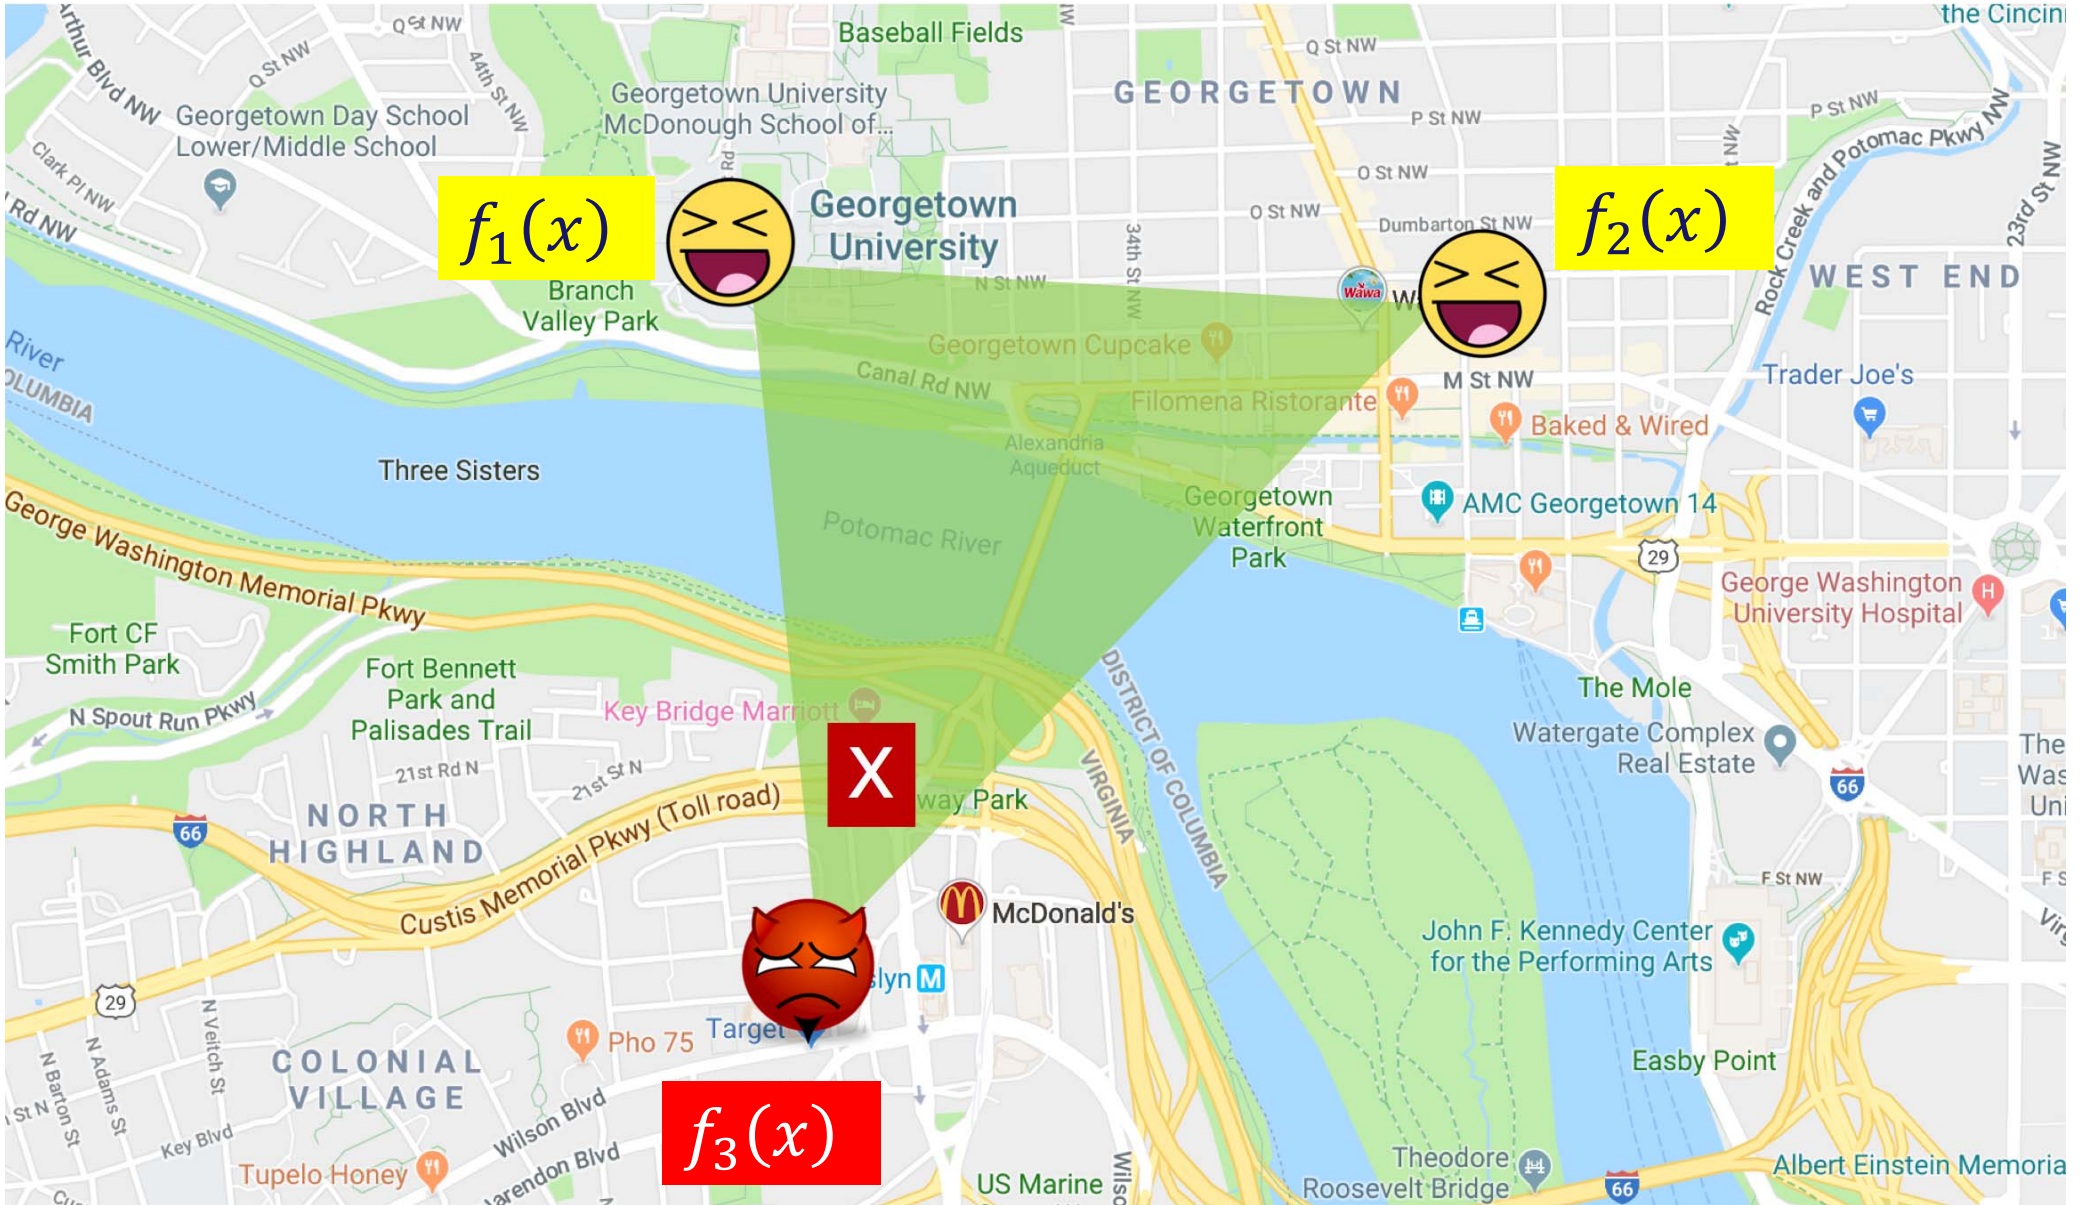
Secure / Fault-Tolerant Optimization

2015 ...

Rendezvous

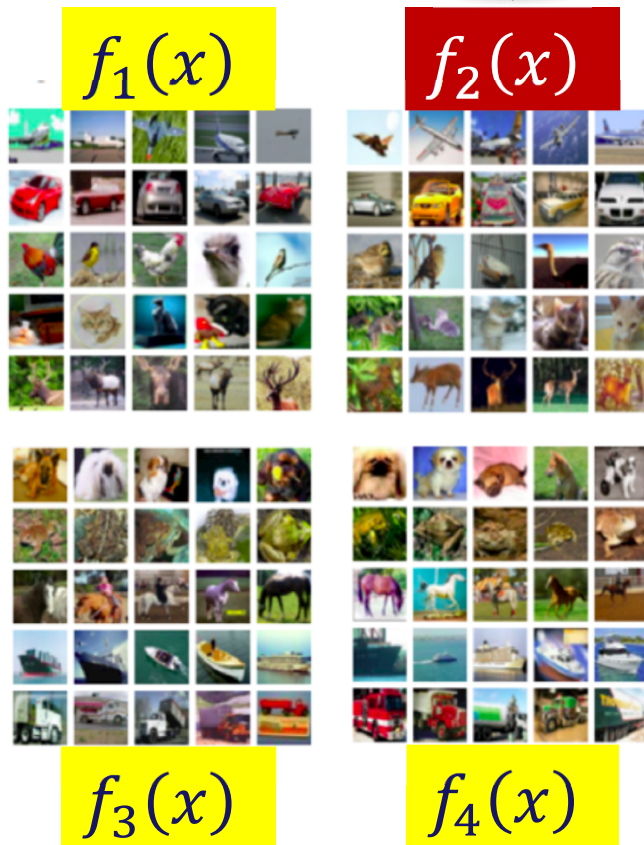


Rendezvous



Machine Learning

Faulty agent can adversely affect model parameters



Minimize global loss


$$\sum f_i(x)$$

Parameter Server

- g Server maintains estimate x_k

In each iteration

- g Agent i
 - i Downloads x_k from server
 - i Uploads gradient $\nabla f_i(x_k)$



- g Server updates estimate

$$x_{k+1} \leftarrow x_k - \lambda \sum \nabla f_i(x_k)$$

Parameter Server


g Server maintains estimate x_k

In each iteration

g Agent i

i Downloads x_k from server

i Uploads gradient $\nabla f_i(x_k)$



g Server updates estimate

$$x_{k+1} \leftarrow x_k - \lambda \text{Filtered-Gradient}$$



But what do we mean by fault-tolerance?

Fault-Tolerance

g Optimize over only good agents ... set G

Fault-Tolerance

- g Optimize over only good agents ... set G

$$\operatorname{argmin} \sum_{i \in G} f_i(x)$$

$$\operatorname{argmin} \sum_{i \in G} f_i(x)$$

Is this achievable?

$$\operatorname{argmin} \sum_{i \in G} f_i(x)$$

Is this achievable?

It Depends

$$\operatorname{argmin} \sum_{i \in G} f_i(x)$$


Is this achievable?

