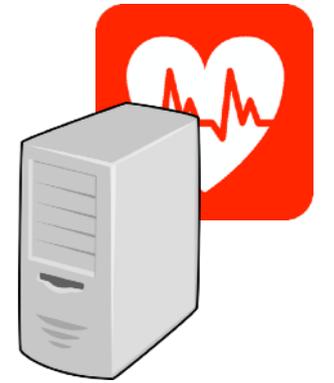


Prio: Private, Robust, and Efficient Computation of Aggregate Statistics

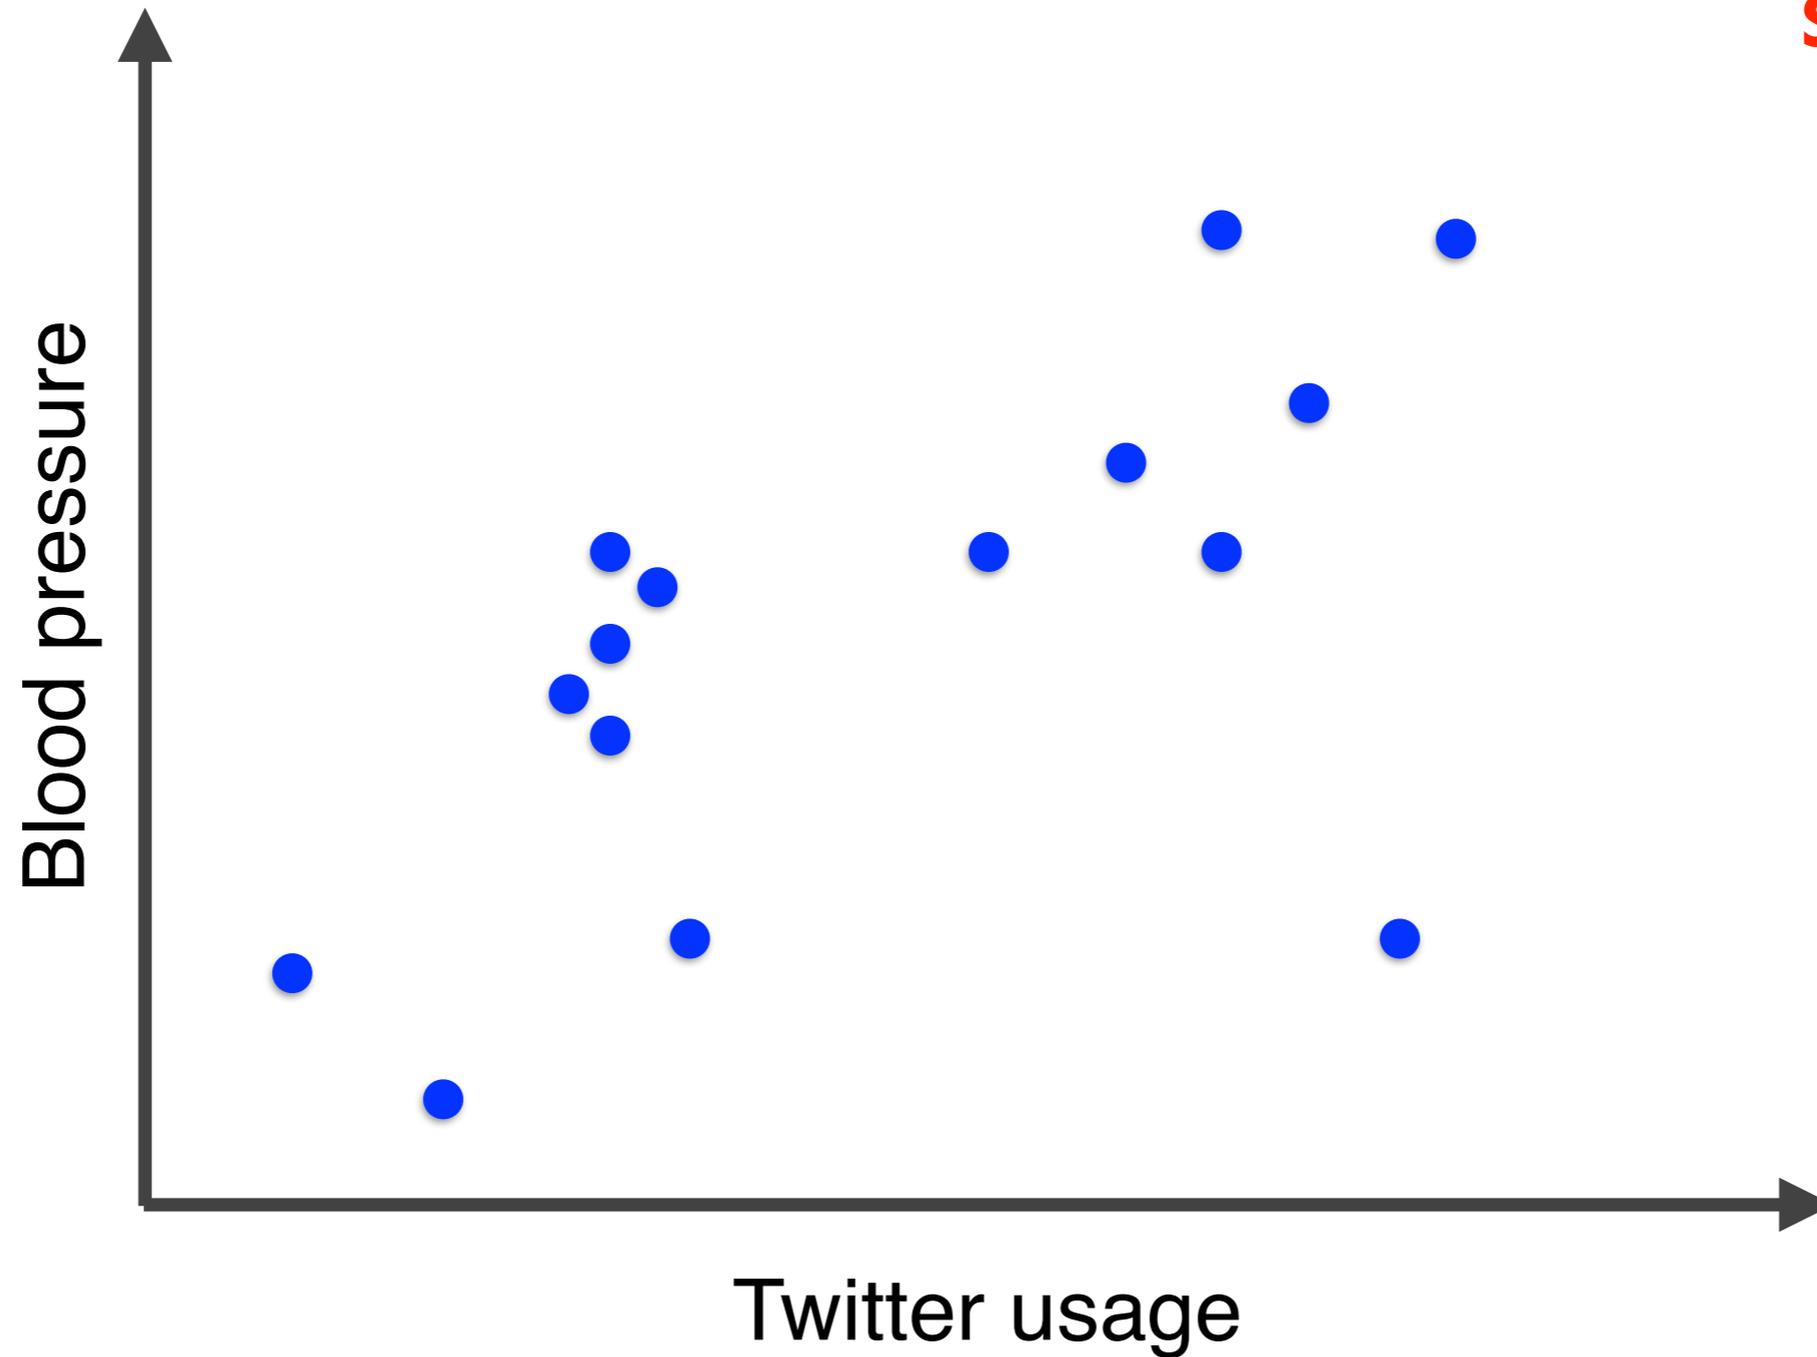
Henry Corrigan-Gibbs and Dan Boneh
Stanford University

NSDI 2017

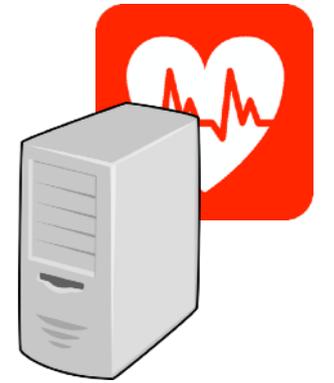
Today: Non-private aggregation



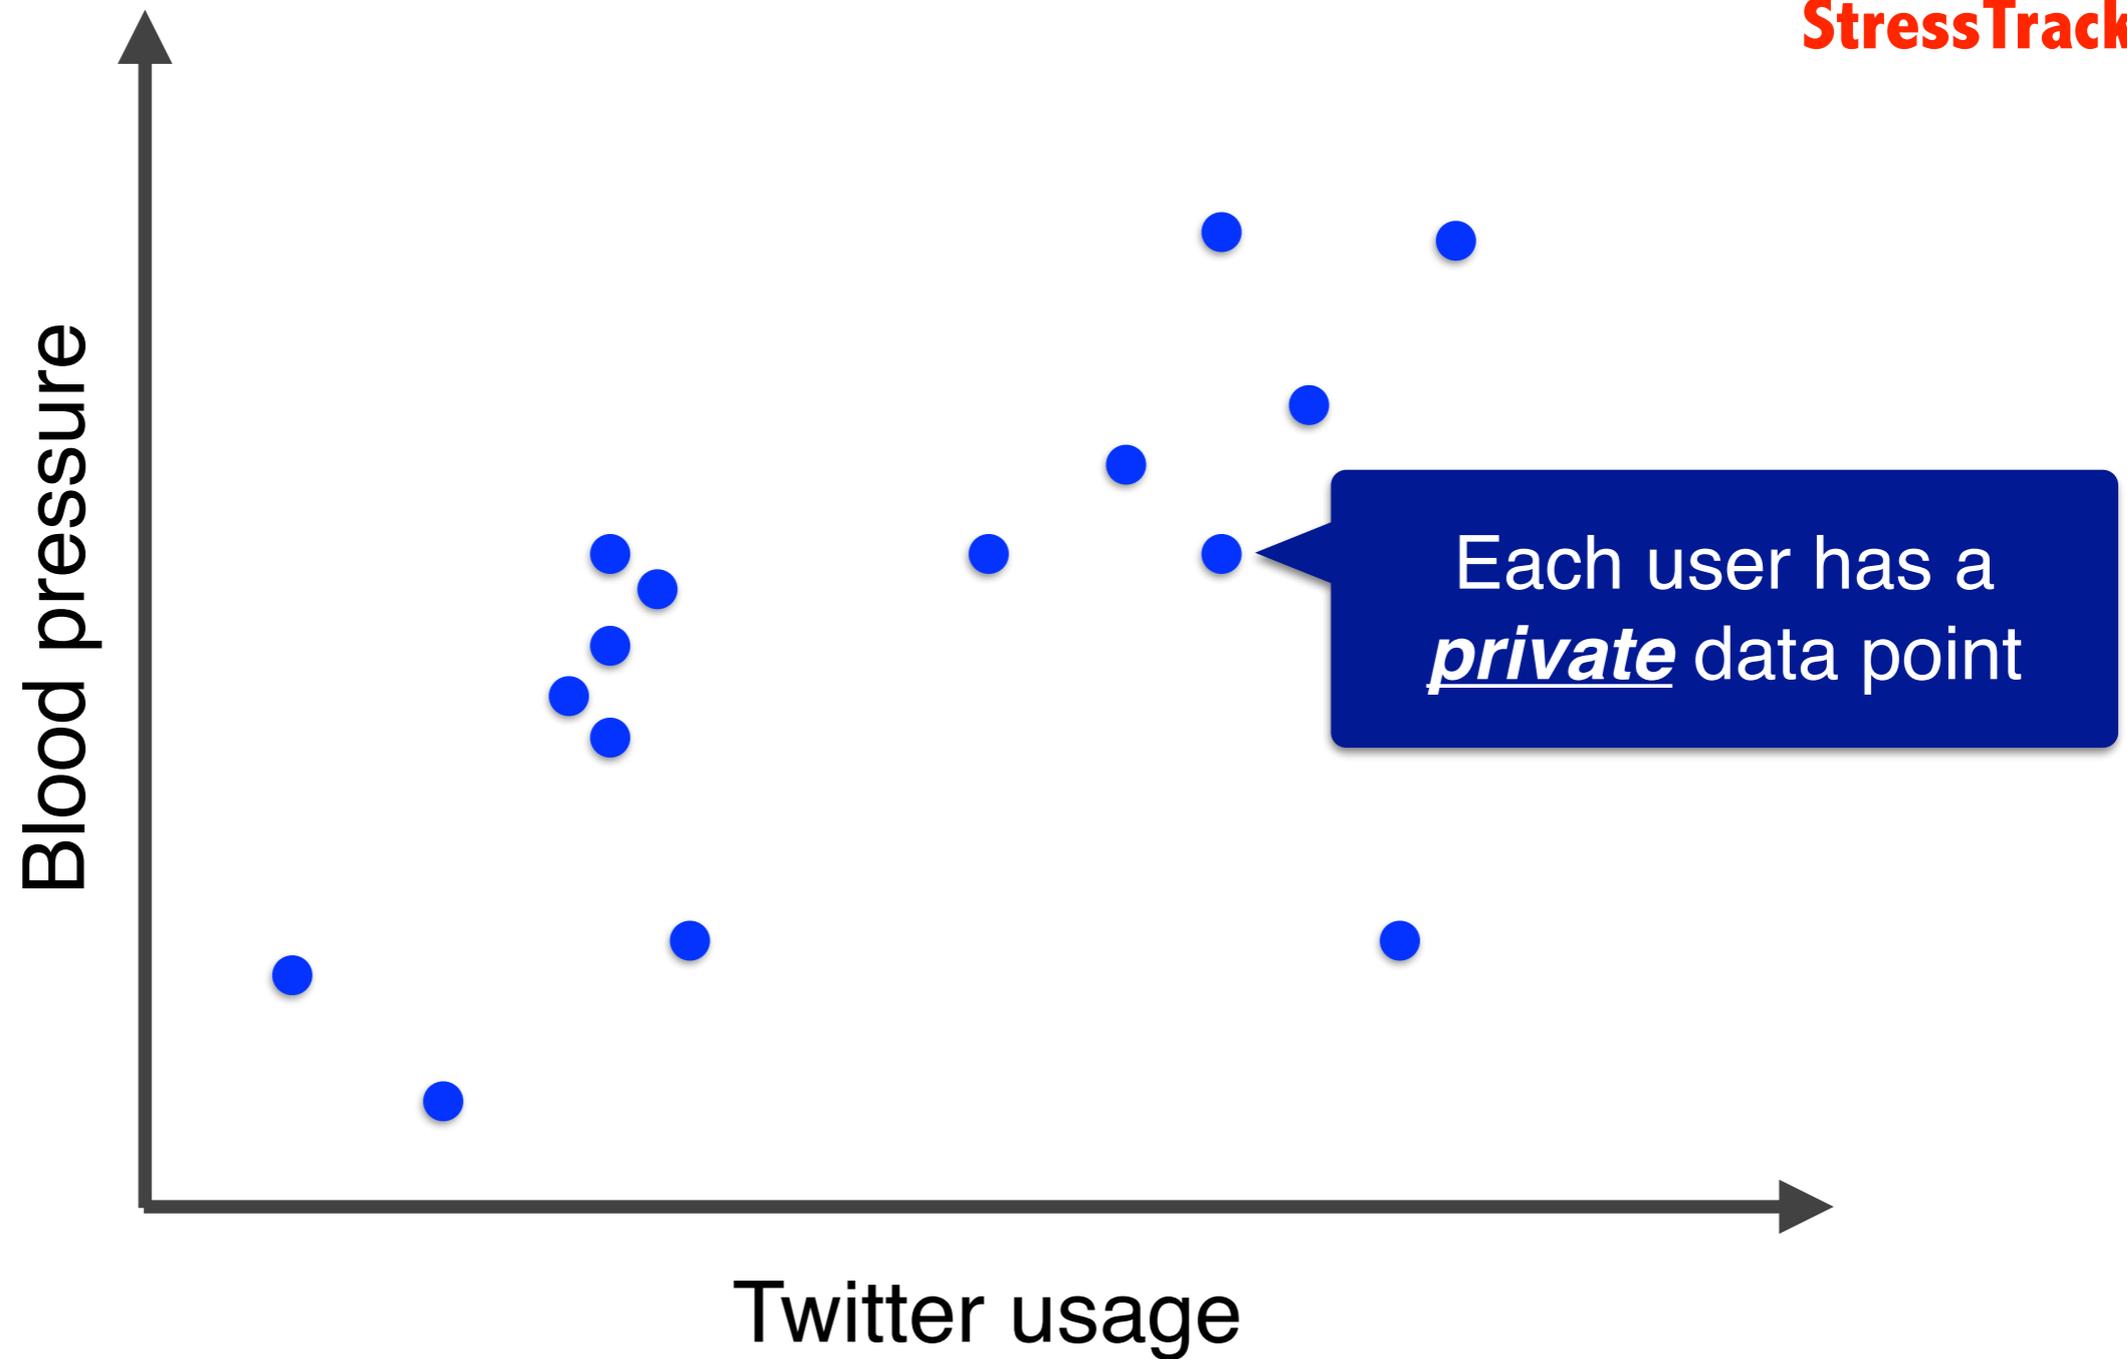
StressTracker



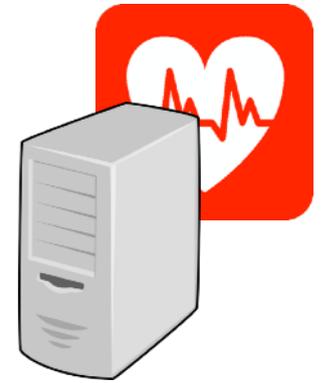
Today: Non-private aggregation



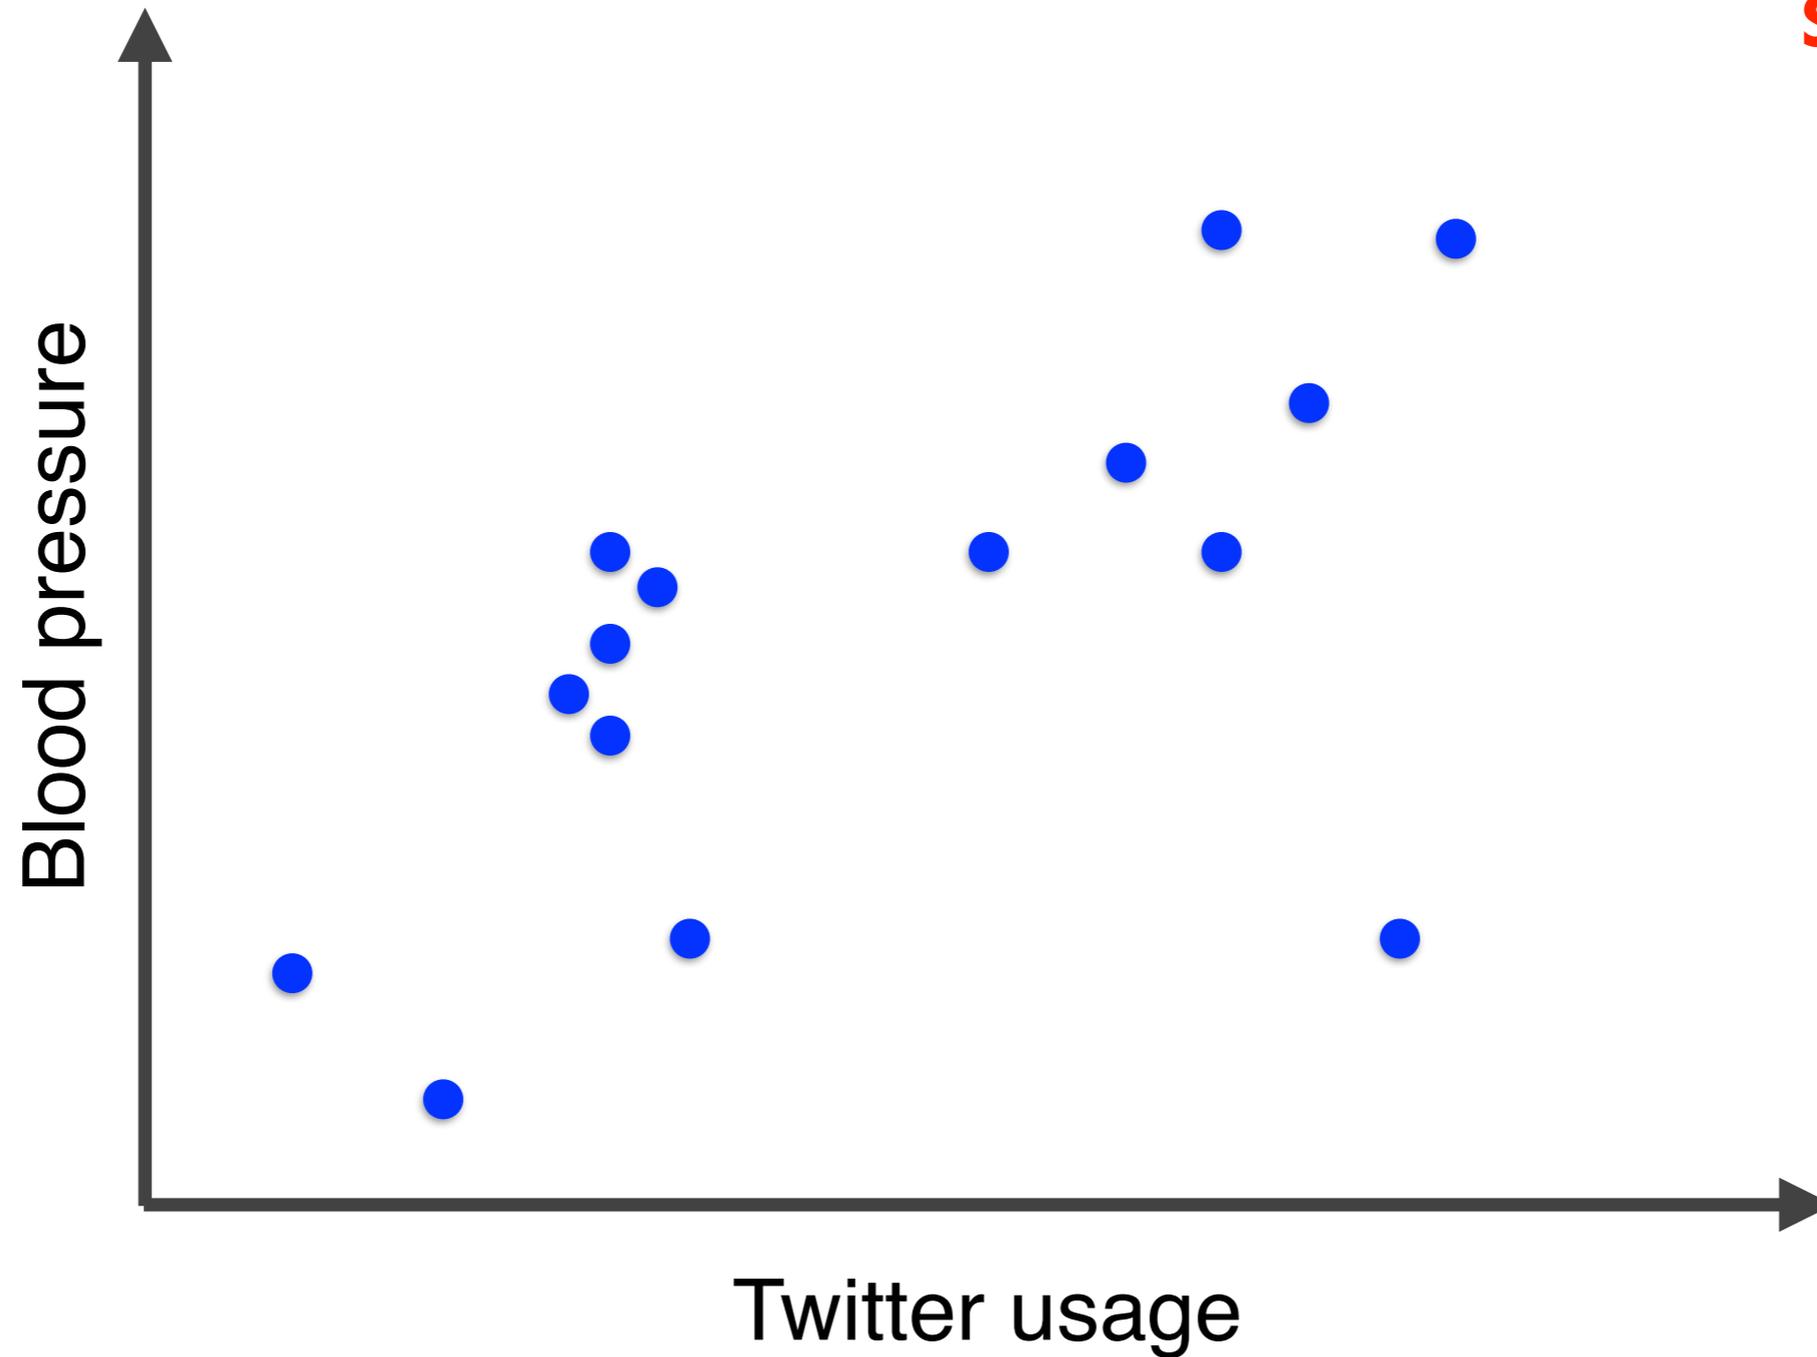
StressTracker



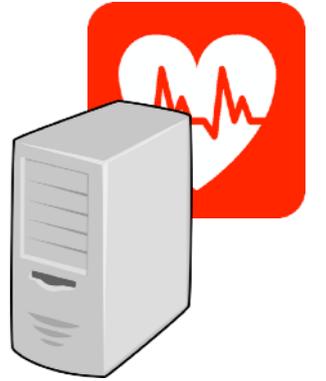
Today: Non-private aggregation



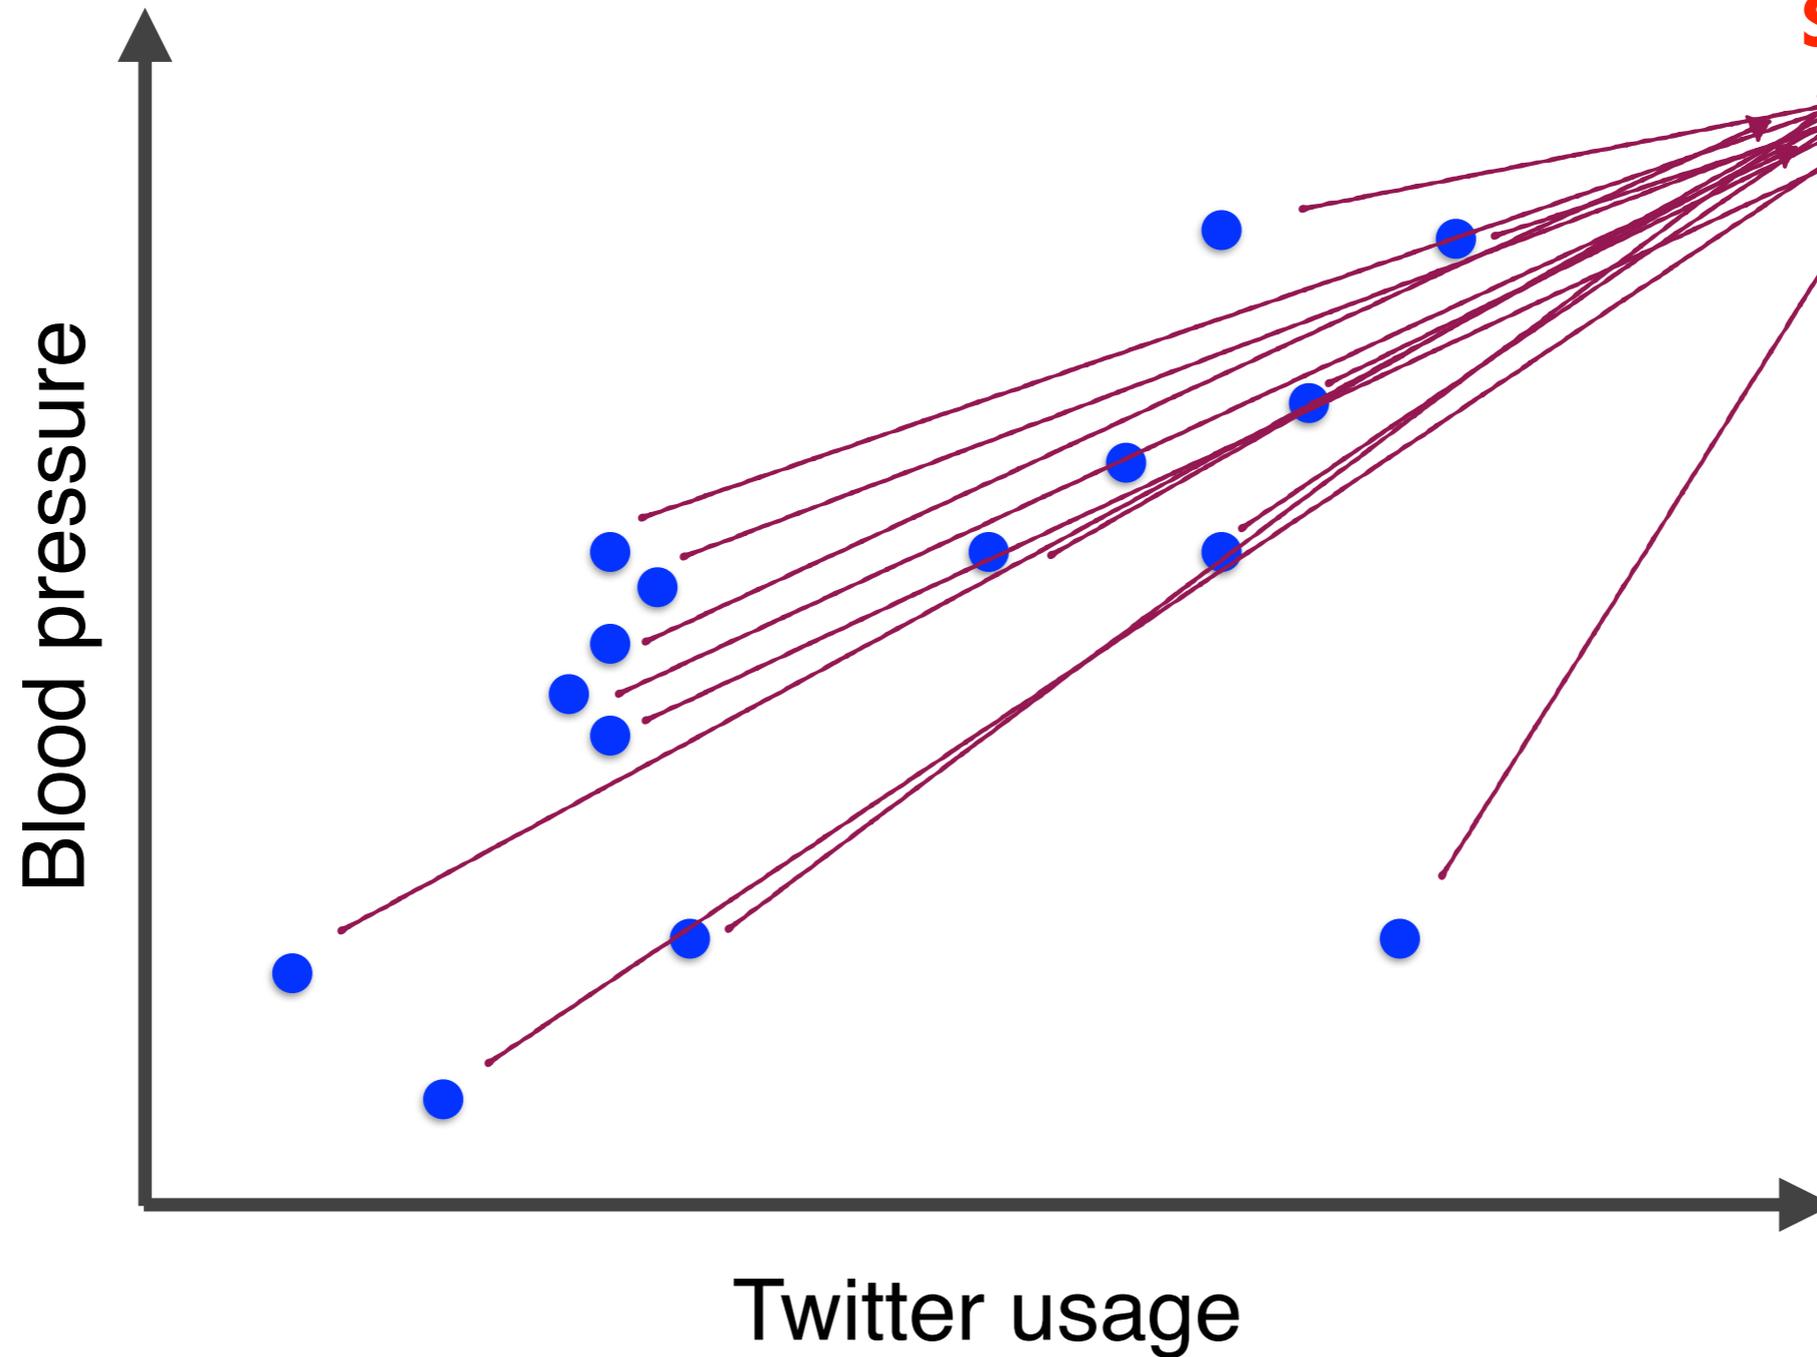
StressTracker



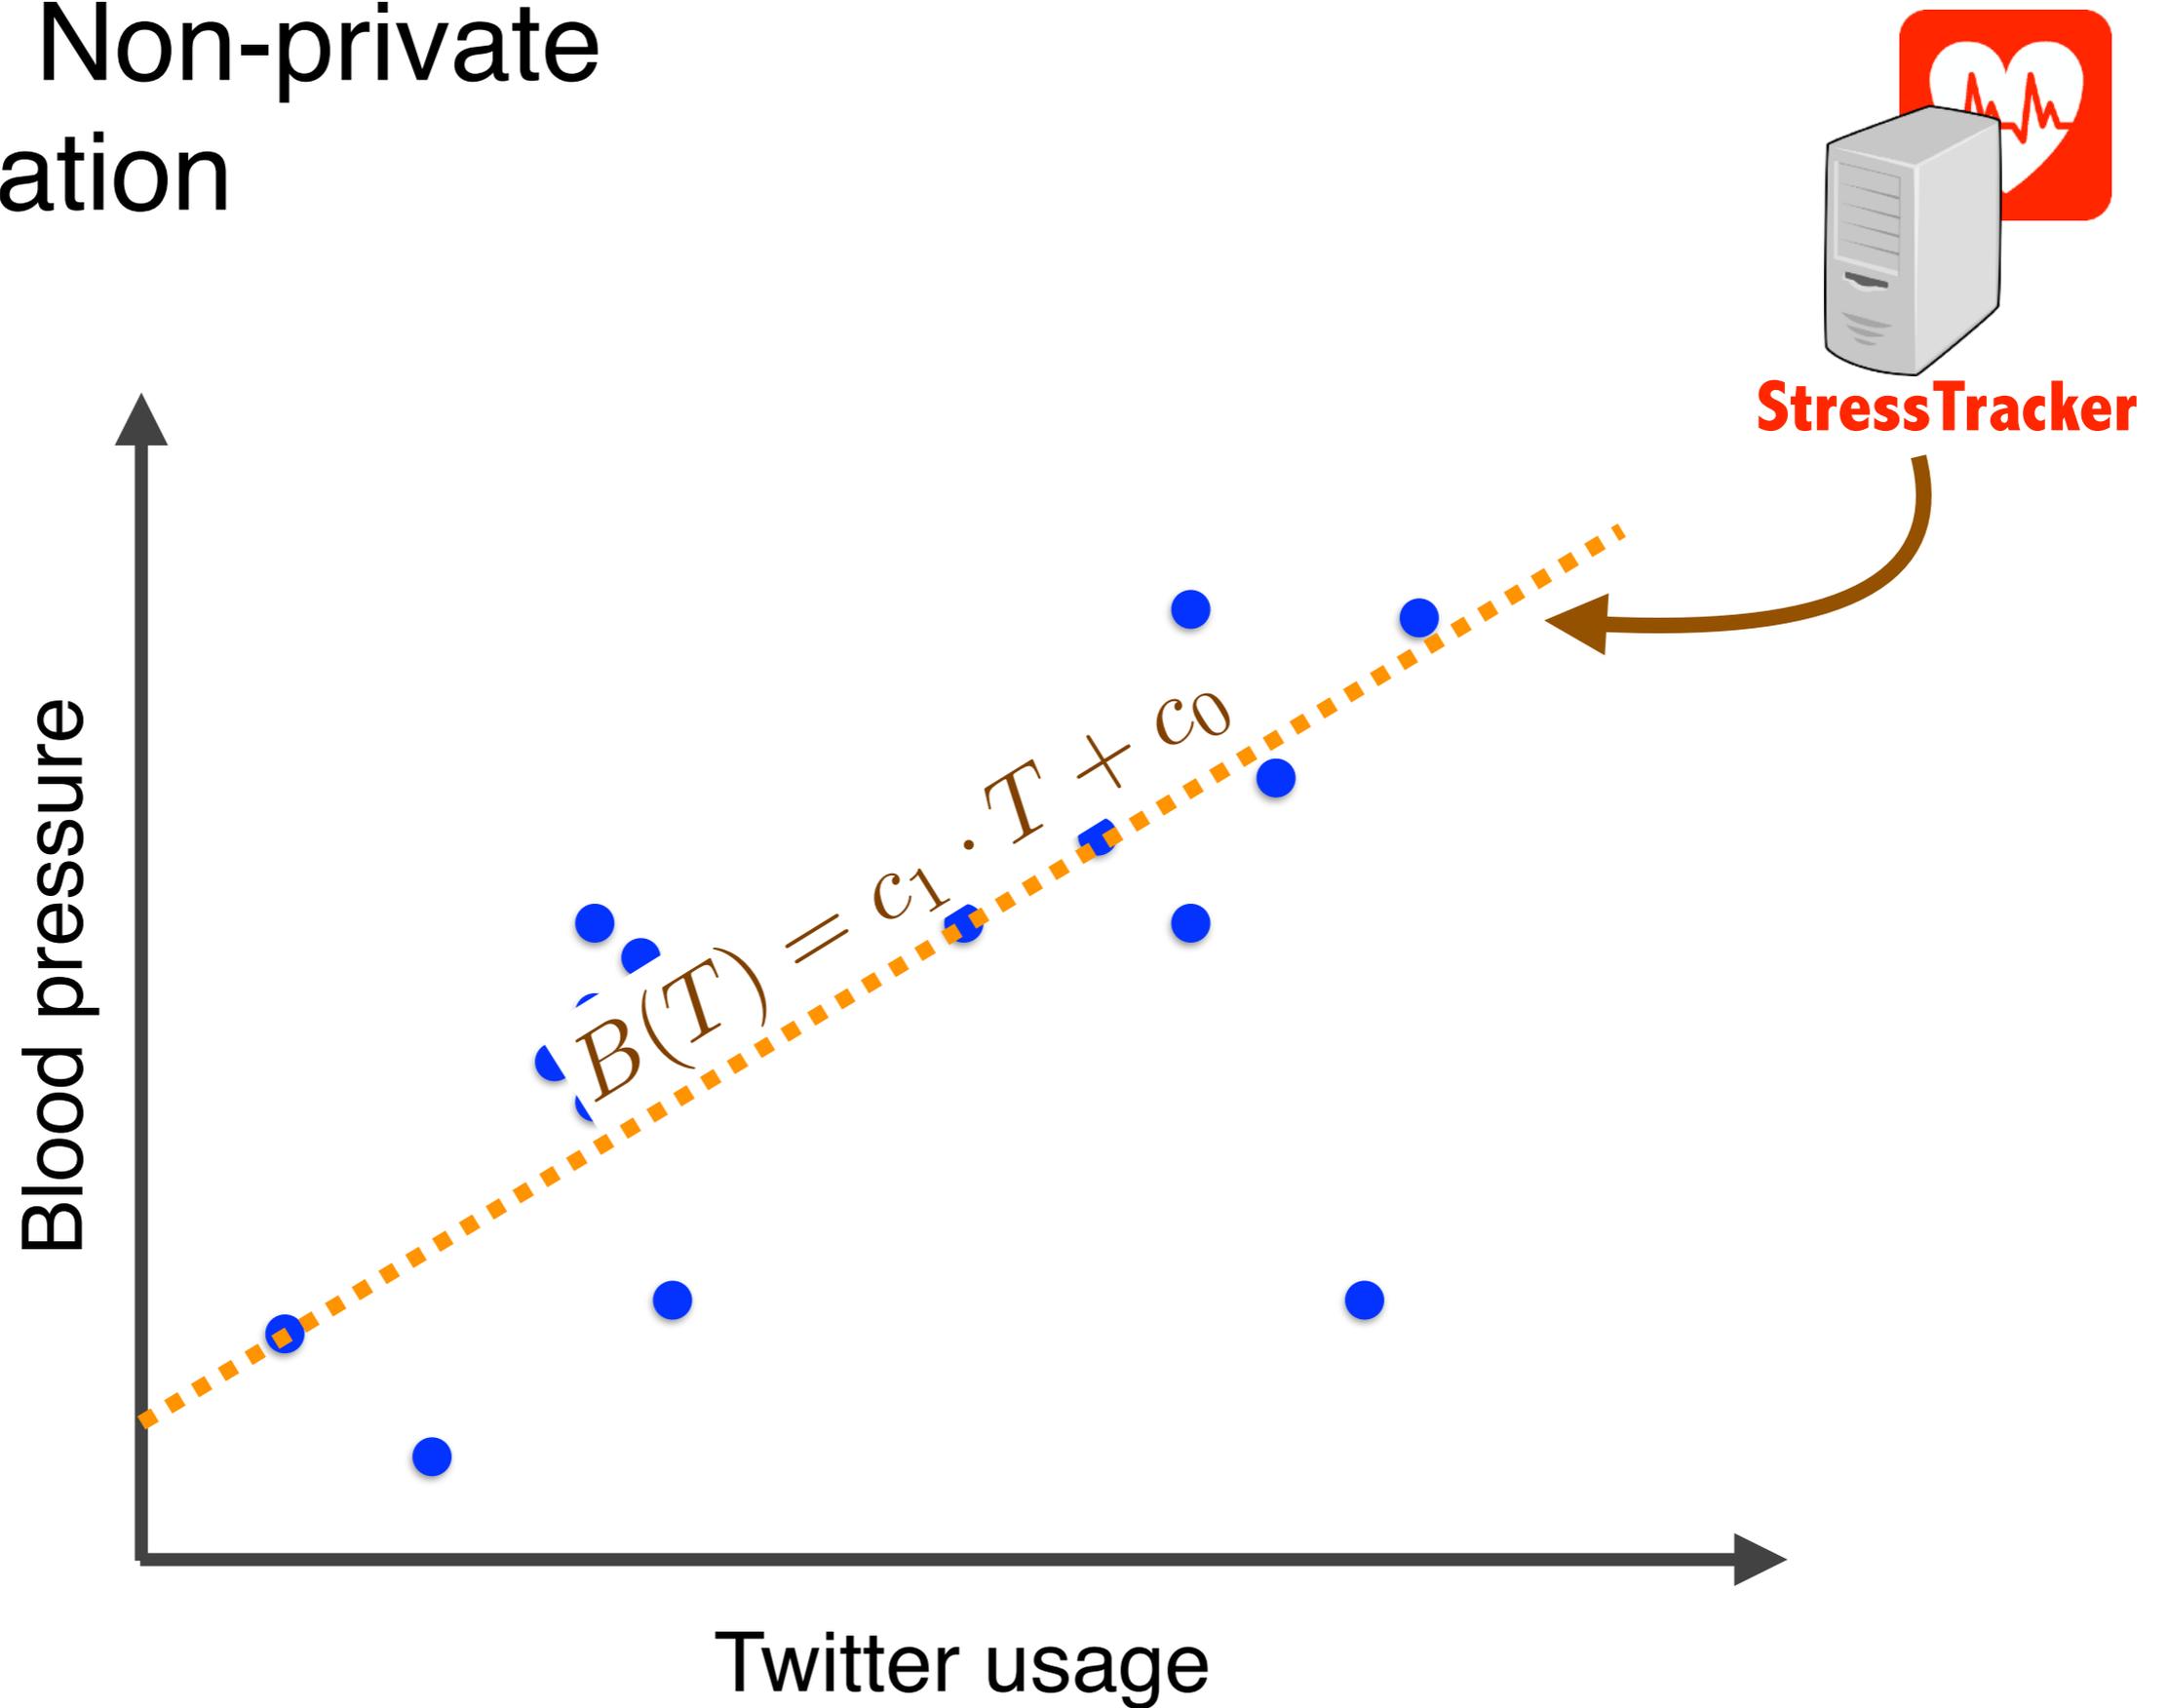
Today: Non-private aggregation



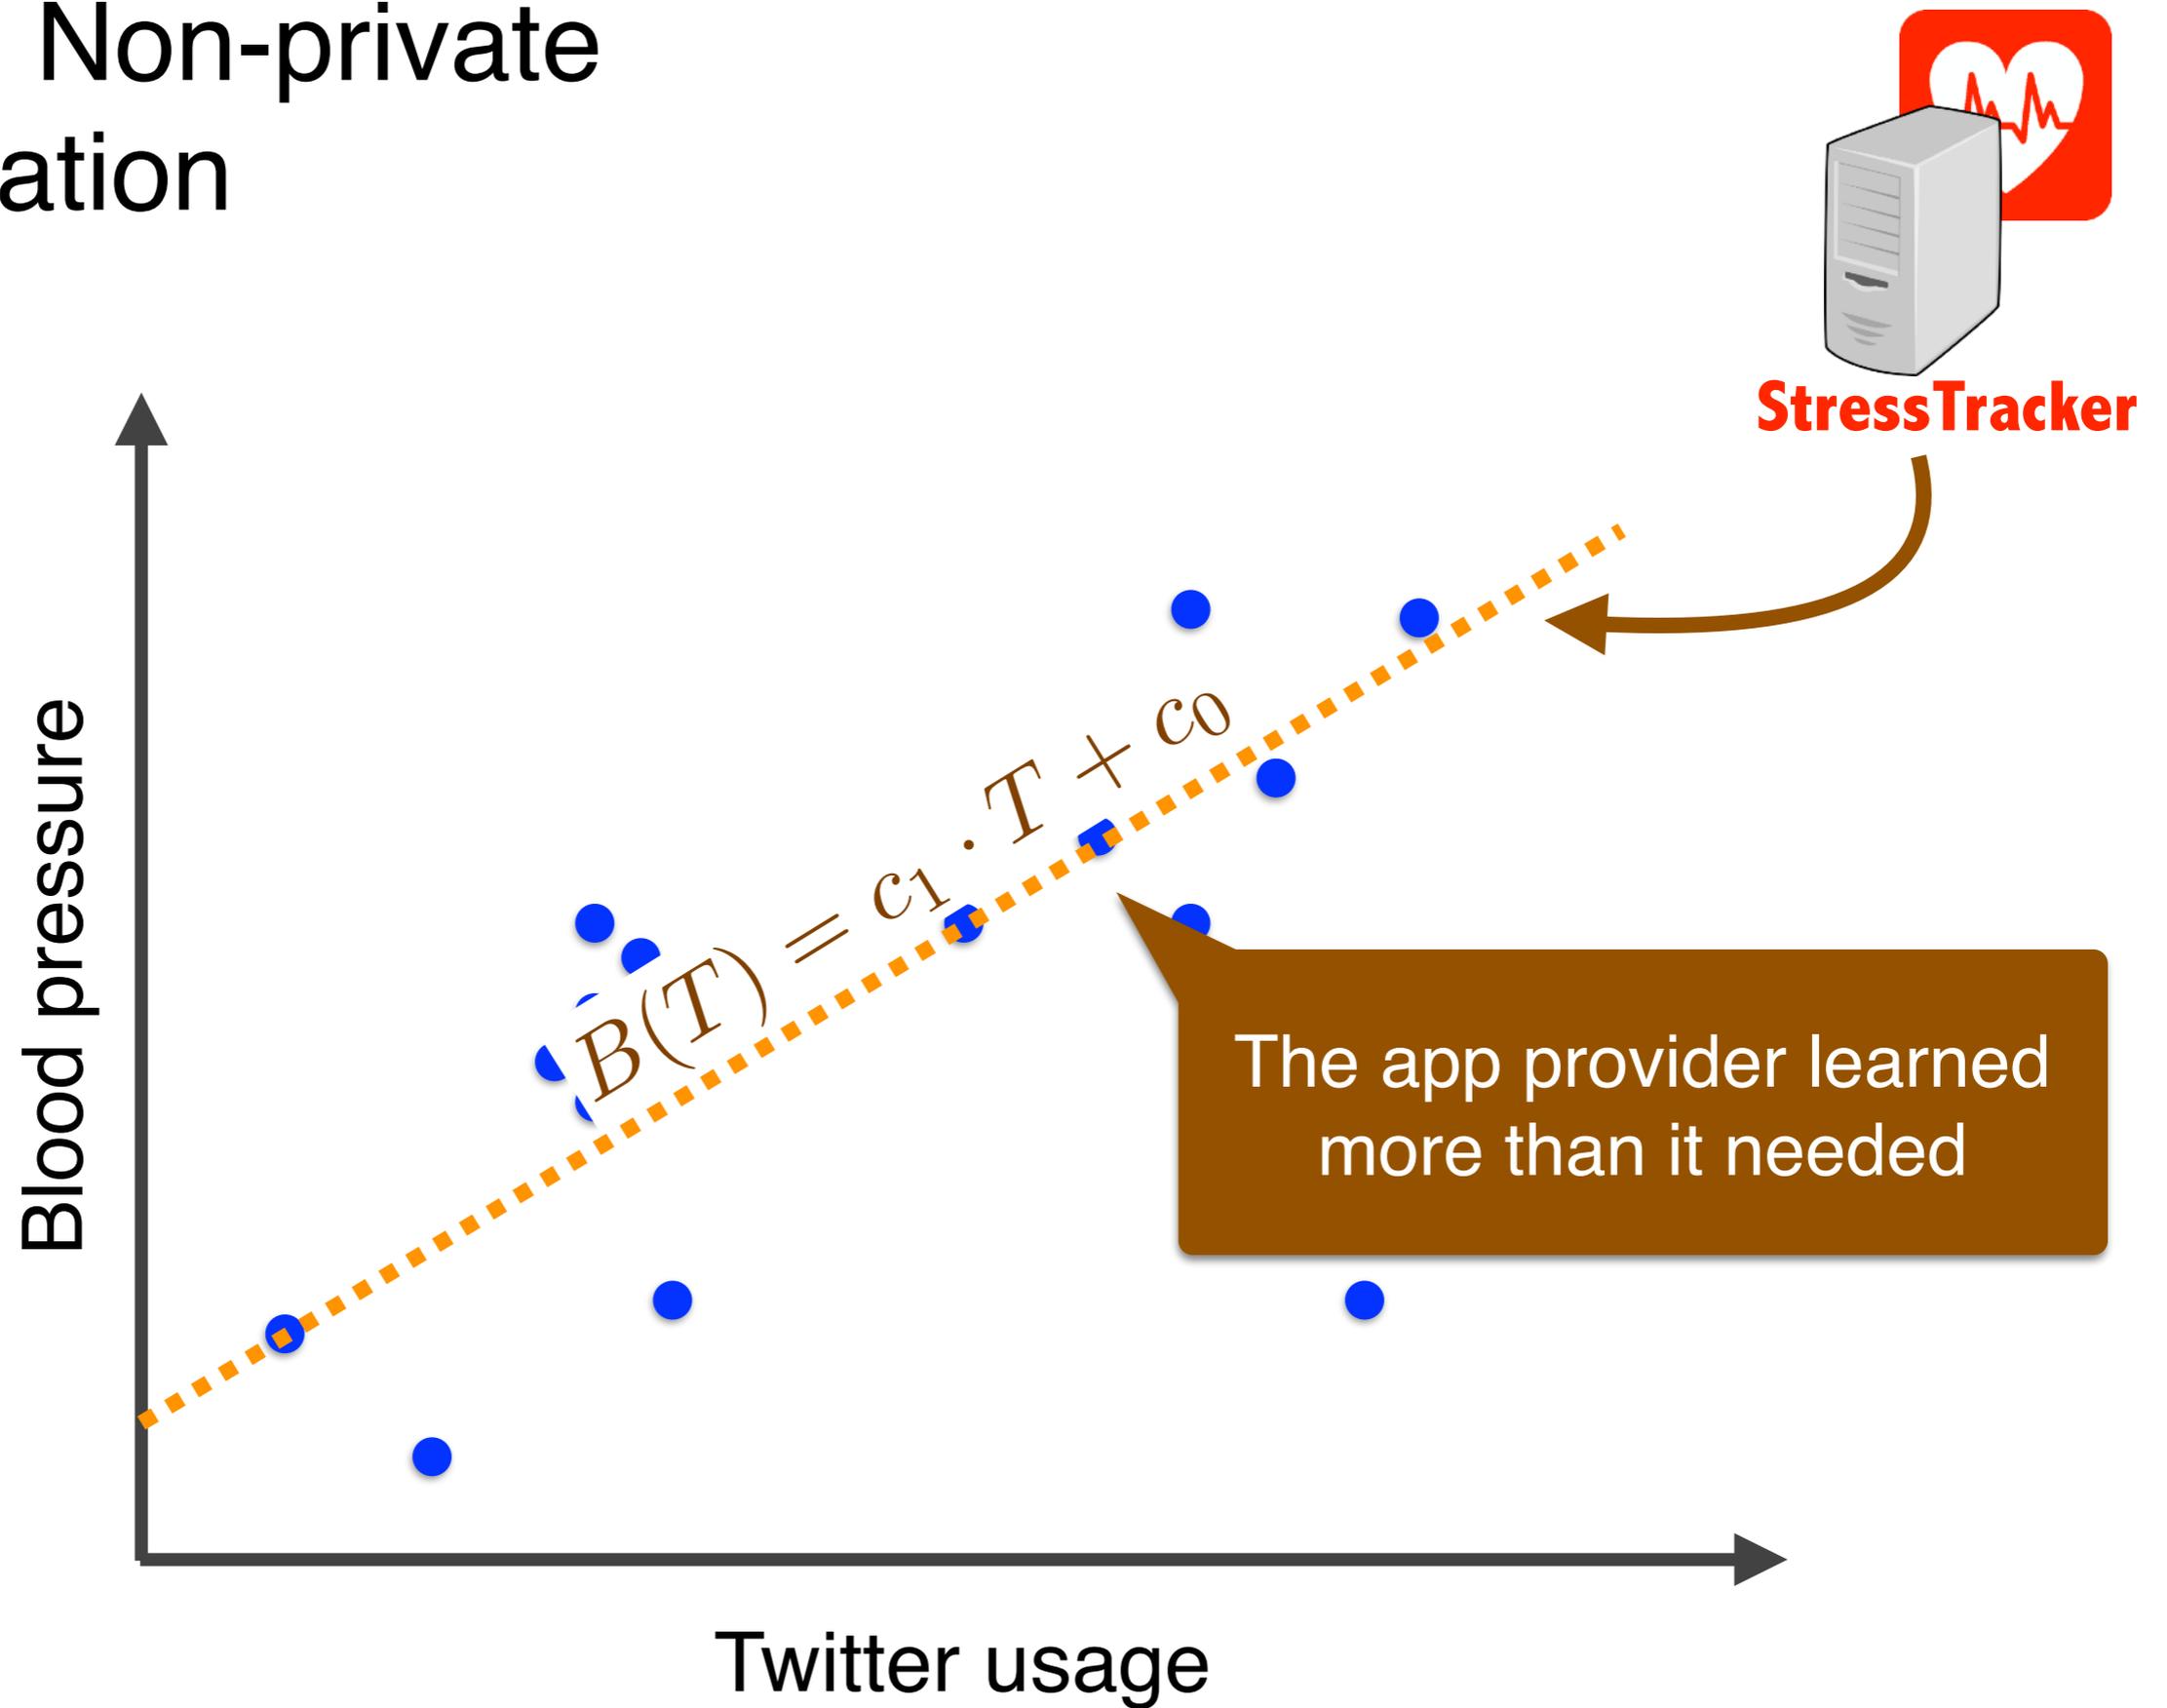
StressTracker



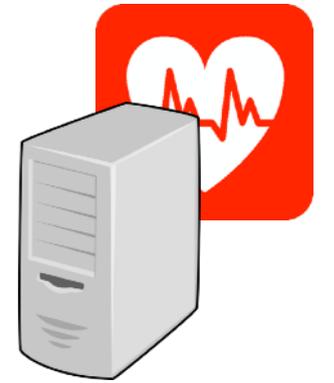
Today: Non-private aggregation



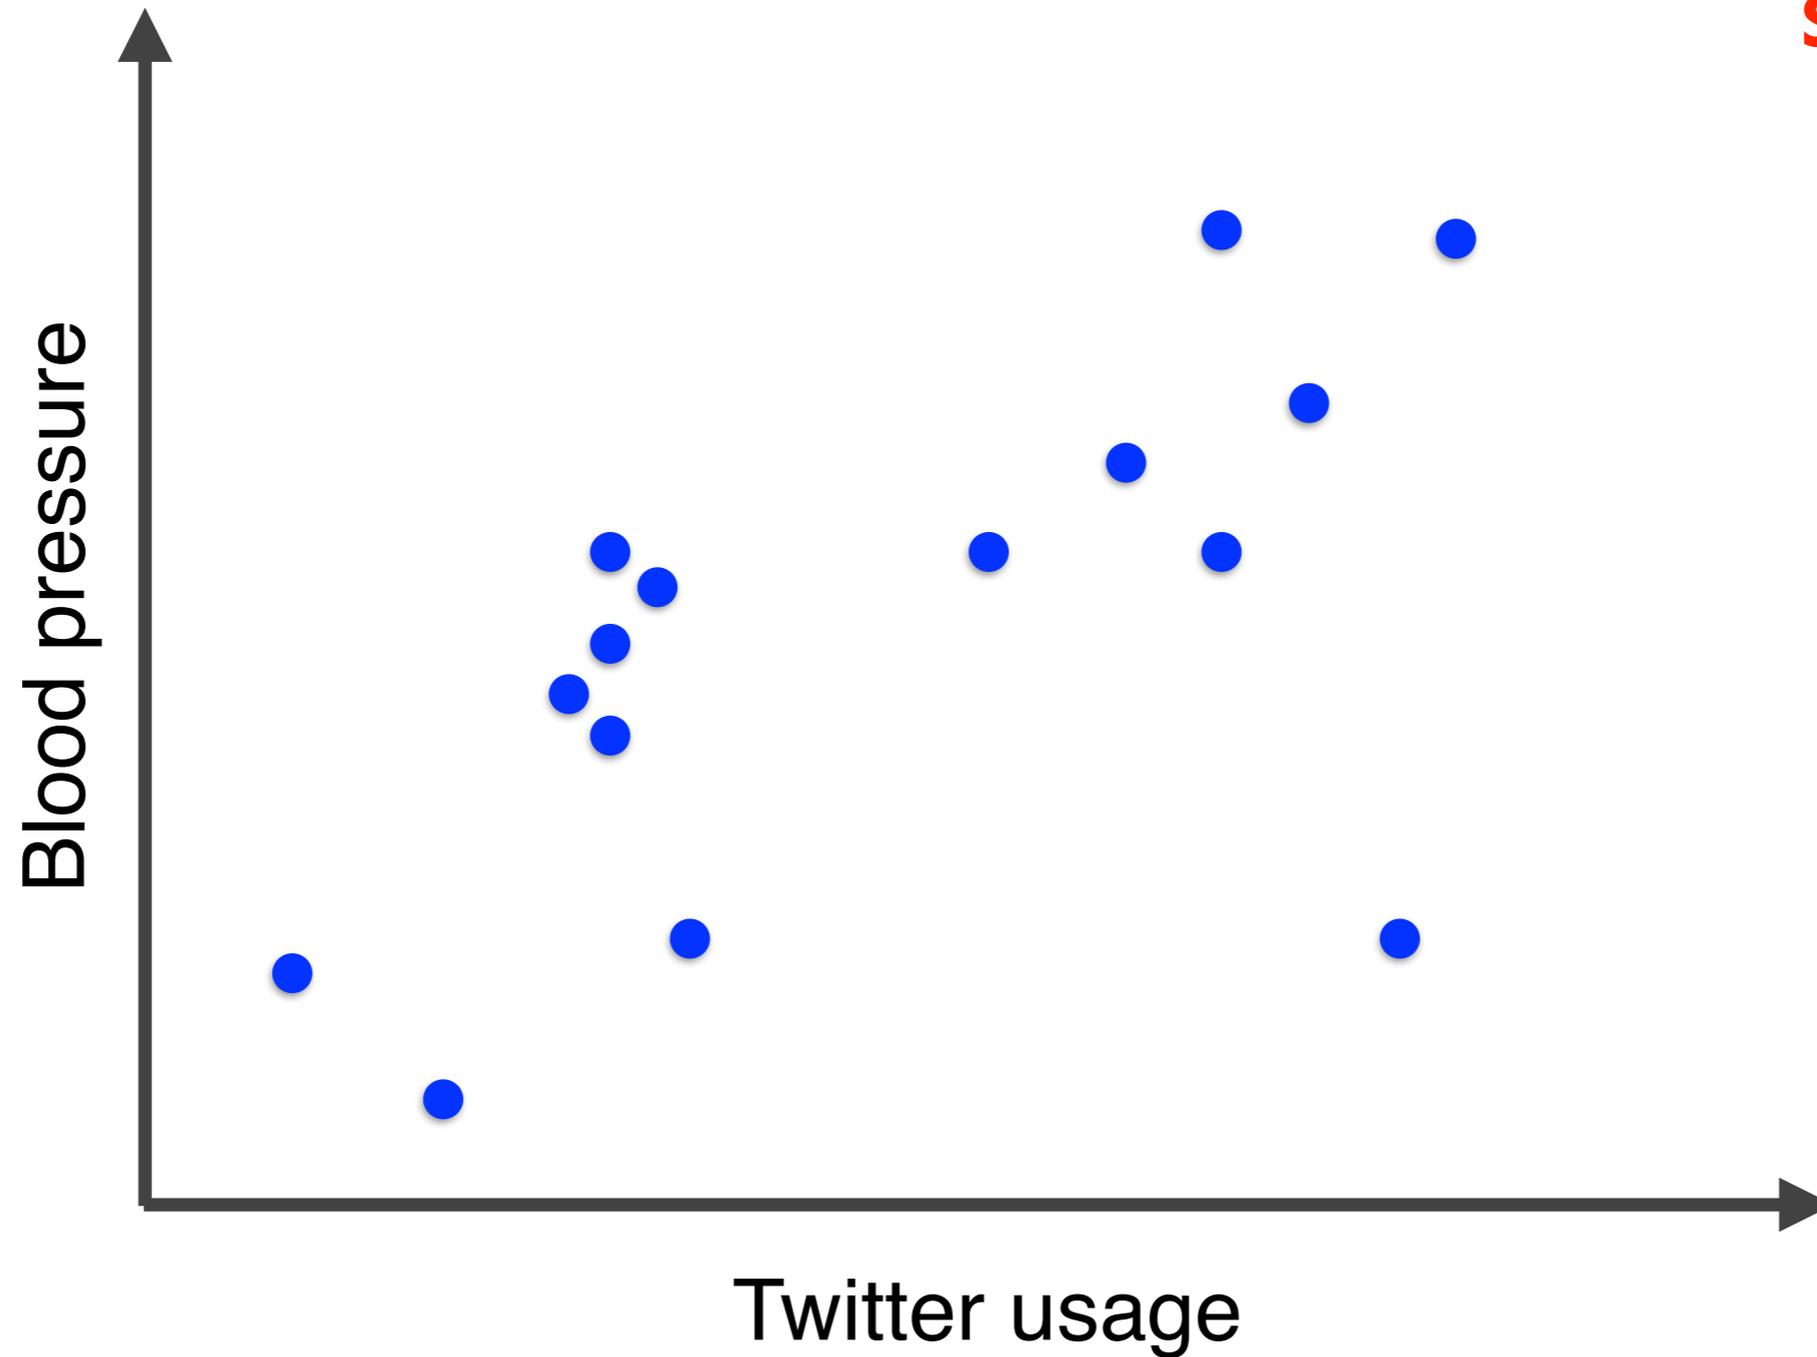
Today: Non-private aggregation



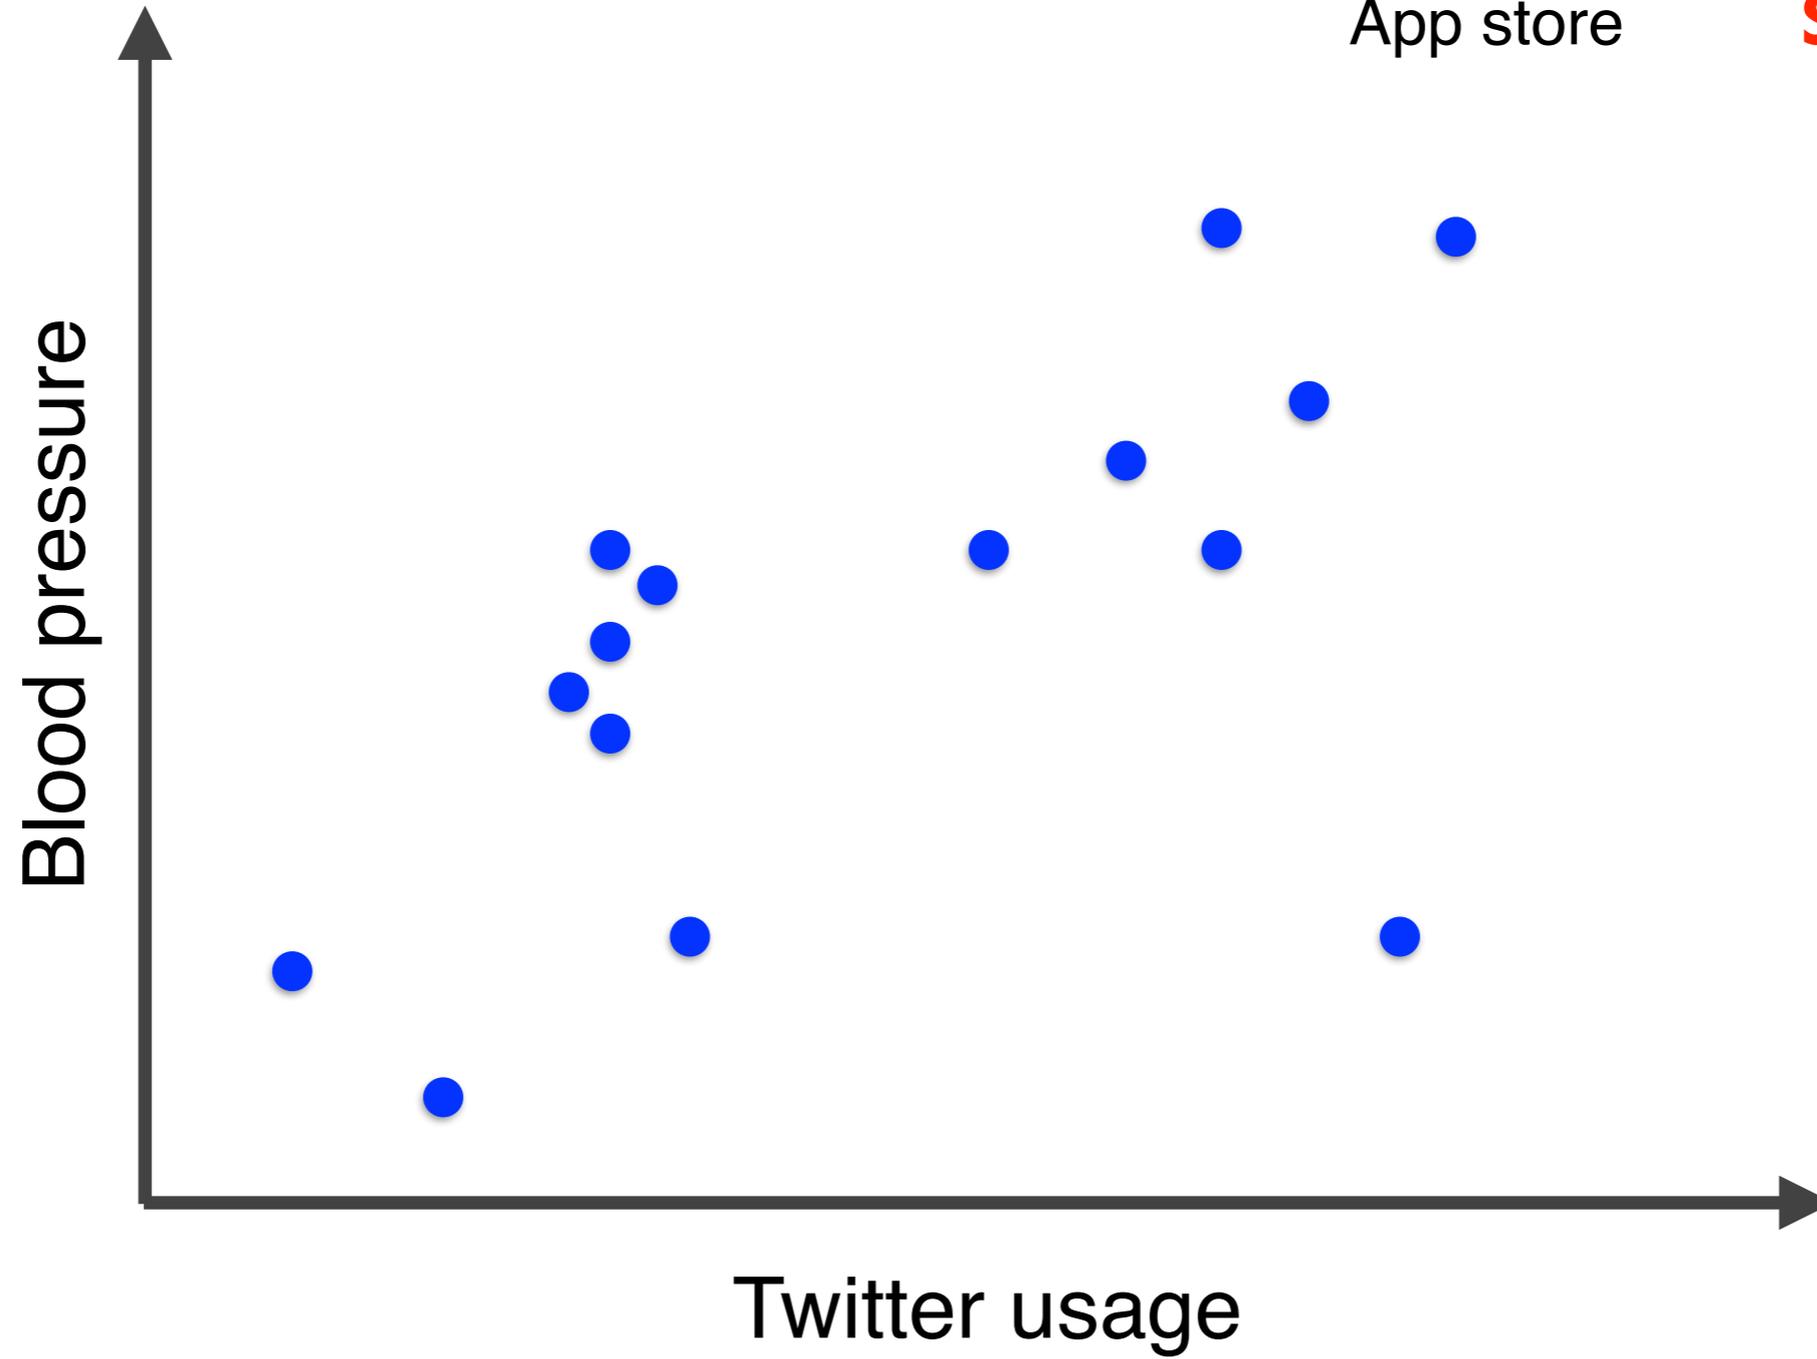
Today: Non-private aggregation



StressTracker



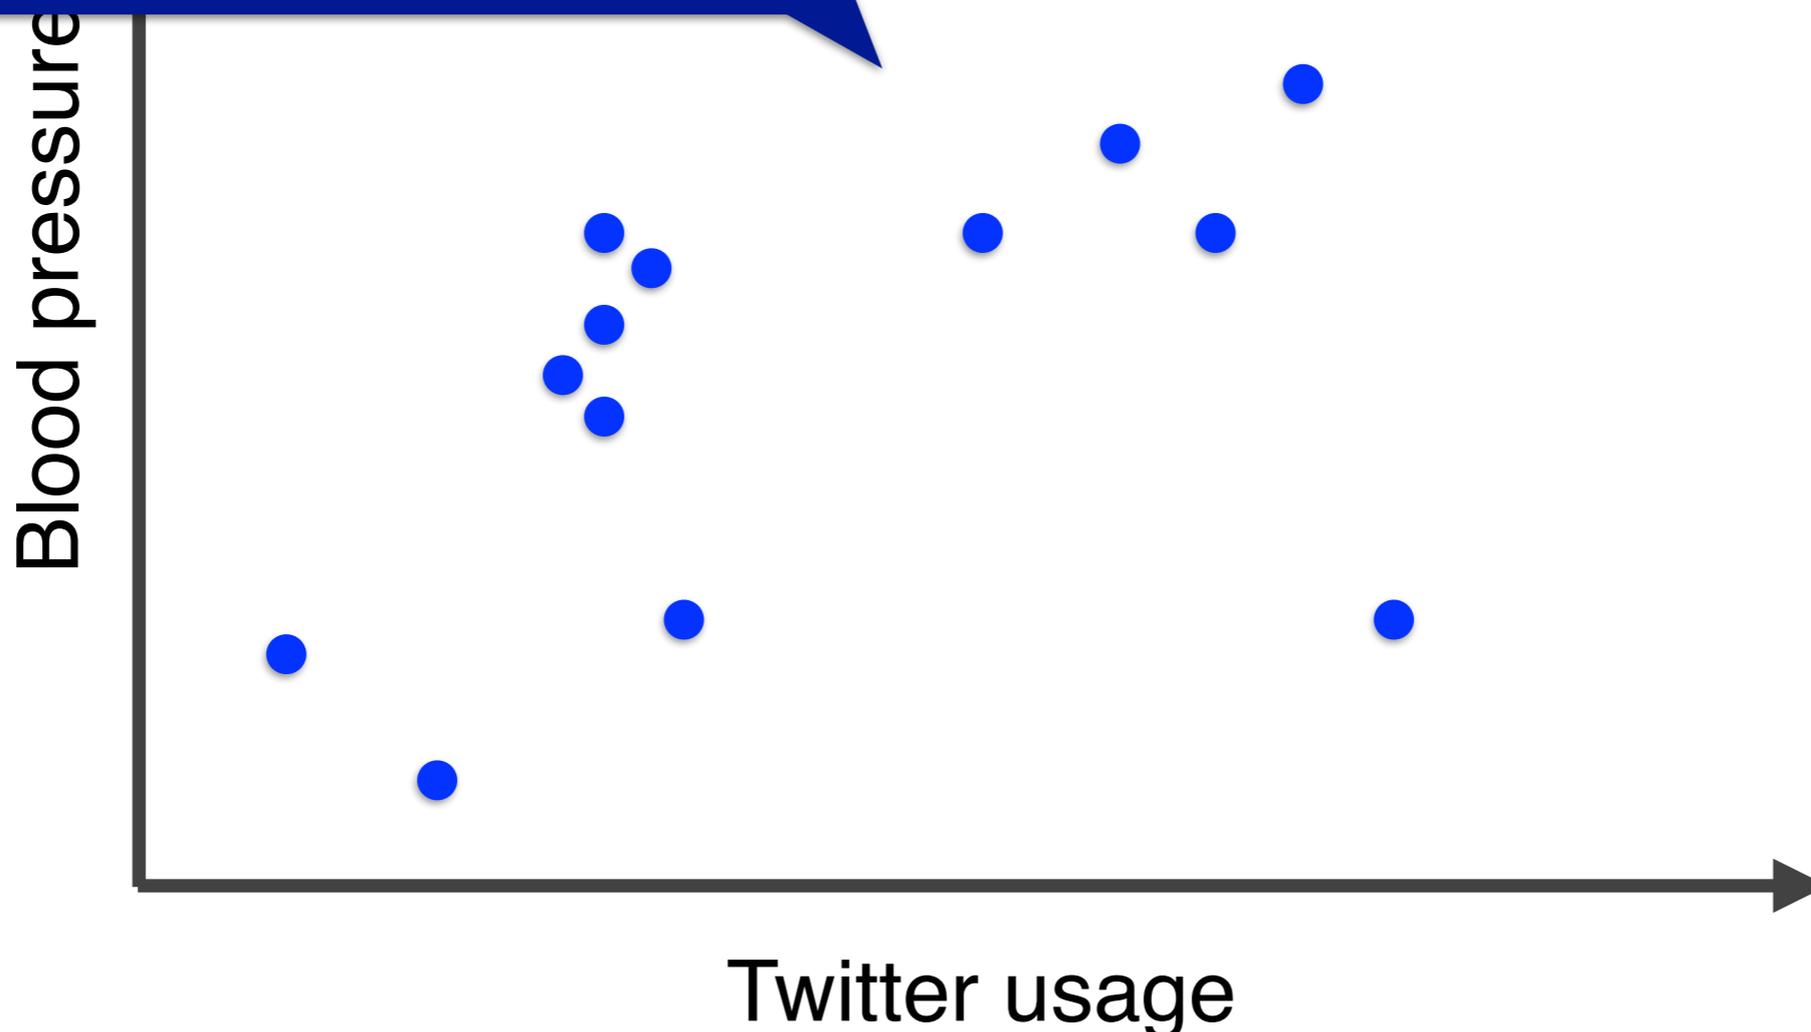
This paper: Private aggregation



This paper: Private aggregation



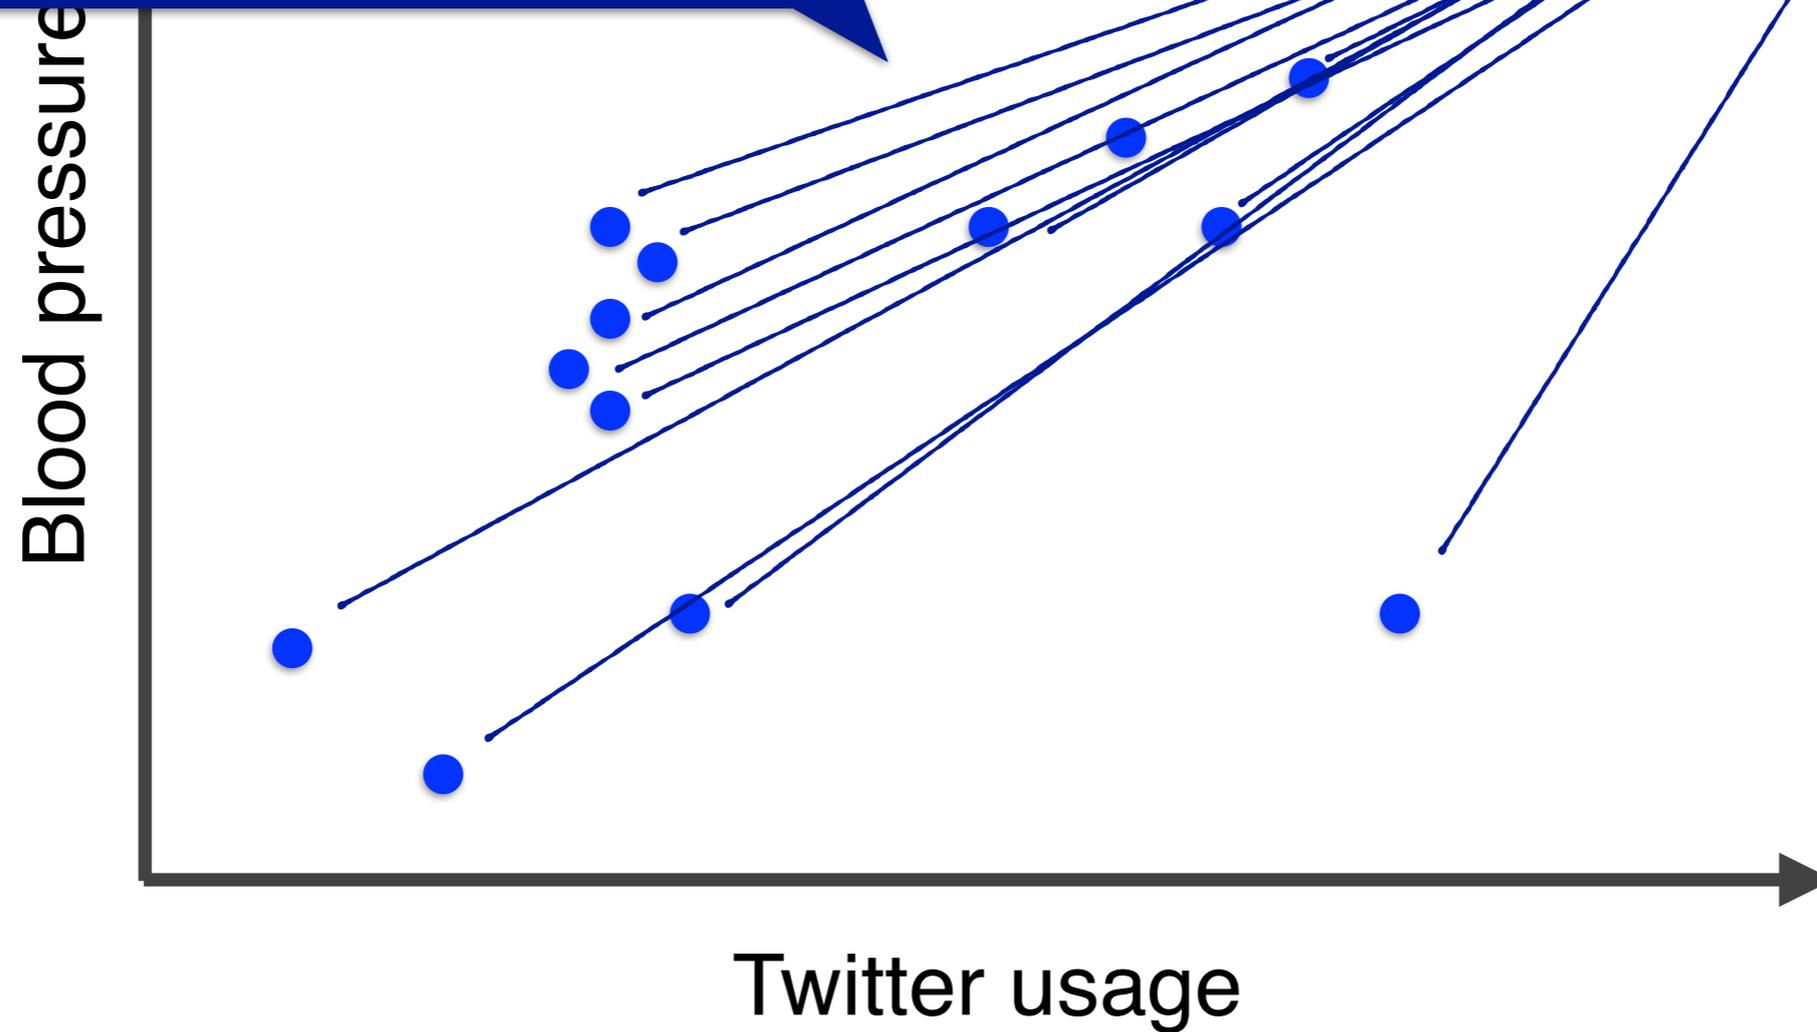
Clients send an *encrypted share* of their data to each aggregator



