


# Prio: Private, Robust, and Efficient Computation of Aggregate Statistics


Henry Corrigan-Gibbs and Dan Boneh  
Stanford University

Appeared at NSDI 2017

# Today: Non-private aggregation



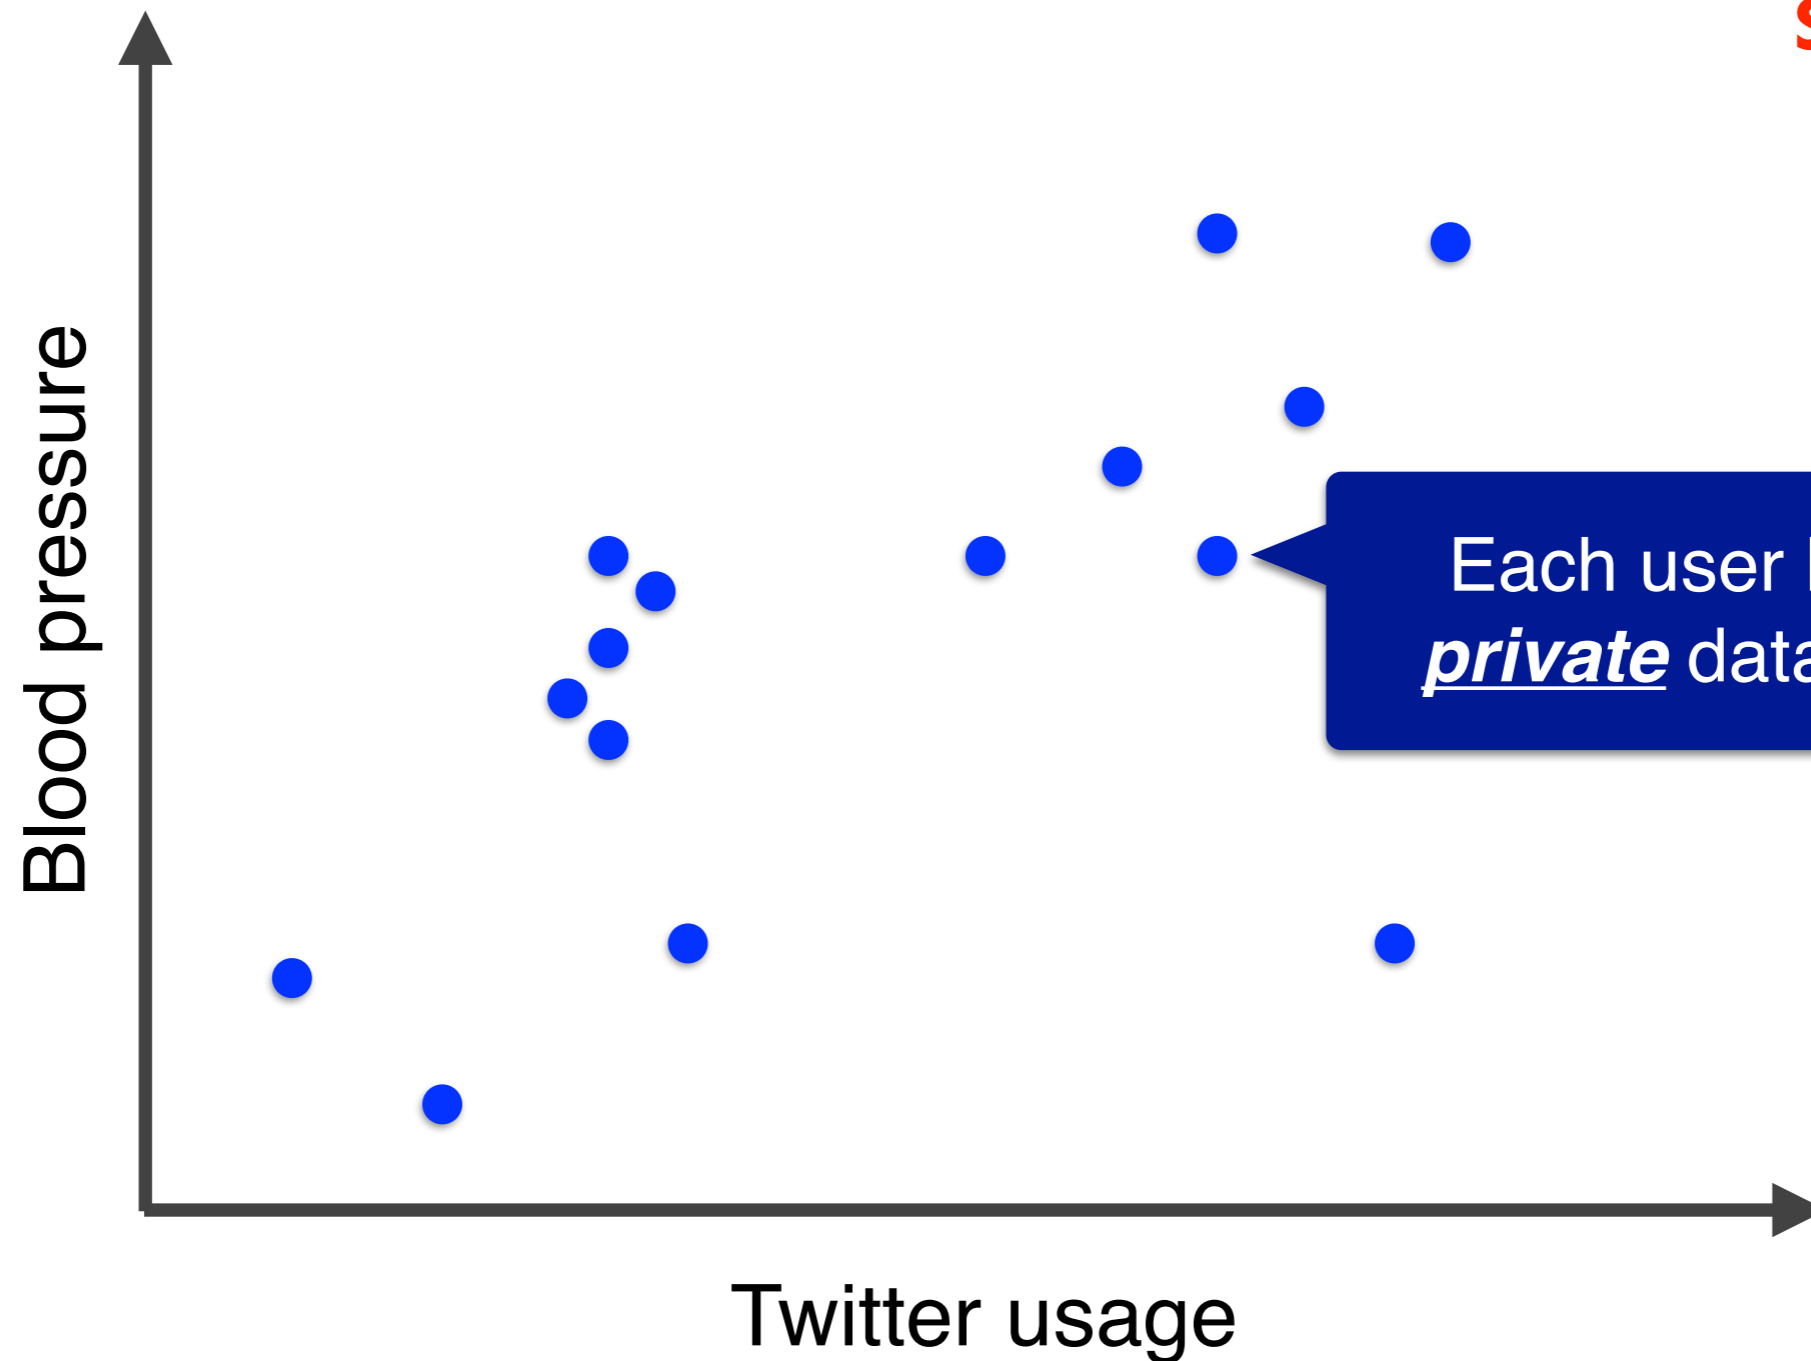
**StressTracker**



# Today: Non-private aggregation




**StressTracker**



# Today: Non-private aggregation




**StressTracker**



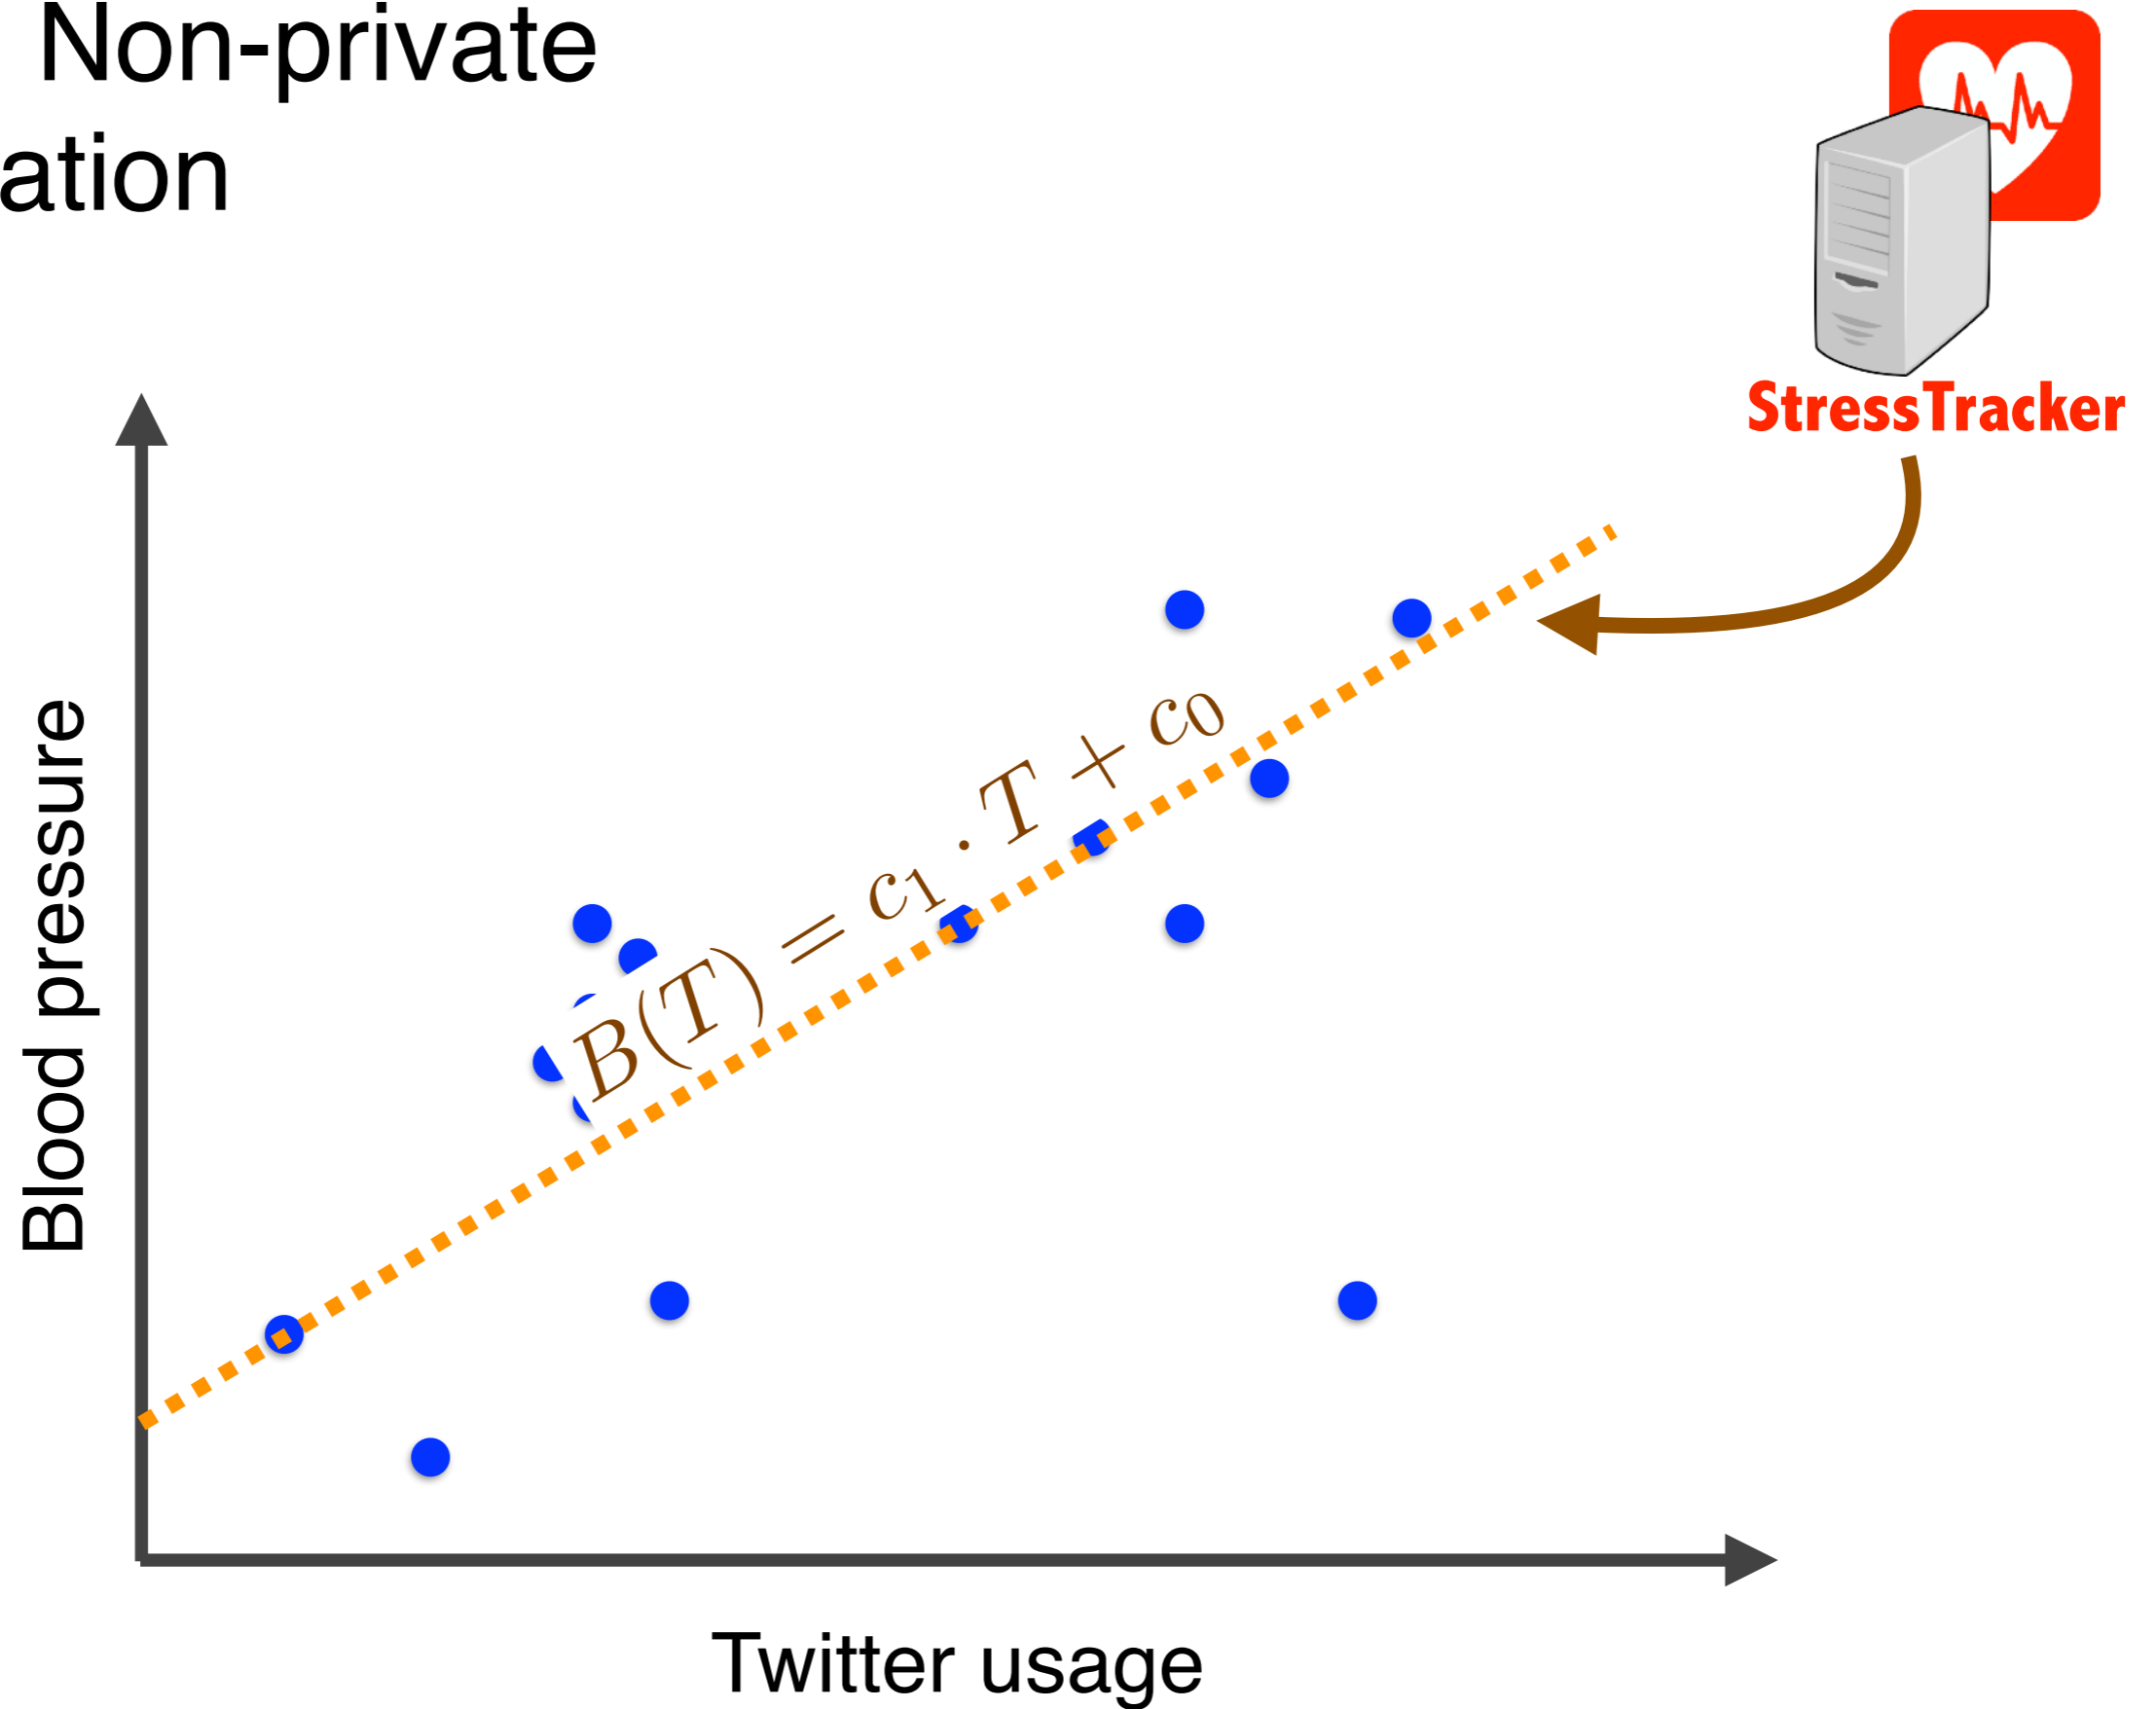
# Today: Non-private aggregation



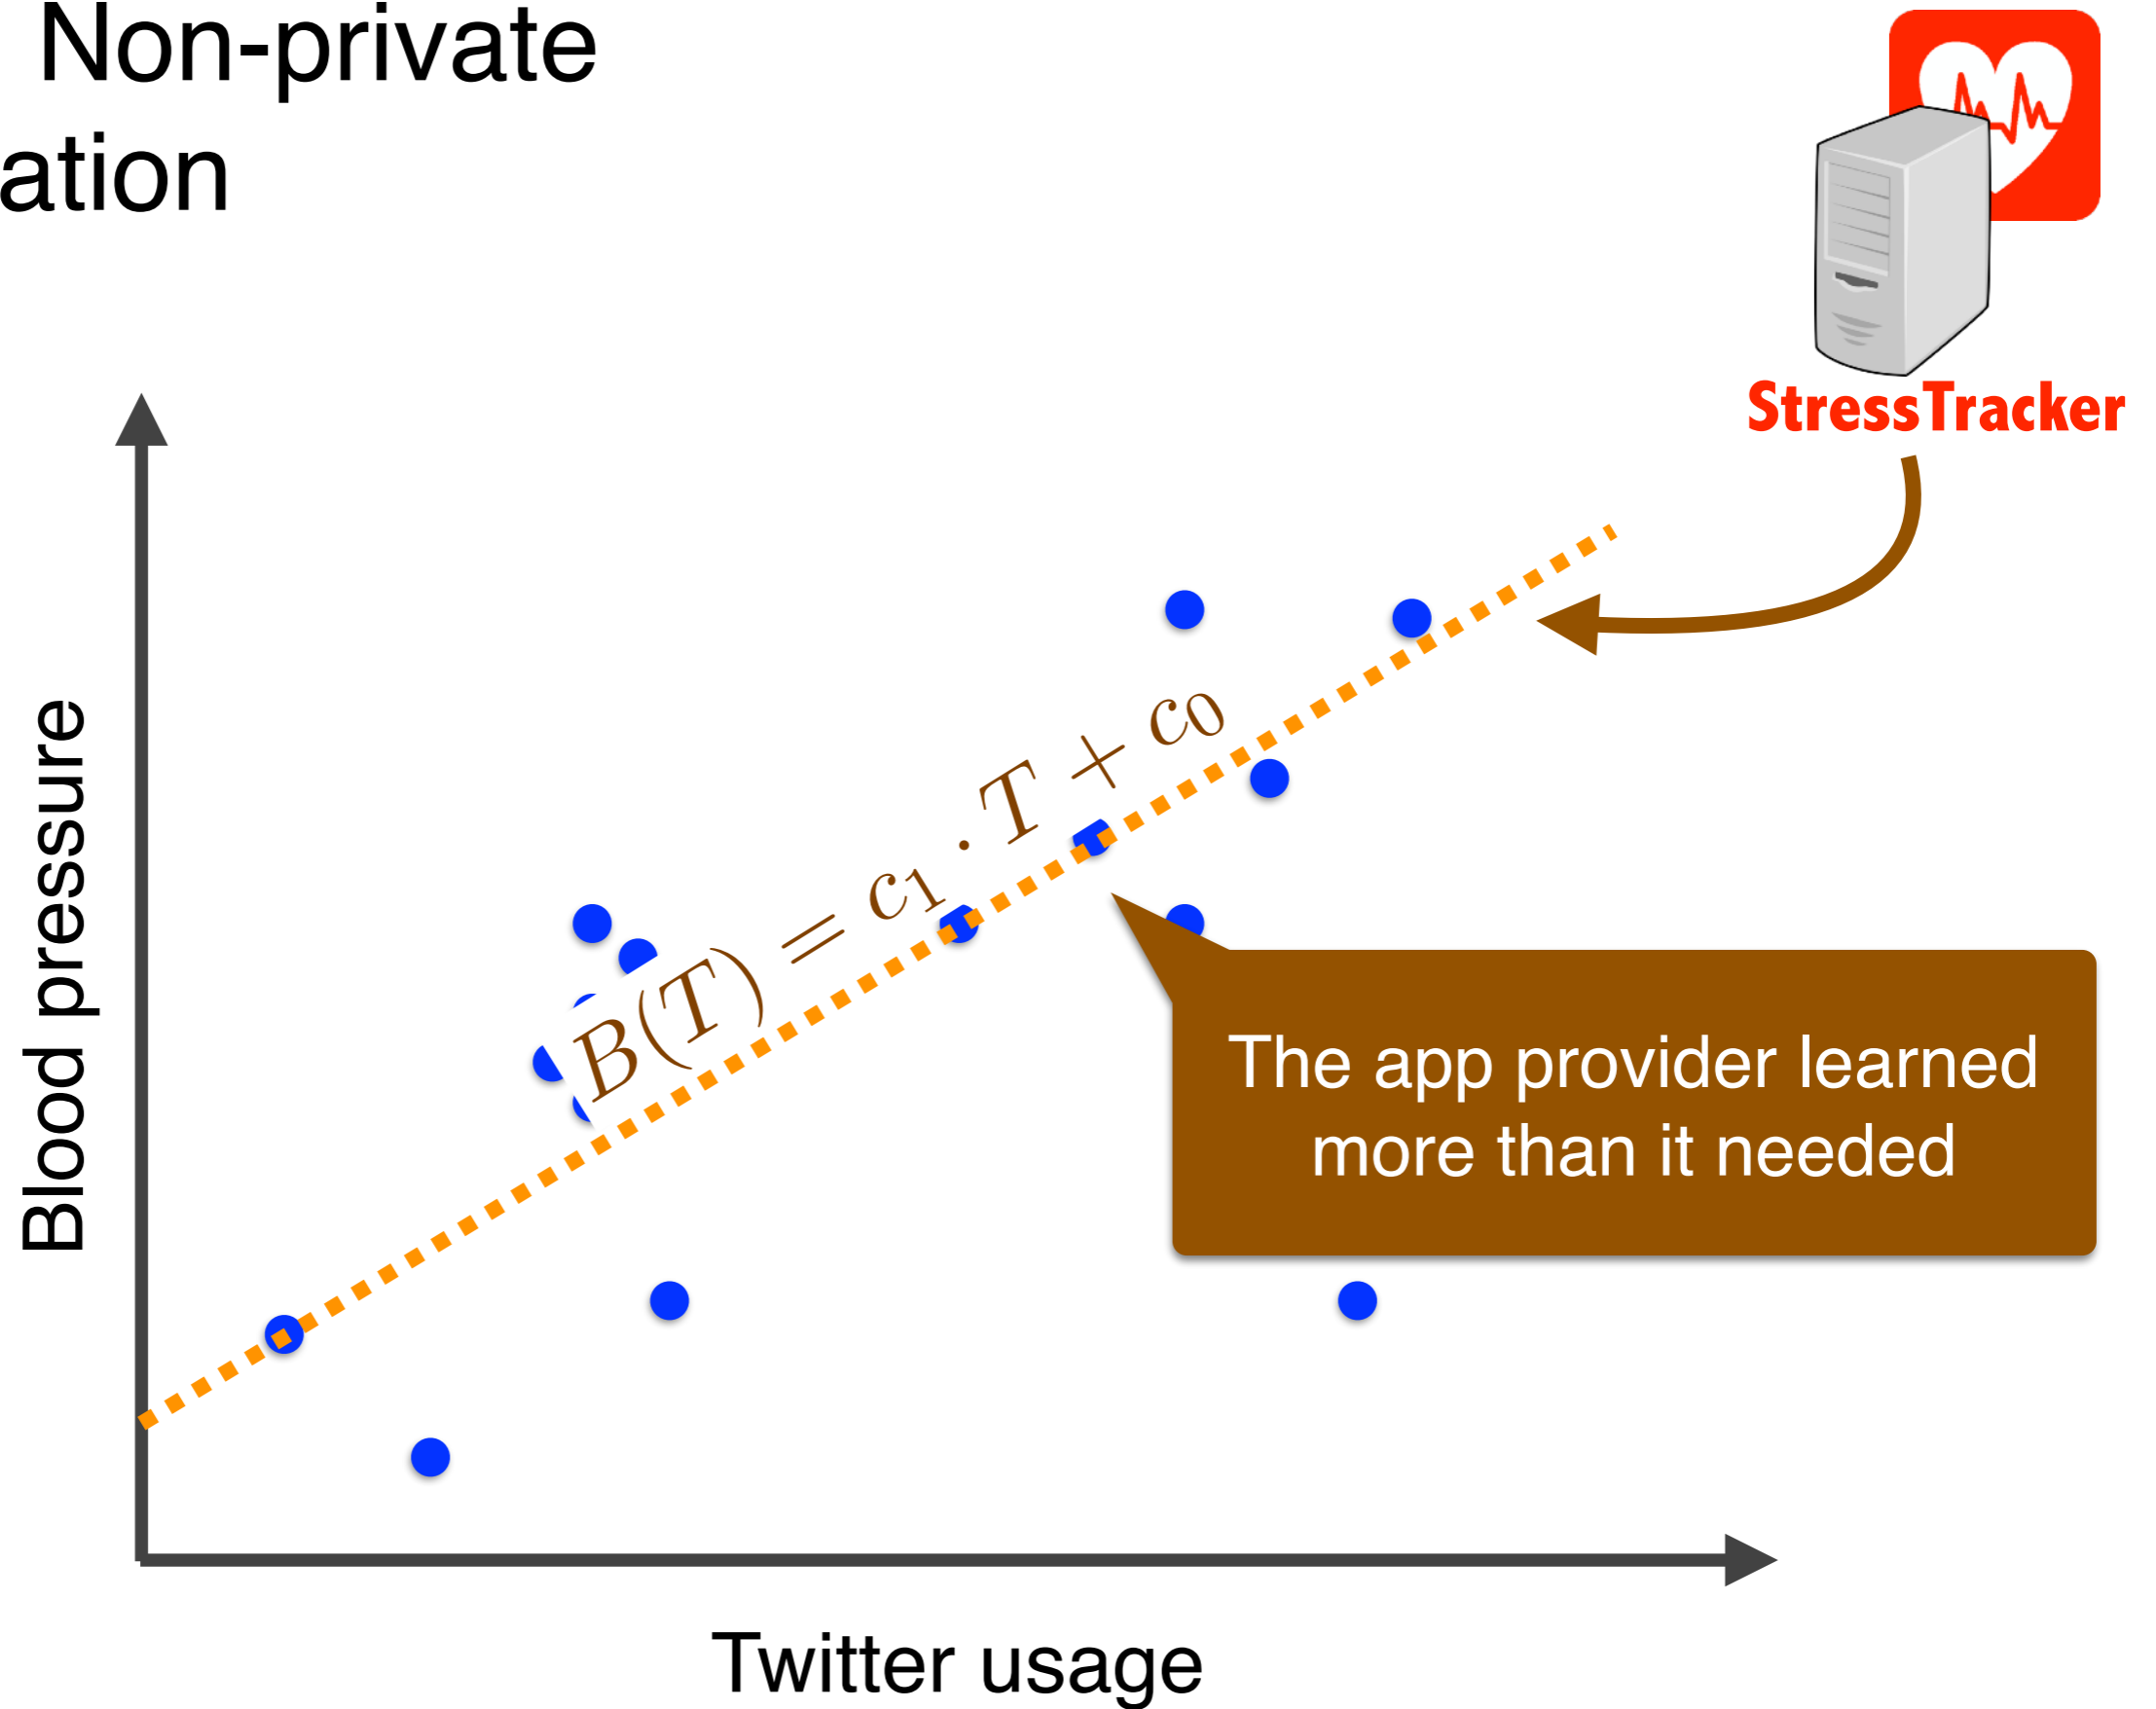
**StressTracker**



# Today: Non-private aggregation




# Today: Non-private aggregation



# Today: Non-private aggregation




**StressTracker**





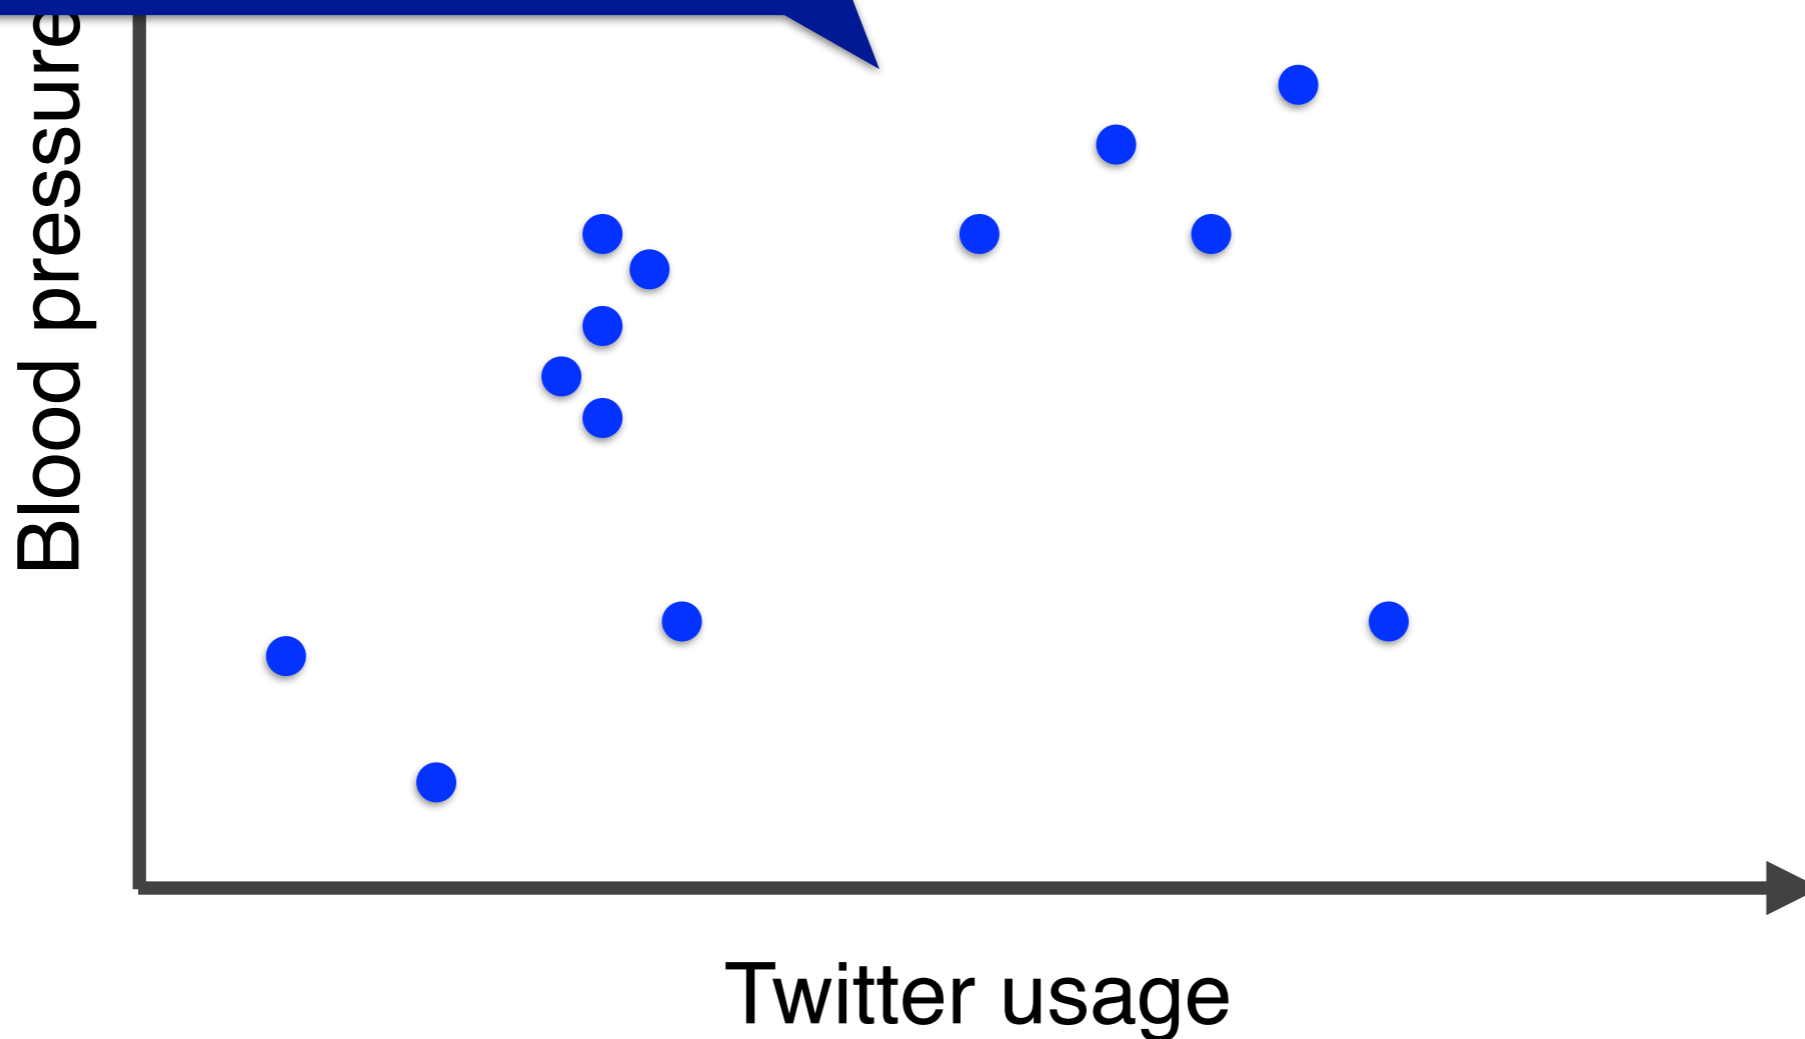
# This work: Private aggregation



# This work: Private aggregation




Clients send an *encrypted share* of their data to each aggregator



# This work: Private aggregation



Clients send an *encrypted share* of their data to each aggregator



# This work: Private aggregation

Clients send an *encrypted share* of their data to each aggregator

Blood pressure

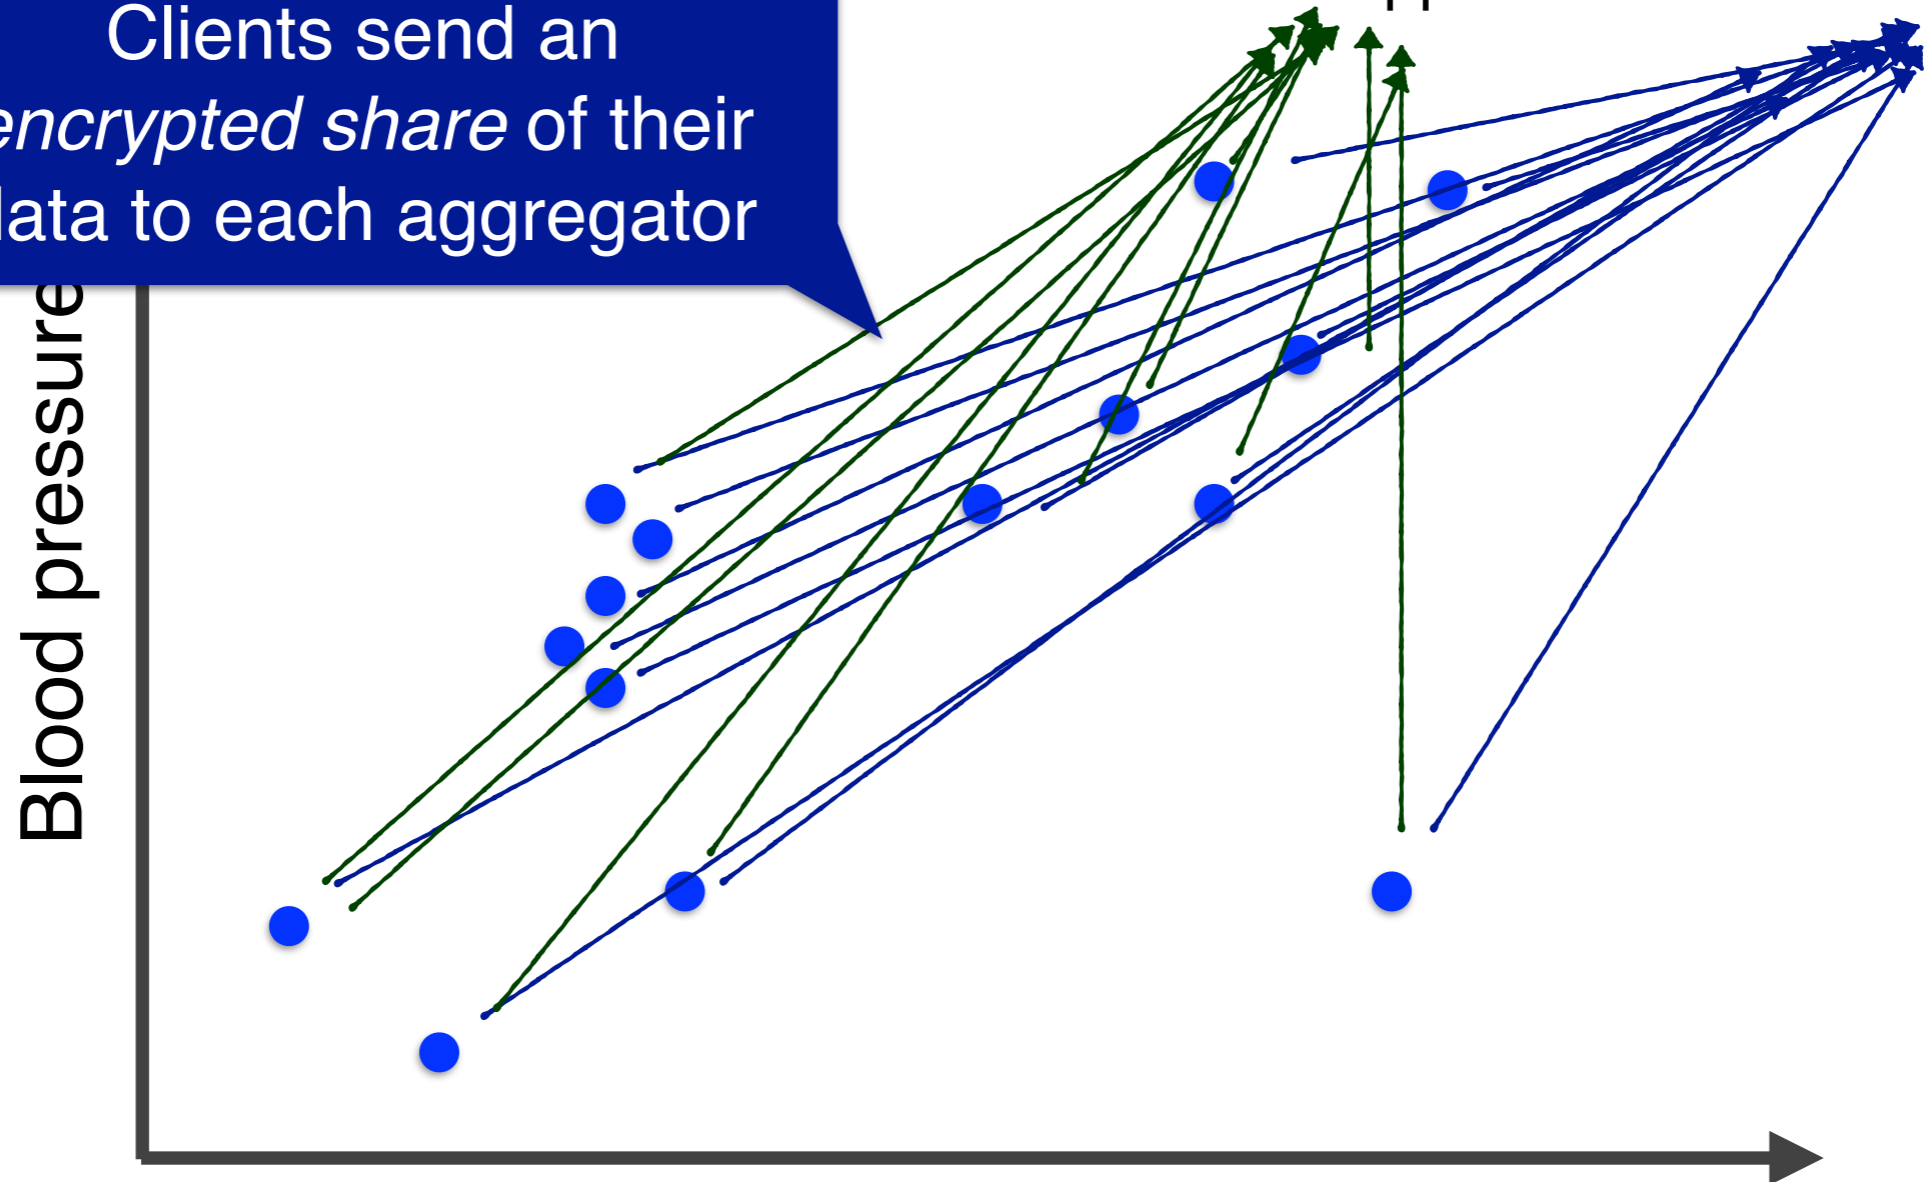
Twitter usage



App store

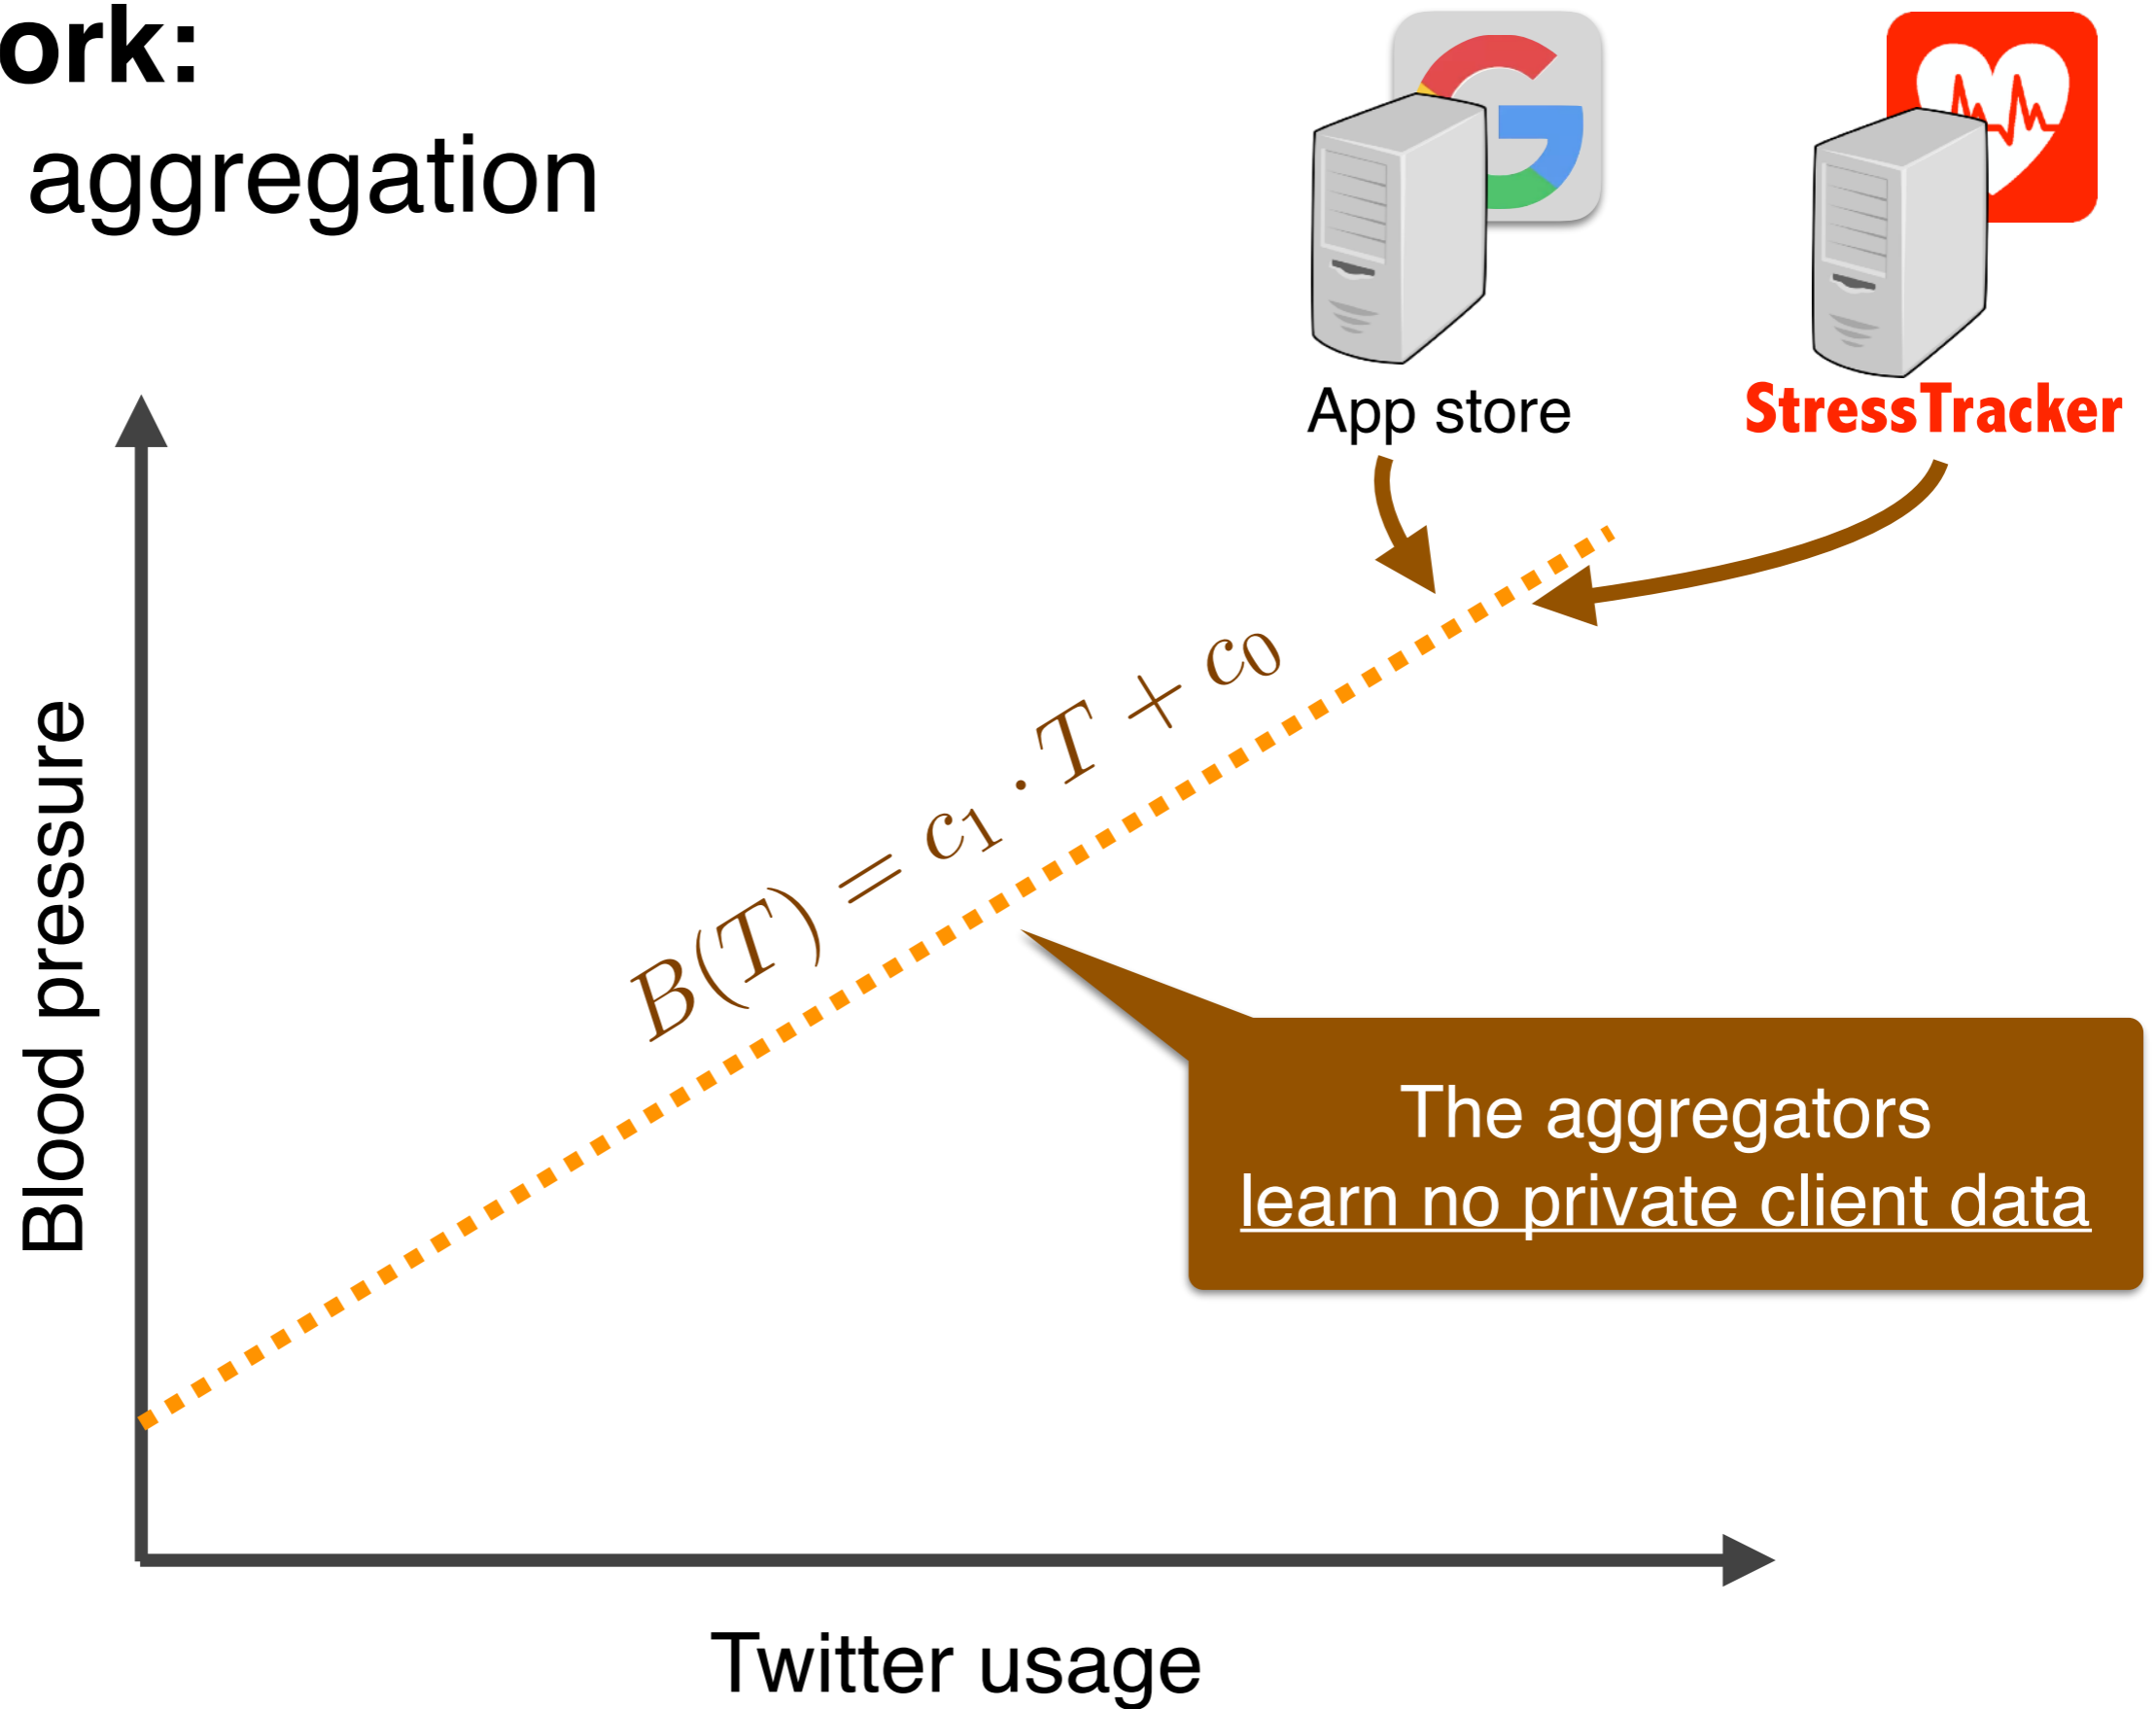


StressTracker




# This work:

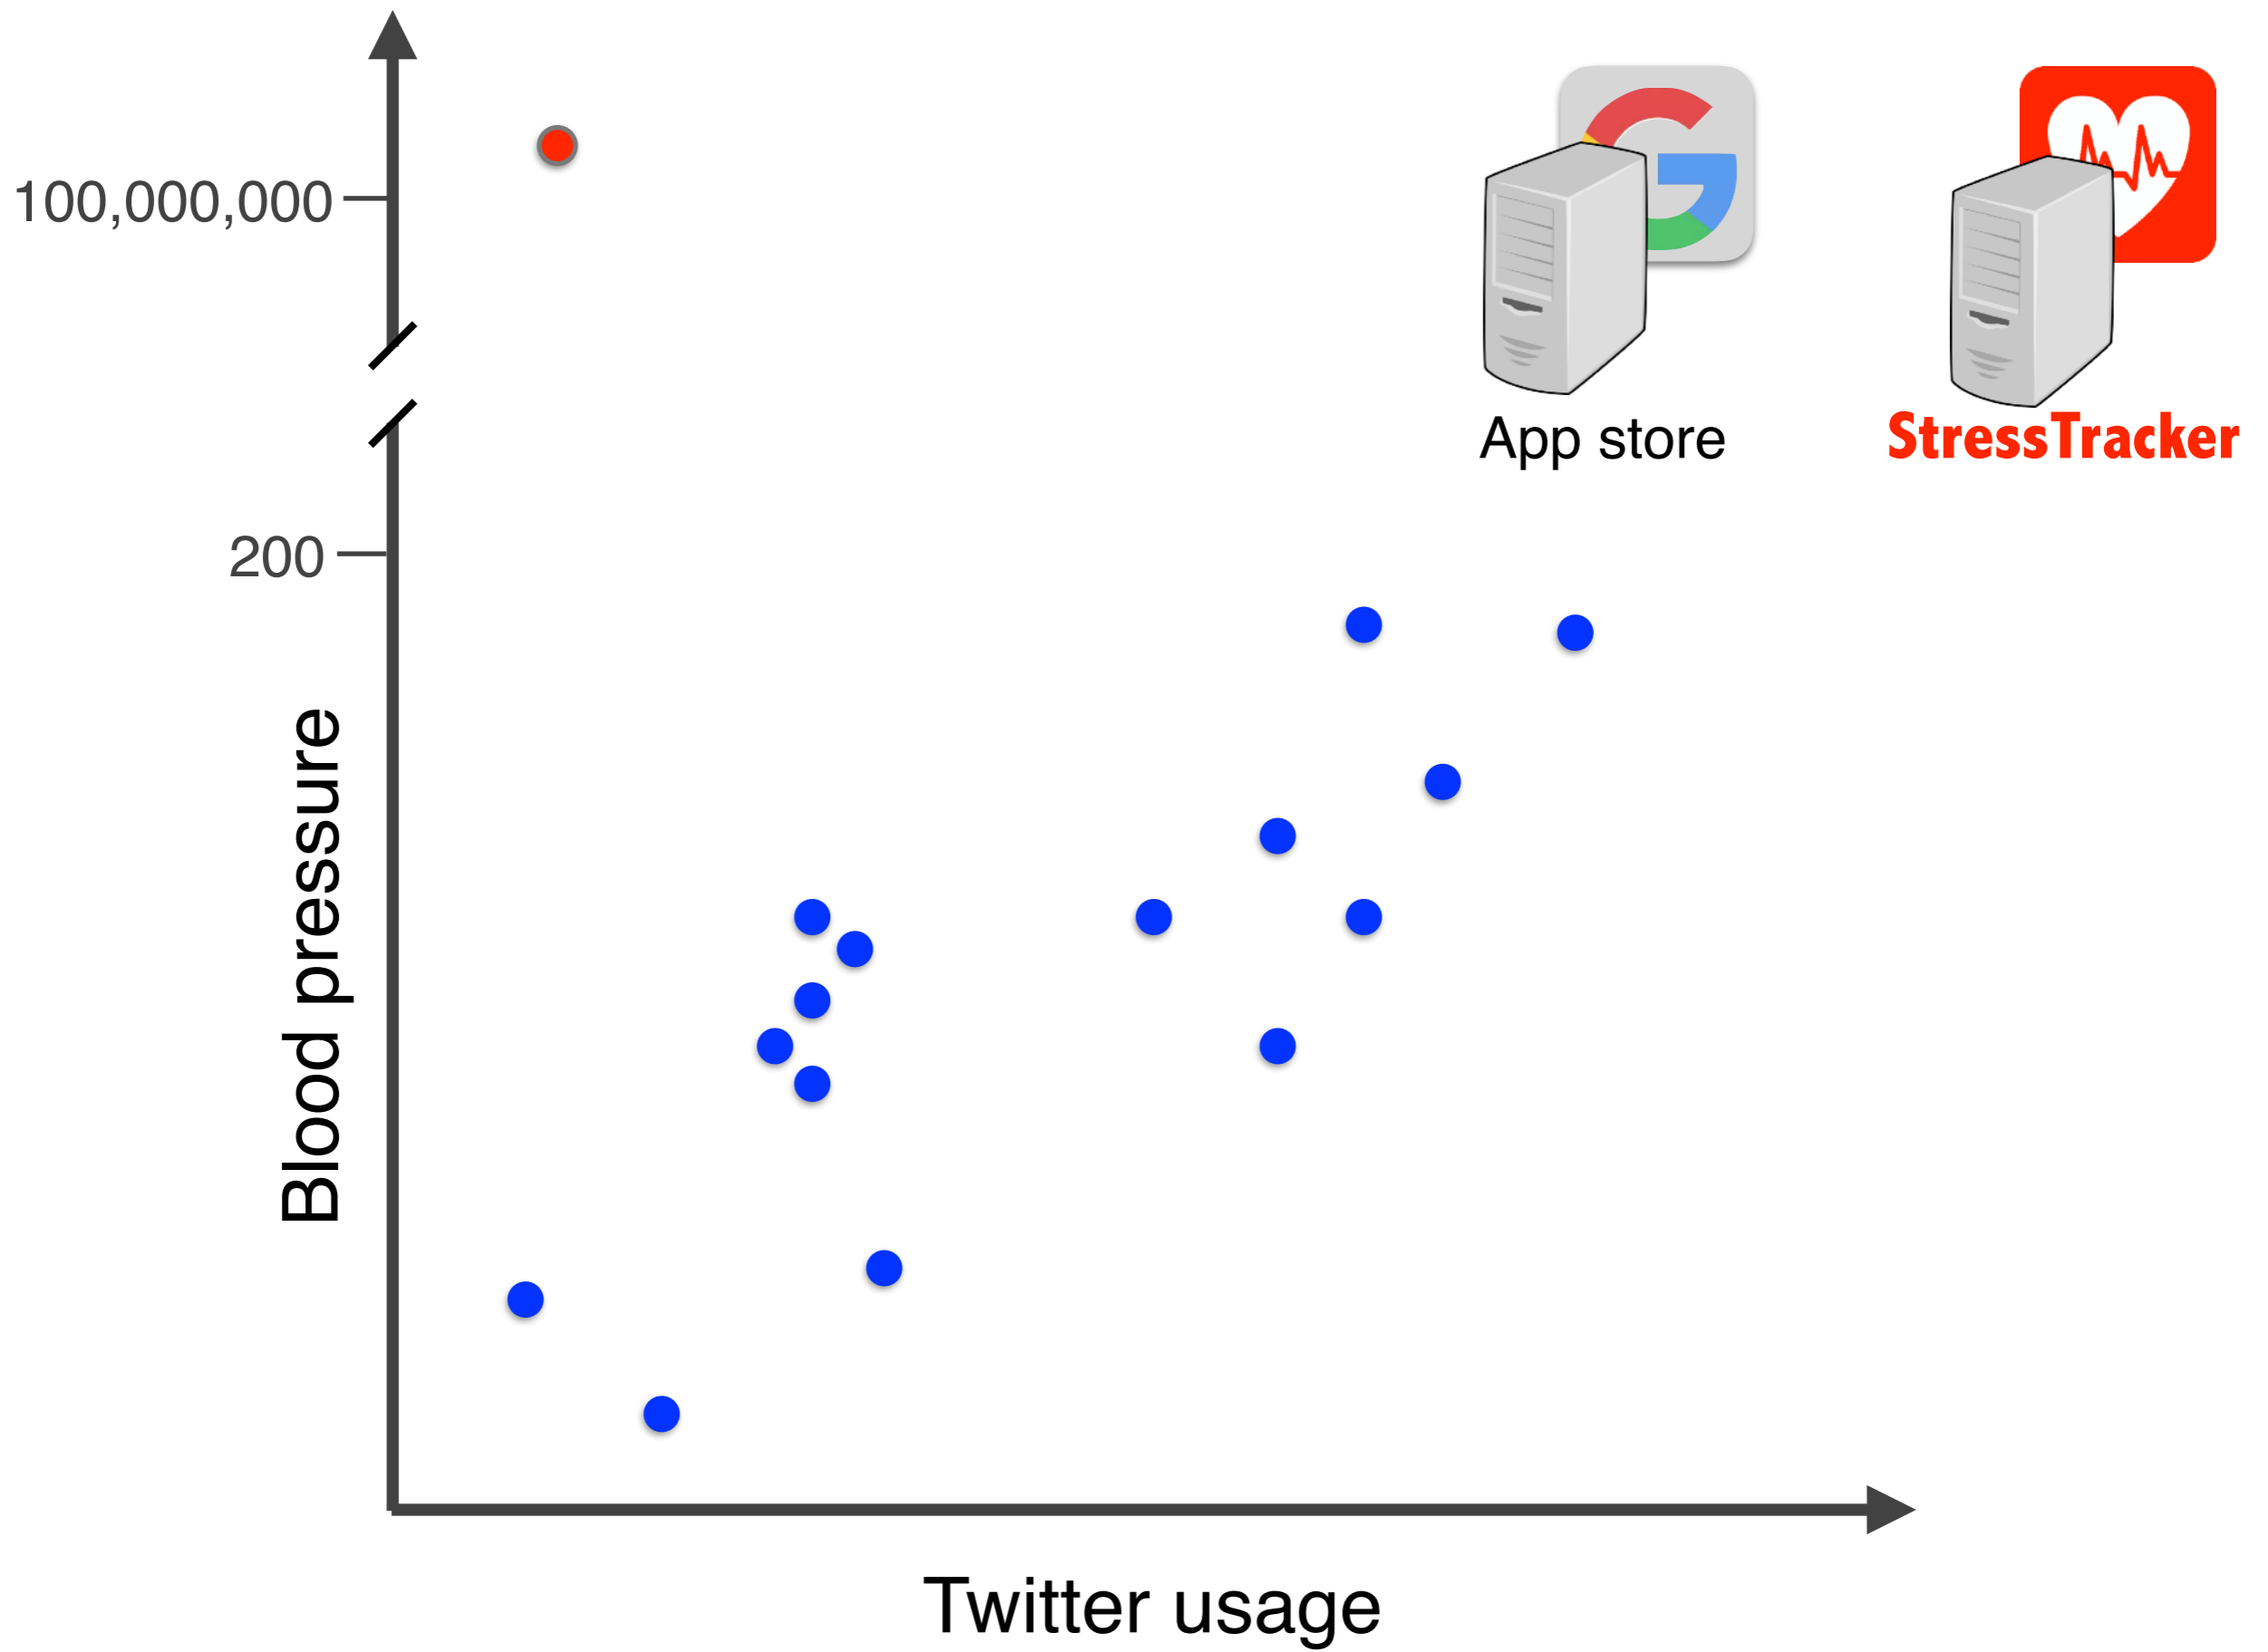
## Private aggregation

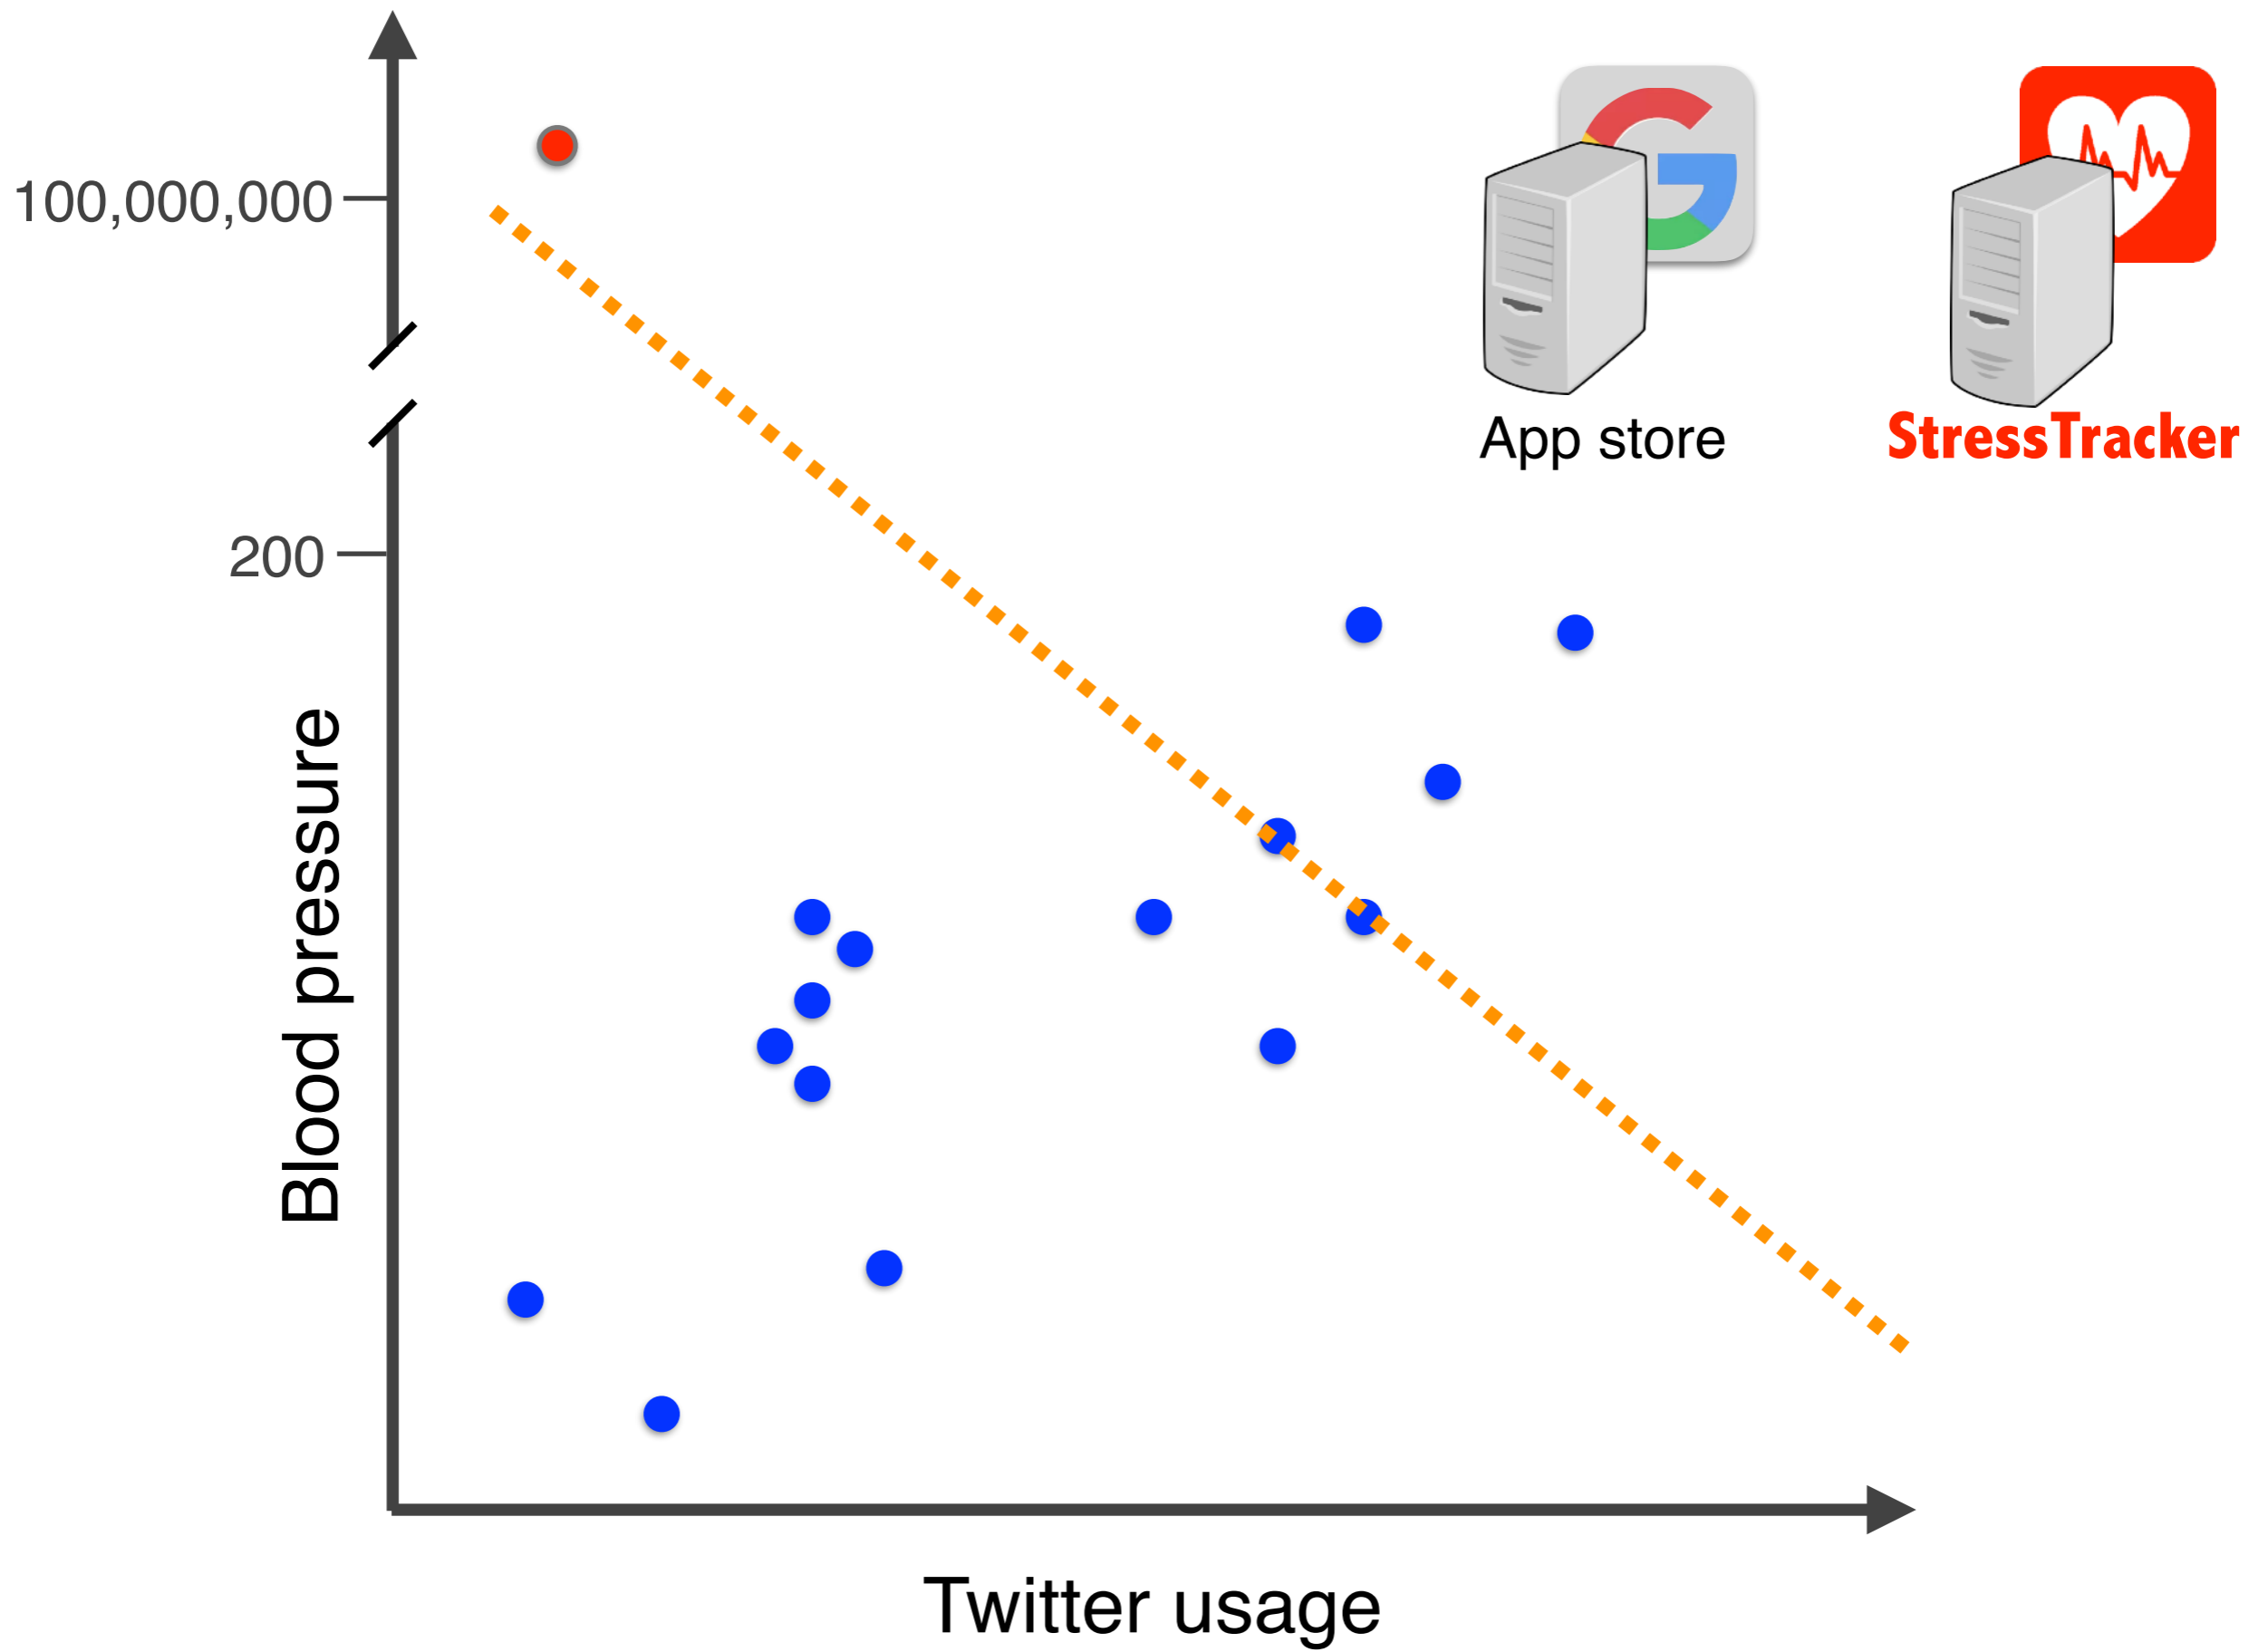


# Private aggregation



- 
- 1. Exact correctness**      If all servers are honest, servers learn  $f(\cdot)$
  - 2. Privacy**                      If one server is honest, servers learn only\*  $f(\cdot)$
  - 3. Robustness**                    Malicious clients have bounded influence
  - 4. Efficiency**                      No public-key crypto (apart from TLS)  
1000s of submissions per second







# Private aggregation



- 
- 1. Exact correctness**      If all servers are honest, servers learn  $f(\cdot)$
  - 2. Privacy**                      If one server is honest, servers learn only\*  $f(\cdot)$
  - 3. Robustness**                      Malicious clients have bounded influence
  - 4. Efficiency**                      No public-key crypto (apart from TLS)

**Prio is the first system to achieve all four.**

# Private aggregation



## 1. Exact correctness

If all servers are honest, servers learn  $f(\cdot)$

## 2. Privacy

If one server is honest, servers learn only\*  $f(\cdot)$

## 3. Robustness

Malicious clients have

## 4. Efficiency

No public-key crypto

...and Prio supports a wide range of aggregation functions  $f(\cdot)$

**Prio is the first system to achieve all four.**

# Private aggregation



- 1. Exact correctness** If all servers are honest, servers learn  $f(\cdot)$
- 2. Privacy** If one server is honest, servers learn only\*  $f(\cdot)$
- 3. Robustness** Malicious clients have bounded influence
- 4. Efficiency** No public-key crypto (apart from TLS)

**Prio is the first system to achieve all four.**

# Contributions

## 1. **Secret-shared non-interactive proofs (SNIPs)**

- Client proves that its encoded submission is well-formed
- We do not need the power of traditional “heavy” crypto tools

## 2. **Aggregatable encodings**

Can compute sums privately  $\implies$  Can compute  $f(\cdot)$  privately  
...for many  $f$ 's of interest

# Contributions

## 1. **Secret-shared non-interactive proofs (SNIPs)**

- Client proves that its encoded submission is well-formed
- We do not need the power of traditional “heavy” crypto tools

## 2. **Aggregatable encodings**

Can compute sums privately  $\implies$  Can compute  $f(\cdot)$  privately  
...for many  $f$ 's of interest

See the paper

# Related systems

- **Additively homomorphic encryption**  