It Depends

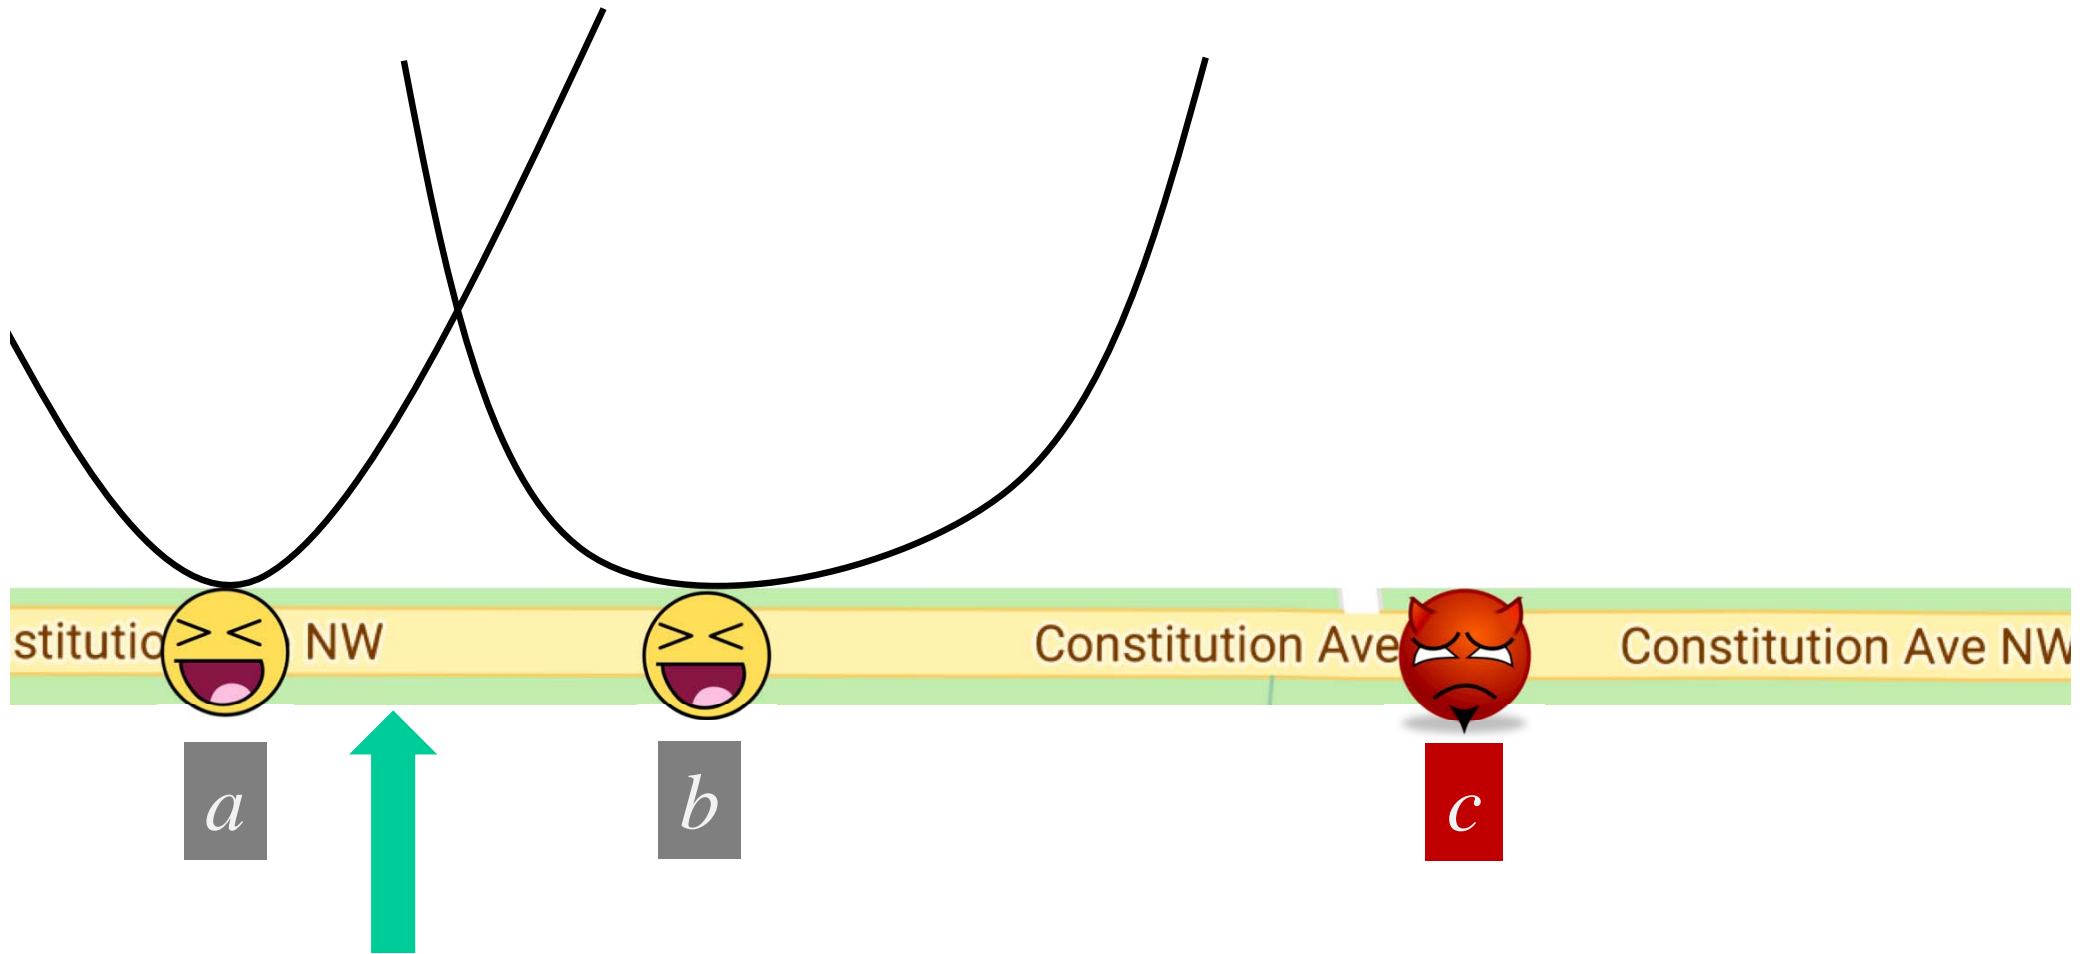
Independent
functions

“Enough”
redundancy


Independent Functions



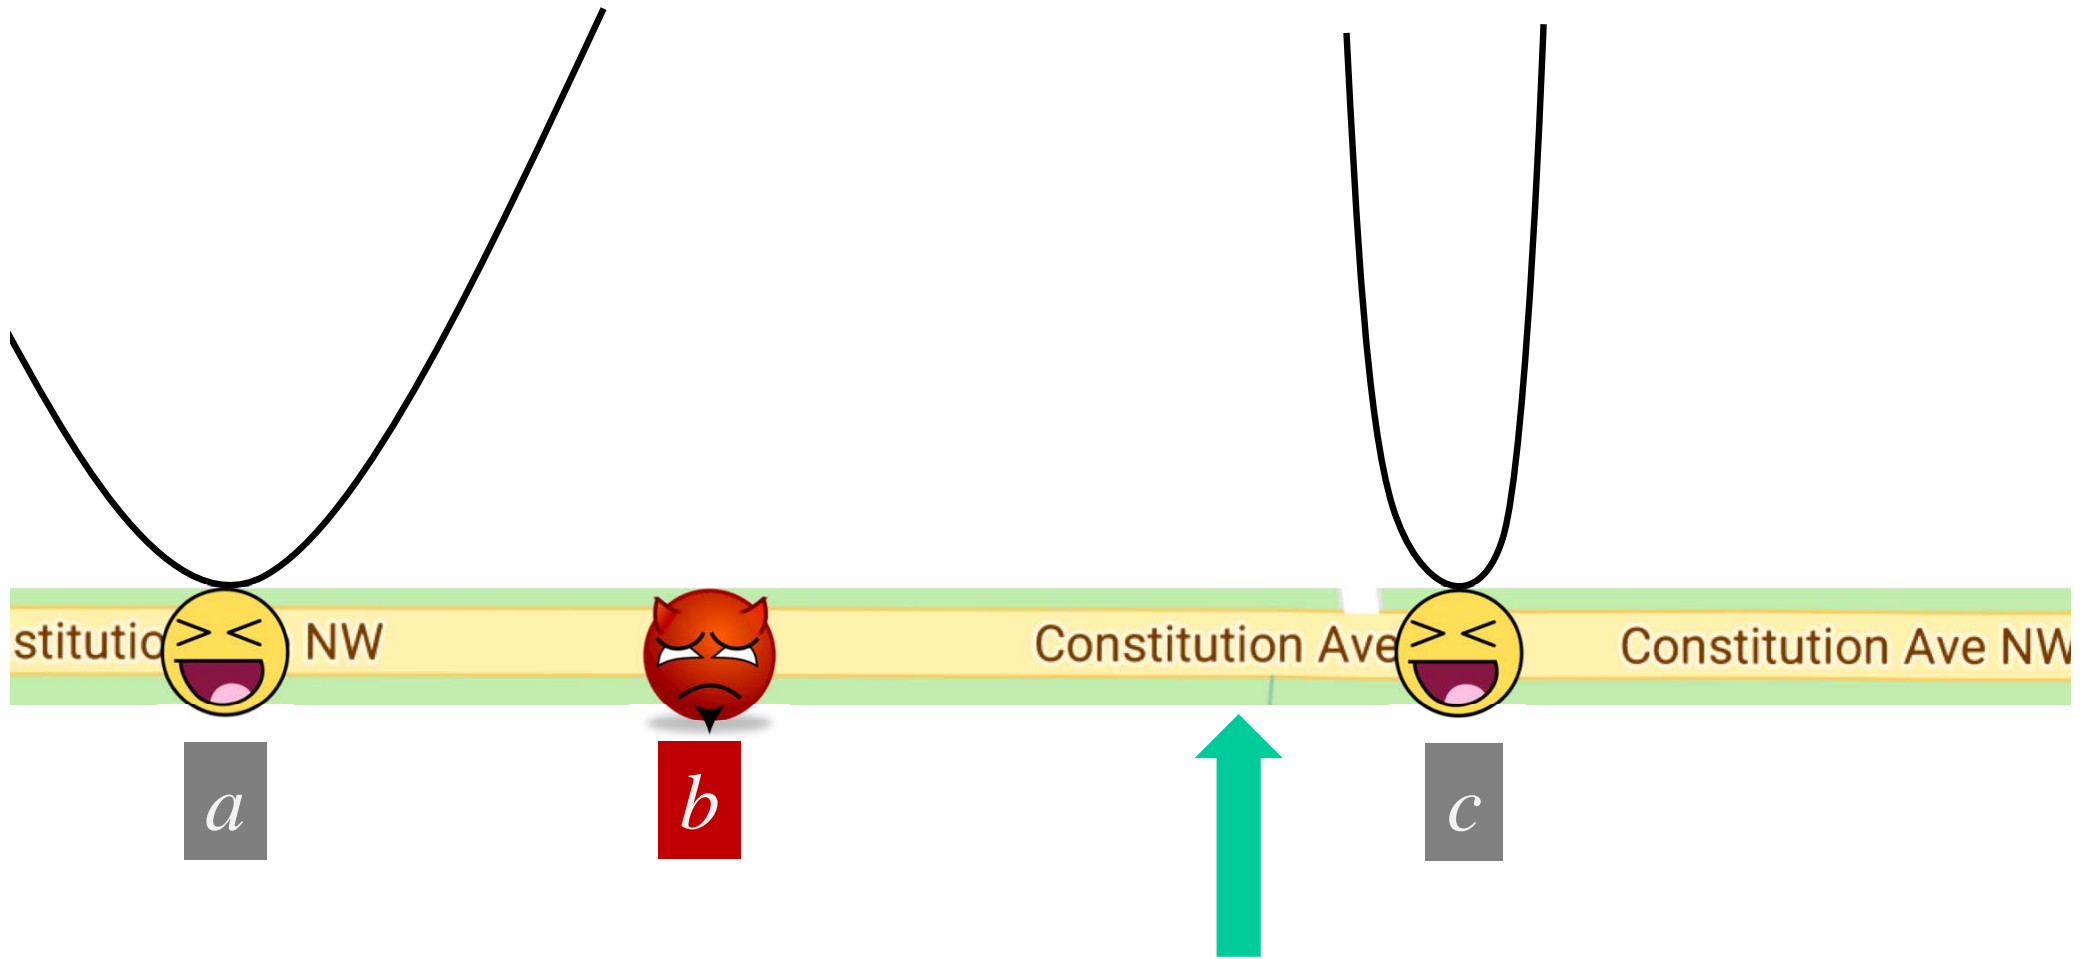
Independent Functions




Independent Functions



Independent Functions



Independent Functions



Provably impossible to compute

$$\operatorname{argmin} \sum_{i \in G} f_i(x)$$

$$\operatorname{argmin} \sum_{i \in G} f_i(x)$$

Is this achievable?

Independent
functions

“Enough”
redundancy

Approximate

Exact

$$\operatorname{argmin} \sum_{i \in G} f_i(x)$$

Is this achievable?

Independent
functions

“Enough”
redundancy

Approximate

Exact

An Example of Redundancy

n agents

t bad agents

g Aggregate cost of **ANY** $n - 2t$ agents has