This paper: Private aggregation



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



This paper: Private aggregation

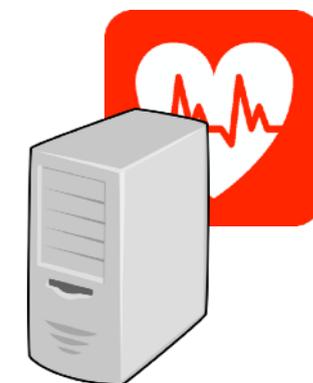
Clients send an *encrypted share* of their data to each aggregator

Blood pressure

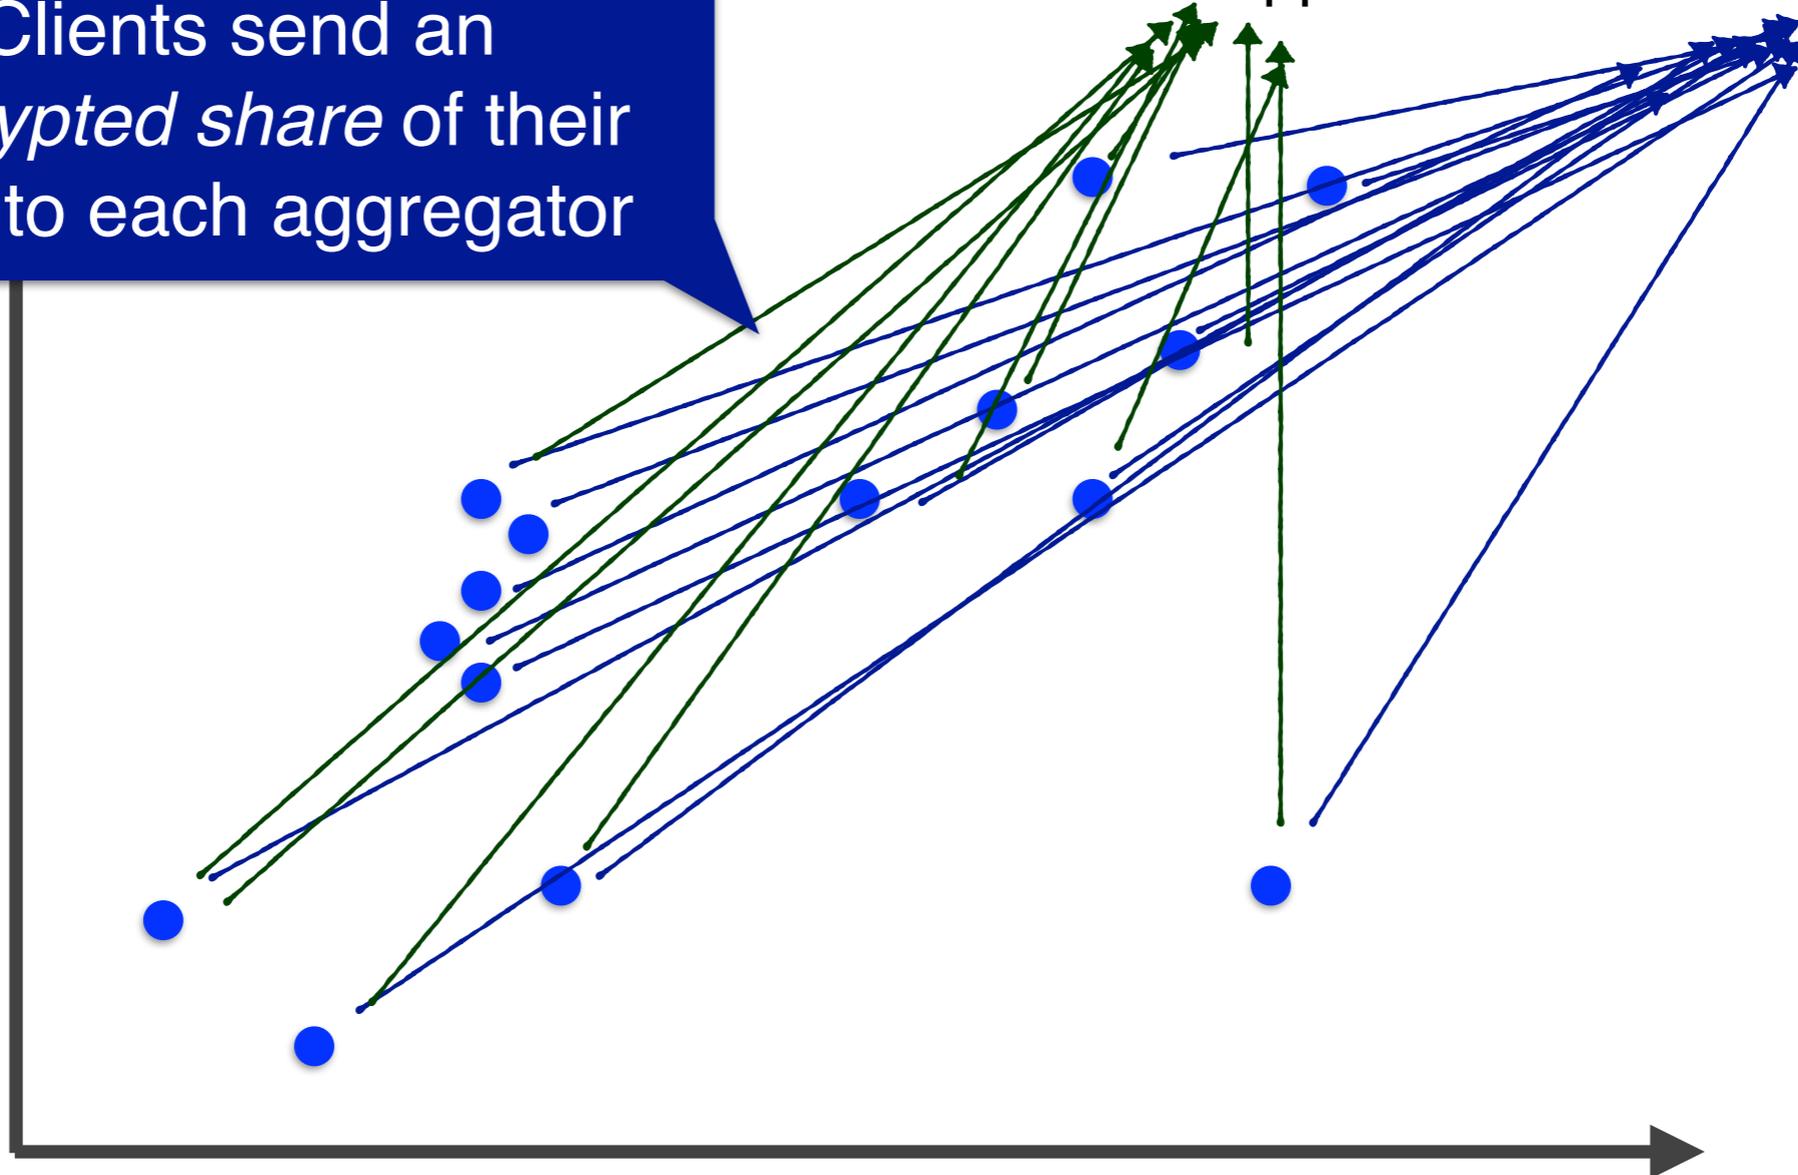
Twitter usage



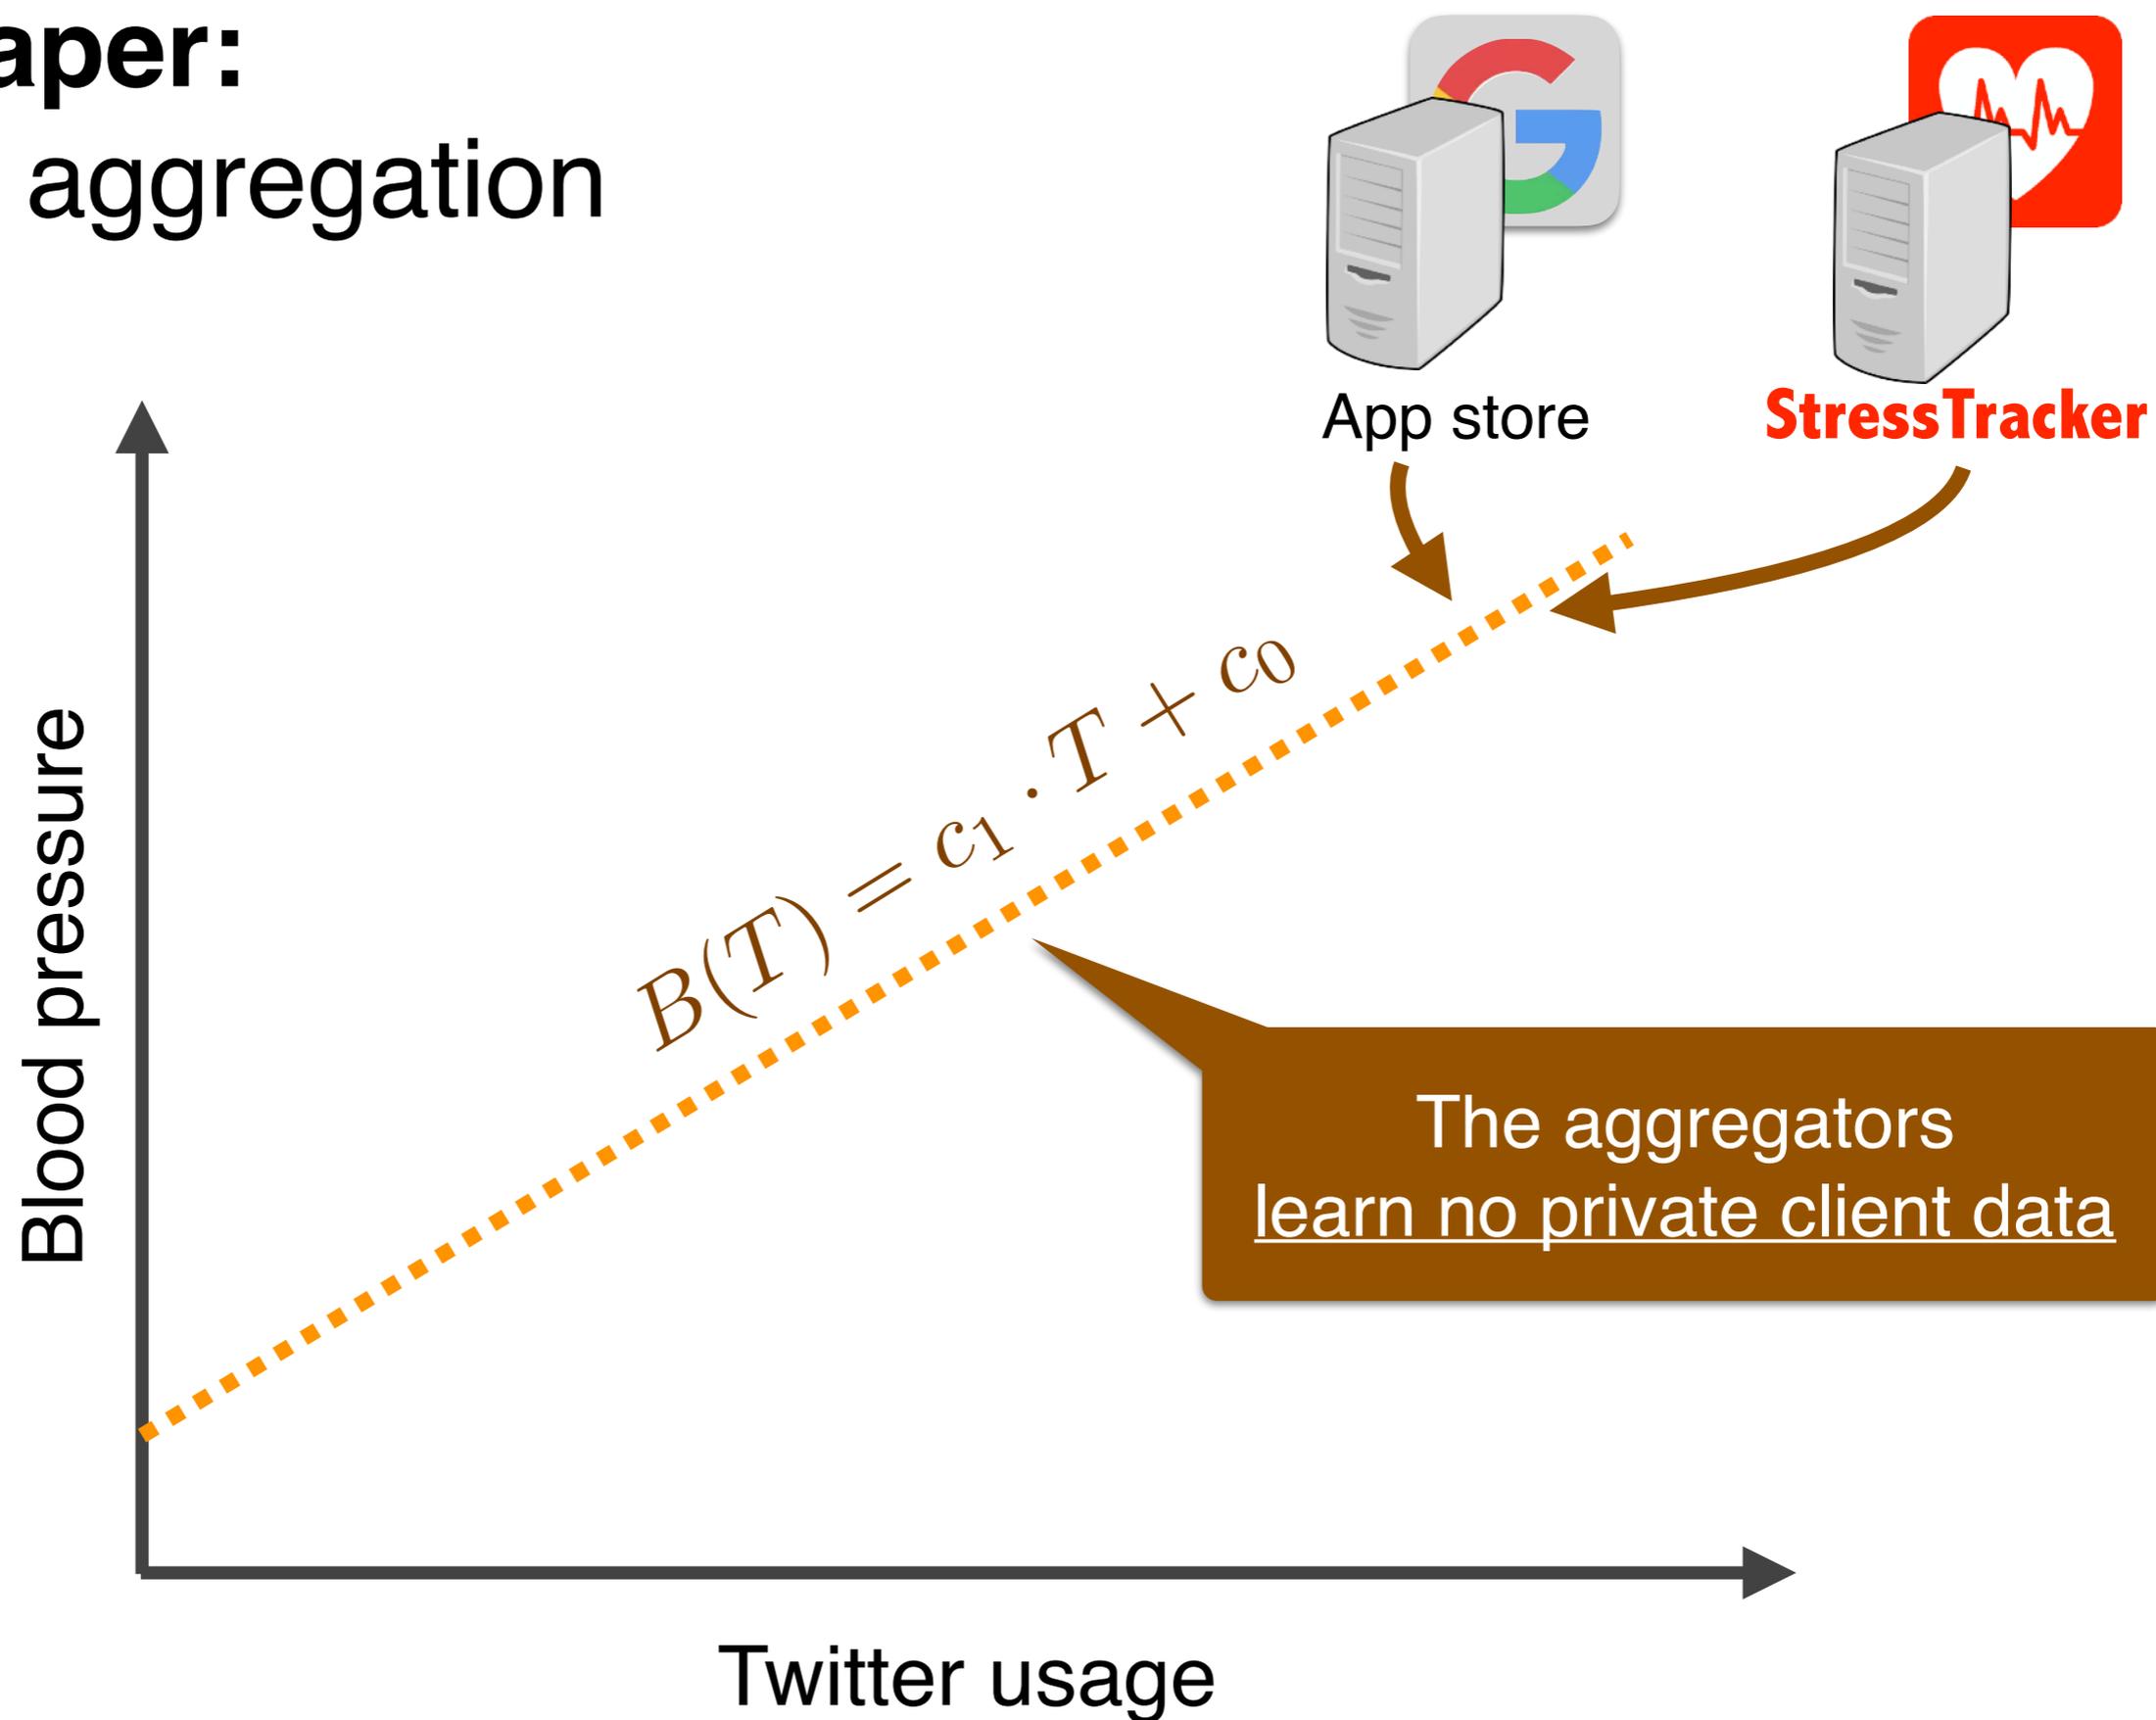
App store



StressTracker



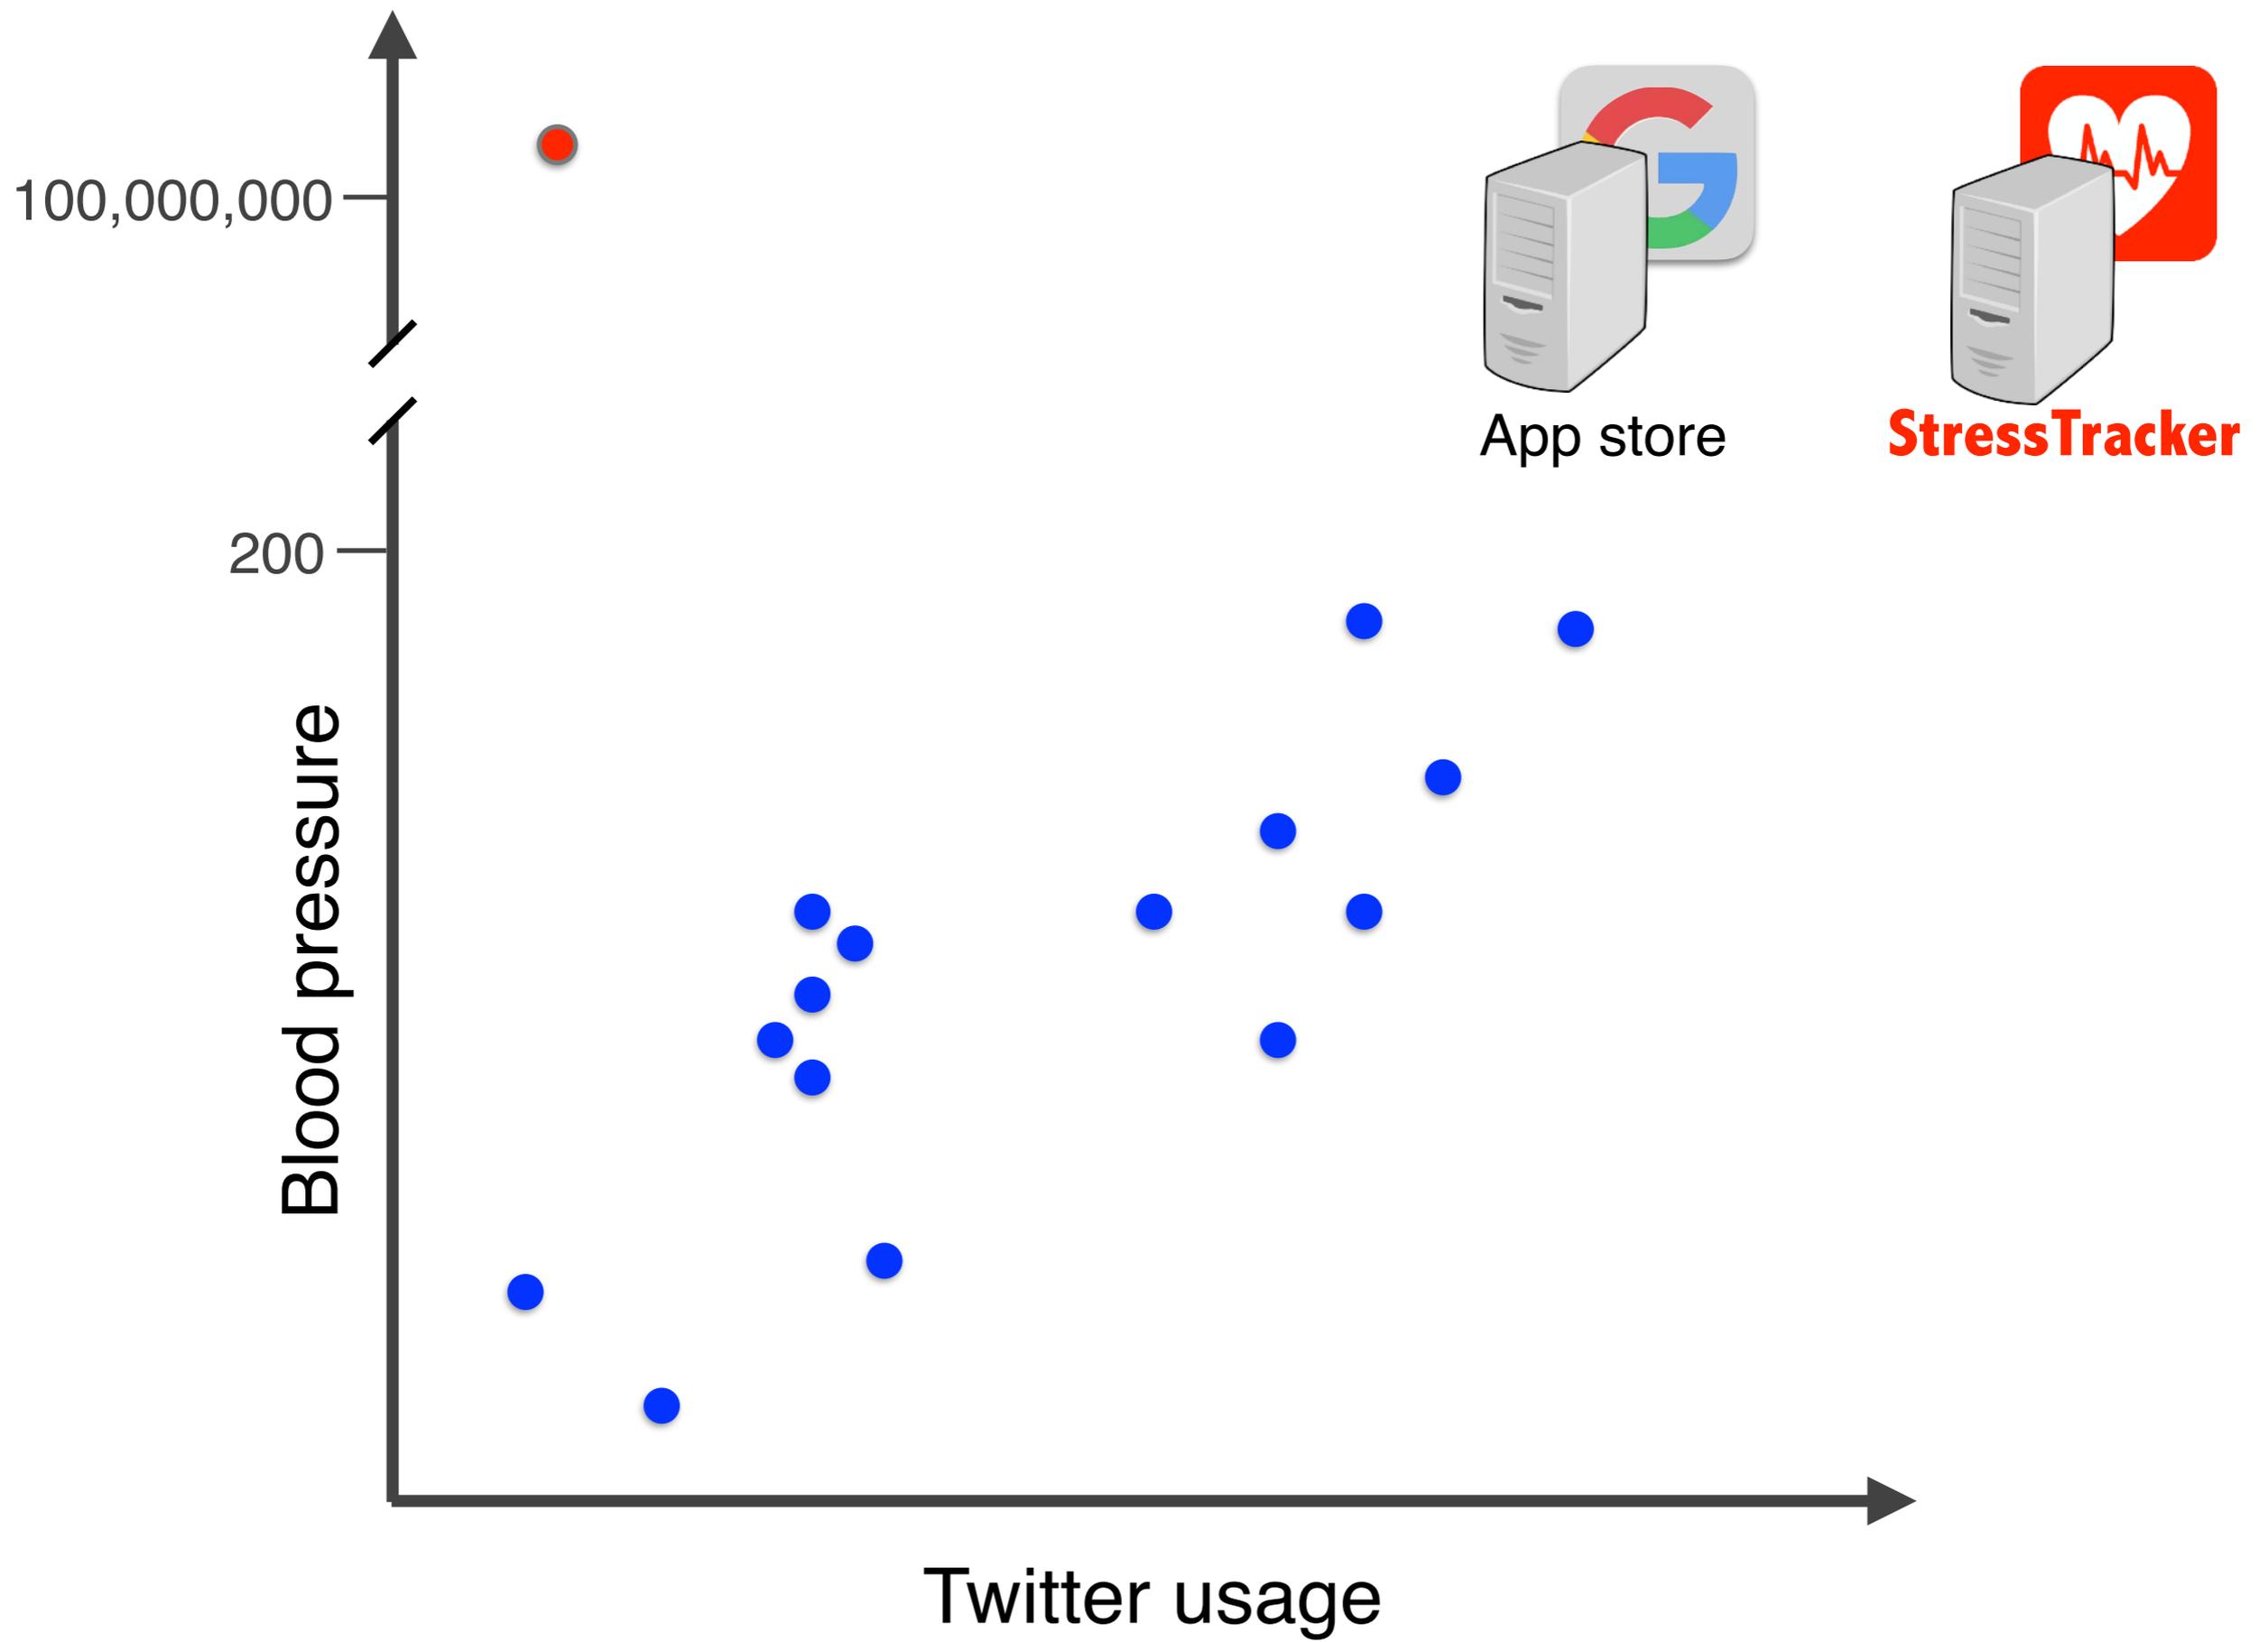
This paper: Private aggregation

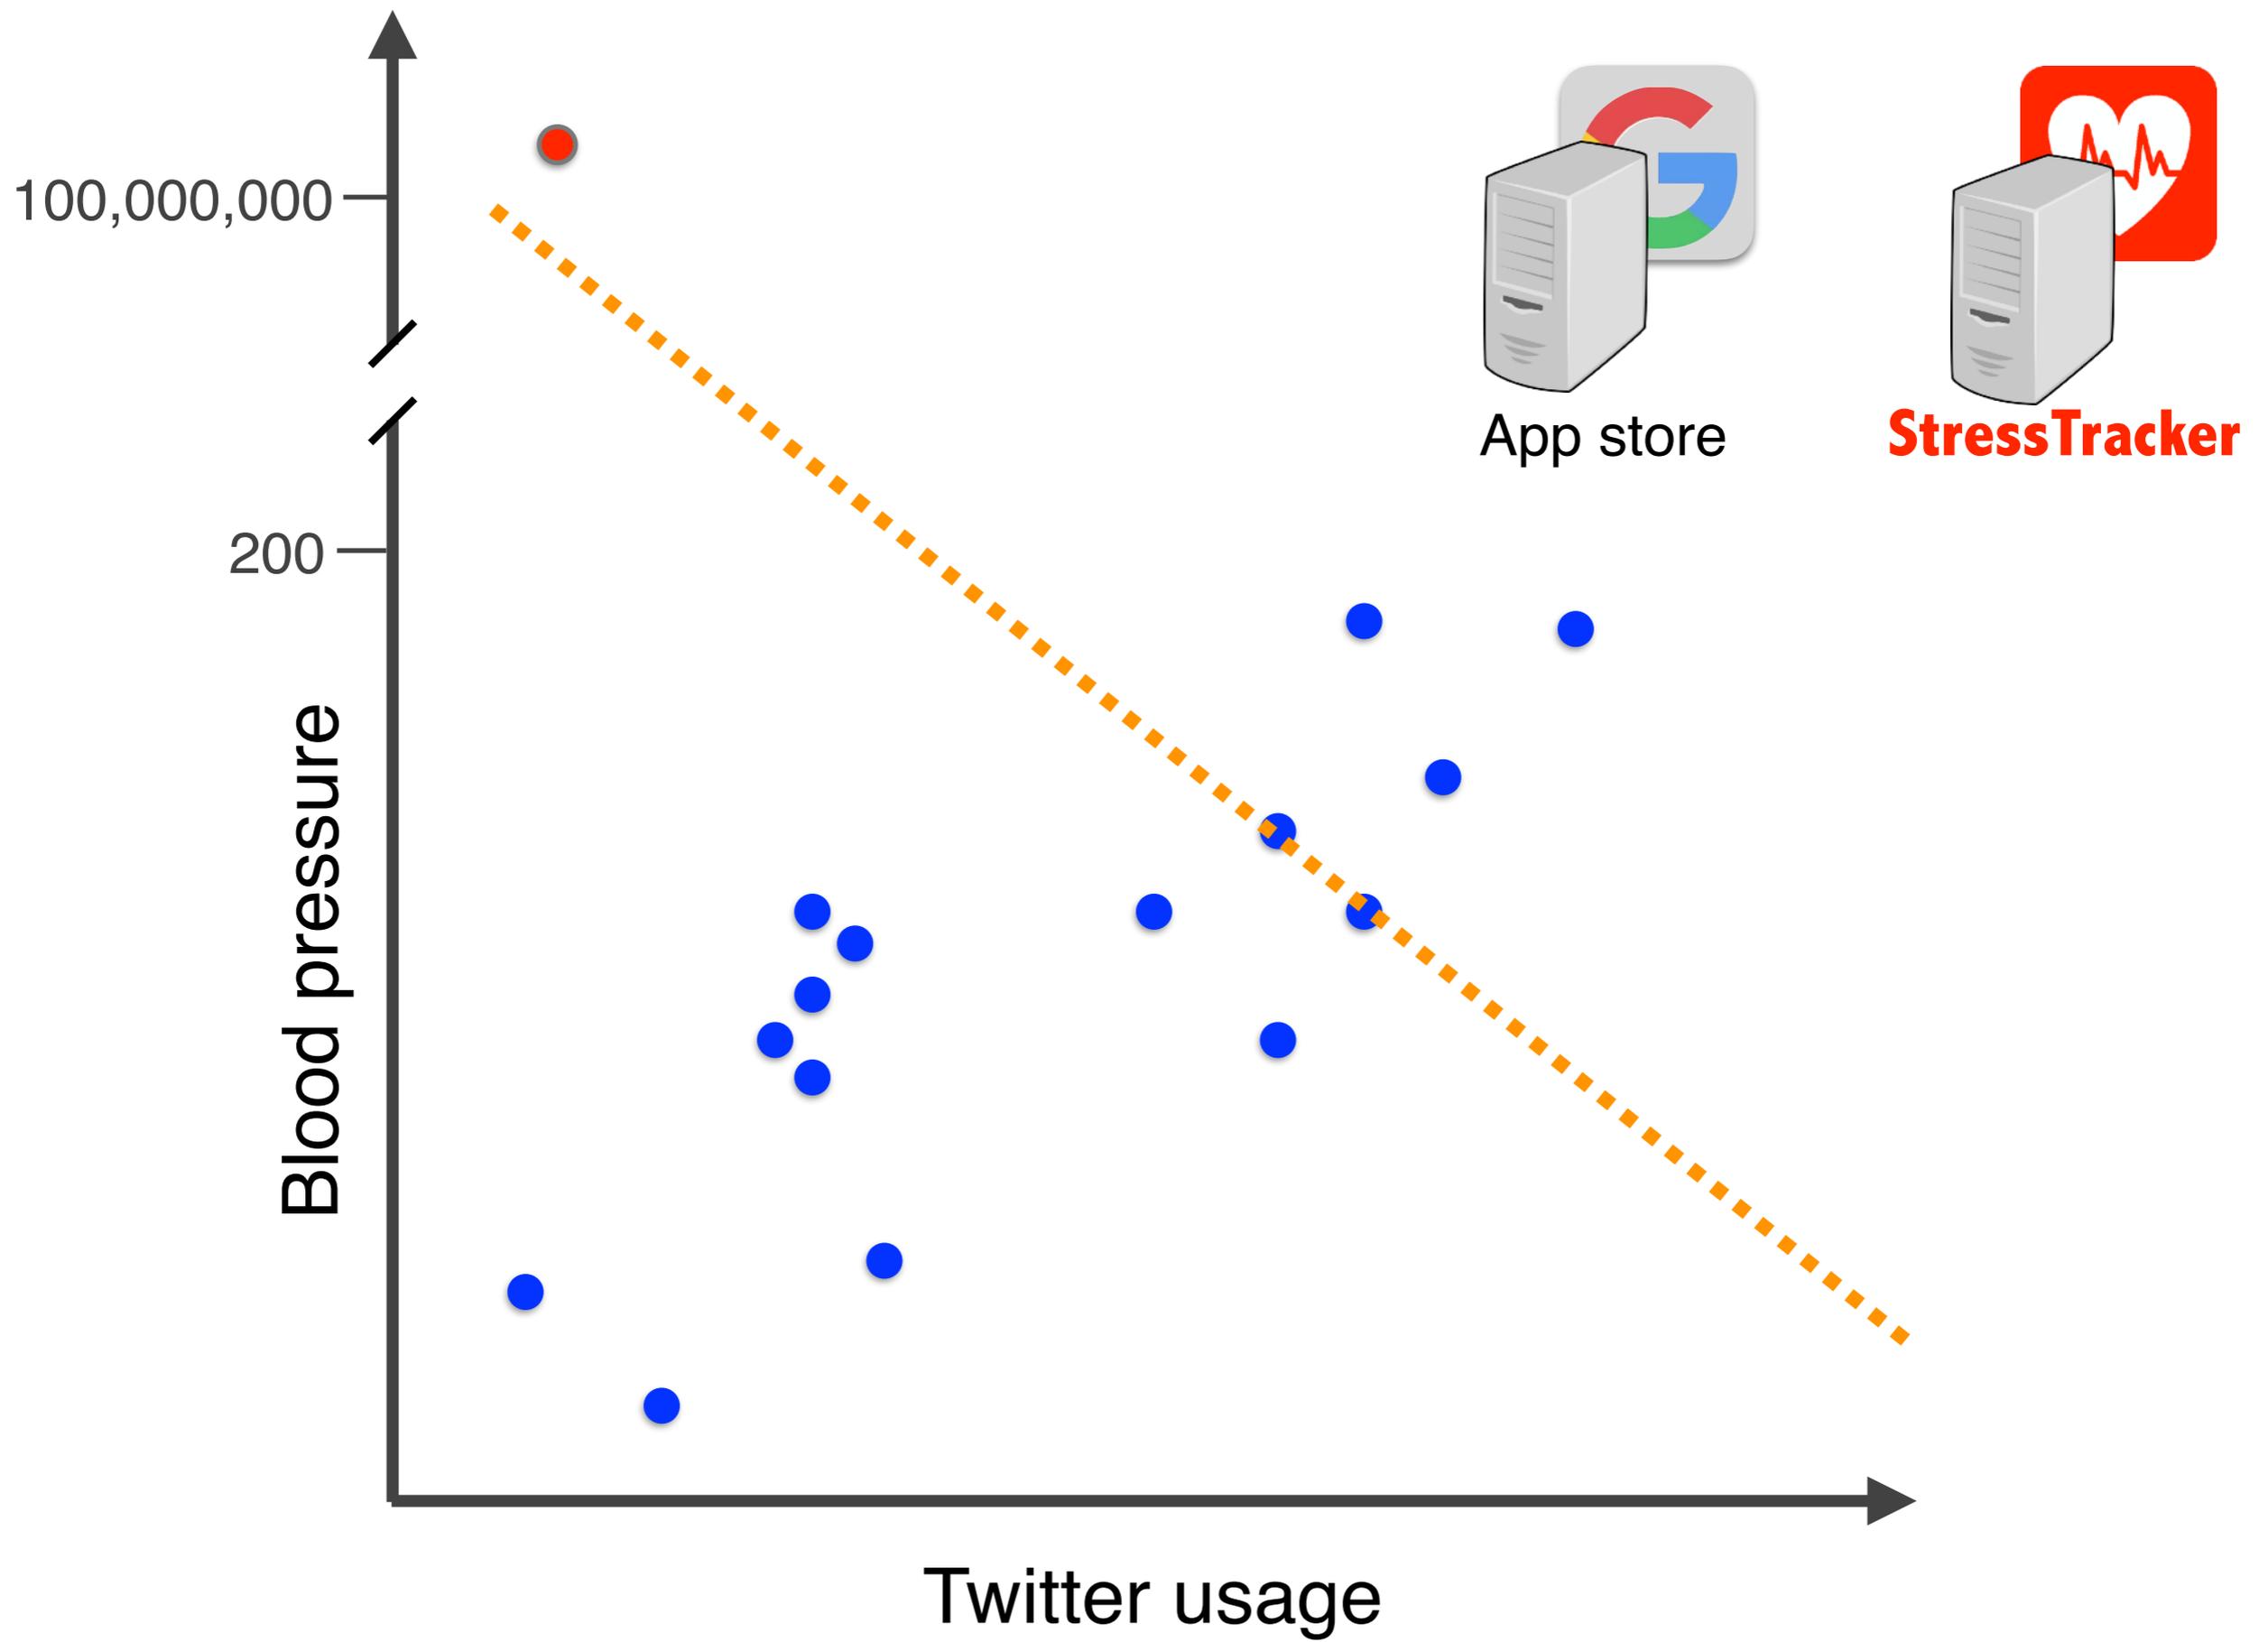


Private aggregation



-
- 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)
1000s of submissions per second

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 cryptography
1000s of submissions

...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)
1000s of submissions per second

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

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
- Encodings for complex aggregates

Outline

- Background: The private aggregation problem
- **A straw-man solution for private sums**
- Providing robustness with SNIPs
- Evaluation
- Encodings for complex aggregates

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] ...

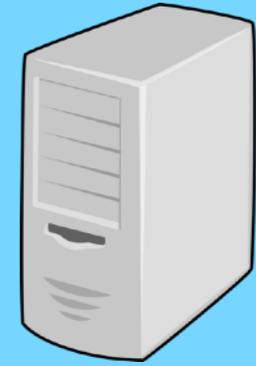
Server A



Server B

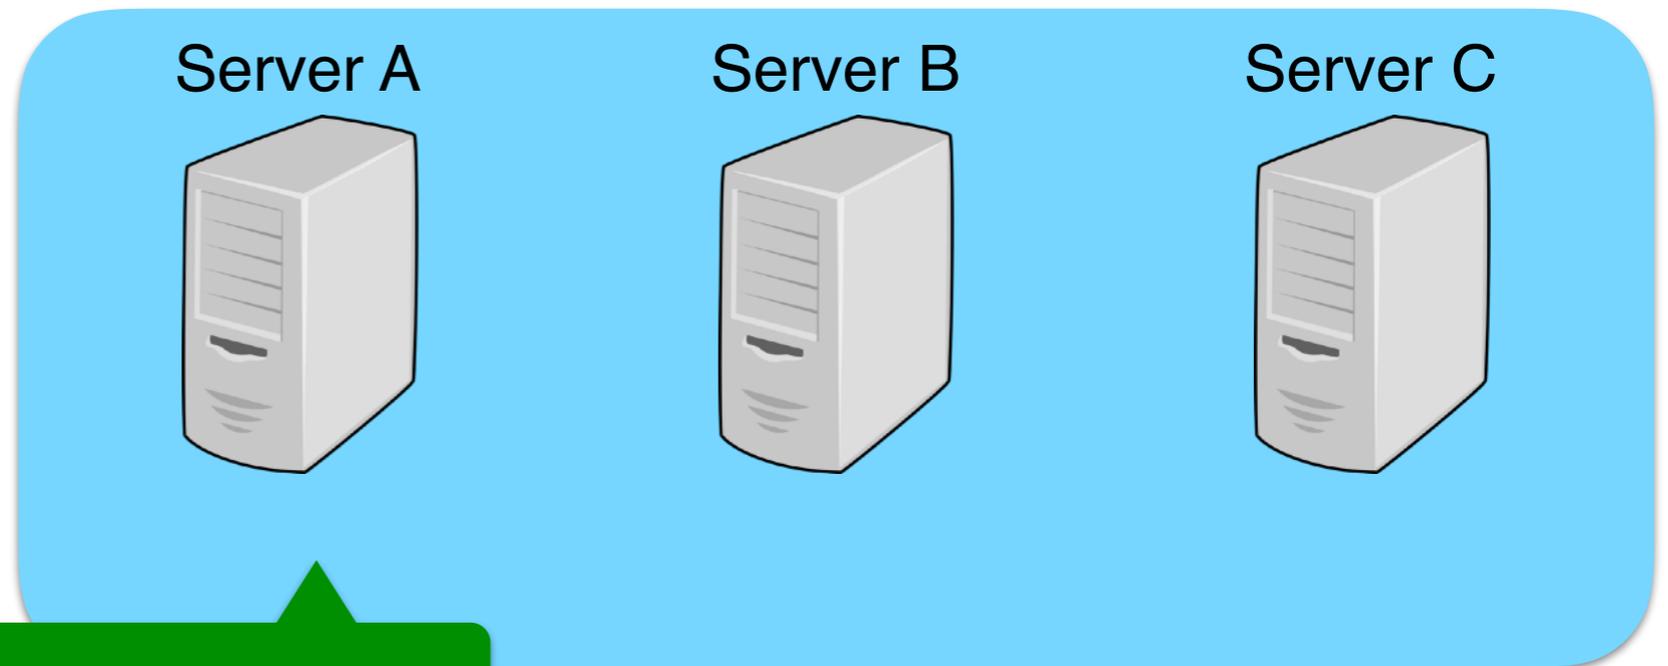


Server C



Private sums: A “straw-man” scheme

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

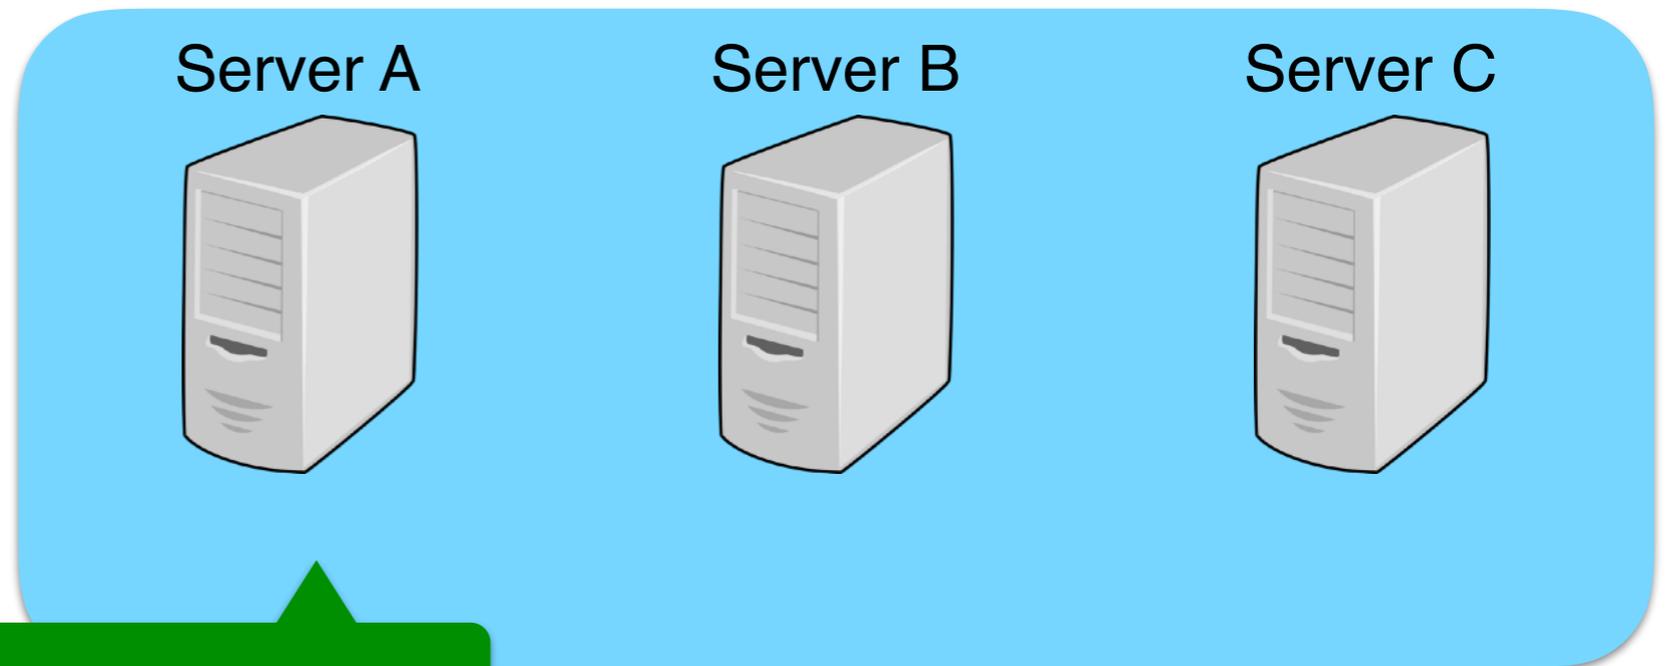


Assume that the servers are
non-colluding.

Equivalently: that at least one
server is honest.

Private sums: A “straw-man” scheme

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



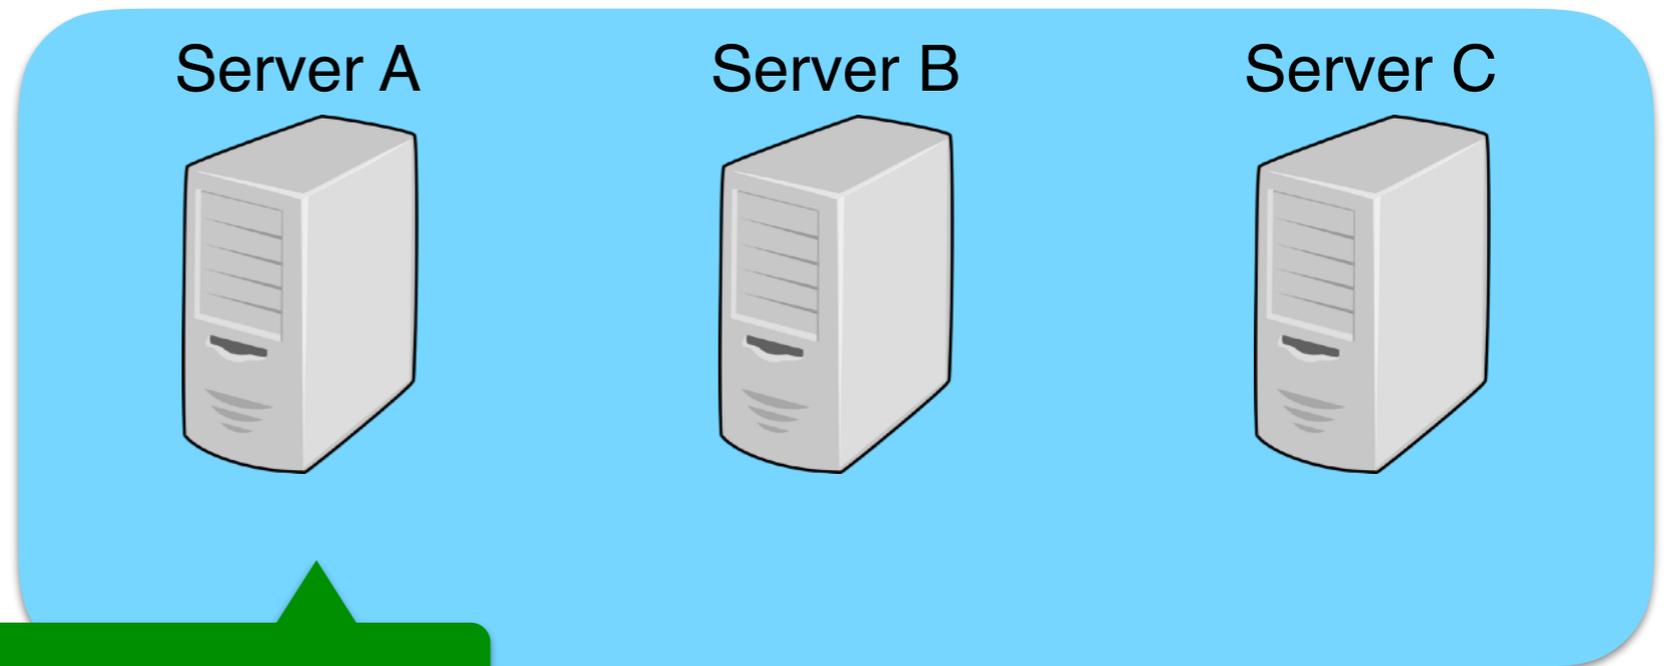
Assume that the servers are
non-colluding.

Equivalently: that at least one
server is honest.



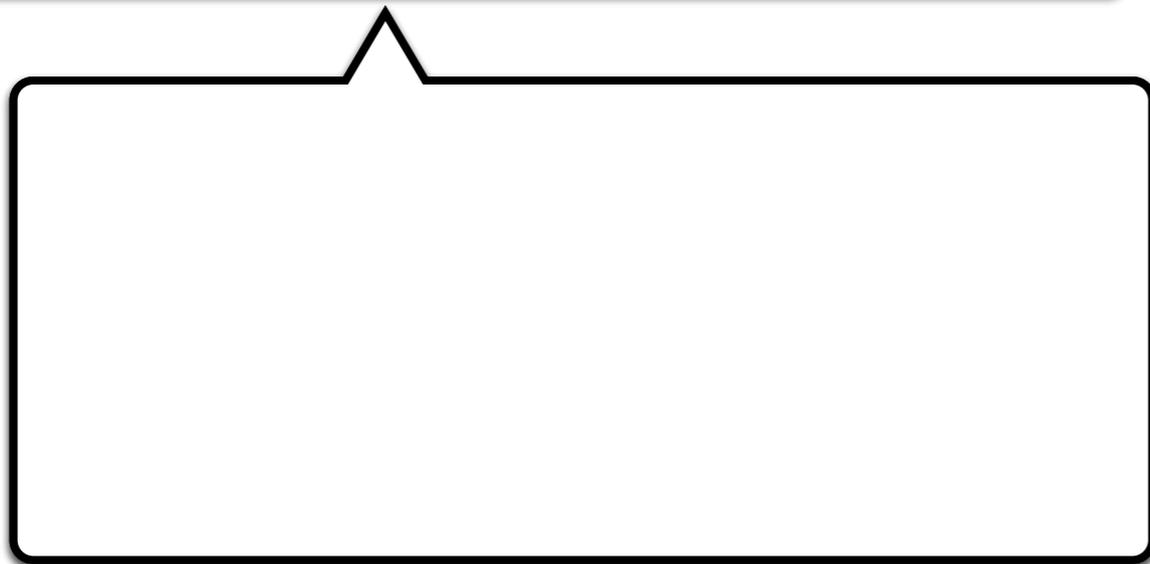
Private sums: A “straw-man” scheme

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



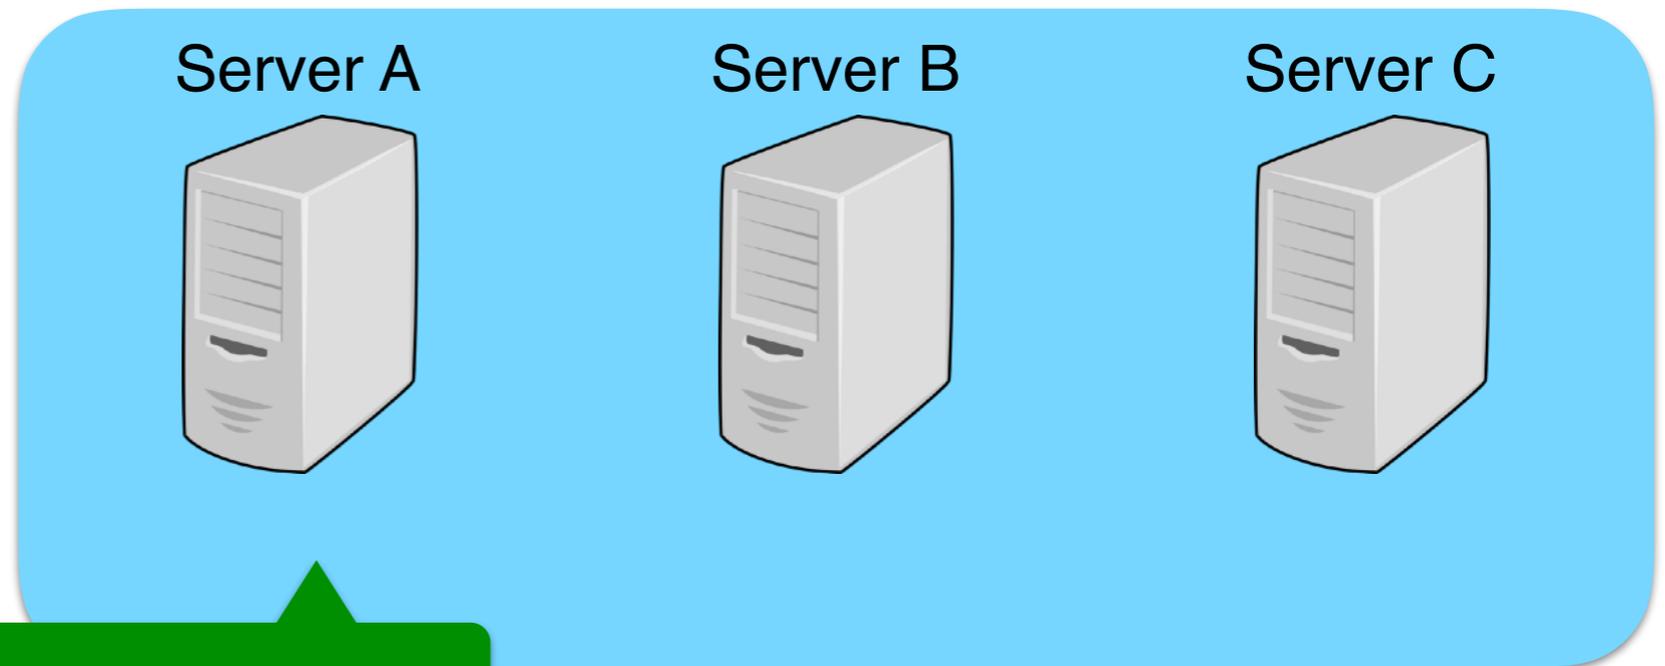
Assume that the servers are
non-colluding.

Equivalently: that at least one
server is honest.

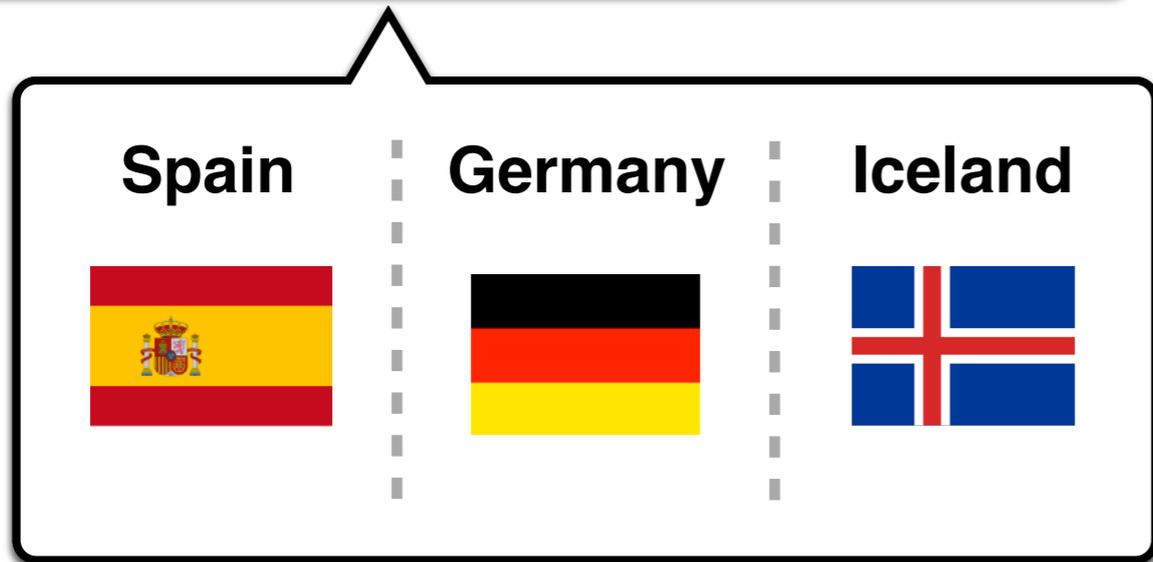


Private sums: A “straw-man” scheme

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



Assume that the servers are
non-colluding.
Equivalently: that at least one
server is honest.