P4P (2010), Private stream aggregation (2011), Grid aggregation (2011), PDDP (2012), SplitX (2013), PrivEx (2014), PrivCount (2016), Succinct sketches (2016), ...
- **Multi-party computation** [GMW87], [BGW88]  
FairPlay (2004), Brickell-Shmatikov (2006), FairplayMP (2008), SEPIA (2010), Private matrix factorization (2013), JustGarble (2013), ...
- **Anonymous credentials/tokens**  
VPriv (2009), PrivStats (2011), ANONIZE (2014), ...
- **Randomized response** [W65], [DMNS06], [D06]  
RAPPOR (2014, 2016), ...

**Prio is the first system to achieve**  
exact correctness, privacy, robustness, efficiency.

# Outline

- **Background: The private aggregation problem**
- A straw-man solution for private sums
- Providing robustness with SNIPs
- Evaluation
- Discussion: Real-world considerations

# Outline

- Background: The private aggregation problem
- **A straw-man solution for private sums**
- Providing robustness with SNIPs
- Evaluation
- Discussion: Real-world considerations



# Warm-up: Computing private sums

# Warm-up: Computing private sums

- Every device  $i$  holds a value  $x_i$

- We want to compute

$$f(x_1, \dots, x_N) = x_1 + \dots + x_N$$

without learning any users' private value  $x_i$ .

# Warm-up: Computing private sums

- Every device  $i$  holds a value  $x_i$

- We want to compute

$$f(x_1, \dots, x_N) = x_1 + \dots + x_N$$

without learning any users' private value  $x_i$ .

---

**Example:** Privately measuring traffic congestion.

$$\begin{aligned} x_i &= 1 && \text{if user } i \text{ is on the Bay Bridge} \\ &= 0 && \text{otherwise} \end{aligned}$$



The sum  $x_1 + \dots + x_N$  yields the number of app users on the Bay Bridge.

# Private sums: A “straw-man” scheme

[Chaum88], [BGW88], ...  
[KDK11] [DFKZ13] [PrivEx14] ...

Server A



Server B




Server C



# Private sums: A “straw-man” scheme

[Chaum88], [BGW88], ...  
[KDK11] [DFKZ13] [PrivEx14] ...




Assume that the servers are  
non-colluding.

Equivalently: that at least one  
server is honest.

# Private sums: A “straw-man” scheme

[Chaum88], [BGW88], ...  
[KDK11] [DFKZ13] [PrivEx14] ...




Assume that the servers are  
non-colluding.

Equivalently: that at least one  
server is honest.




# Private sums: A “straw-man” scheme

[Chaum88], [BGW88], ...  
[KDK11] [DFKZ13] [PrivEx14] ...




Assume that the servers are  
non-colluding.

Equivalently: that at least one  
server is honest.



# Private sums: A “straw-man” scheme

[Chaum88], [BGW88], ...  
[KDK11] [DFKZ13] [PrivEx14] ...



Assume that the servers are  
non-colluding.

Equivalently: that at least one  
server is honest.

Spain



Germany



Iceland





# Private sums: A “straw-man” scheme

[Chaum88], [BGW88], ...  
[KDK11] [DFKZ13] [PrivEx14] ...

Server A




Server B



Server C



# Private sums: A “straw-man” scheme



1





## Secret sharing

Pick three random “shares” that sum to 1.

$$1 = 15 + (-12) + (-2) \quad (\text{mod } 31)$$

Need all three shares to recover the shared value.

# Private sums: A “straw-man” scheme




## Secret sharing

Pick three random “shares” that sum to 1.

$$1 = 15 + (-12) + (-2) \quad (\text{mod } 31)$$

Need all three shares to recover the shared value.

# Private sums: A “straw-man” scheme



15    -12    -2

1



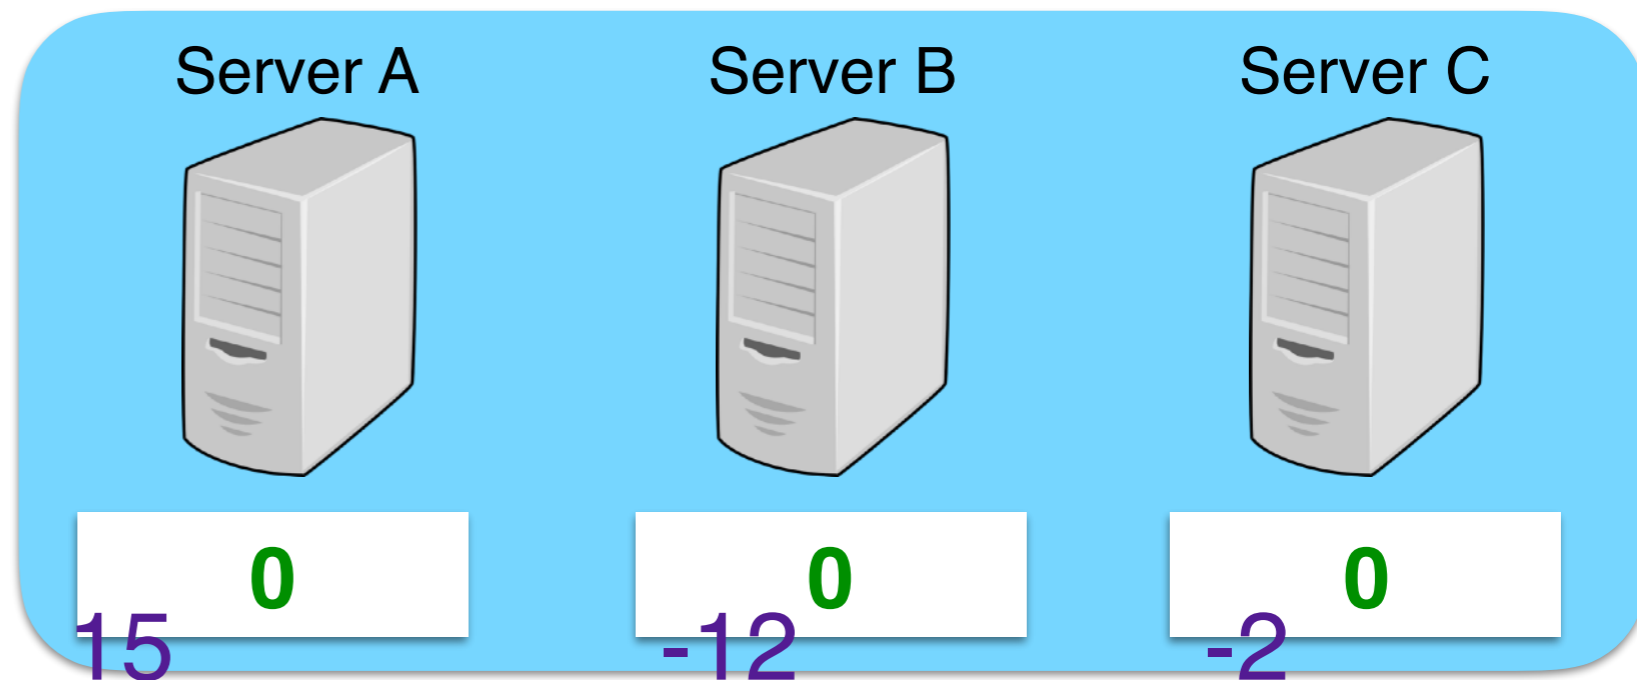
## Secret sharing

Pick three random “shares” that sum to 1.

$$1 = 15 + (-12) + (-2) \quad (\text{mod } 31)$$

Need all three shares to recover the shared value.


# Private sums: A “straw-man” scheme



1




# Private sums: A “straw-man” scheme



1




# Private sums: A “straw-man” scheme



0