argmin identical to **desired argmin** $\sum_{i \in G} f_i(x)$

Parameter Server

- g Server maintains estimate x_k

In each iteration

- g Agent i
 - i Downloads x_k from server
 - i Uploads gradient $\nabla f_i(x_k)$



- g Server updates estimate

$$x_{k+1} \leftarrow x_k - \lambda \text{Filtered-Gradient}$$

Norm Filter

g Clip the largest t norms to equal $t + 1^{\text{th}}$ norm

$$|\nabla f_1(x_k)| = 1$$

$$|\nabla f_2(x_k)| = 3$$

$$|\nabla f_3(x_k)| = 2$$

Norm Filter

g Clip the largest t norms to equal $t + 1^{\text{th}}$ norm

$$|\nabla f_1(x_k)| = 1$$

$$|\nabla f_2(x_k)| = 3$$

$$|\nabla f_3(x_k)| = 2$$

$$\text{Filtered gradient} = \nabla f_1(x_k) + \frac{2}{3} \nabla f_2(x_k) + \nabla f_3(x_k)$$

Norm Filter

- g Clip the largest t norms to equal $t + 1^{\text{th}}$ norm

$$|\nabla f_1(x_k)| = 1$$

$$|\nabla f_2(x_k)| = 3$$

$$|\nabla f_3(x_k)| = 2$$

$$\text{Filtered gradient} = \nabla f_1(x_k) + \frac{2}{3} \nabla f_2(x_k) + \nabla f_3(x_k)$$

Exact optimum computed despite faulty agents

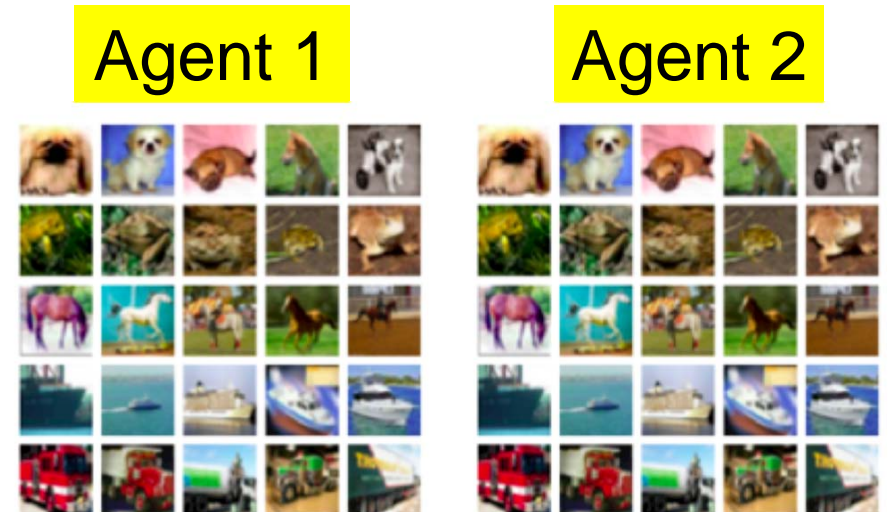
Another Example of Redundancy

Another Example of Redundancy

g Machine learning

g Agents draw samples from **identical** data distribution

g Filter on stochastic gradients



Relaxing Redundancy Requirements

Ideas also extend to a relaxed notion of

“enough redundancy”

$$\operatorname{argmin} \sum_{i \in G} f_i(x)$$

Is this achievable?