Private sums: A “straw-man” scheme

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

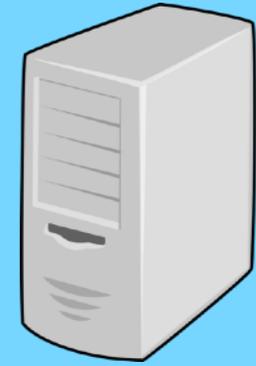
Server A



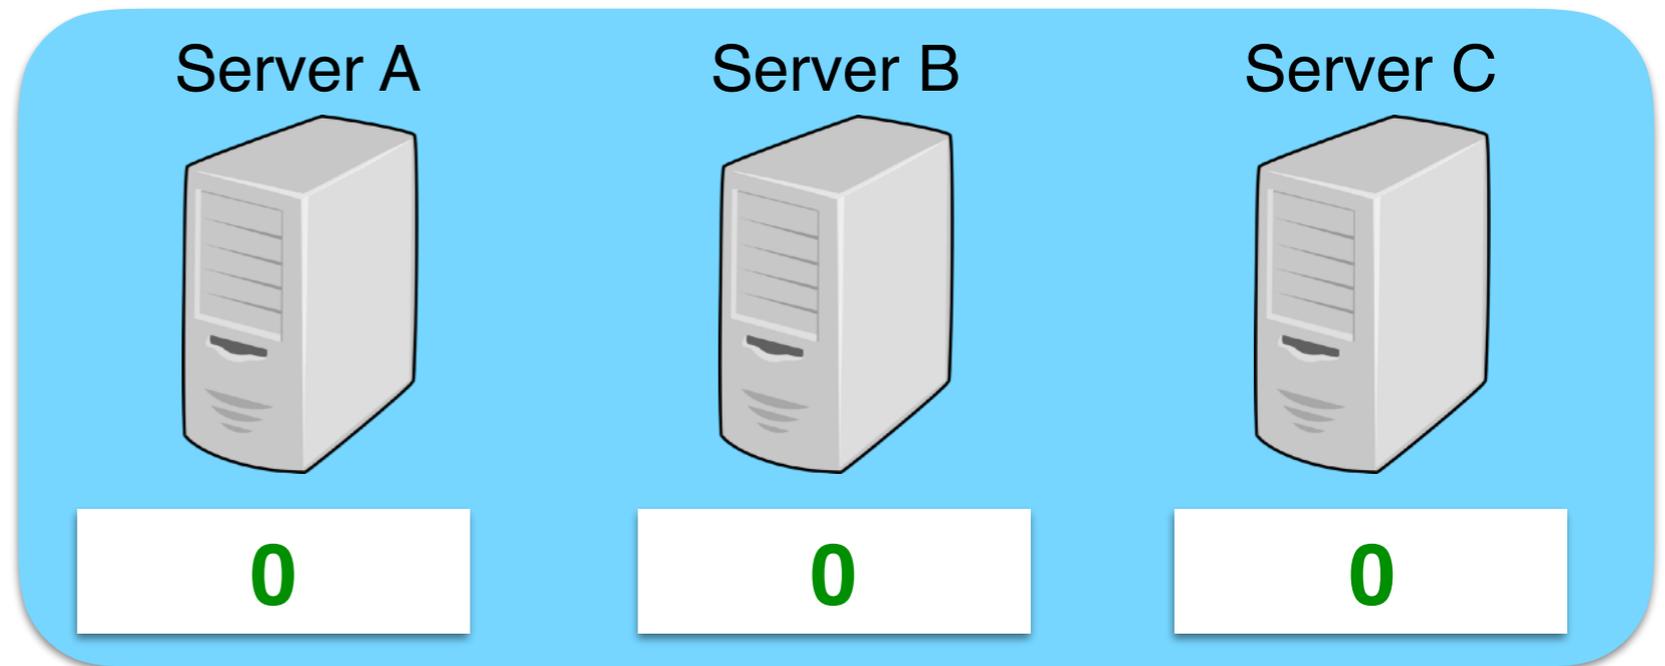
Server B



Server C



Private sums: A “straw-man” scheme



1



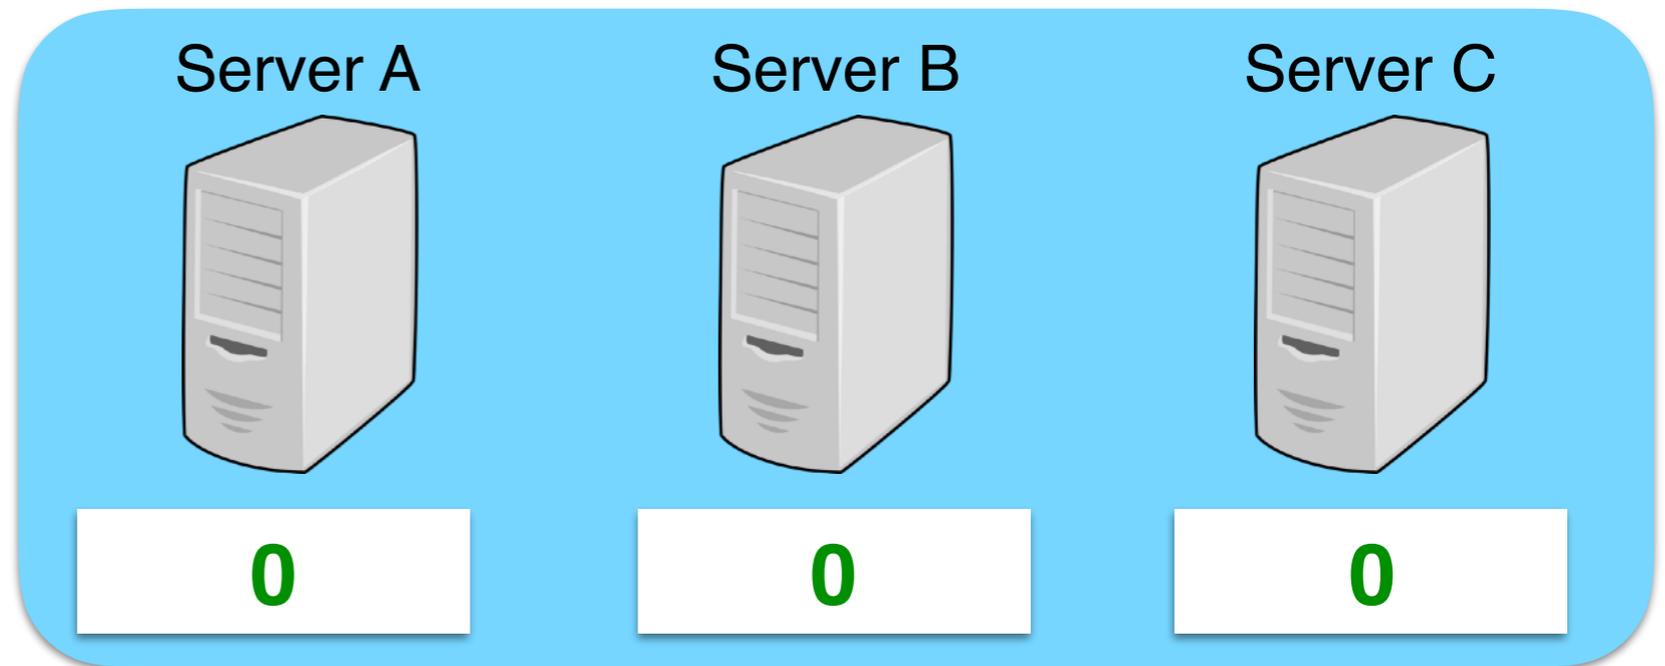
Secret sharing

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

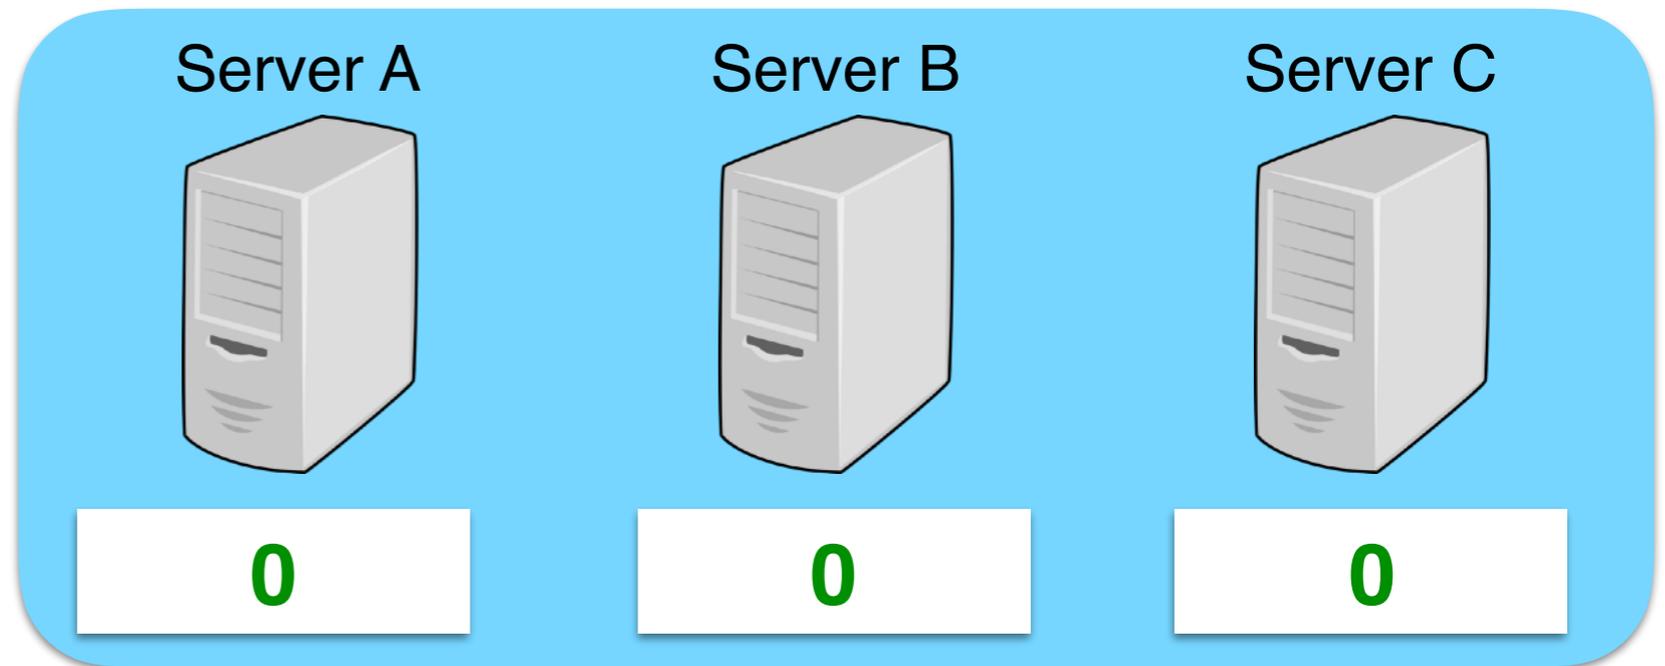
Pick three random “shares” that sum to 1.

$$1 = 15 + (-12) + (-2) \pmod{31}$$

In real system, we
use a “big” prime

Need all three shares to recover the shared value.

Private sums: A “straw-man” scheme



1



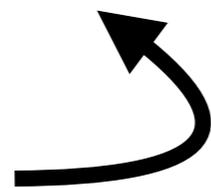
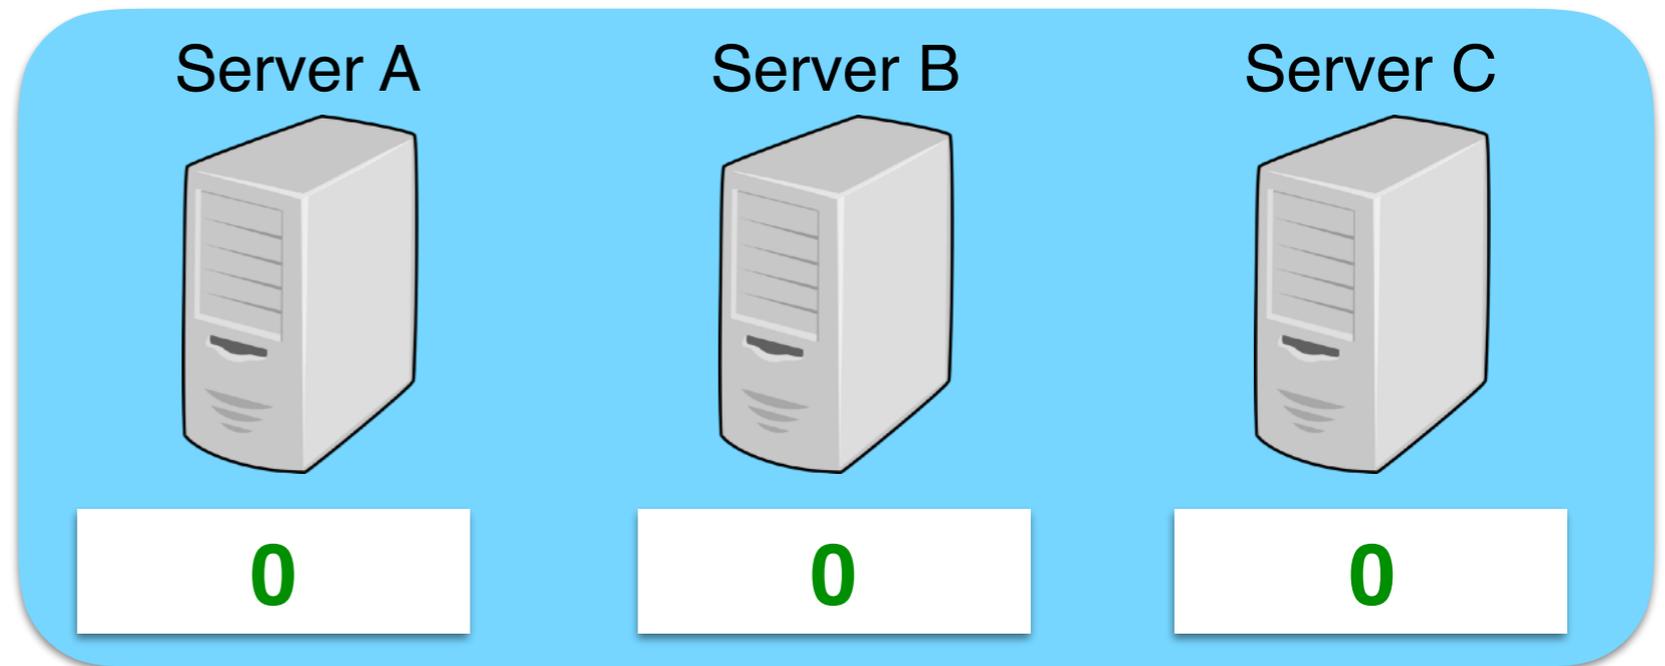
Secret sharing

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



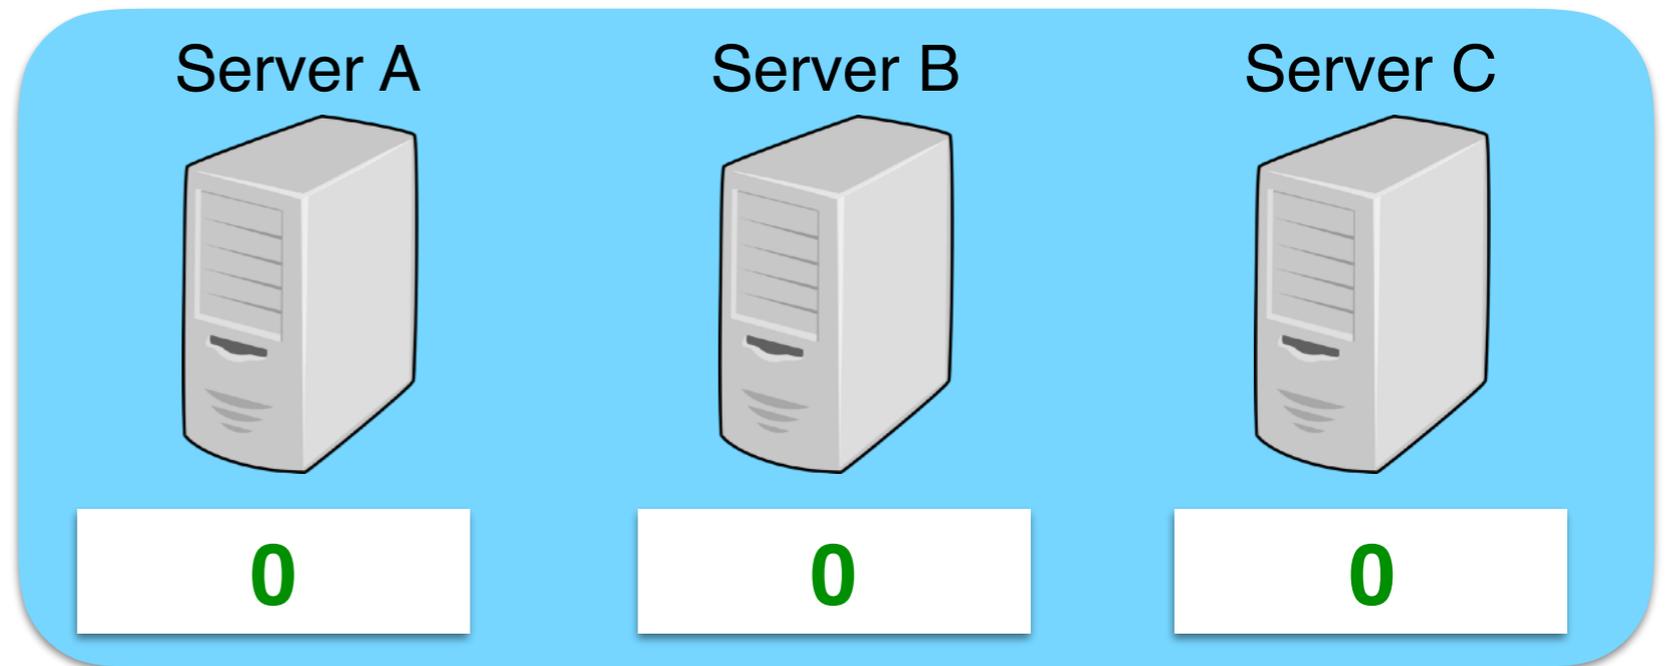
Secret sharing

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



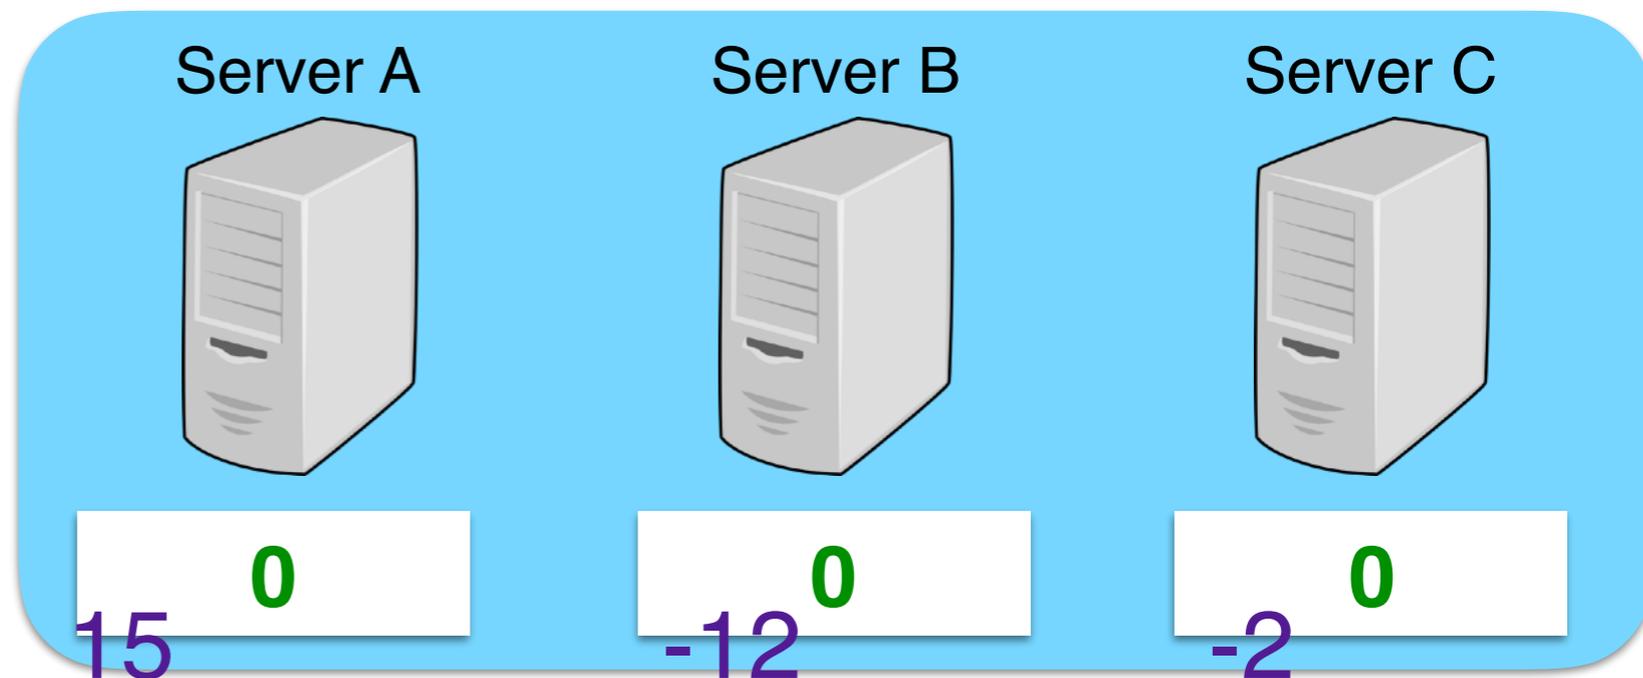
Secret sharing

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.

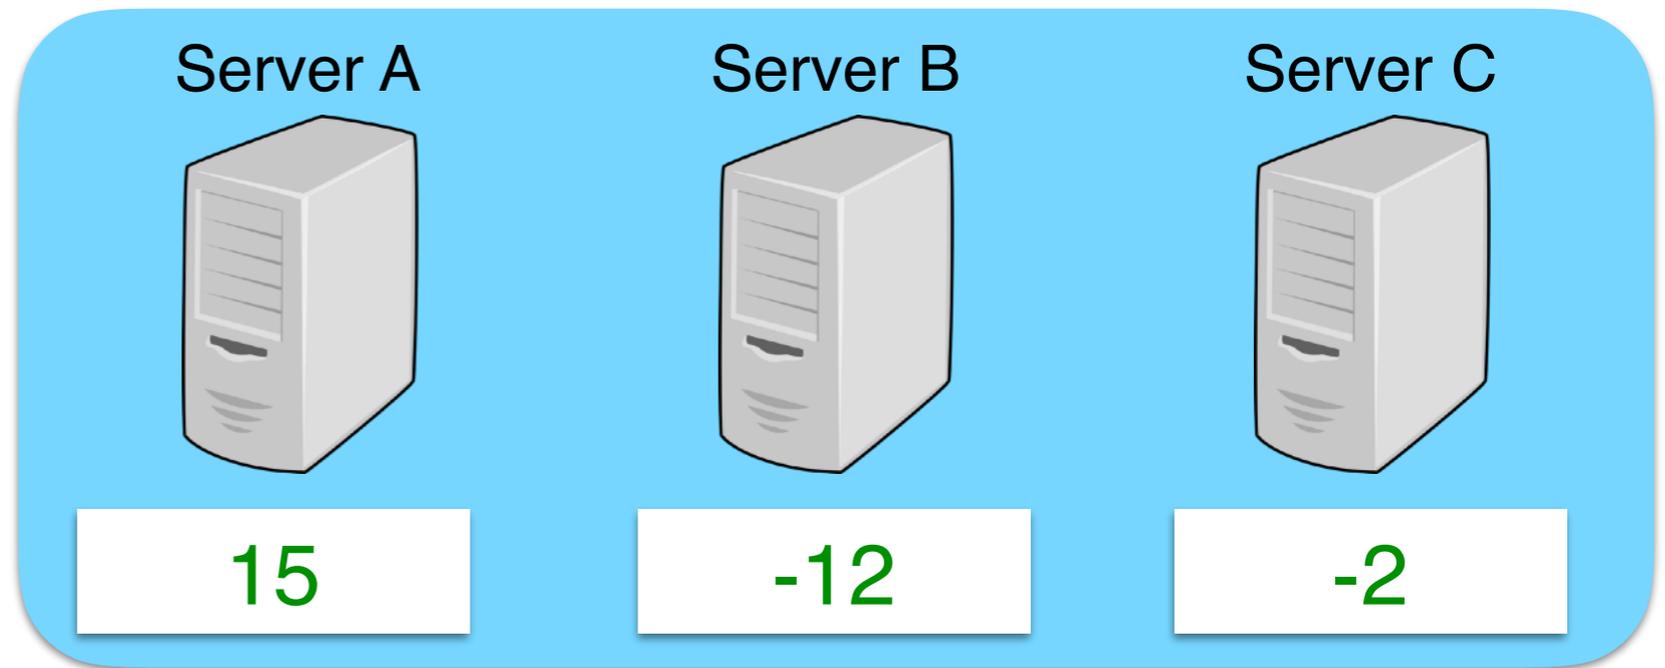
Private sums: A “straw-man” scheme



1



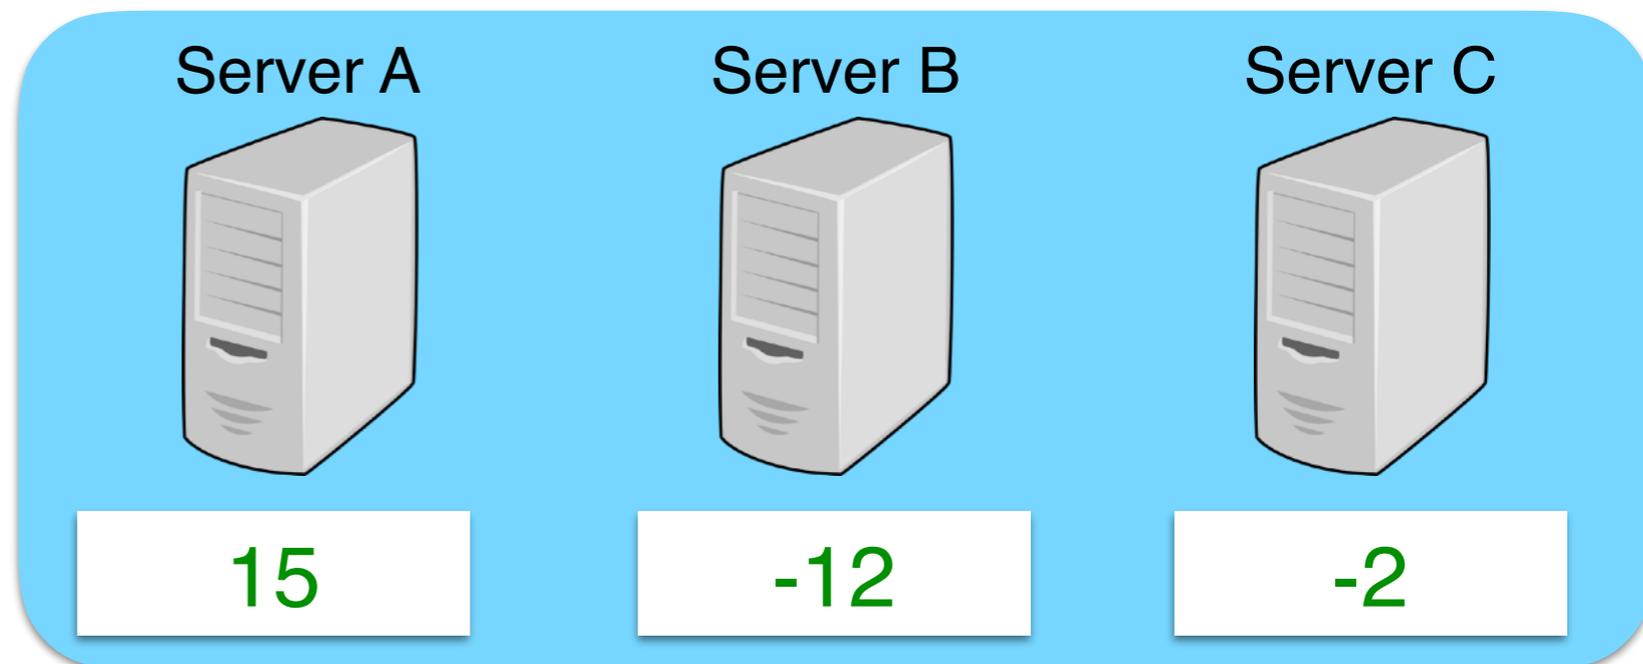
Private sums: A “straw-man” scheme



1



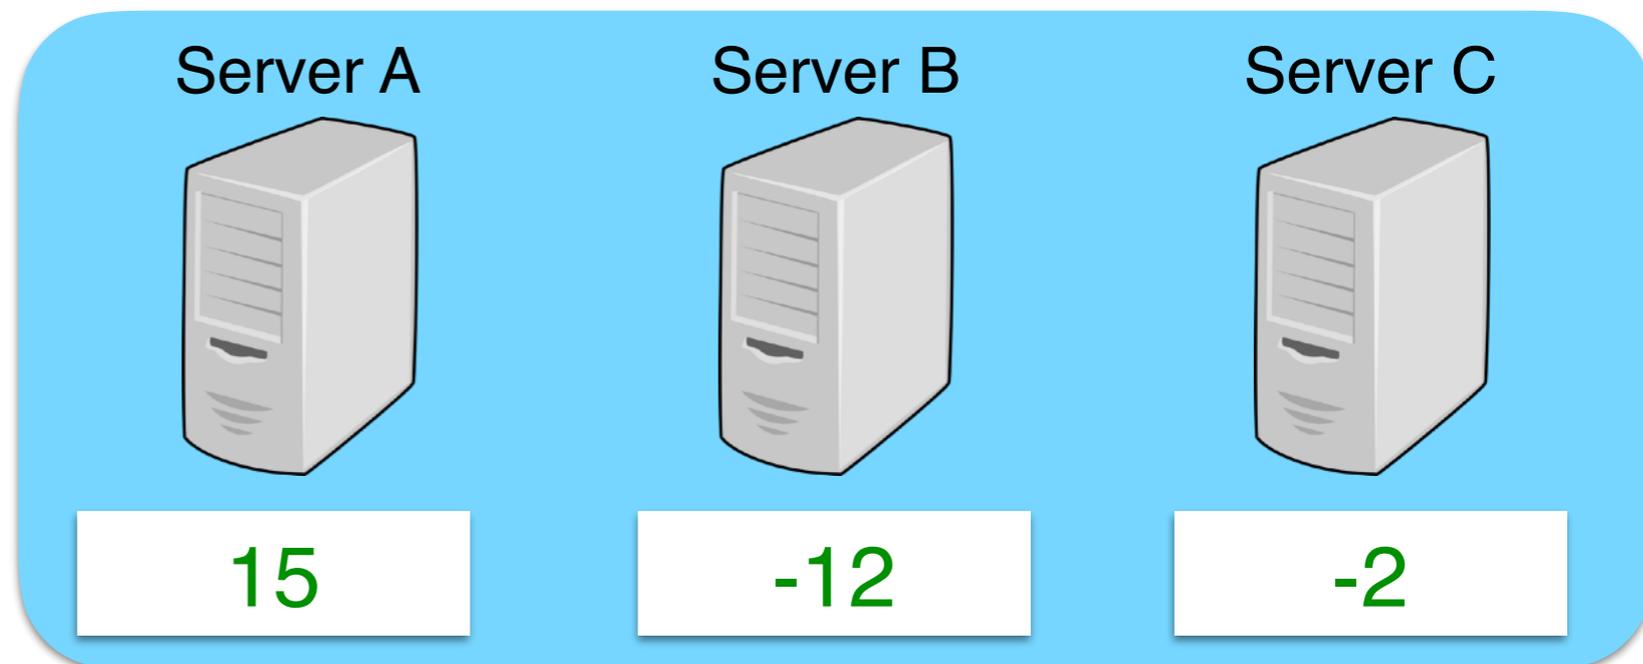
Private sums: A “straw-man” scheme



0



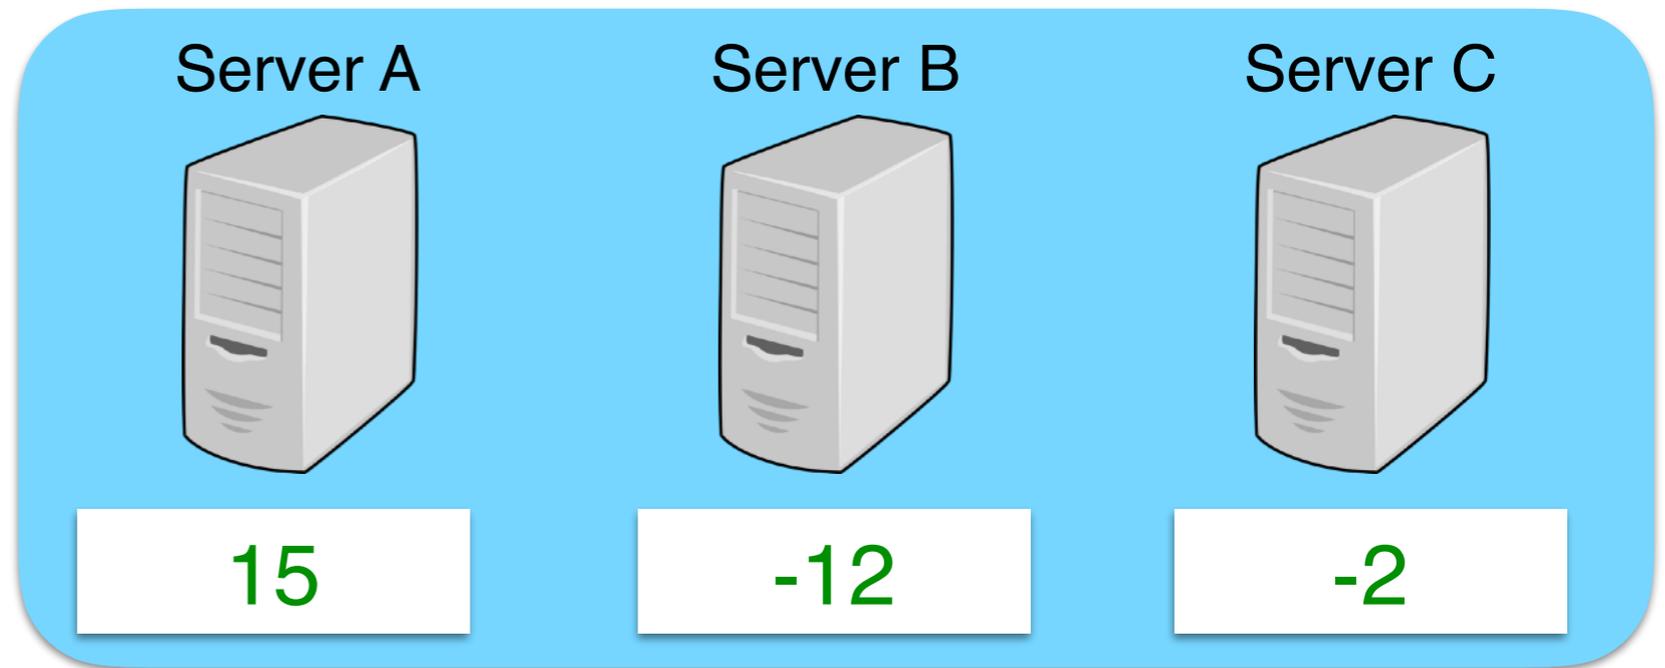
Private sums: A “straw-man” scheme



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



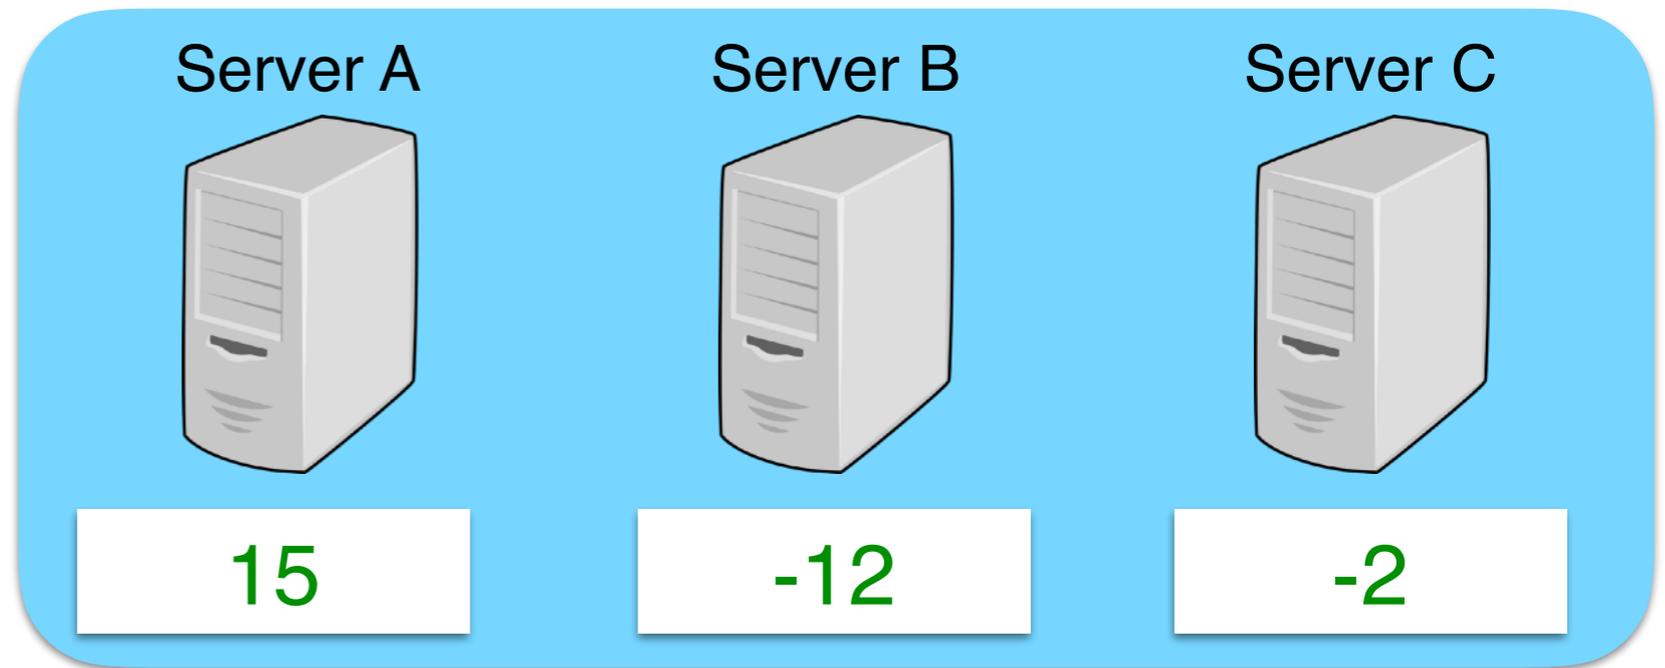
Private sums: A “straw-man” scheme



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



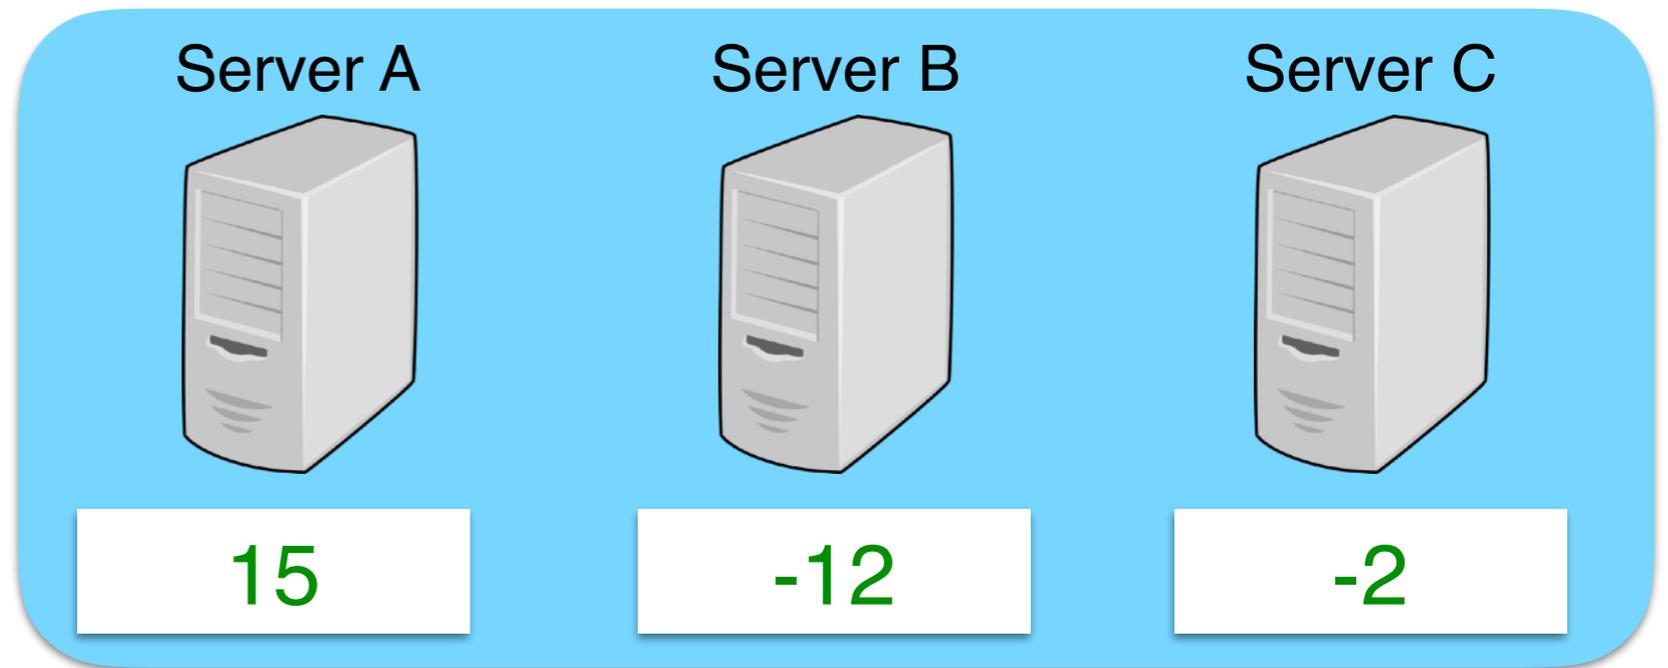
Private sums: A “straw-man” scheme



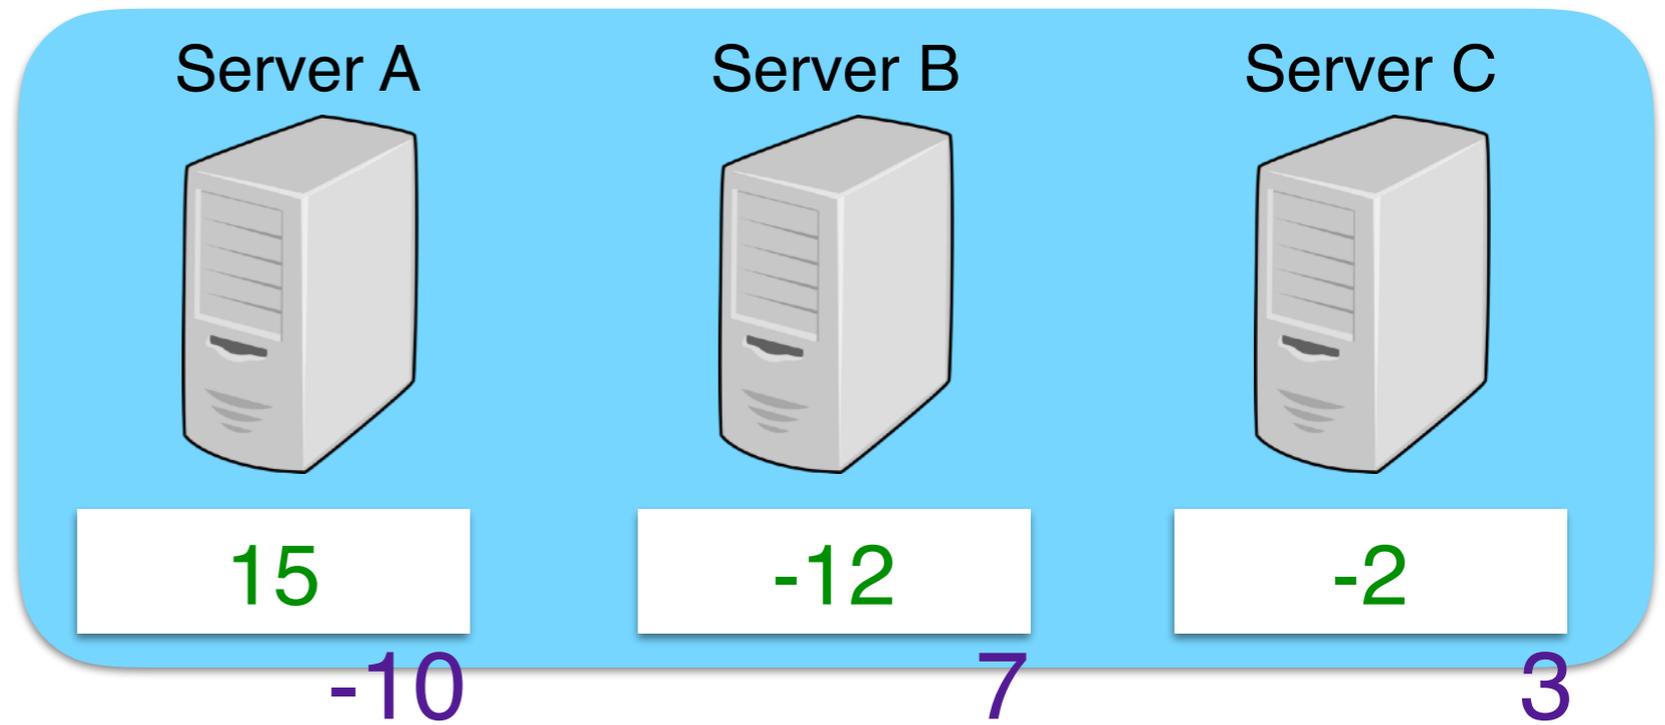
-10 7 3
↖ 0 = (-10) + 7 + 3



Private sums: A “straw-man” scheme



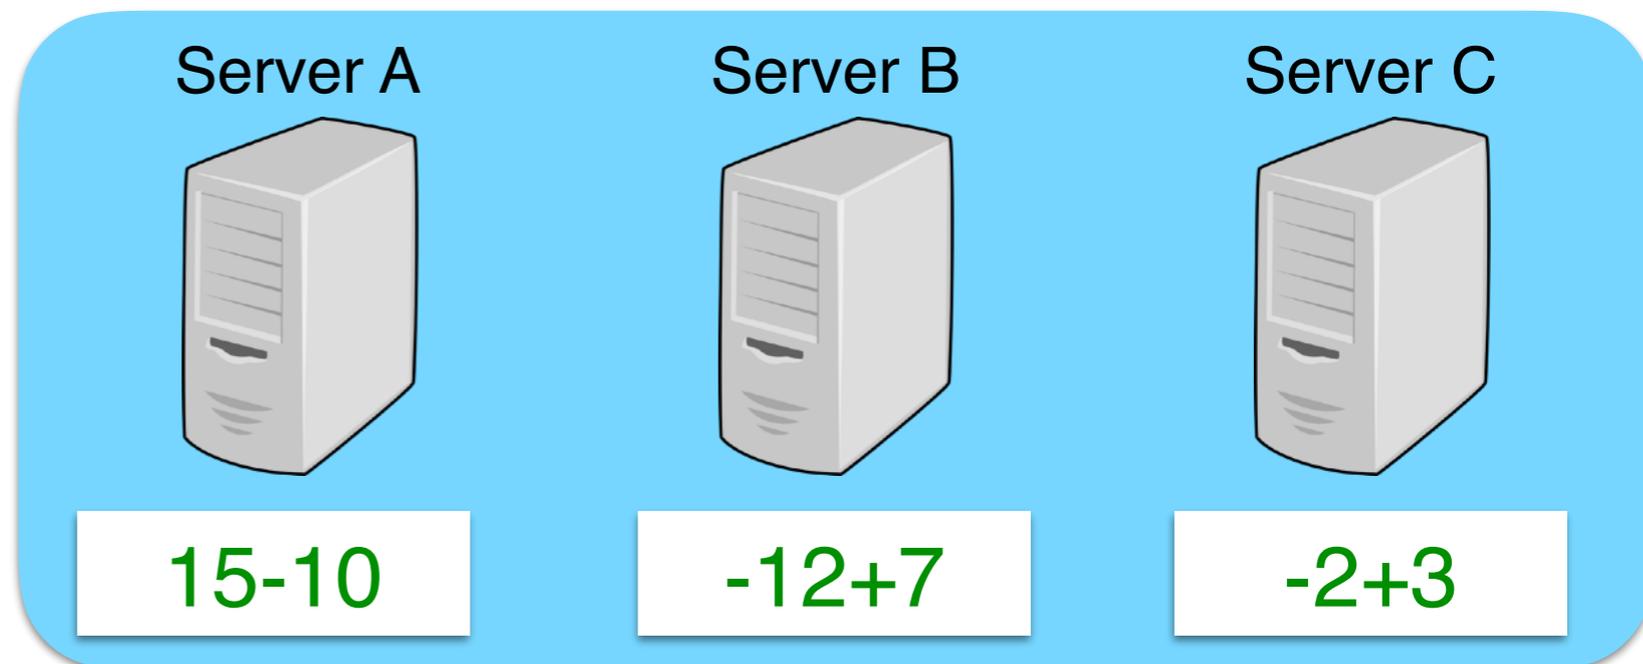
Private sums: A “straw-man” scheme



0



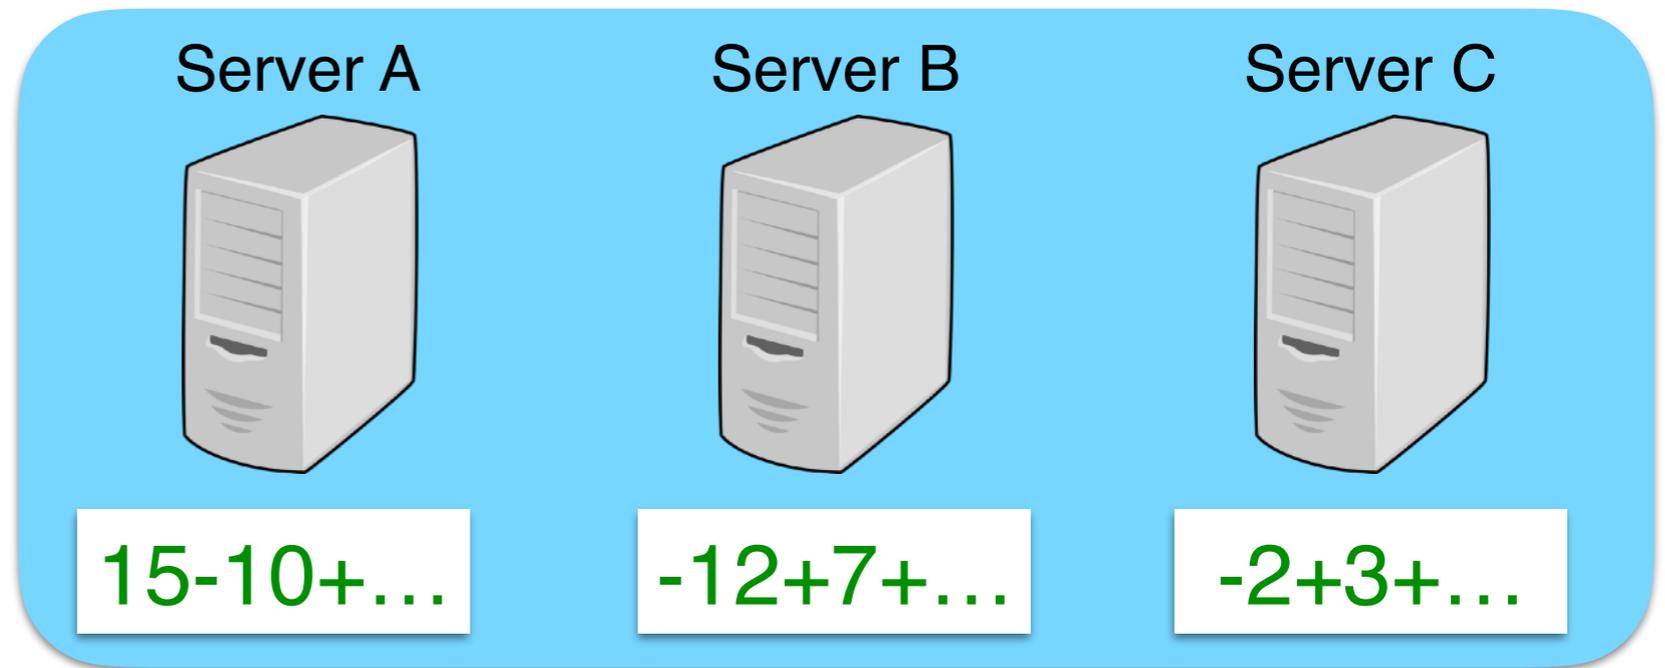
Private sums: A “straw-man” scheme



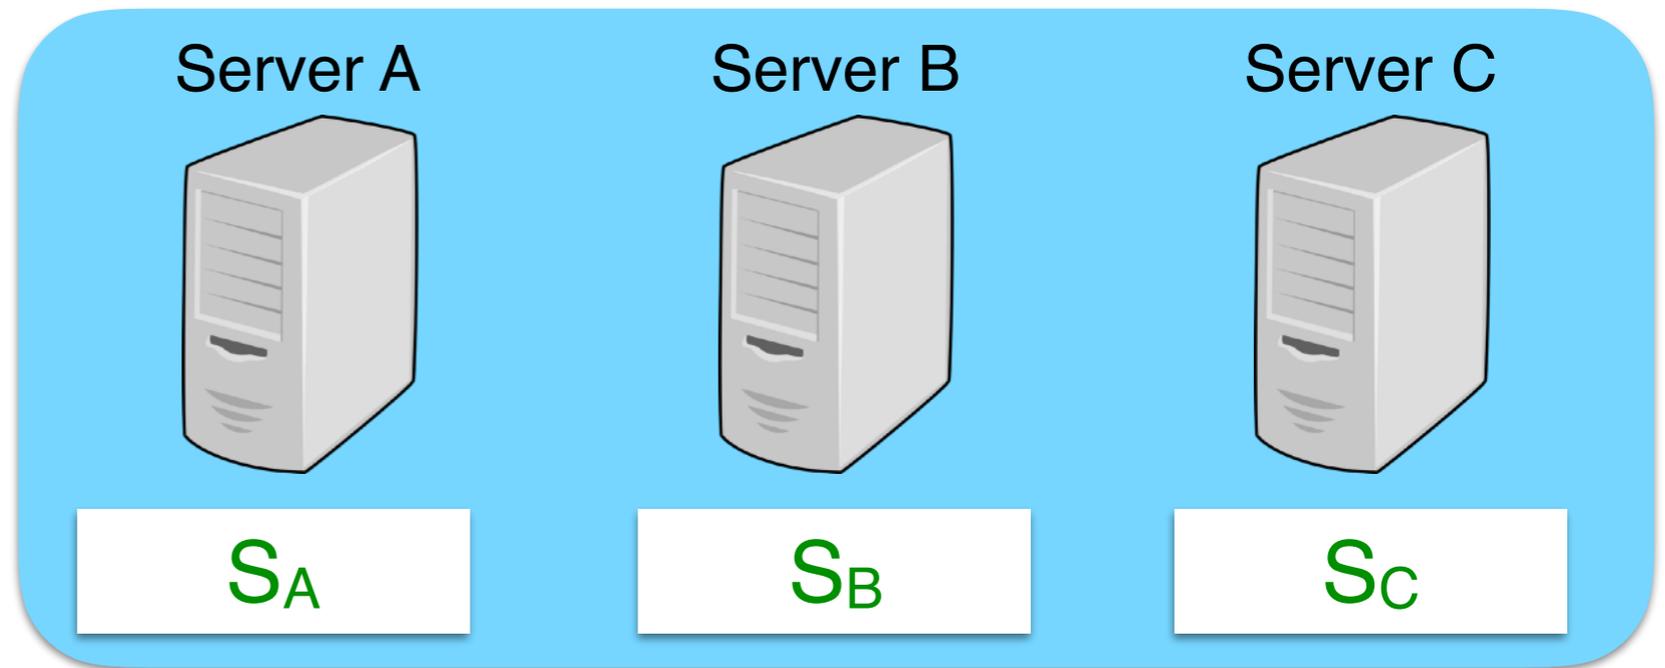
0



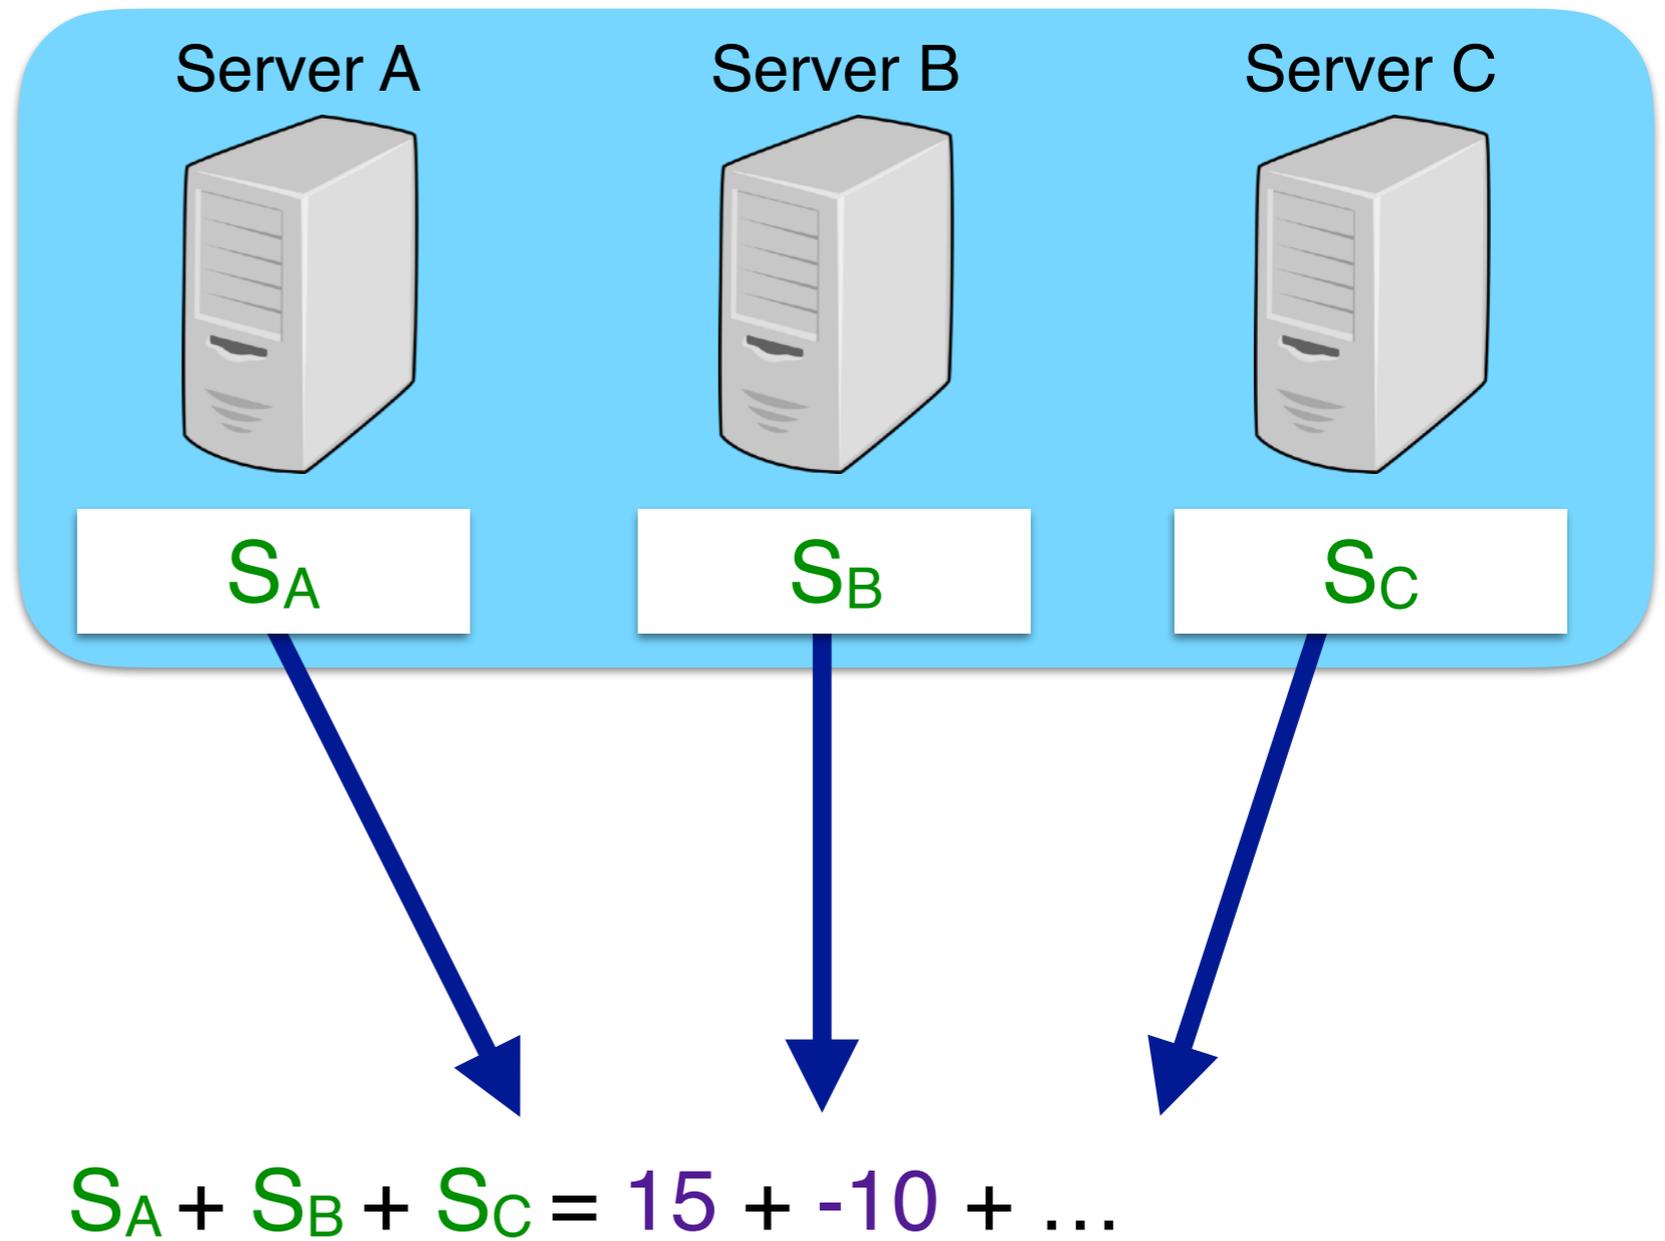
Private sums: A “straw-man” scheme



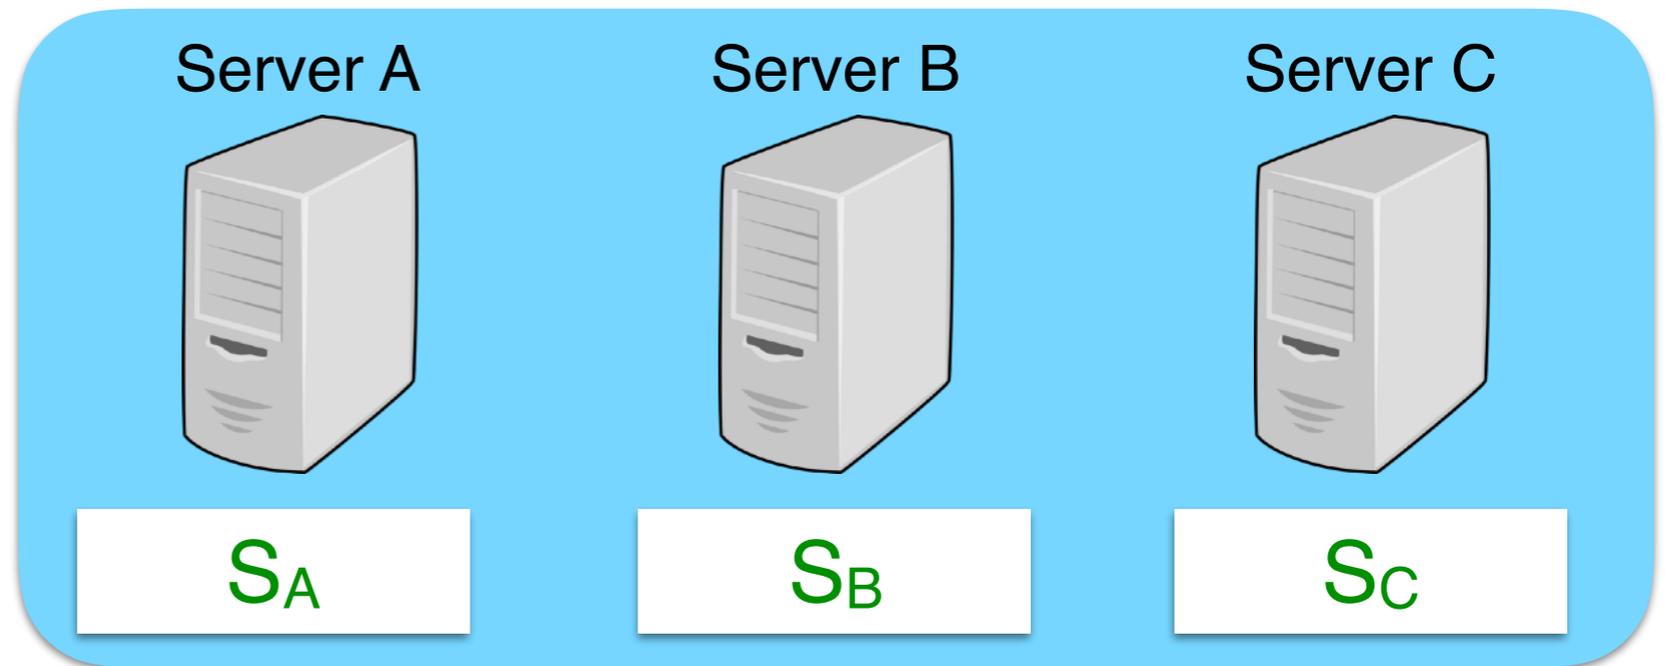
Private sums: A “straw-man” scheme



Private sums:
A “straw-man”
scheme

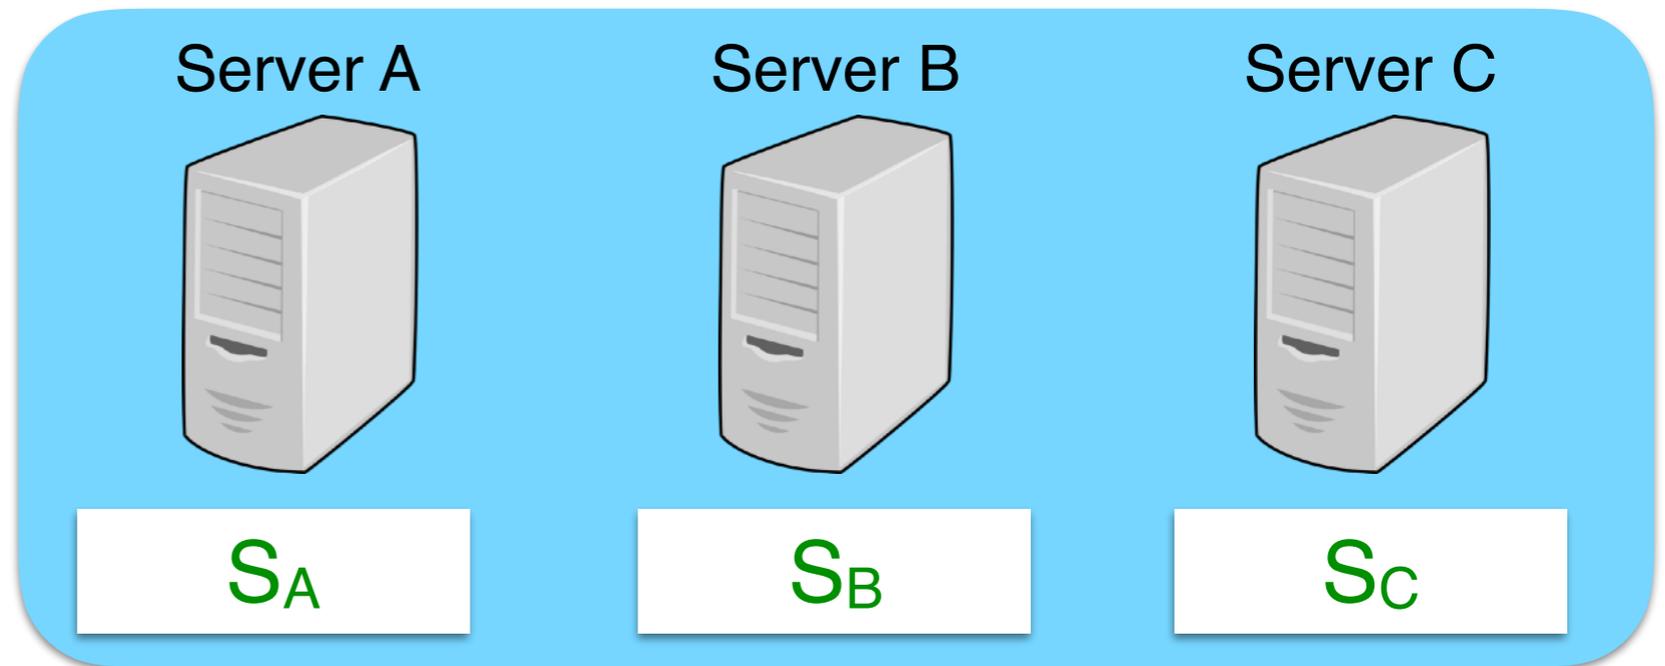


Private sums: A “straw-man” scheme



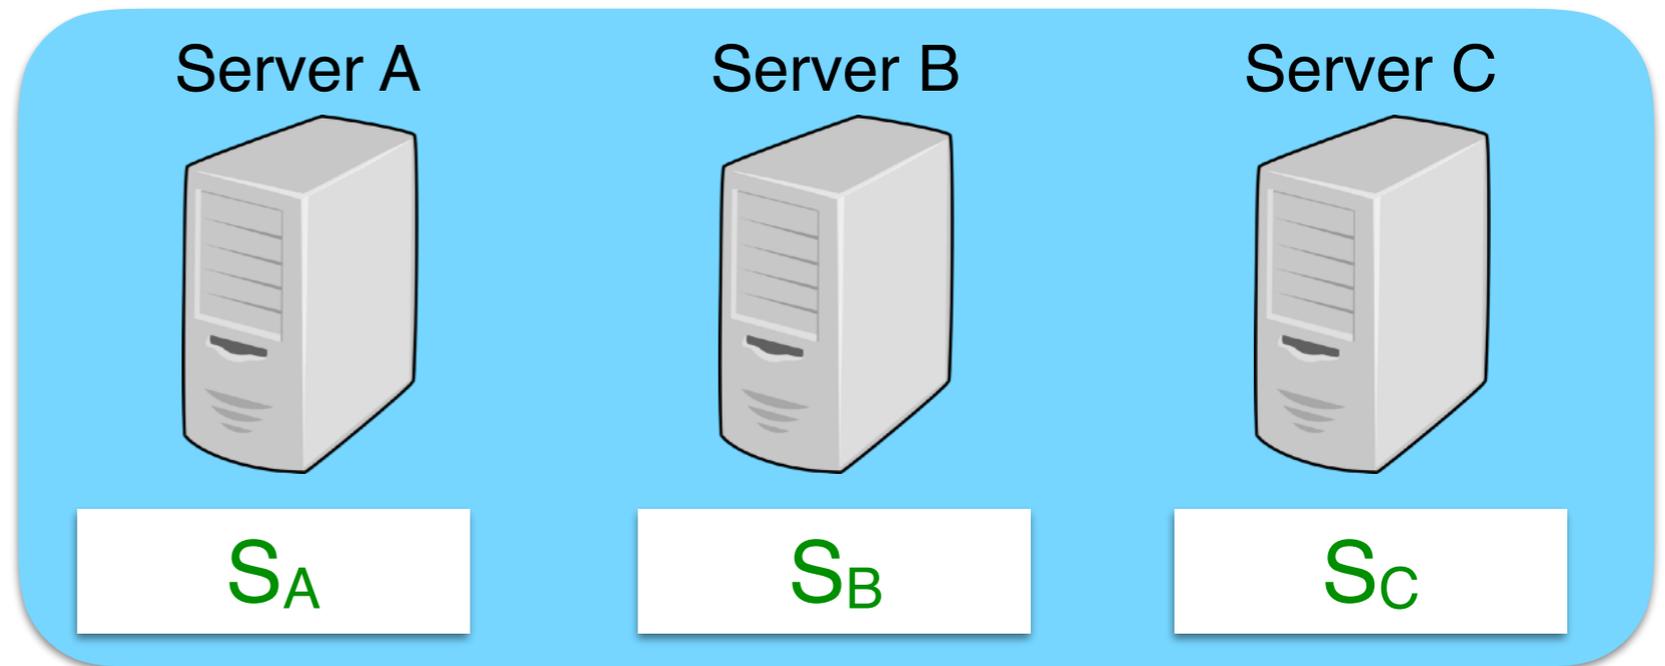
$$S_A + S_B + S_C = 15 + -10 + \dots$$

Private sums: A “straw-man” scheme



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

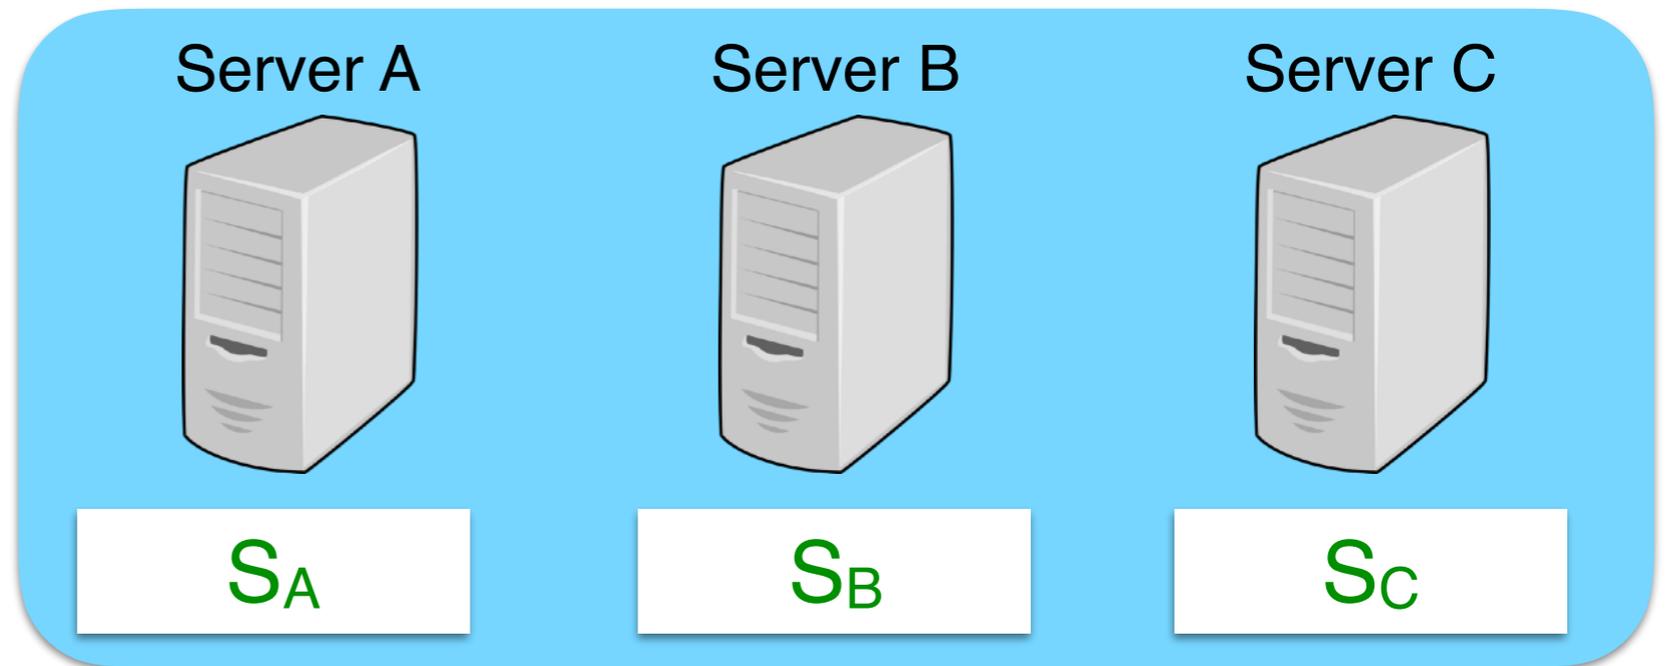
Private sums: A “straw-man” scheme



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

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

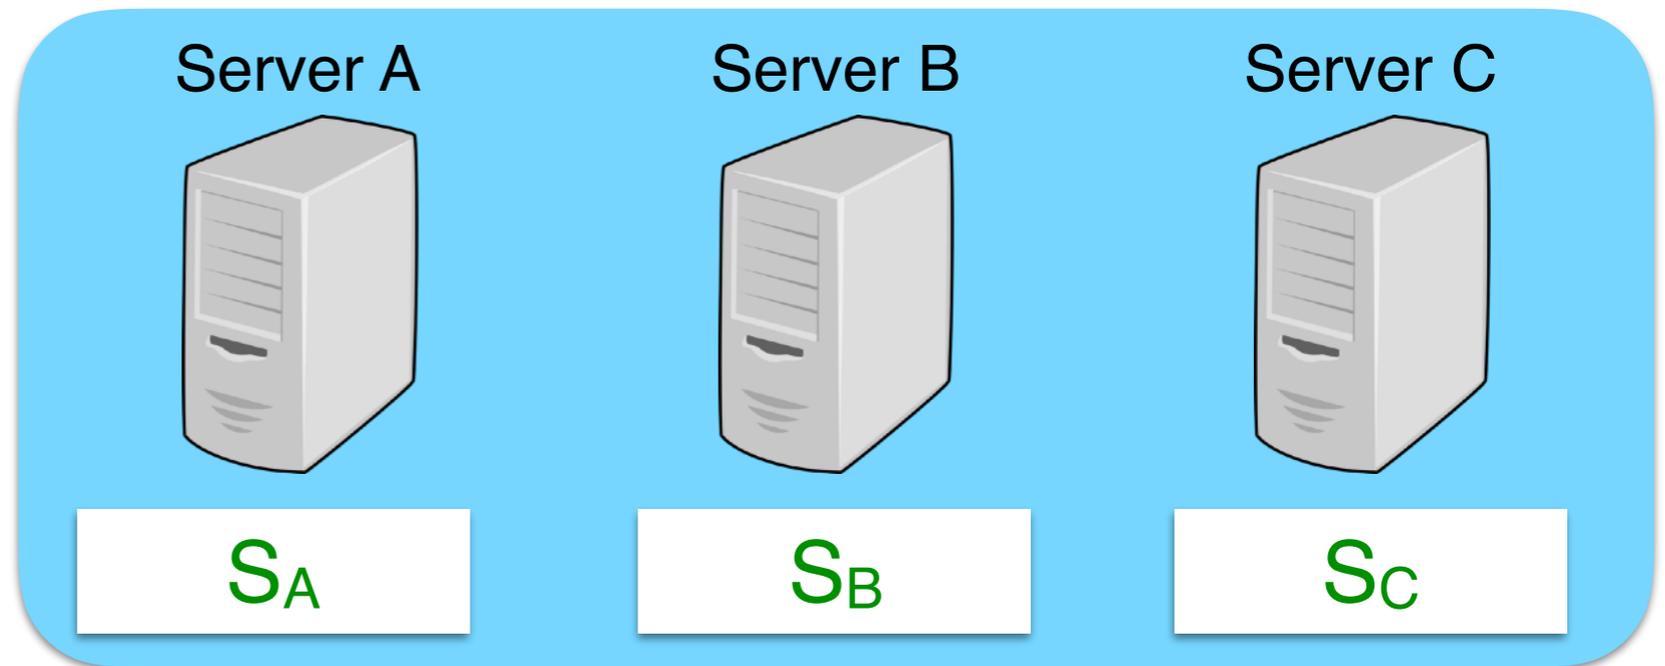
Private sums: A “straw-man” scheme



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

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

Private sums: A “straw-man” scheme



$$S_A + S_B + S_C = 15 + -10 + \dots$$
$$= 1 + 0 + \dots + 1$$

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

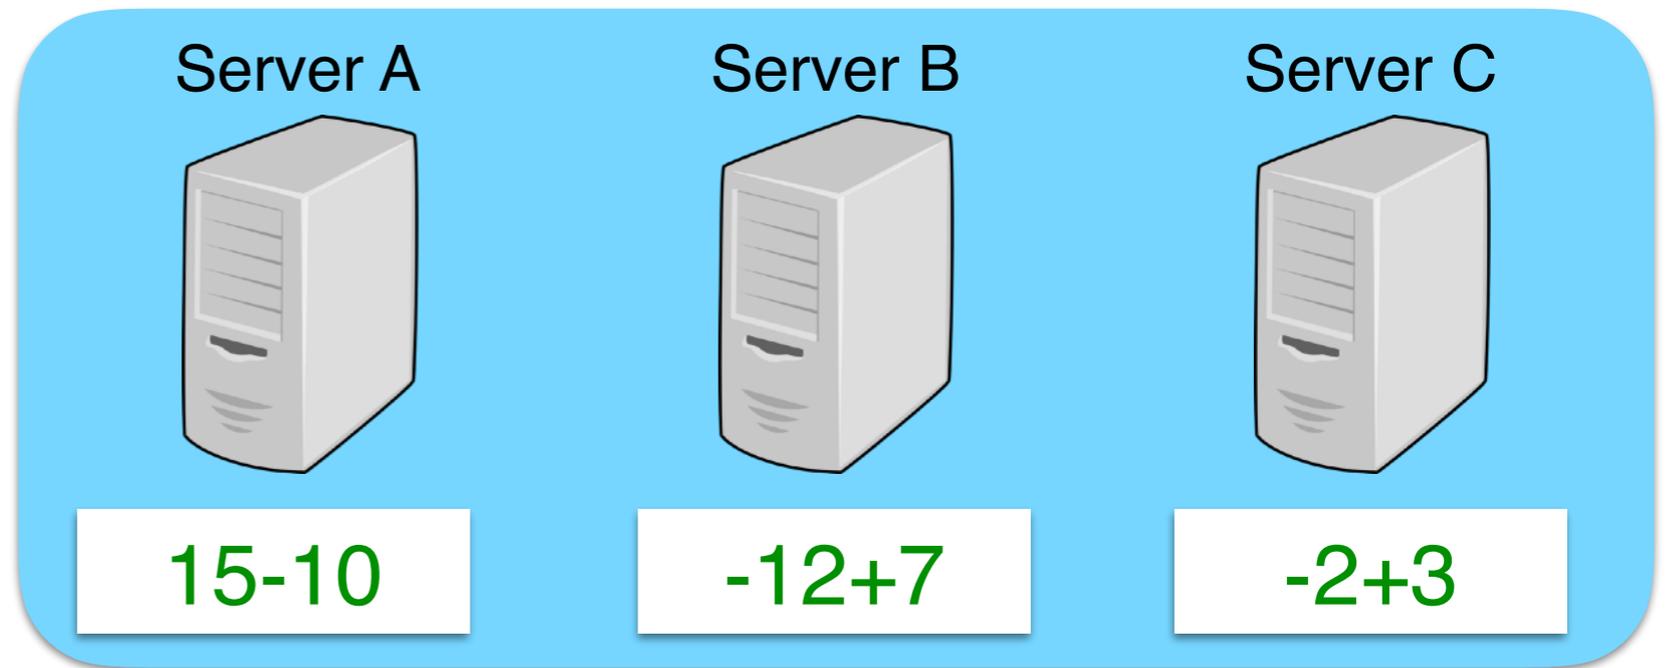
Exact correctness: If everyone follows the protocol, servers compute the sum of all x_i s.