# Private sums: A “straw-man” scheme




$$0 = (-10) + 7 + 3$$






# Private sums: A “straw-man” scheme



$$0 = (-10) + 7 + 3$$




# Private sums: A “straw-man” scheme




-10    7    3  
↖    0 = (-10) + 7 + 3



# Private sums: A “straw-man” scheme




# Private sums: A “straw-man” scheme



0




# Private sums: A “straw-man” scheme




0




# Private sums: A “straw-man” scheme



# Private sums: A “straw-man” scheme




Private sums:  
A “straw-man”  
scheme






# Private sums: A “straw-man” scheme




$$S_A + S_B + S_C = 15 + -10 + \dots$$

# Private sums: A “straw-man” scheme



$$\begin{aligned} S_A + S_B + S_C &= 15 + -10 + \dots \\ &= 1 + 0 + \dots + 1 \end{aligned}$$


# Private sums: A “straw-man” scheme



$$\begin{aligned} S_A + S_B + S_C &= 15 + -10 + \dots \\ &= 1 + 0 + \dots + 1 \end{aligned}$$

Servers learn the  
sum of client values  
and learn *nothing else*.


# Private sums: A “straw-man” scheme



$$\begin{aligned} S_A + S_B + S_C &= 15 + -10 + \dots \\ &= 1 + 0 + \dots + 1 \end{aligned}$$

Servers learn the  
sum of client values  
and learn *nothing else*.

# Private sums: A “straw-man” scheme



$$\begin{aligned} S_A + S_B + S_C &= 15 + -10 + \dots \\ &= 1 + 0 + \dots + 1 \end{aligned}$$

Learn that three phones  
are on the Bay Bridge—  
don't know which three

# Computing private sums

# Computing private sums

**Exact correctness:** If everyone follows the protocol, servers compute the sum of all  $x_i$ s.

**Privacy:** Any proper subset of the servers learns nothing but the sum of the  $x_i$ s.

**Efficiency:** Follows by inspection.

# Computing private sums

**Exact correctness:** If everyone follows the protocol, servers compute the sum of all  $x_i$ s.


**Privacy:** Any proper subset of the servers learns nothing but the sum of the  $x_i$ s.

**Efficiency:** Follows by inspection.

**Robustness:** ???




# Private sums: A “straw-man” scheme



X

# Private sums: A “straw-man” scheme




x is supposed to be  
a 0/1 value




x


# Private sums: A “straw-man” scheme



# Private sums: A “straw-man” scheme




# Private sums: A “straw-man” scheme



An evil client needn't  
follow the rules!



# Private sums: A “straw-man” scheme




An evil client needn't follow the rules!


$$10 + 4 + 7 = 21$$



# Private sums: A “straw-man” scheme




# Private sums: A “straw-man” scheme





# Private sums: A “straw-man” scheme



A single bad client  
can undetectably  
corrupt the sum

Users have  
incentives to cheat

Typical defenses  
(NIZKs) are costly



# Outline


- Background: The private aggregation problem
- **A straw-man solution for private sums**
- Providing robustness with SNIPs
- Evaluation
- Discussion: Real-world considerations

# Outline

- Background: The private aggregation problem
- A straw-man solution for private sums
- **Providing robustness with SNIPs**
- Evaluation
- Discussion: Real-world considerations

# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)




$$x = 1$$




# Contribution 1

Secret-shared  
non-interactive  
proofs (SNIPs)

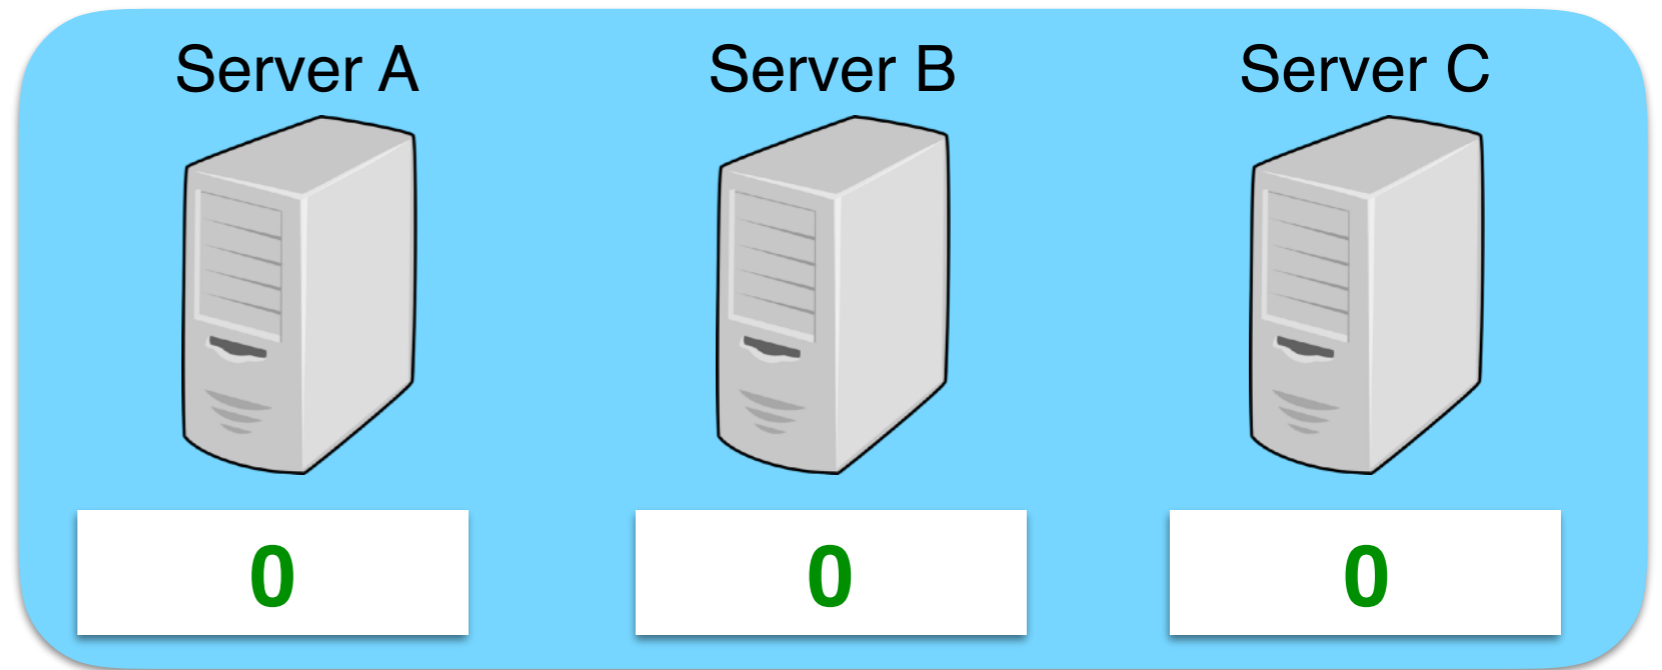


$x = 1$




# Contribution 1

Secret-shared  
non-interactive  
proofs (SNIPs)



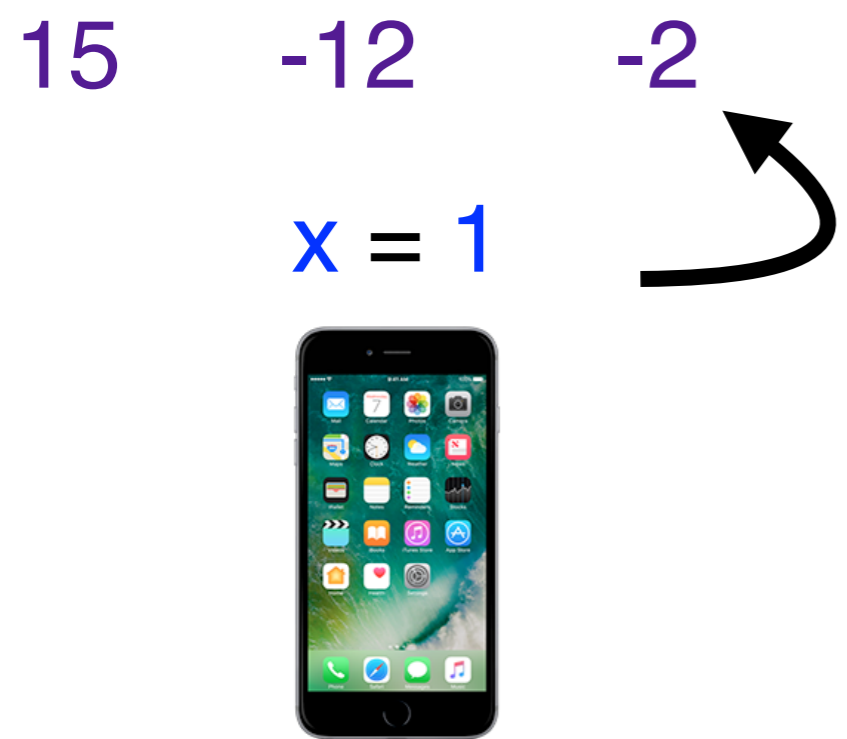

$$15 + (-12) + (-2) = 1$$

$$x = 1$$



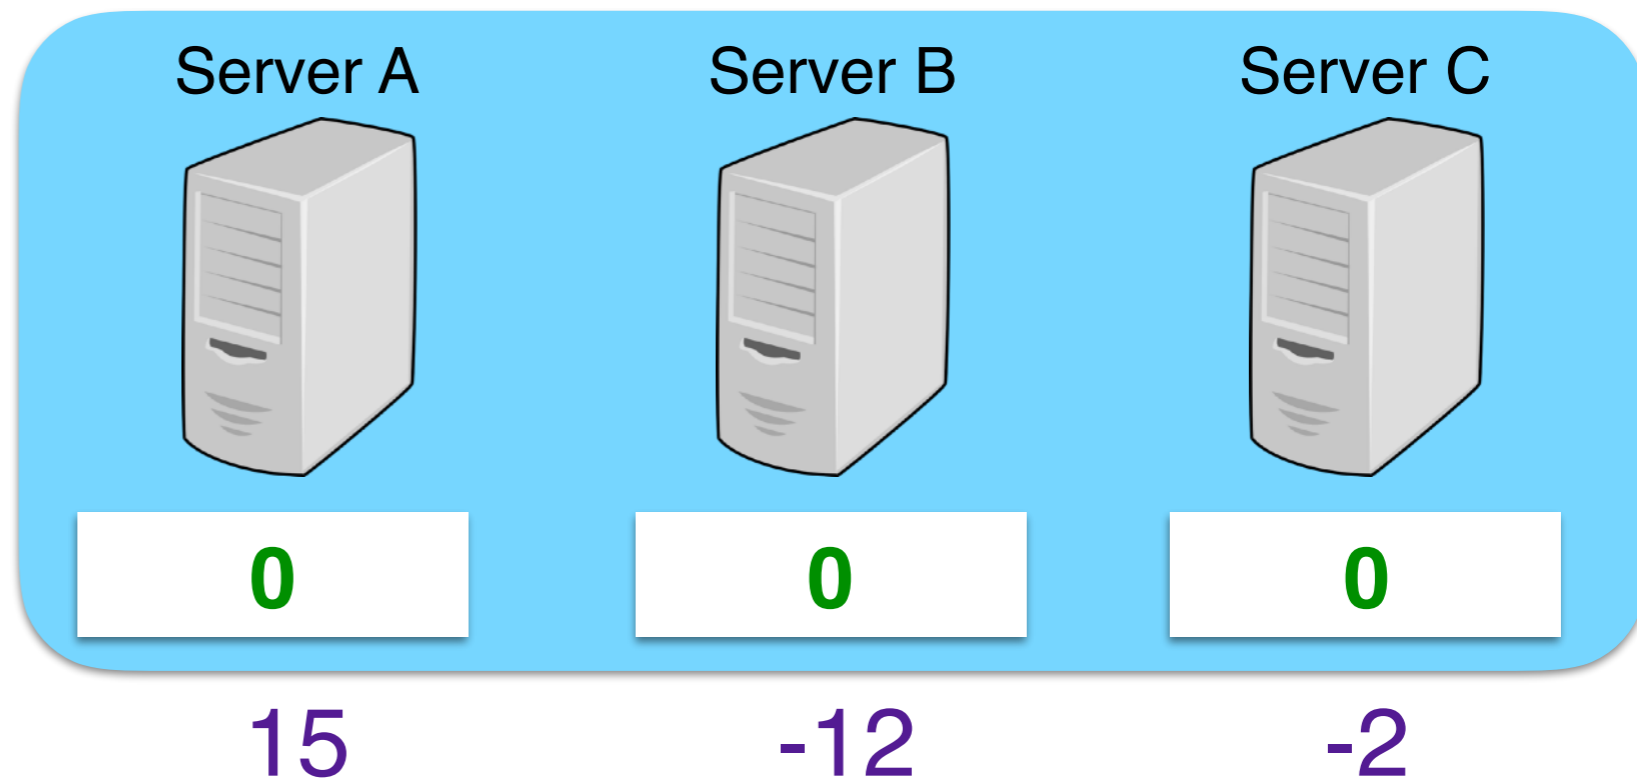
# Contribution 1

Secret-shared  
non-interactive  
proofs (SNIPs)



# Contribution 1

Secret-shared  
non-interactive  
proofs (SNIPs)



$$x = 1$$






# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)

$$x = 1$$



15


-12

-2

In this example, the servers want to ensure that their shares sum to 0 or 1 ...without learning  $x$ .

# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)



$$x = 1$$




### More generally, servers

- hold shares of the client's private value  $x$
- hold an arbitrary public predicate  $\text{Valid}(\cdot)$ 
  - expressed as an arithmetic circuit
- want to test if “ $\text{Valid}(x)$ ” holds, without leaking  $x$

# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)



$x = 1$




### More generally, servers

- hold shares of the client's private value  $x$
- hold an arbitrary public predicate  $\text{Valid}(\cdot)$ 
  - expressed as an arithmetic circuit
- want to test if “ $\text{Valid}(x)$ ”

For our running example:  
 $\text{Valid}(x) = “x \in \{0, 1\}”$

# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)



$$x = 1$$




### More generally, servers

- hold shares of the client's private value  $x$
- hold an arbitrary public predicate  $\text{Valid}(\cdot)$ 
  - expressed as an arithmetic circuit
- want to test if “ $\text{Valid}(x)$ ” holds, without leaking  $x$

# Contribution 1

Secret-shared  
non-interactive  
proofs (SNIPs)




$x = 1$



# Contribution 1


## Secret-shared non-interactive proofs (SNIPs)



$x = 1$

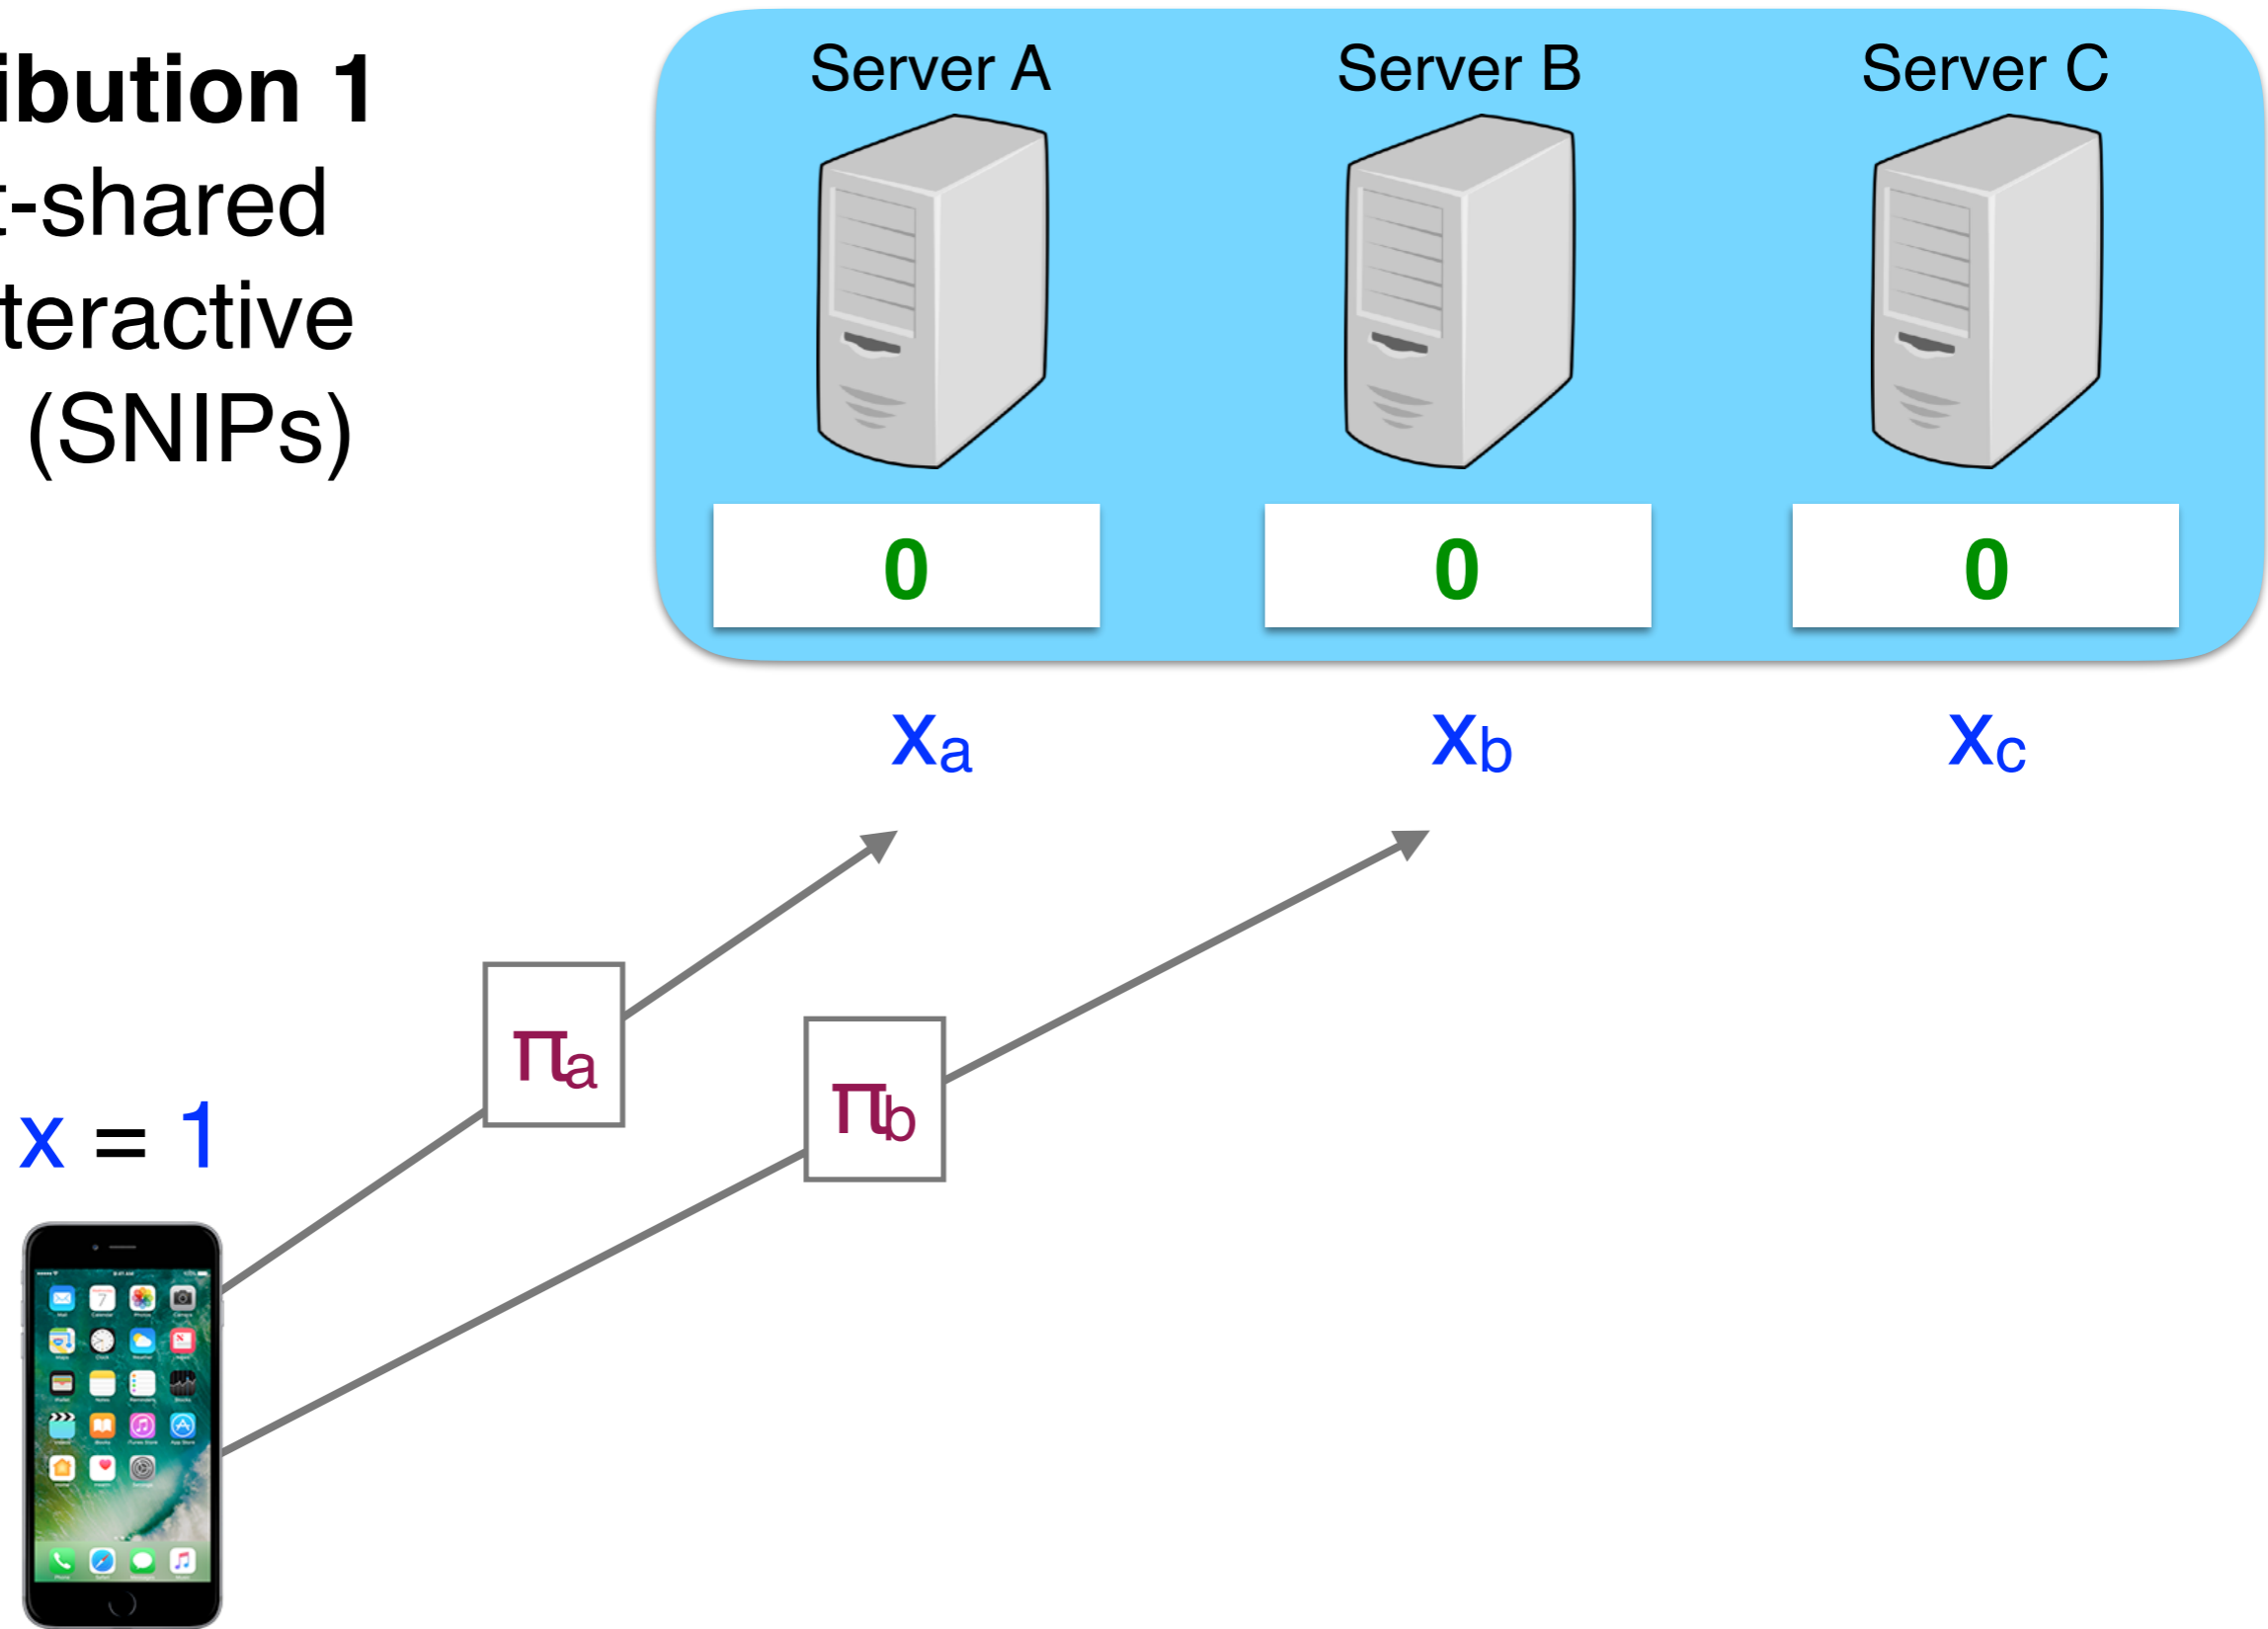


$\pi_a$



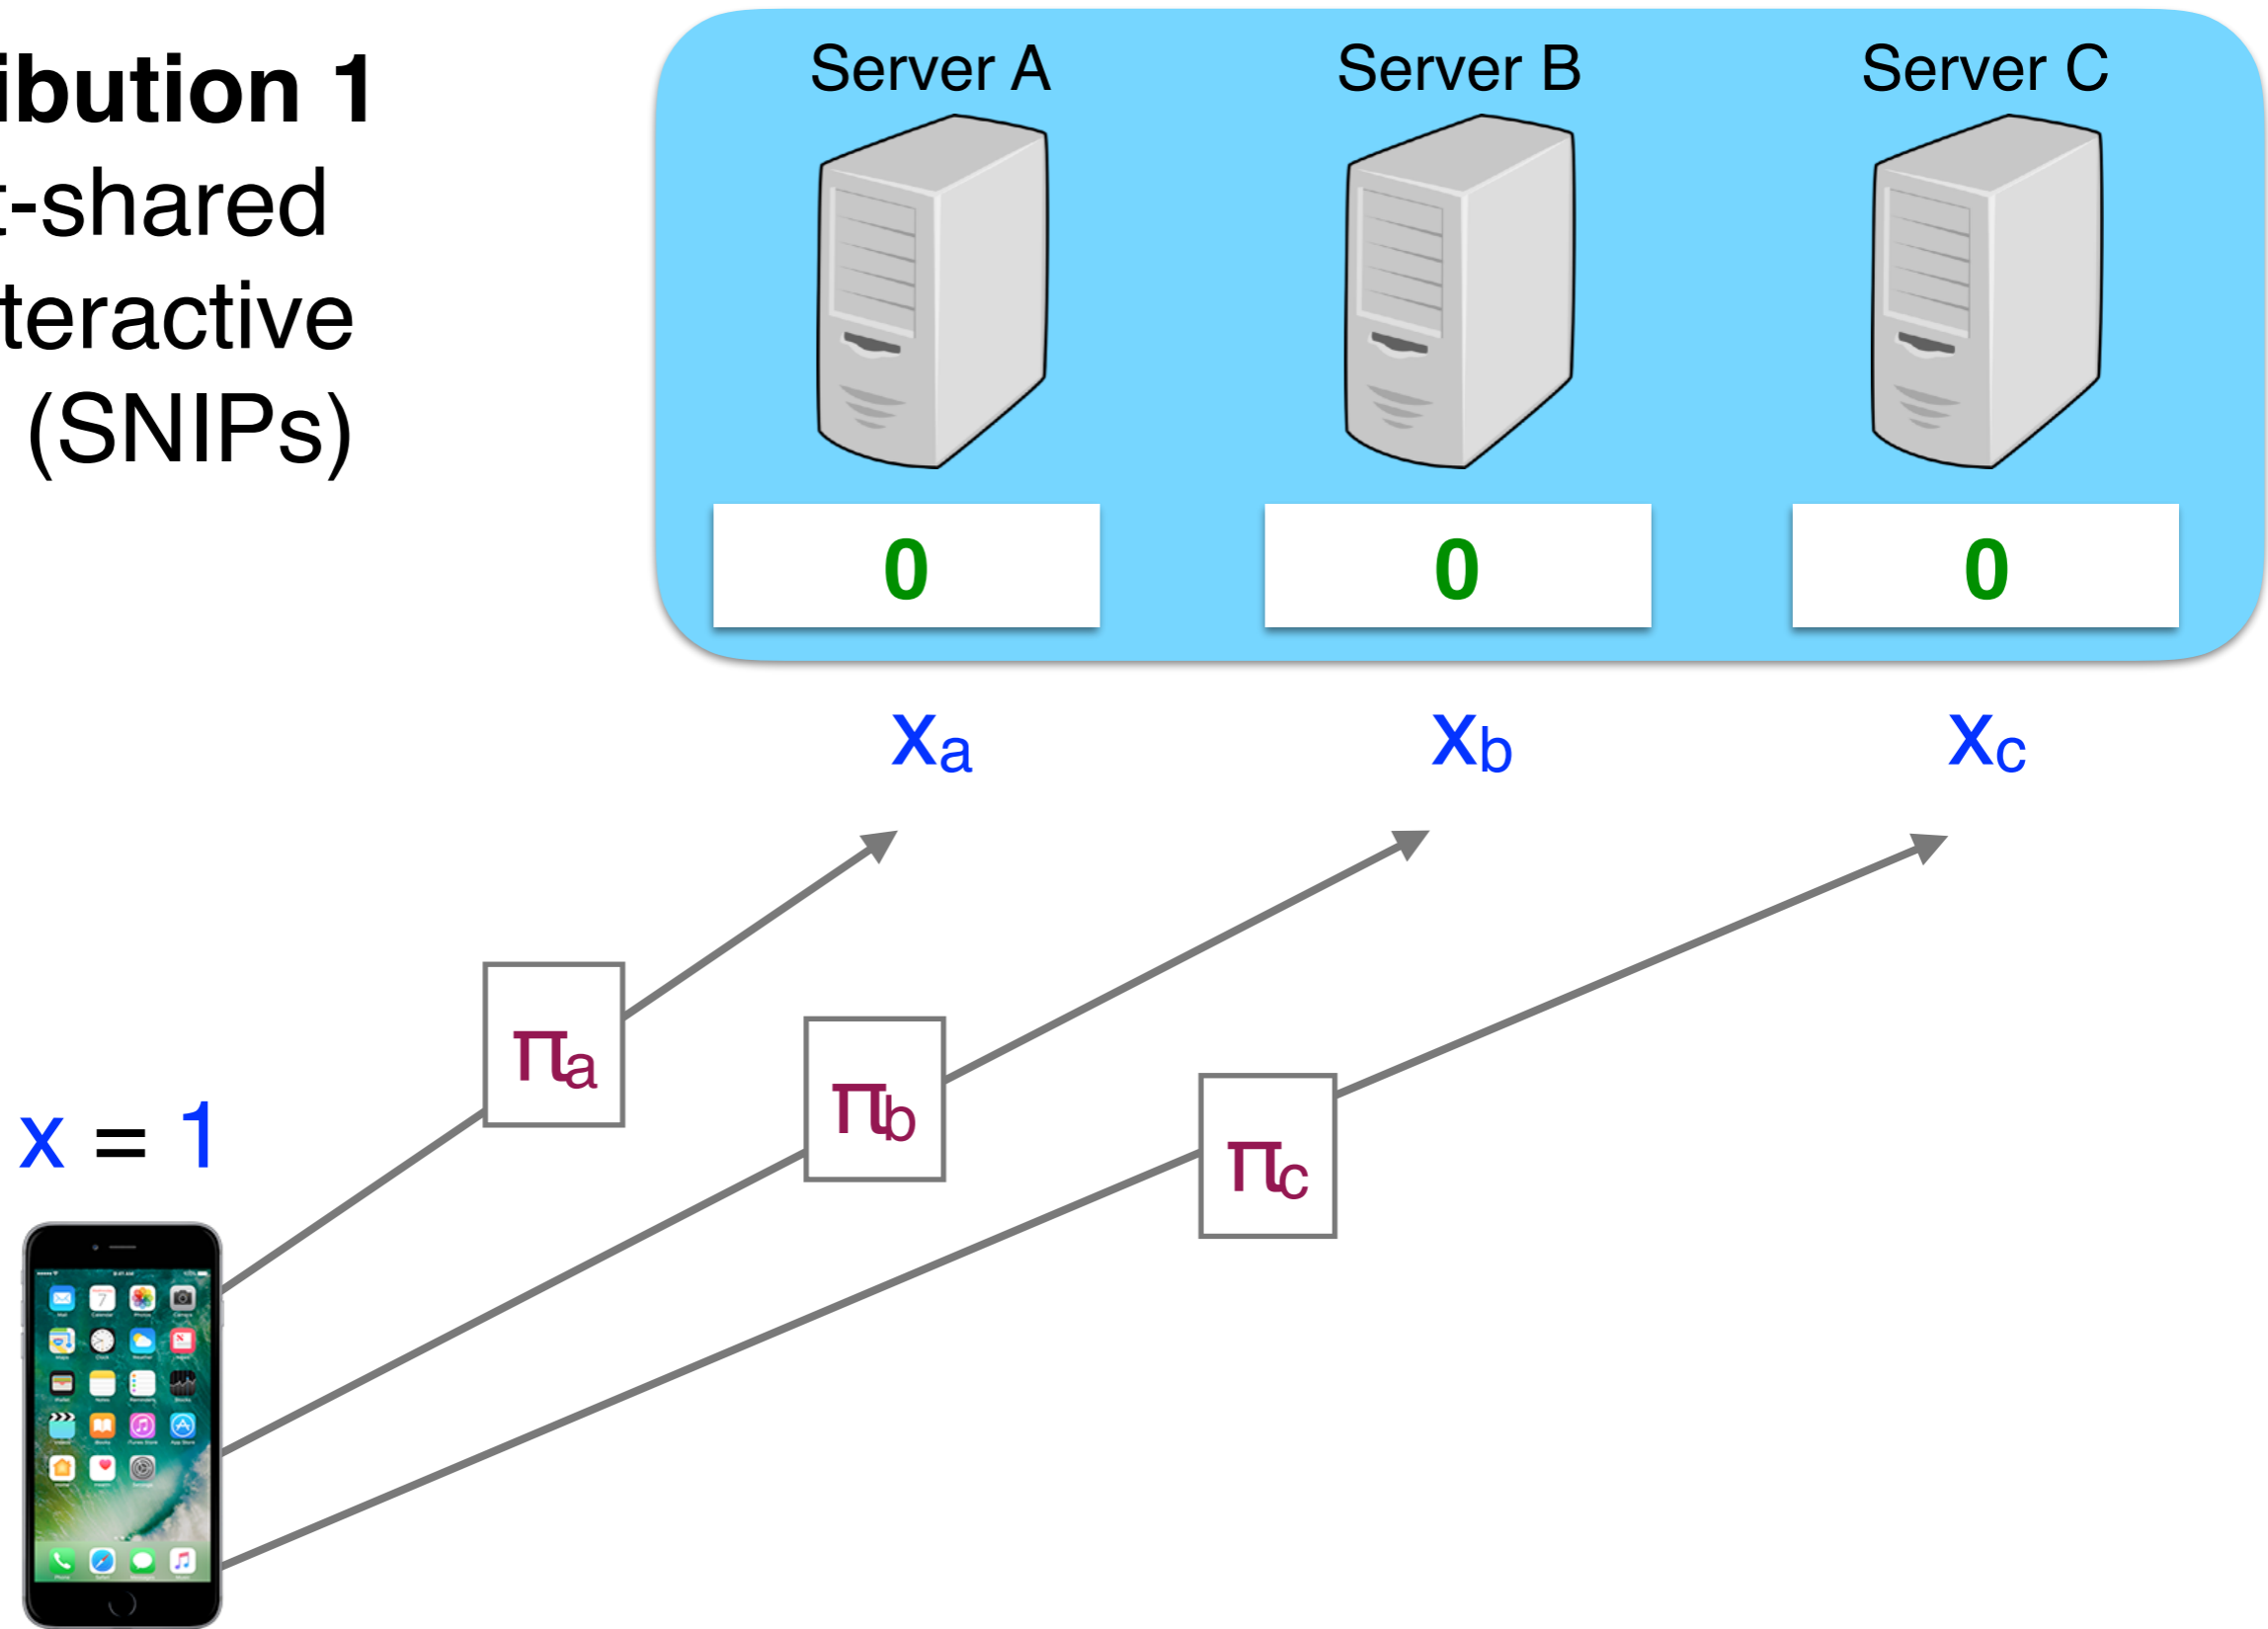
# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)



# Contribution 1


## Secret-shared non-interactive proofs (SNIPs)





# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)

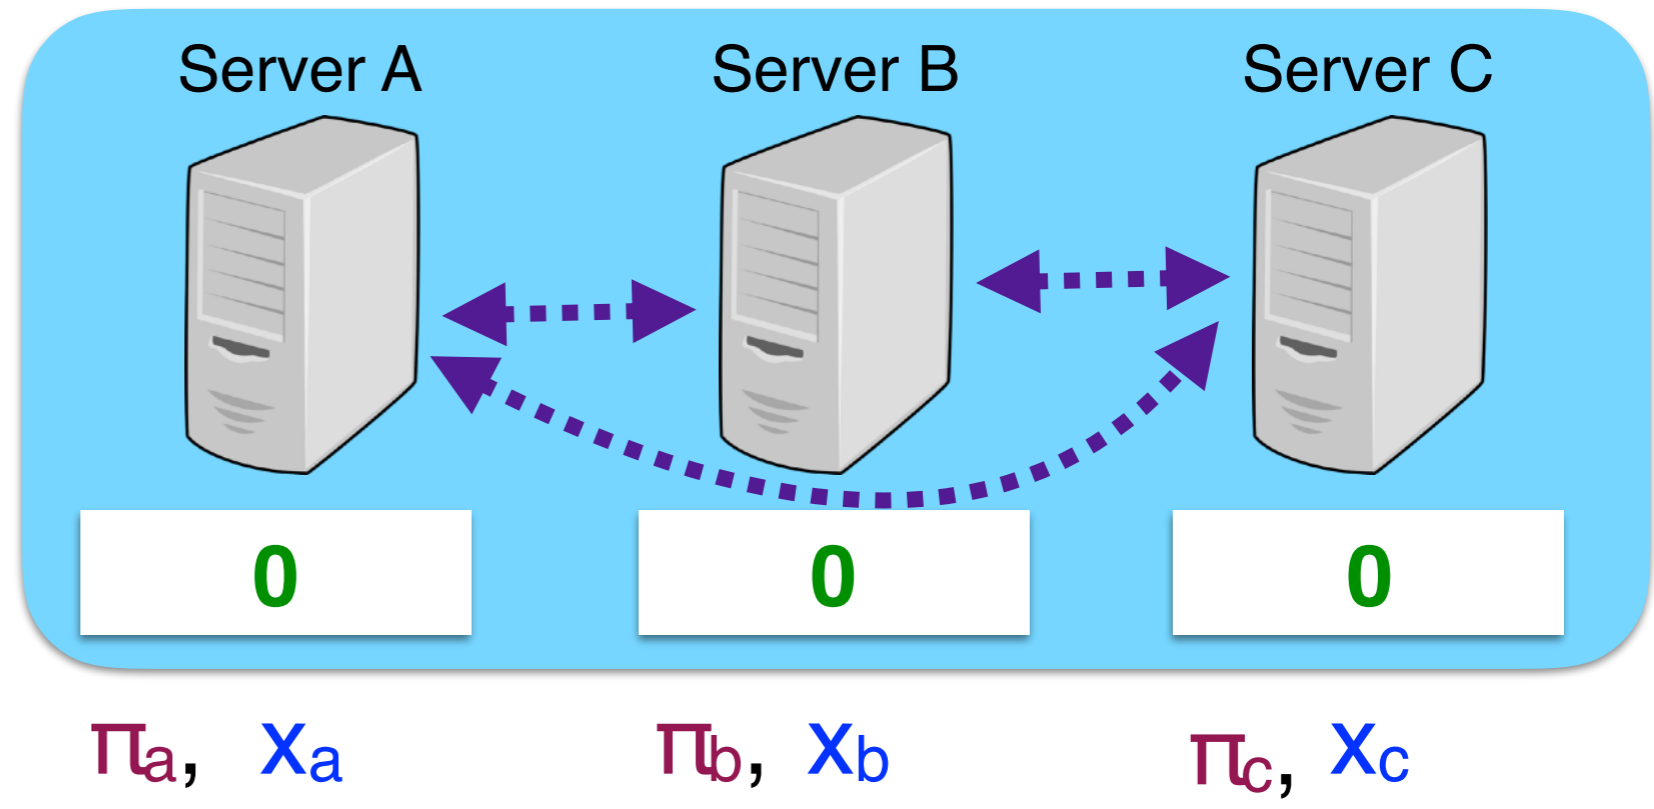


$x = 1$



# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)

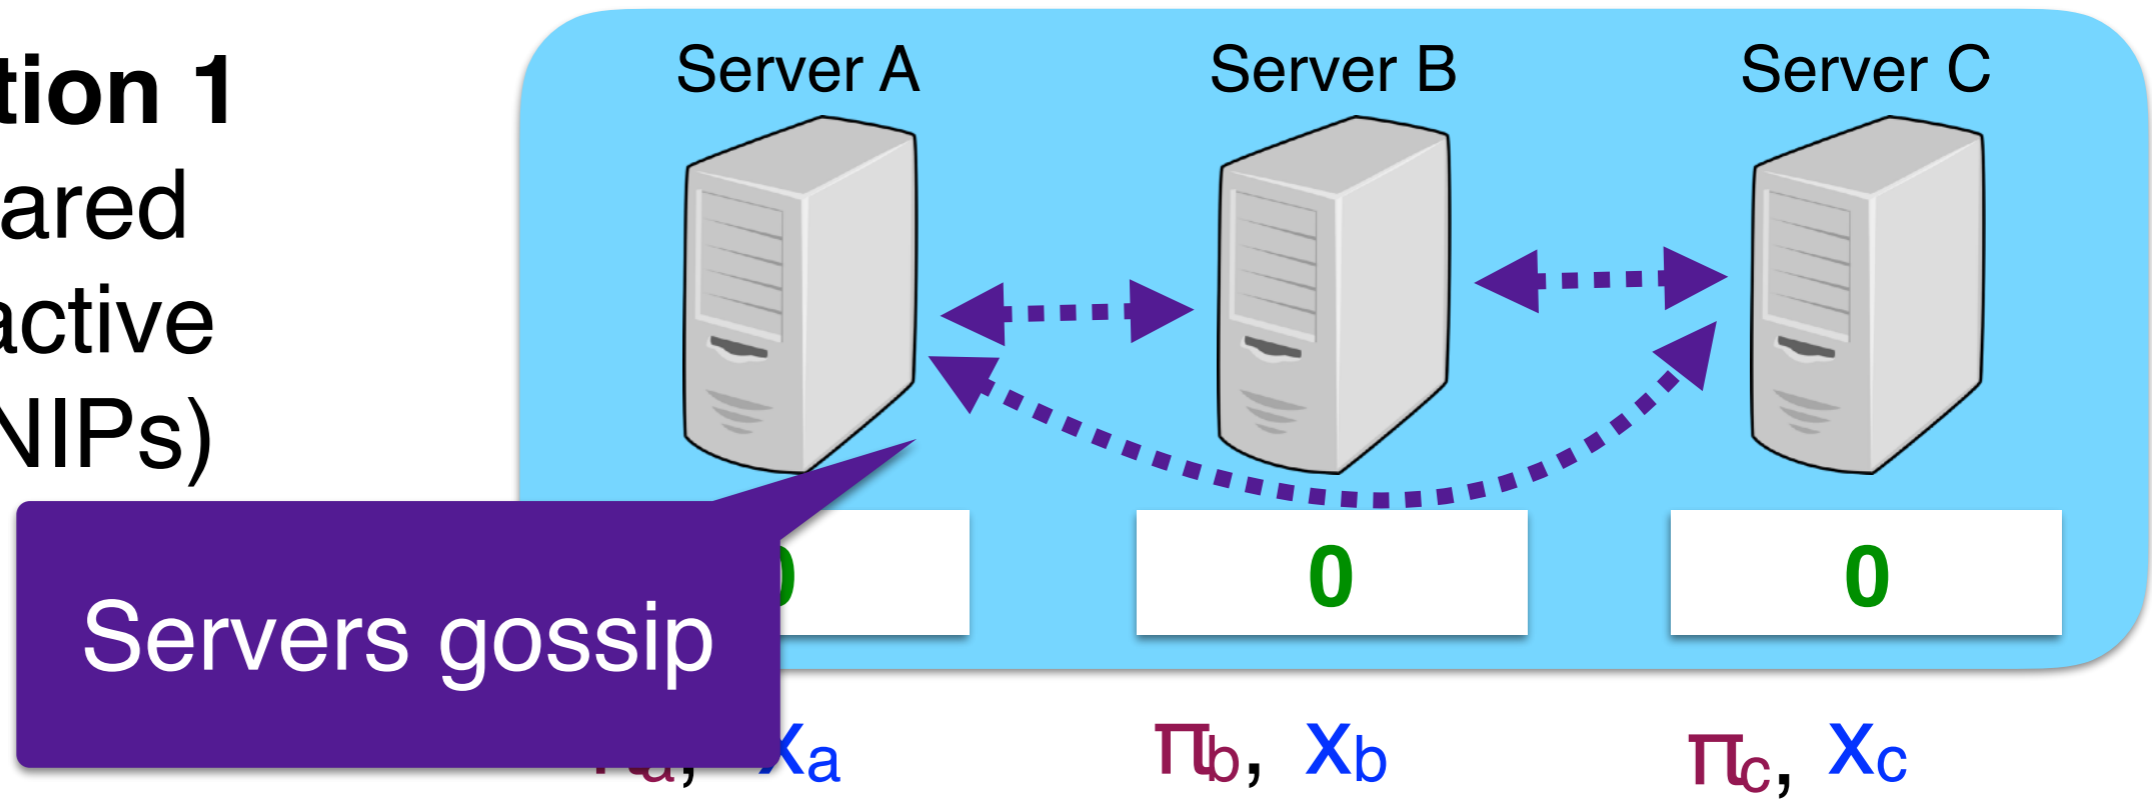


$x = 1$



# Contribution 1

Secret-shared  
non-interactive  
proofs (SNIPs)




$x = 1$



# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)