Independent
functions

“Enough”
redundancy

Approximate

Exact

How to Approximate $\operatorname{argmin} \sum_{i \in G} f_i(x)$?

How to Approximate $\operatorname{argmin} \sum_{i \in G} f_i(x)$?

Several alternatives for approximation

How to Approximate $\operatorname{argmin} \sum_{i \in G} f_i(x)$?

$$\operatorname{argmin} \sum_{i \in G} f_i(x) = \operatorname{argmin} \sum_{i \in G} \frac{1}{|G|} f_i(x)$$

How to Approximate $\operatorname{argmin} \sum_{i \in G} f_i(x)$?

$$\operatorname{argmin} \sum_{i \in G} f_i(x) = \operatorname{argmin} \sum_{i \in G} \frac{1}{|G|} f_i(x)$$

g **Ideal goal:** Equal weight for all non-faulty agents

How to Approximate $\operatorname{argmin} \sum_{i \in G} f_i(x)$?

$$\operatorname{argmin} \sum_{i \in G} f_i(x) = \operatorname{argmin} \sum_{i \in G} \frac{1}{|G|} f_i(x)$$

- g **Ideal goal:** Equal weight for all non-faulty agents
- g **Approximation:** Unequal weights

Results

- g For each faulty agent,
a good agent may be ignored → Weight 0

Results

- g For each faulty agent,
a good agent may be ignored → Weight 0
- g But remaining $(n - 2t)$ good agents get
“almost” uniform weight

n agents,
up to t faulty

Results

- g For each faulty agent,
a good agent may be ignored → Weight 0
- g But remaining $(n - 2t)$ good agents get
“almost” uniform weight
- g Bad agents all get weight = 0


n agents,
up to t faulty

Parameter Server

- g Server maintains estimate x_k

In each iteration

- g Agent i
 - i Downloads x_k from server
 - i Uploads gradient $\nabla f_i(x_k)$




- g Server updates estimate

$$x_{k+1} \leftarrow x_k - \lambda \text{Filtered-Gradient}$$


Privacy-Preserving Optimization

2016 ...

Communication Leaks Information

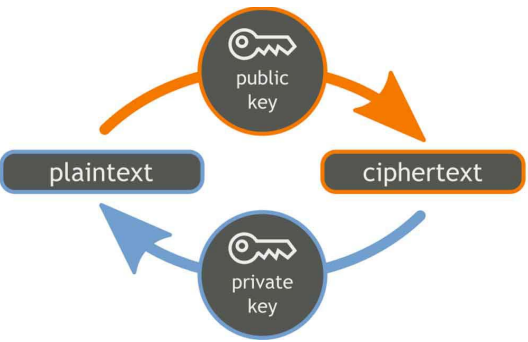
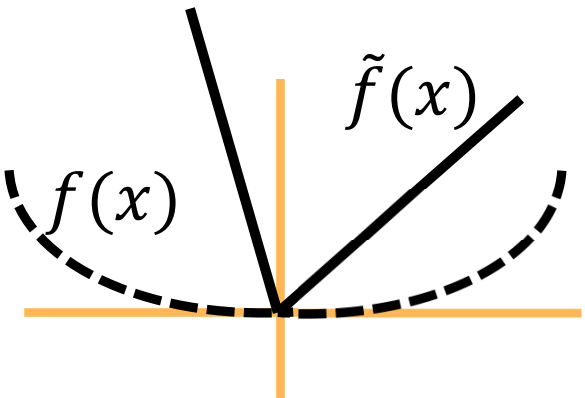
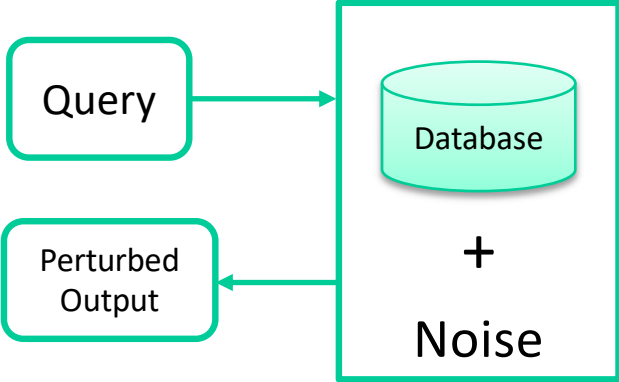


Communication Leaks Information



Server can use gradients to infer polynomial cost functions (up to a constant)