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

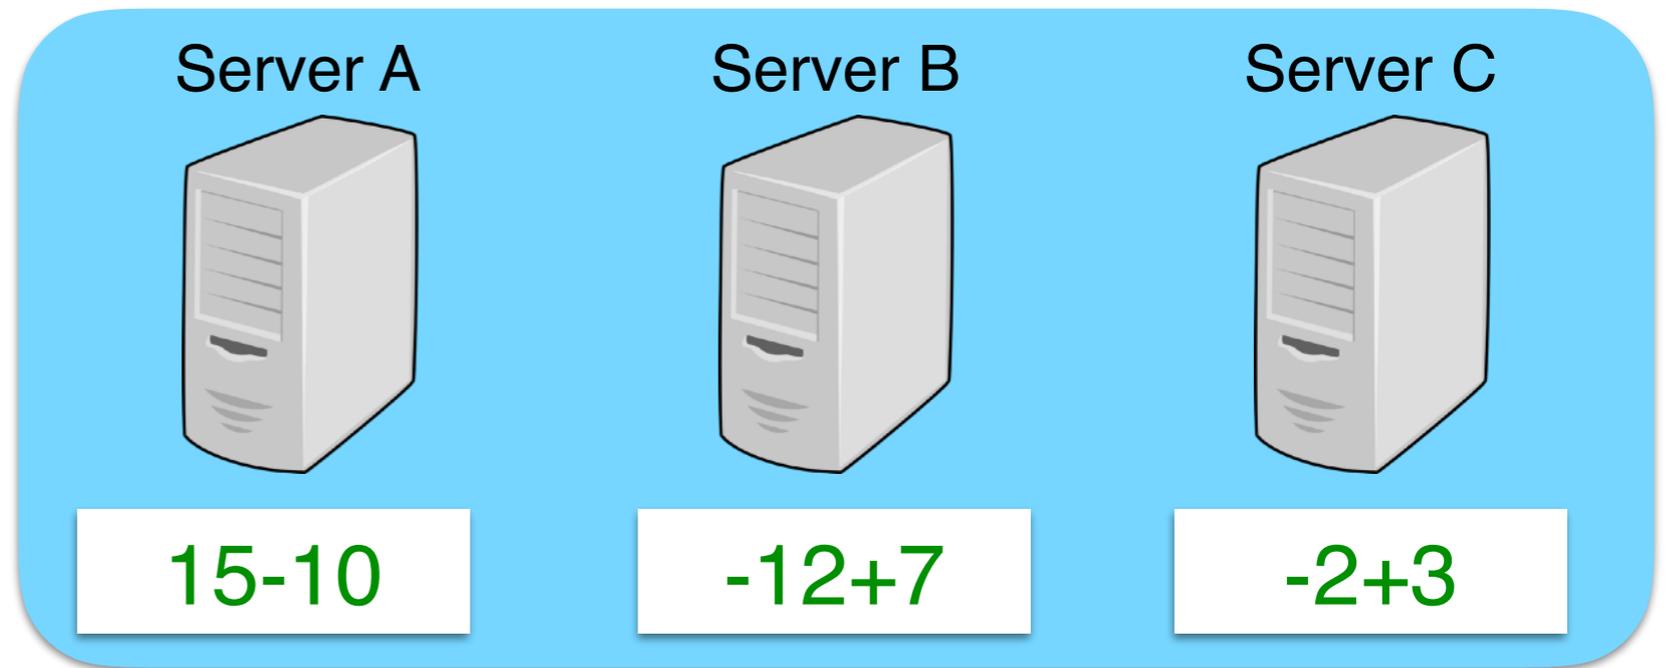
Efficiency: Follows by inspection.

Robustness: ???

Private sums: A “straw-man” scheme



Private sums: A “straw-man” scheme

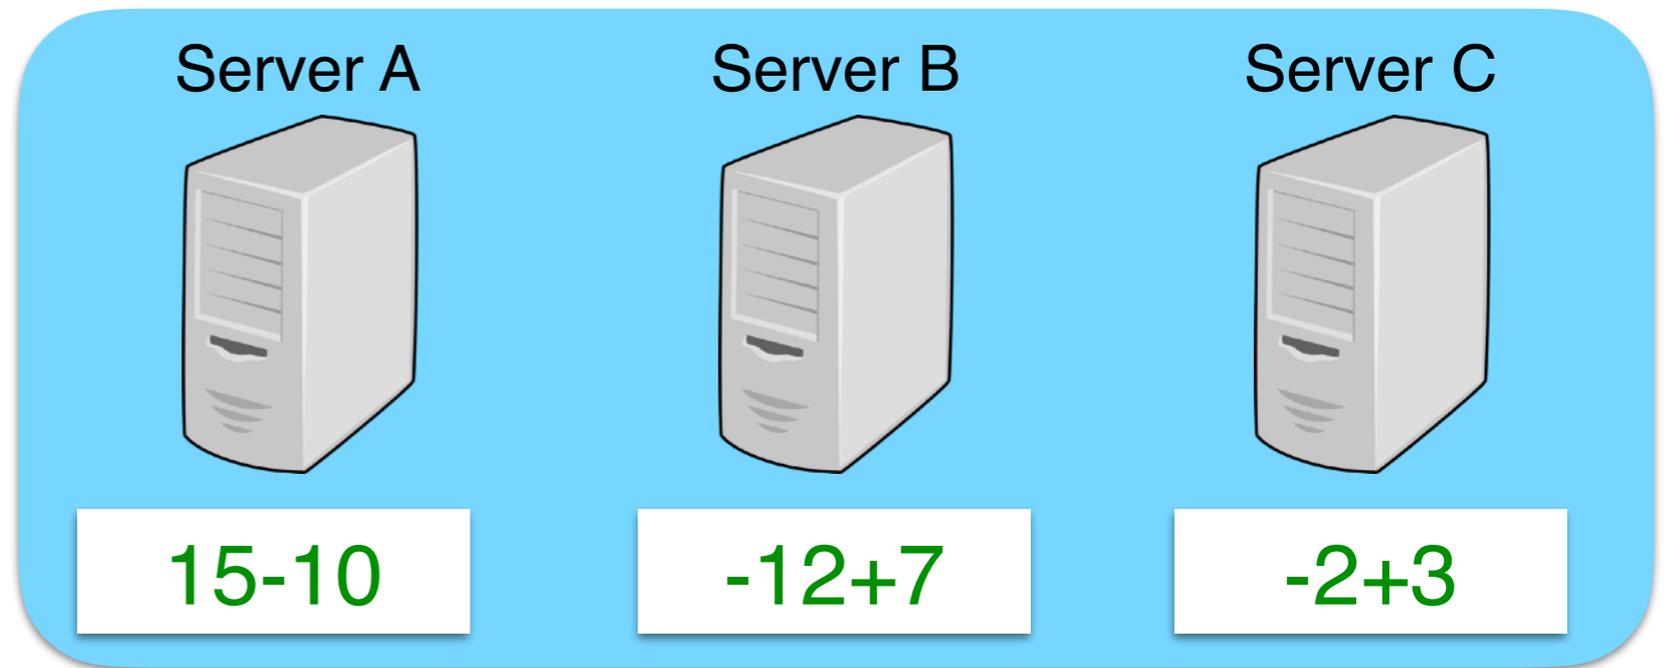


x is supposed to be a 0/1 value



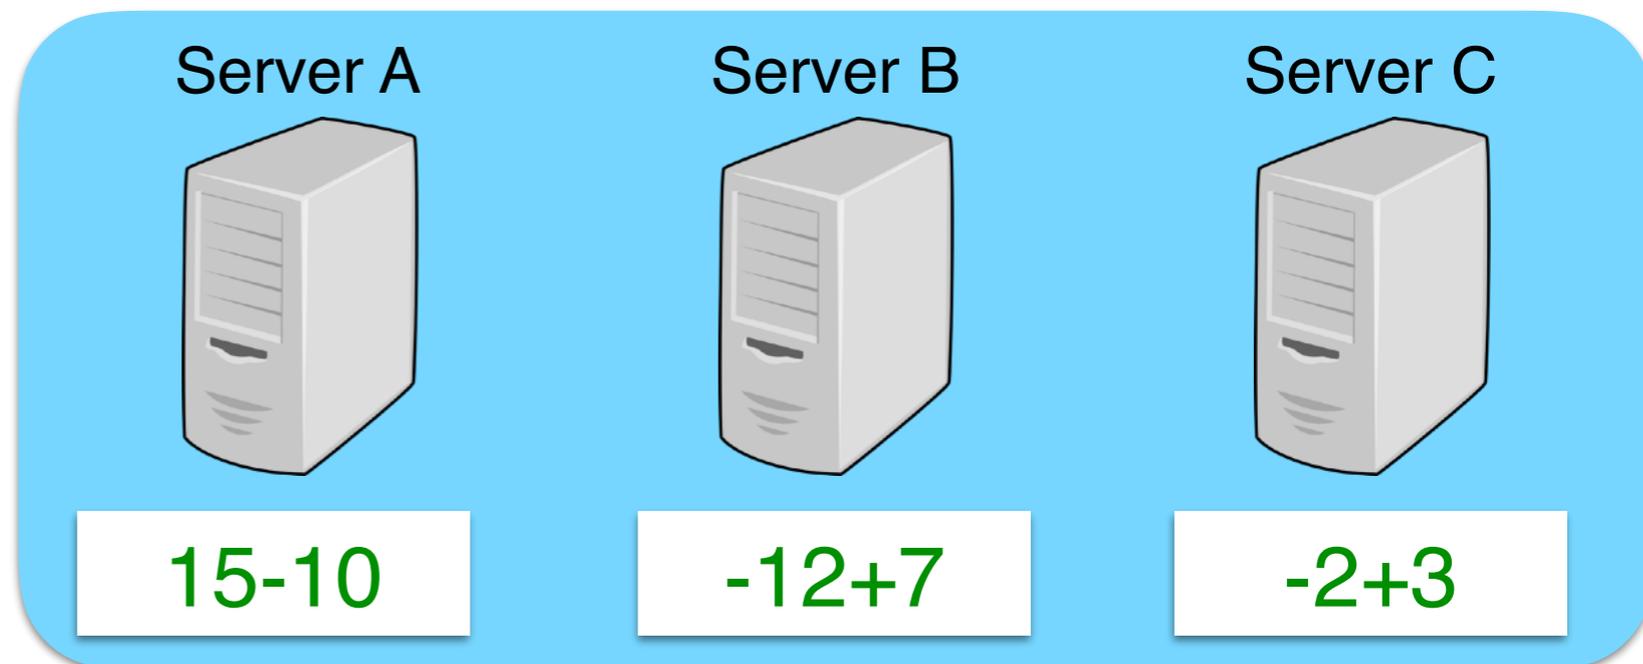
x

Private sums: A “straw-man” scheme

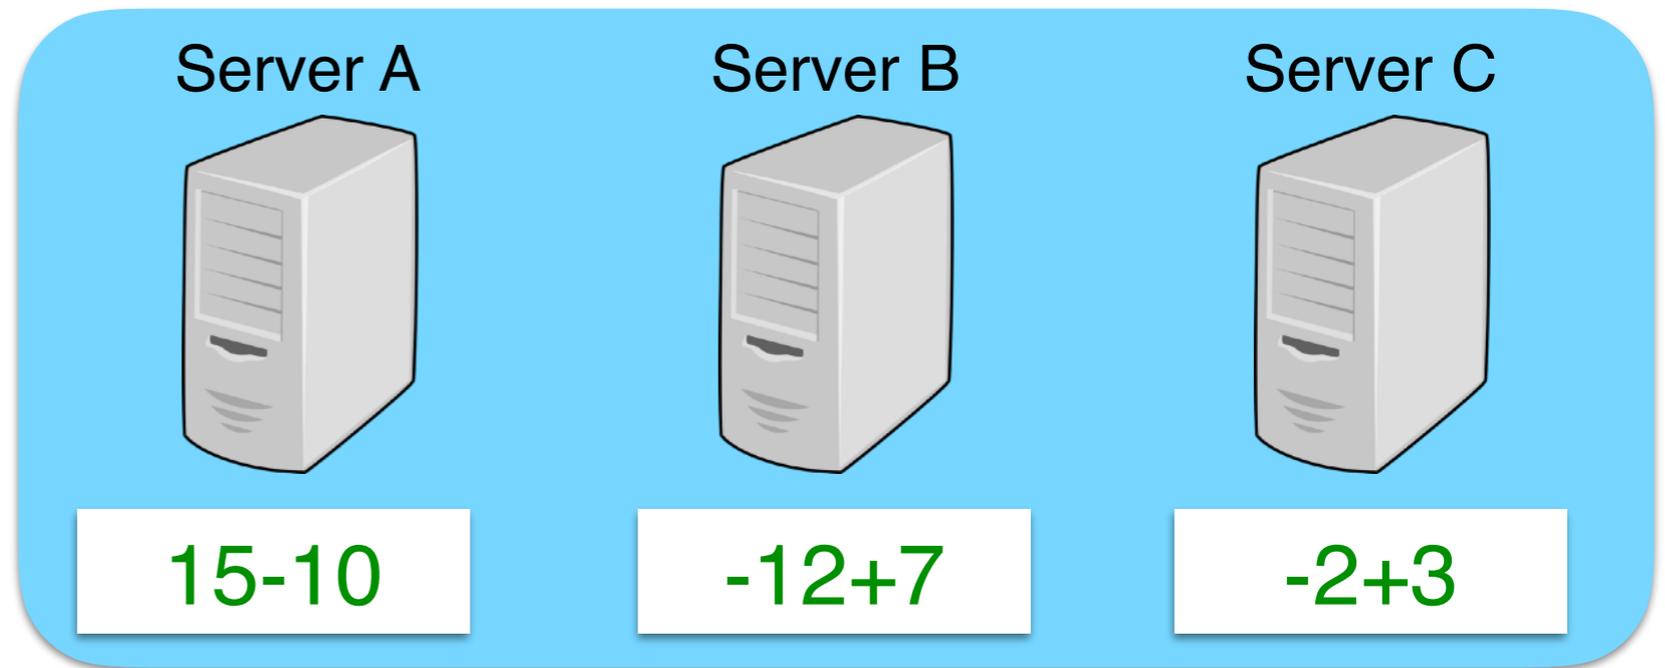


X

Private sums: A “straw-man” scheme



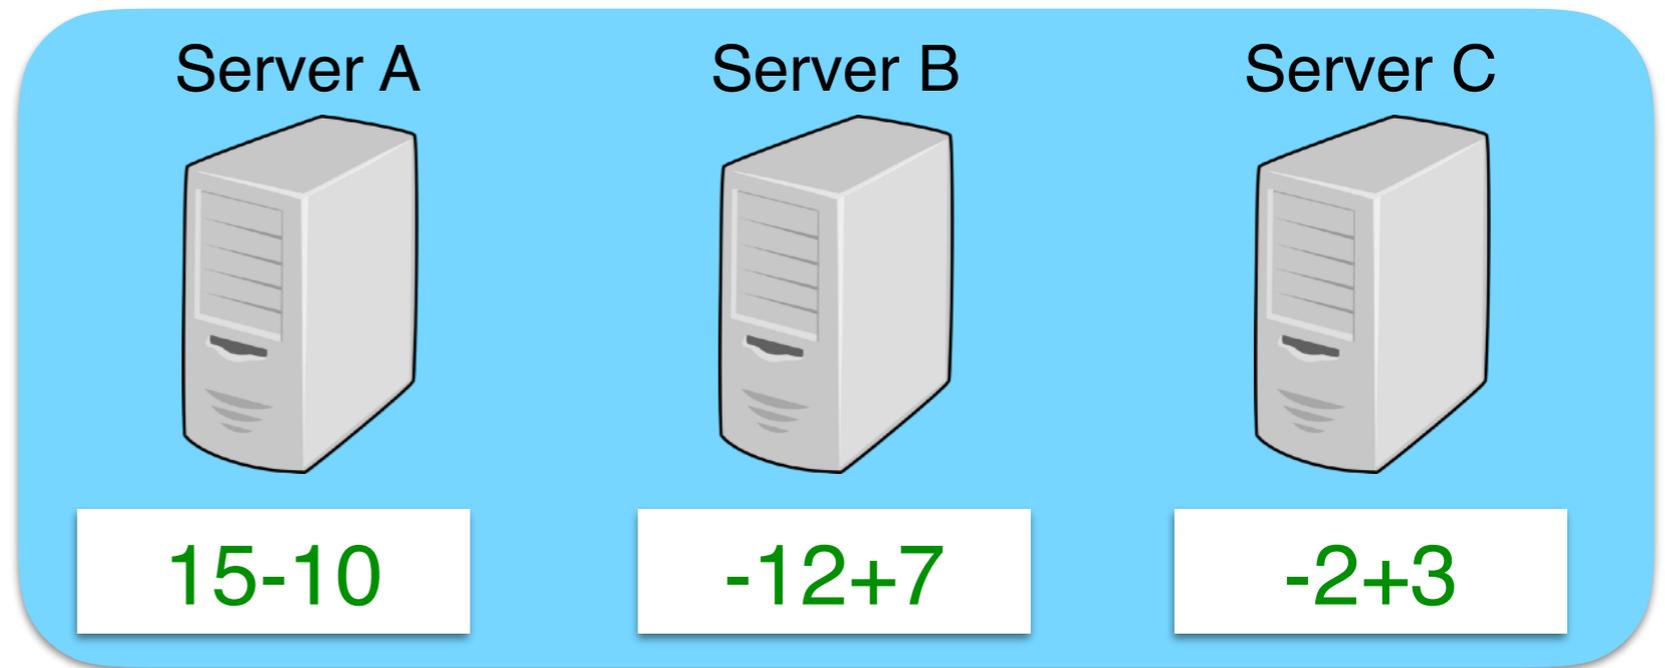
Private sums: A “straw-man” scheme



An evil client needn't
follow the rules!



Private sums: A “straw-man” scheme

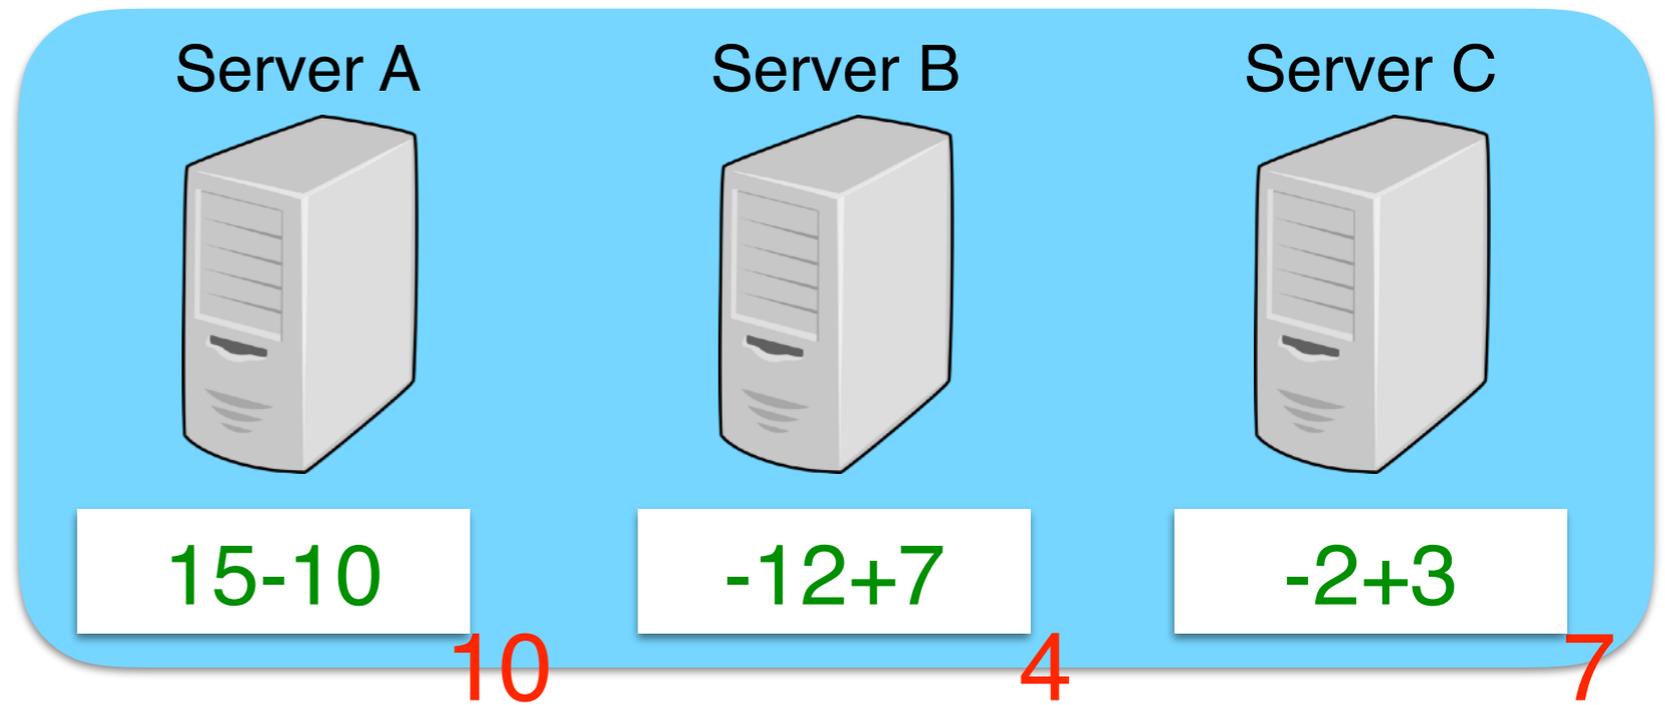


An evil client needn't
follow the rules!

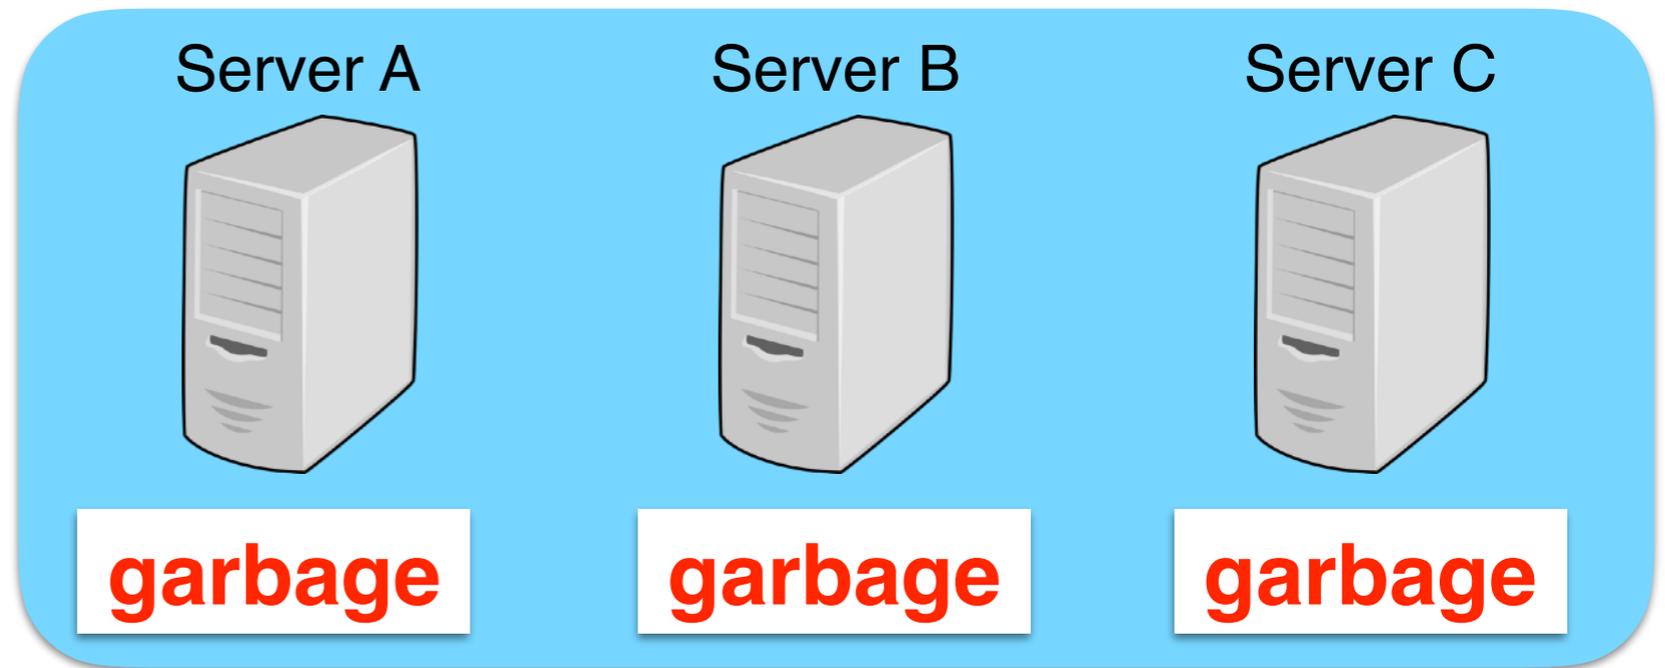
$$10 + 4 + 7 = 21$$



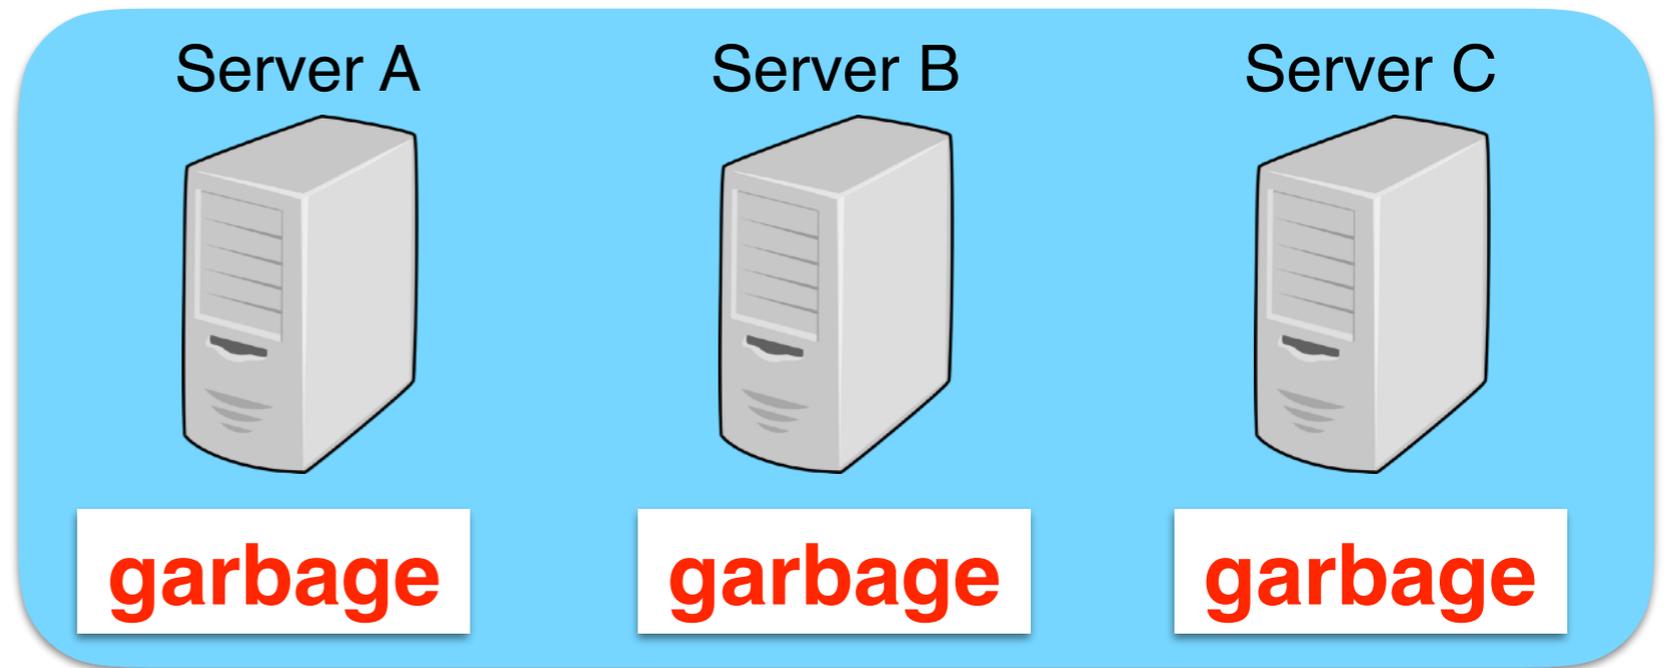
Private sums: A “straw-man” scheme



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

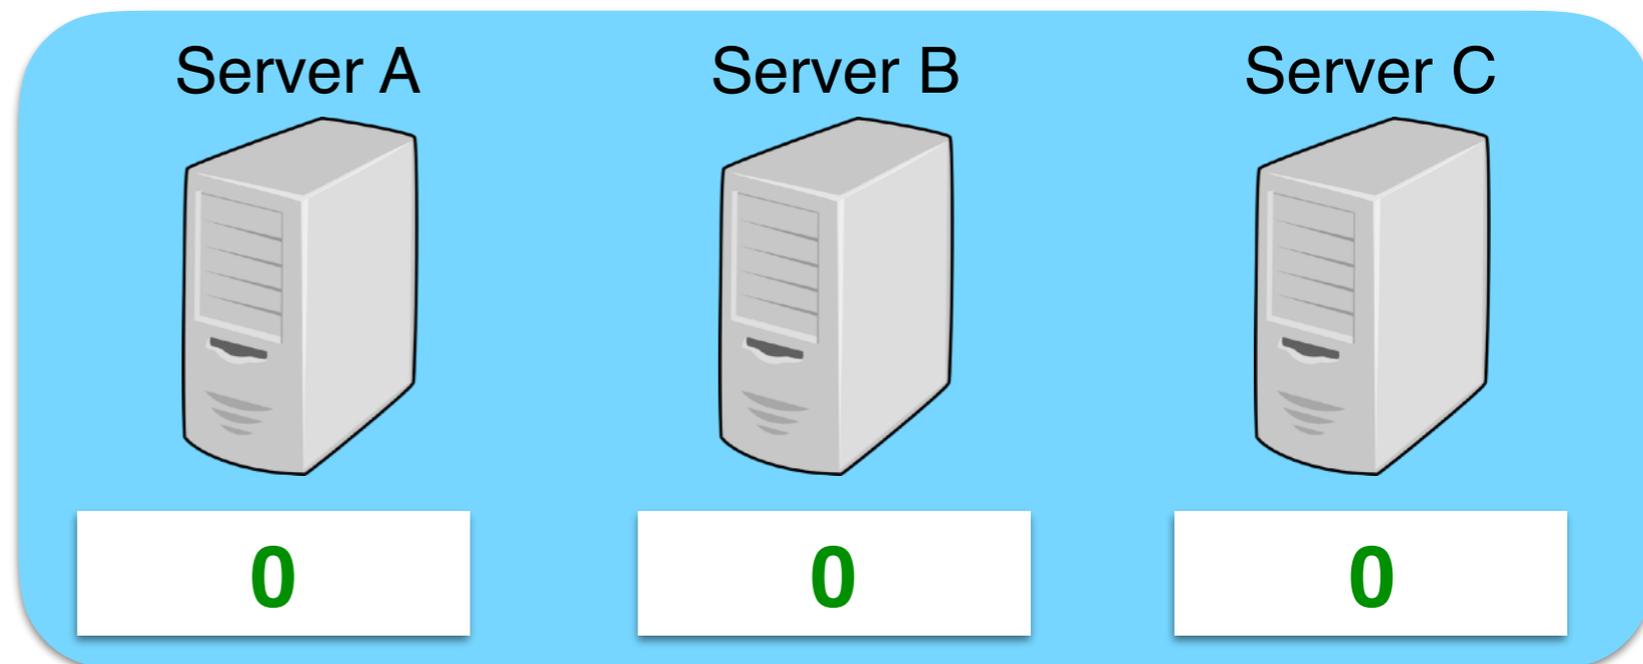
- Background: The private aggregation problem
- **A straw-man solution for private sums**
- Providing robustness with SNIPs
- Evaluation
- Encodings for complex aggregates

Outline

- Background: The private aggregation problem
- A straw-man solution for private sums
- **Providing robustness with SNIPs**
- Evaluation
- Encodings for complex aggregates

Contribution 1

Secret-shared non-interactive proofs (SNIPs)

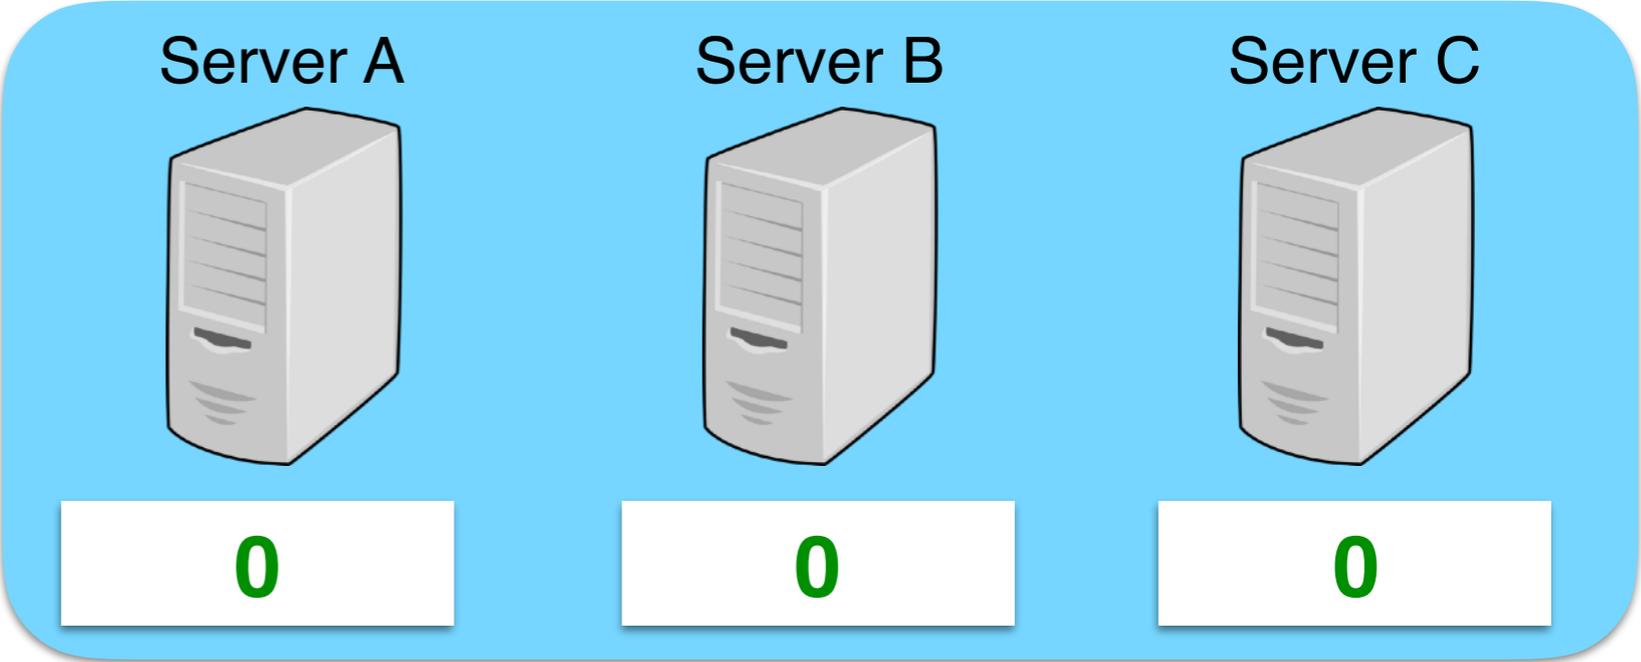


$$x = 1$$

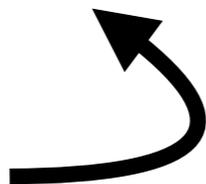


Contribution 1

Secret-shared
non-interactive
proofs (SNIPs)

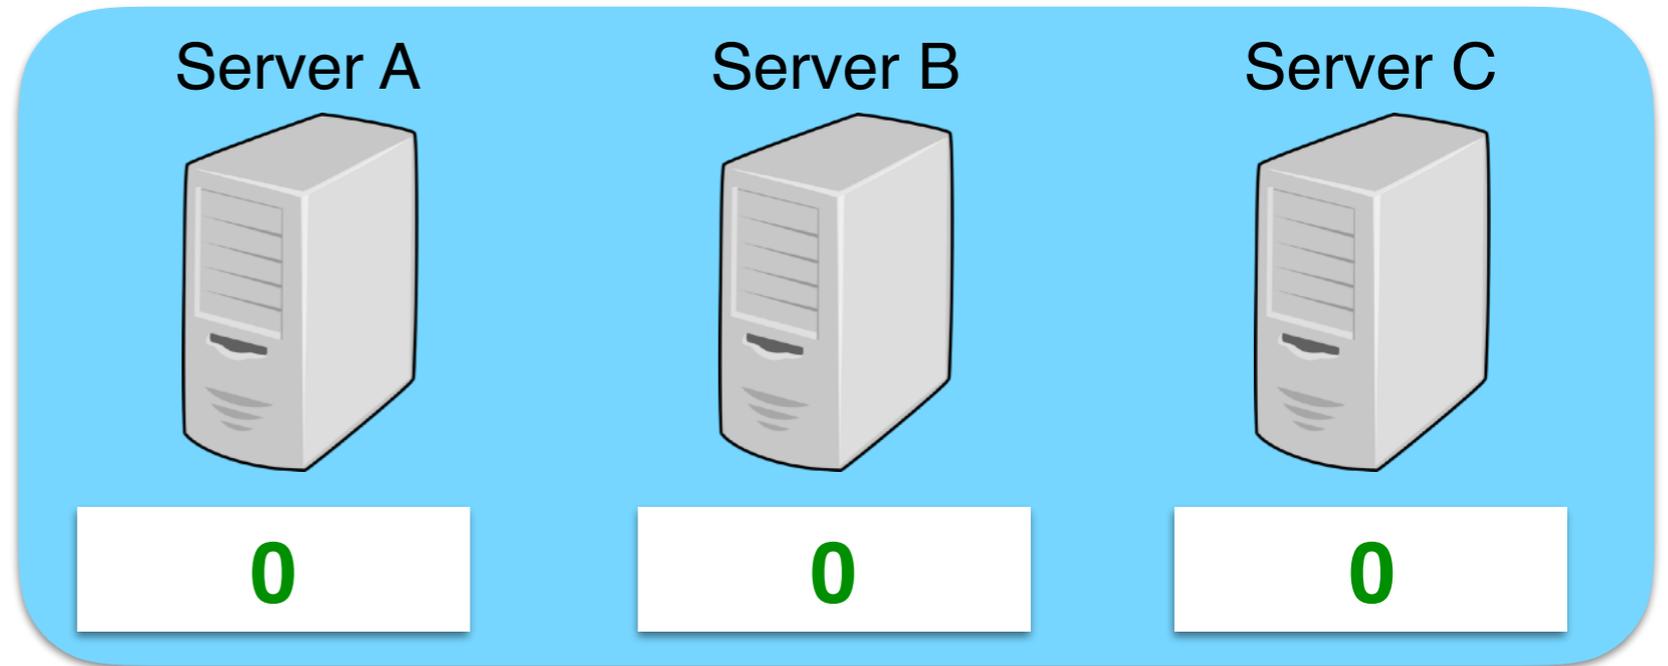


$x = 1$



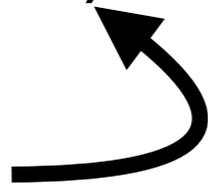
Contribution 1

Secret-shared
non-interactive
proofs (SNIPs)



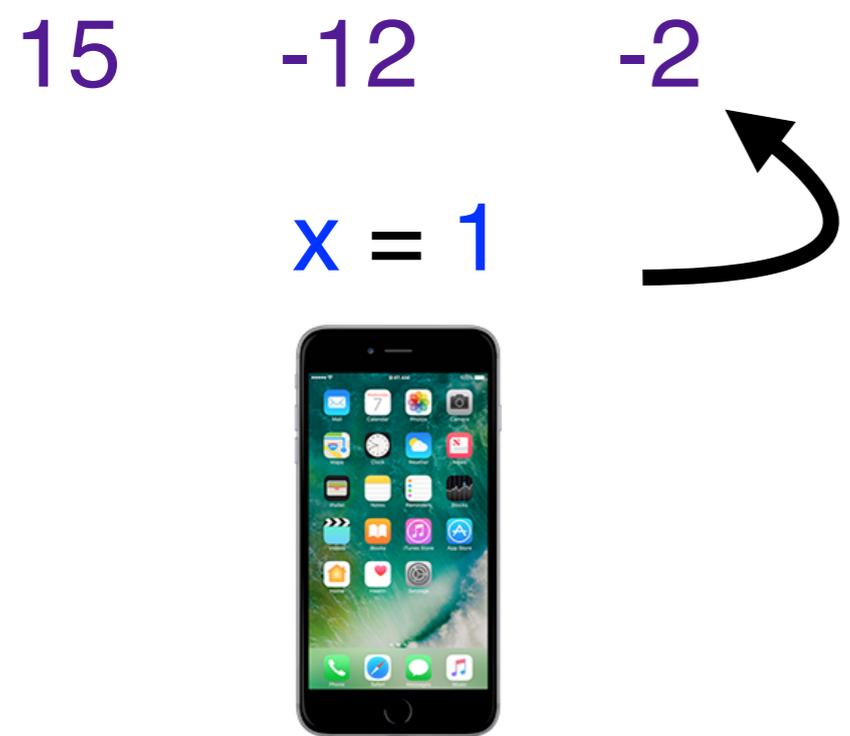
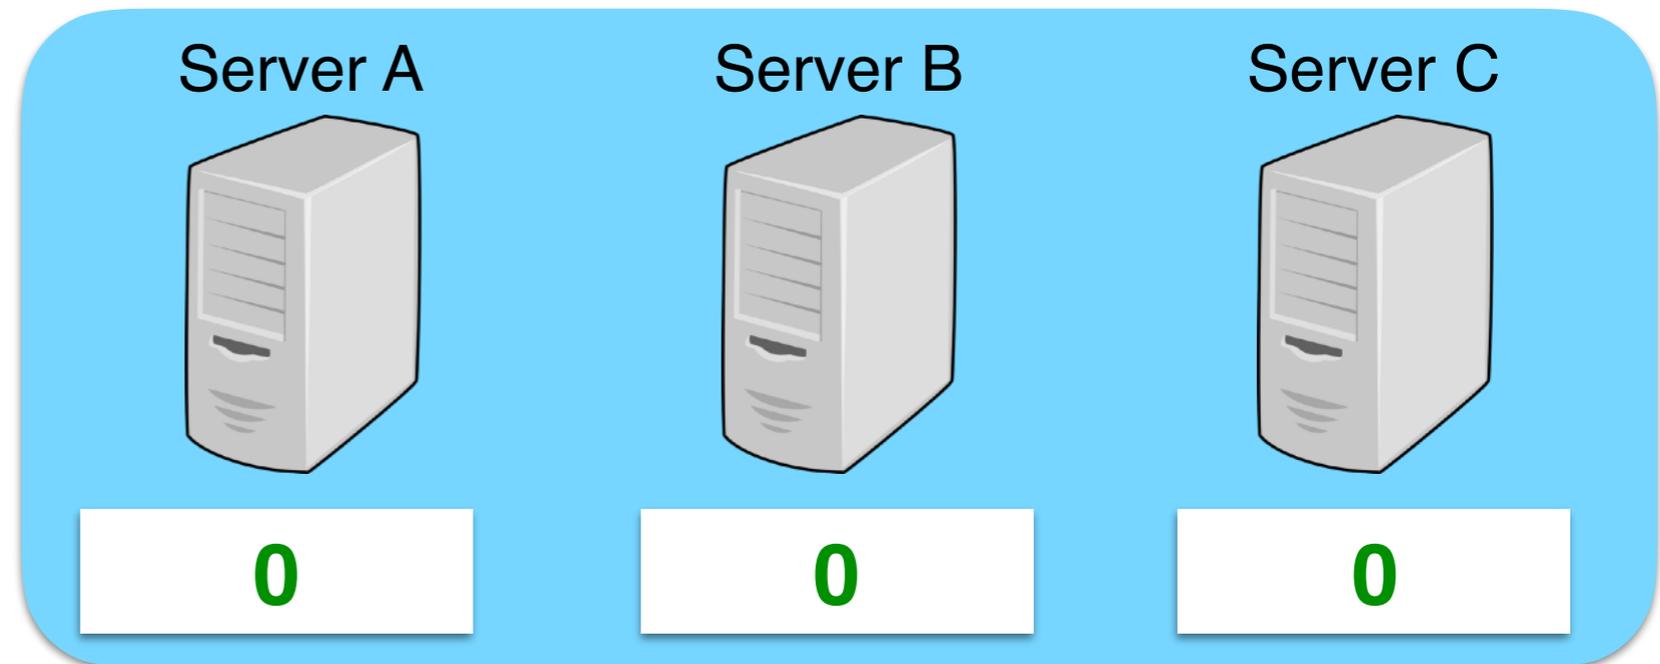
$$15 + (-12) + (-2) = 1$$

$$x = 1$$



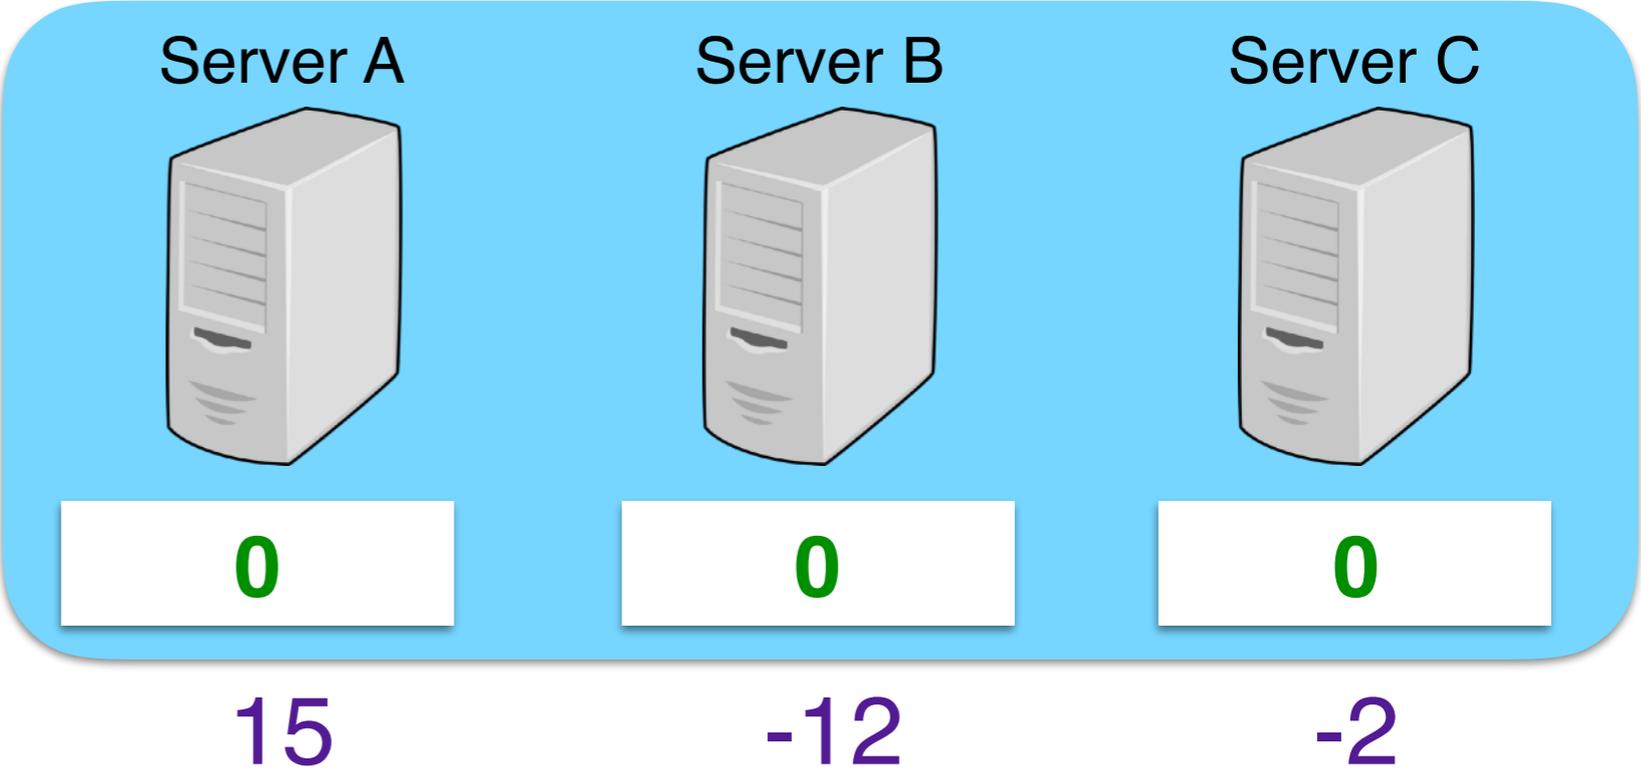
Contribution 1

Secret-shared
non-interactive
proofs (SNIPs)



Contribution 1

Secret-shared
non-interactive
proofs (SNIPs)



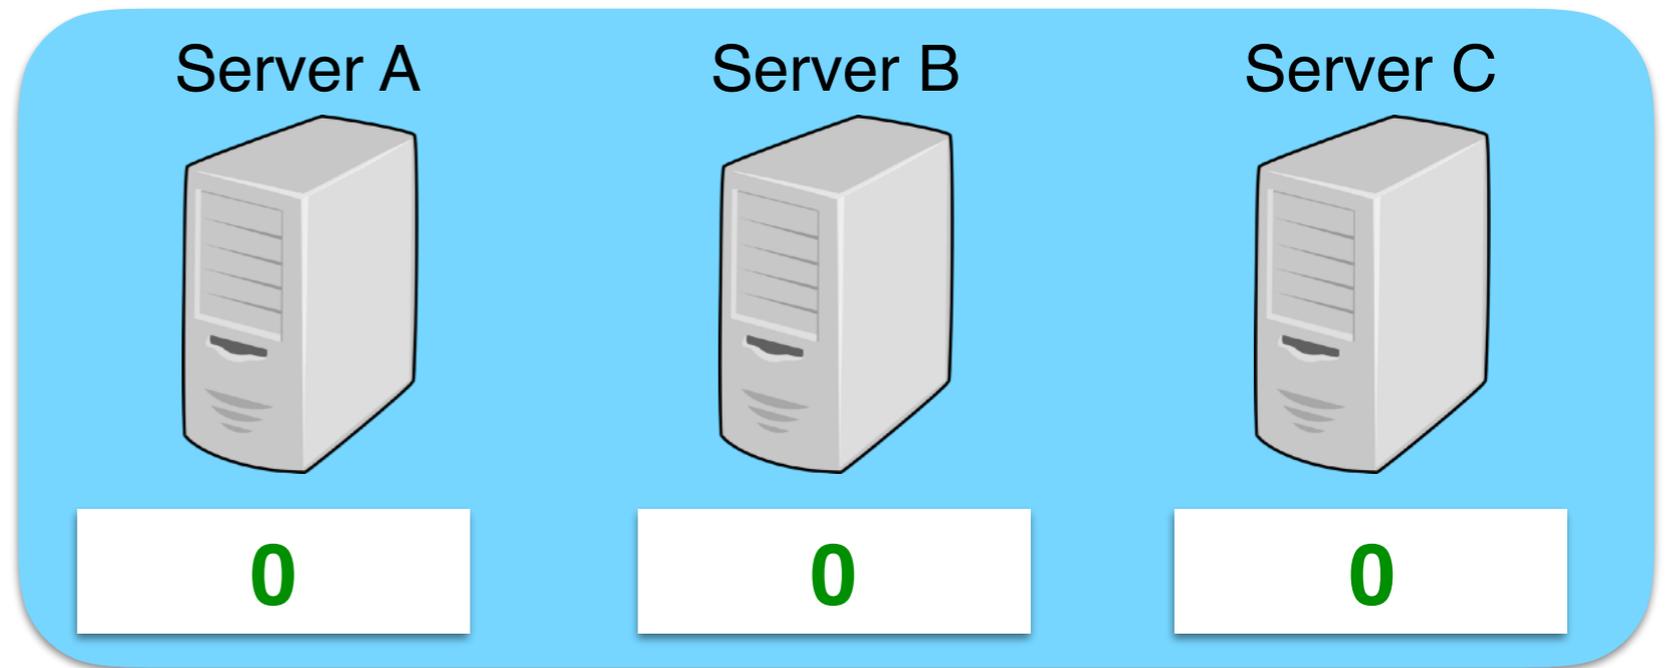
$$x = 1$$



Contribution 1

Secret-shared
non-interactive
proofs (SNIPs)

$$x = 1$$



15

-12

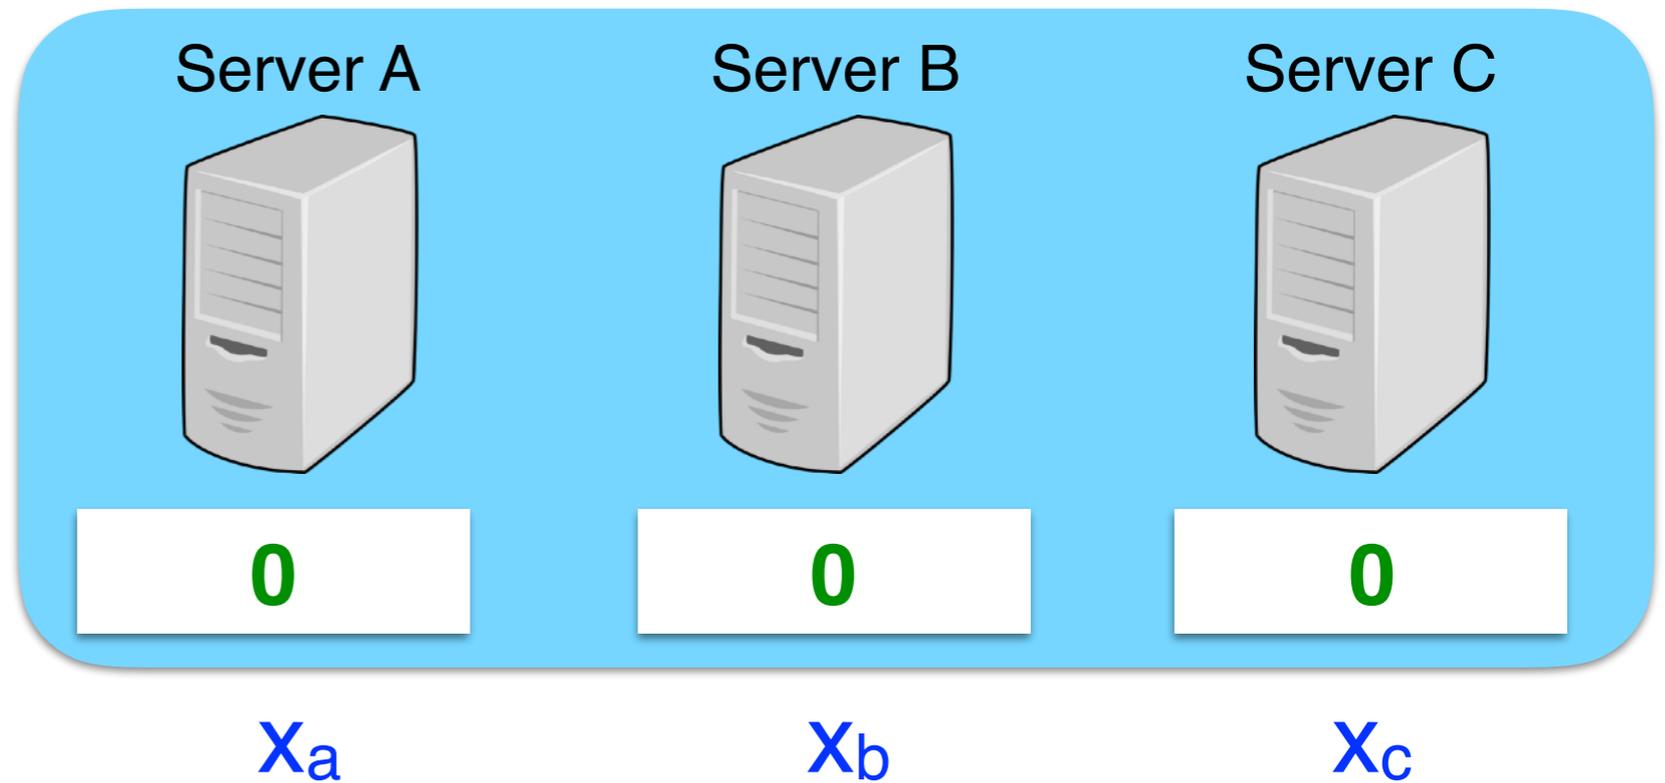
-2

The servers want to ensure that their shares sum to 0 or 1

...without learning x .