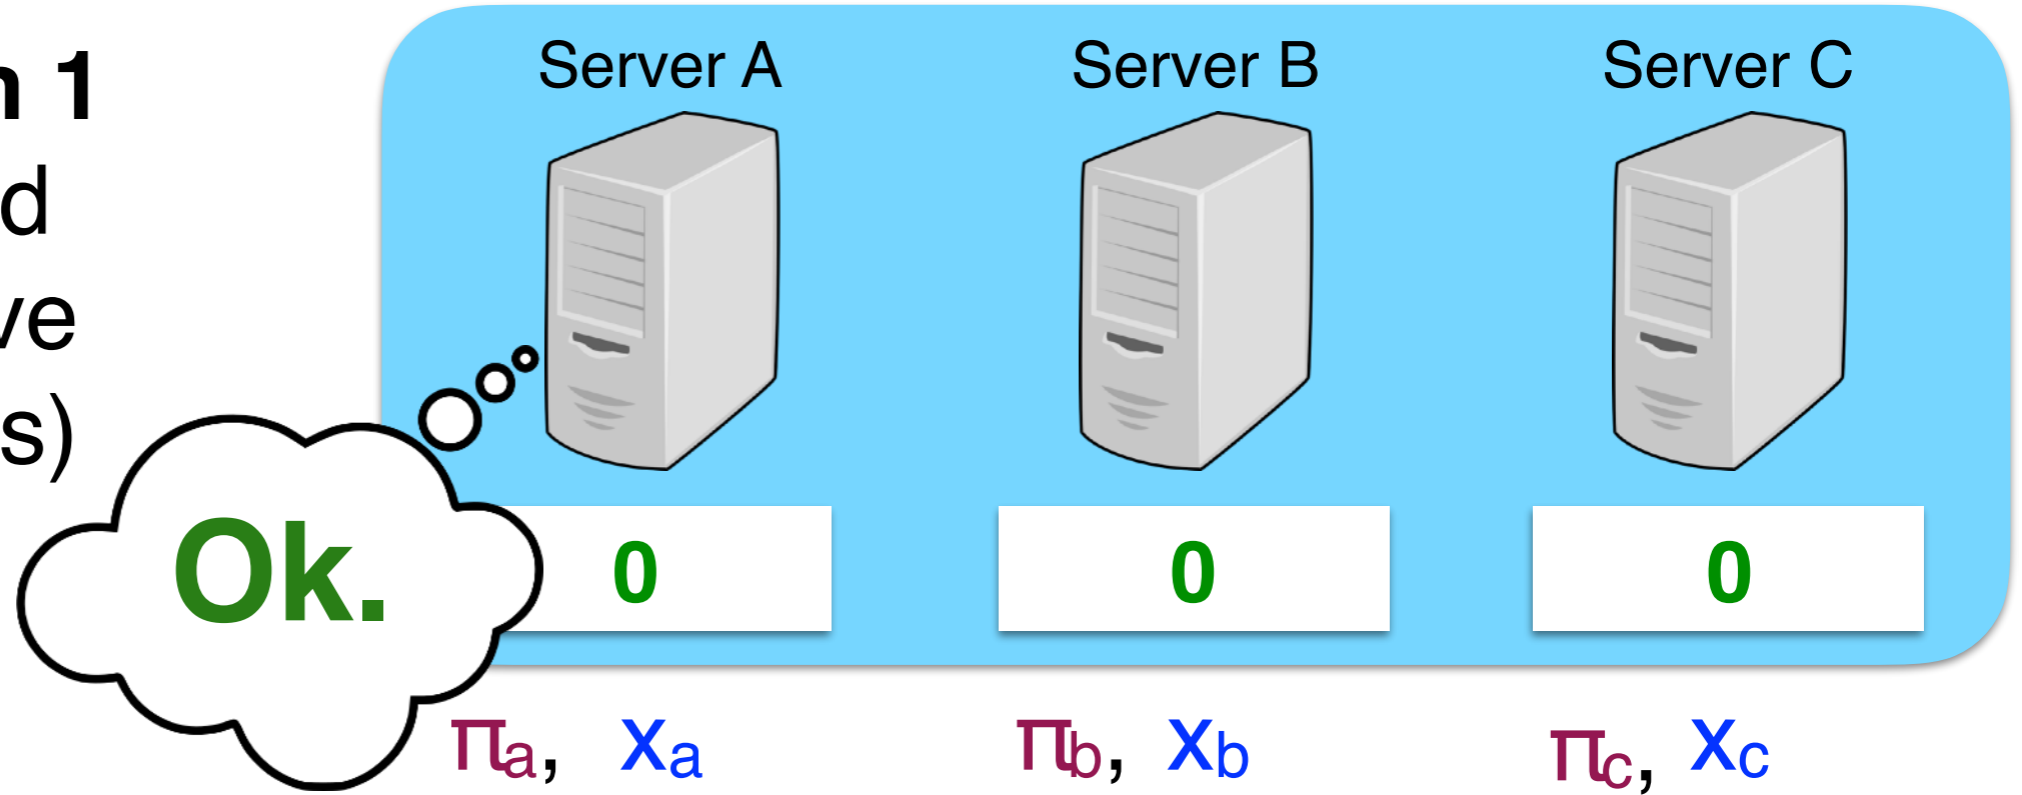


$x = 1$



# Contribution 1

Secret-shared  
non-interactive  
proofs (SNIPs)

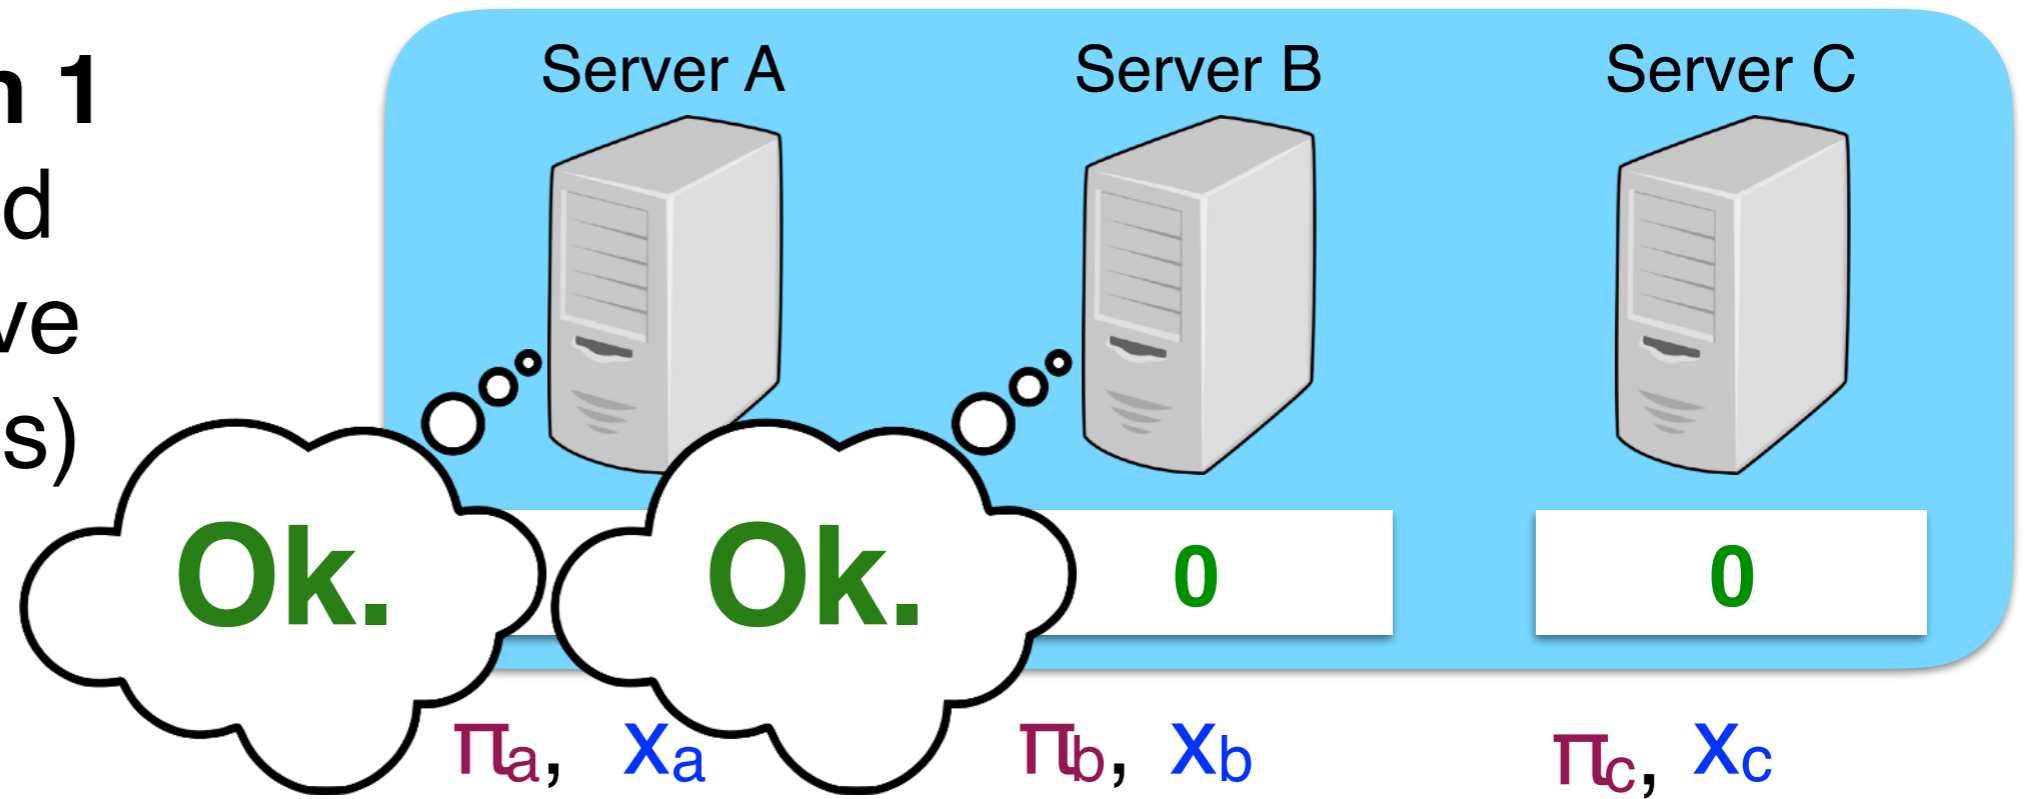


$x = 1$



# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)

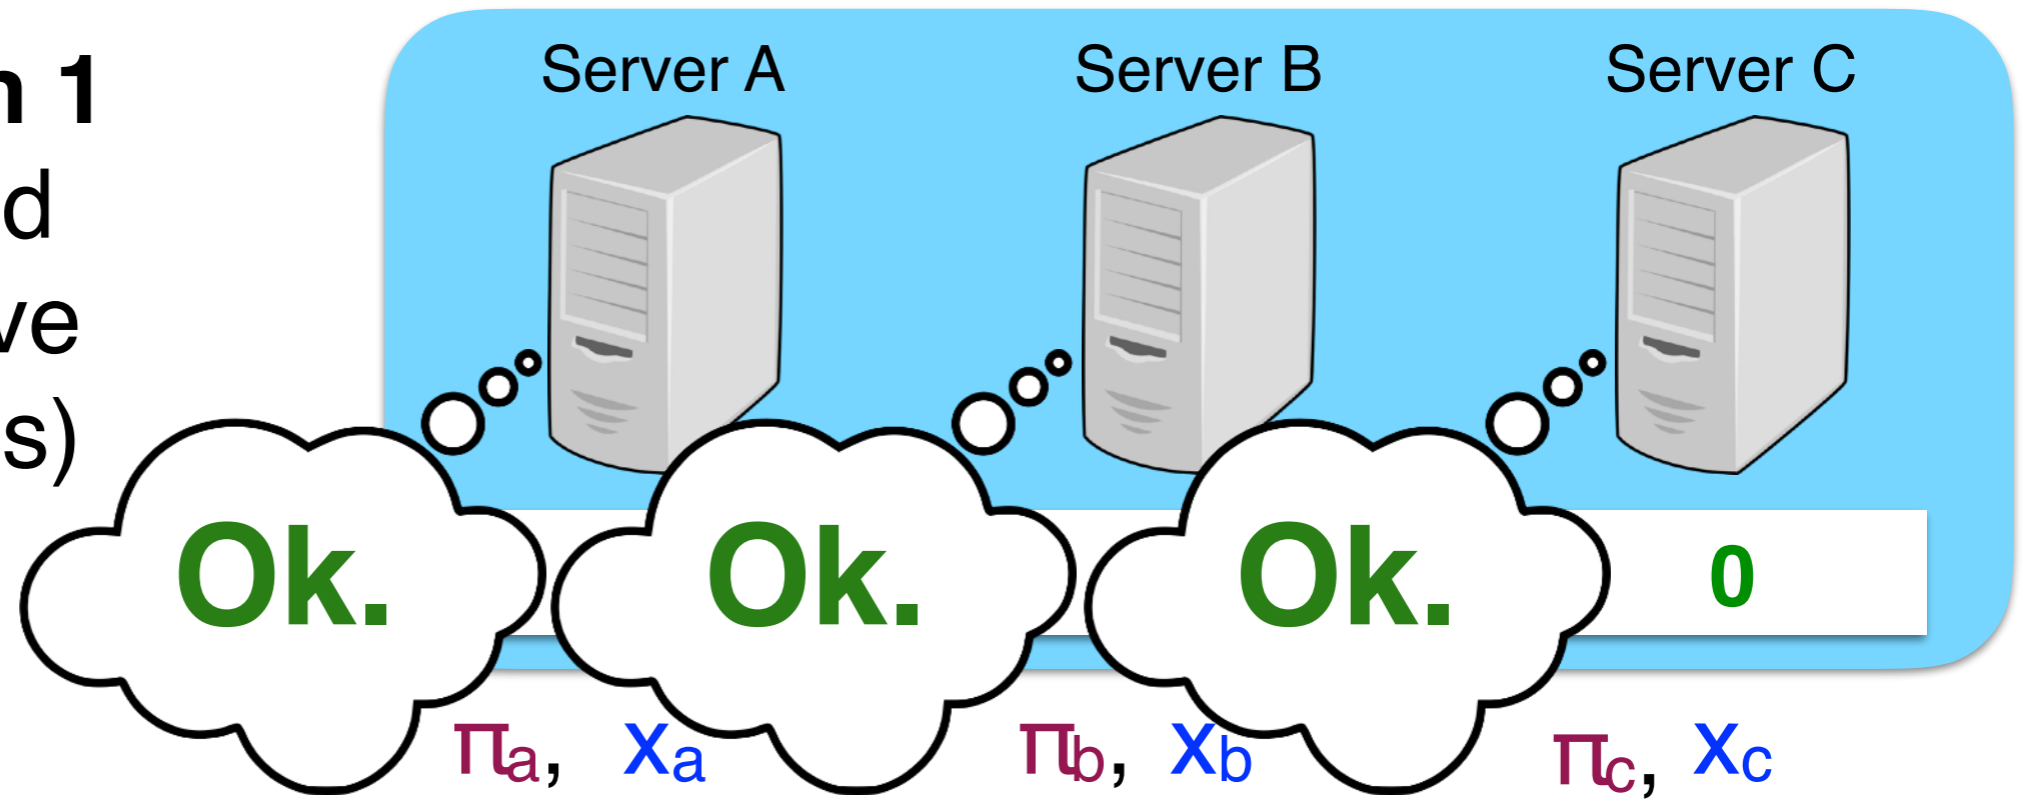


$$x = 1$$



# Contribution 1

Secret-shared  
non-interactive  
proofs (SNIPs)




$$x = 1$$



# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)




$x = 1$





# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)

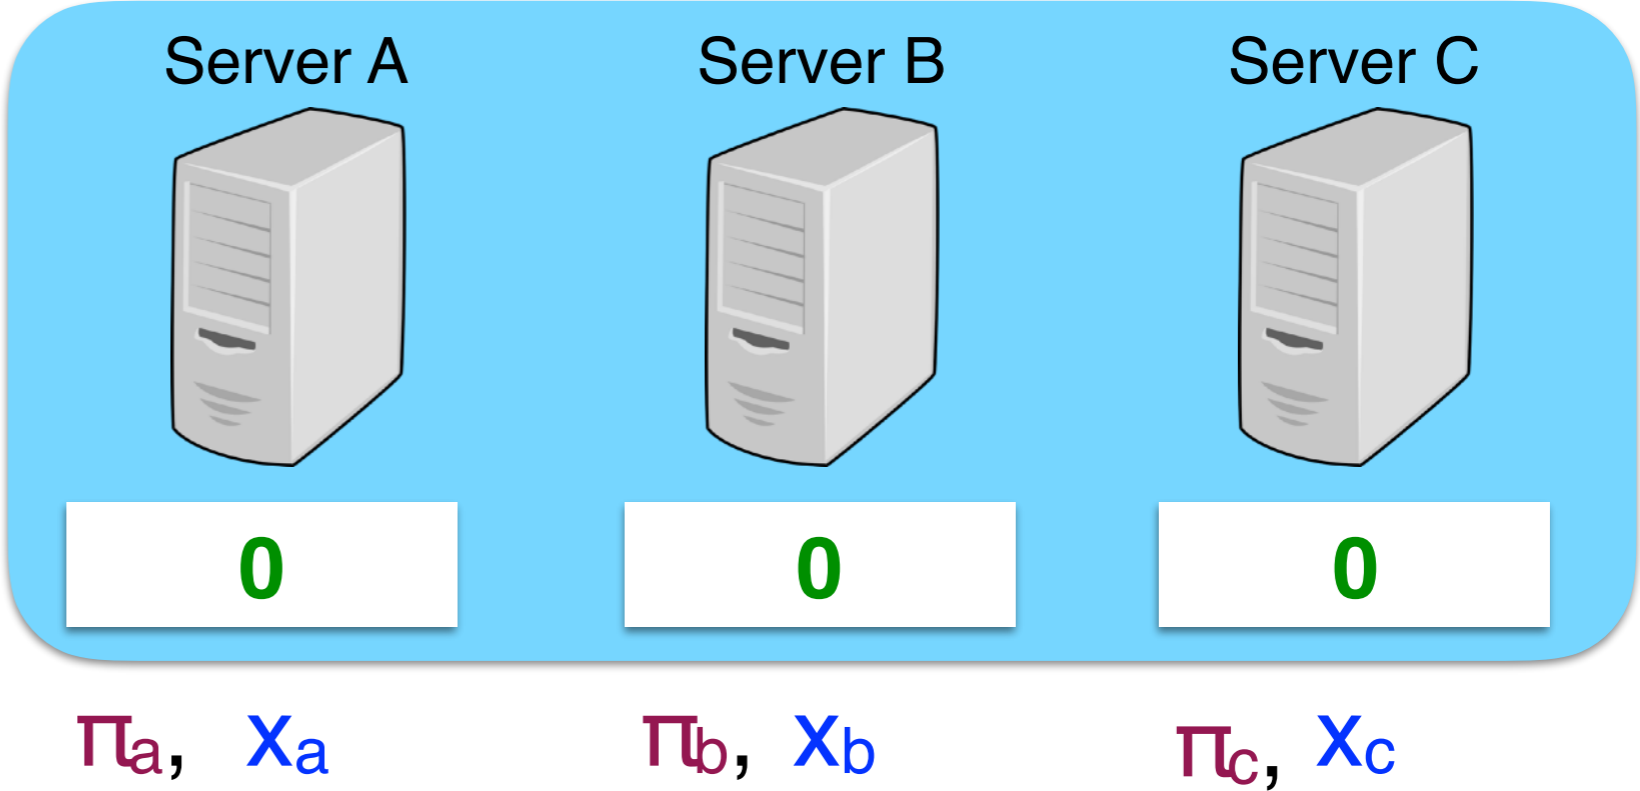


$$x = 1$$



# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)

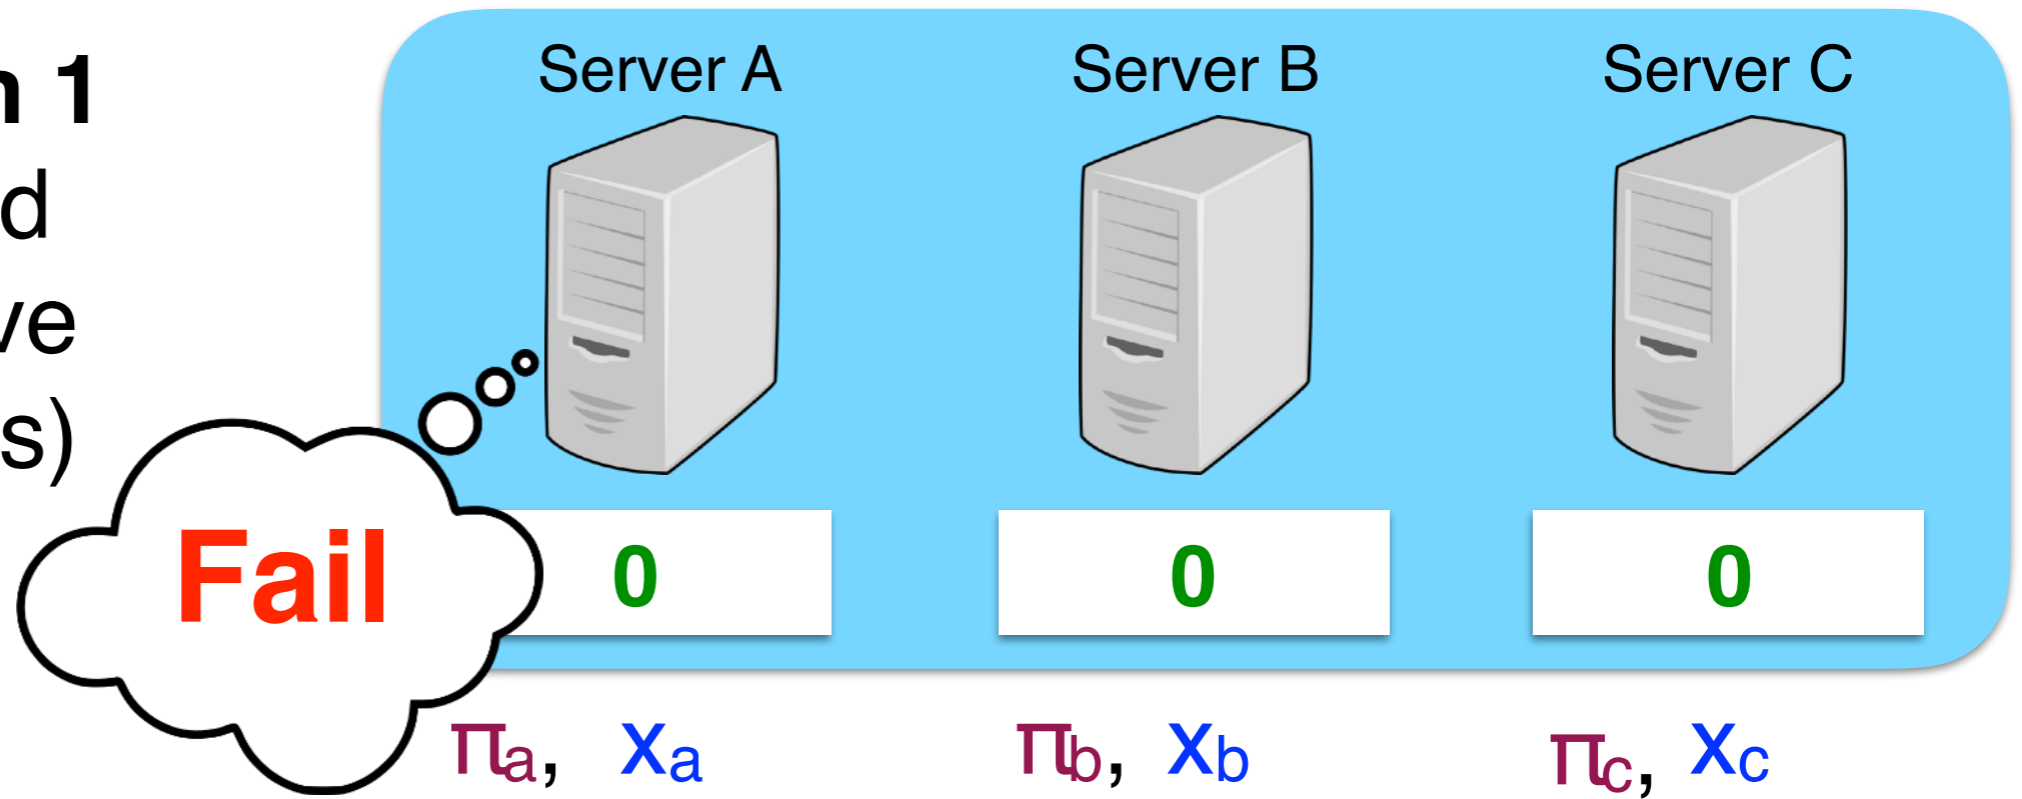


$x = 1$



# Contribution 1

Secret-shared  
non-interactive  
proofs (SNIPs)




$x = 1$



# Contribution 1

Secret-shared  
non-interactive  
proofs (SNIPs)

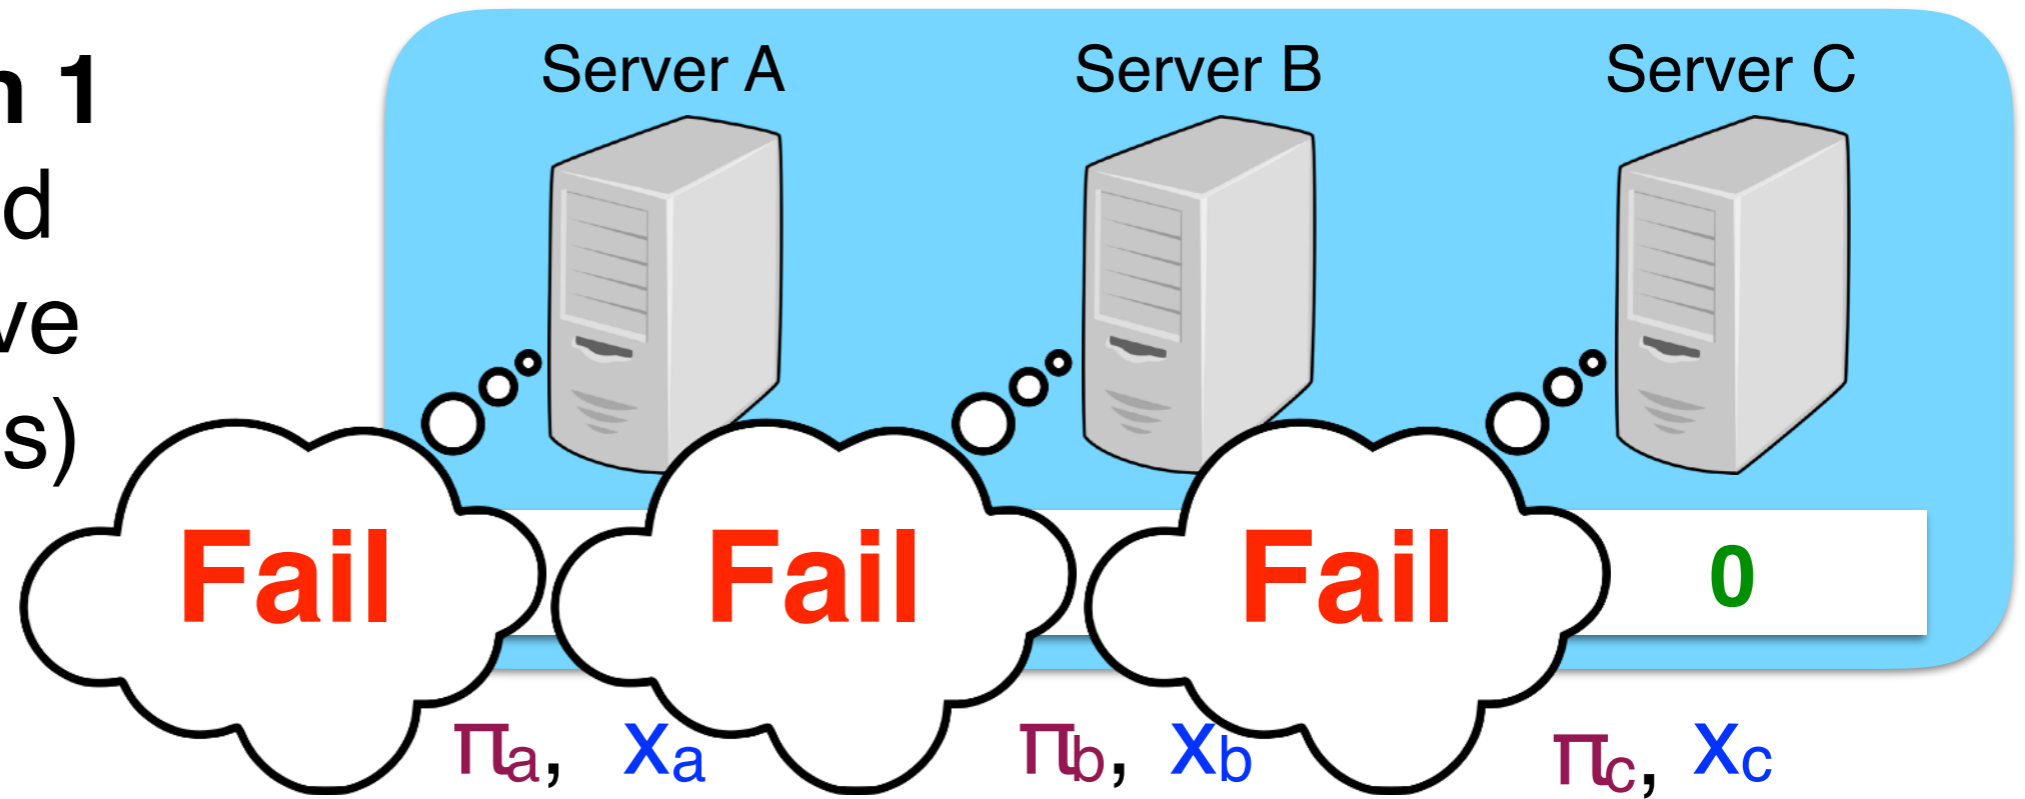


$$x = 1$$



# Contribution 1

Secret-shared  
non-interactive  
proofs (SNIPs)




$$x = 1$$



# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)




$x = 1$



# Contribution 1

## Secret-shared non-interactive proofs (SNIPs)




$$x = 1$$



- Prio servers *detect and reject malformed client submissions*
- In this example, each client can influence the aggregate statistic by  $\pm 1$ , at most

# We need a proof system





# We need a proof system

**Prover**

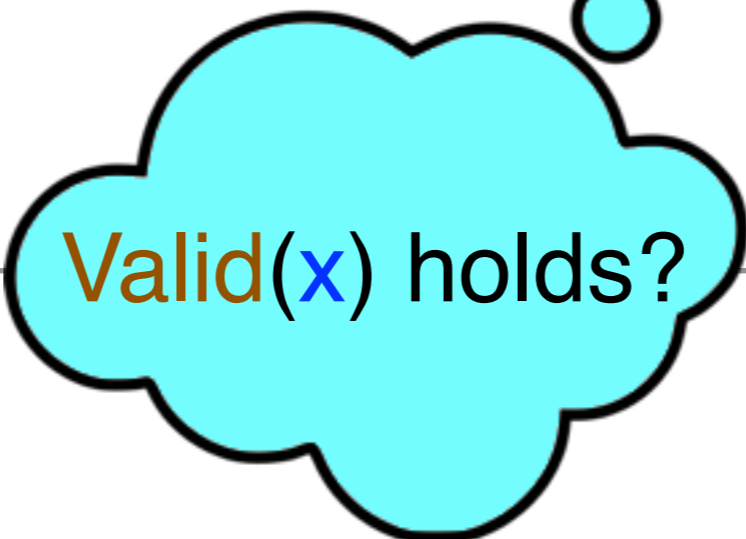



A “valid”  $x$

$\pi_a, \pi_b, \pi_c$




**Verifiers**




Valid( $x$ ) holds?

# We need a proof system



# We need a proof system



---

**Completeness.**

Honest prover convinces honest verifiers.

**Soundness.**

Dishonest prover rarely convinces honest verifiers.

**Zero knowledge.**

Any proper subset of the verifiers learns nothing about  $x$ , except that  $x$  is valid.

# Traditional techniques

- **Non-interactive proofs in ROM**

[FS86], [BFM88], [BDMP91], [CP92], [CS97], [M00], ...

- **zkSNARKs and KOE-based proofs**

[G10], [L12], [GGPR13], [BCGTV13], [PGHR13], ...

- **Multi-party computation**

[Y82], [GMW87], [BGW88], [CCD88], [CLOS02], [DPSZ12], [DKLPSS13], ...

**In our setting, SNIPs are a more efficient solution.**

# Traditional techniques

- **Non-interactive proofs in ROM**

[FS86], [BFM88], [BDMP91], [CP92], [CS97], [M00], ...

- **zkSNARKs and KOE-based proofs**


[G10], [L12], [GGPR13], [BCGTV13], [PGHR13], ...

- **Multi-party computation**


[Y82], [GMW87], [BGW88], [CCD88], [CLOS02], [DPSZ12], [DKLPSS13], ...