Related Work

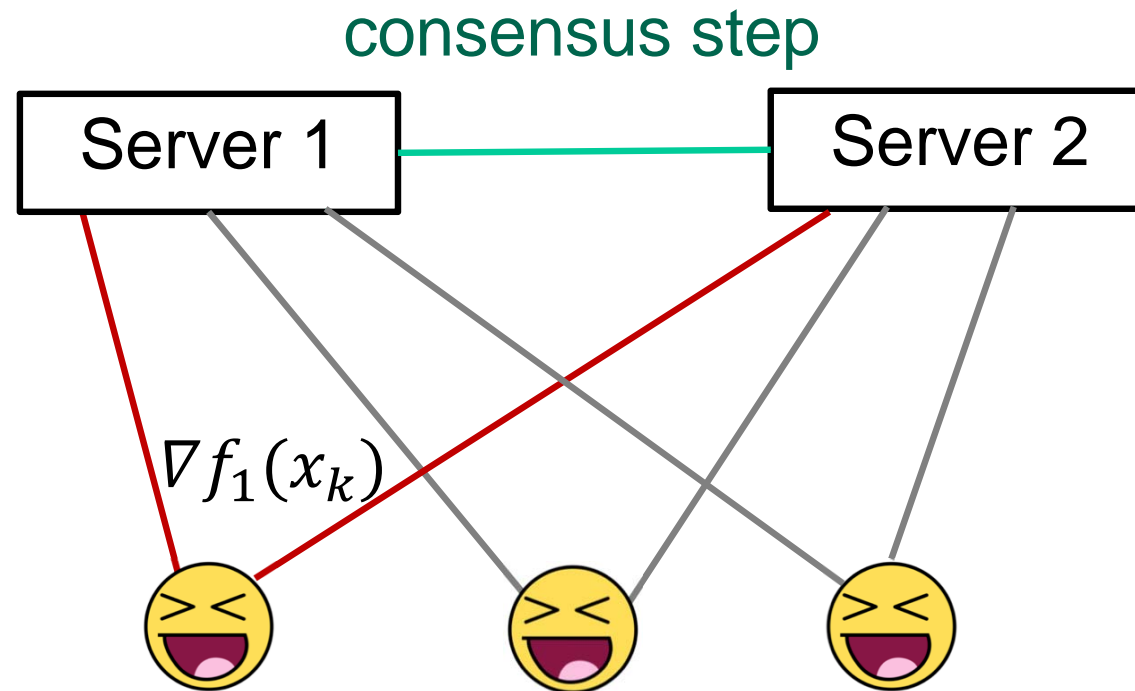
Cryptographic Methods	Transformation Methods	Differential Privacy
 <p>The diagram illustrates the cryptographic process. It shows a cycle: 'plaintext' (in a blue rounded rectangle) is converted to 'ciphertext' (in an orange rounded rectangle) using a 'public key' (represented by a key icon in a grey circle). Conversely, 'ciphertext' is converted back to 'plaintext' using a 'private key' (represented by a key icon in a blue circle).</p>	 <p>The diagram shows a coordinate system with a vertical orange axis. A dashed black curve represents the function $f(x)$. A solid black line represents the perturbed function $\tilde{f}(x)$. The perturbation is shown as a sharp peak at the origin, indicating a significant change in the function's value at that point.</p>	 <p>The diagram illustrates the differential privacy process. A 'Query' (in a rounded rectangle) is sent to a 'Database' (represented by a cylinder icon). The database output is then combined with 'Noise' (indicated by a plus sign) to produce a 'Perturbed Output' (in a rounded rectangle).</p>

Our Approach

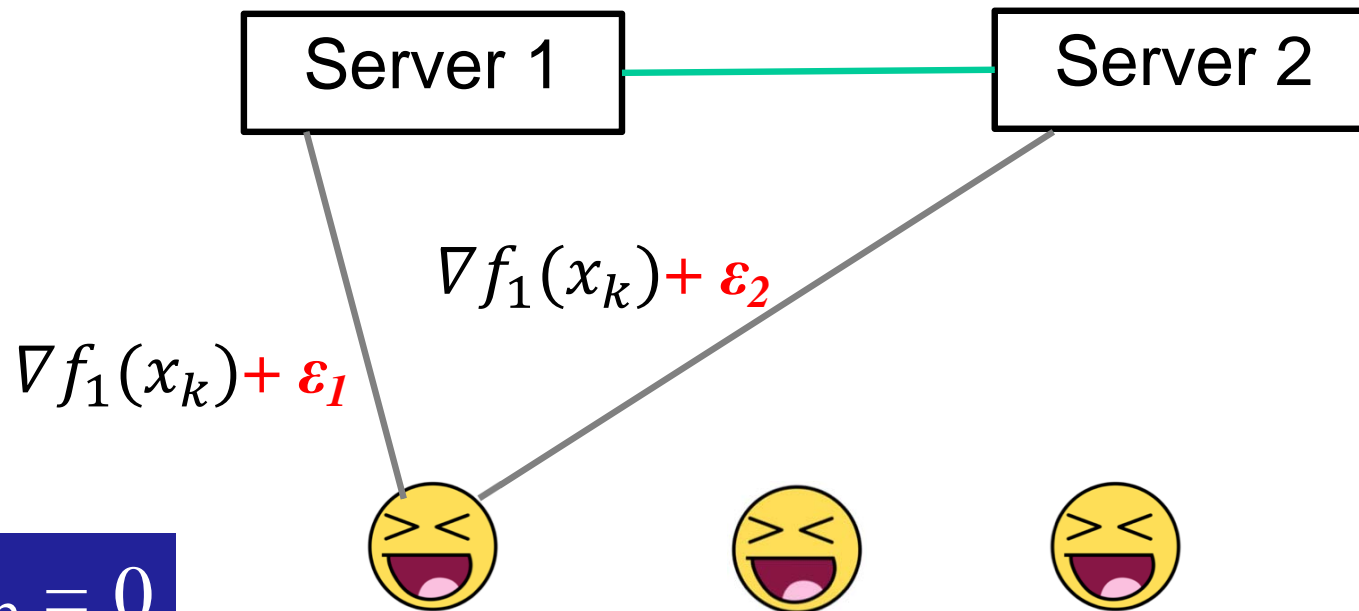
g Motivated by secret sharing & differential privacy