Contribution 1

Secret-shared non-interactive proofs (SNIPs)



$$x = 1$$

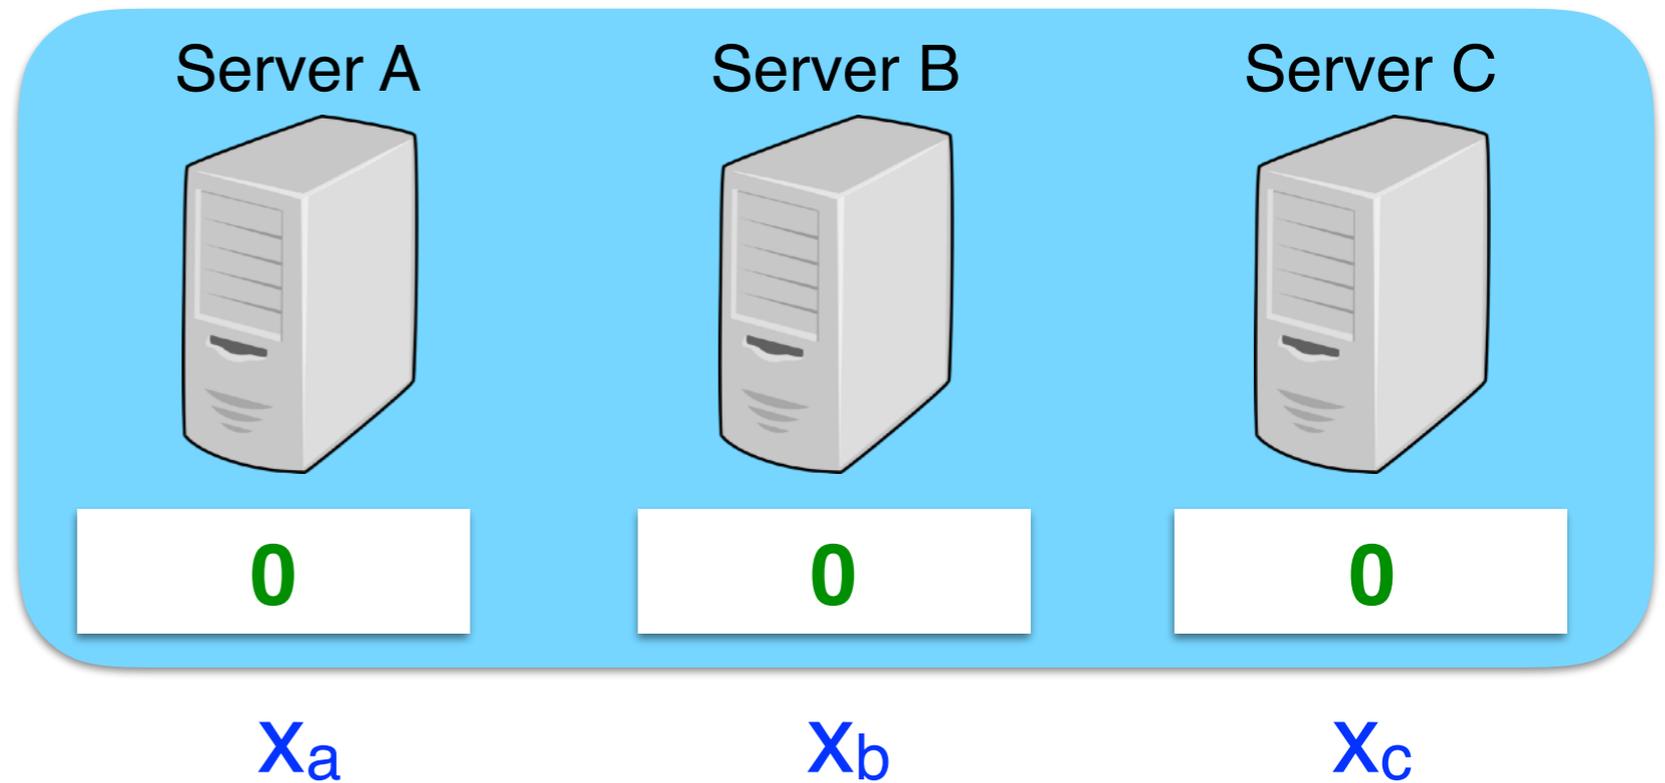


More generally, servers

- 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)



$x = 1$



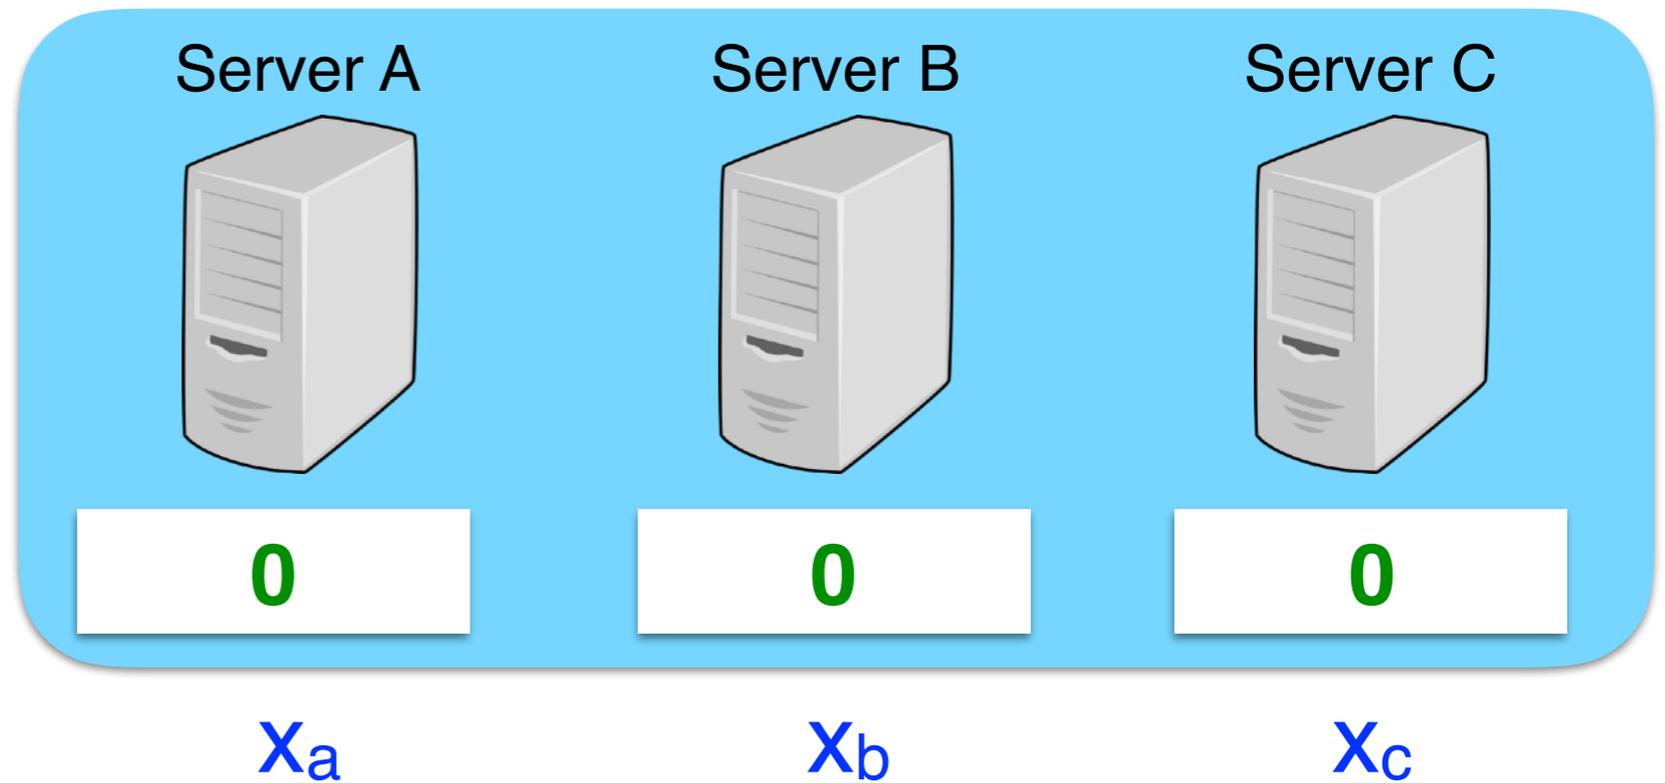
More generally, servers

- 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)



$$x = 1$$

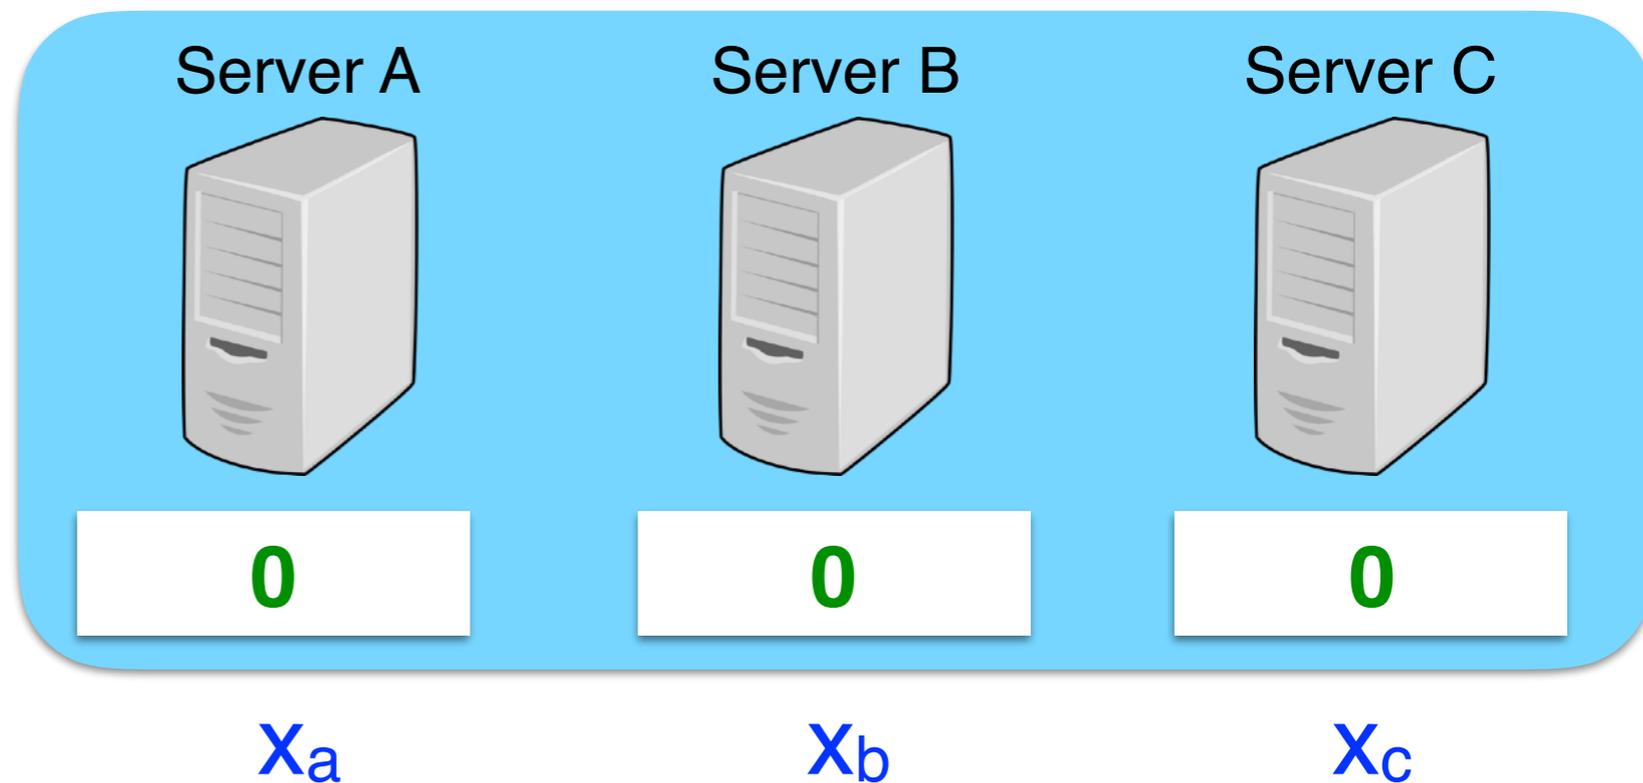


More generally, servers

- 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)

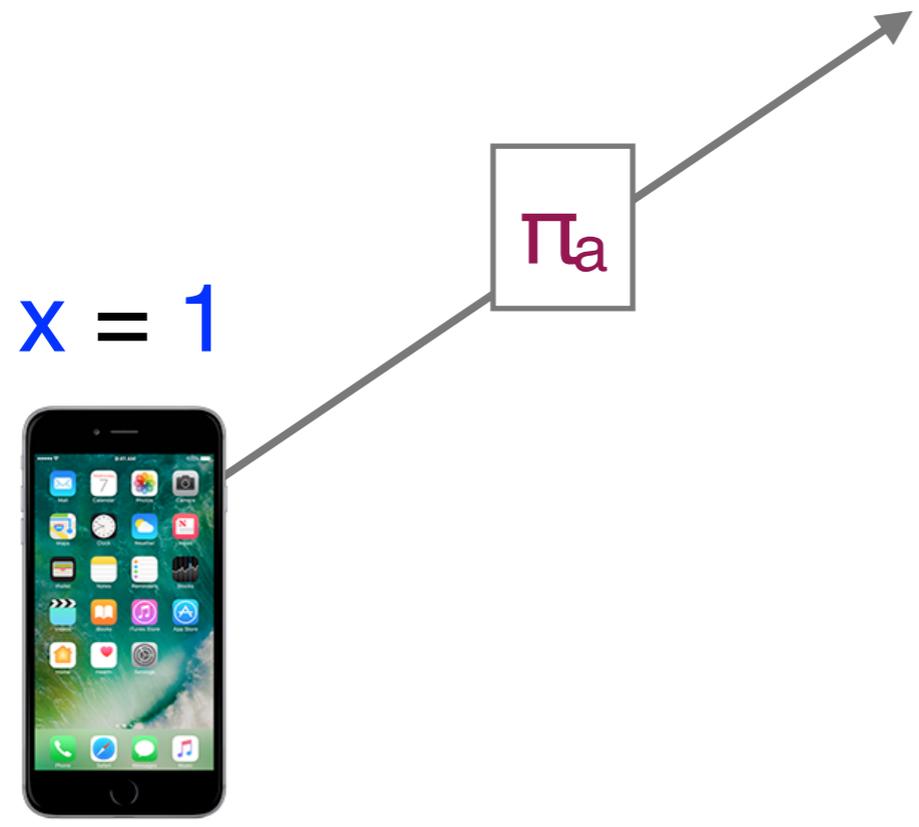
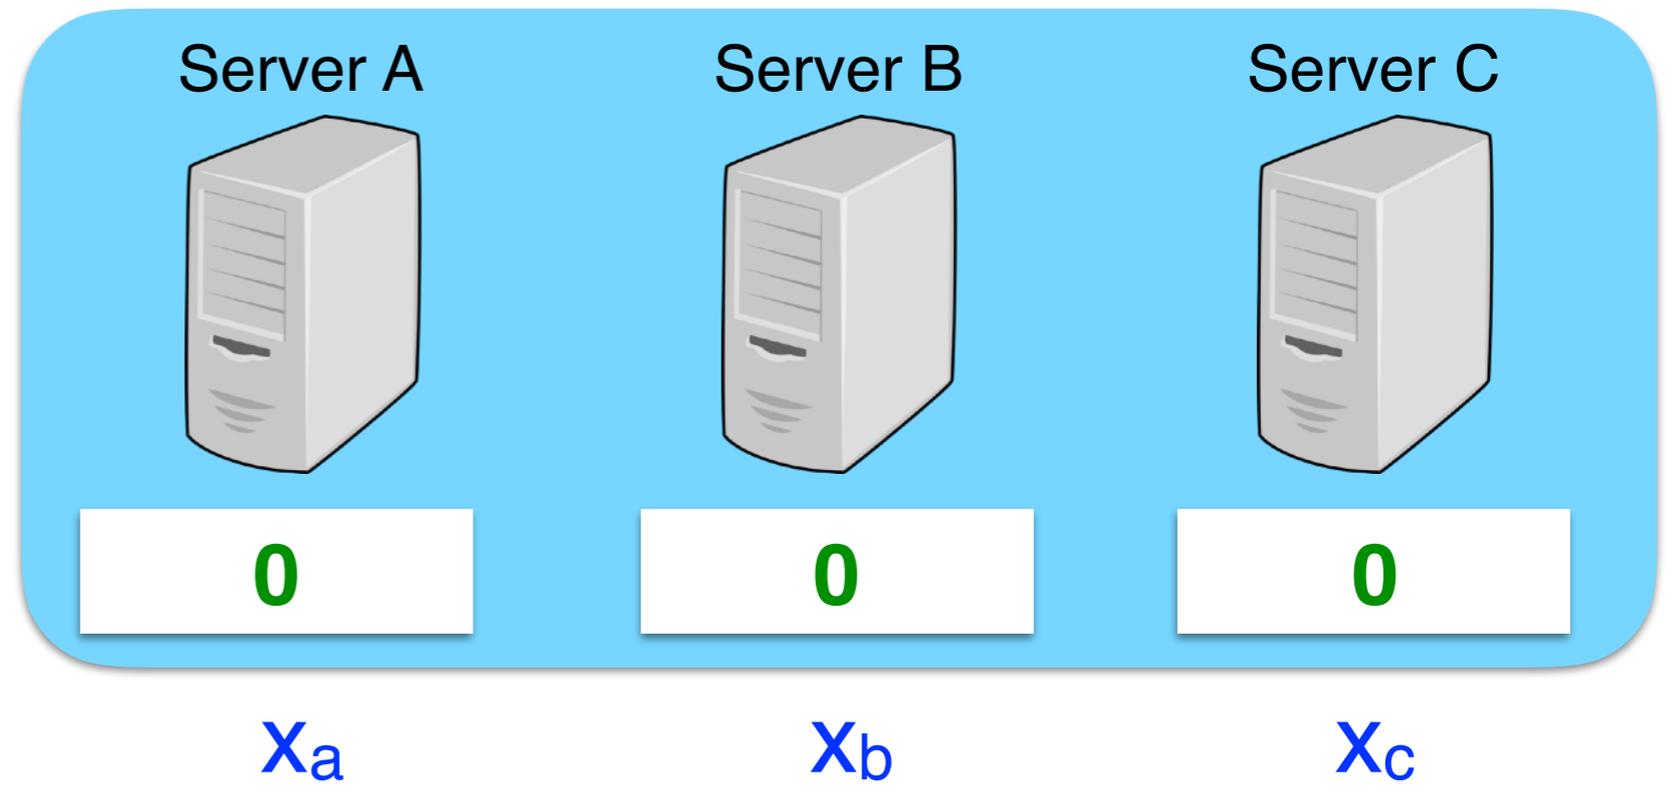


$$x = 1$$



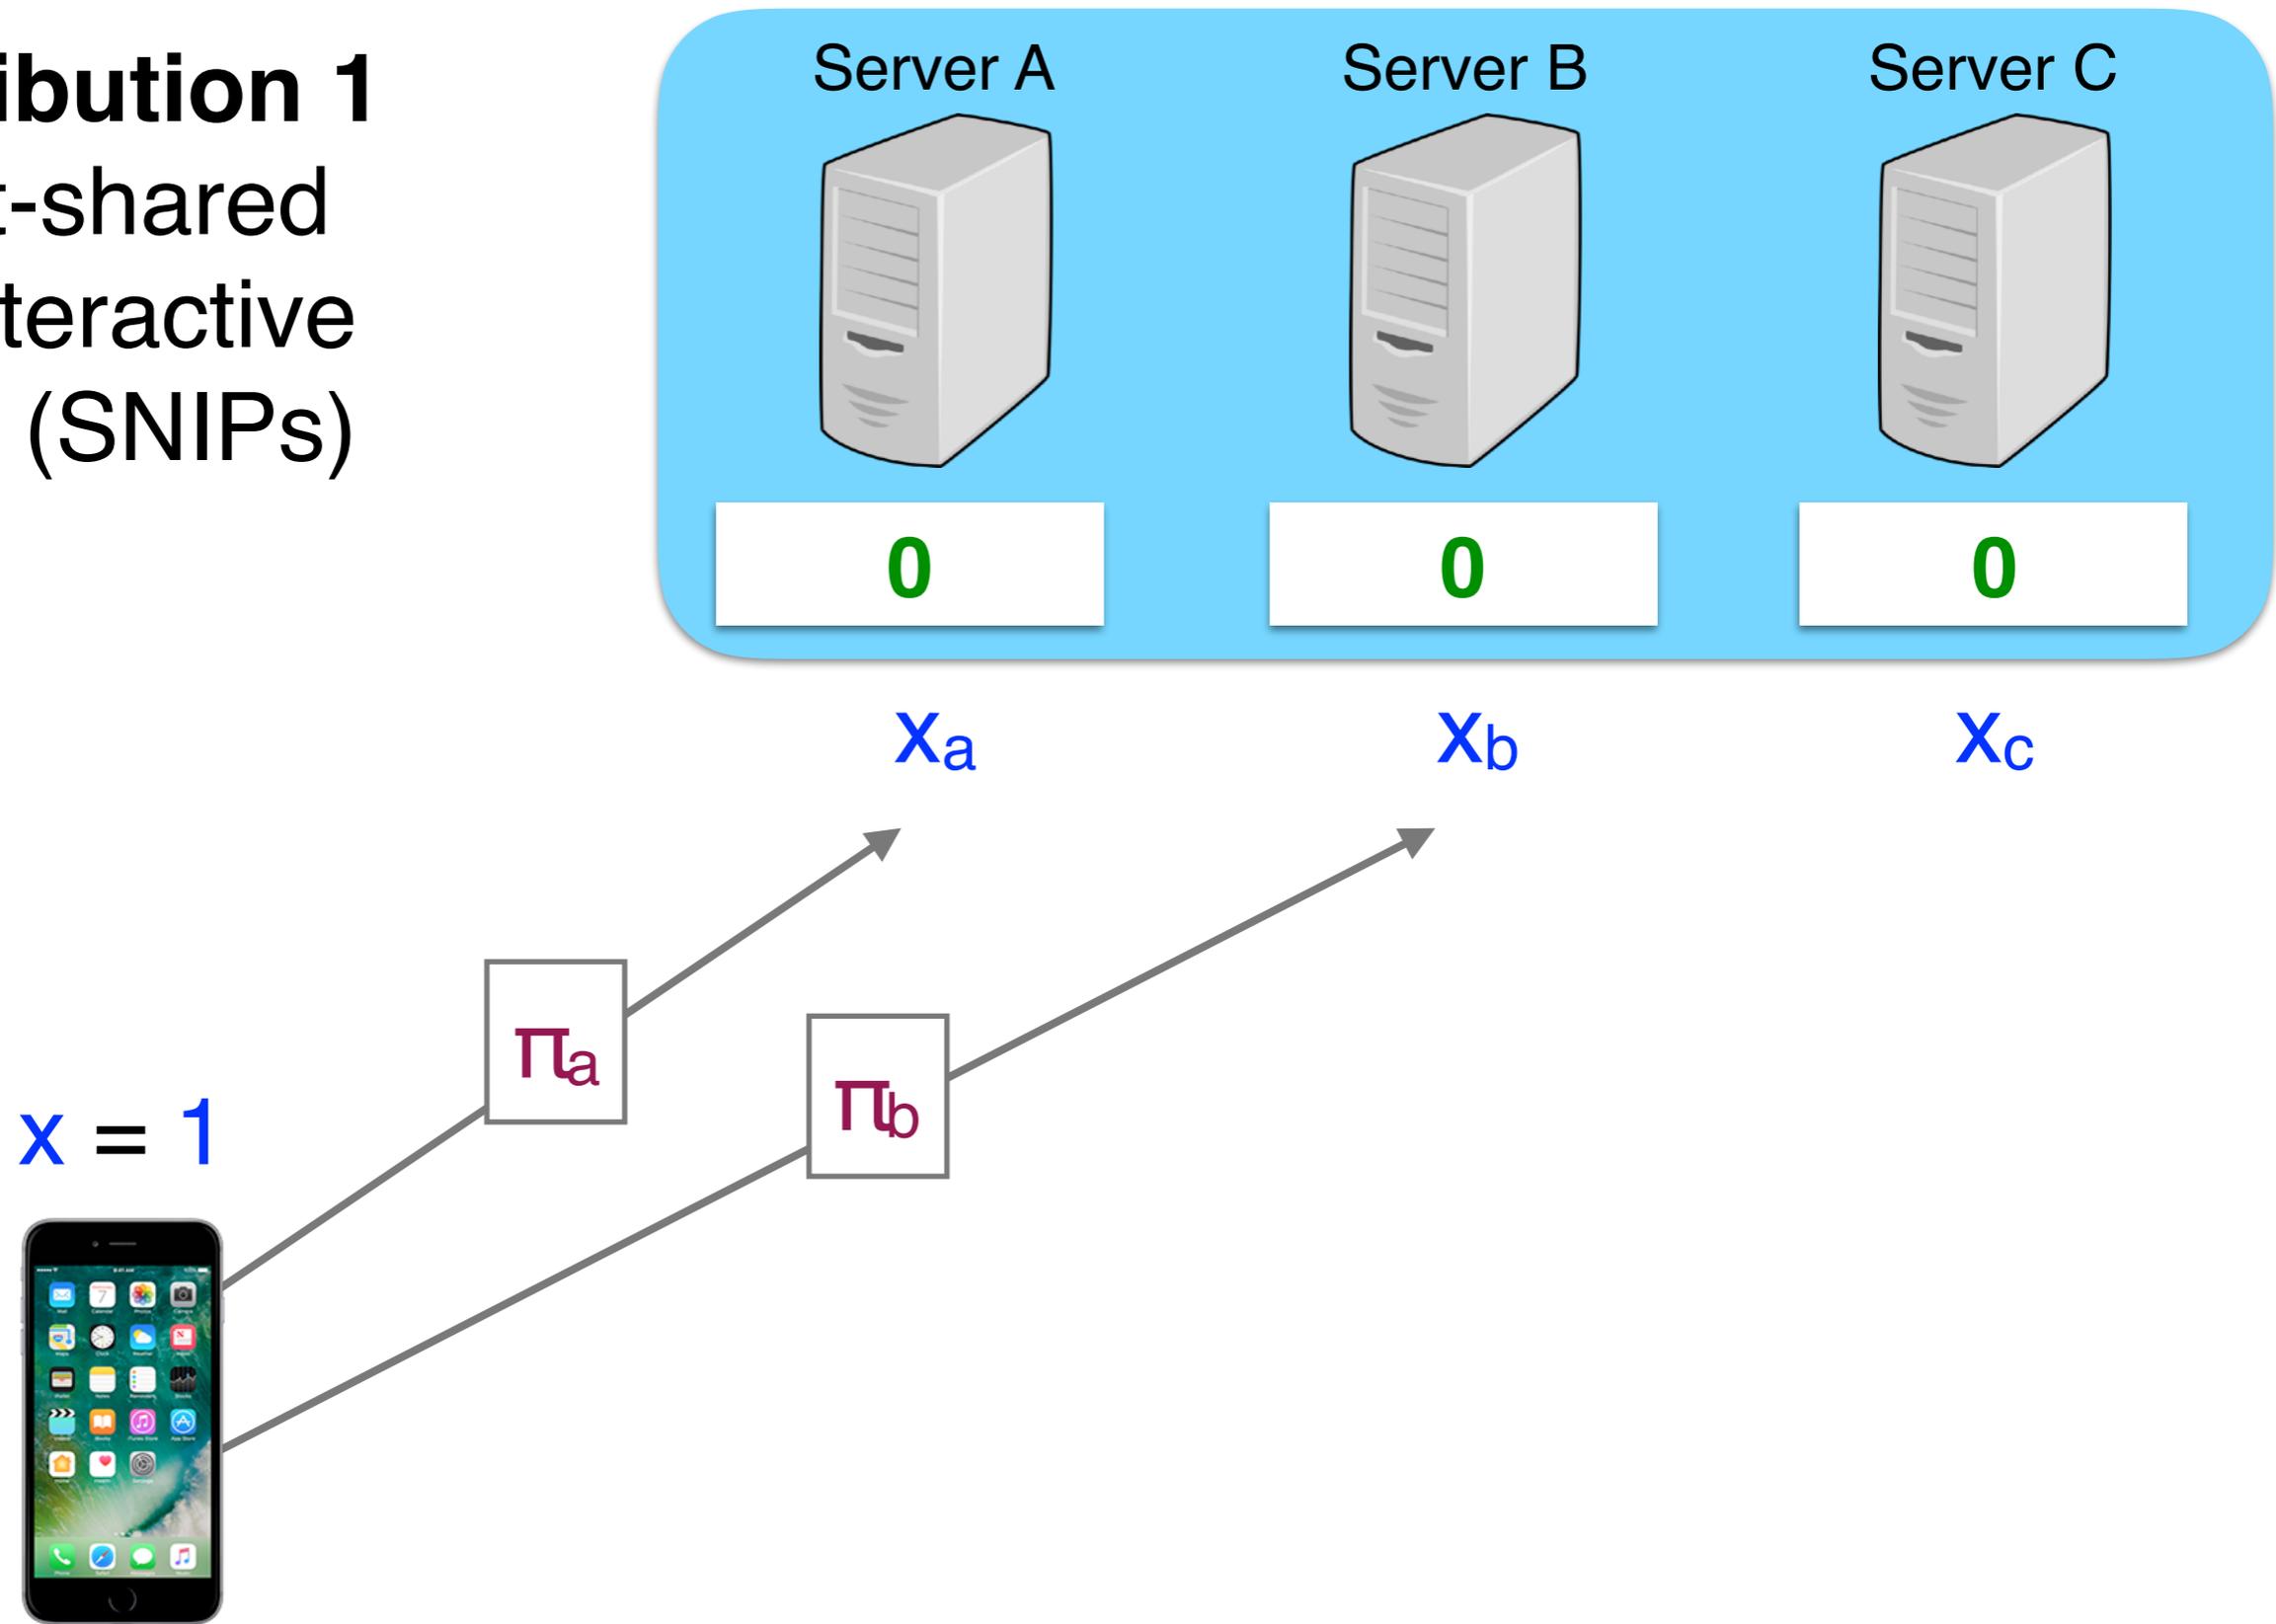
Contribution 1

Secret-shared non-interactive proofs (SNIPs)



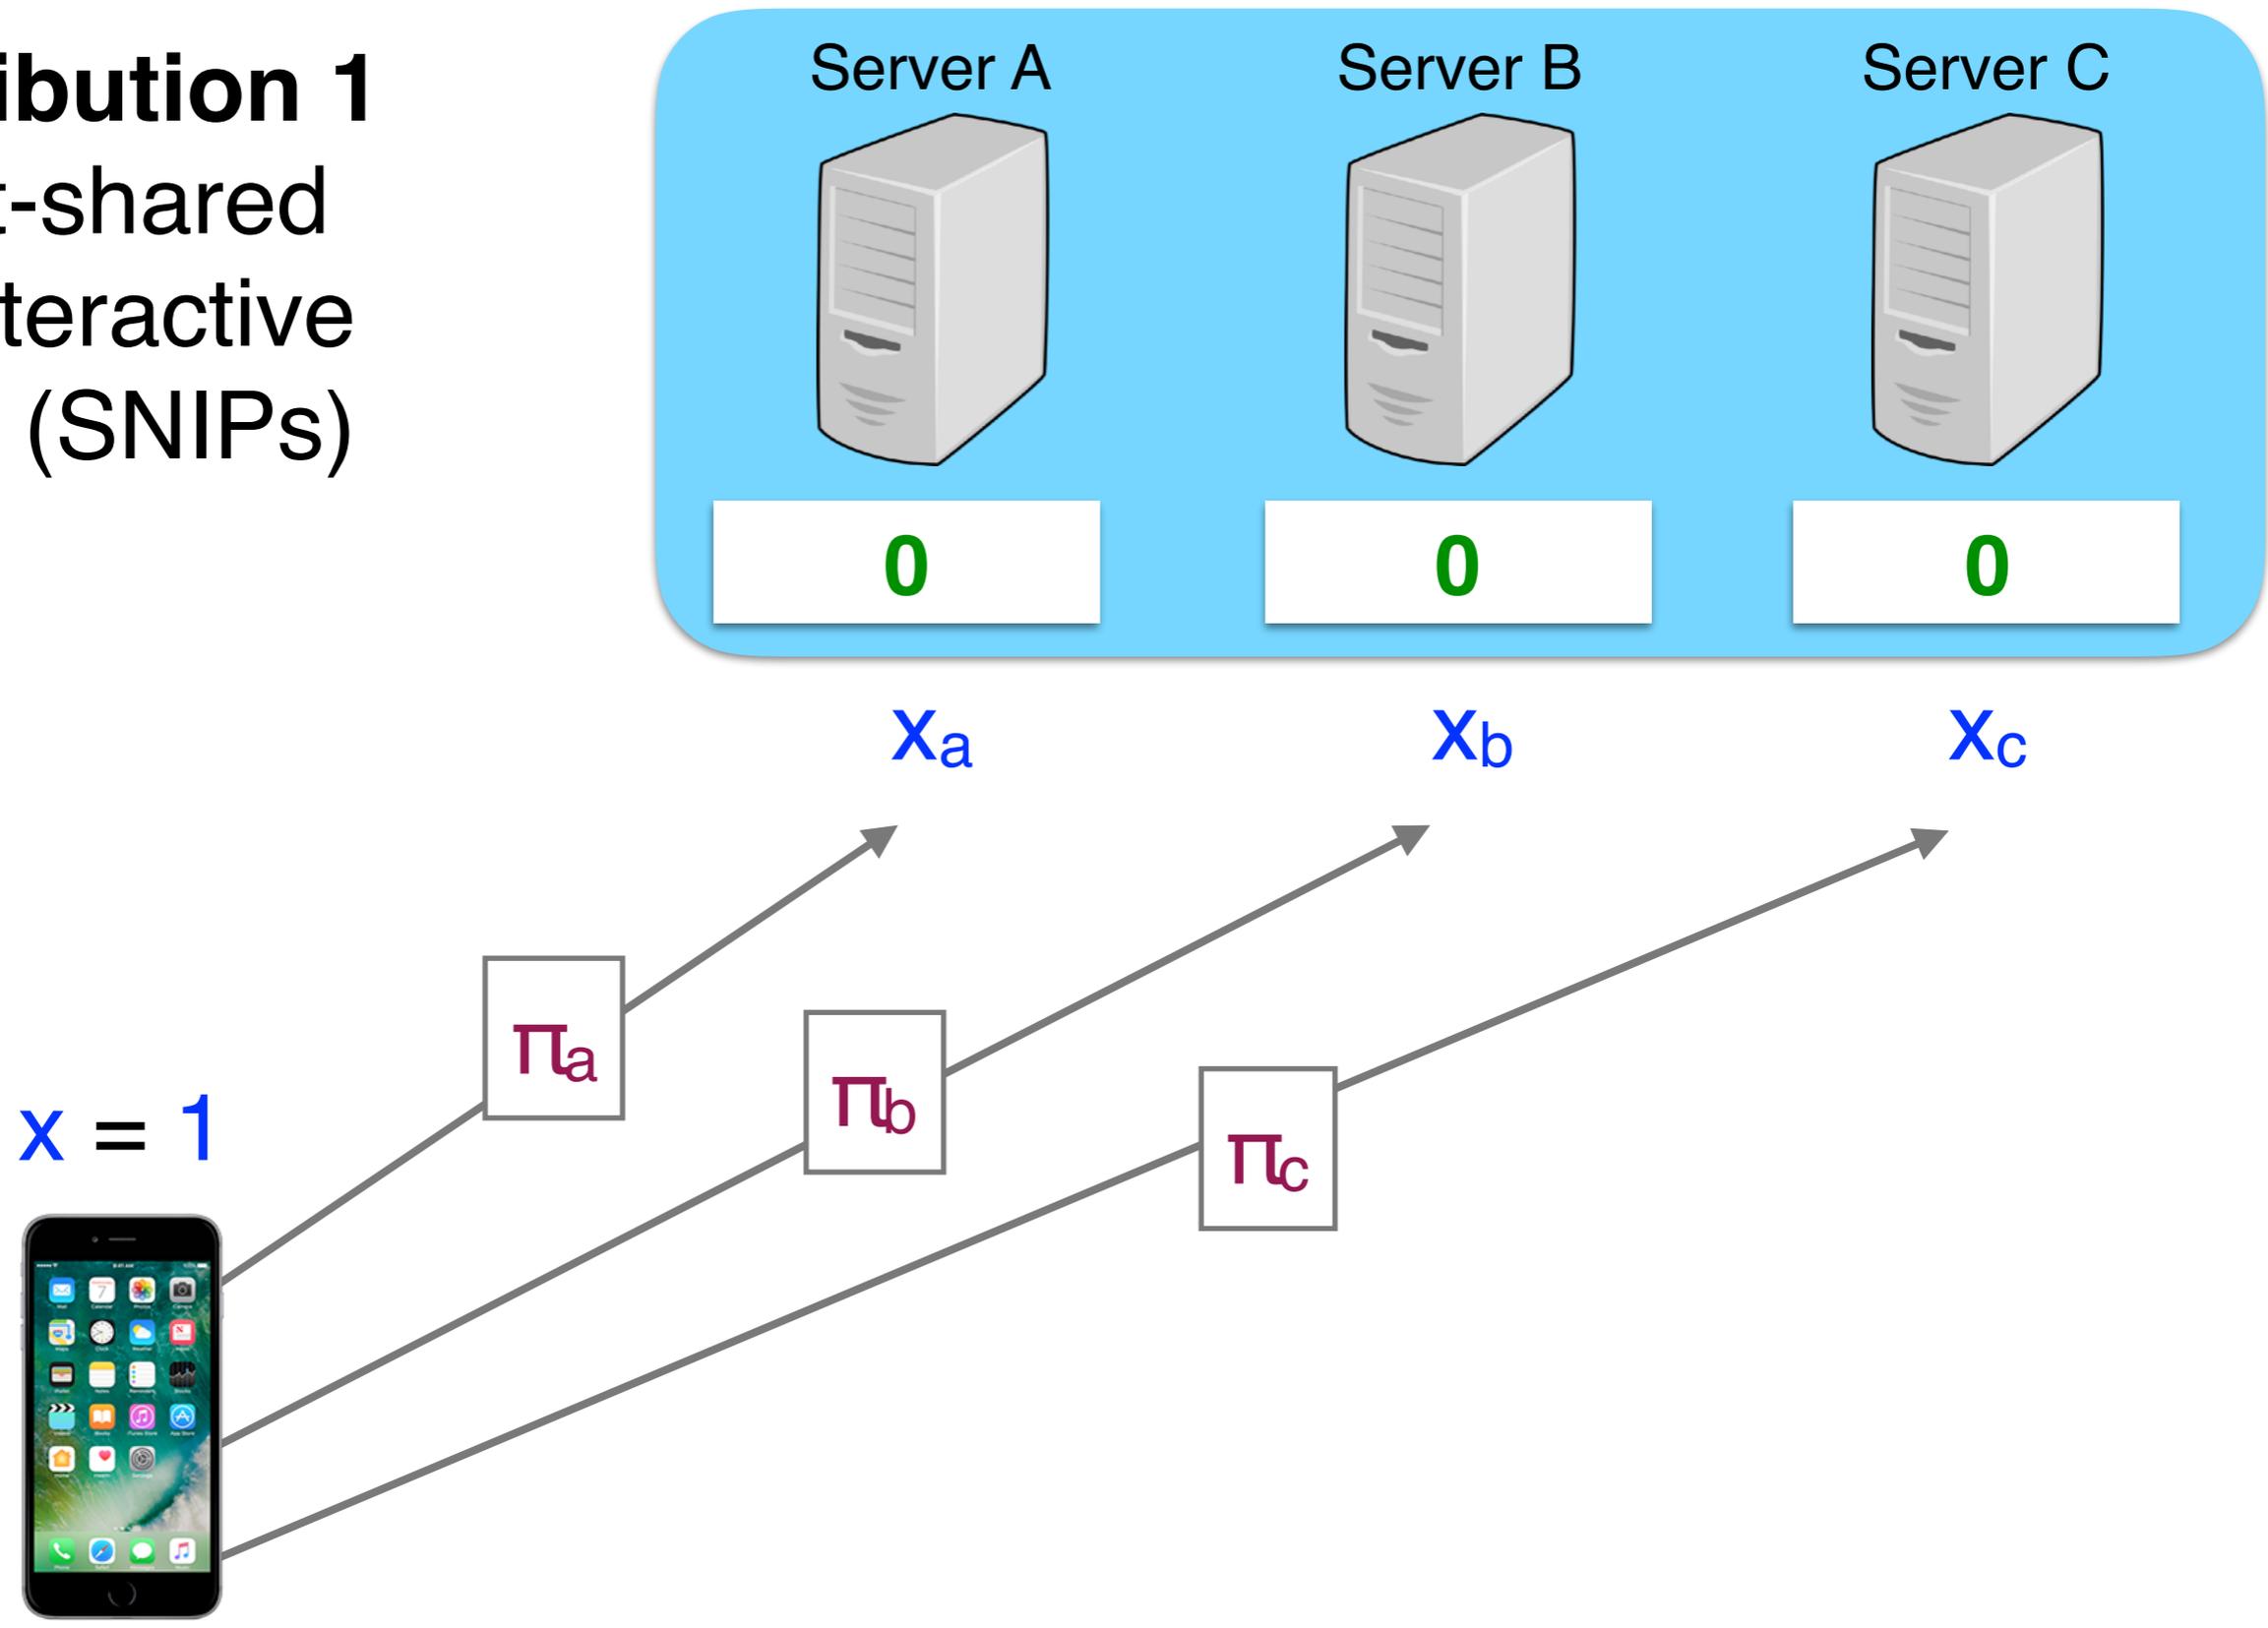
Contribution 1

Secret-shared non-interactive proofs (SNIPs)



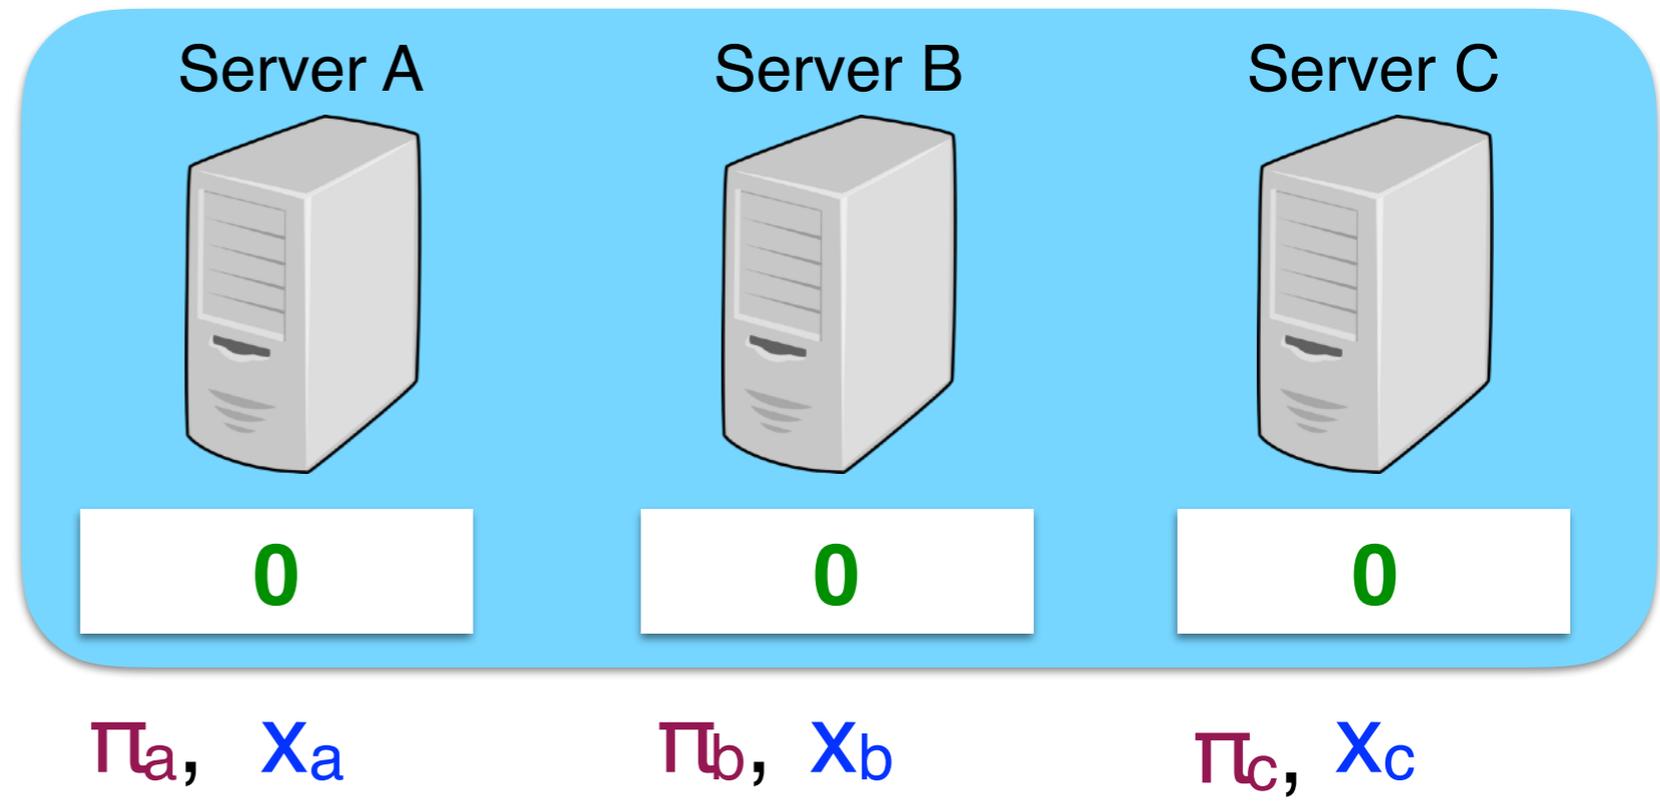
Contribution 1

Secret-shared non-interactive proofs (SNIPs)



Contribution 1

Secret-shared non-interactive proofs (SNIPs)

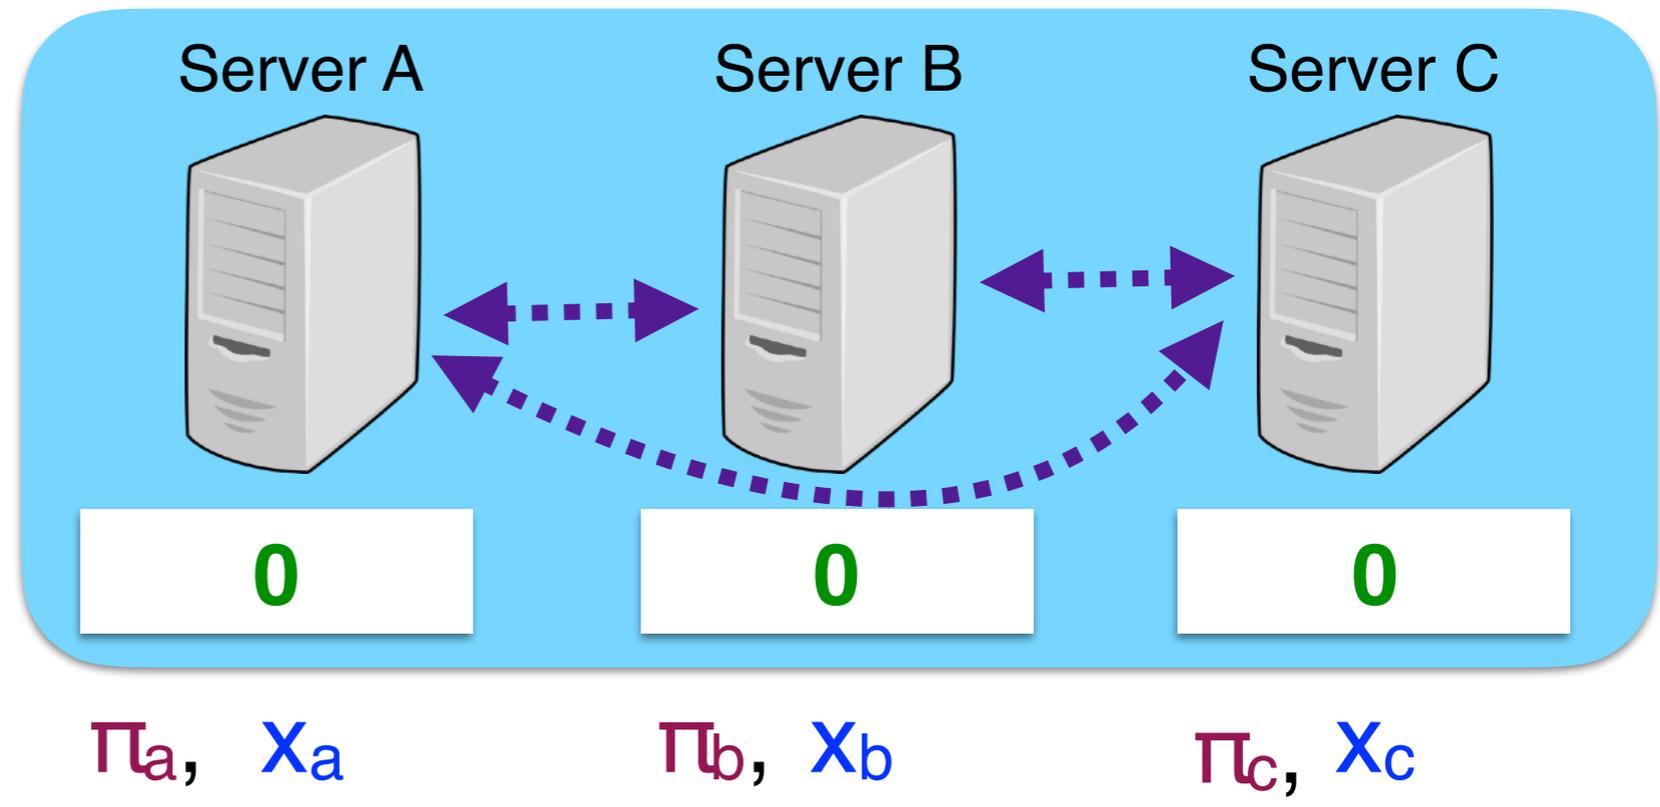


$x = 1$



Contribution 1

Secret-shared non-interactive proofs (SNIPs)

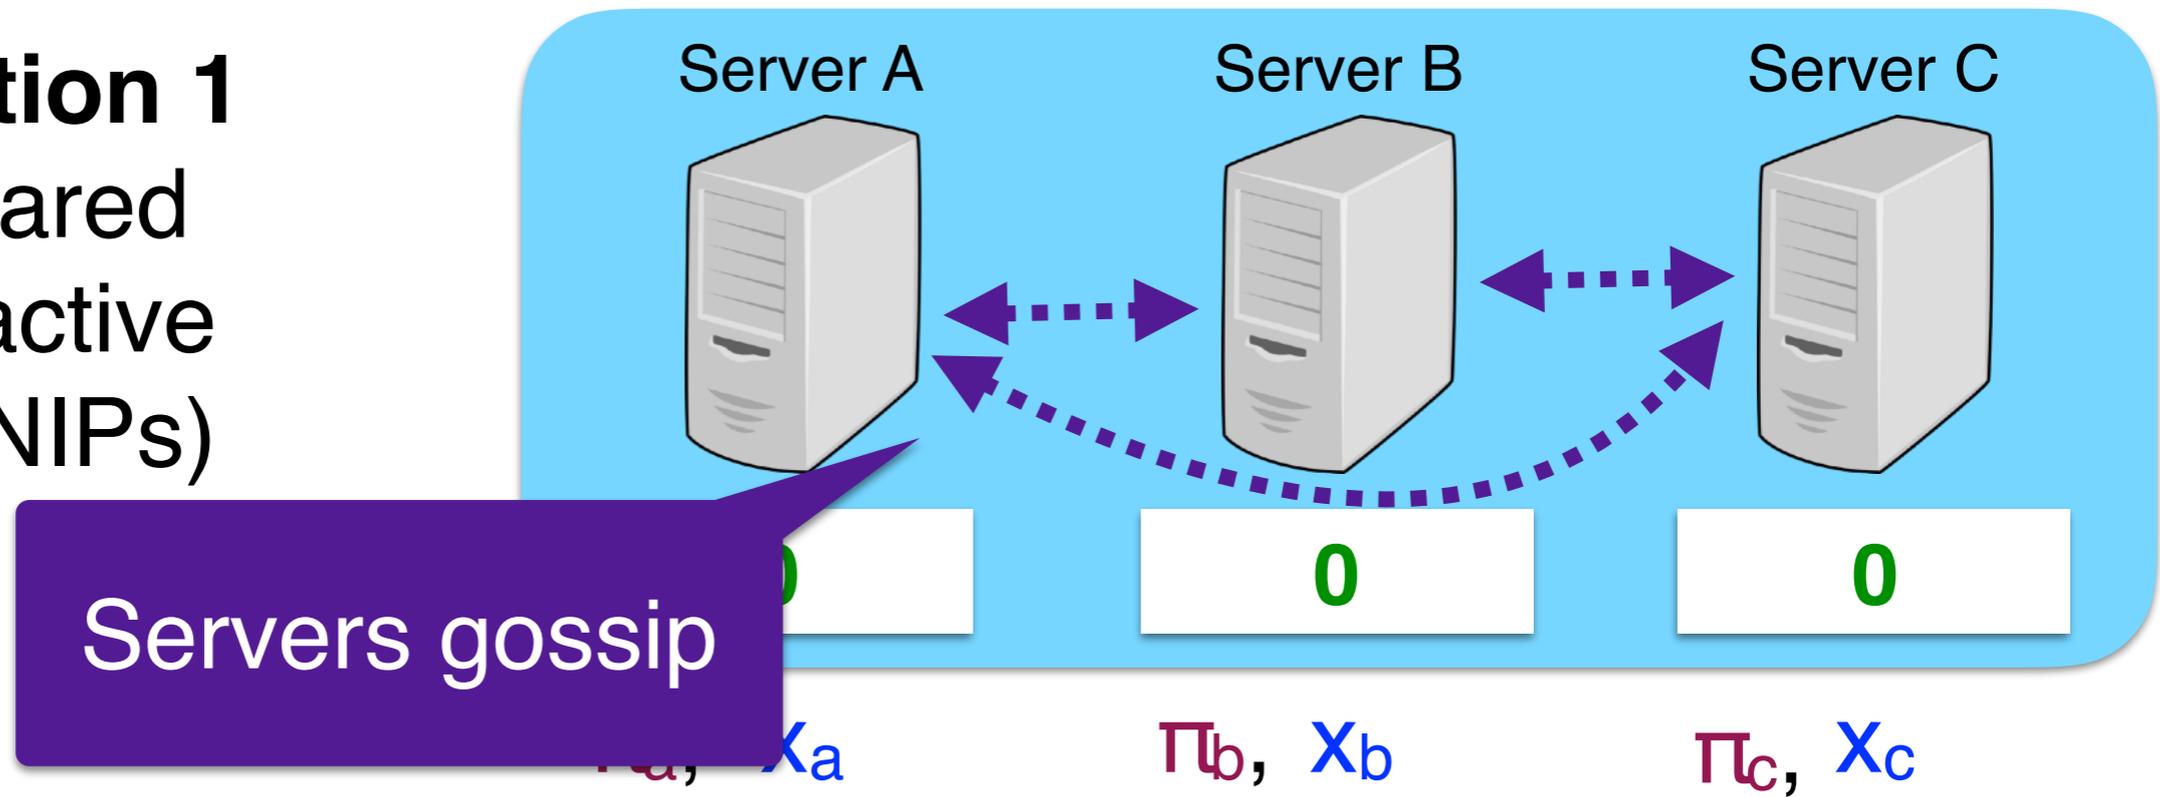


$x = 1$



Contribution 1

Secret-shared
non-interactive
proofs (SNIPs)

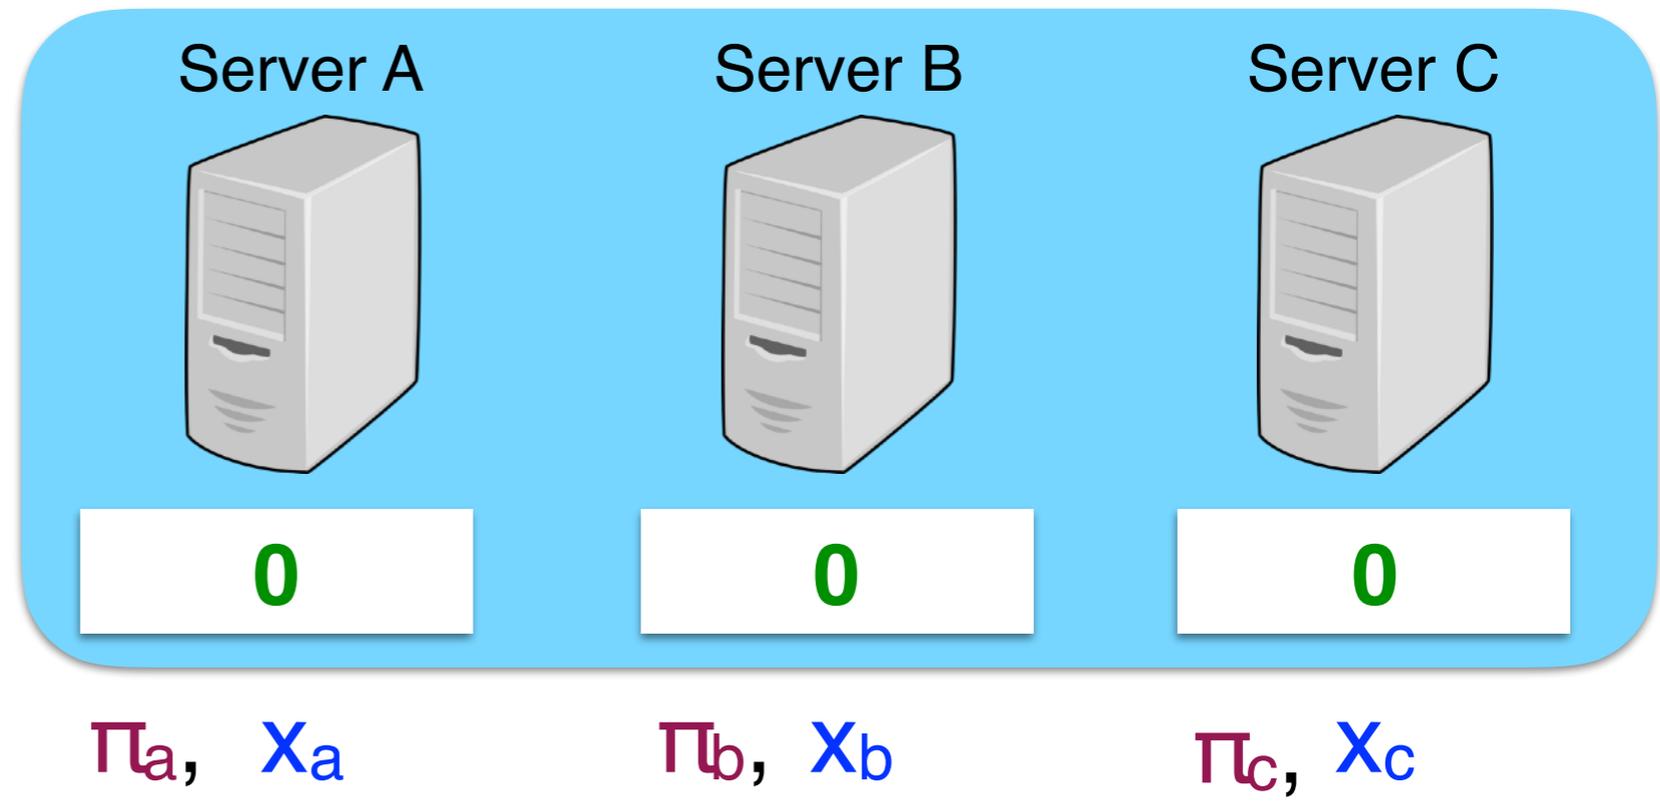


$x = 1$



Contribution 1

Secret-shared non-interactive proofs (SNIPs)

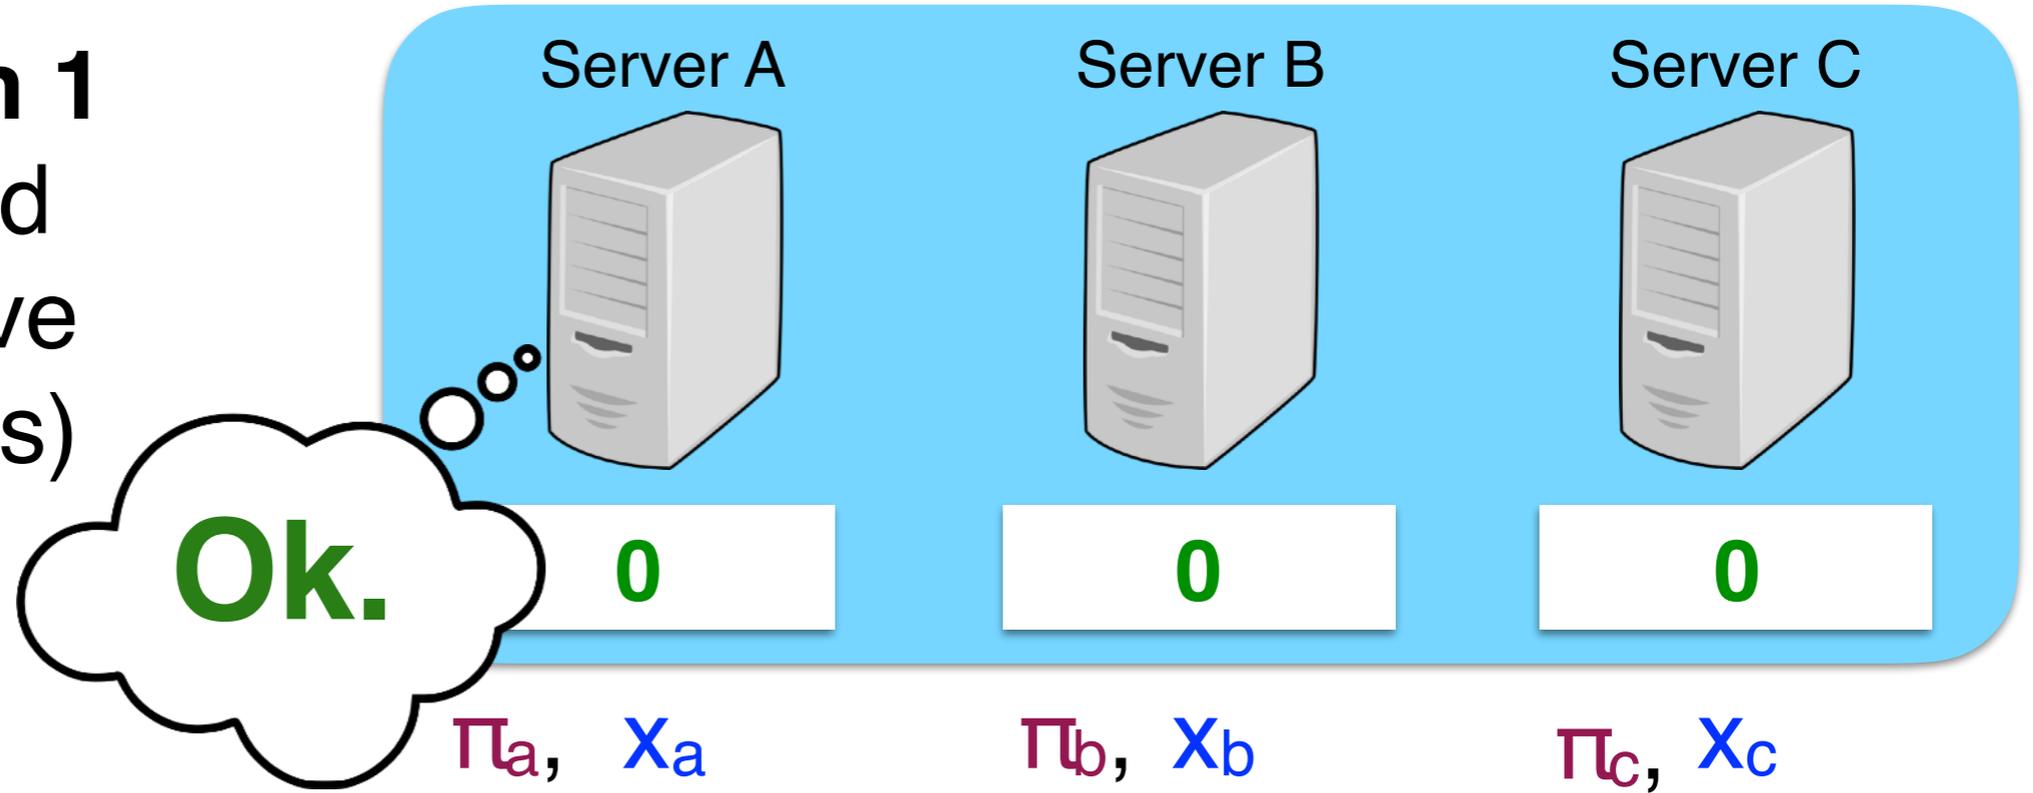


$x = 1$



Contribution 1

Secret-shared
non-interactive
proofs (SNIPs)

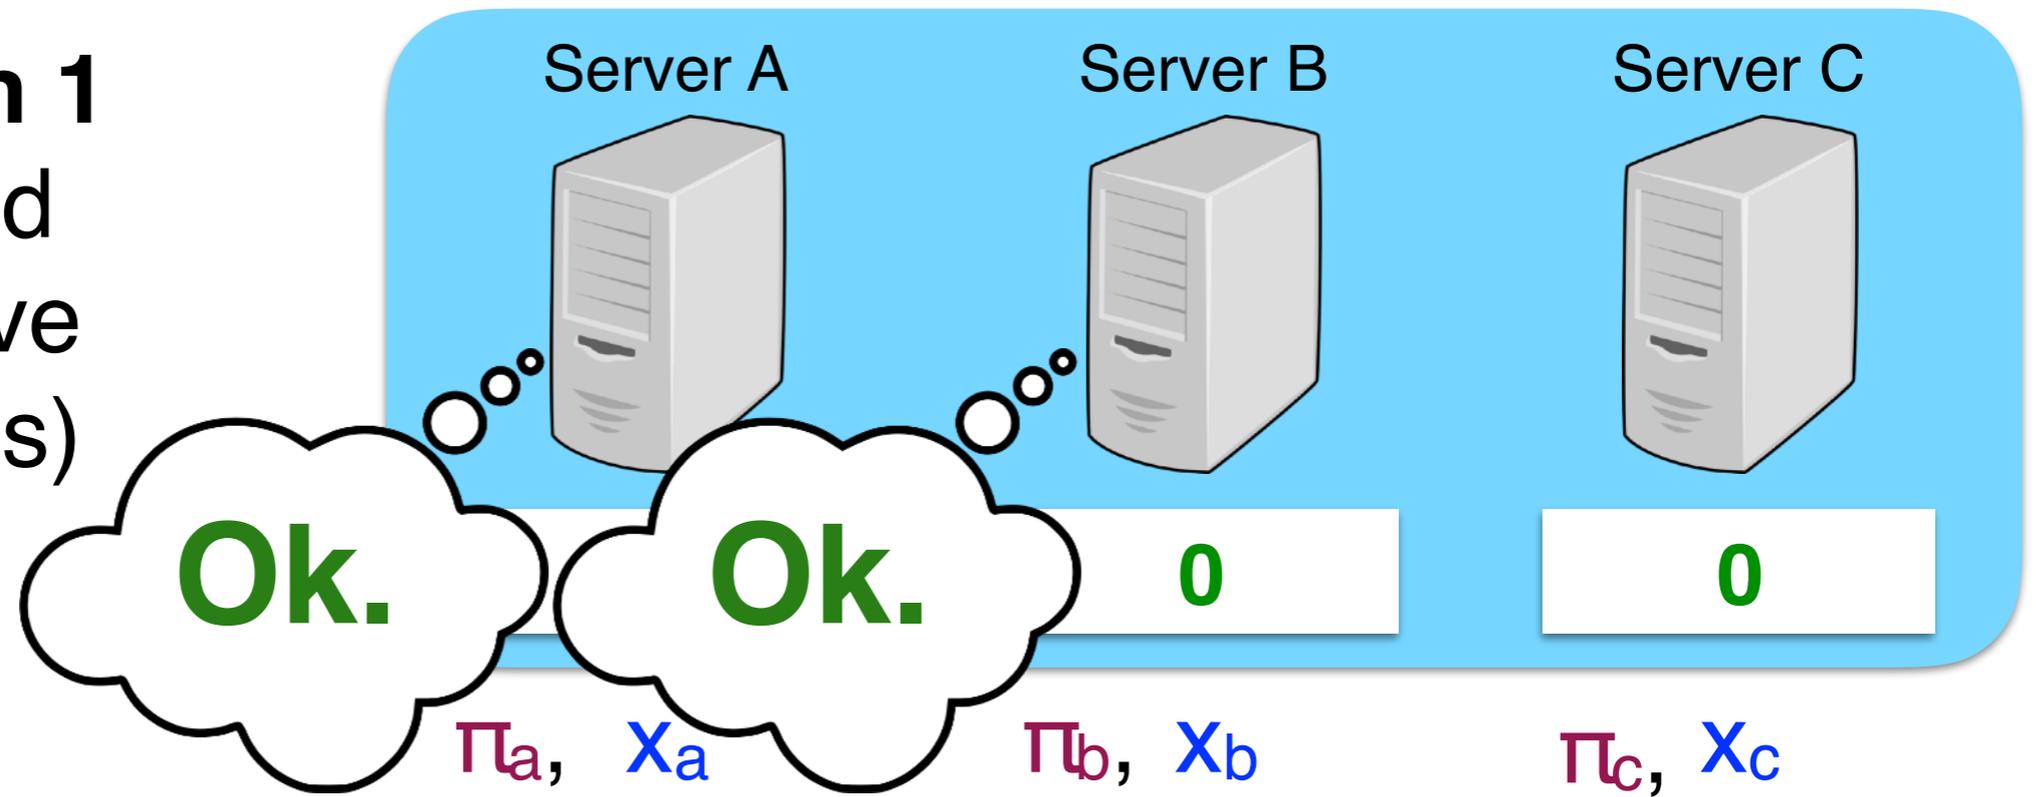


$x = 1$



Contribution 1

Secret-shared
non-interactive
proofs (SNIPs)

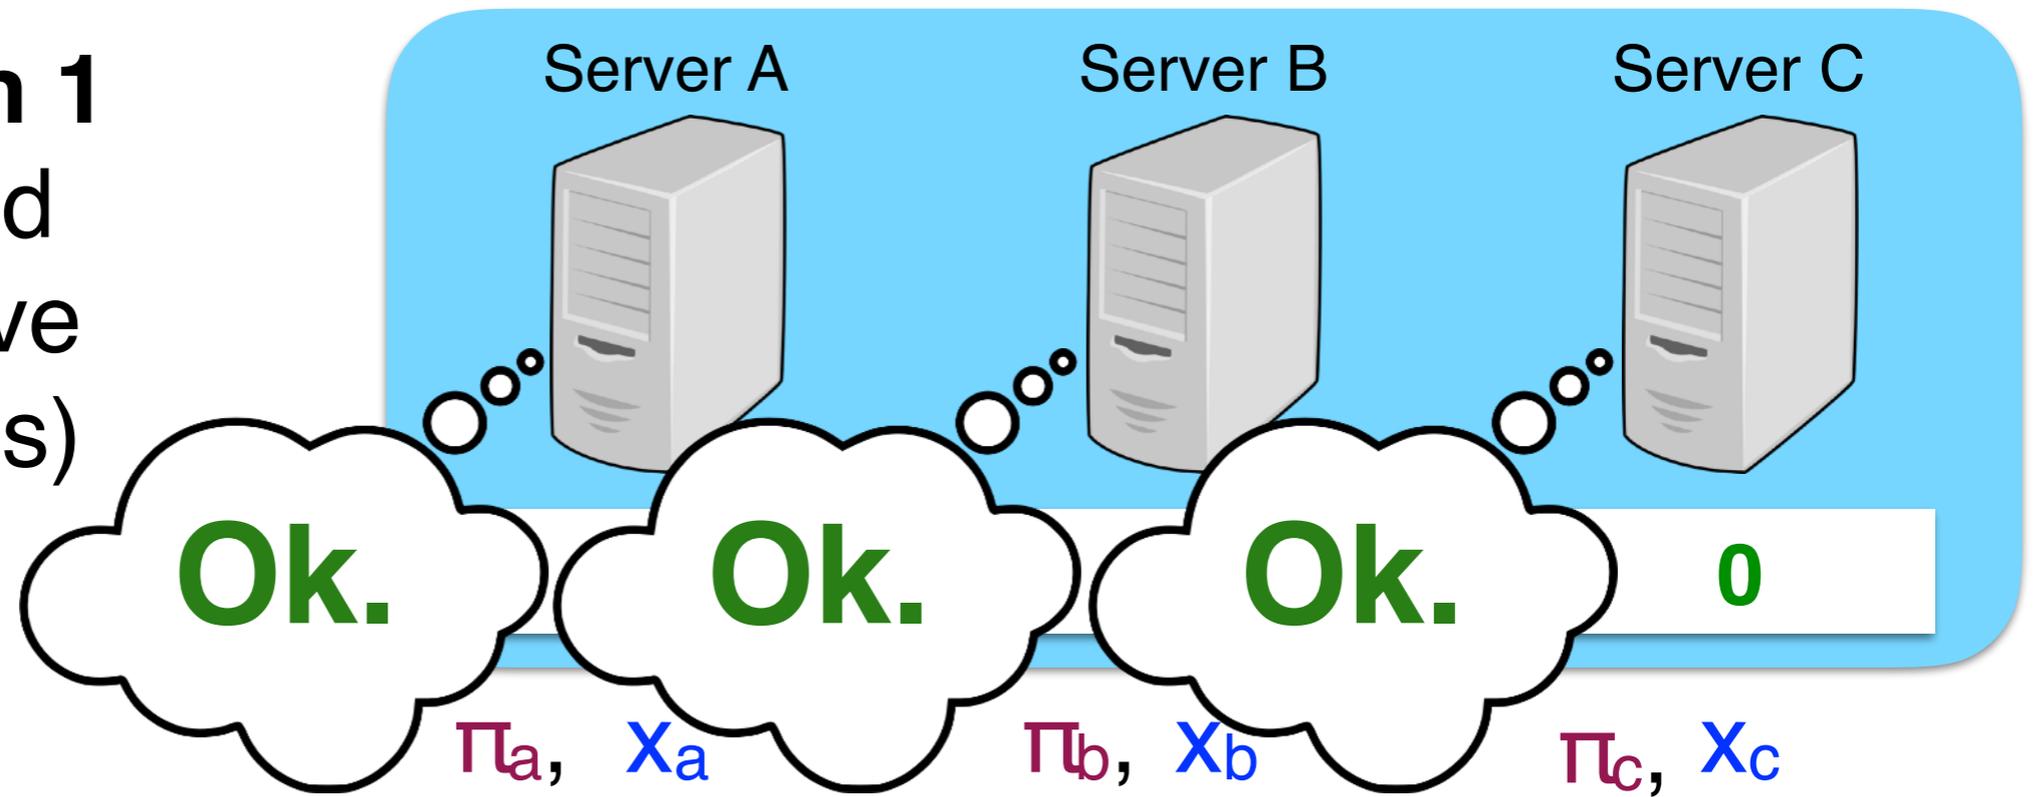


$$x = 1$$



Contribution 1

Secret-shared
non-interactive
proofs (SNIPs)

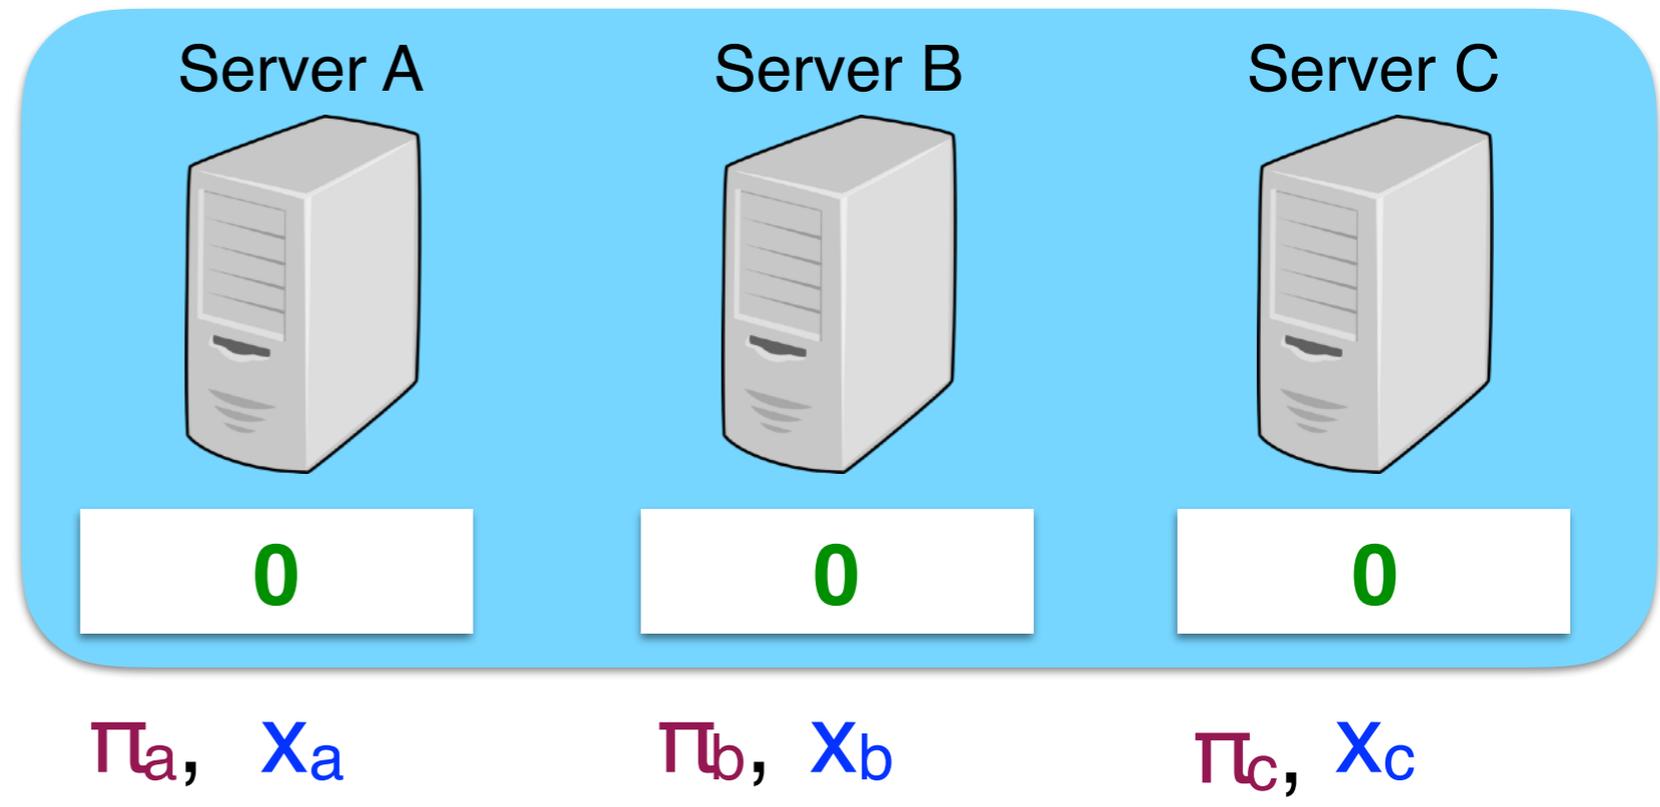


$$x = 1$$



Contribution 1

Secret-shared non-interactive proofs (SNIPs)