**In our setting, SNIPs are a more efficient solution.**

# How SNIPs work

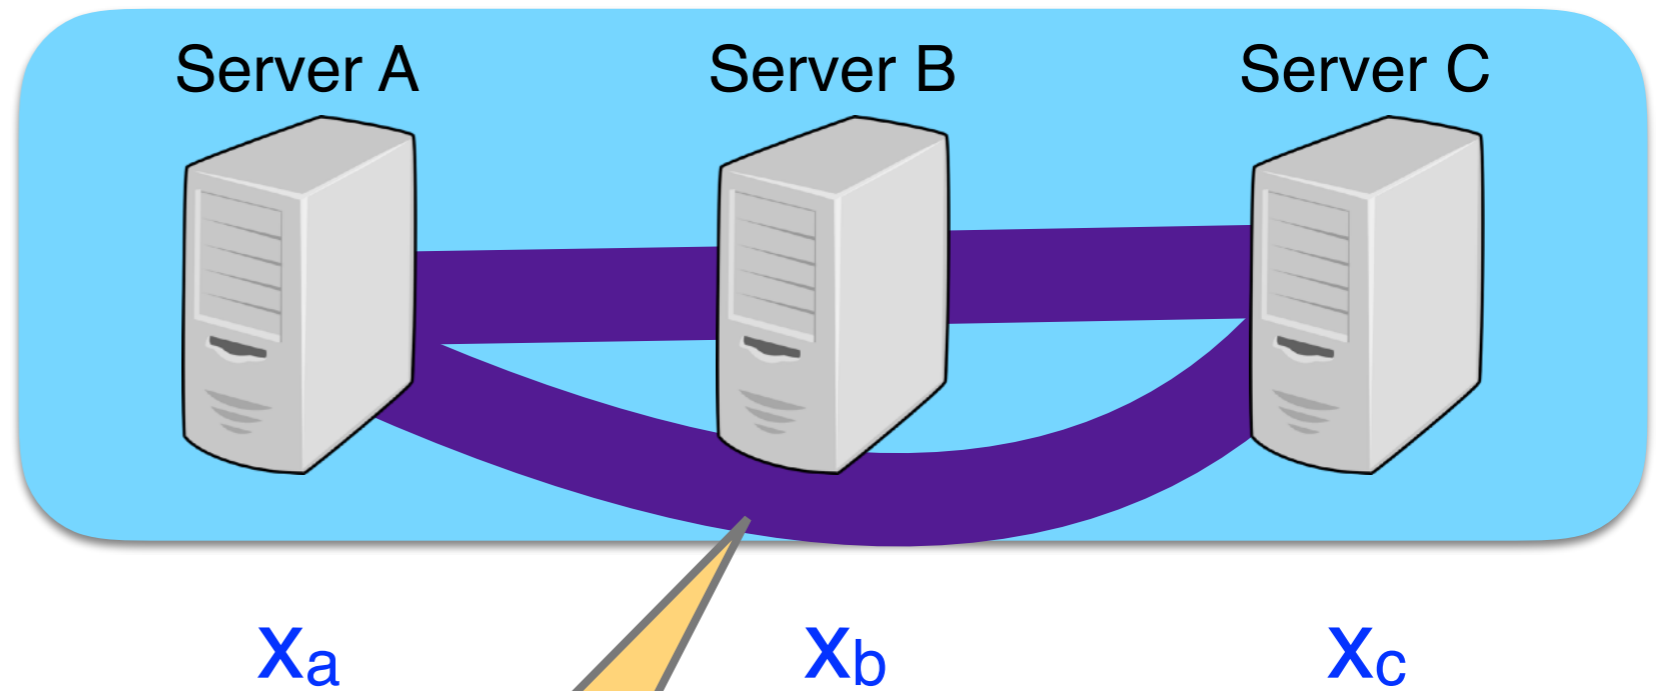


# How SNIPs work



Could run **secure multiparty computation** to check that  $\text{Valid}(x) = 1$ .  
[GMW87], [BGW88]


# How SNIPs work




Could run **secure multiparty computation** to check that  $\text{Valid}(x) = 1$ .  
[GMW87], [BGW88]



# How SNIPs work




# How SNIPs work



$X$




# How SNIPs work



$X_a$

$X_b$


$X_c$



**Idea:** Client generates the transcripts that servers *would* have observed in a multi-party computation

See also [IKOS07]


# How SNIPs work



$X_a$

$X_b$


$X_c$



**Idea:** Client generates the transcripts that servers *would* have observed in a multi-party computation

See also [IKOS07]

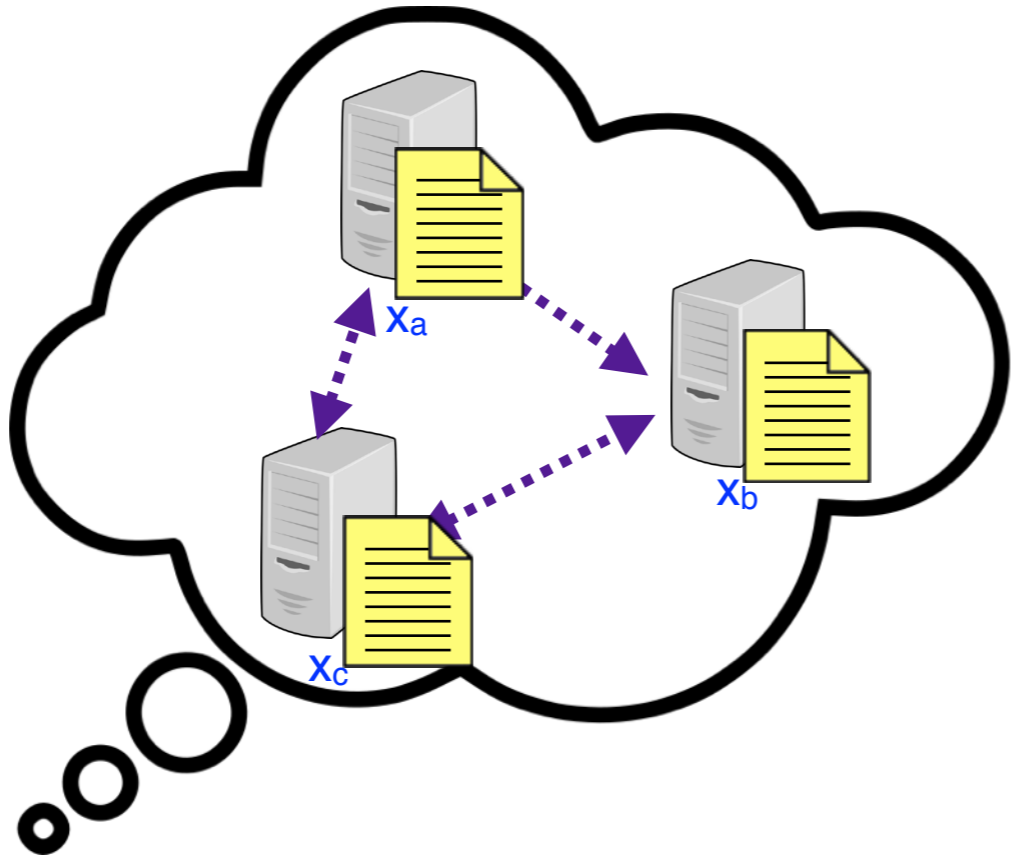
# How SNIPs work




Xa

Xb

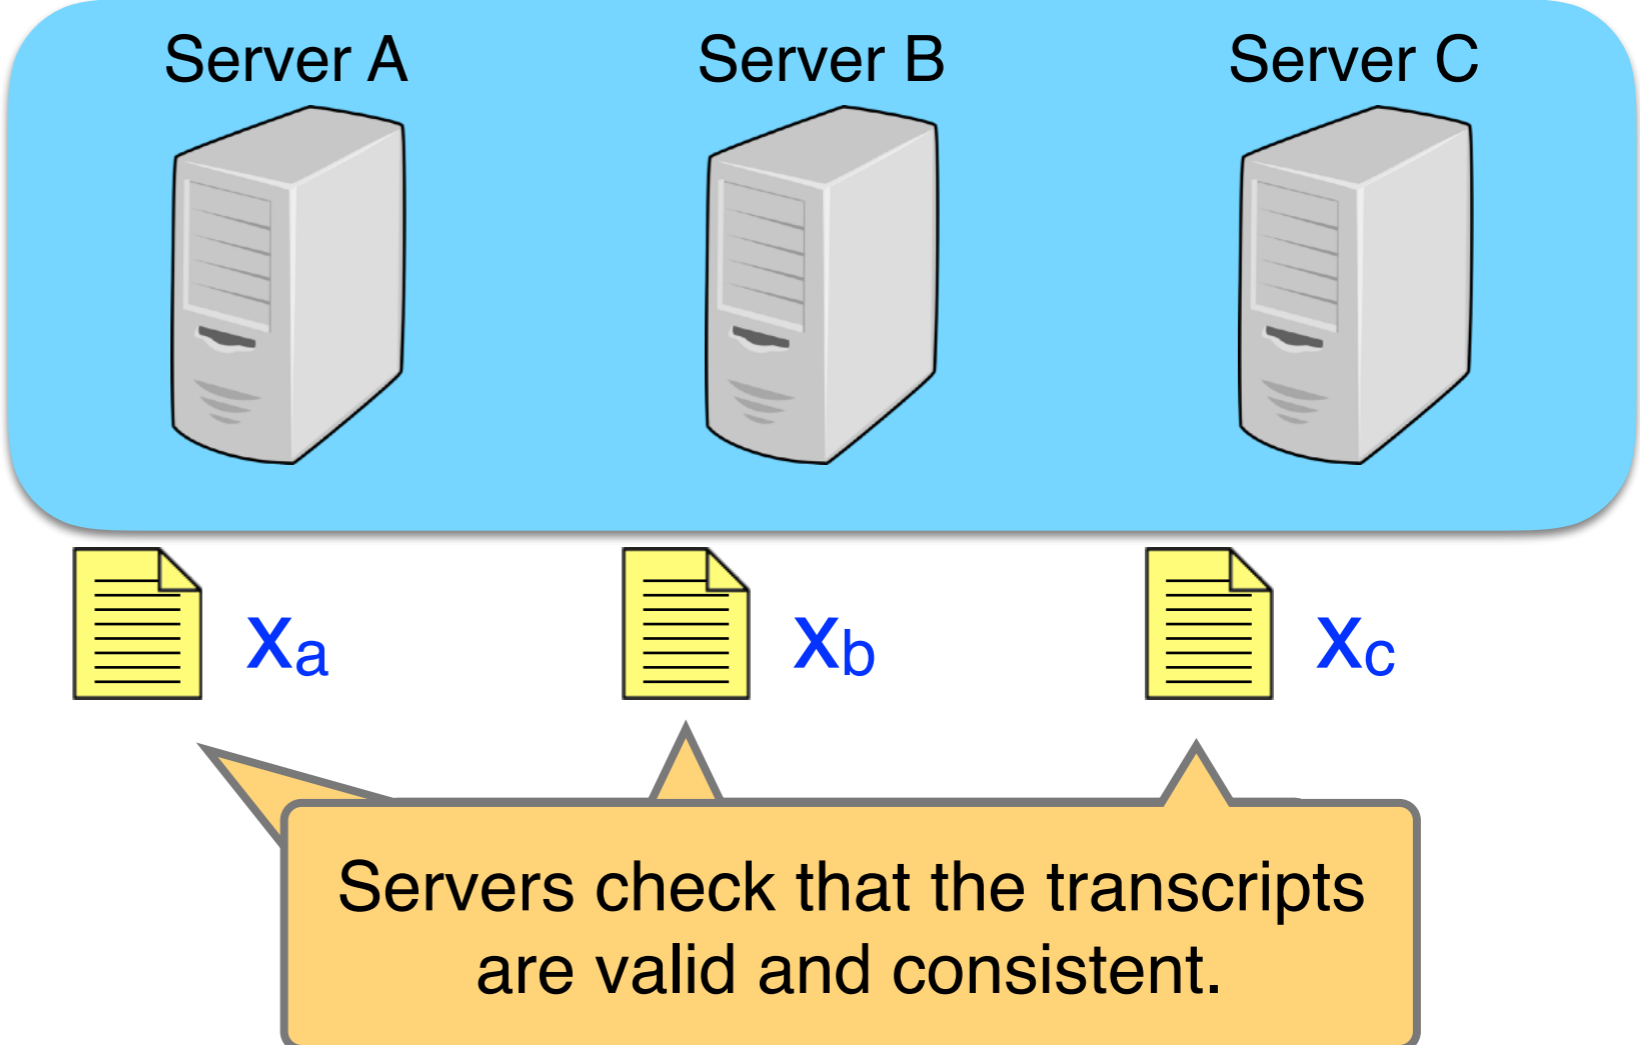
Xc



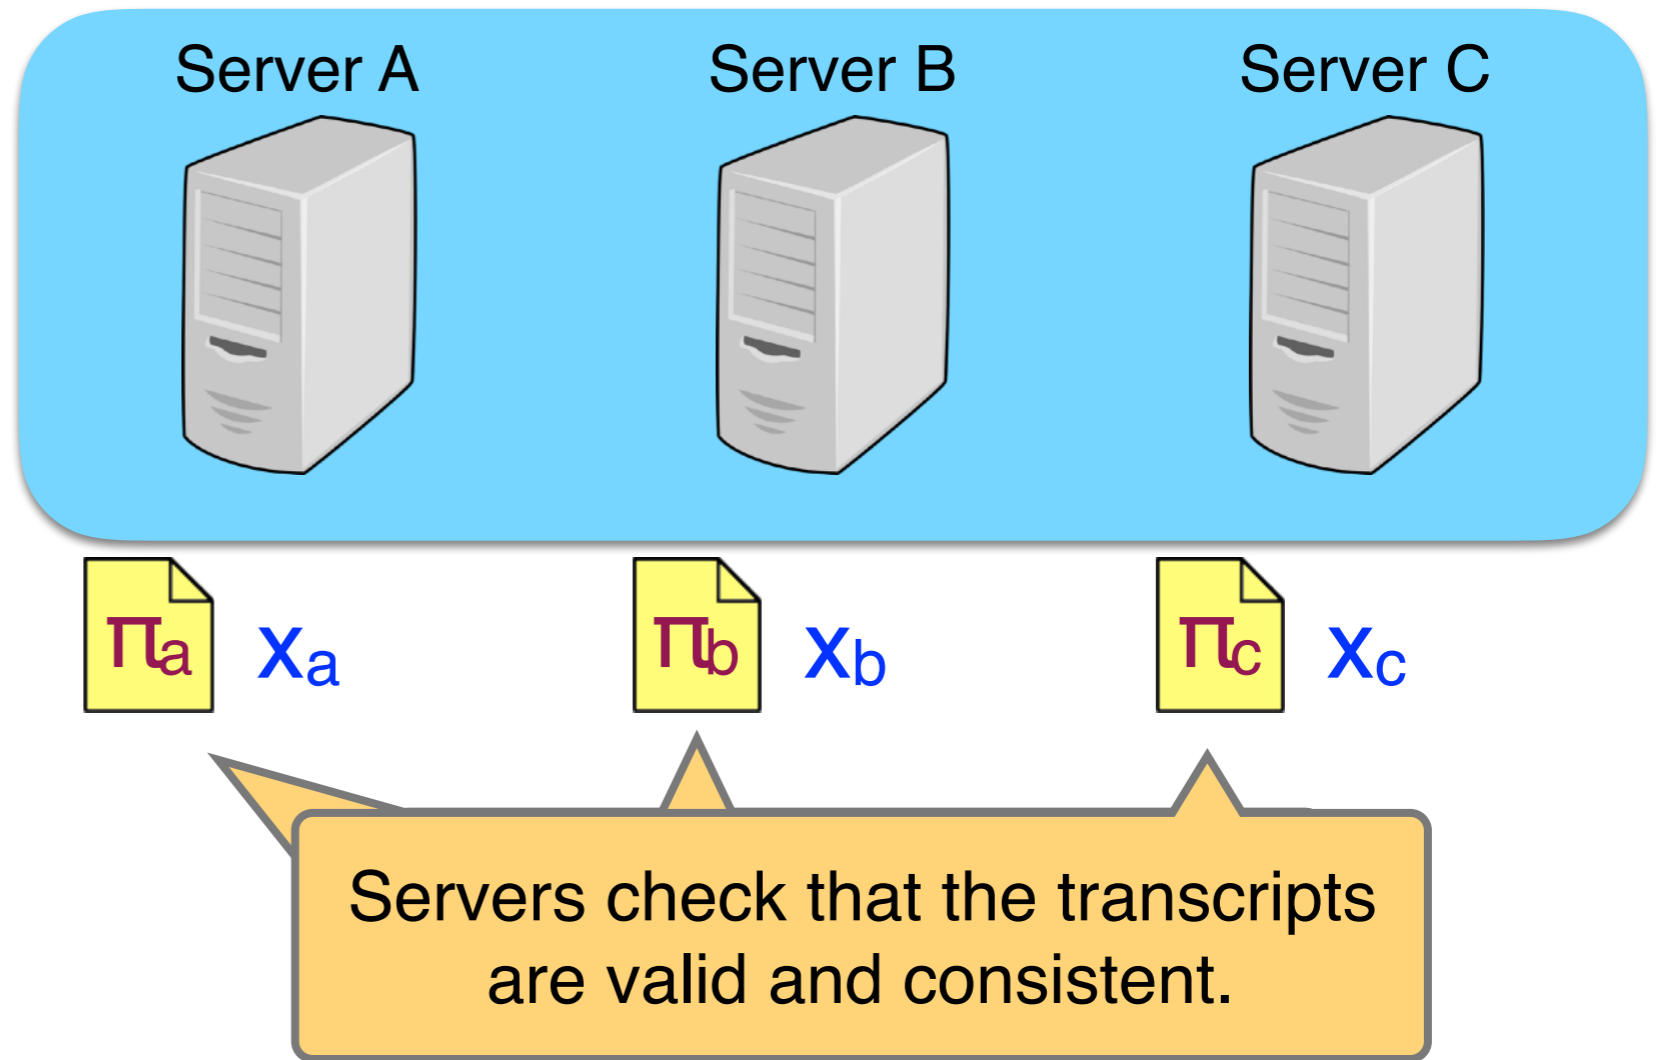
# How SNIPs work



# How SNIPs work



# How SNIPs work





X






# How SNIPs work



$X_a$



$X_b$



$X_c$

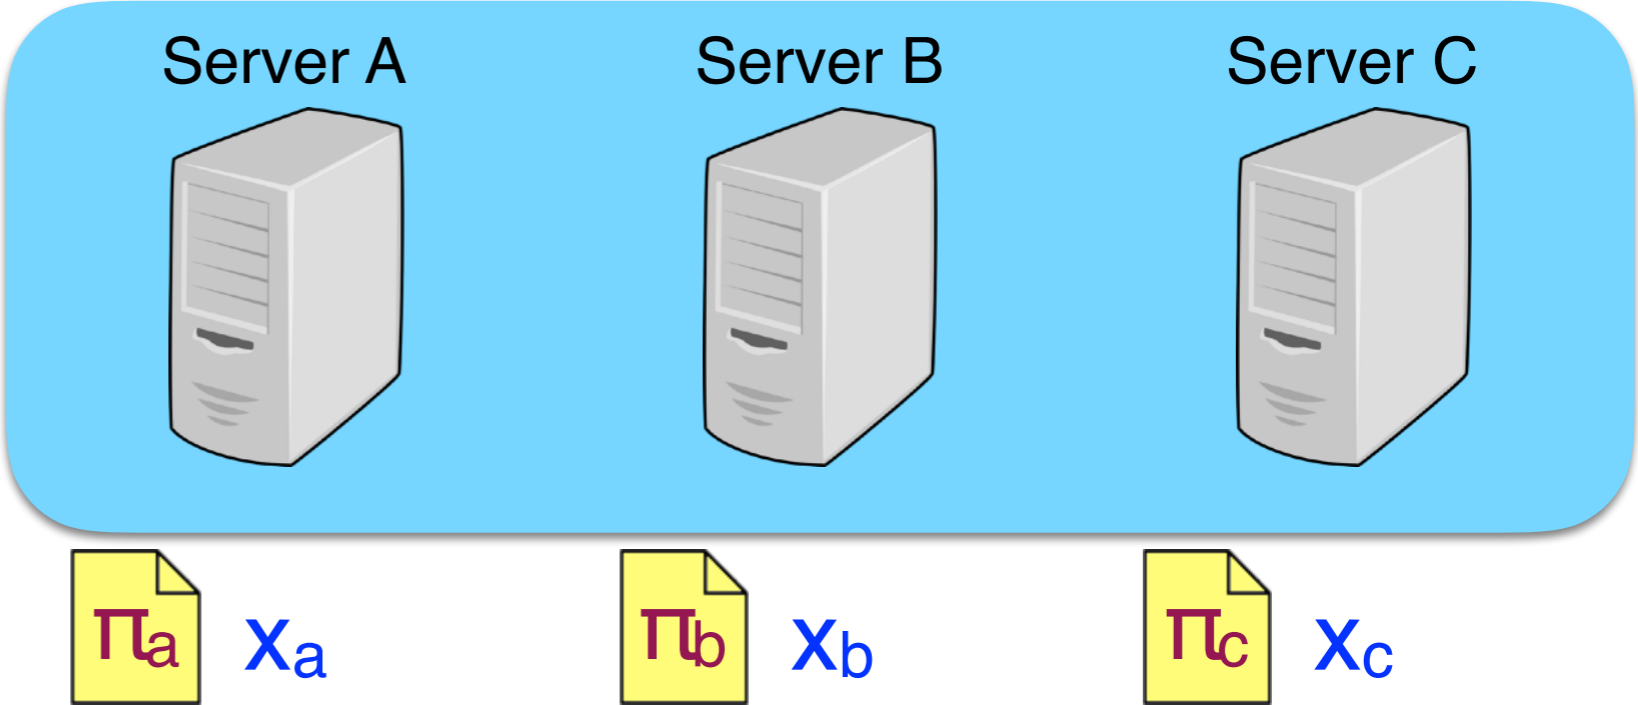
Servers check that the transcripts are valid and consistent.

Checking a transcript is **much easier** than generating it!


X



# How SNIPs work



# How SNIPs work



$X_a$



$D_a$



$X_b$



$D_b$

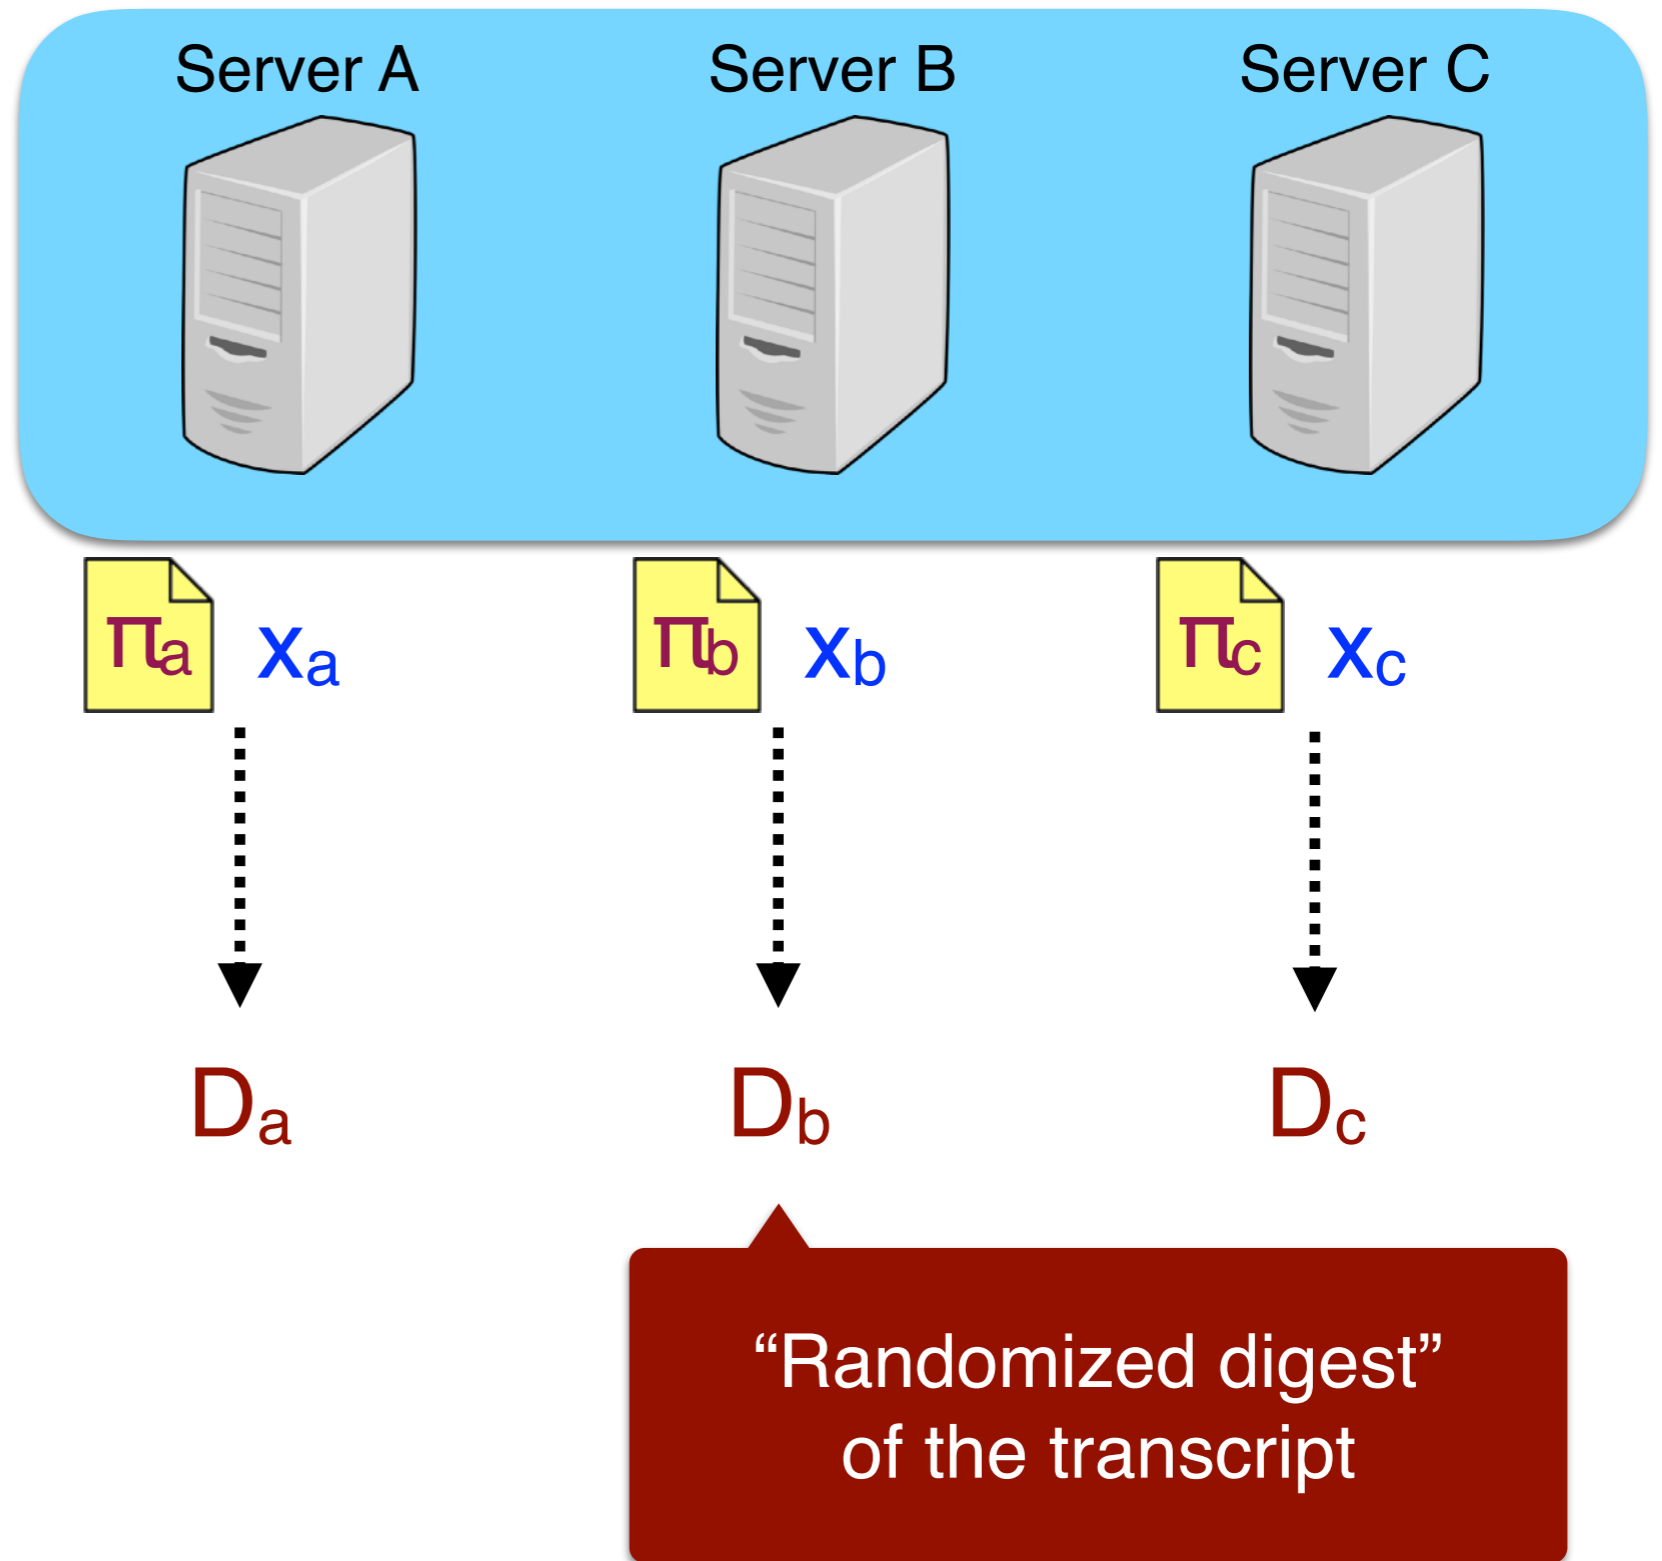


$X_c$




$D_c$

# How SNIPs work



# How SNIPs work



$X_a$



$D_a$



$X_b$



$D_b$




$X_c$




$D_c$


# How SNIPs work



# How SNIPs work



# How SNIPs work



- 
- If  $x$  is valid,  $D_a + D_b + D_c = 0$
  - If  $x$  is invalid,  $D_a + D_b + D_c \neq 0$  with high probability

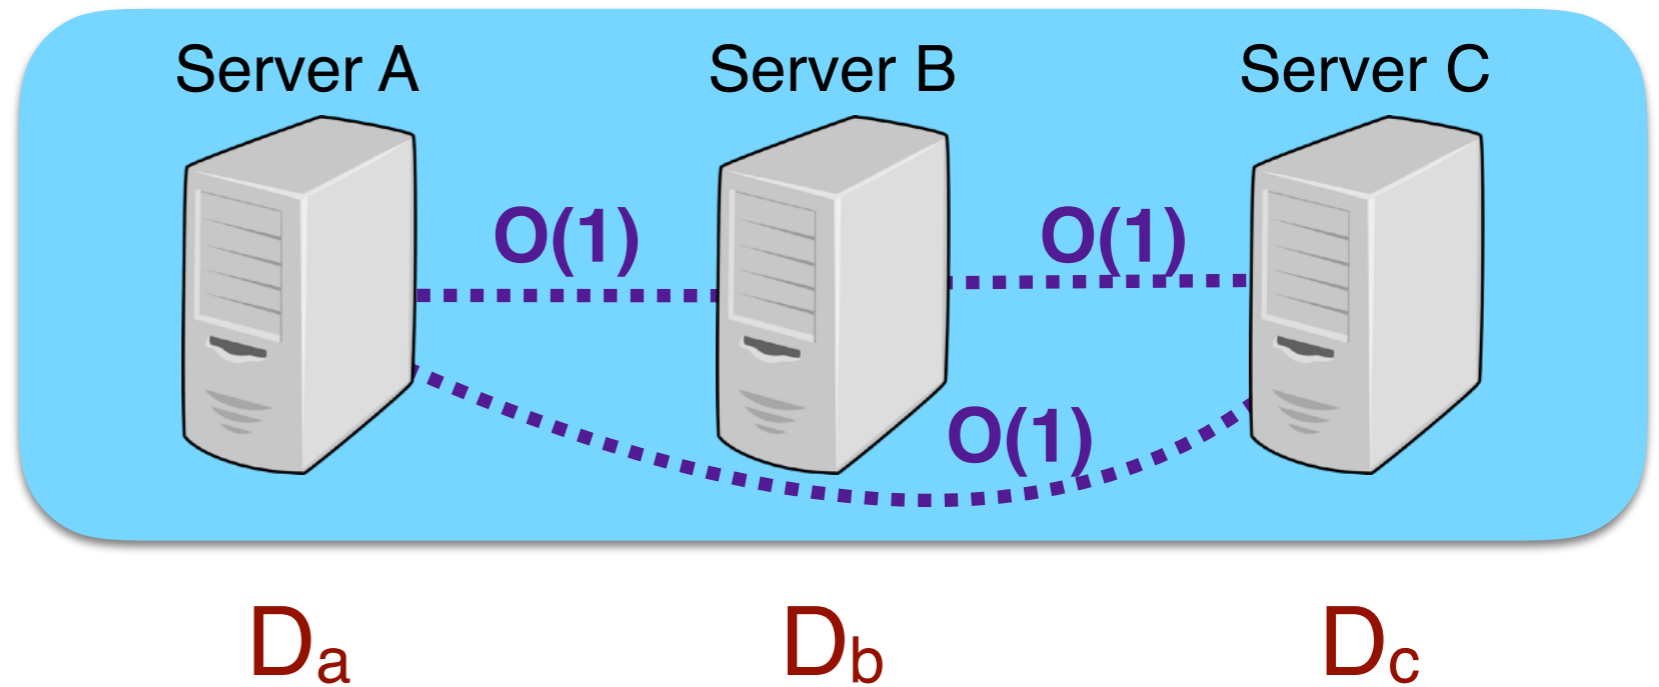
Servers run lightweight multi-party computation to check that

$$D_a + D_b + D_c = 0$$

If so, servers accept  $x$  is valid.



# How SNIPs work



- 
- If  $x$  is valid,  $D_a + D_b + D_c = 0$
  - If  $x$  is invalid,  $D_a + D_b + D_c \neq 0$  with high probability

Servers run lightweight multi-party computation to check that

$$D_a + D_b + D_c = 0$$

If so, servers accept  $x$  is valid.



M = # of multiplication gates in **Valid**(·) circuit

---

<b>Public-key ops.</b>		<b>Communication</b>		<b>Slow-down</b>
Client	Server	C-to-S	S-to-S	

M = # of multiplication gates in **Valid**(·) circuit

	Public-key ops.		Communication		Slow-down
	Client	Server	C-to-S	S-to-S	
Dishonest-maj. MPC <small>[CLOS02], [DPSZ12], ...</small>	0	$\Theta(M)$	0	$\Theta(M)$	5,000x at server

M = # of multiplication gates in **Valid**(·) circuit

	Public-key ops.		Communication		Slow-down
	Client	Server	C-to-S	S-to-S	
Dishonest-maj. MPC <small>[CLOS02], [DPSZ12], ...</small>	0	$\Theta(M)$	0	$\Theta(M)$	5,000x at server
Commits + NIZKs <small>[FS86], [CP92], [CS97], ...</small>	$\Theta(M)$	$\Theta(M)$	$\Theta(M)$	$\Theta(M)$	50x at server

	Public-key ops.		Communication		Slow-down
	Client	Server	C-to-S	S-to-S	
M = # of multiplication gates in <b>Valid</b> (·) circuit Dishonest-maj. MPC <small>[CLOS02], [DPSZ12], ...</small>	0	$\Theta(M)$	0	$\Theta(M)$	5,000x at server
Commits + NIZKs <small>[FS86], [CP92], [CS97], ...</small>	$\Theta(M)$	$\Theta(M)$	$\Theta(M)$	$\Theta(M)$	50x at server
Commits + SNARKs <small>[GGPR13], [BCGTV13], ...</small>	$\Theta(M)$	$O(1)$	$O(1)$	$O(1)$	500x at client

	Public-key ops.		Communication		Slow-down
	Client	Server	C-to-S	S-to-S	
M = # of multiplication gates in <b>Valid</b> (·) circuit Dishonest-maj. MPC <small>[CLOS02], [DPSZ12], ...</small>	0	$\Theta(M)$	0	$\Theta(M)$	5,000x at server
Commits + NIZKs <small>[FS86], [CP92], [CS97], ...</small>	$\Theta(M)$	$\Theta(M)$	$\Theta(M)$	$\Theta(M)$	50x at server
Commits + SNARKs <small>[GGPR13], [BCGTV13], ...</small>	$\Theta(M)$	$O(1)$	$O(1)$	$O(1)$	500x at client
<b>This work: SNIPs</b>	0	0	$\Theta(M)$	$O(1)$	1x

	Public-key ops.		Communication		Slow-down
	Client	Server	C-to-S	S-to-S	
M = # of multiplication gates in <b>Valid</b> (·) circuit					
Dishonest-maj. MPC <small>[CLOS02], [DPSZ12], ...</small>	0	$\Theta(M)$	0	$\Theta(M)$	5,000x at server
Commits + NIZKs <small>[FS86], [CP92], [CS97], ...</small>	$\Theta(M)$	$\Theta(M)$	$\Theta(M)$	$\Theta(M)$	50x at server
Commits + SNARKs <small>[GGPR13], [BCGTV13], ...</small>	$\Theta(M)$	$O(1)$	$O(1)$	$O(1)$	500x at client
<b>This work: SNIPs</b>	0	0	$\Theta(M)$	$O(1)$	1x

For specific Valid() circuits, it is possible to eliminate this cost [BGI16]



	Public-key ops.		Communication		Slow-down
	Client	Server	C-to-S	S-to-S	
M = # of multiplication gates in <b>Valid</b> (·) circuit					
Dishonest-maj. MPC <small>[CLOS02], [DPSZ12], ...</small>	0	$\Theta(M)$	0	$\Theta(M)$	5,000x at server
Commits + NIZKs <small>[FS86], [CP92], [CS97], ...</small>	$\Theta(M)$	$\Theta(M)$	$\Theta(M)$	$\Theta(M)$	50x at server
Commits + SNARKs <small>[GGPR13], [BCGTV13], ...</small>	$\Theta(M)$	$O(1)$	$O(1)$	$O(1)$	500x at client
<b>This work: SNIPs</b>	0	0	$\Theta(M)$	$O(1)$	1x

# From sums to more complex aggregates

If you can compute private sums, you can compute many other interesting aggregates

[PrivStats11], [KDK11], [DFKZ13], [PrivEx14], [MDD16], ...