→ Add **cancellable** noise

Multiple Parameter Servers

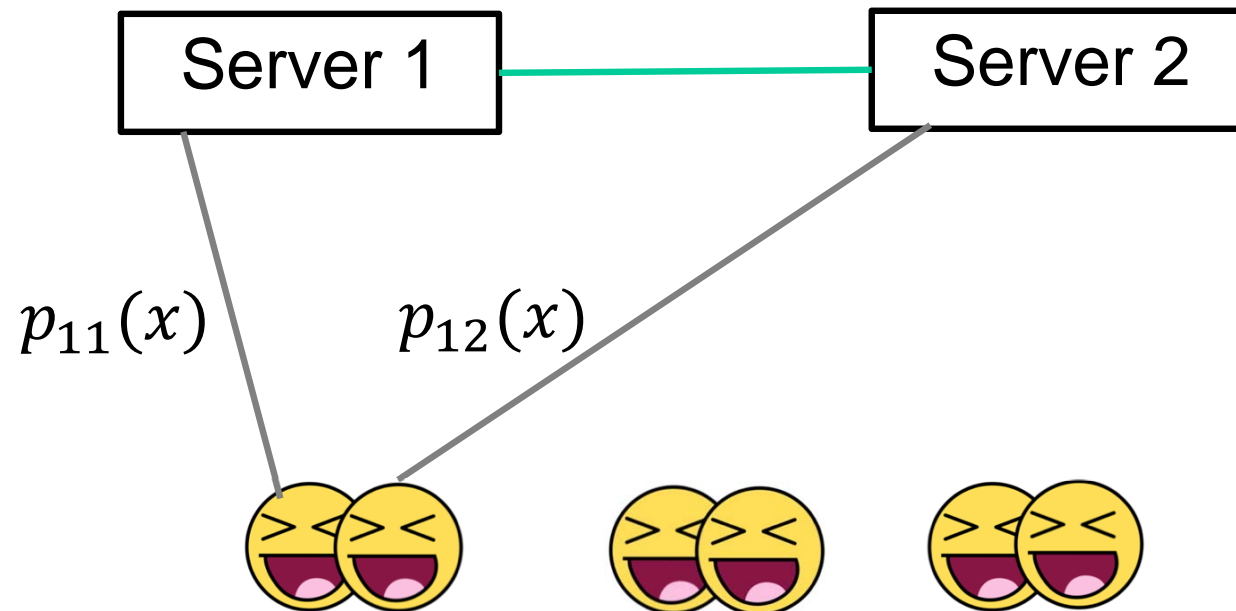


Improving Privacy



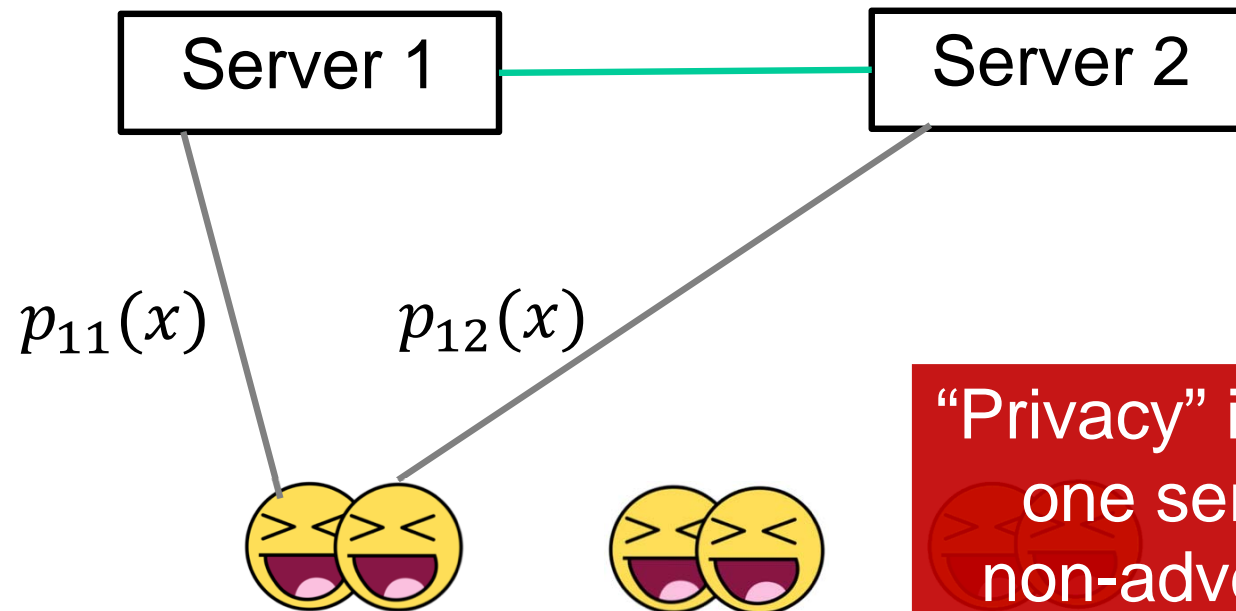
$\epsilon_1 + \epsilon_2 = 0$
over time
and space

Convex Sum of Non-Convex Functions



$$p_{11}(x) + p_{12}(x) = f_1(x)$$

Convex Sum of Non-Convex Functions



$$p_{11}(x) + p_{12}(x) = f_1(x)$$



Summary: $\operatorname{argmin} \sum f_i(x)$

Fault-tolerance

→ Gradient filters

Privacy

→ Cancellable noise



Acknowledgments

- g Lili Su
- g Shripad Gade
- g Nirupam Gupta
- g Shuo Liu
- g Connor Lu
- g Dimitrios Pylorof



Thanks!

A longer tutorial at

disc.georgetown.domains → Talks

Results (Scalar x)

Can output

$$\operatorname{argmin} \sum_{i \in G} \alpha_i f_i(x)$$

for some weights (α_i)

where at least $(n - 2t)$ good agents have weight

$$\alpha_i \geq \frac{1}{2(n - 2t)}$$