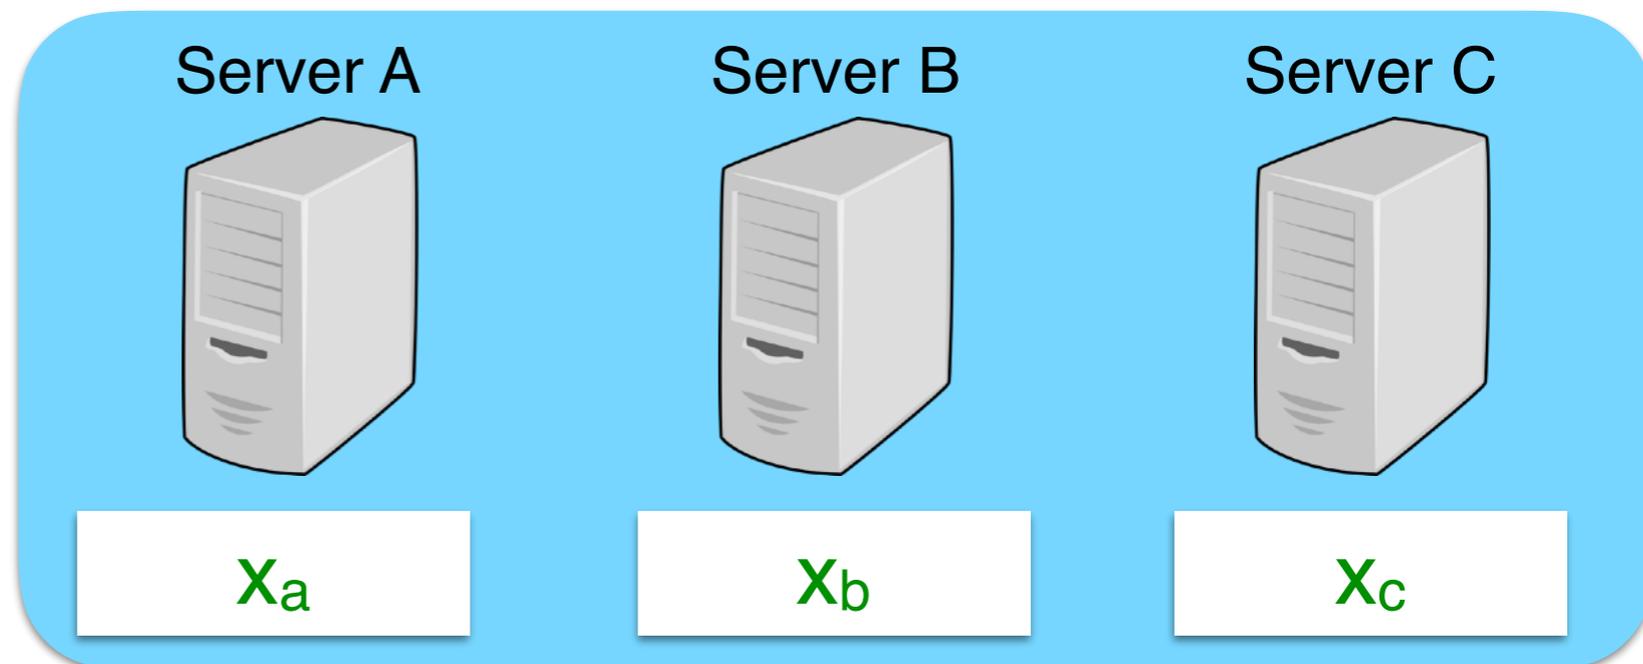


$x = 1$



Contribution 1

Secret-shared non-interactive proofs (SNIPs)

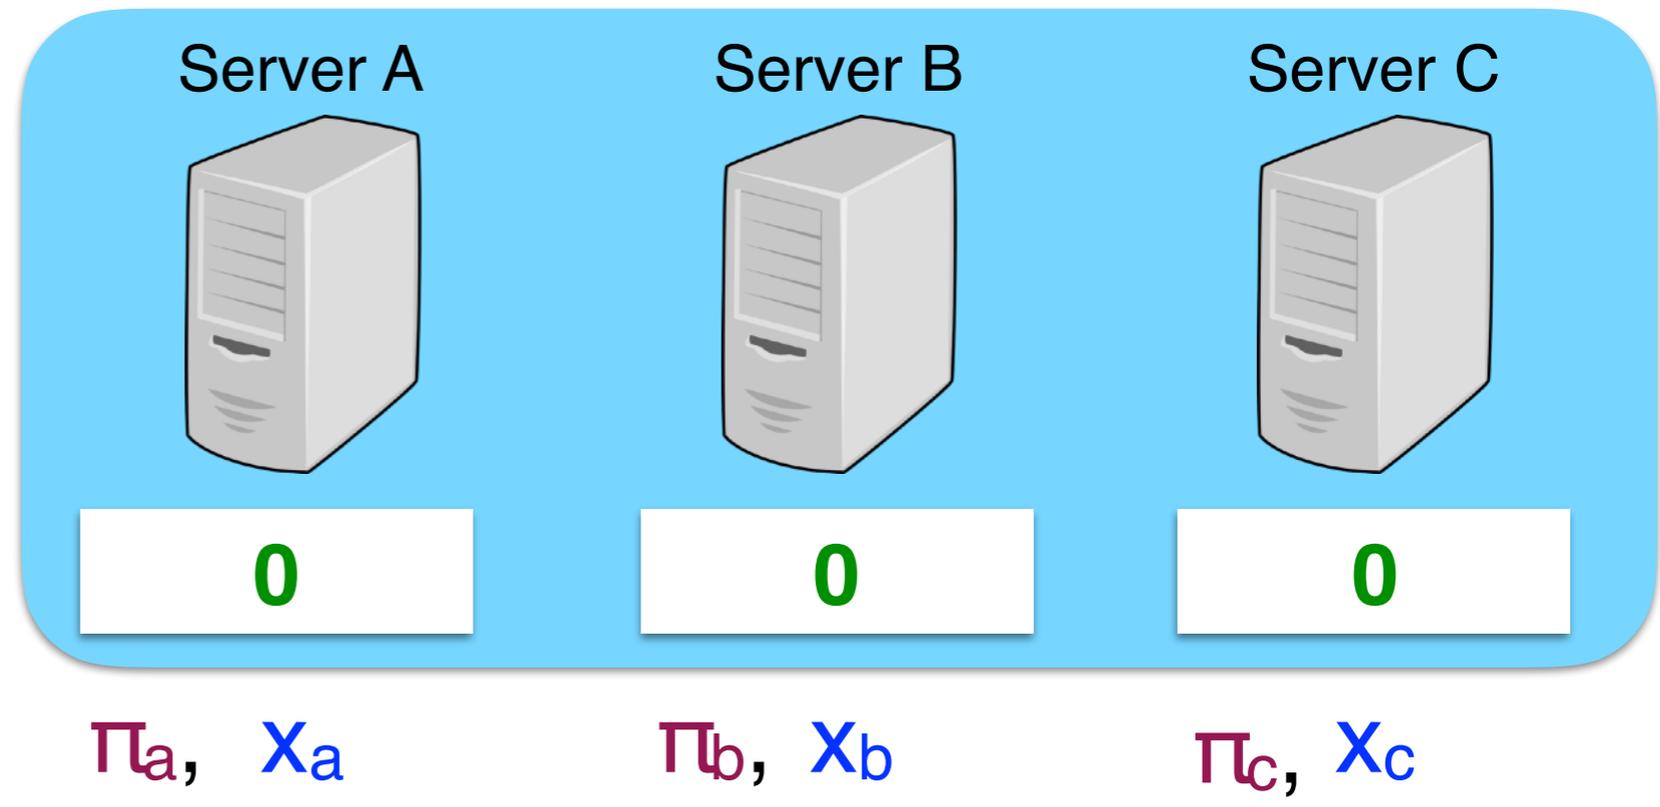


$$x = 1$$



Contribution 1

Secret-shared non-interactive proofs (SNIPs)

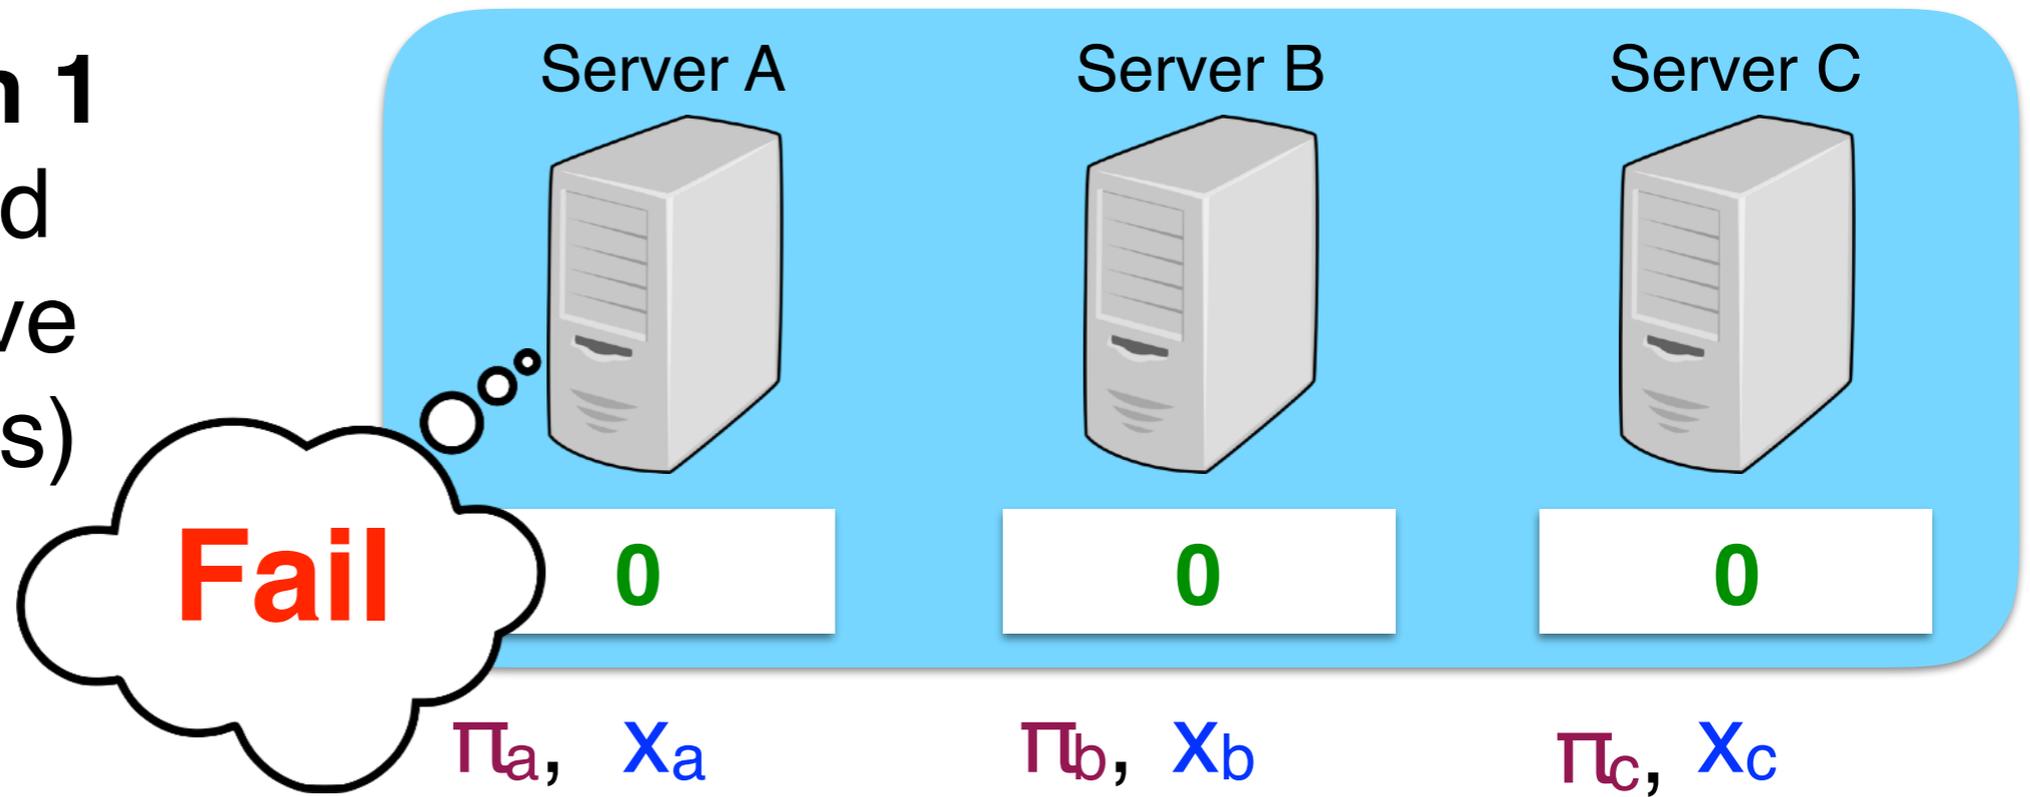


$x = 1$



Contribution 1

Secret-shared
non-interactive
proofs (SNIPs)

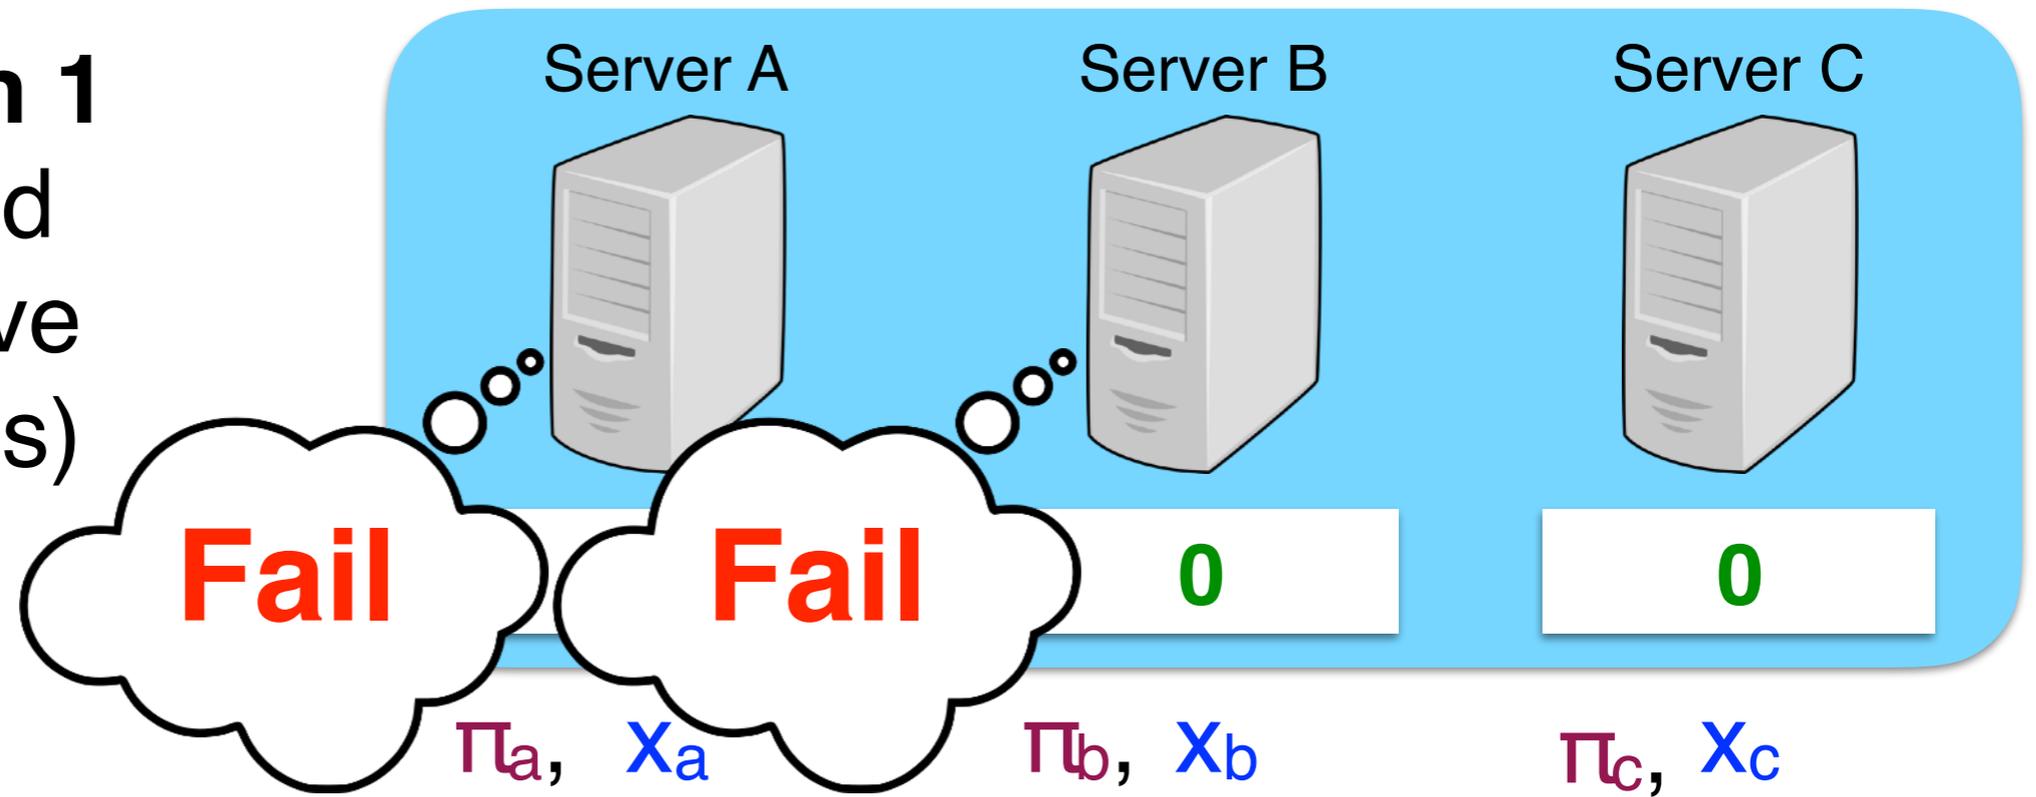


$x = 1$



Contribution 1

Secret-shared
non-interactive
proofs (SNIPs)

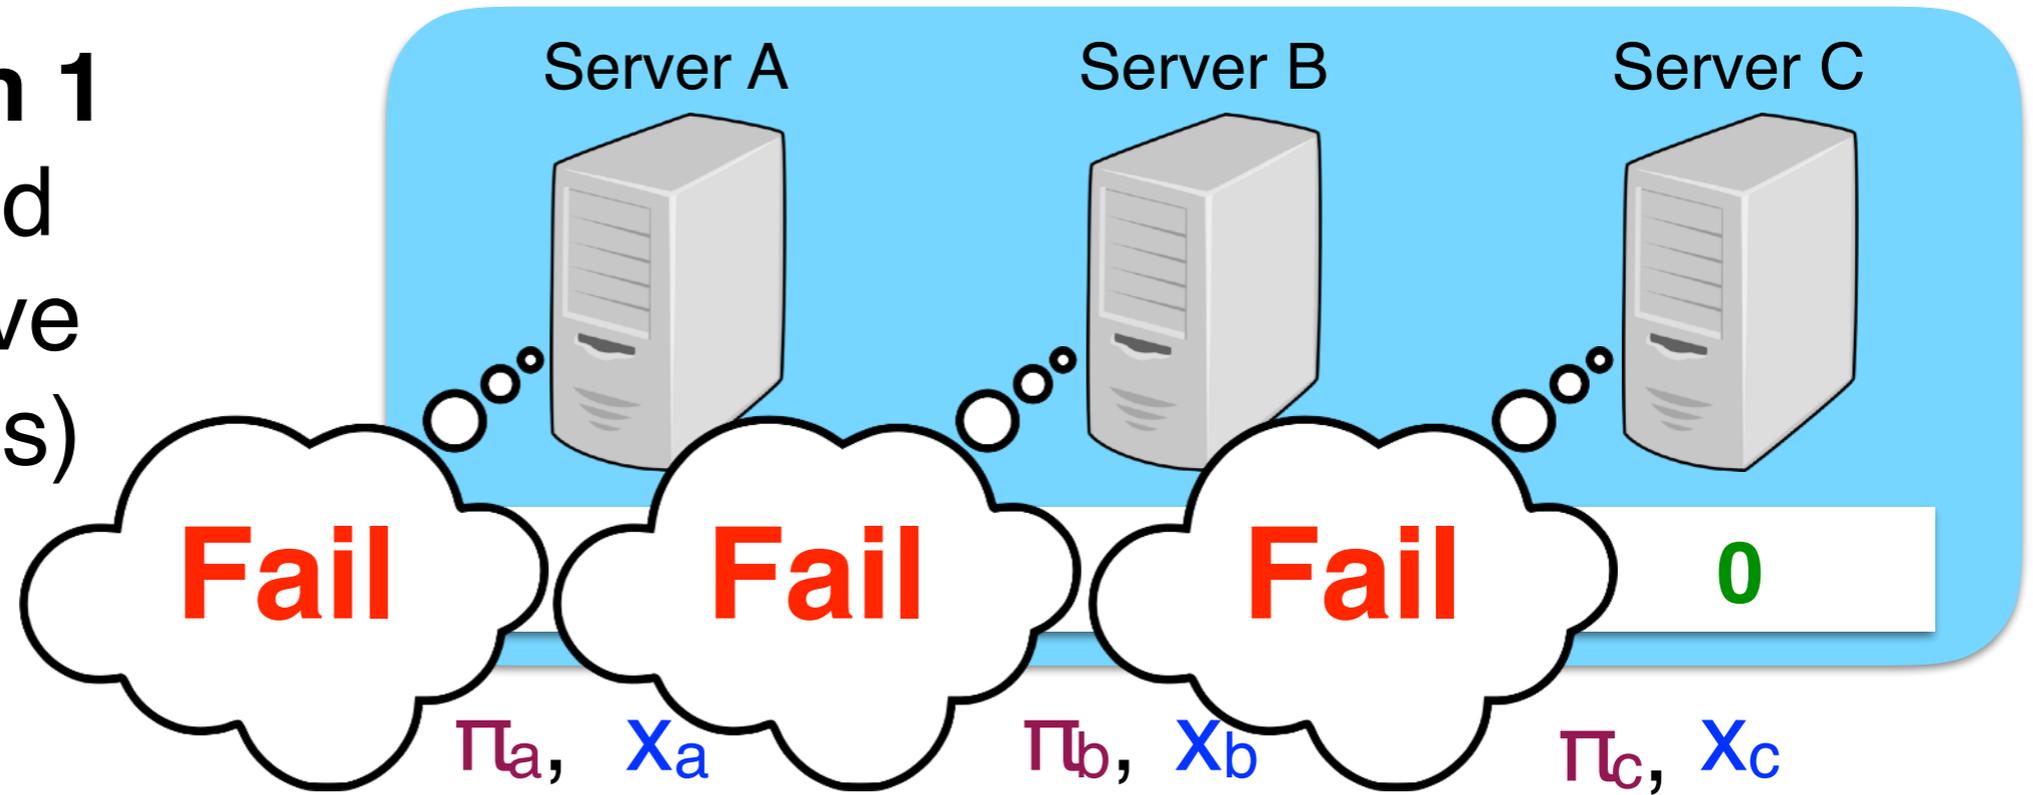


$x = 1$



Contribution 1

Secret-shared
non-interactive
proofs (SNIPs)

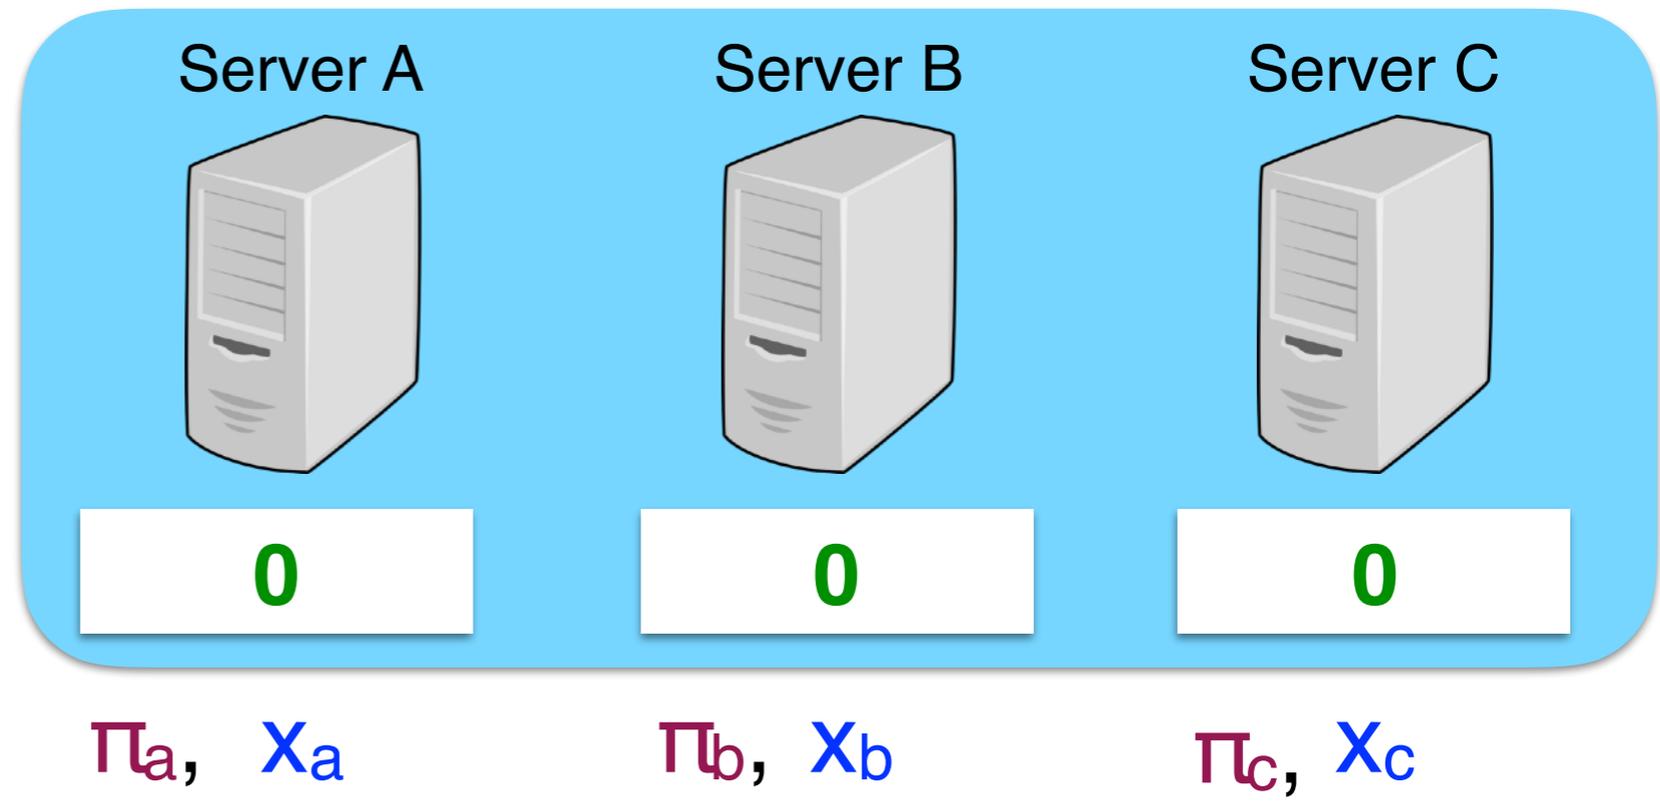


$$x = 1$$



Contribution 1

Secret-shared non-interactive proofs (SNIPs)

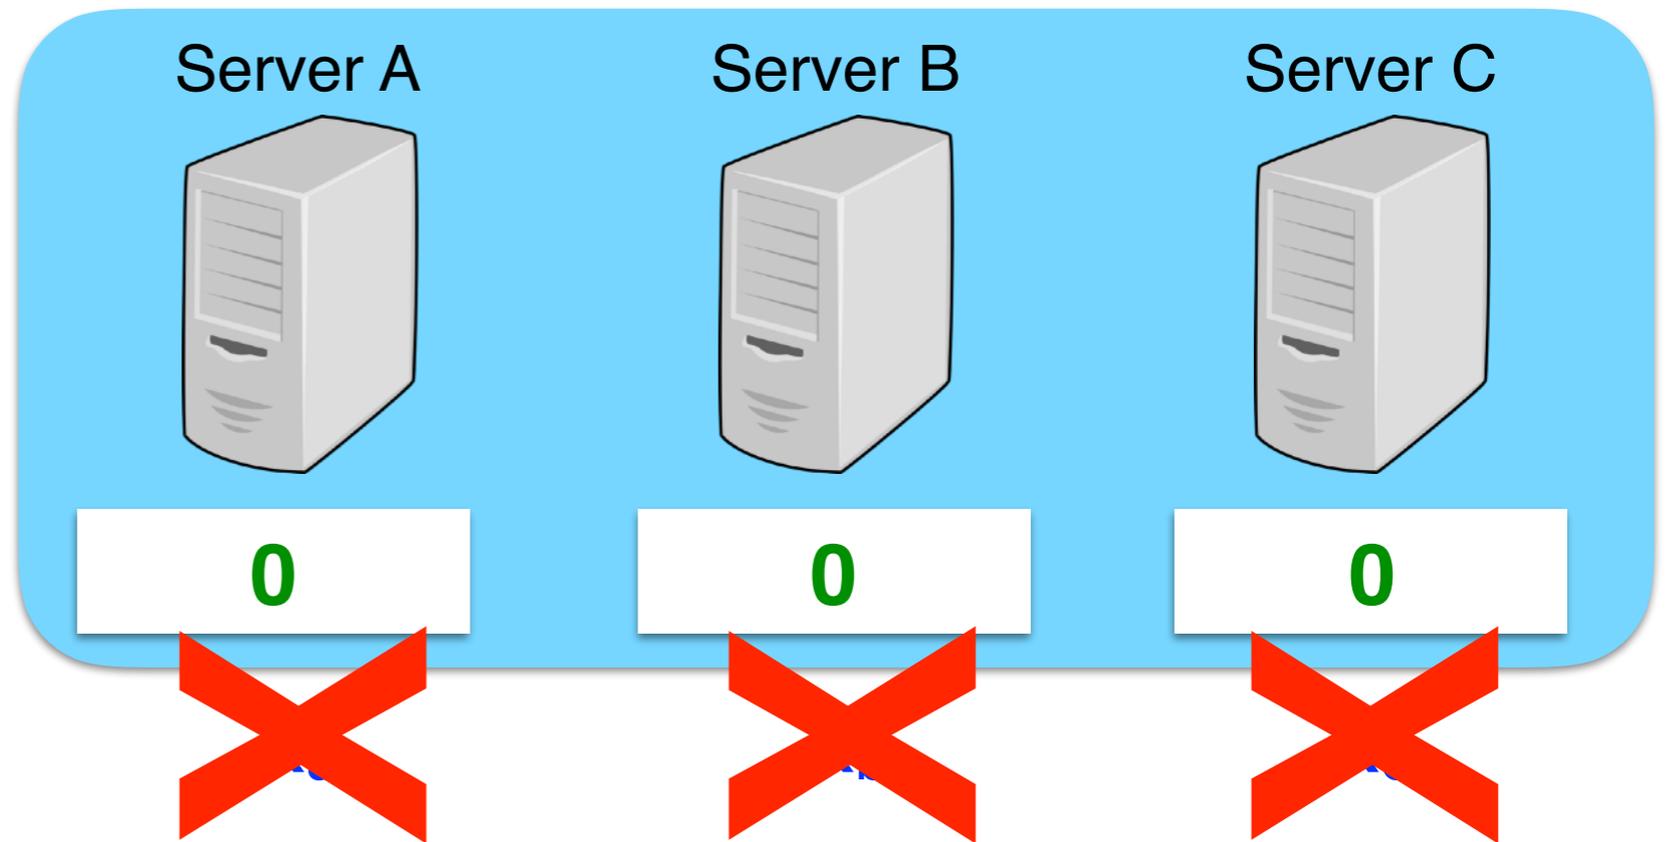


$x = 1$



Contribution 1

Secret-shared non-interactive proofs (SNIPs)

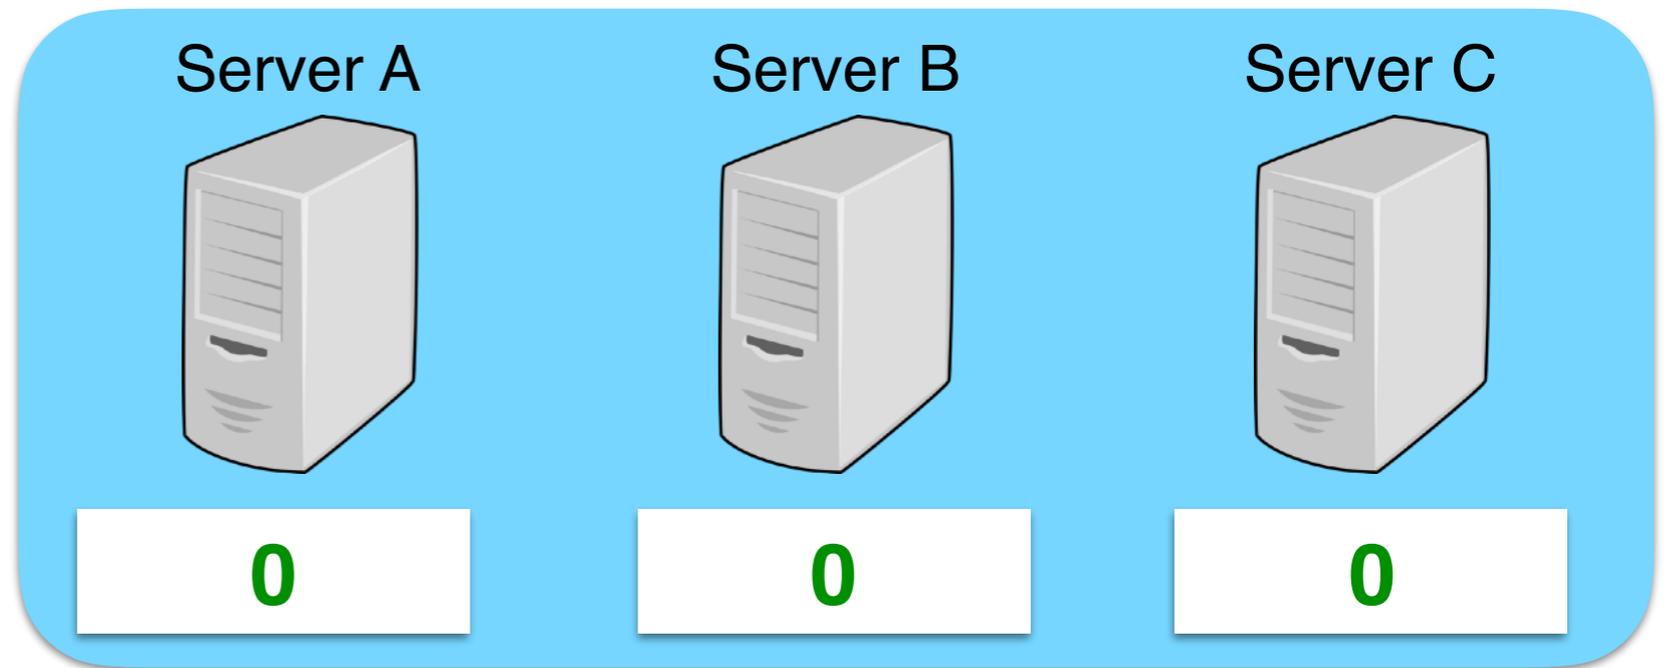


$$x = 1$$



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

How SNIPs work



x_a

x_b

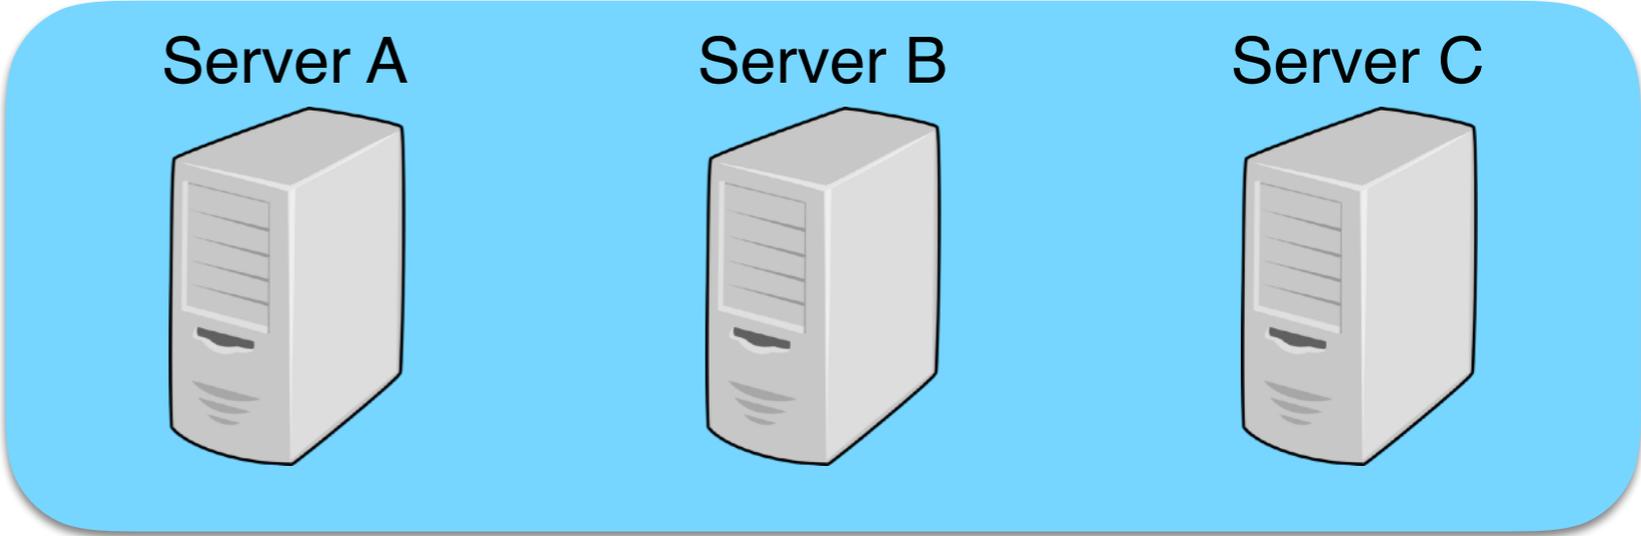
x_c

The servers want to ensure that
 $\text{Valid}(x) = \text{Valid}(x_a + x_b + x_c) = 1$
...without learning x .

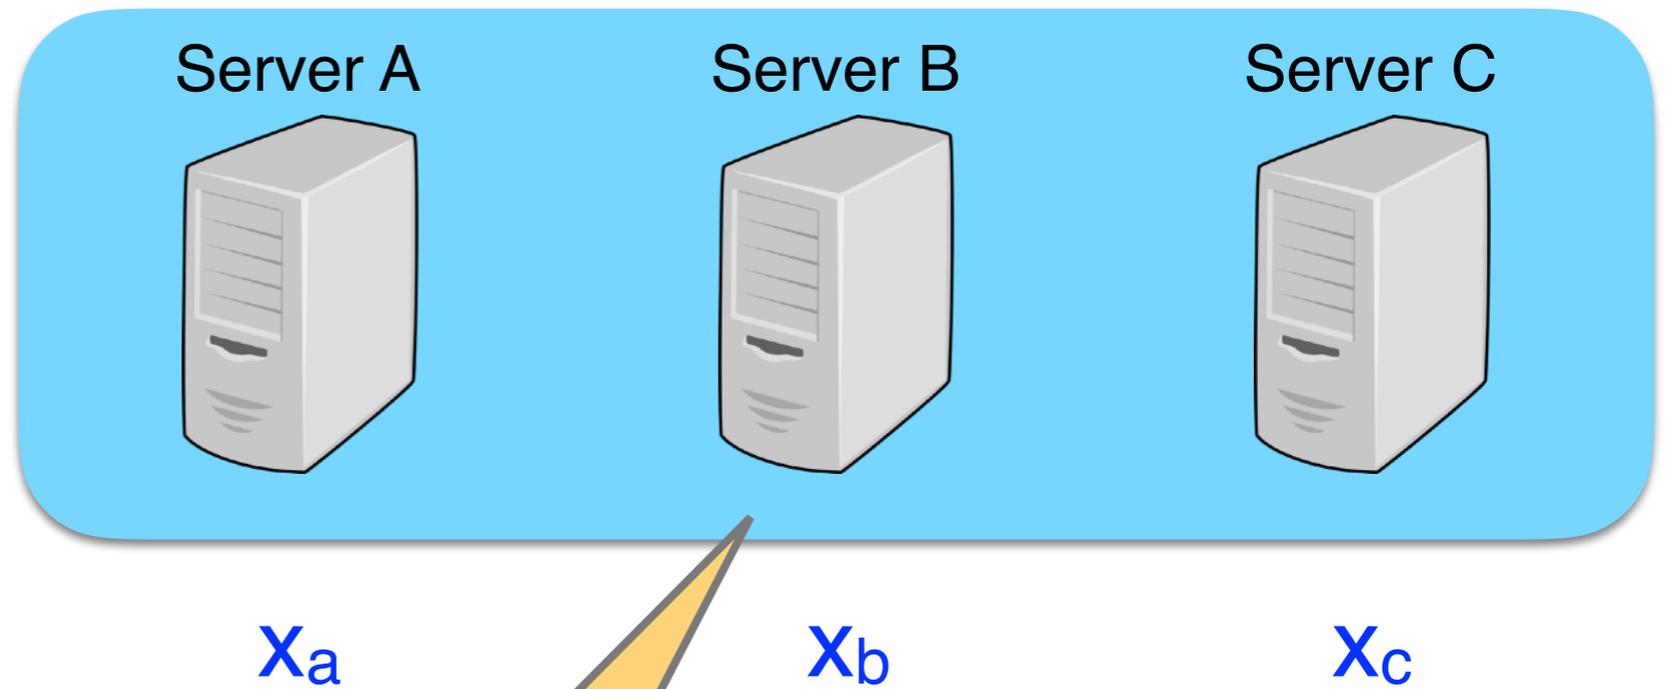
$x = 1$



How SNIPs work

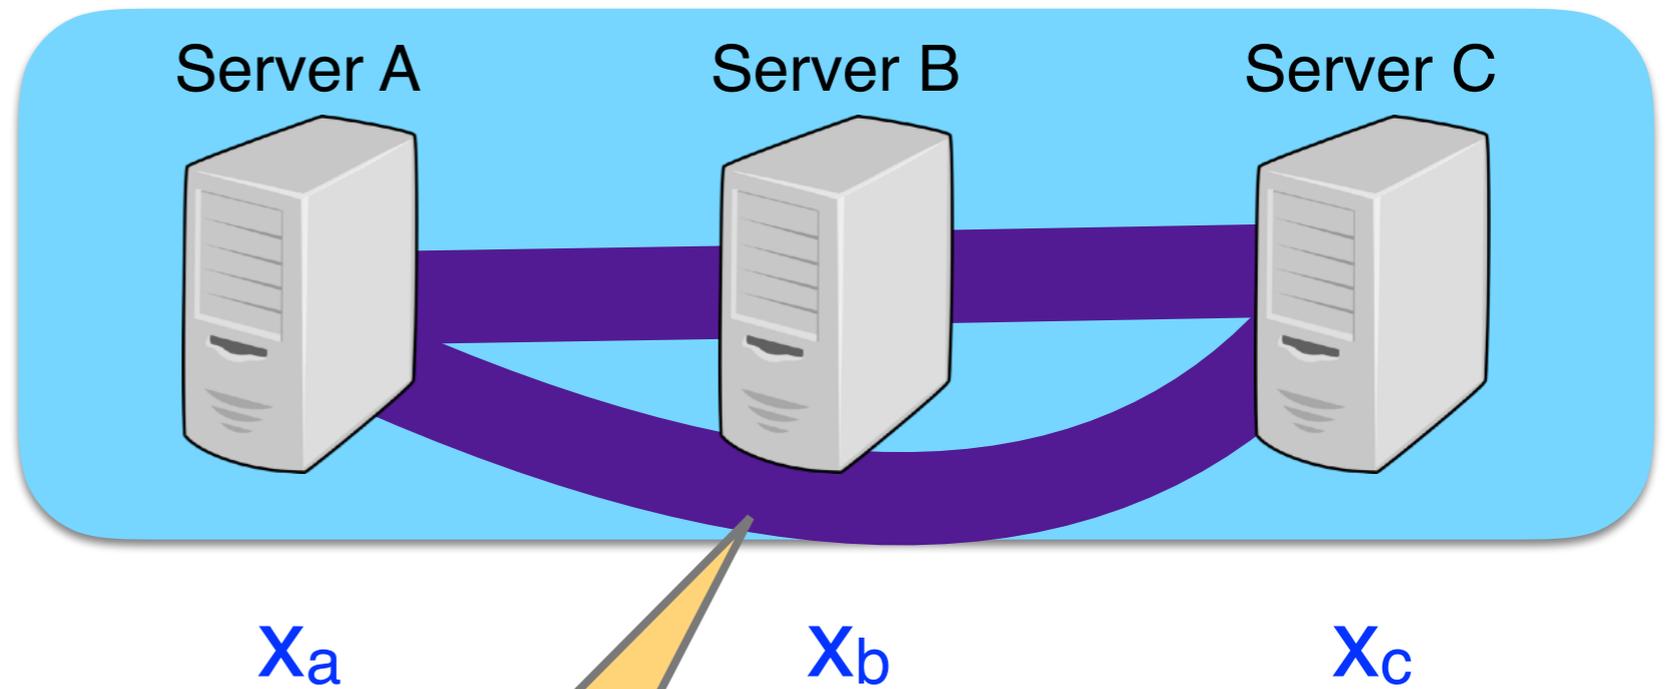


How SNIPs work



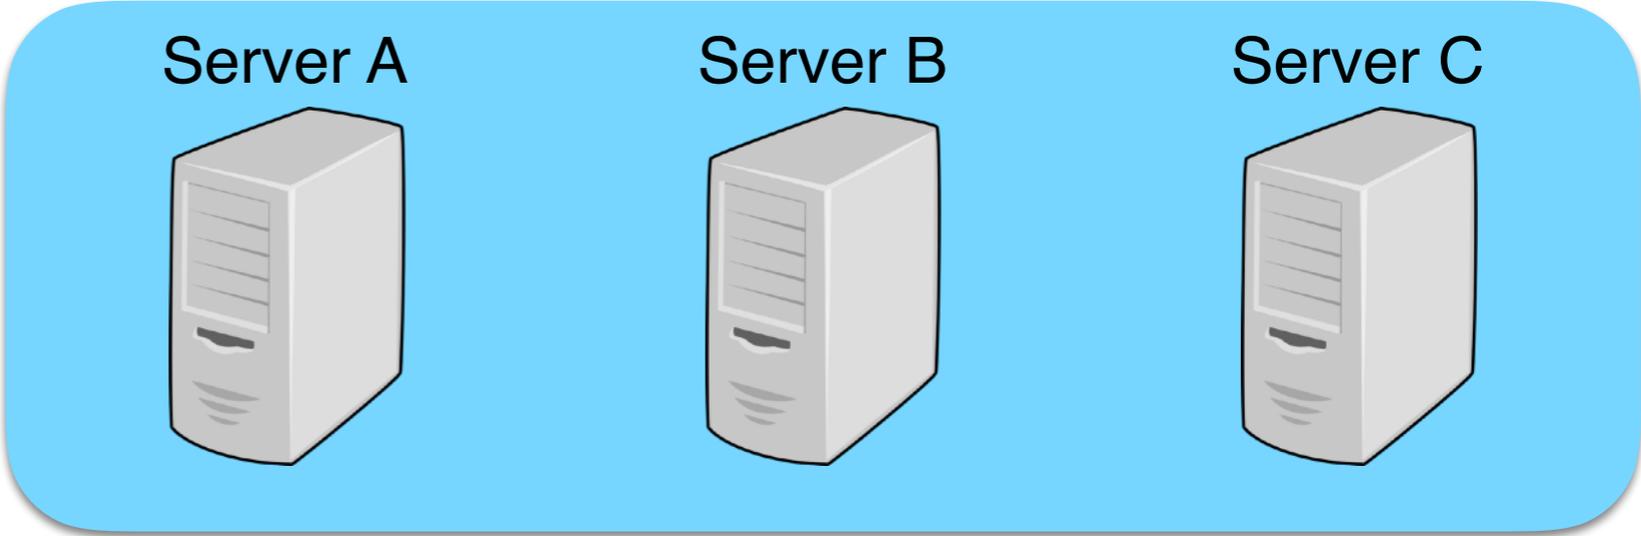
Could run **secure multiparty computation** to check that $\text{Valid}(x) = 1$.
[GMW87], [BGW88]

How SNIPs work

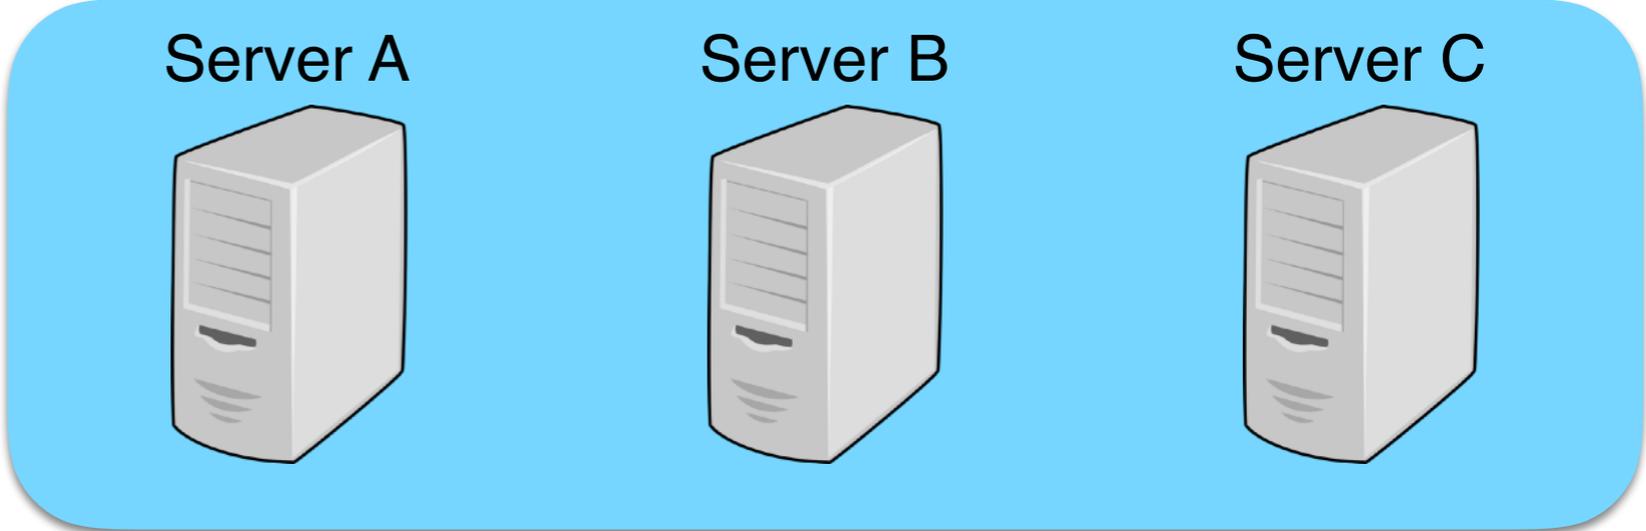


Could run **secure multiparty computation** to check that $\text{Valid}(x) = 1$.
[GMW87], [BGW88]

How SNIPs work



How SNIPs work



X_a

X_b

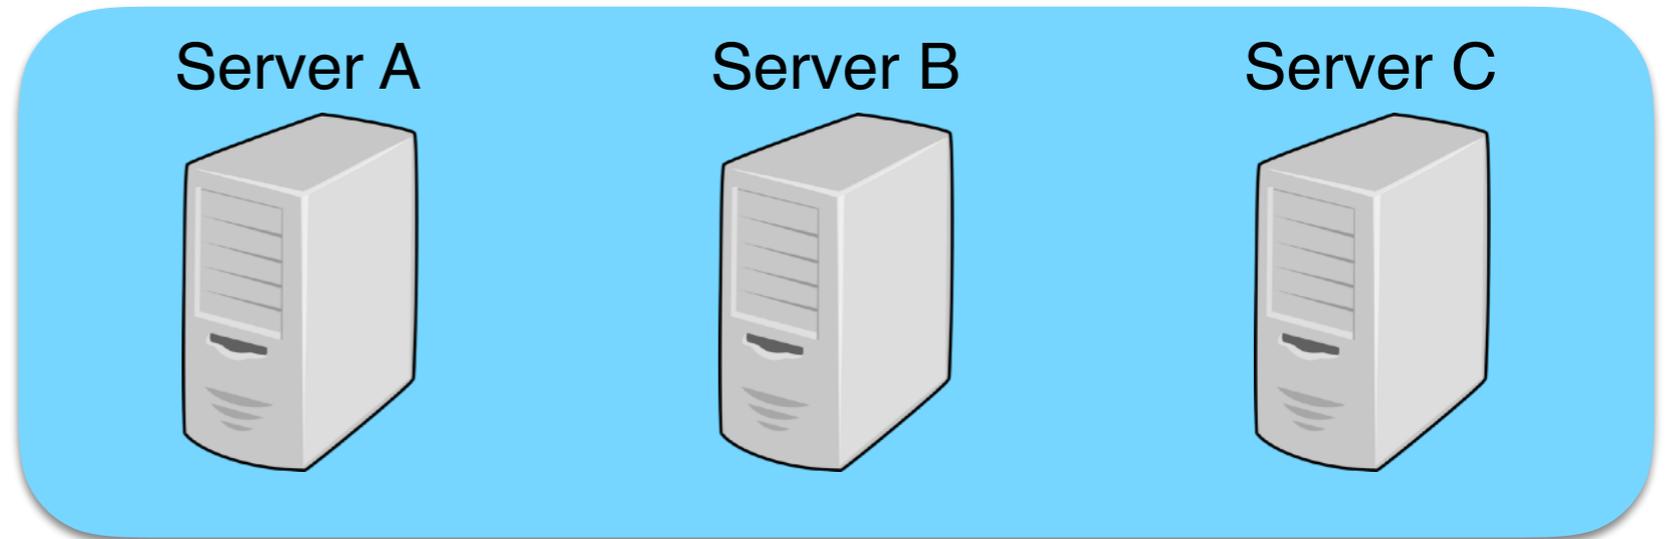
X_c



X



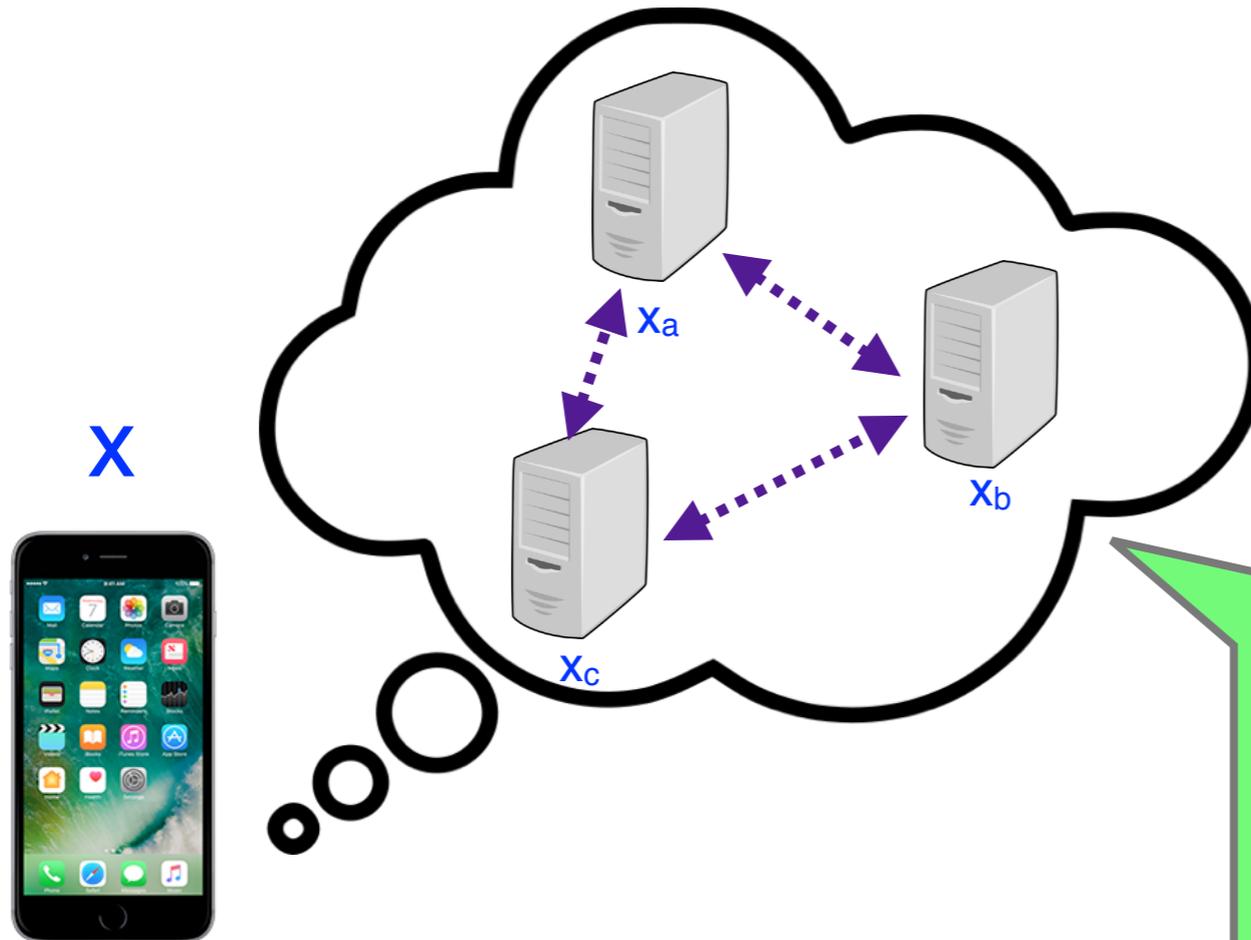
How SNIPs work



X_a

X_b

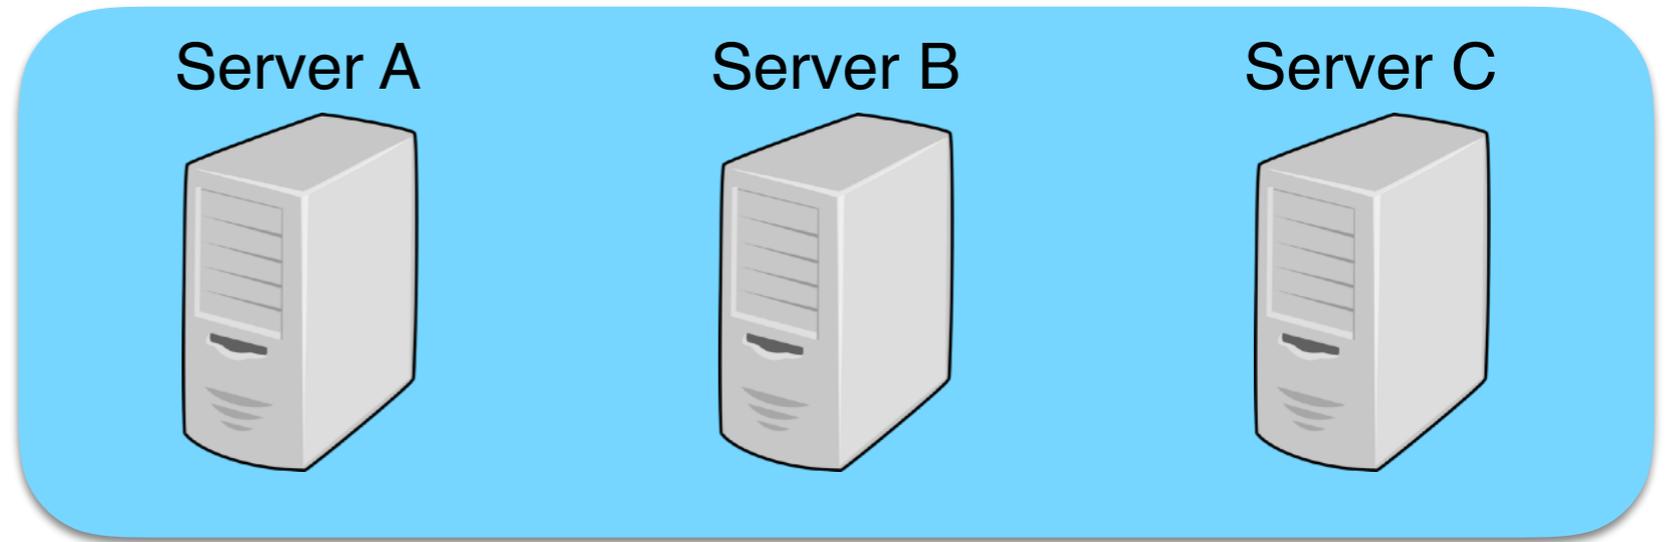
X_c



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

See also [IKOS07]

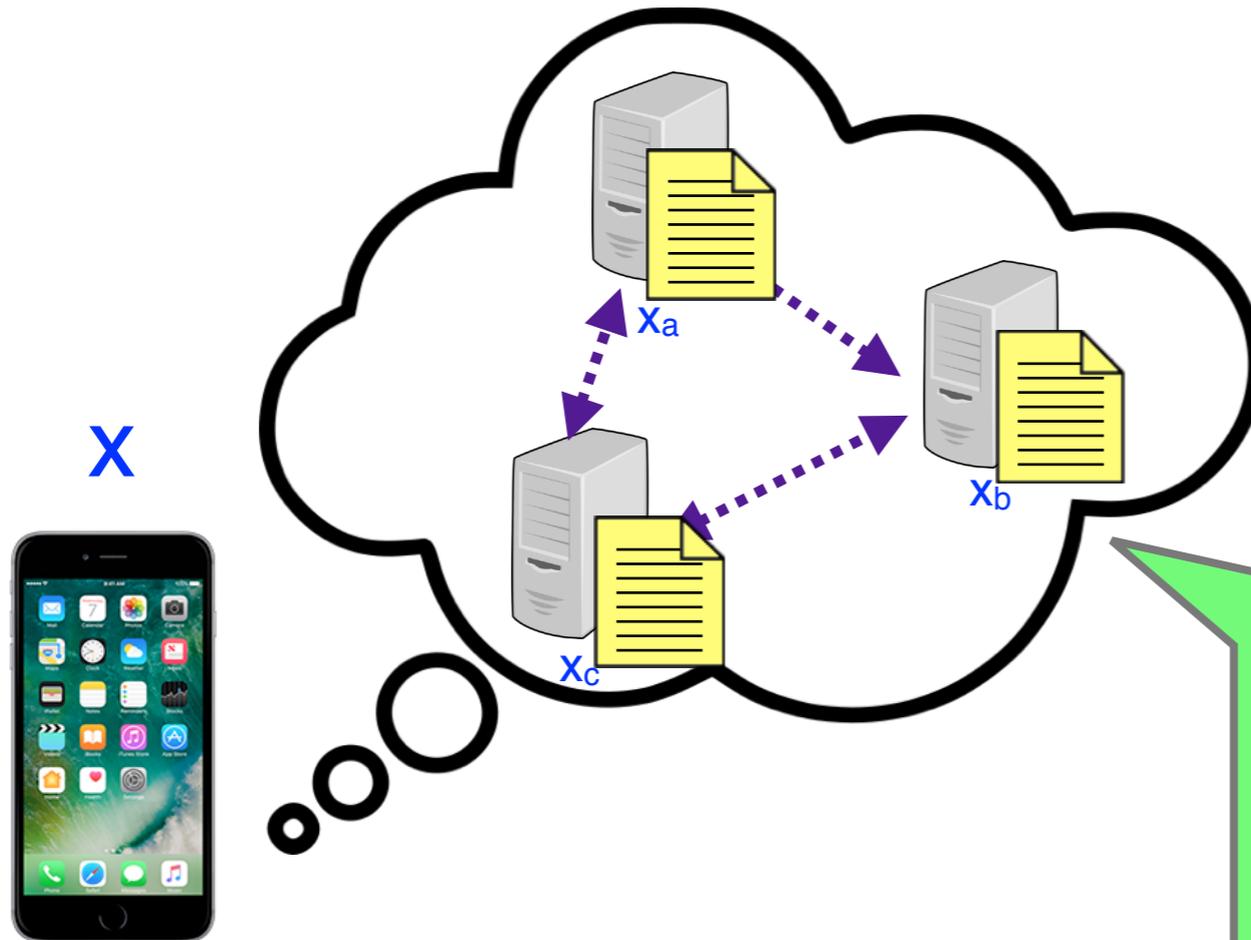
How SNIPs work



X_a

X_b

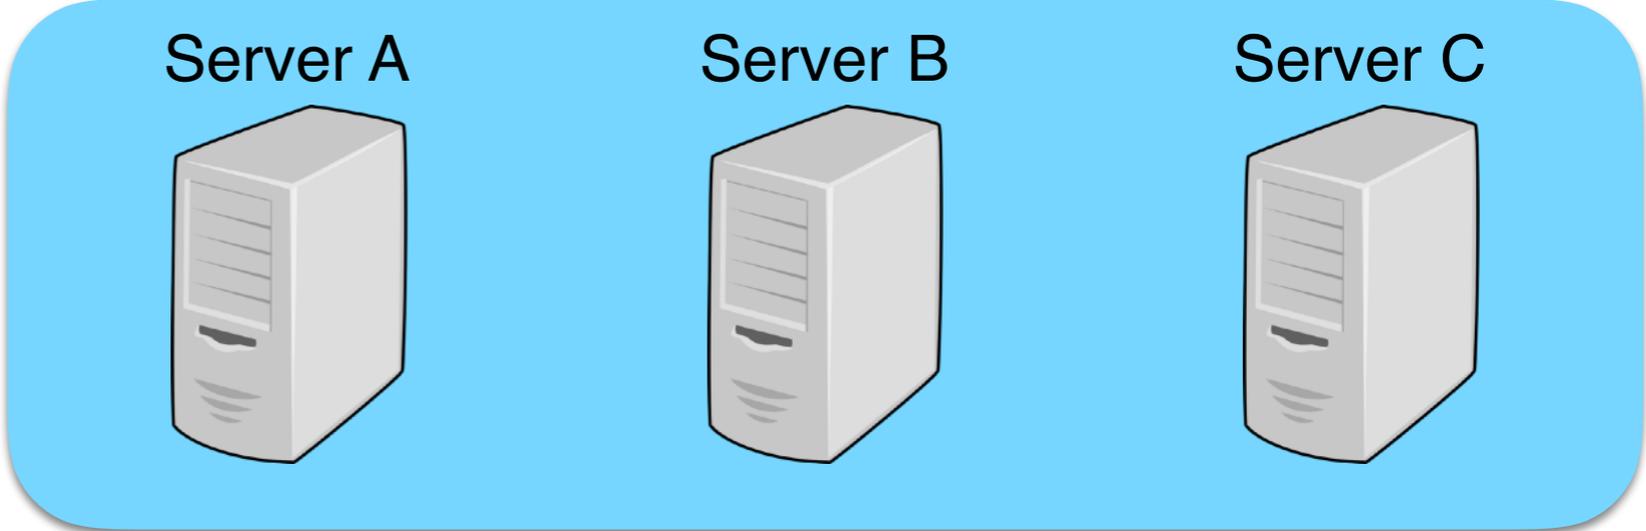
X_c



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

See also [IKOS07]