- Average
- Variance
- Standard deviation
- Most popular (approx)
- “Heavy hitters” (approx)
- Min and max (approx)
- Quality of arbitrary regression model ( $R^2$ )
- Least-squares regression
- Stochastic gradient descent [Bonawitz et al. 2016]

# Outline

- Background: The private aggregation problem
- A straw-man solution for private sums
- **Providing robustness with SNIPs**
- Evaluation
- Discussion: Real-world considerations

# Outline

- Background: The private aggregation problem
- A straw-man solution for private sums
- Providing robustness with SNIPs
- **Evaluation**
- Discussion: Real-world considerations

# Evaluation

- Implemented Prio in Go  
(see optimizations described in paper)
- Five-server cluster in EC2
- System collects the sum of “N” 0/1 values

## Four variants

1. No privacy
2. No robustness (“straw man”)
3. Prio (privacy + robustness)
4. NIZK (privacy + robustness)



# Evaluation

- Implemented Prio in Go  
(see optimizations described in paper)
- Five-server cluster in EC2
- System collects the sum of “N” 0/1 values



E.g., for privately measuring telemetry data.

## Four variants

1. No privacy
2. No robustness (“straw man”)
3. Prio (privacy + robustness)
4. NIZK (privacy + robustness)

# Evaluation

- Implemented Prio in Go  
(see optimizations described in paper)
- Five-server cluster in EC2
- System collects the sum of “N” 0/1 values

## Four variants

1. No privacy
2. No robustness (“straw man”)
3. Prio (privacy + robustness)
4. NIZK (privacy + robustness)



# Evaluation

- Implemented Prio in Go  
(see optimizations described in paper)
- Five-server cluster in EC2
- System collects the sum of “N” 0/1 values

## Four variants

1. No privacy
2. No robustness (“straw man”)
3. Prio (privacy + robustness)
4. NIZK (privacy + robustness)





# Evaluation


- Implemented Prio in Go  
(see optimizations described in paper)
- Five-server cluster in EC2
- System collects the sum of “N” 0/1 values

## Four variants


1. No privacy
2. No robustness (“straw man”)
3. Prio (privacy + robustness)
4. NIZK (privacy + robustness)




Five-server cluster in five  
Amazon data centers



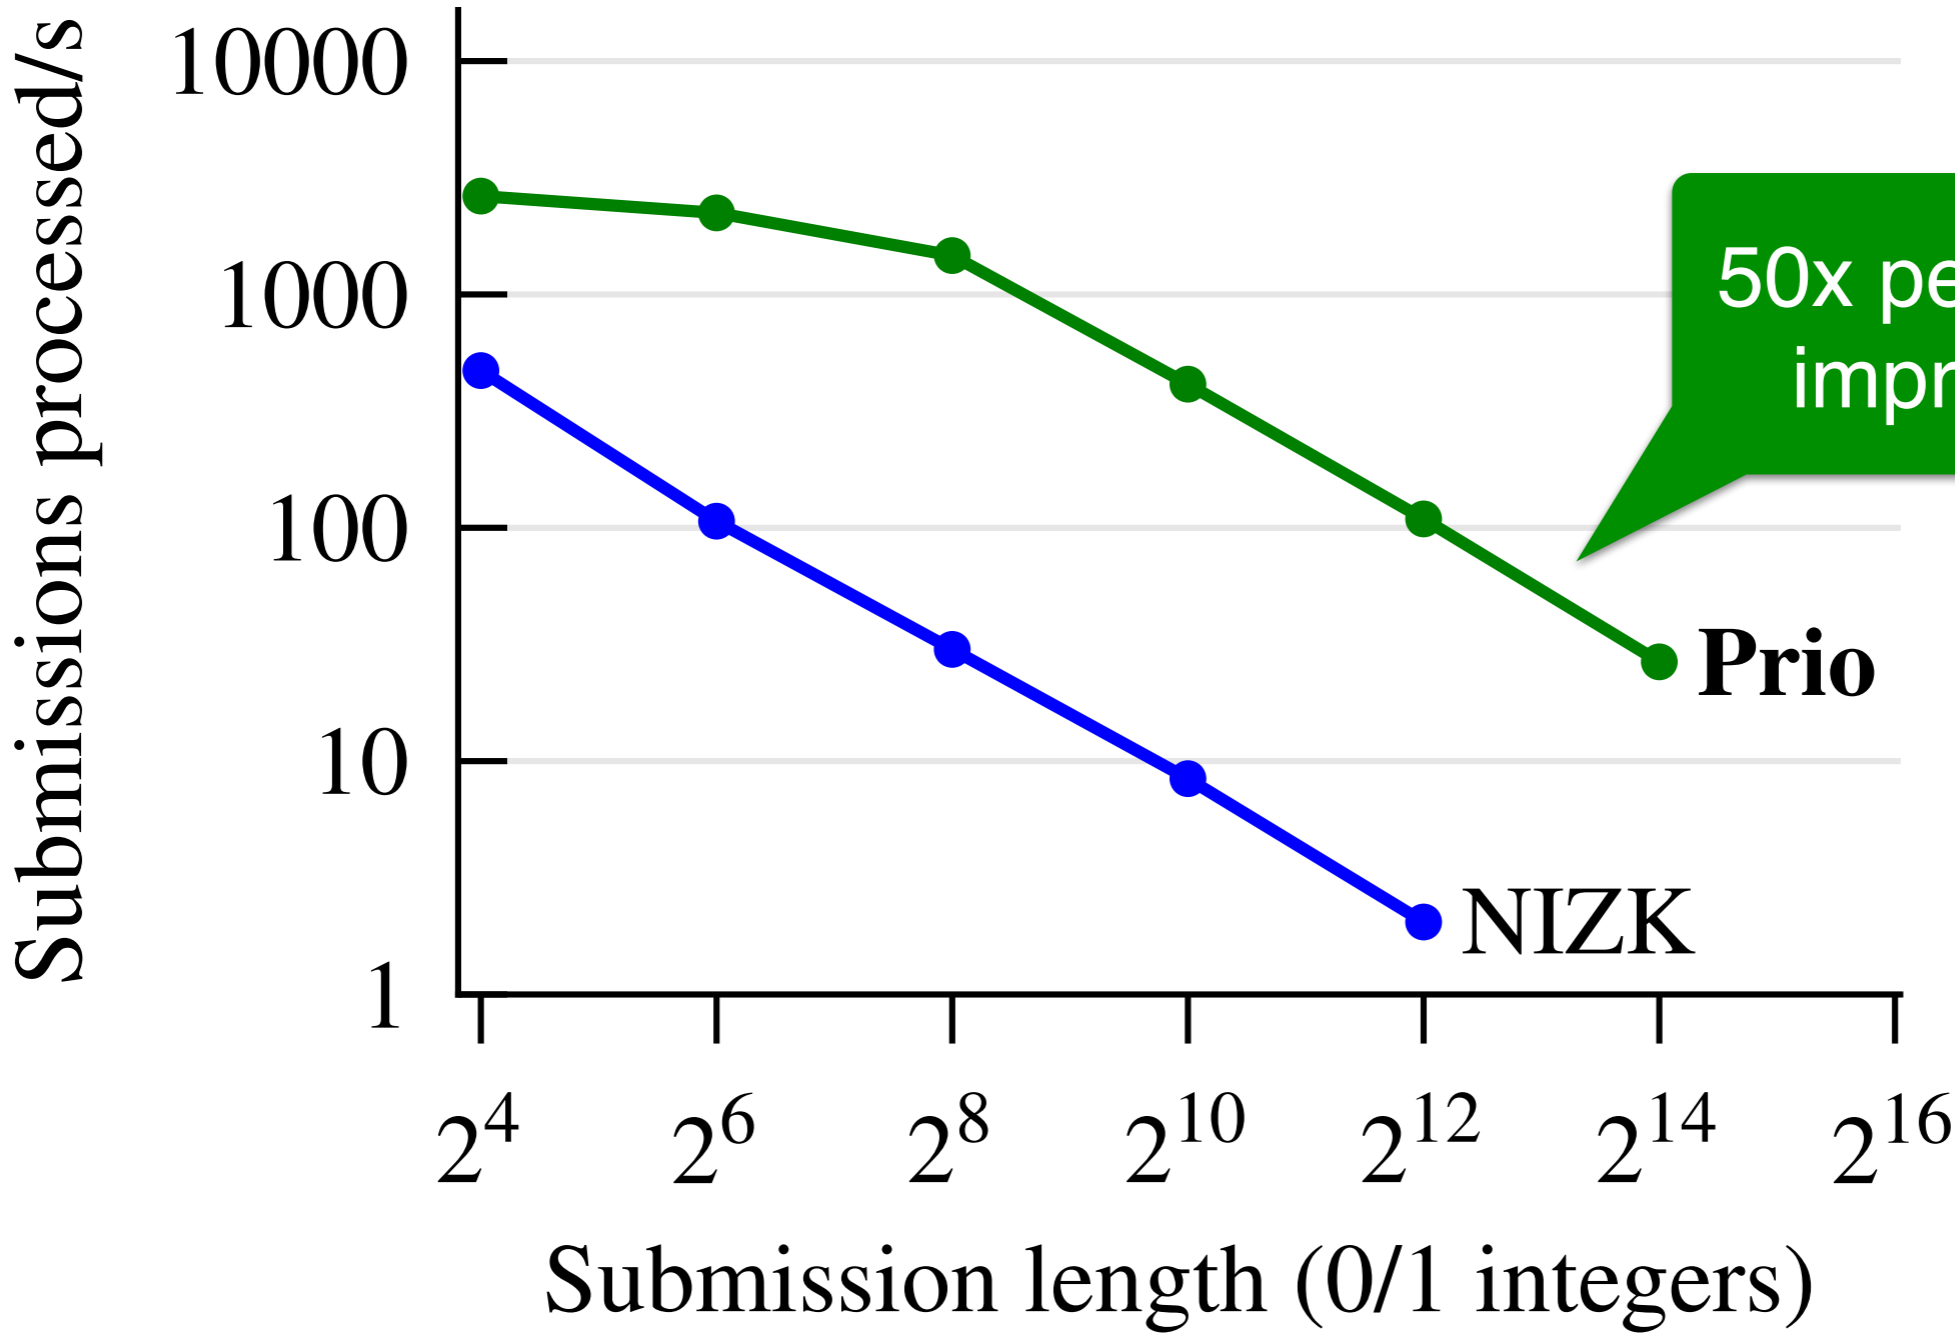
Five-server cluster in five  
Amazon data centers



Five-server cluster in five Amazon data centers



Five-server cluster in five Amazon data centers




50x performance improvement

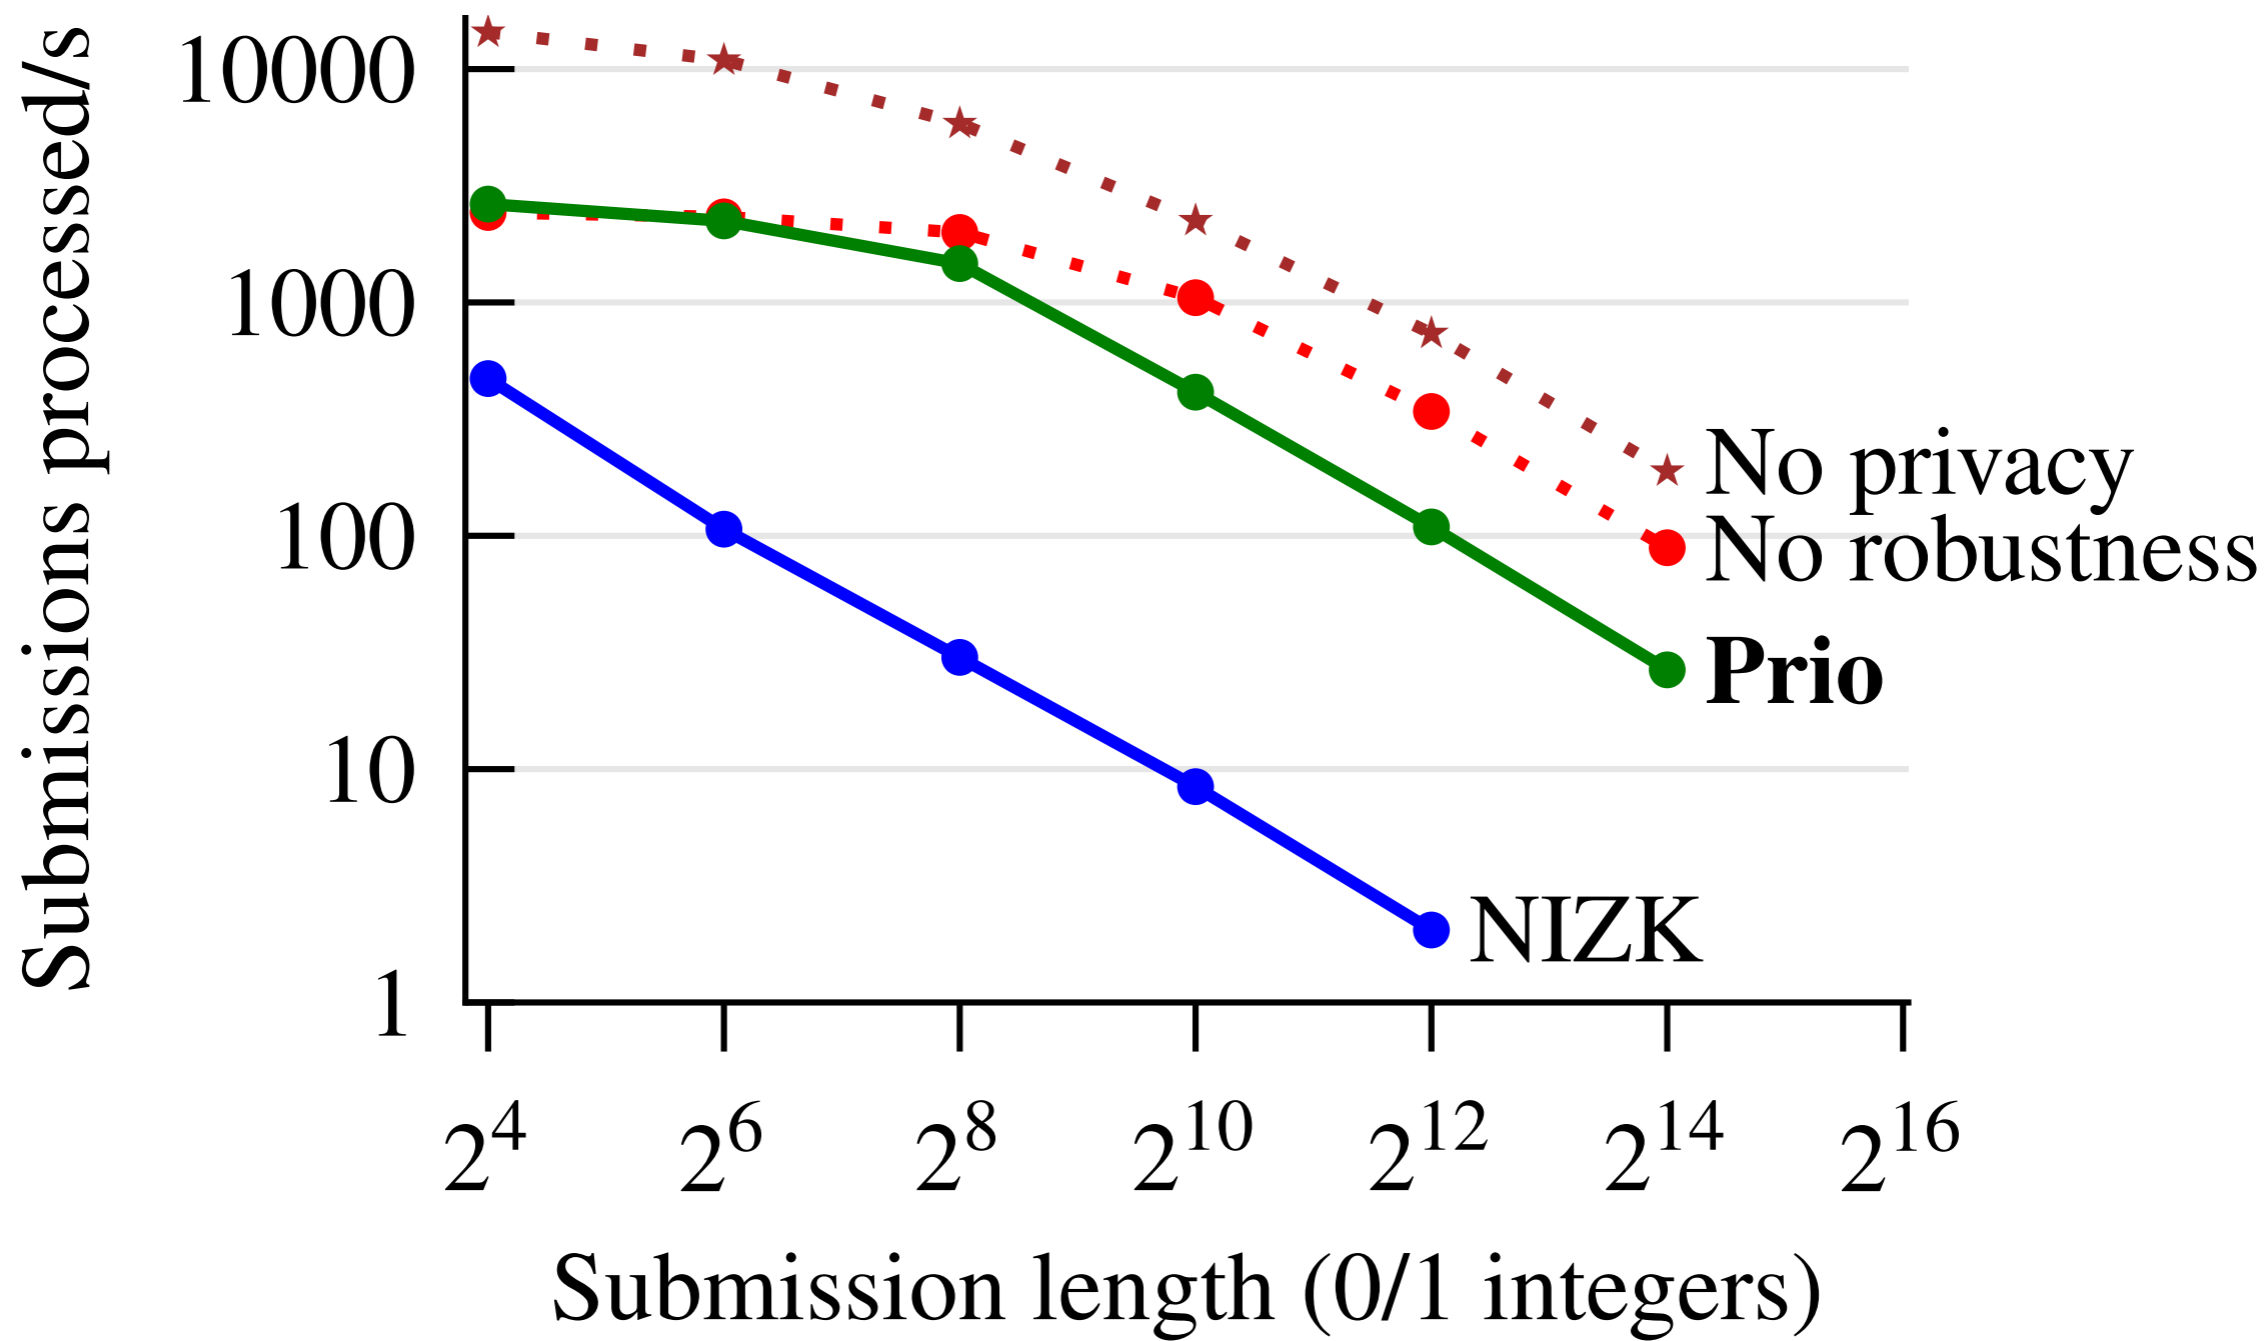
**Prio**

**NIZK**

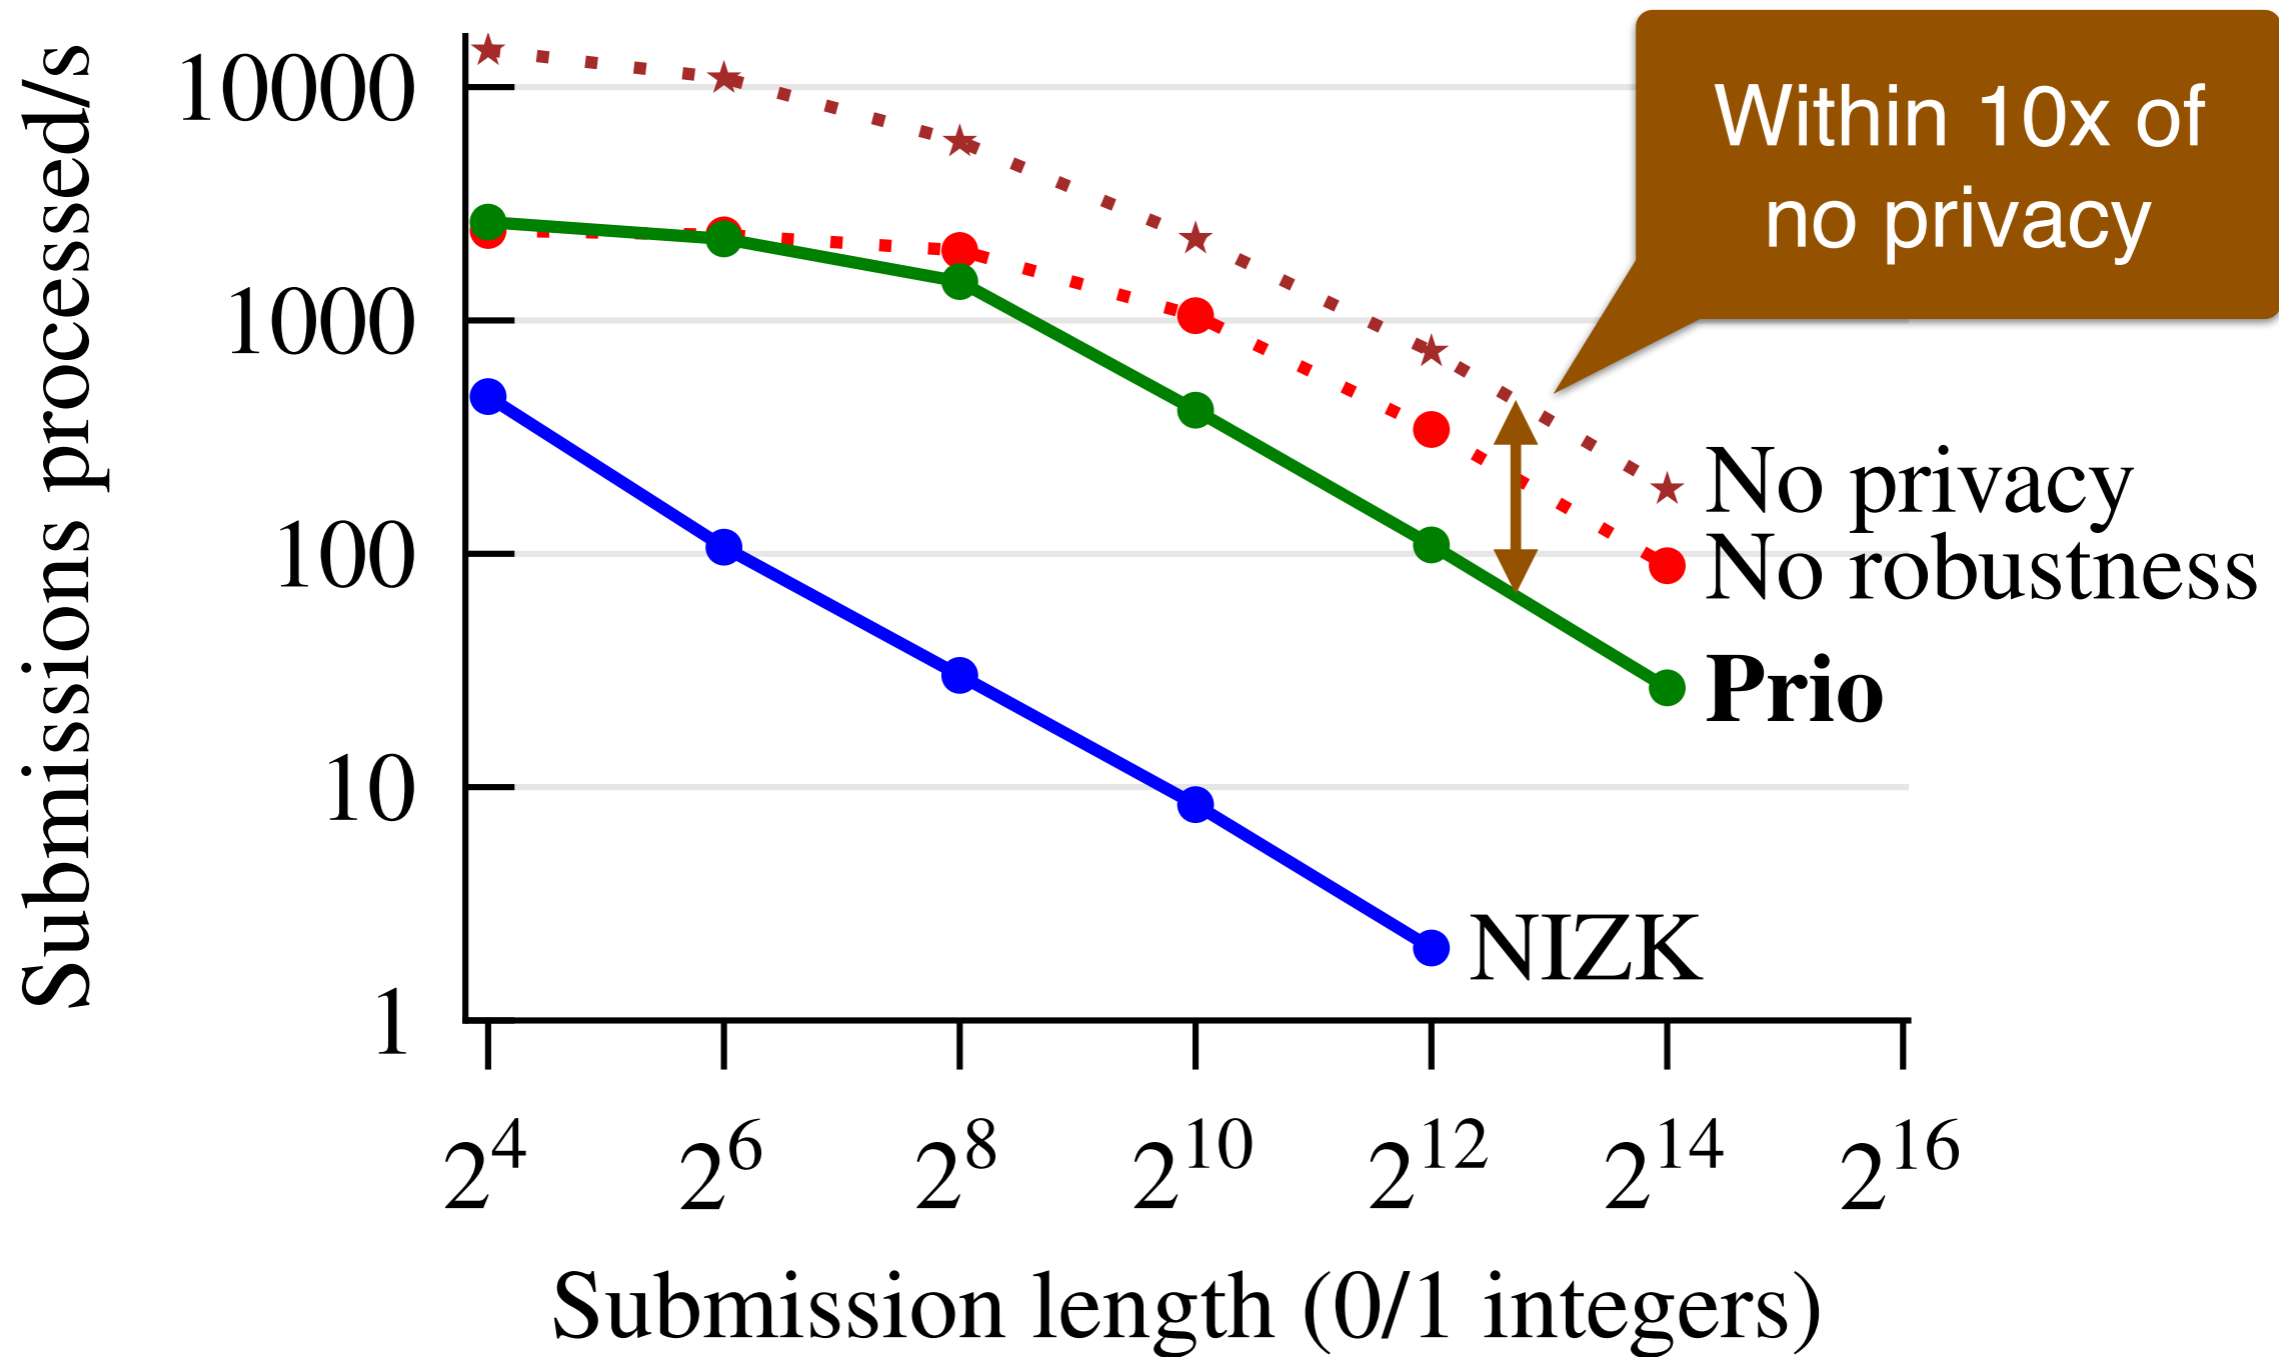
Five-server cluster in five Amazon data centers



Five-server cluster in five  
Amazon data centers



Five-server cluster in five Amazon data centers





# Outline

- Background: The private aggregation problem
- A straw-man solution for private sums
- Providing robustness with SNIPs
- **Evaluation**
- Discussion: Real-world considerations

# Outline

- Background: The private aggregation problem
- A straw-man solution for private sums
- Providing robustness with SNIPs
- Evaluation
- **Discussion: Real-world considerations**

# Real-world considerations: Technical

Every company we spoke with said:

- Server resources are cheap, client resources are not
- Client bandwidth usage is the important quantity to minimize
- Need some defense against faulty/disruptive clients
- Privately collecting popular URLs is the interesting application
  - Existing solutions are good, but not great

Areas of vehement disagreement between companies:

- Non-colluding servers—realistic?
- Does SGX obviate the need for these cryptographic protocols?

# Real-world considerations: Technical

Every company we spoke with said:

- Server resources are cheap, client resources are not
- Client bandwidth usage is the important quantity to minimize
- Need some defense against faulty/disruptive clients
- Privately collecting popular URLs is the interesting application
  - Existing solutions are good, but not great

Areas of vehement disagreement between companies:

- Non-colluding servers—realistic?
- Does SGX obviate the need for these cryptographic protocols?

# Real-world considerations: Non-technical

# Real-world considerations: Non-technical

*“If we’re already collecting \_\_\_\_\_ without privacy protections, why bother adding privacy protections?”*

- Pitch: Collect new statistics that you couldn’t collect before

*“We don’t yet know what aggregates we want to collect.”*

- Pitch: It’s possible to retain some flexibility  
(e.g., can later break out statistics by geographic area)

# Real-world considerations: Non-technical

*“If we’re already collecting \_\_\_\_\_ without privacy protections, why bother adding privacy protections?”*

- Pitch: Collect new statistics that you couldn’t collect before

*“We don’t yet know what aggregates we want to collect.”*

- Pitch: It’s possible to retain some flexibility  
(e.g., can later break out statistics by geographic area)

# Real-world considerations: Non-technical

*“If we’re already collecting \_\_\_\_\_ without privacy protections, why bother adding privacy protections?”*

- Pitch: Collect new statistics that you couldn’t collect before

*“We don’t yet know what aggregates we want to collect.”*

- Pitch: It’s possible to retain some flexibility  
(e.g., can later break out statistics by geographic area)

*“Please sign this stack of non-disclosure agreements...”*



# Real-world considerations: Non-technical

*“If we’re already collecting \_\_\_\_\_ without privacy protections, why bother adding privacy protections?”*

- Pitch: Collect new statistics that you couldn’t collect before

*“We don’t yet know what aggregates we want to collect.”*

- Pitch: It’s possible to retain some flexibility  
(e.g., can later break out statistics by geographic area)

*“Please sign this stack of non-disclosure agreements...  
and before we show you the non-disclosure agreements...”*

# Real-world considerations: Non-technical

*“If we’re already collecting \_\_\_\_\_ without privacy protections, why bother adding privacy protections?”*

- Pitch: Collect new statistics that you couldn’t collect before

*“We don’t yet know what aggregates we want to collect.”*

- Pitch: It’s possible to retain some flexibility  
(e.g., can later break out statistics by geographic area)

*“Please sign this stack of non-disclosure agreements...  
and before we show you the non-disclosure agreements...  
you have to sign a preliminary non-disclosure agreement.”*

# Real-world considerations: Non-technical

*“If we’re already collecting \_\_\_\_\_ without privacy protections, why bother adding privacy protections?”*

- Pitch: Collect new statistics that you couldn’t collect before

*“We don’t yet know what aggregates we want to collect.”*

- Pitch: It’s possible to retain some flexibility (e.g., can later break out statistics by geographic area)


*“Please sign this stack of non-disclosure agreements...  
and before we show you the non-disclosure agreements...  
you have to sign a preliminary non-disclosure agreement.”*

- ???