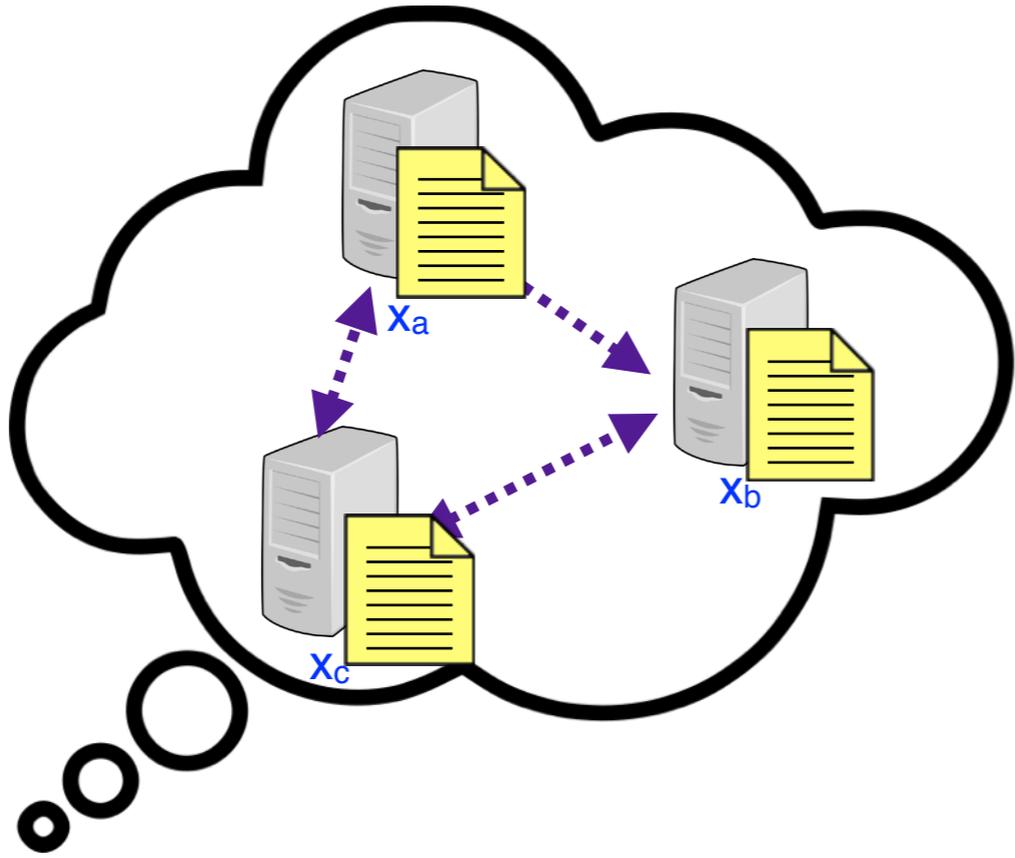
How SNIPs work



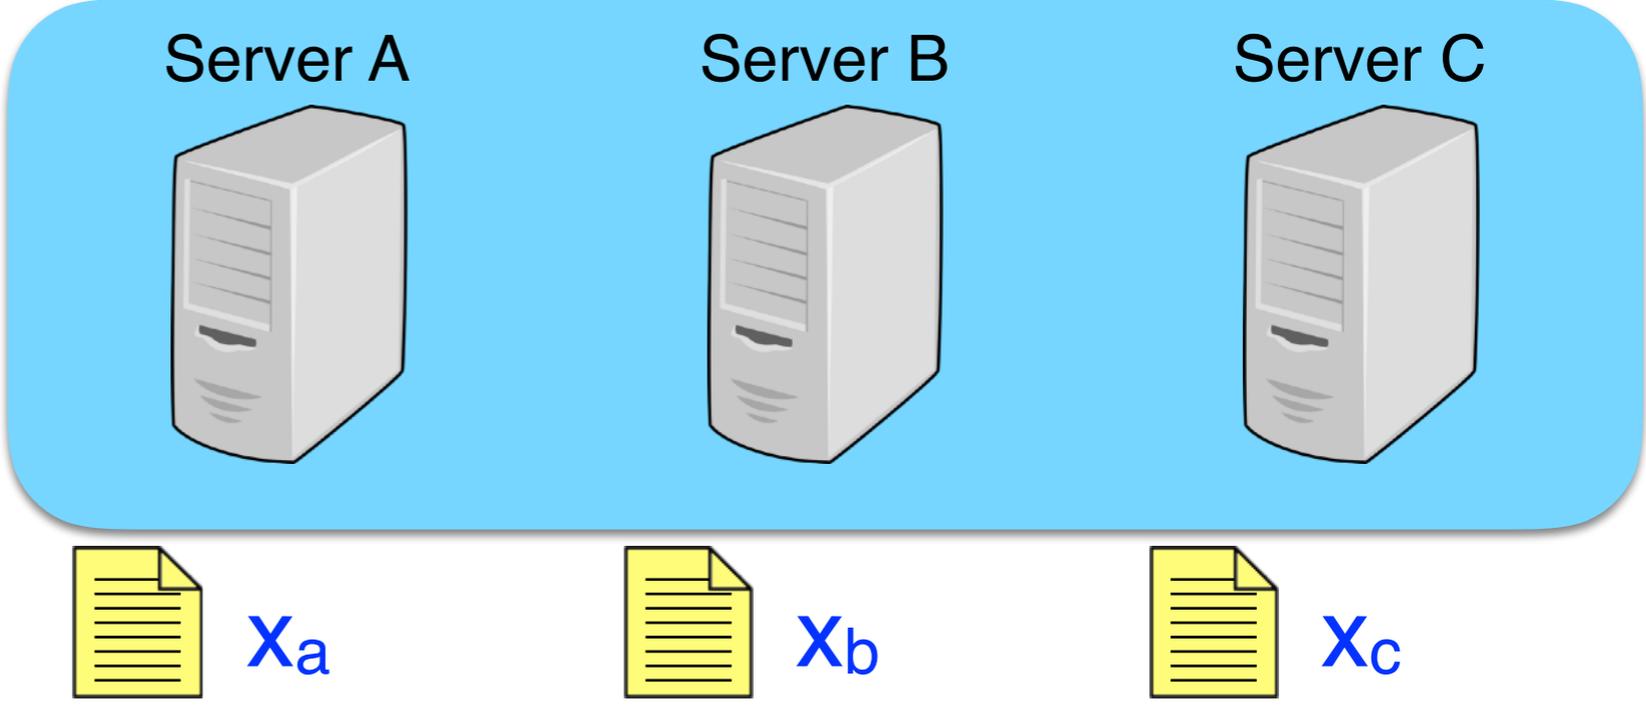
X_a

X_b

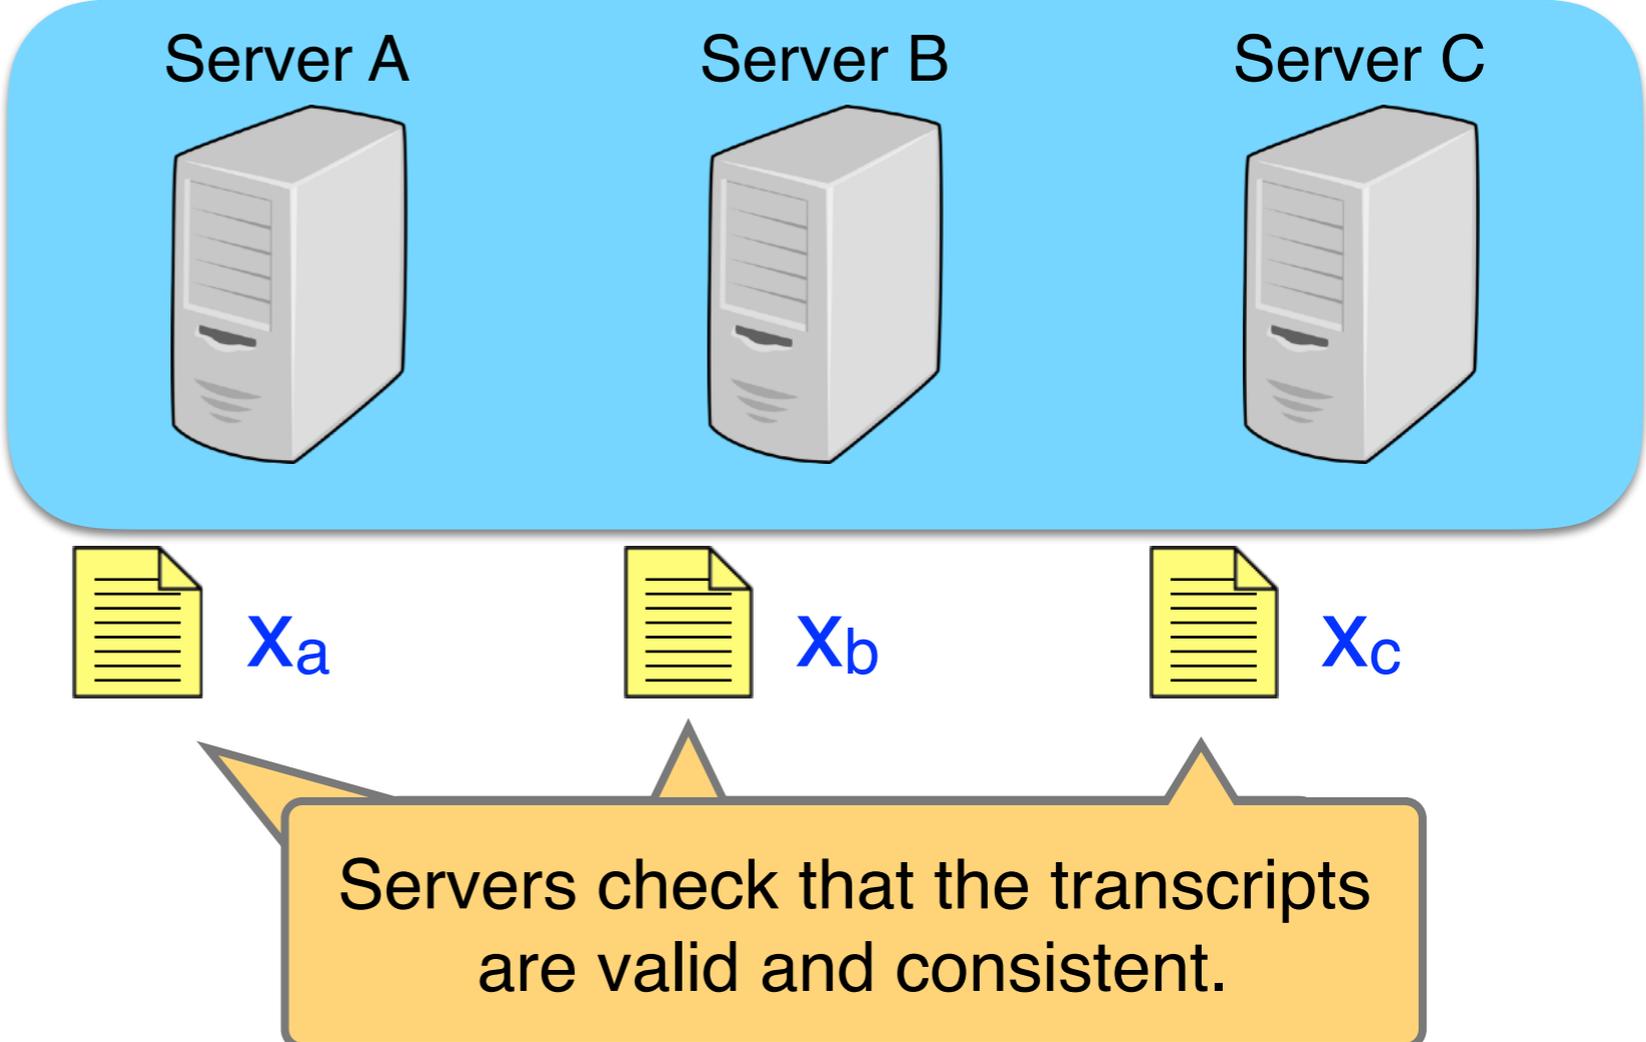
X_c



How SNIPs work



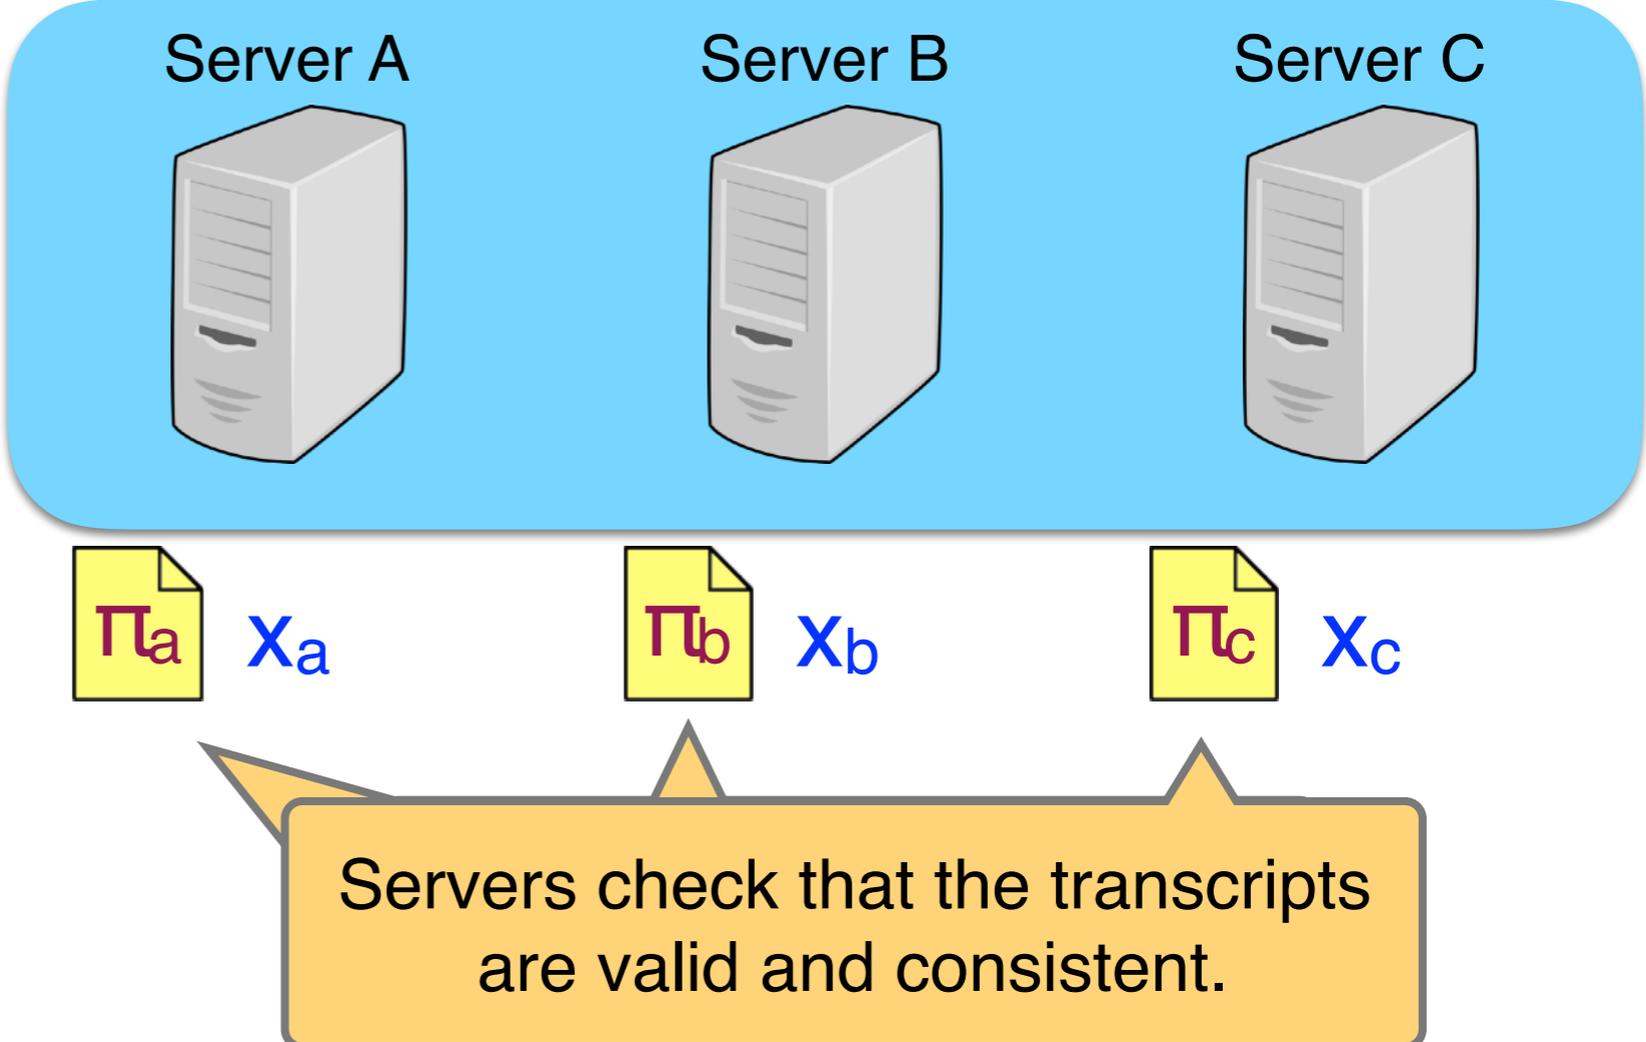
How SNIPs work



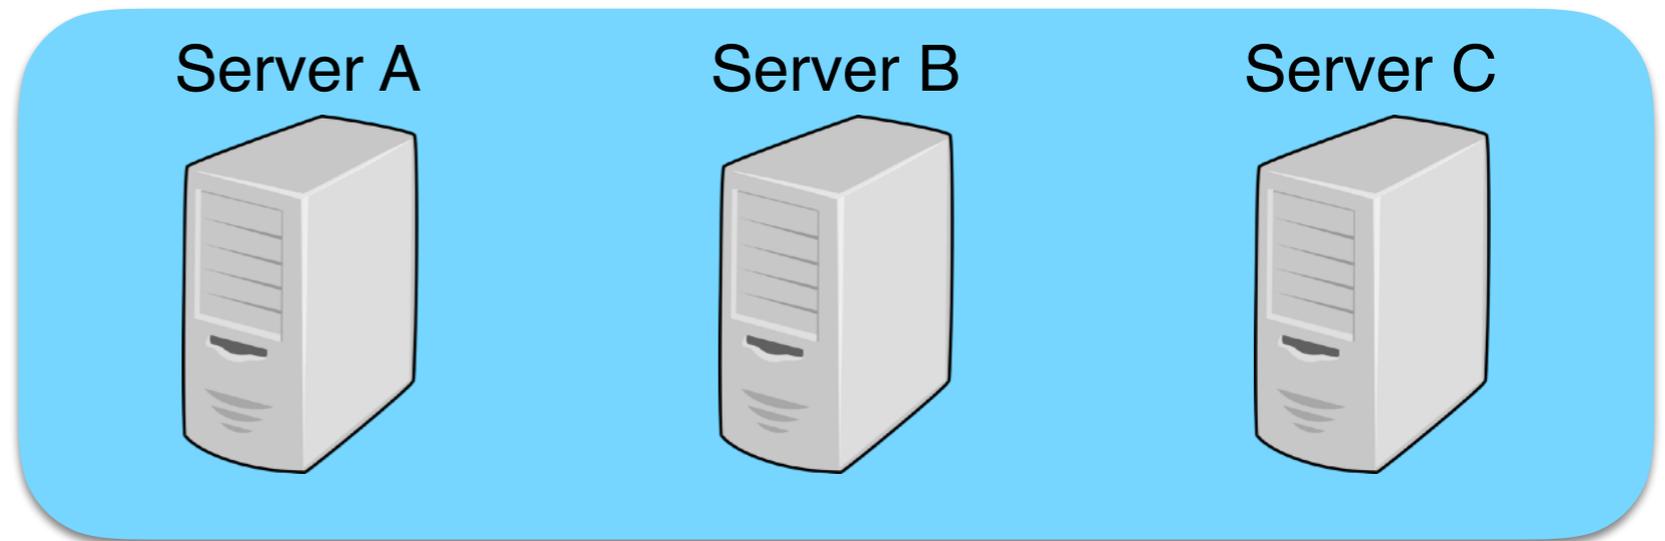
X



How SNIPs work



How SNIPs work



X_a



X_b



X_c

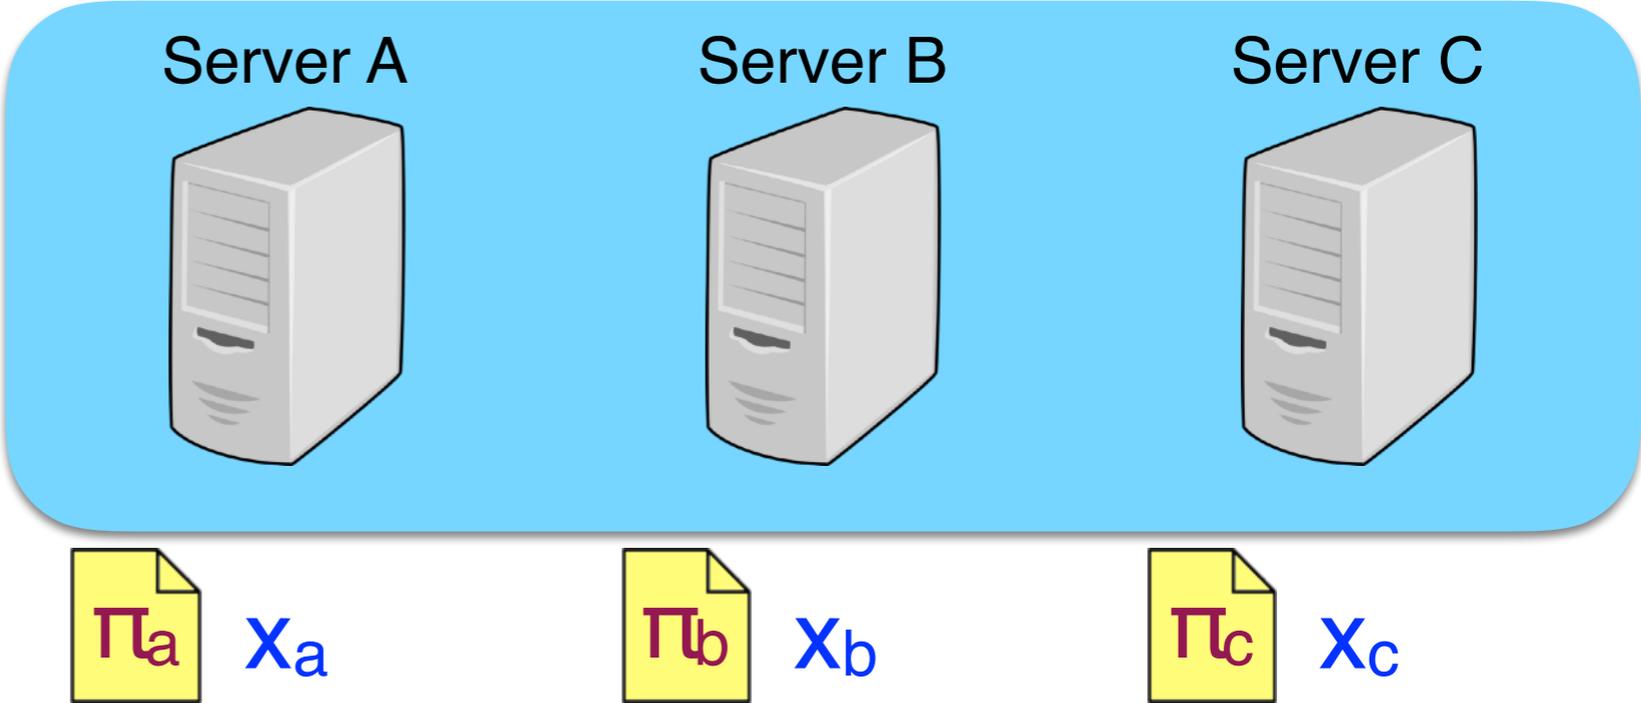
Servers check that the transcripts are valid and consistent.

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

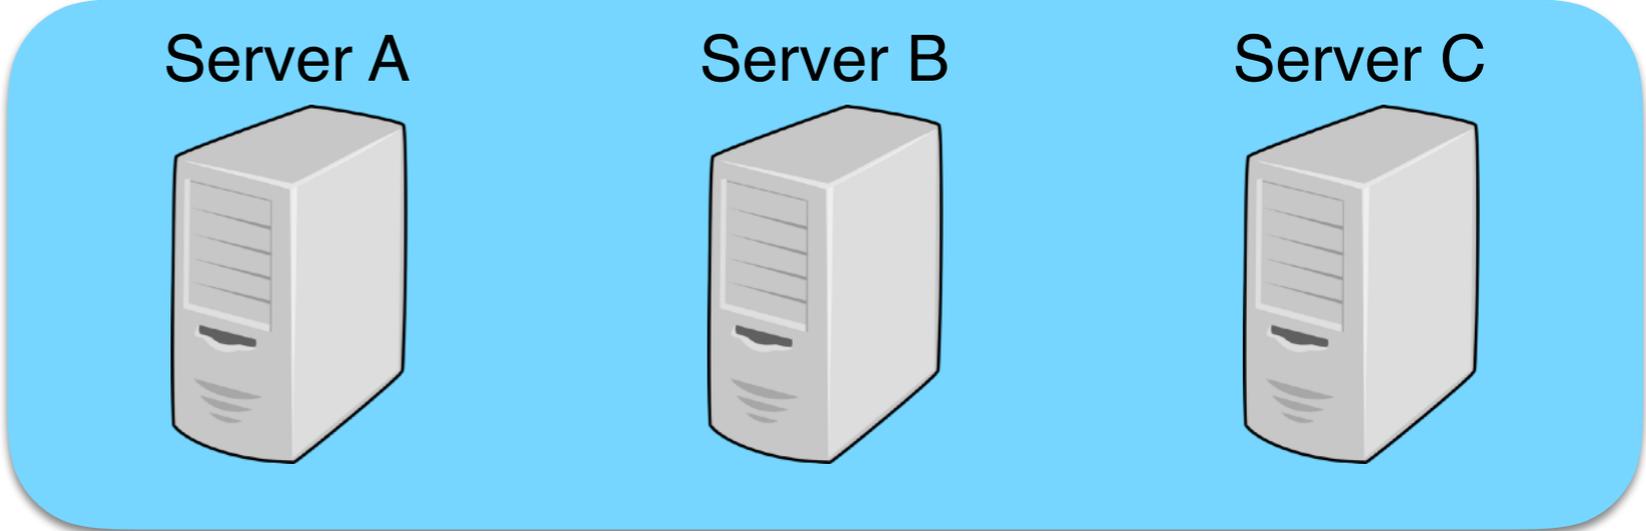
X



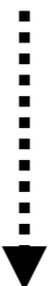
How SNIPs work



How SNIPs work



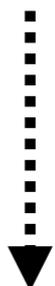
X_a



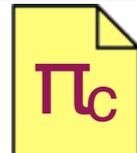
D_a



X_b



D_b

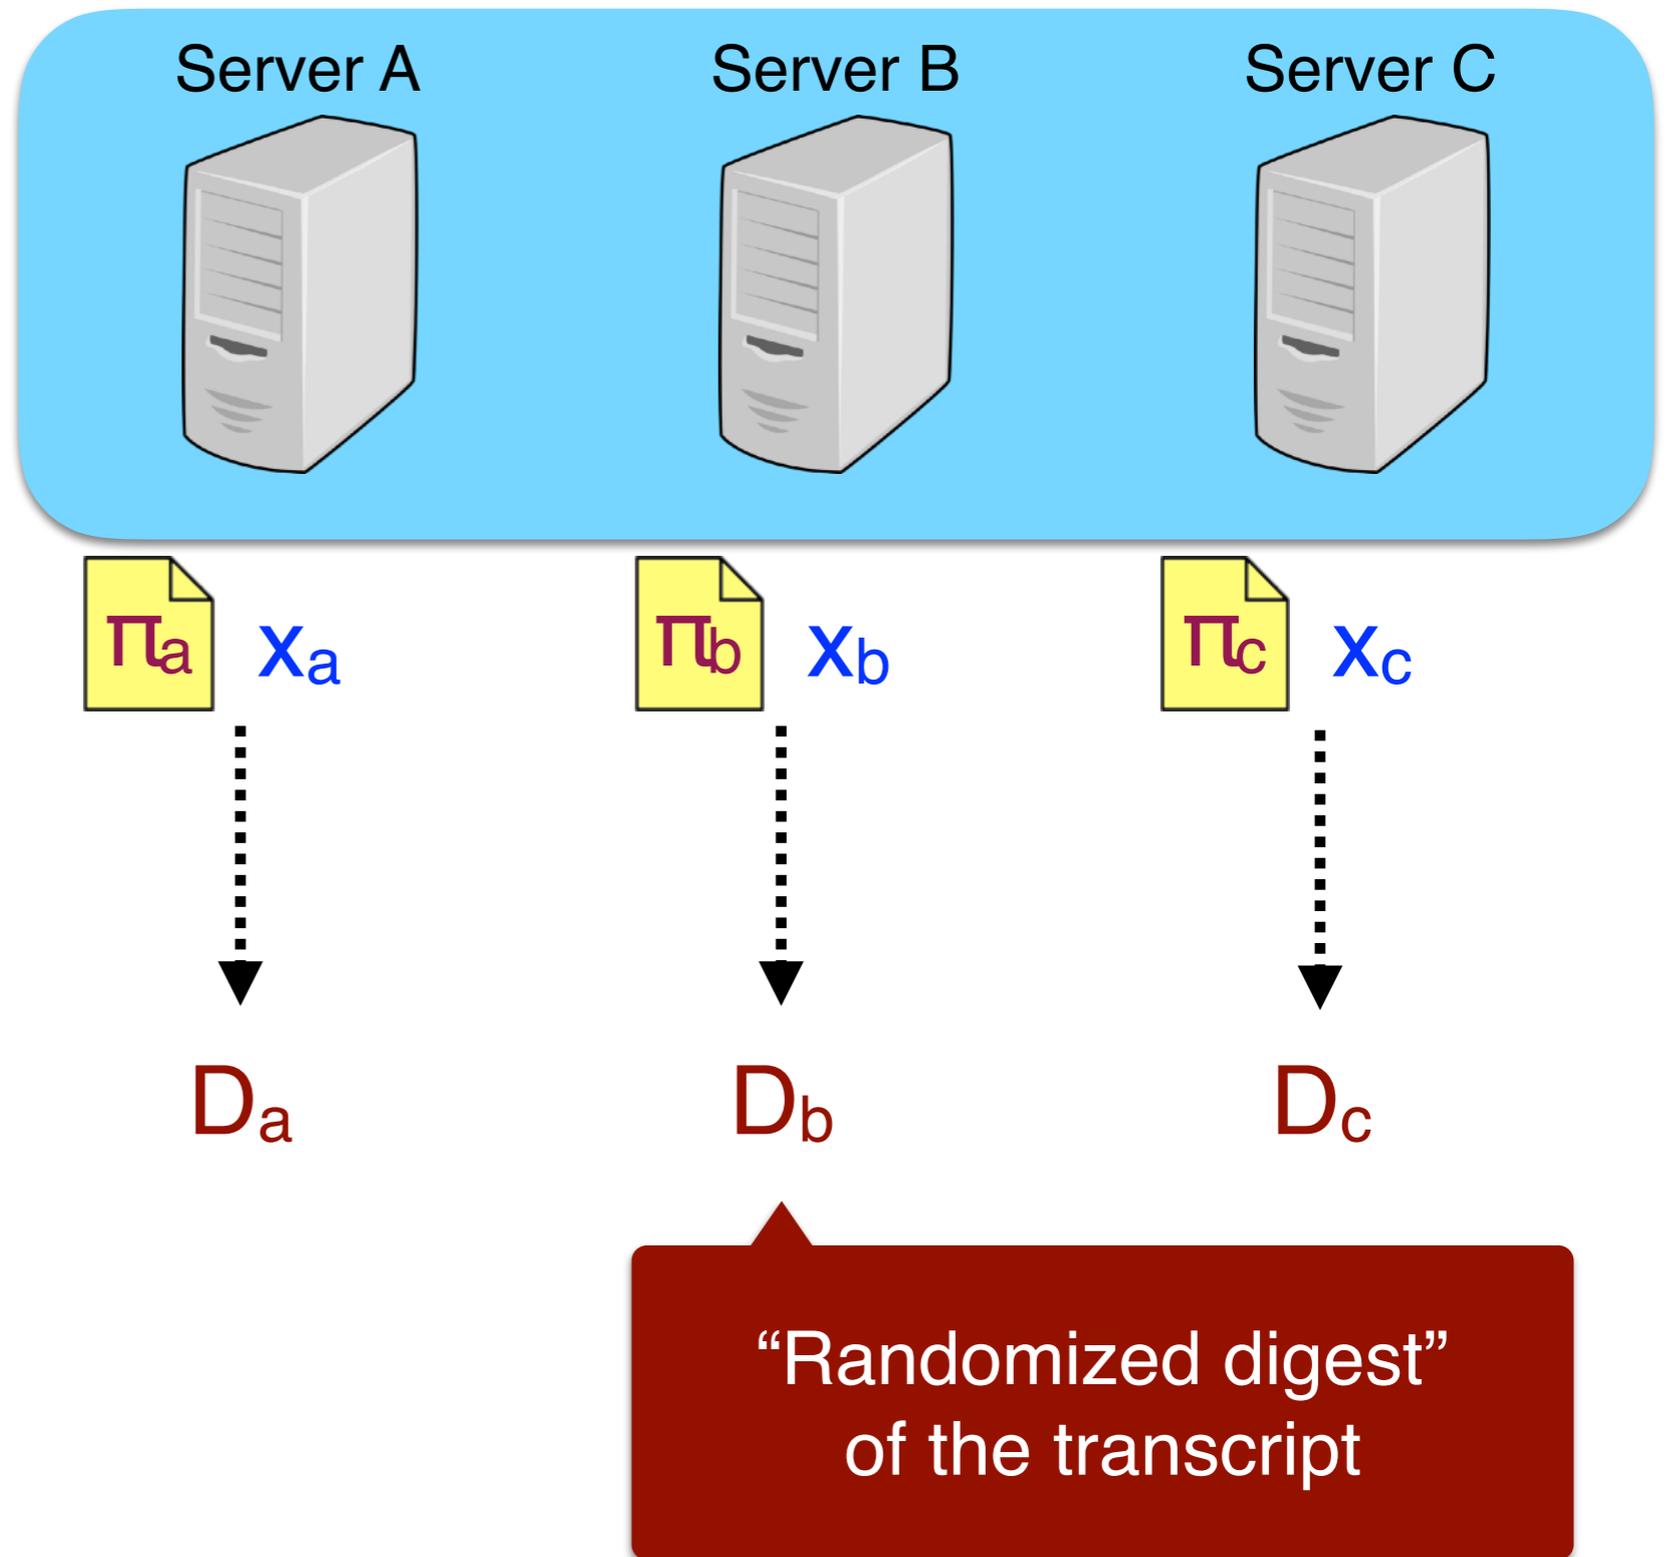


X_c

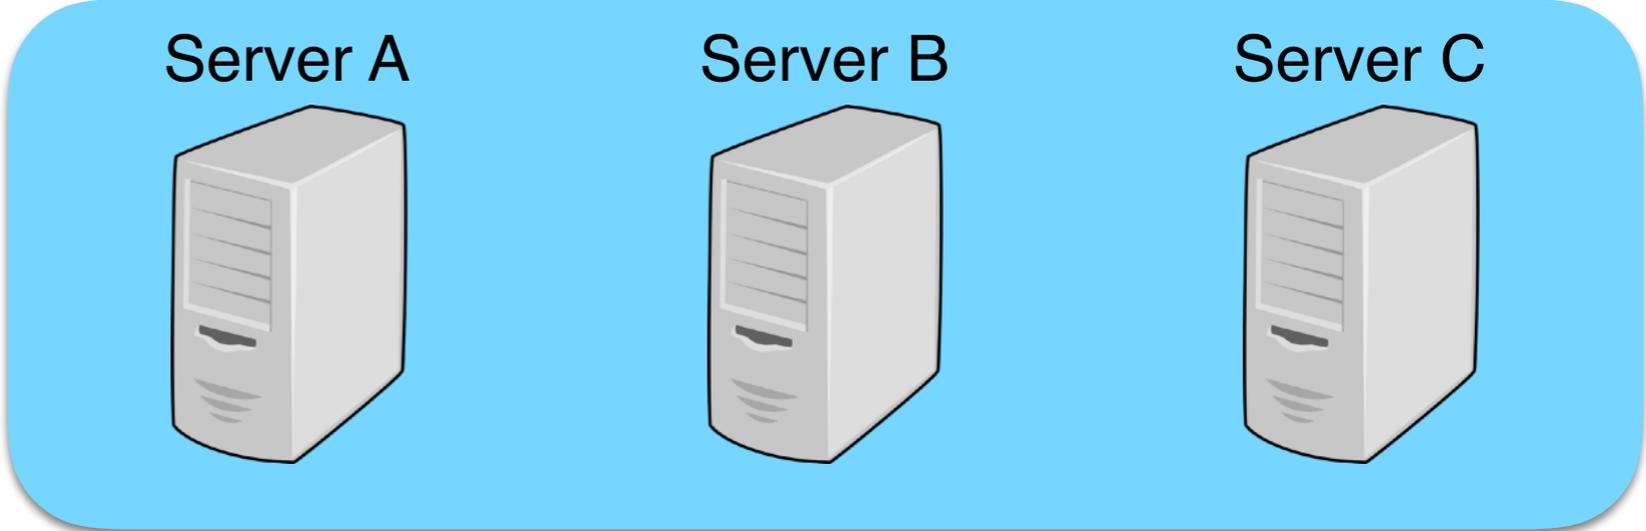


D_c

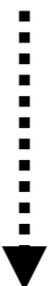
How SNIPs work



How SNIPs work



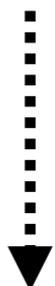
X_a



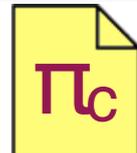
D_a



X_b



D_b

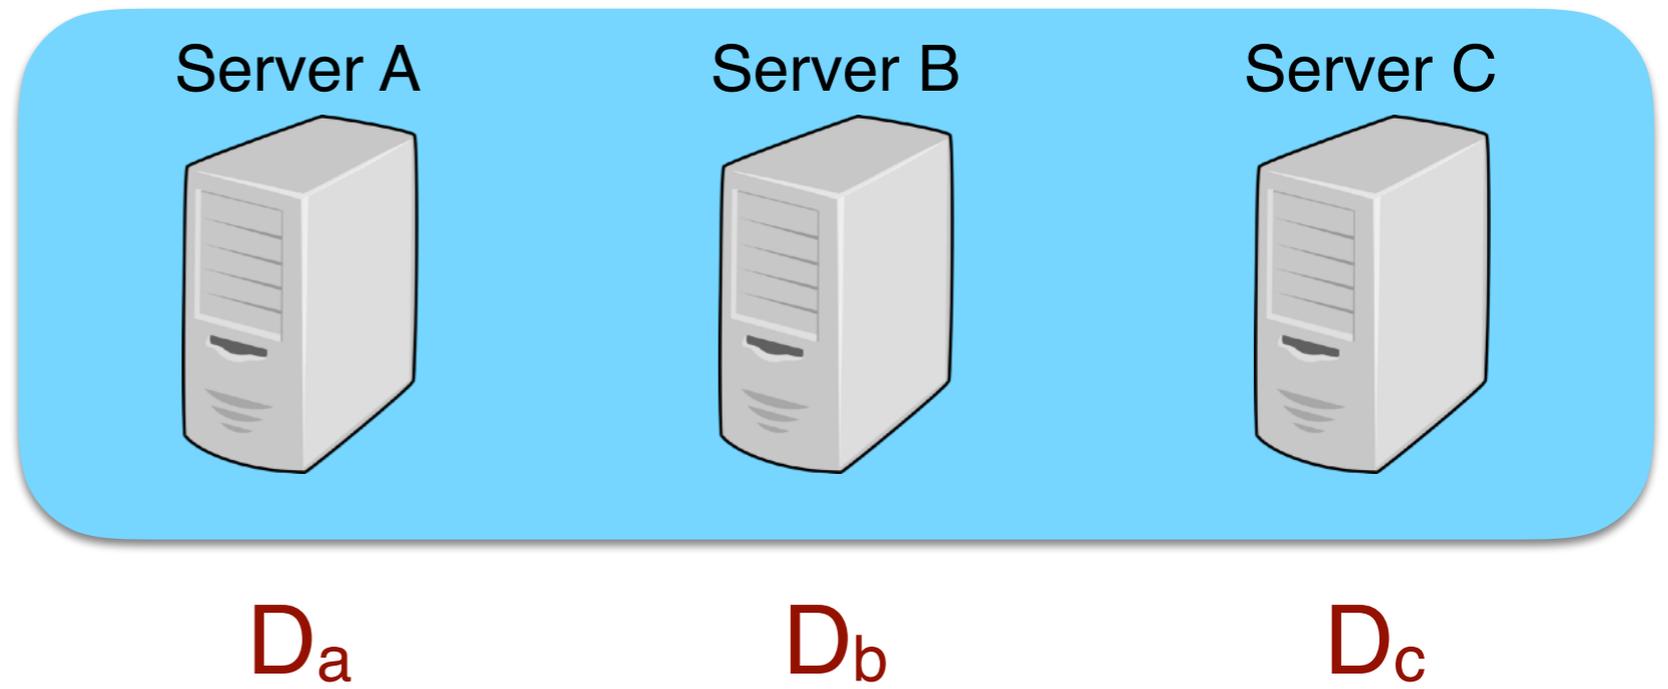


X_c

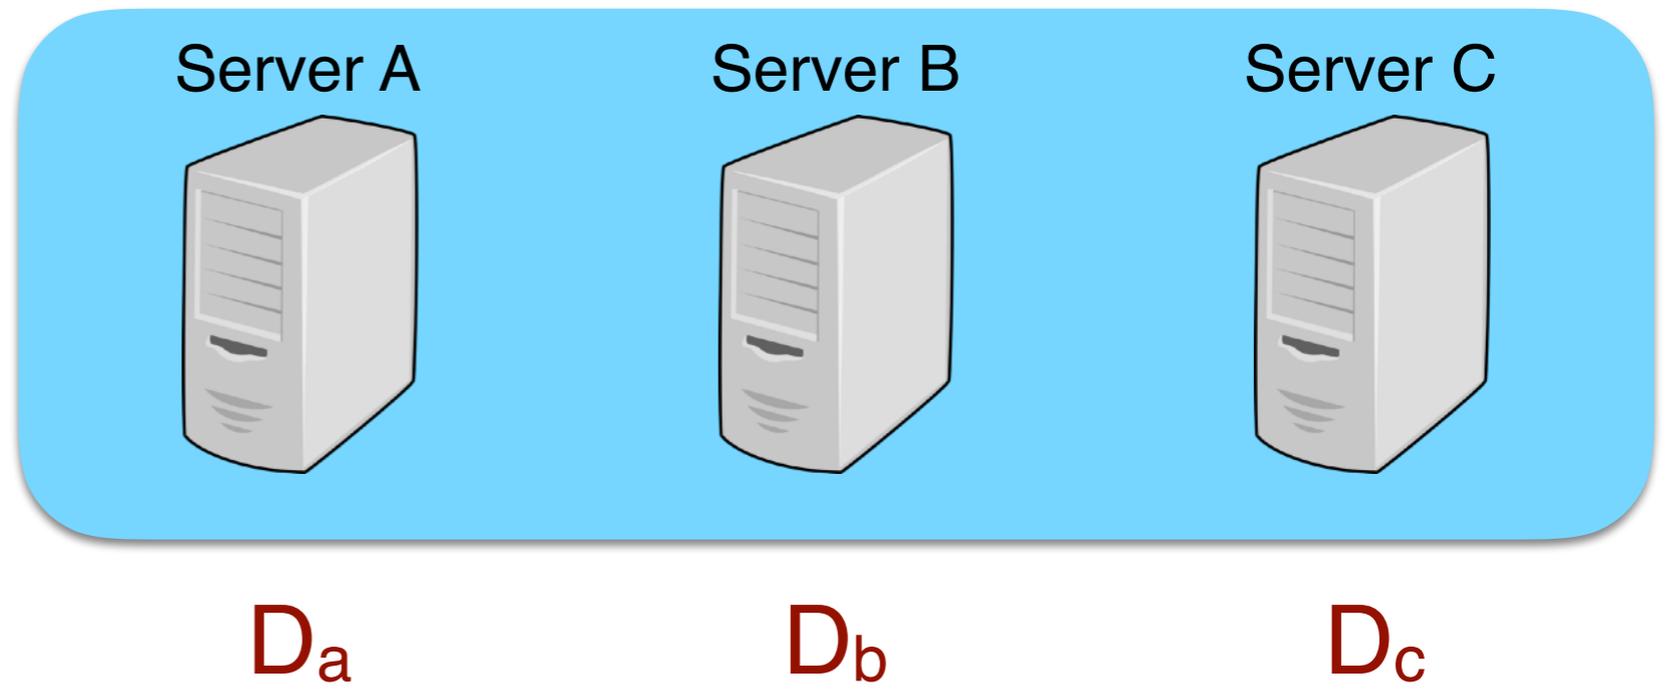


D_c

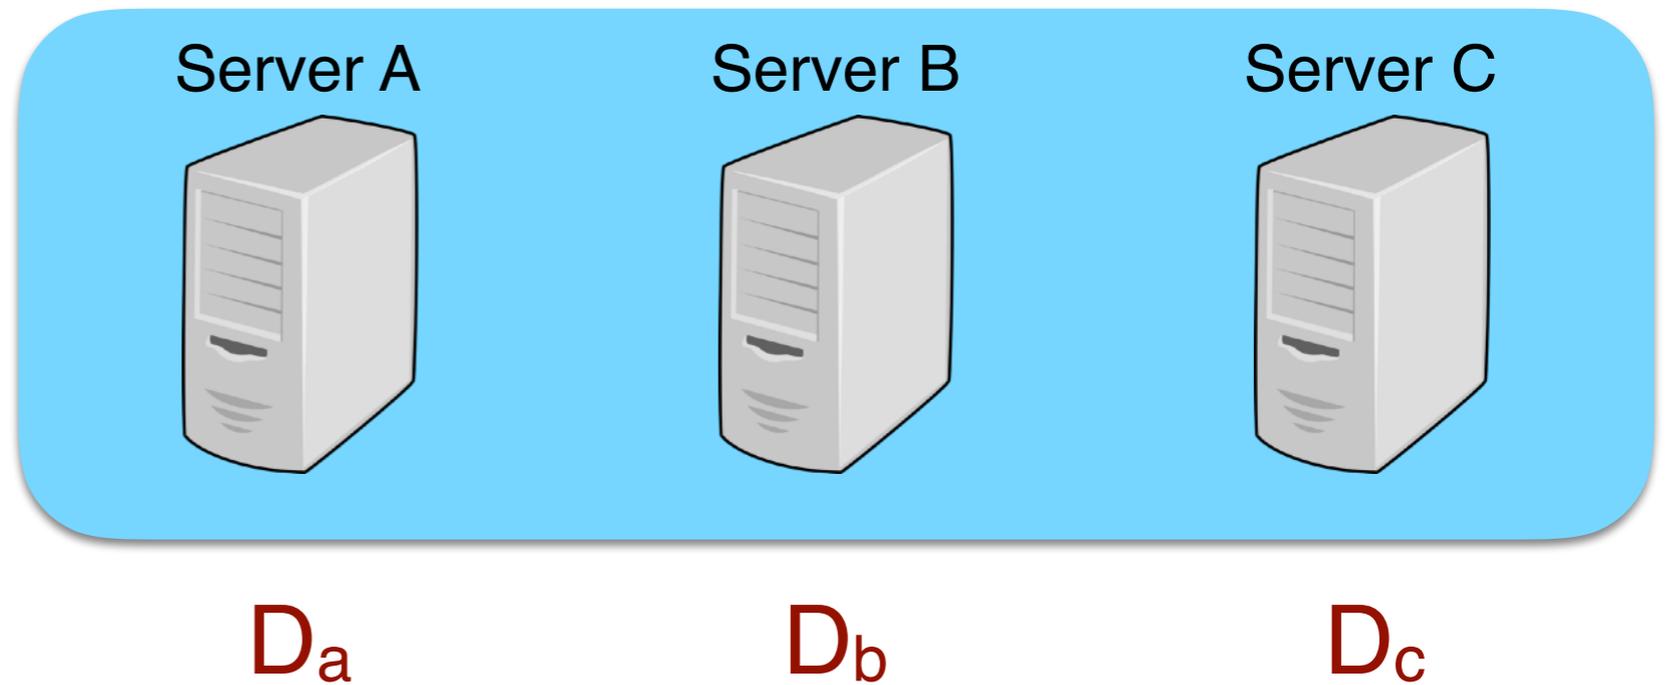
How SNIPs work



How SNIPs work



How SNIPs work



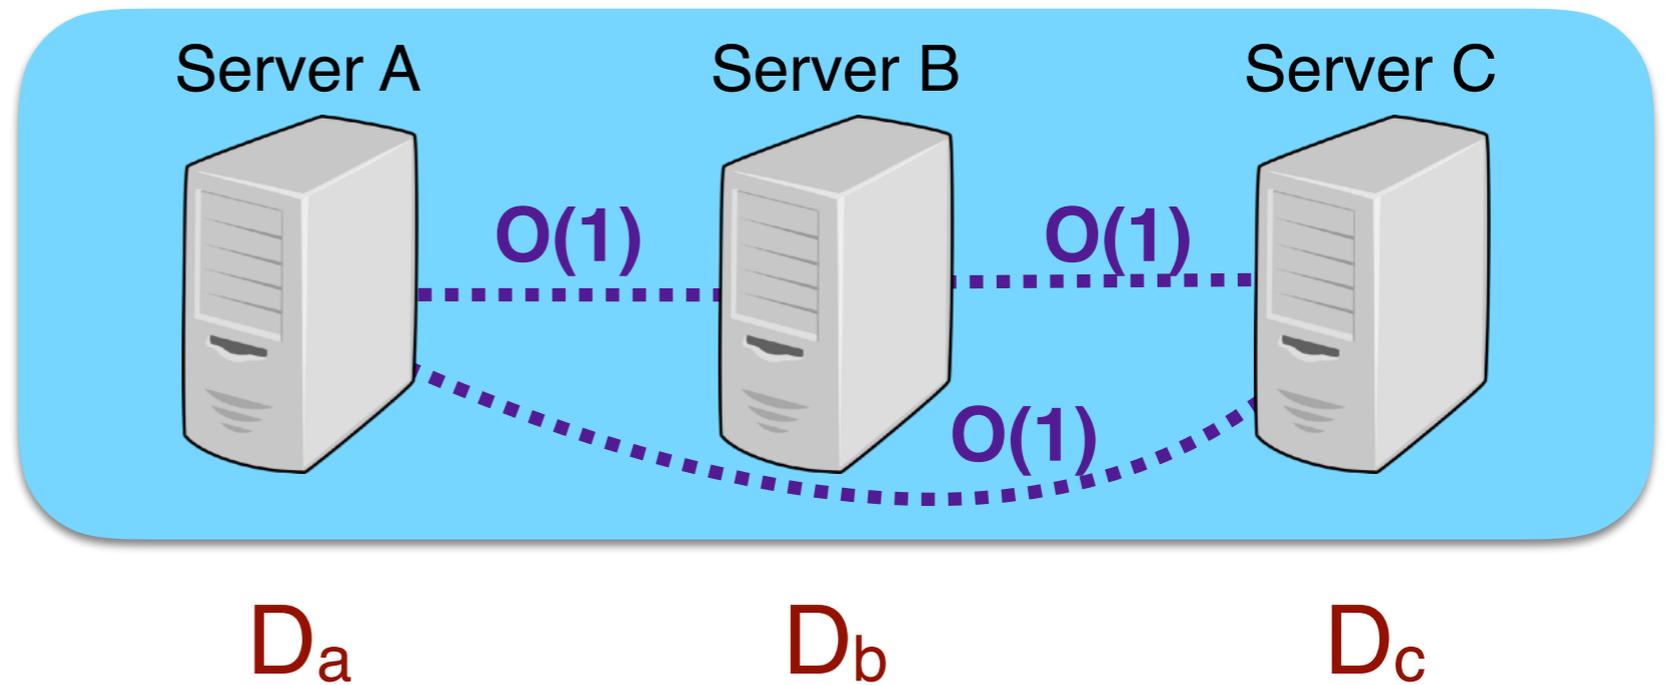
-
- If x is valid, $D_a + D_b + D_c = 0$
 - If x is invalid, $D_a + D_b + D_c \neq 0$ with high probability

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

Public-key ops.		Communication		Slow-down
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
[CLOS02], [DPSZ12], ...

0

$\Theta(M)$

0

$\Theta(M)$

5,000x
at server

Commits + NIZKs
[FS86], [CP92], [CS97], ...

$\Theta(M)$

$\Theta(M)$

$\Theta(M)$

$\Theta(M)$

50x
at server

	Public-key ops.		Communication		Slow-down
	Client	Server	C-to-S	S-to-S	
M = # of multiplication gates in Valid (·) 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 Valid (·) 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
This work: SNIPs	0	0	$\Theta(M)$	$O(1)$	1x

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
Commits + SNARKs <small>[GGPR13], [BCGTV13], ...</small>	$\Theta(M)$	$O(1)$	$O(1)$	$O(1)$	500x at client
This work: SNIPs	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 Valid (·) 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
This work: SNIPs	0	0	$\Theta(M)$	$O(1)$	1x

Outline

- Background: The private aggregation problem
- A straw-man solution for private sums
- **Providing robustness with SNIPs**
- Evaluation
- Encodings for complex aggregates

Outline

- Background: The private aggregation problem
- A straw-man solution for private sums
- Providing robustness with SNIPs
- **Evaluation**
- Encodings for complex aggregates

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

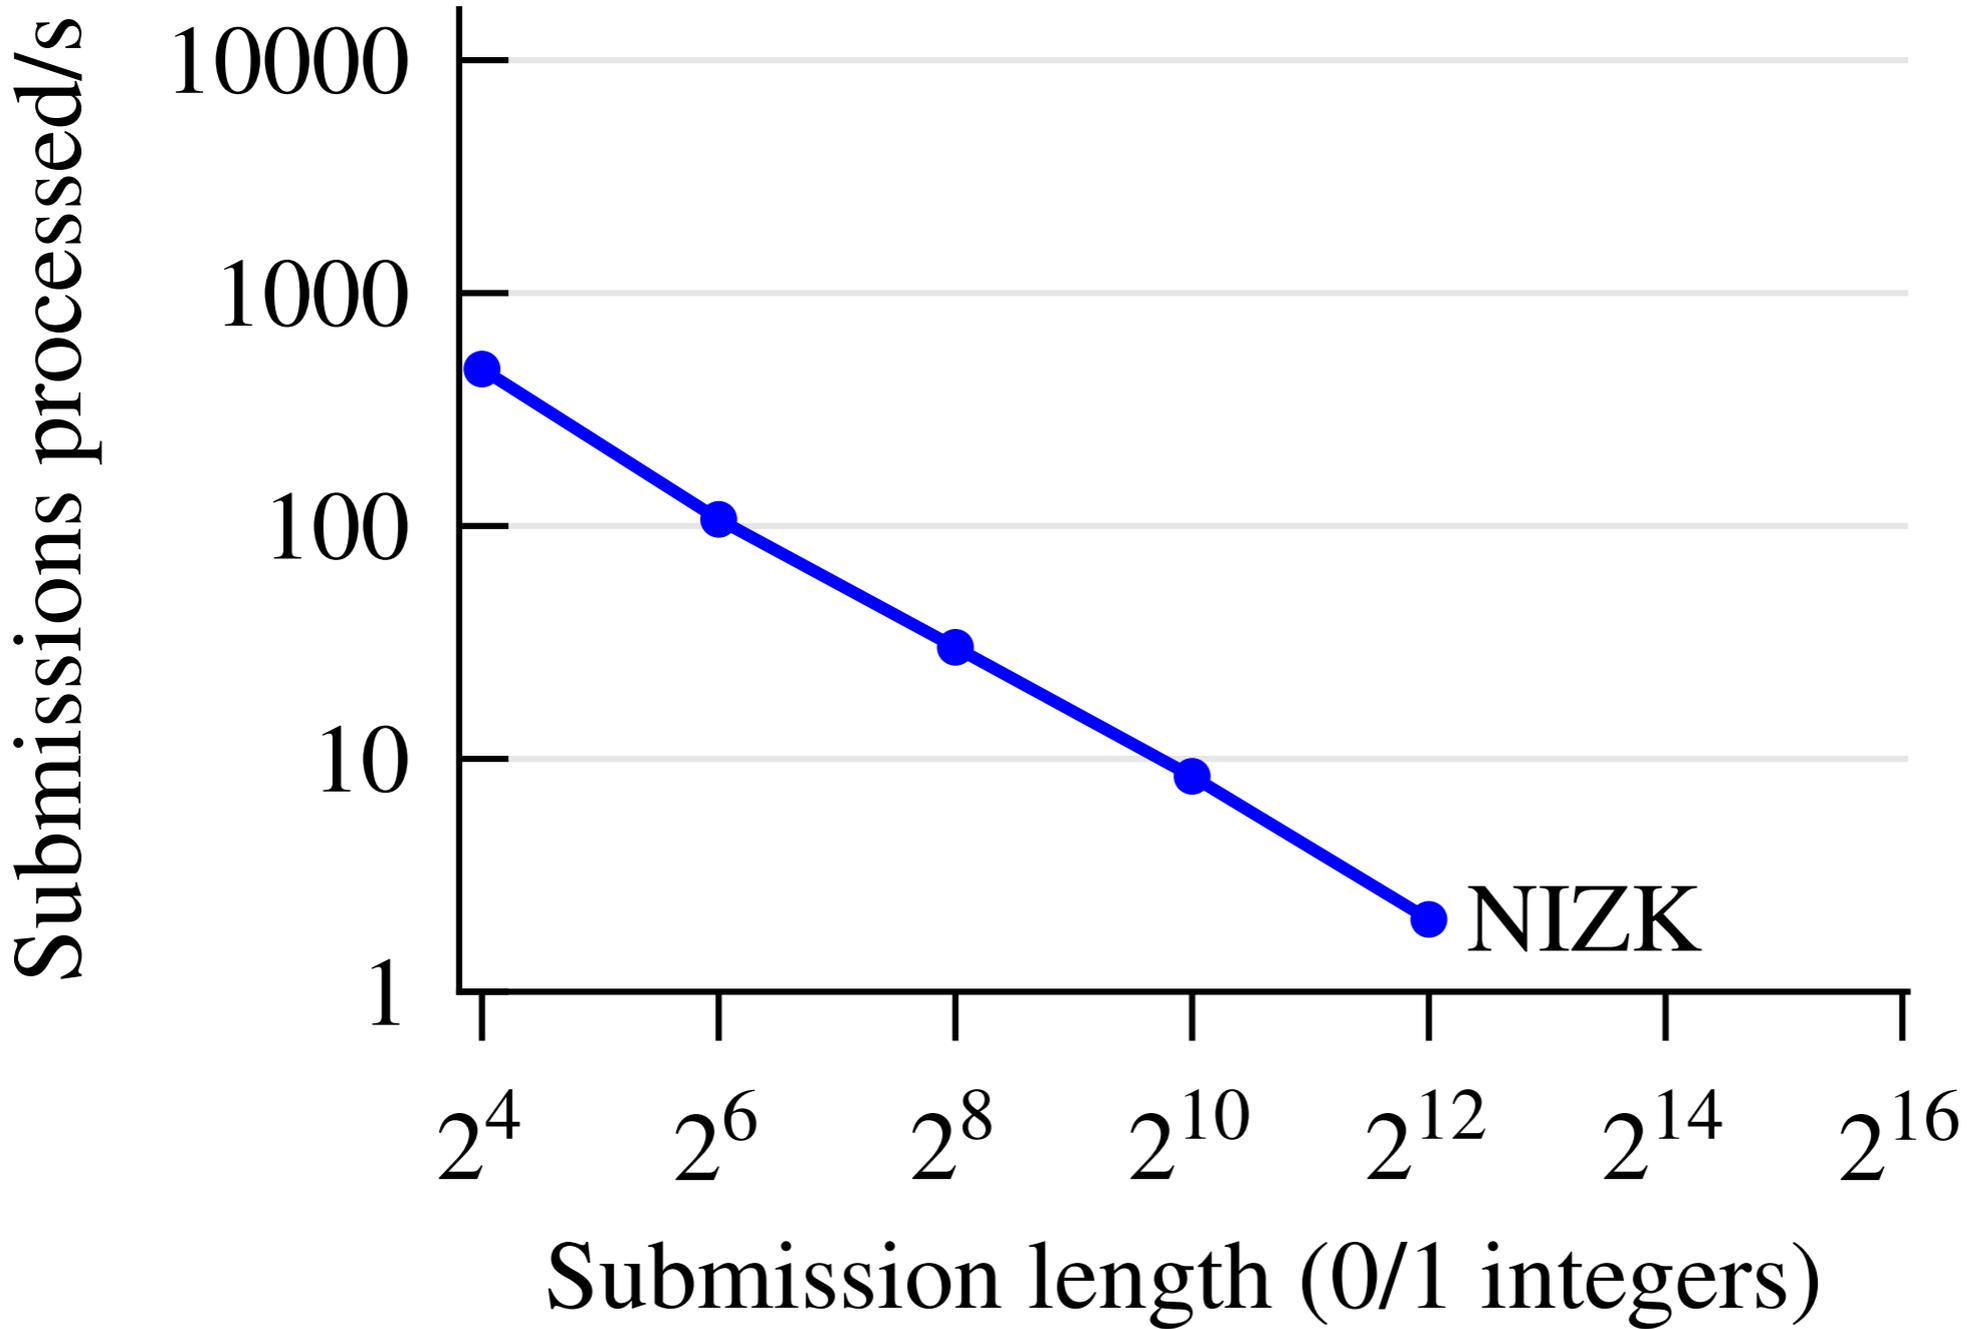
- Implemented Prio in Go
(see optimizations described in paper)
- Five-server cluster in EC2
- System collects the sum of “N” 0/1 values

Four variants

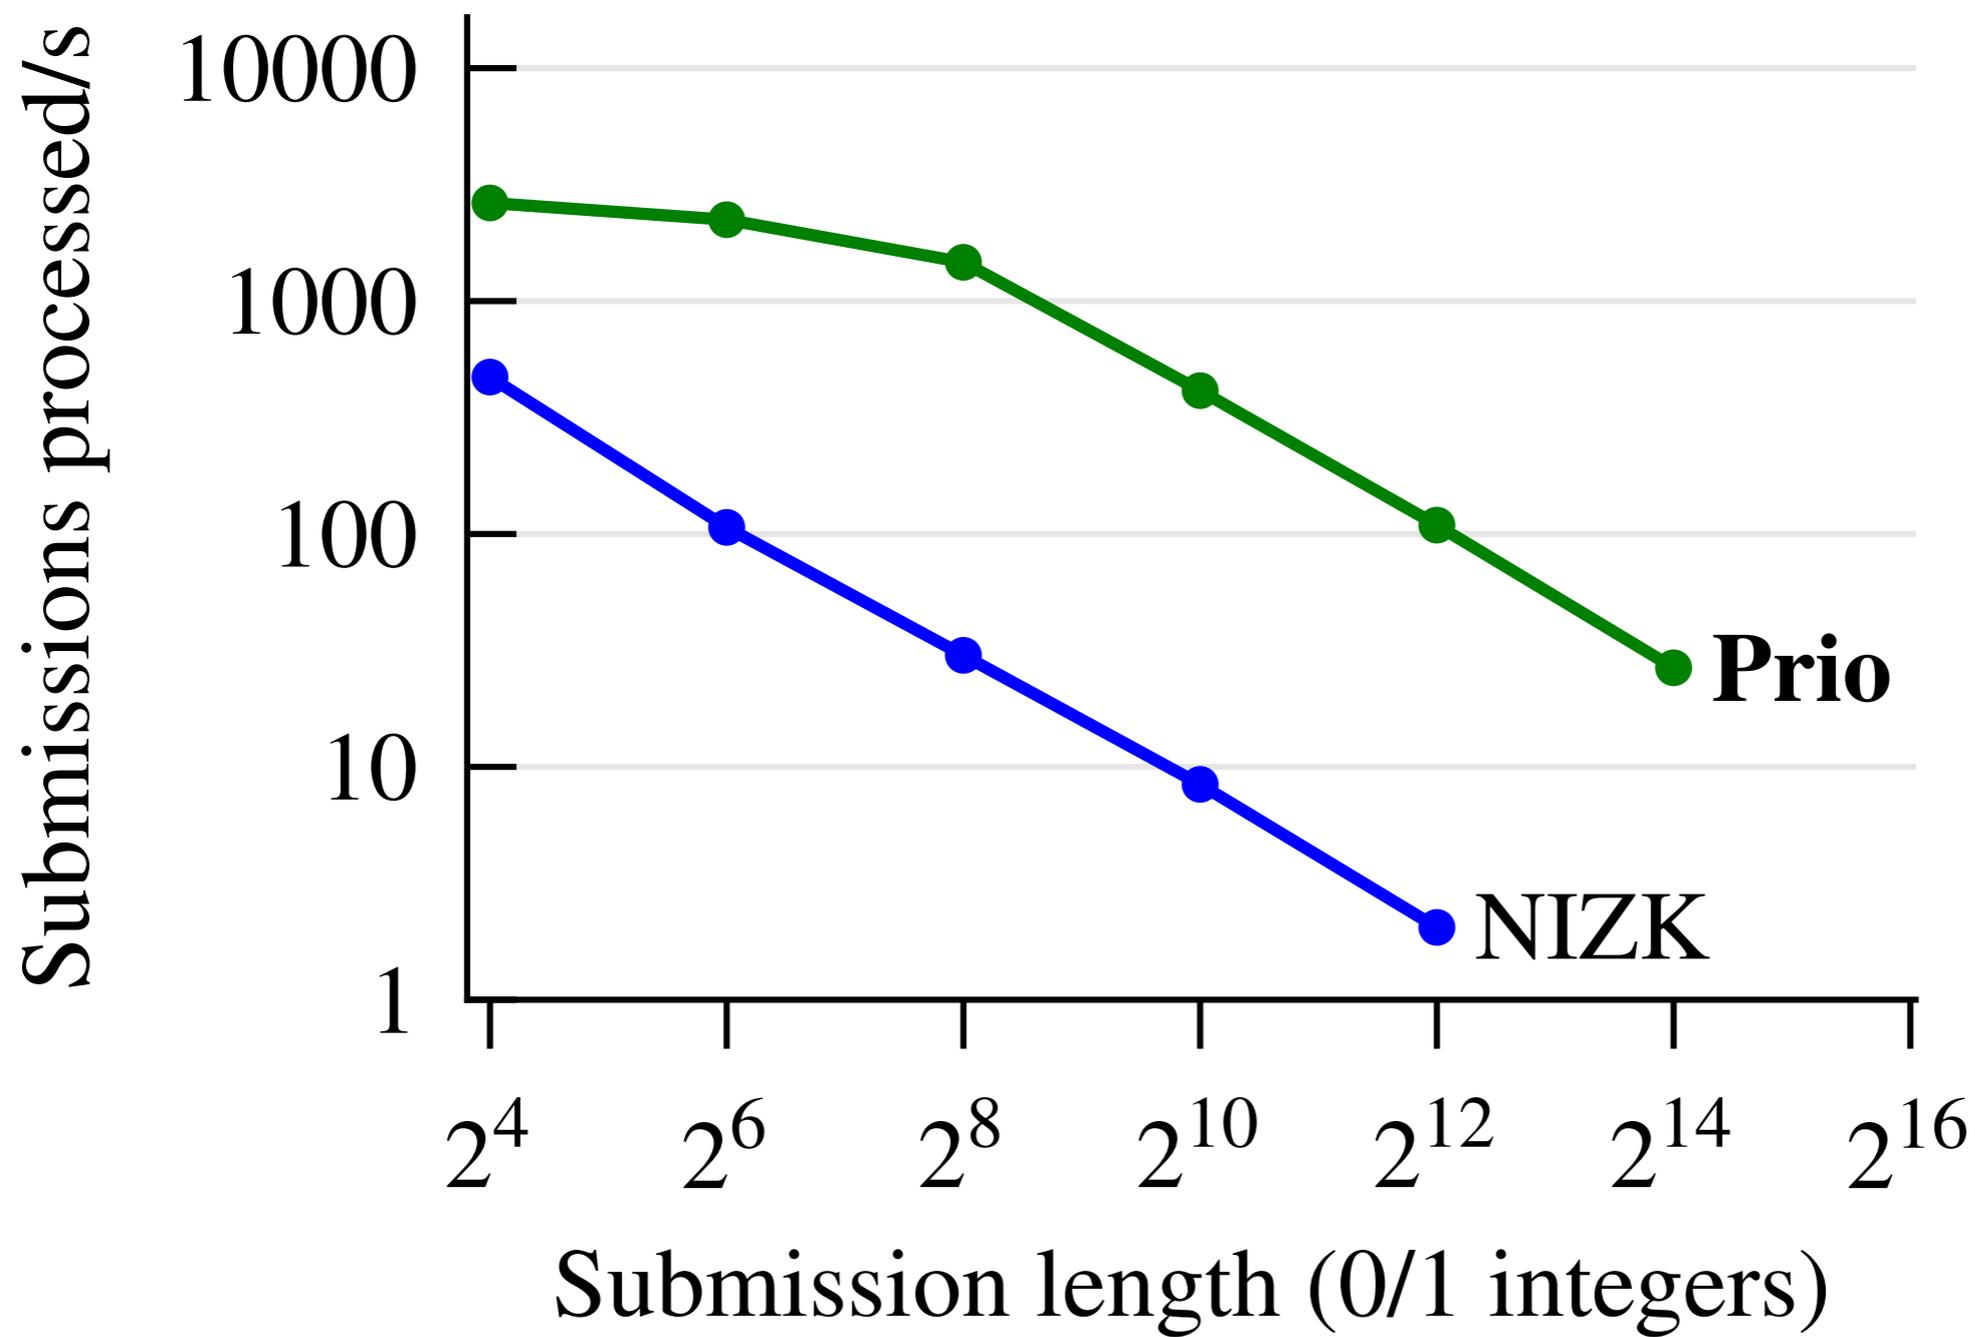
1. No privacy
2. No robustness (“straw man”)
3. Prio (privacy + robustness)
4. NIZK (privacy + robustness)



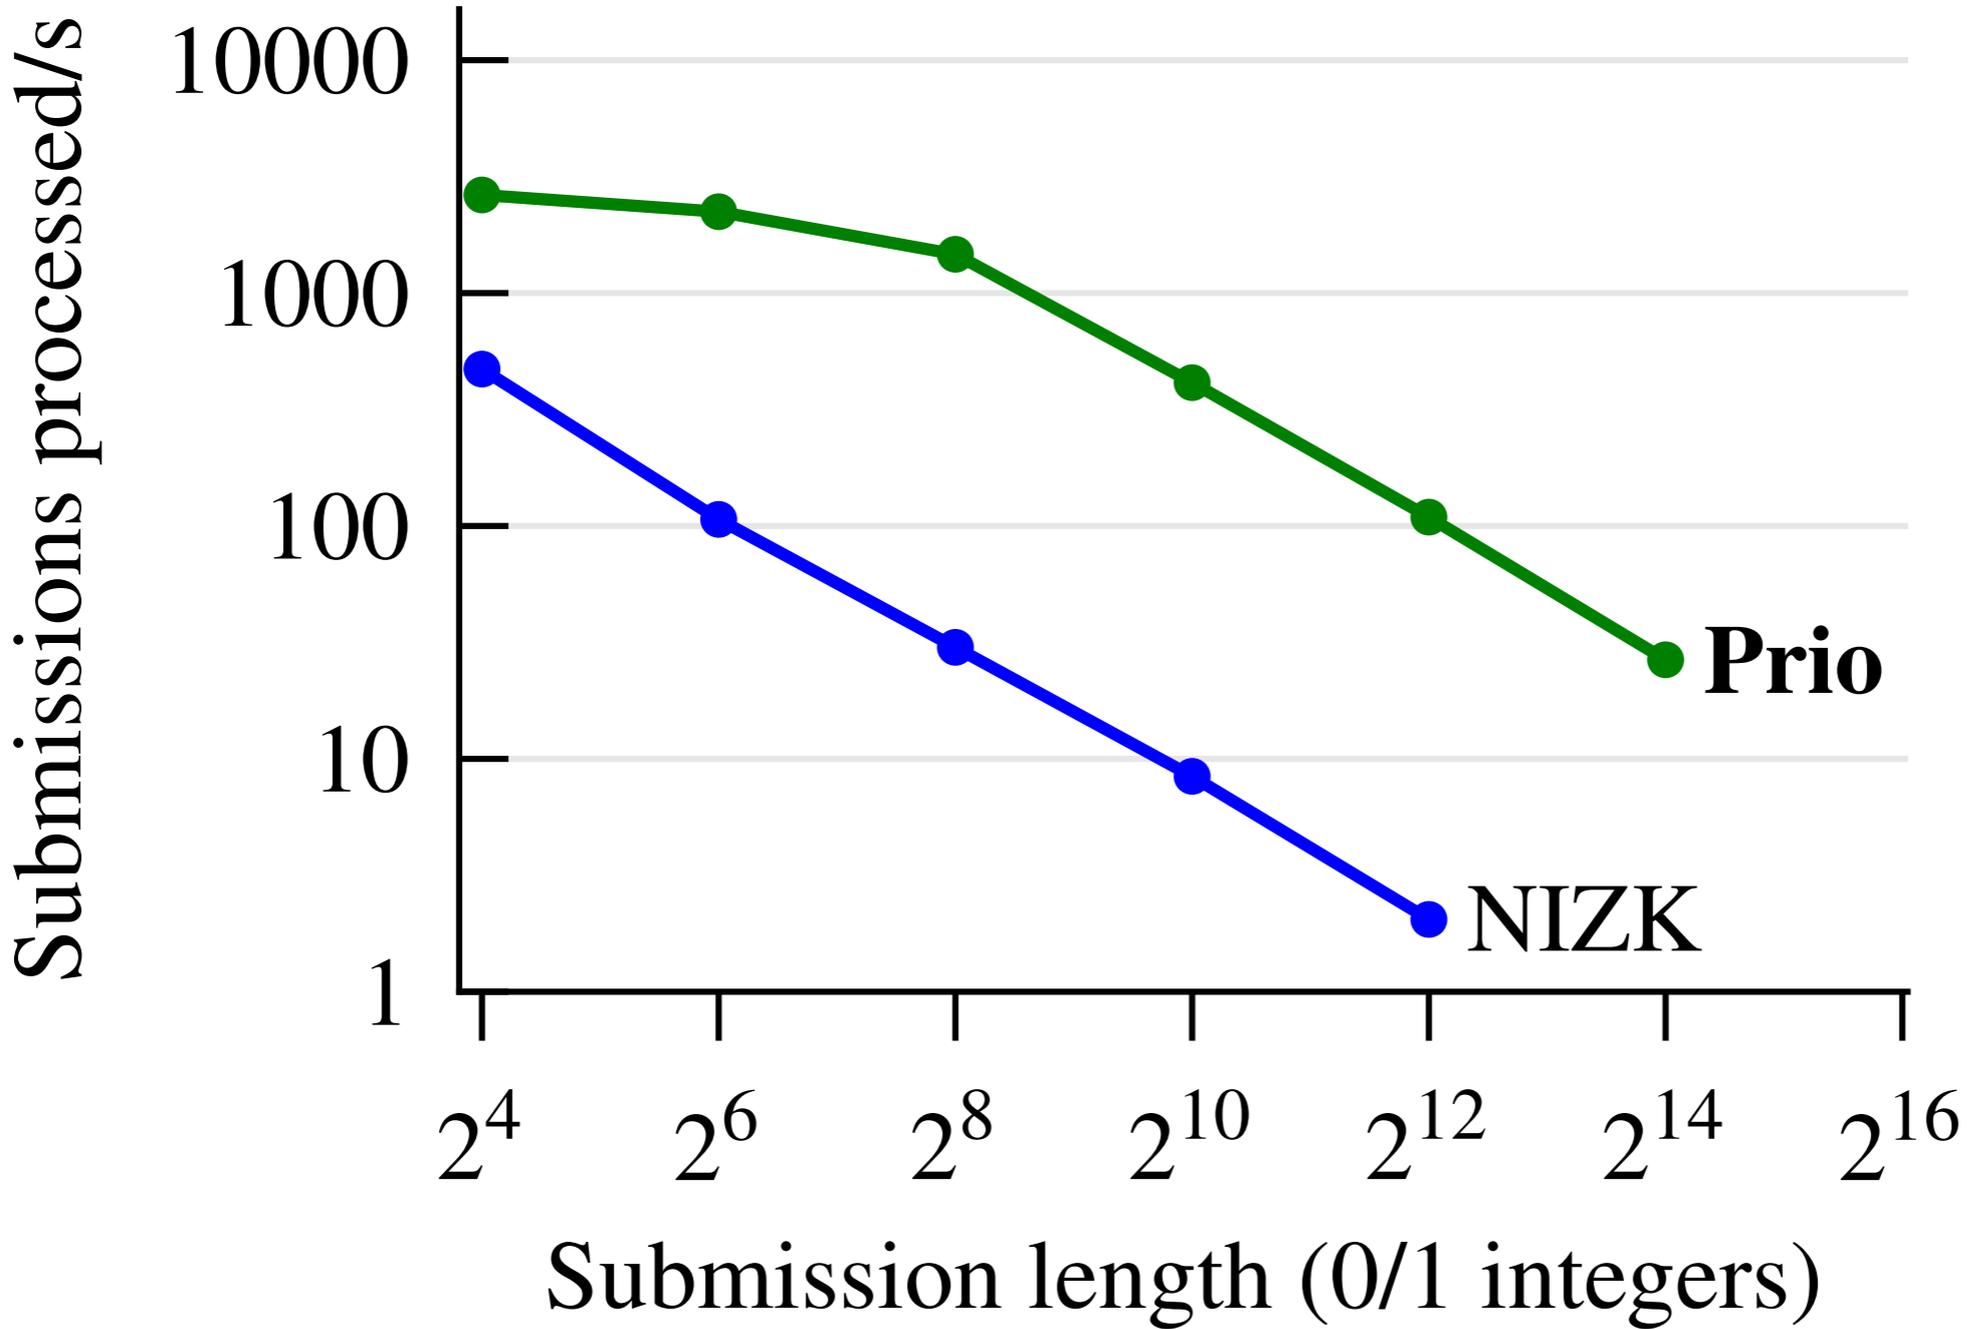
Five-server cluster in five
Amazon data centers