# Today




**StressTracker**



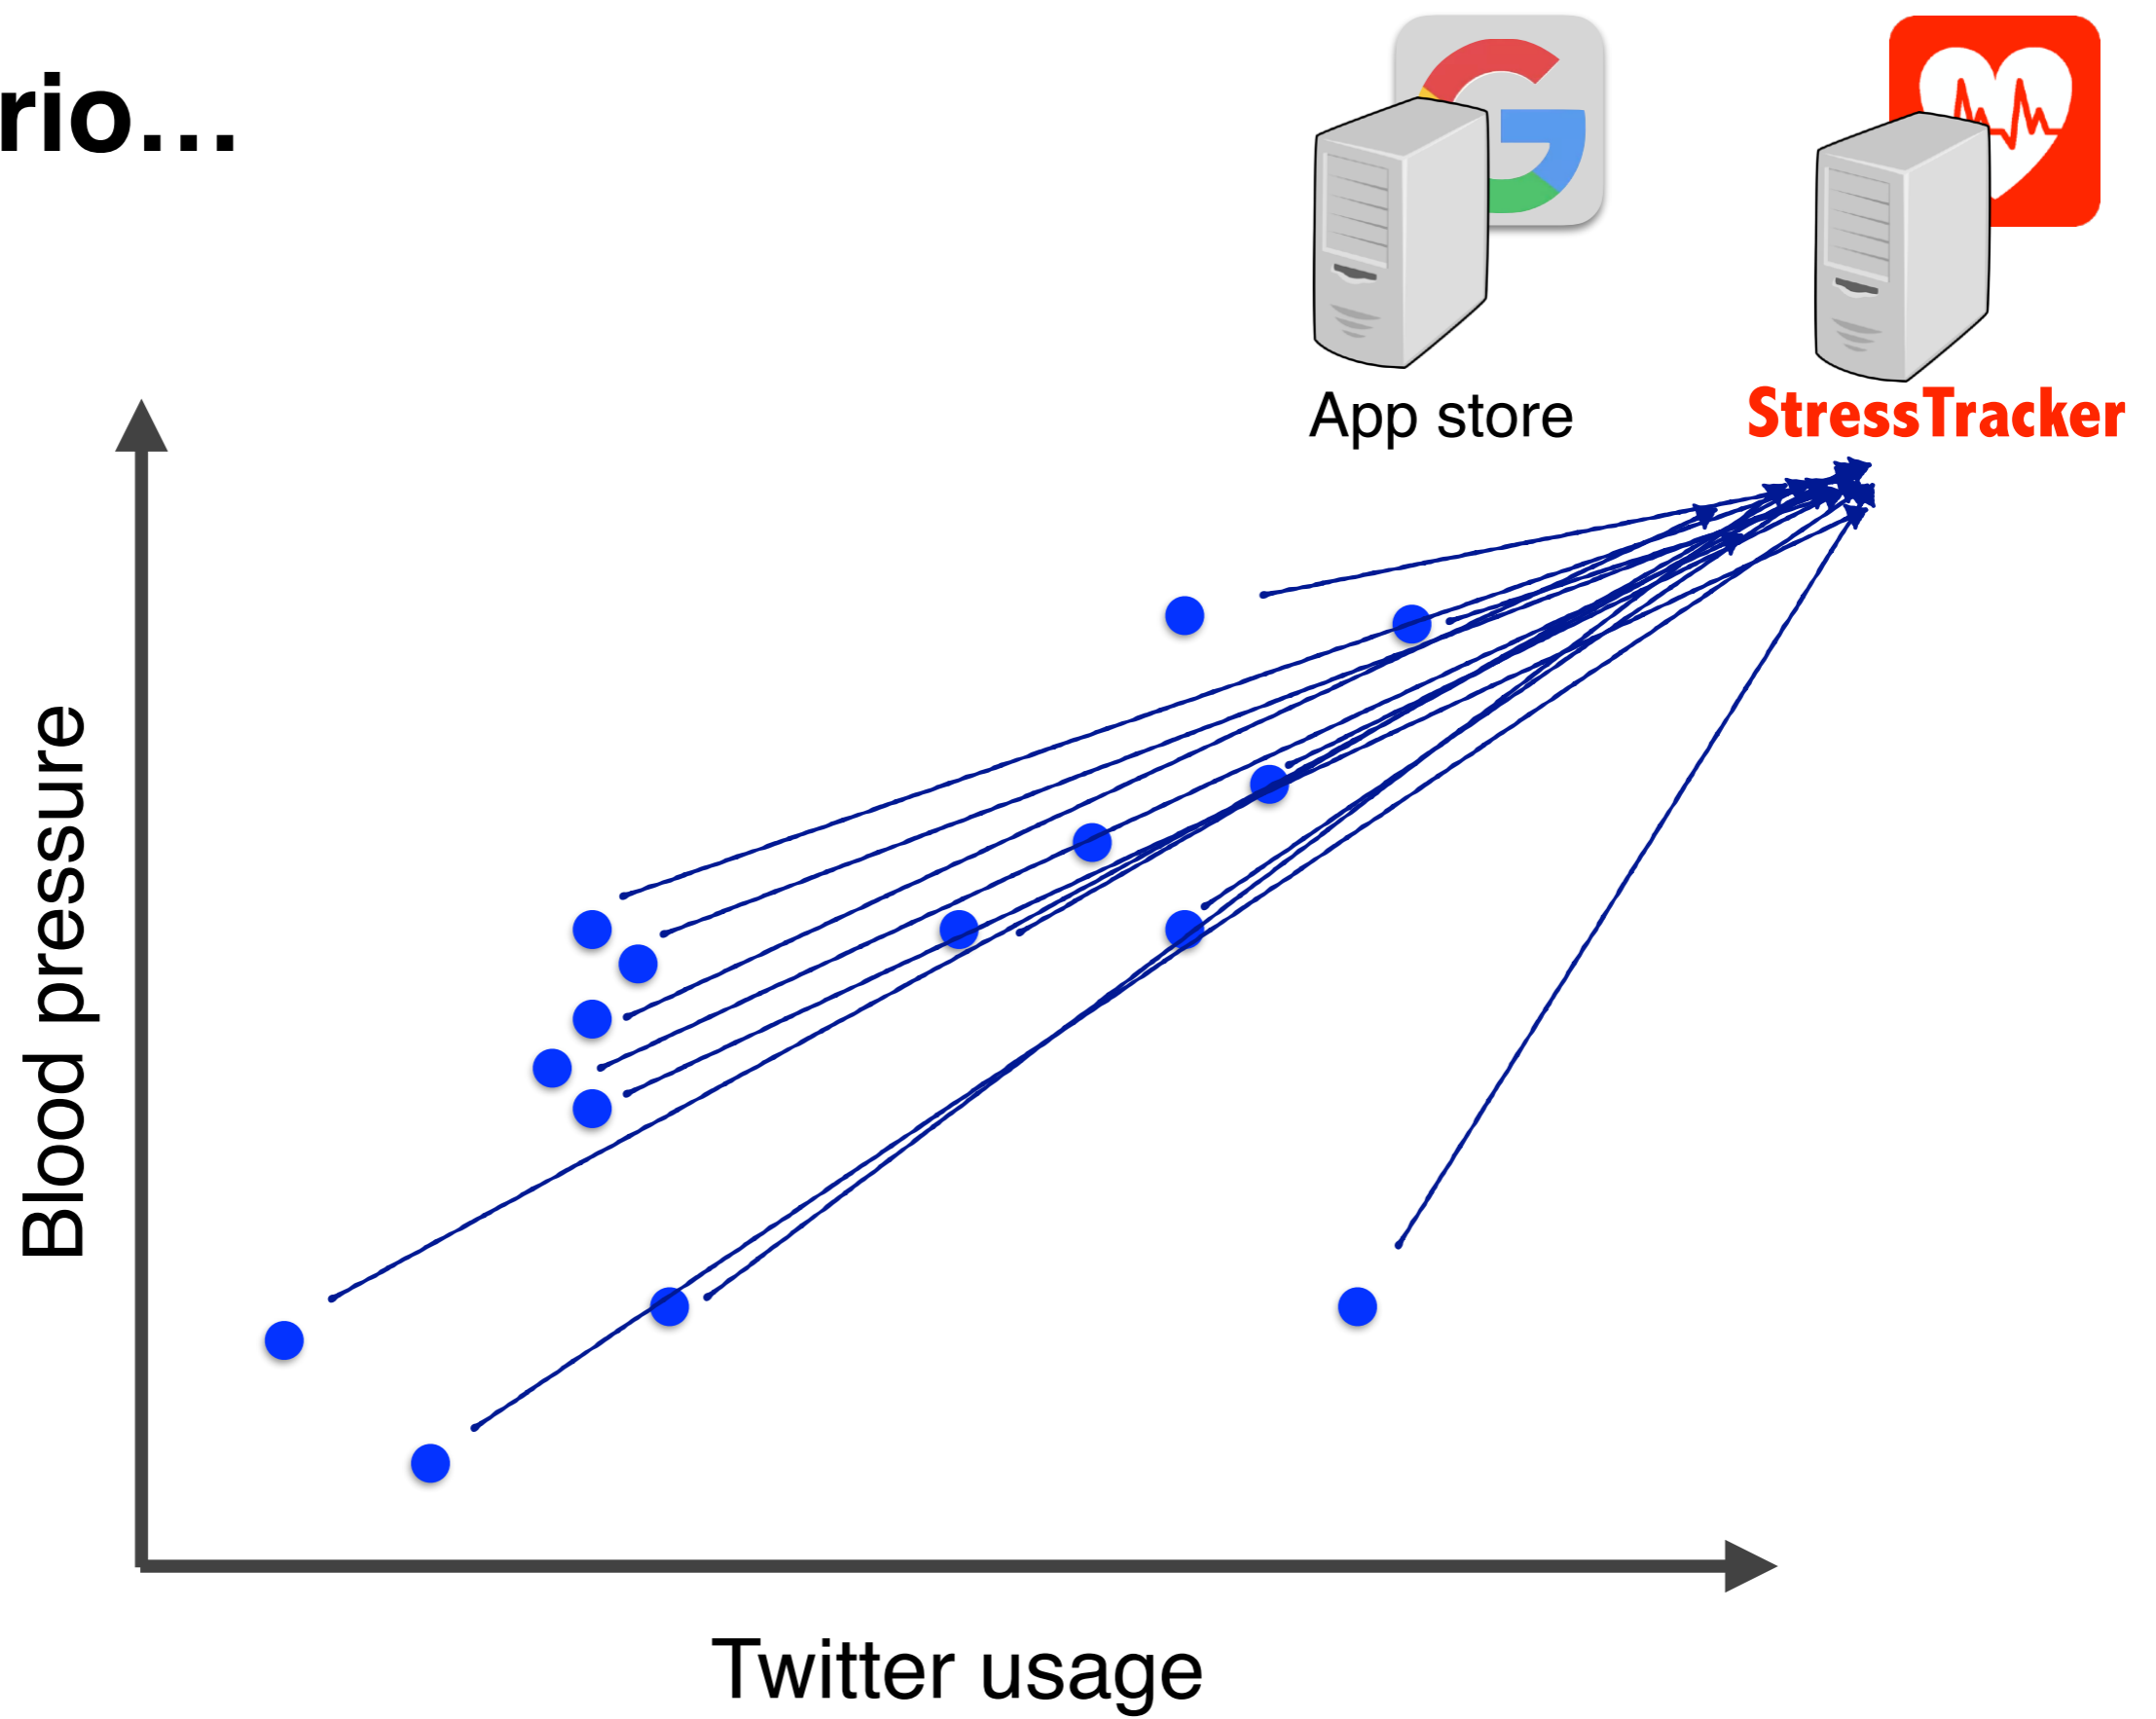
# Today



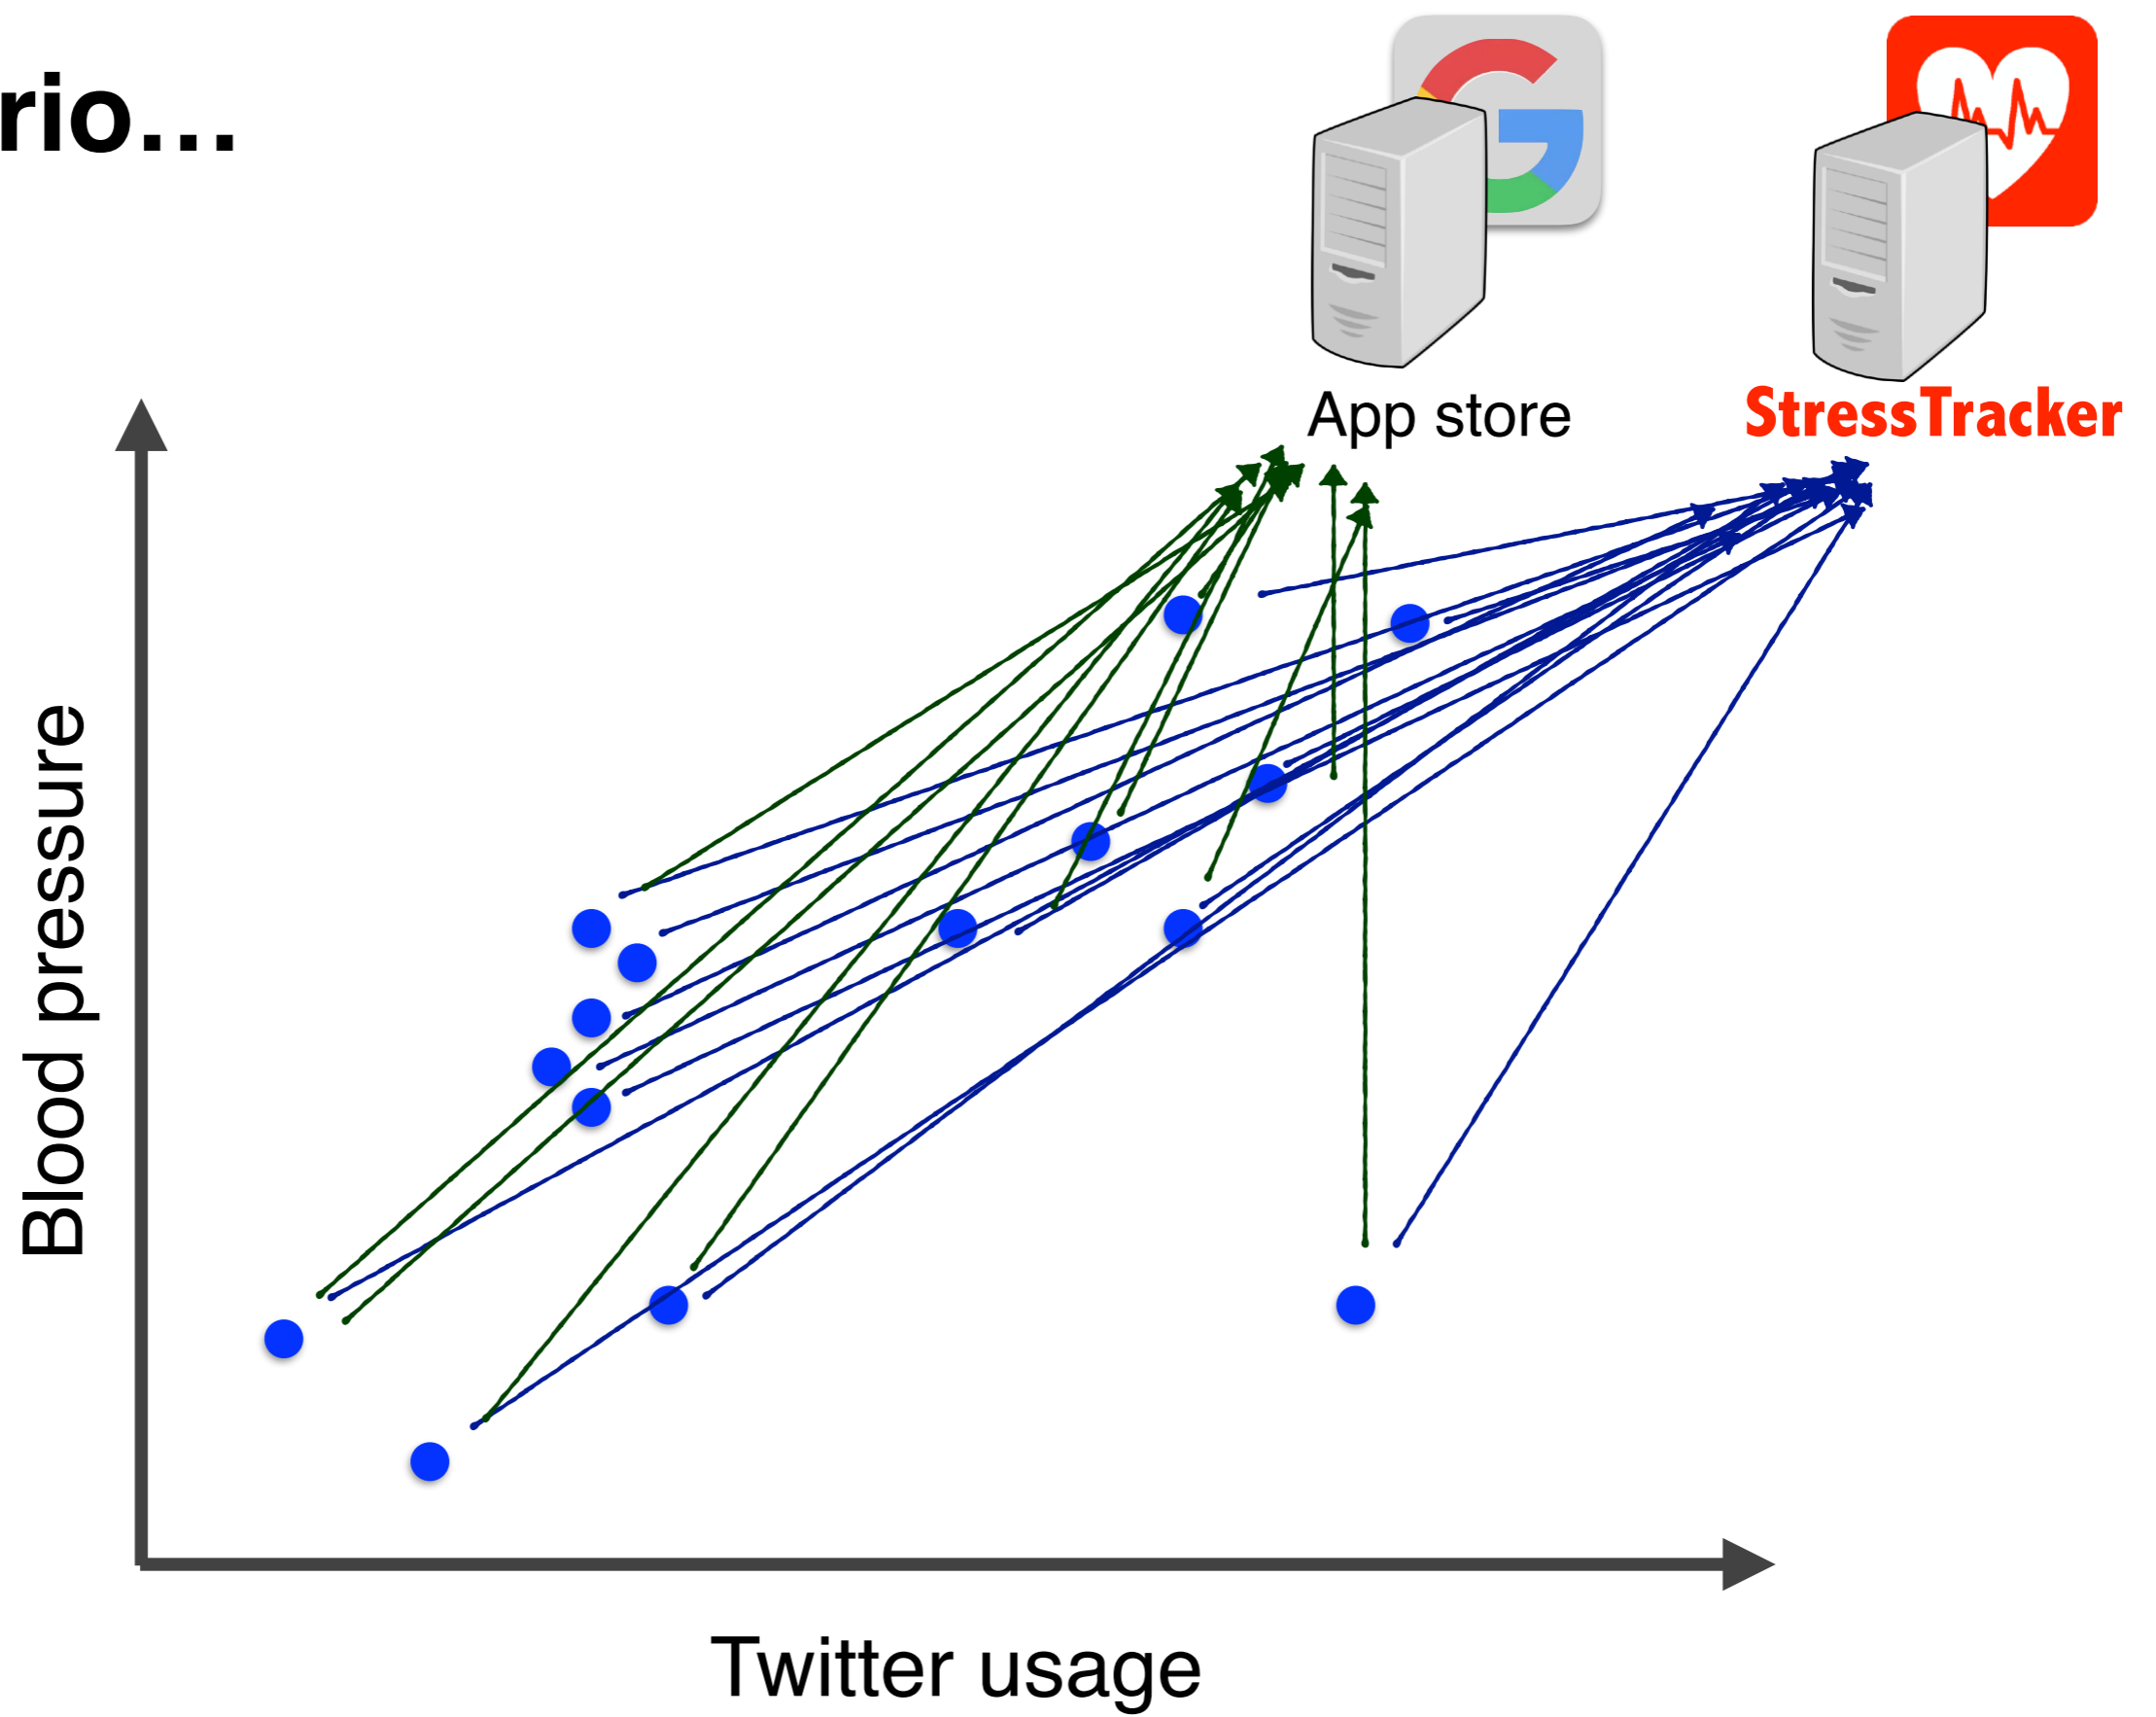
**With Prio...**



**With Prio...**

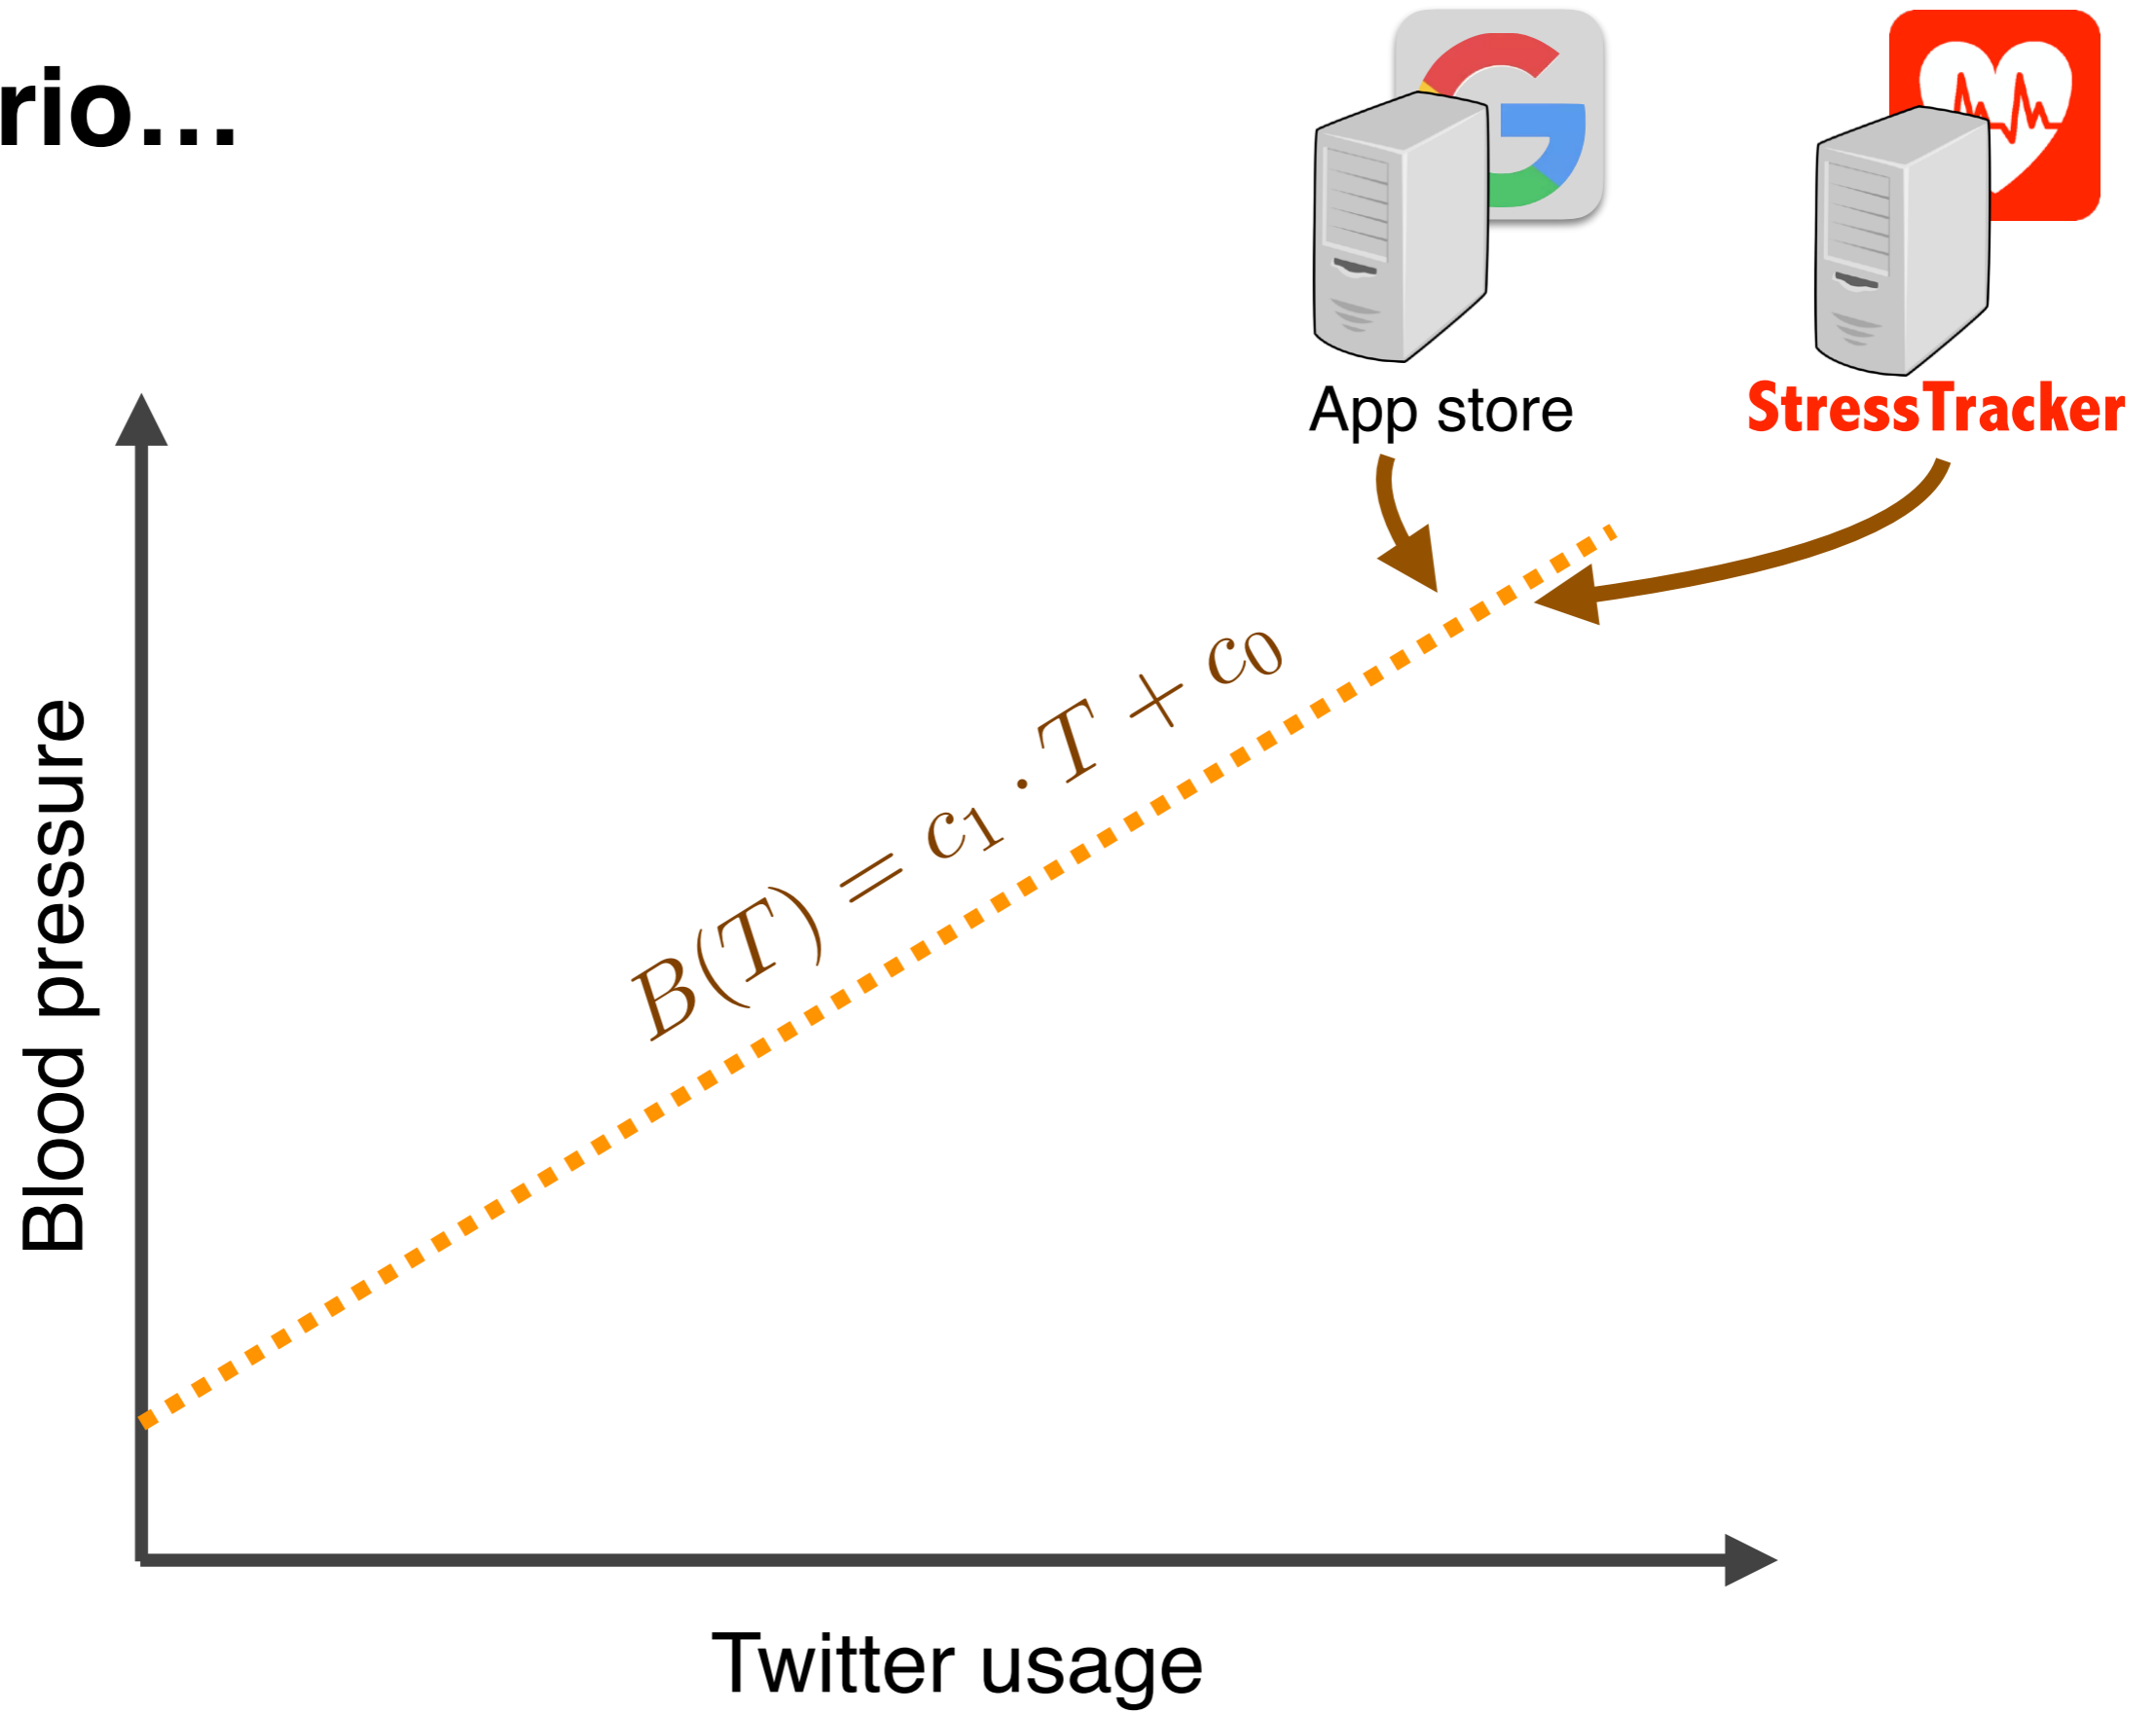


**With Prio...**





**With Prio...**



# Conclusions

- Wholesale collection of sensitive user data puts our security at risk.
- Prio is the first system for aggregation that provides:
  - **exact correctness**,
  - **privacy**,
  - **robustness**, and
  - **efficiency**.
- To do so, Prio uses **SNIPs** and **aggregatable encodings**.
- These techniques together bring private aggregation closer to practical.

**Thank you!**