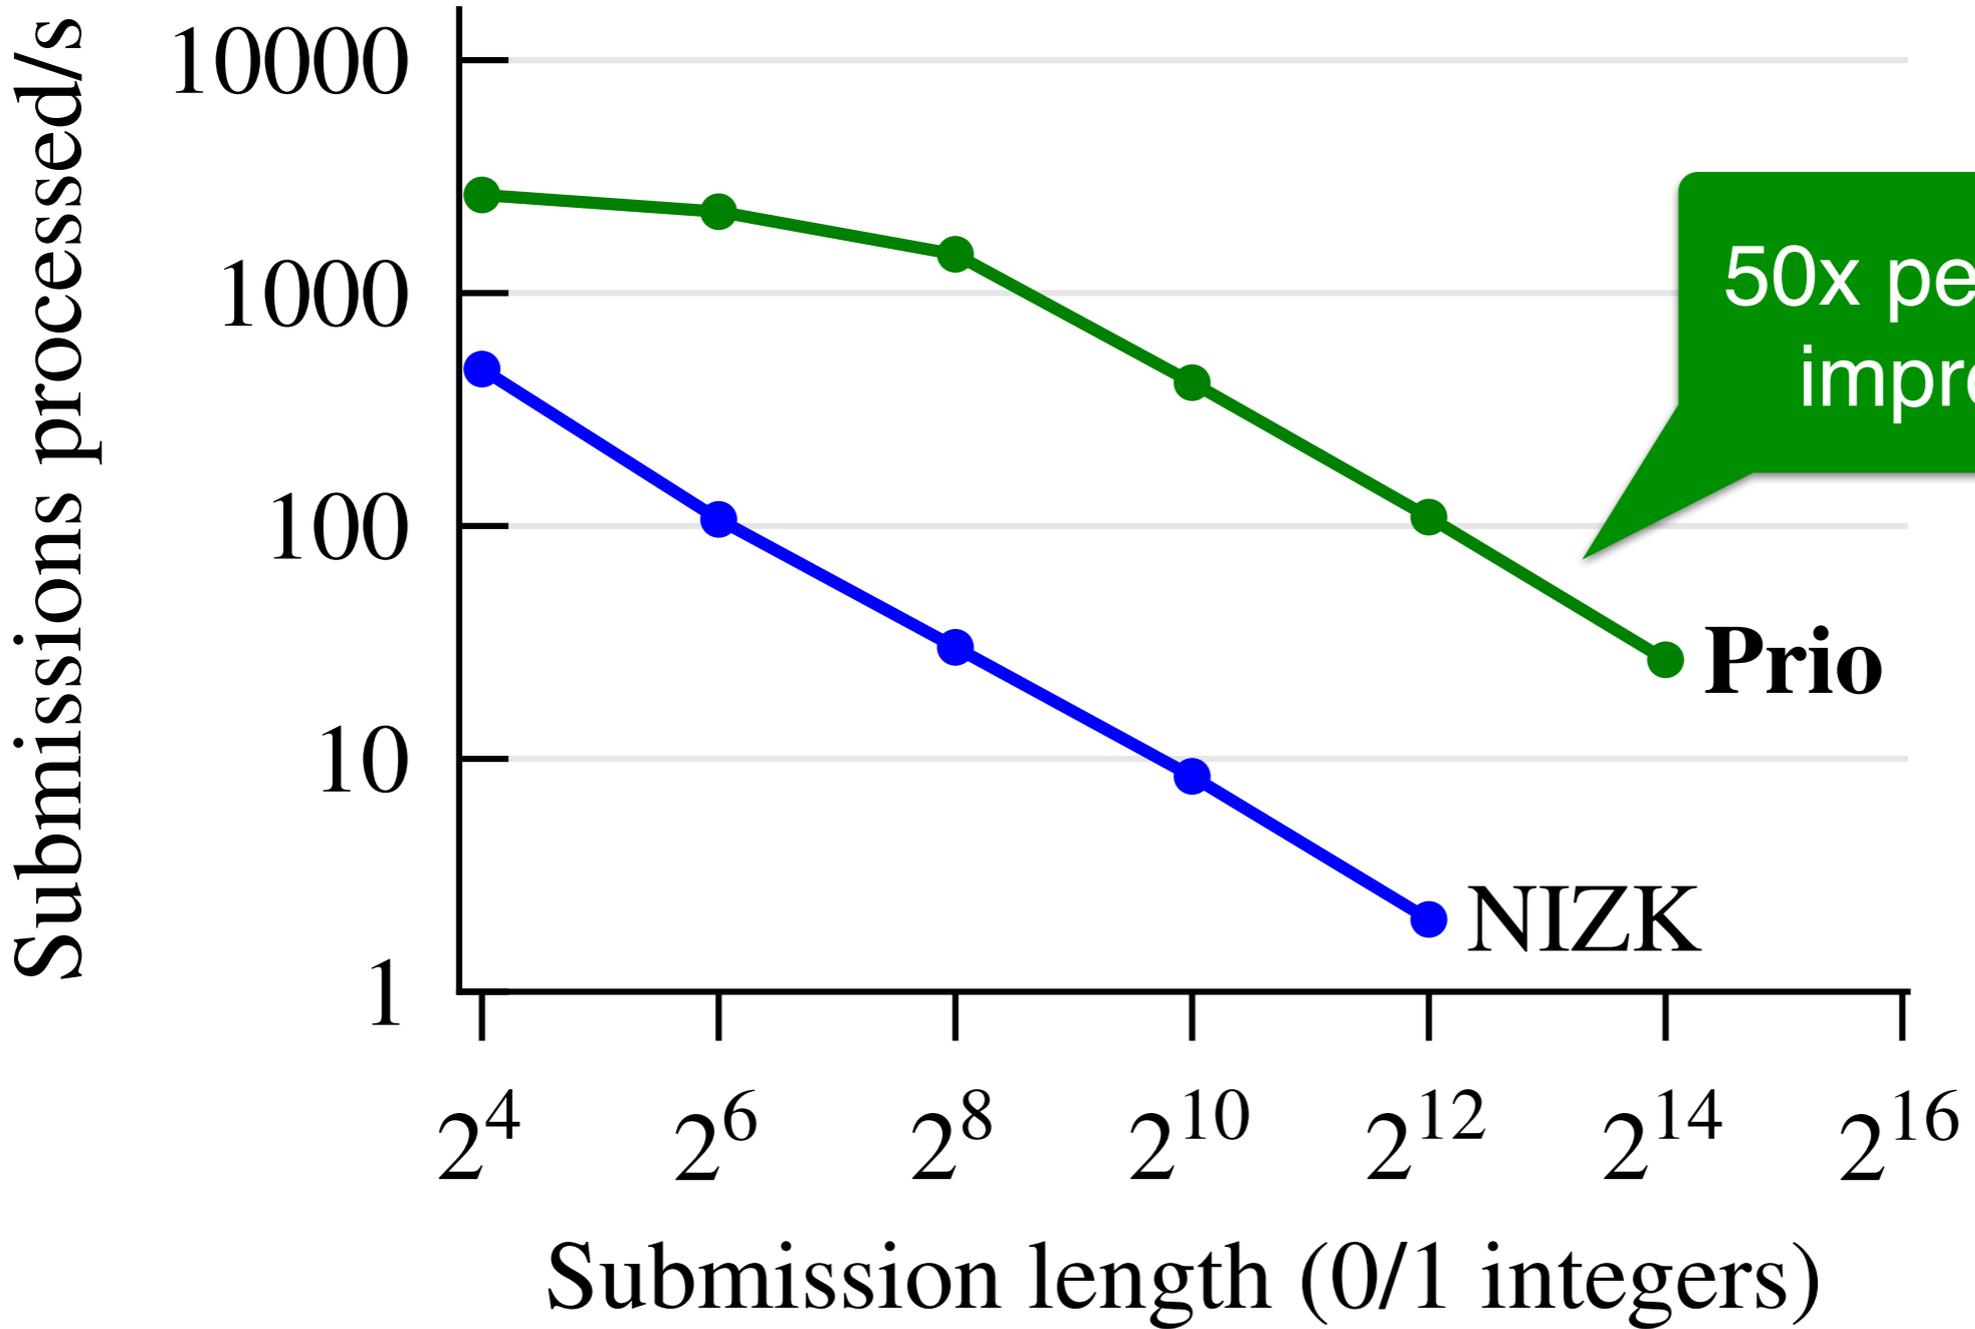
Five-server cluster in five
Amazon data centers



Five-server cluster in five Amazon data centers



Five-server cluster in five Amazon data centers

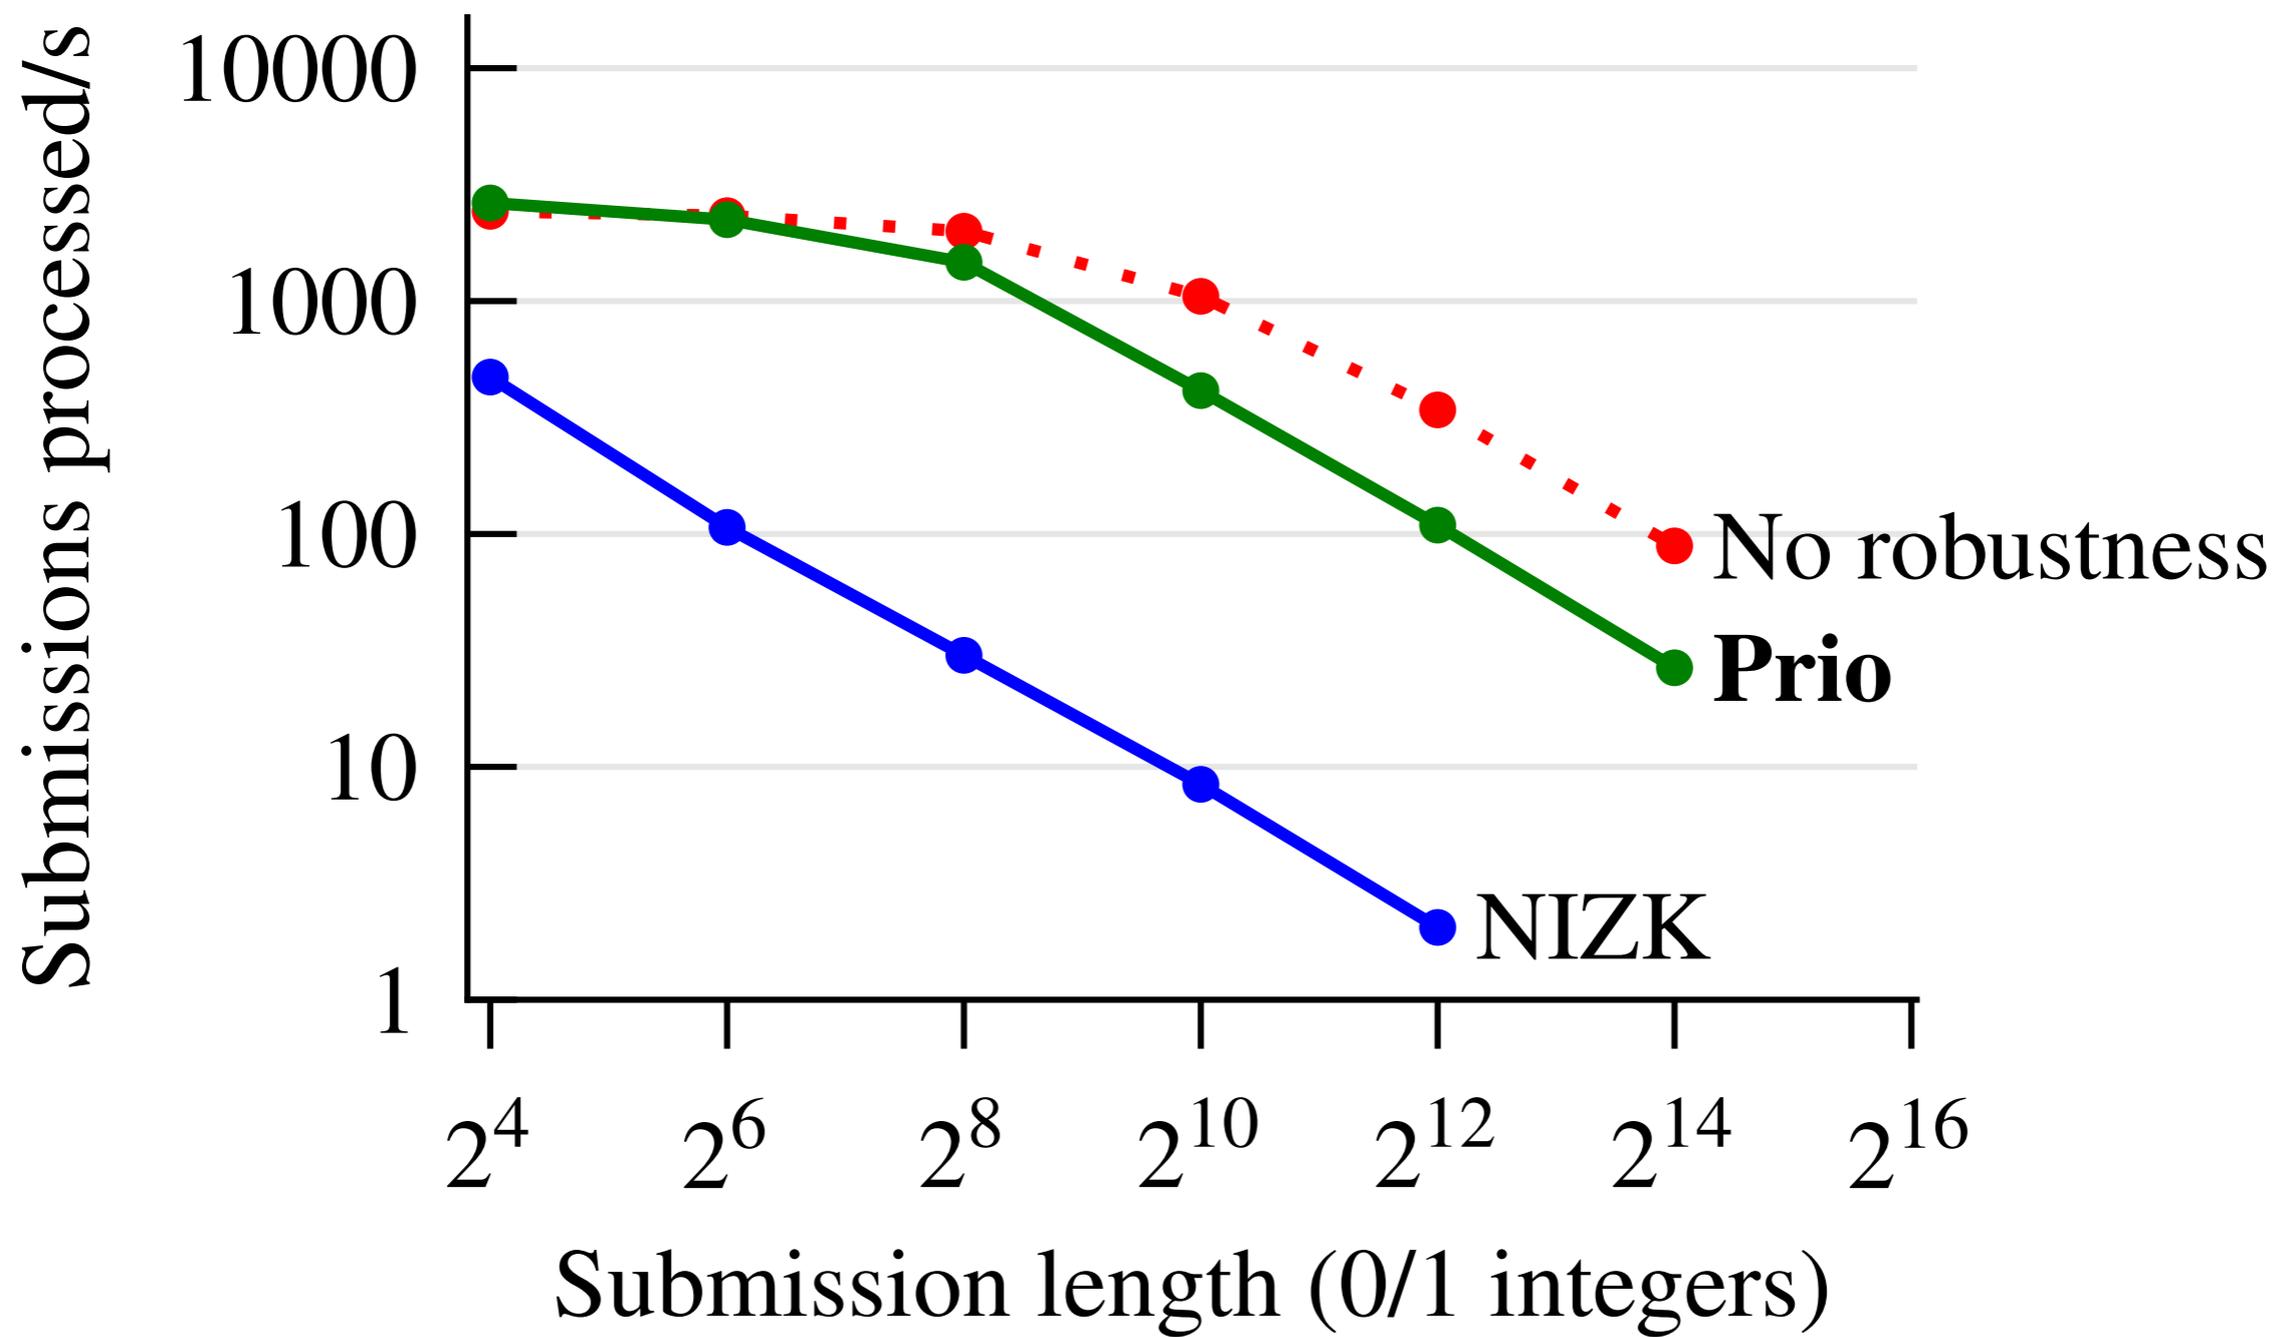


50x performance improvement

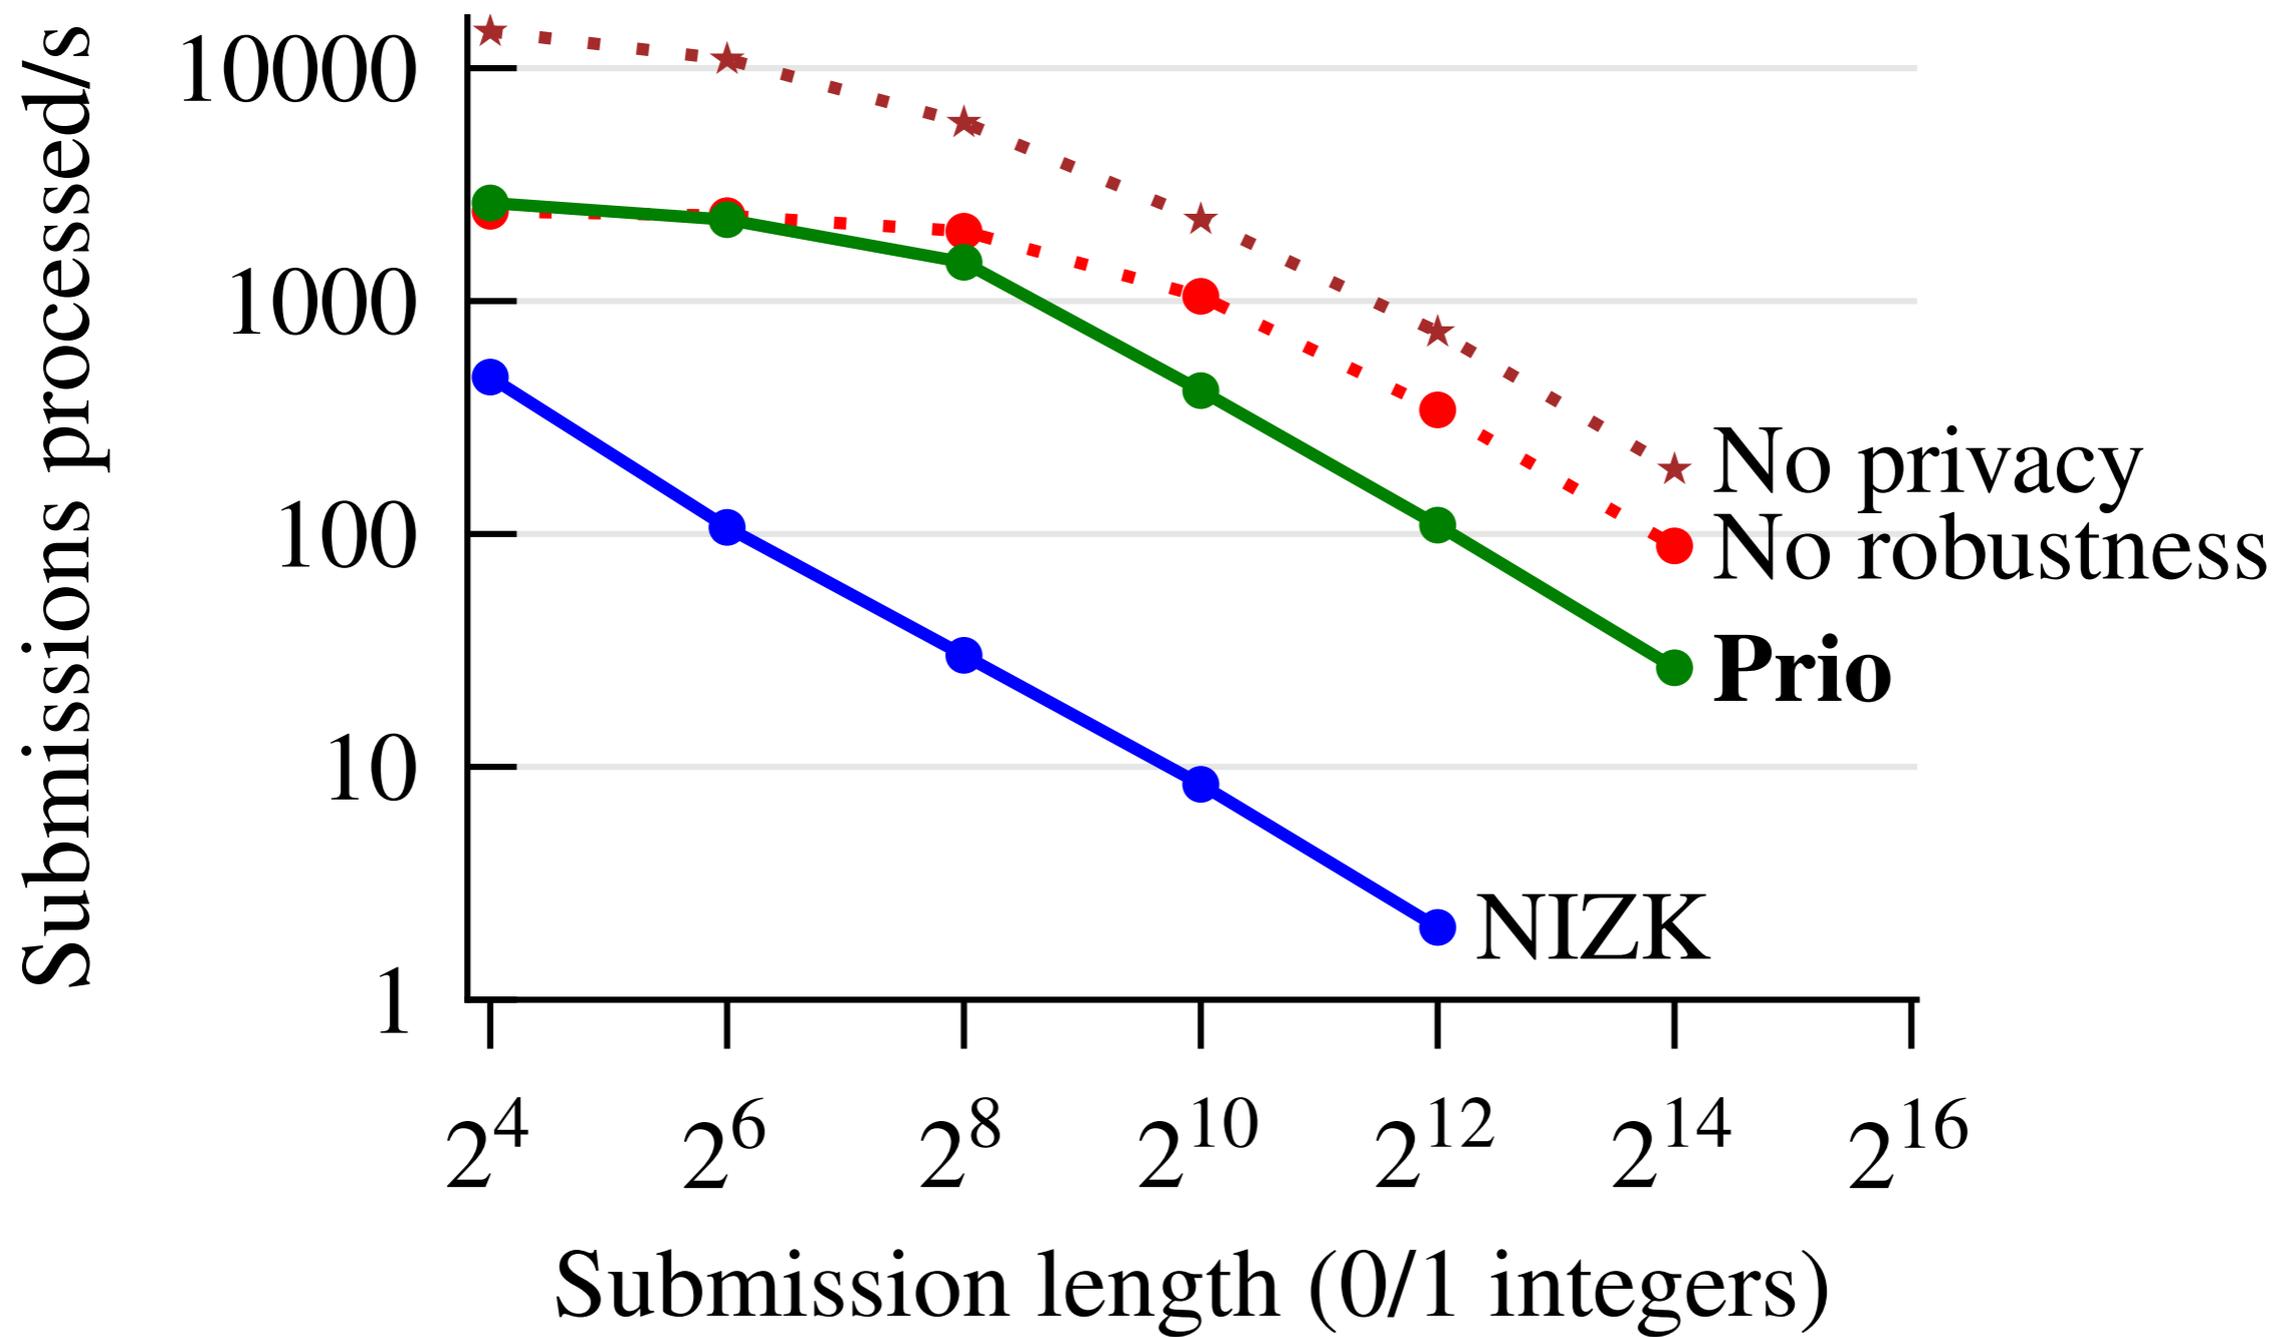
Prio

NIZK

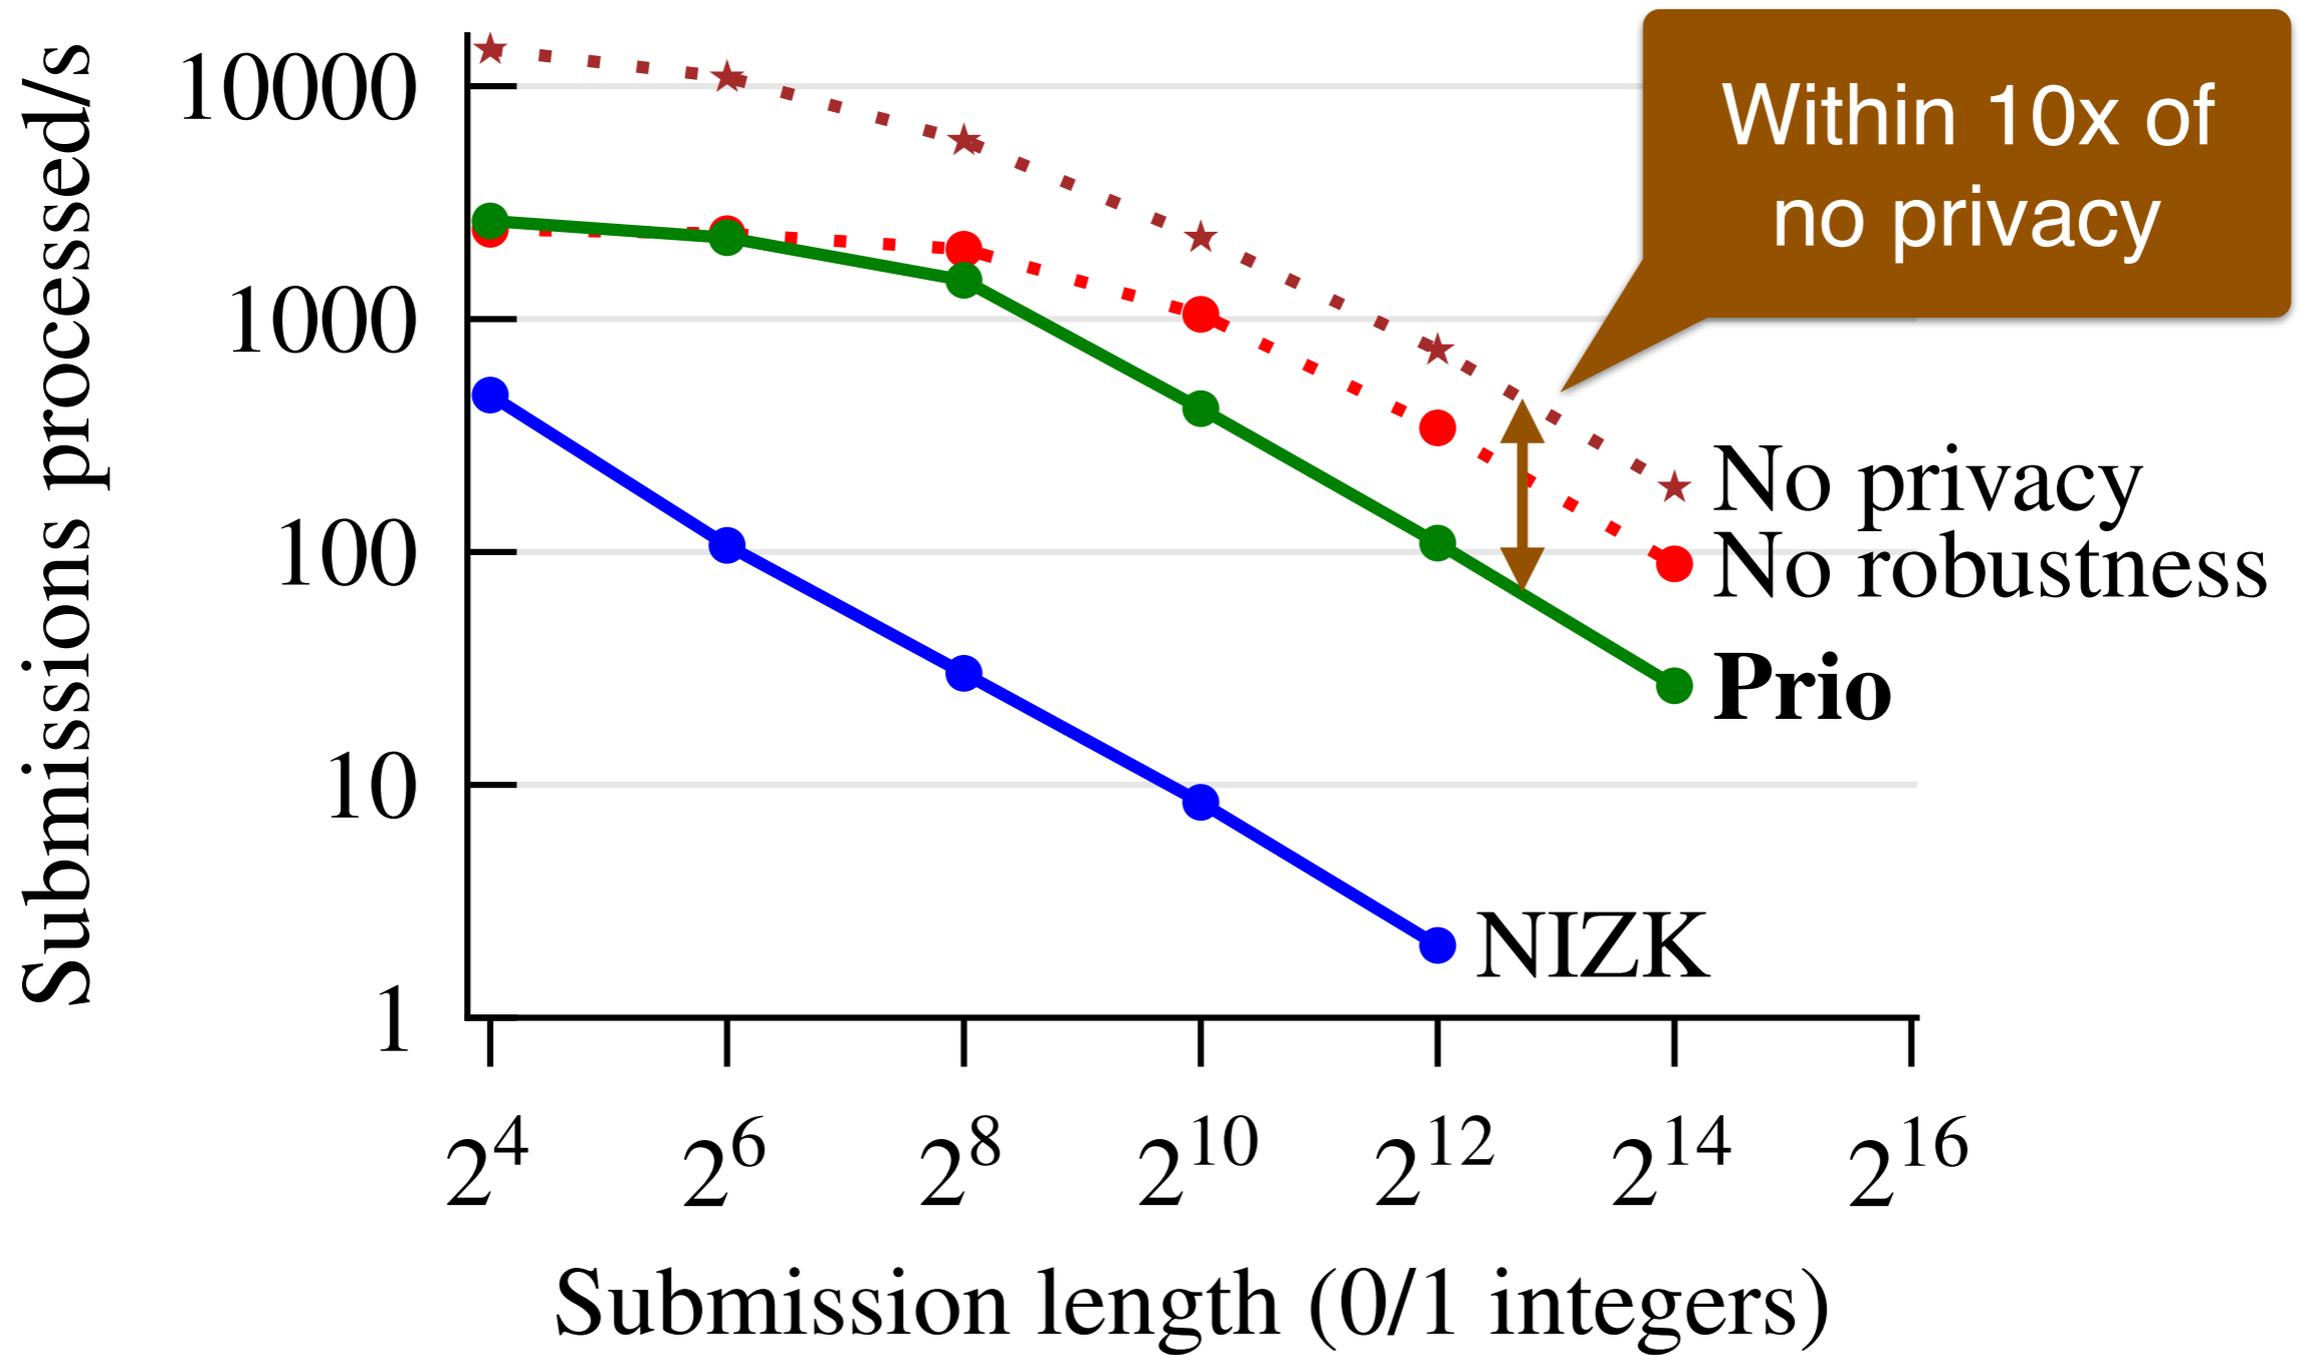
Five-server cluster in five
Amazon data centers

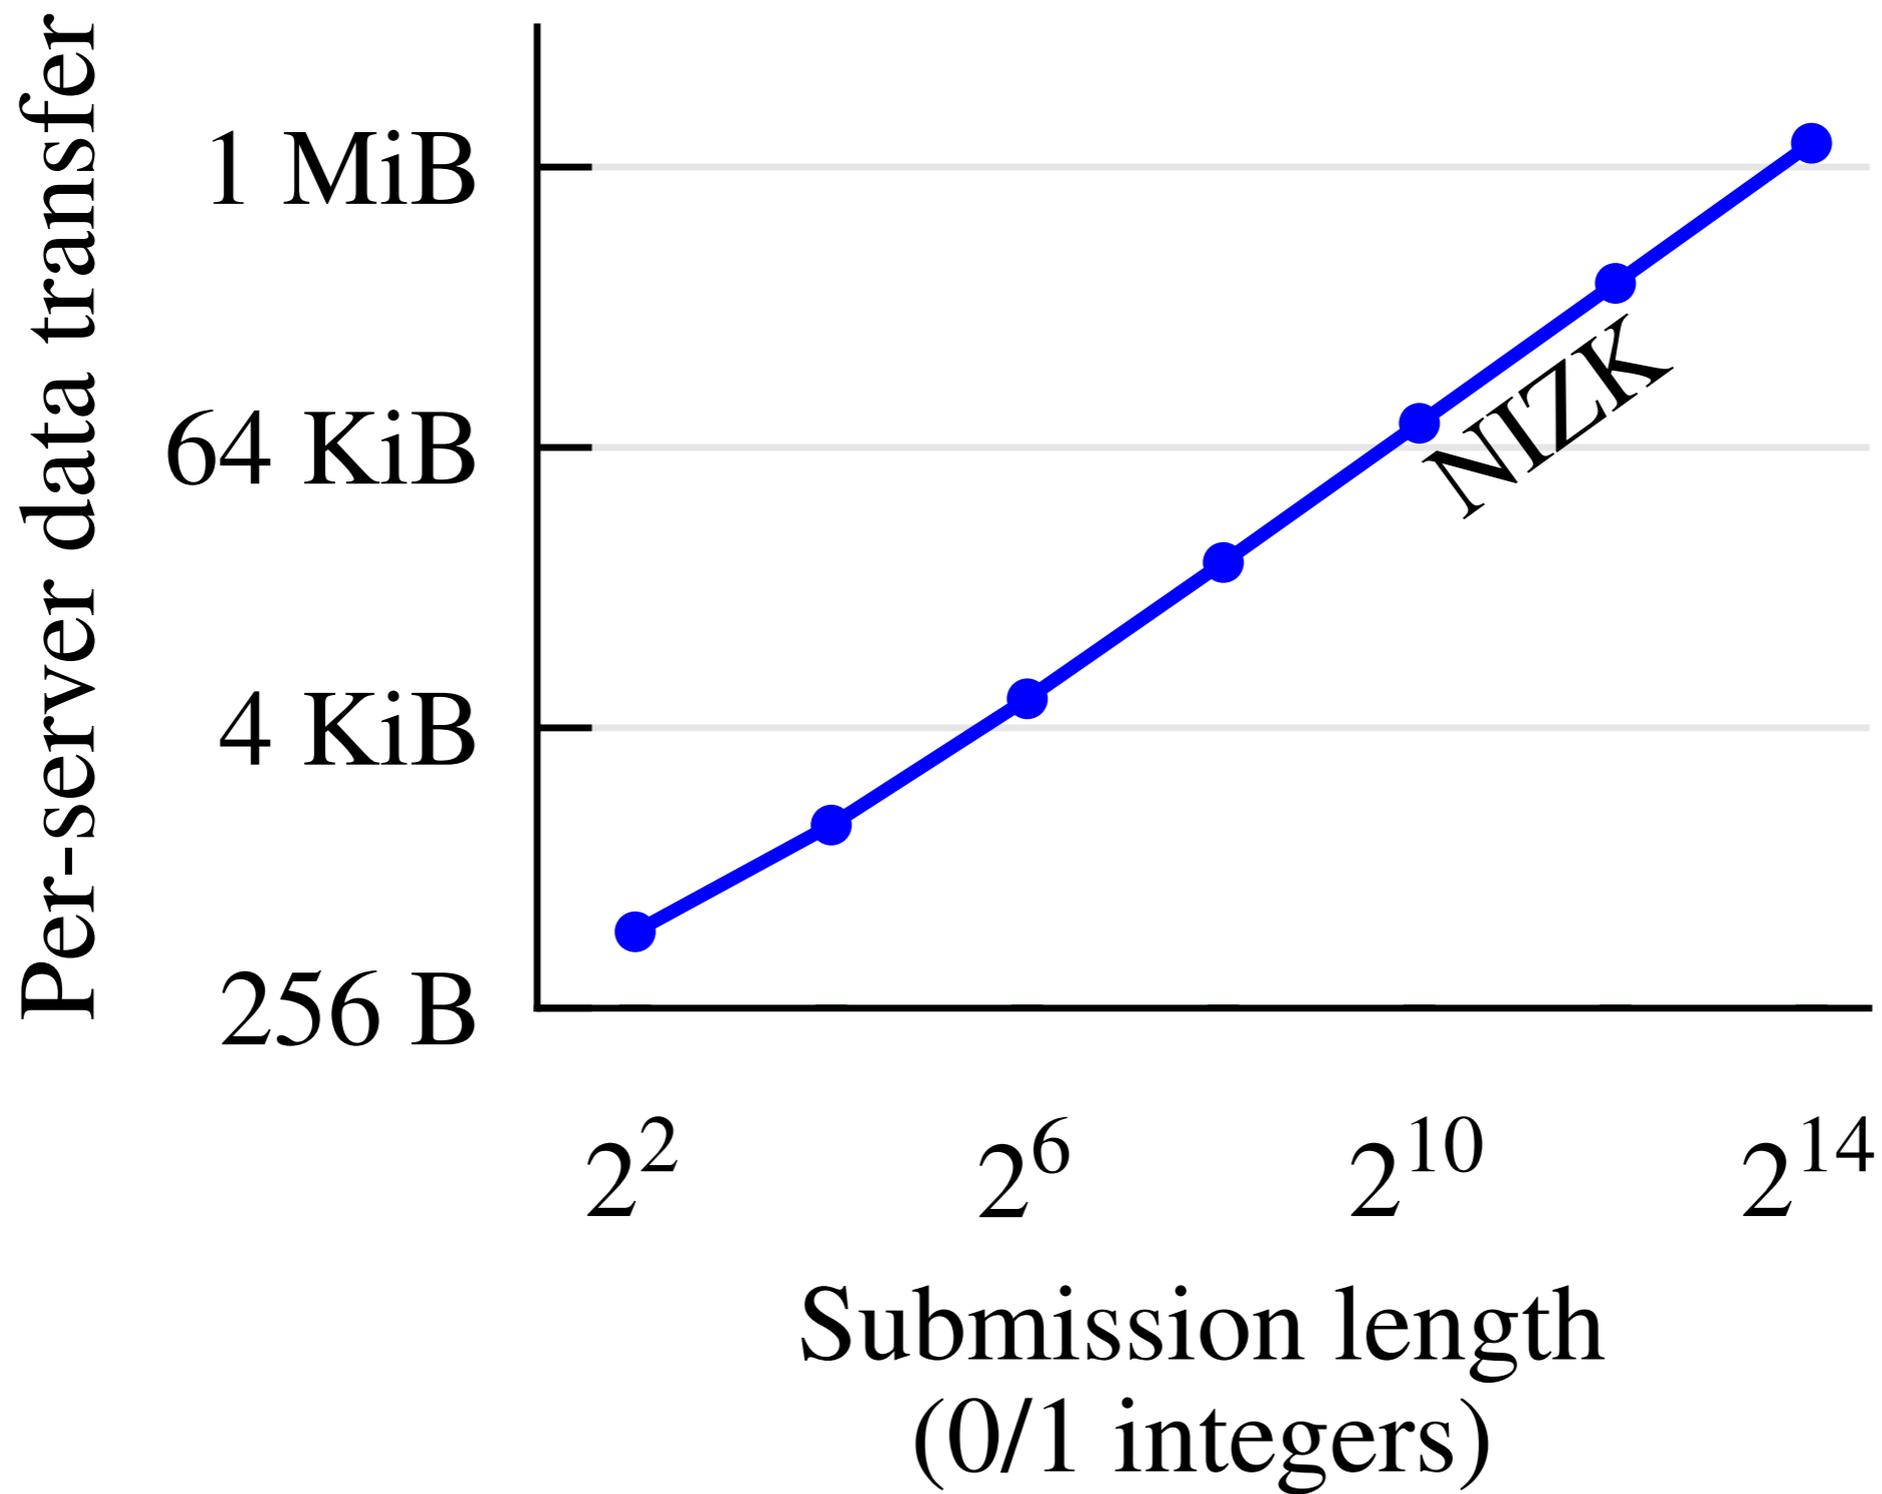


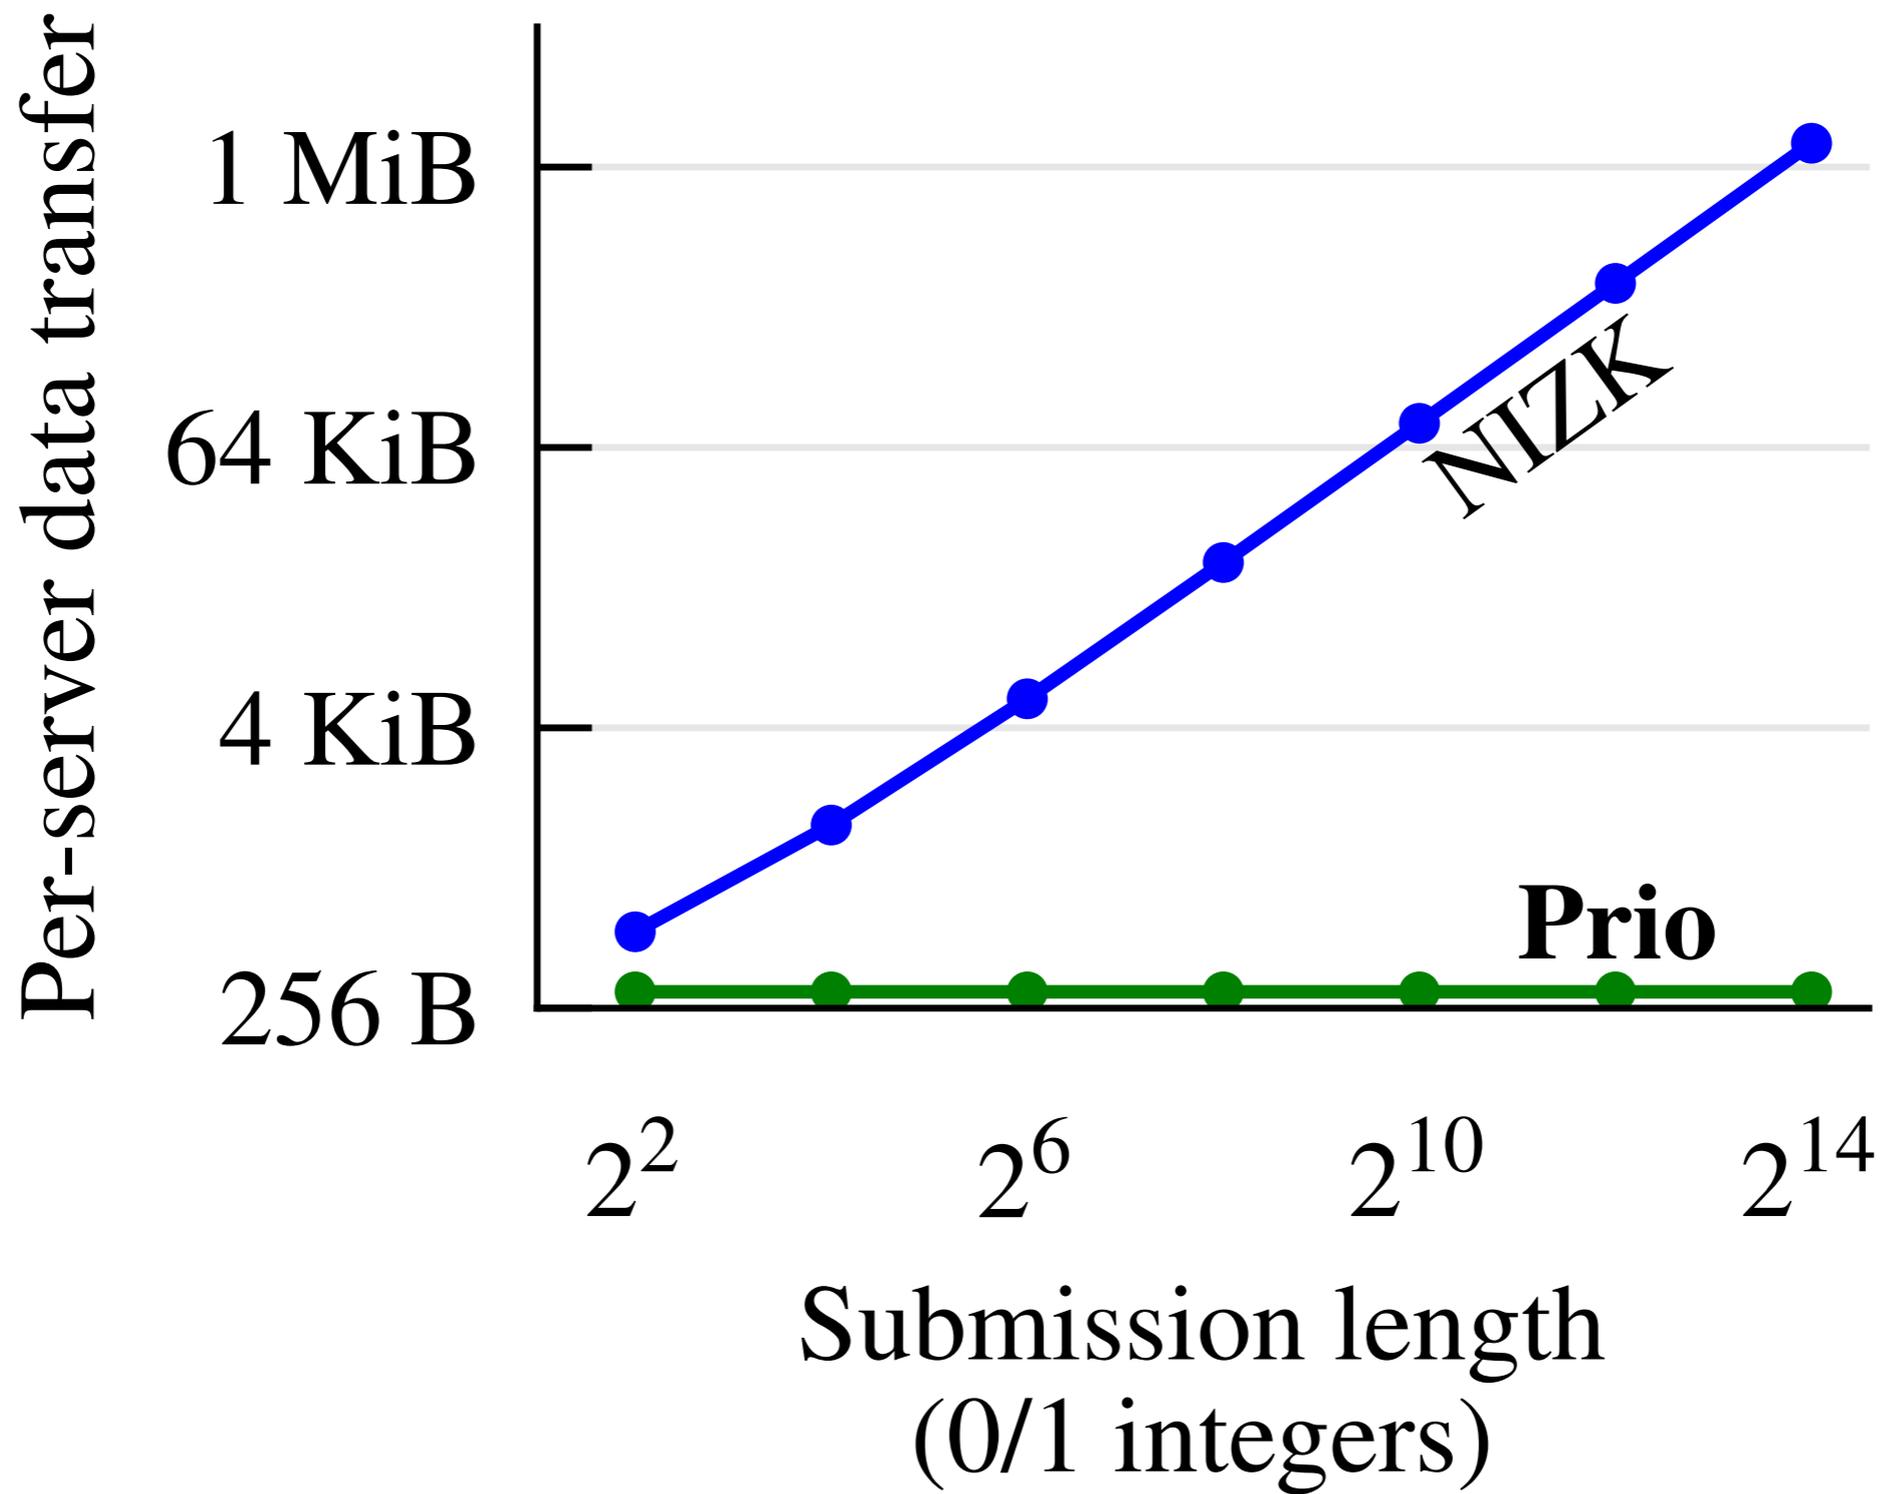
Five-server cluster in five
Amazon data centers

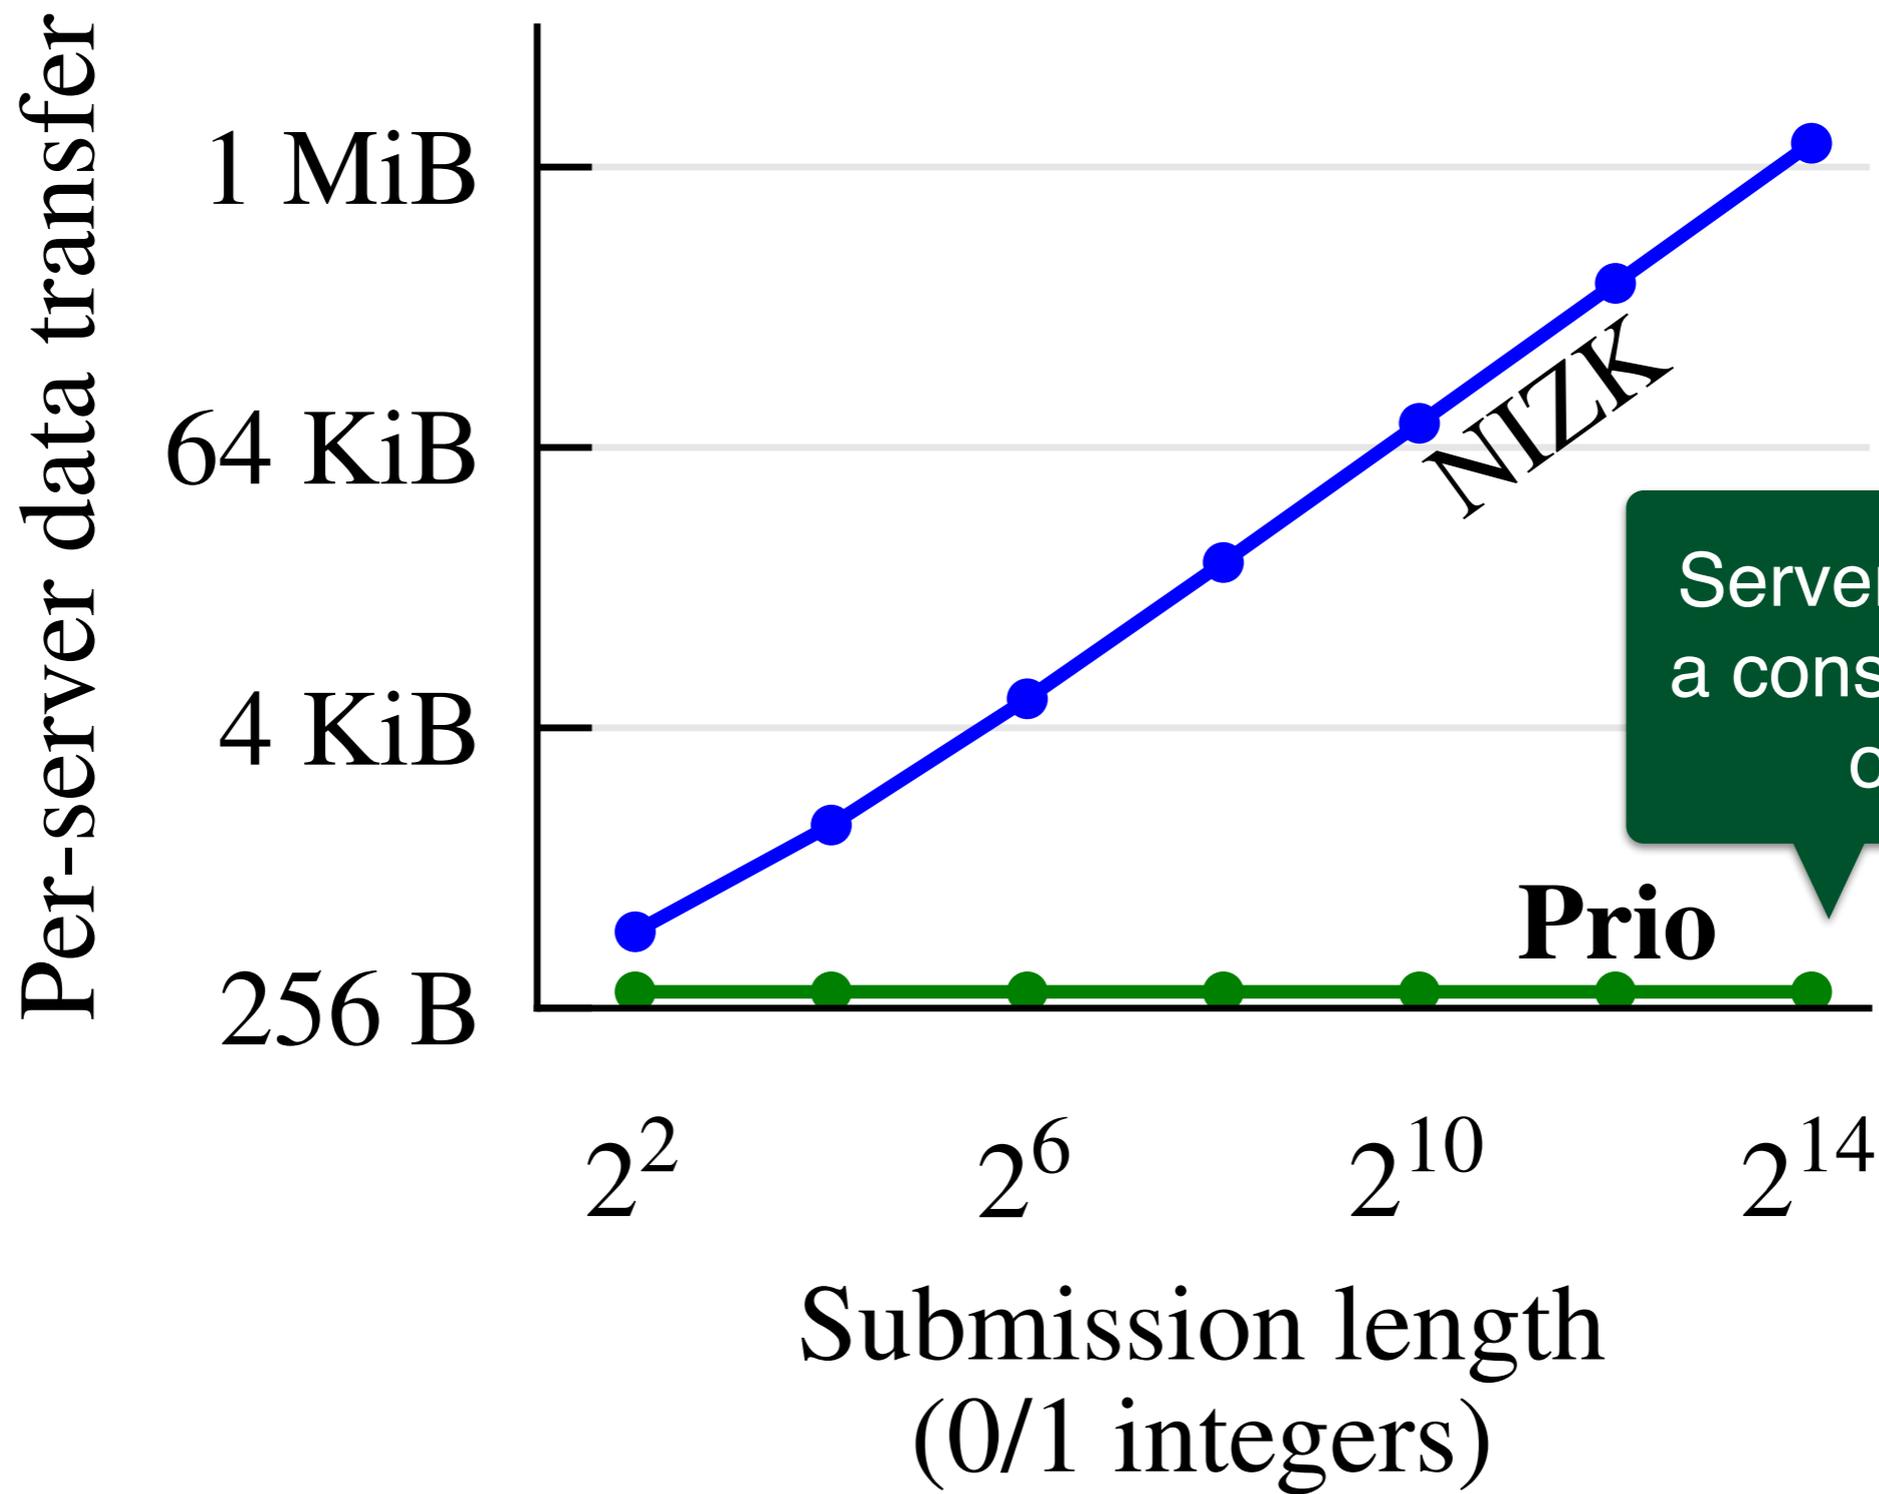


Five-server cluster in five Amazon data centers









Servers exchange a constant number of bytes

Prio

Outline

- Background: The private aggregation problem
- A straw-man solution for private sums
- Providing robustness with SNIPs
- **Evaluation**
- Encodings for complex aggregates

Outline

- Background: The private aggregation problem
- A straw-man solution for private sums
- Providing robustness with SNIPs
- Evaluation
- **Encodings for complex aggregates**

Known techniques: Complex statistics

If you can compute private sums, you can compute many other interesting aggregates using known techniques

[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]

Known techniques: Complex statistics

If you can compute private sums, you can compute many other interesting aggregates using known techniques

[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]

Contribution 2:
SNIP-friendly encodings
for these statistics

Known techniques: Complex statistics

If you can compute private sums, you can compute many other interesting aggregates using known techniques

[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]

Contribution 2:
SNIP-friendly encodings
for these statistics

Prio can't compute all
statistics efficiently

Known techniques: Complex statistics

If you can compute private sums, you can compute many other interesting aggregates using known techniques

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

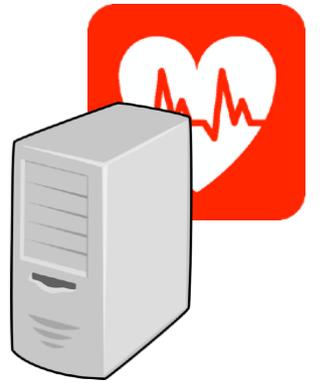
- Average
- Variance
- Standard deviation
- Most popular (approx)
- “Heavy hitters” (approx)
- Min and max (approx)
- Quality of arbitrary regression model (R^2)
- Least-squares regression
- Stochastic gradient descent [Bonawitz et al. 2016]

Contribution 2:
SNIP-friendly encodings
for these statistics

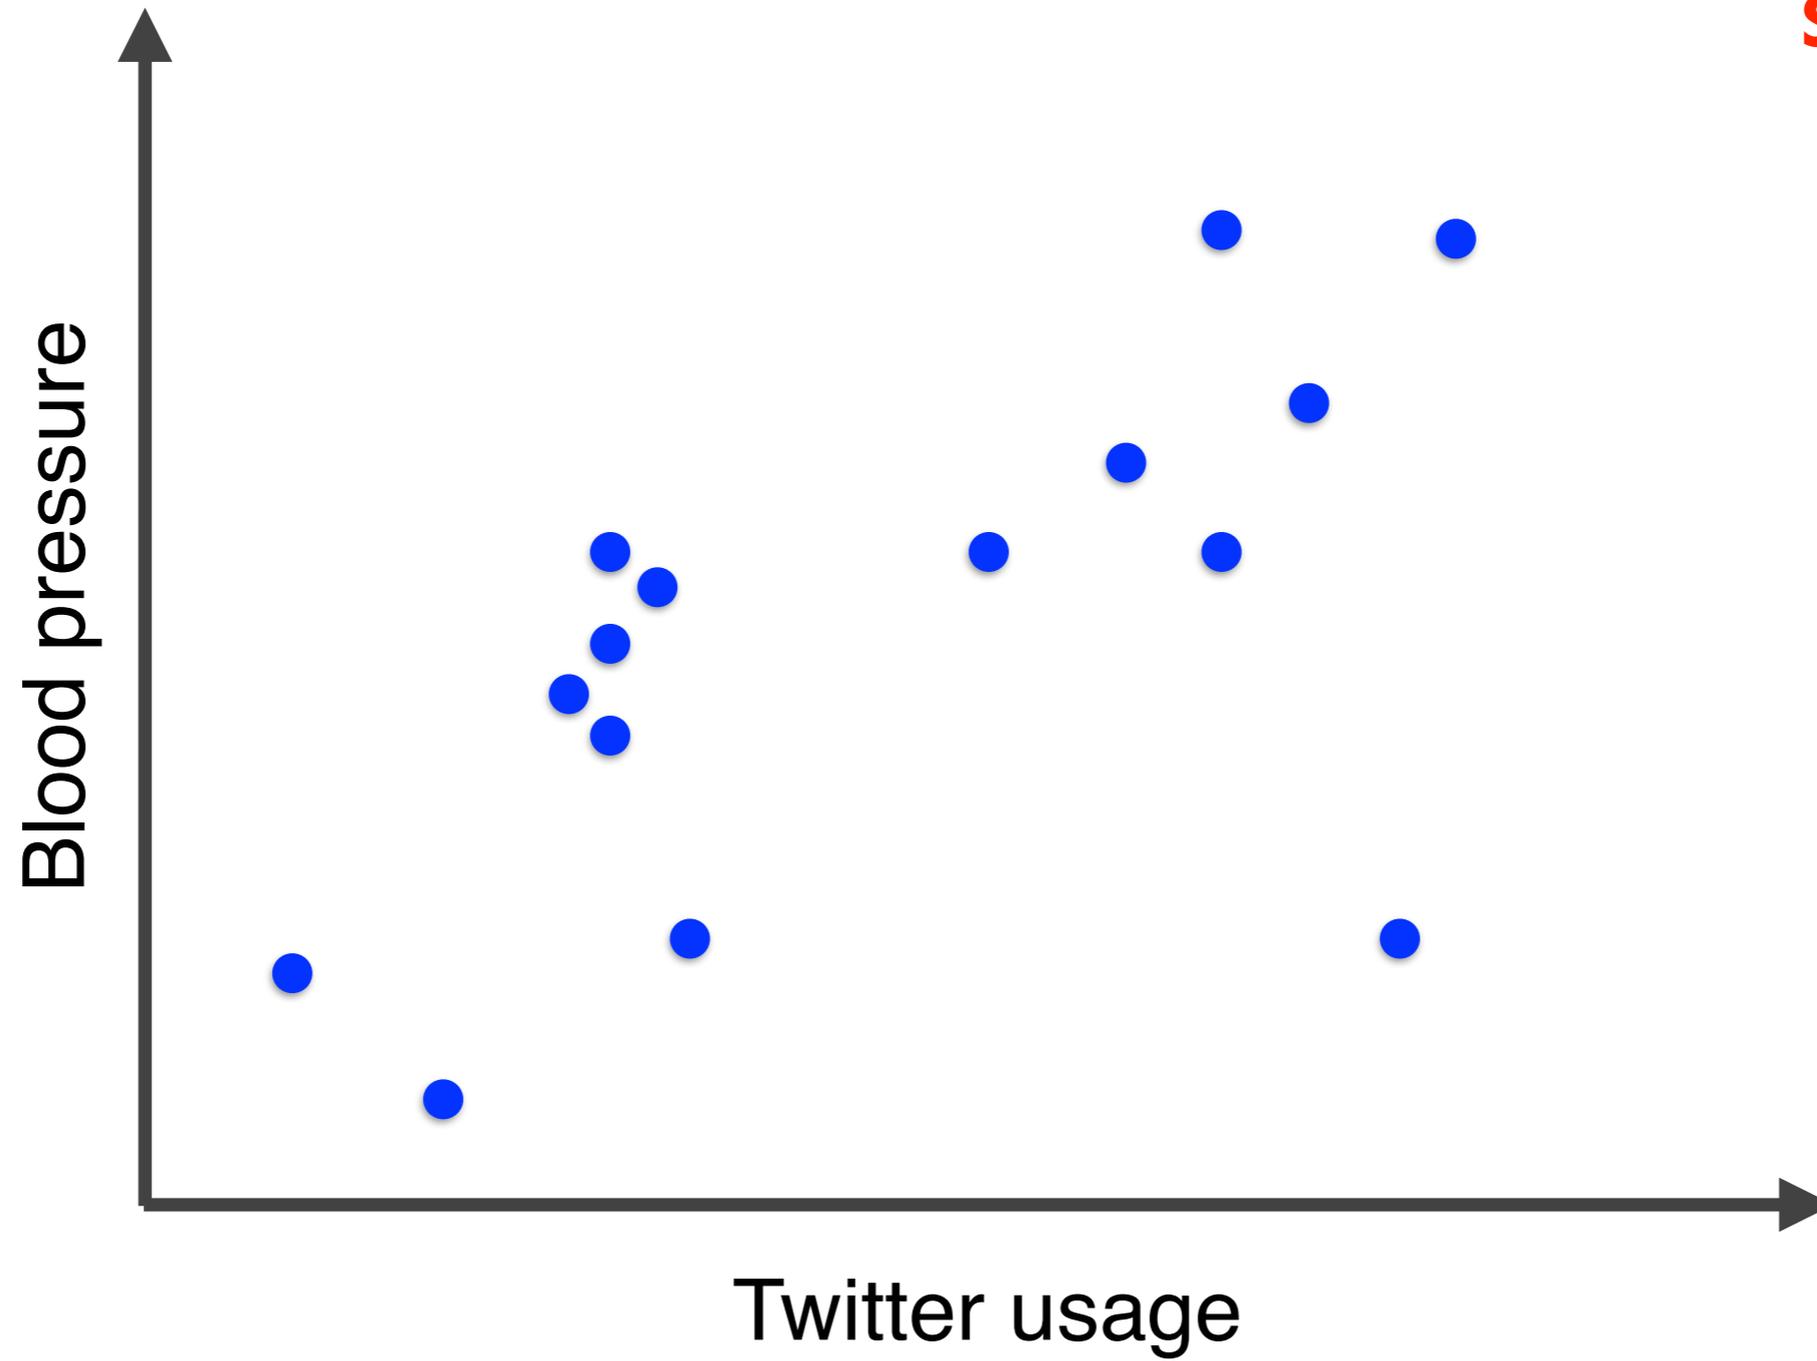
Prio can't compute all
statistics efficiently

See the paper for
the details

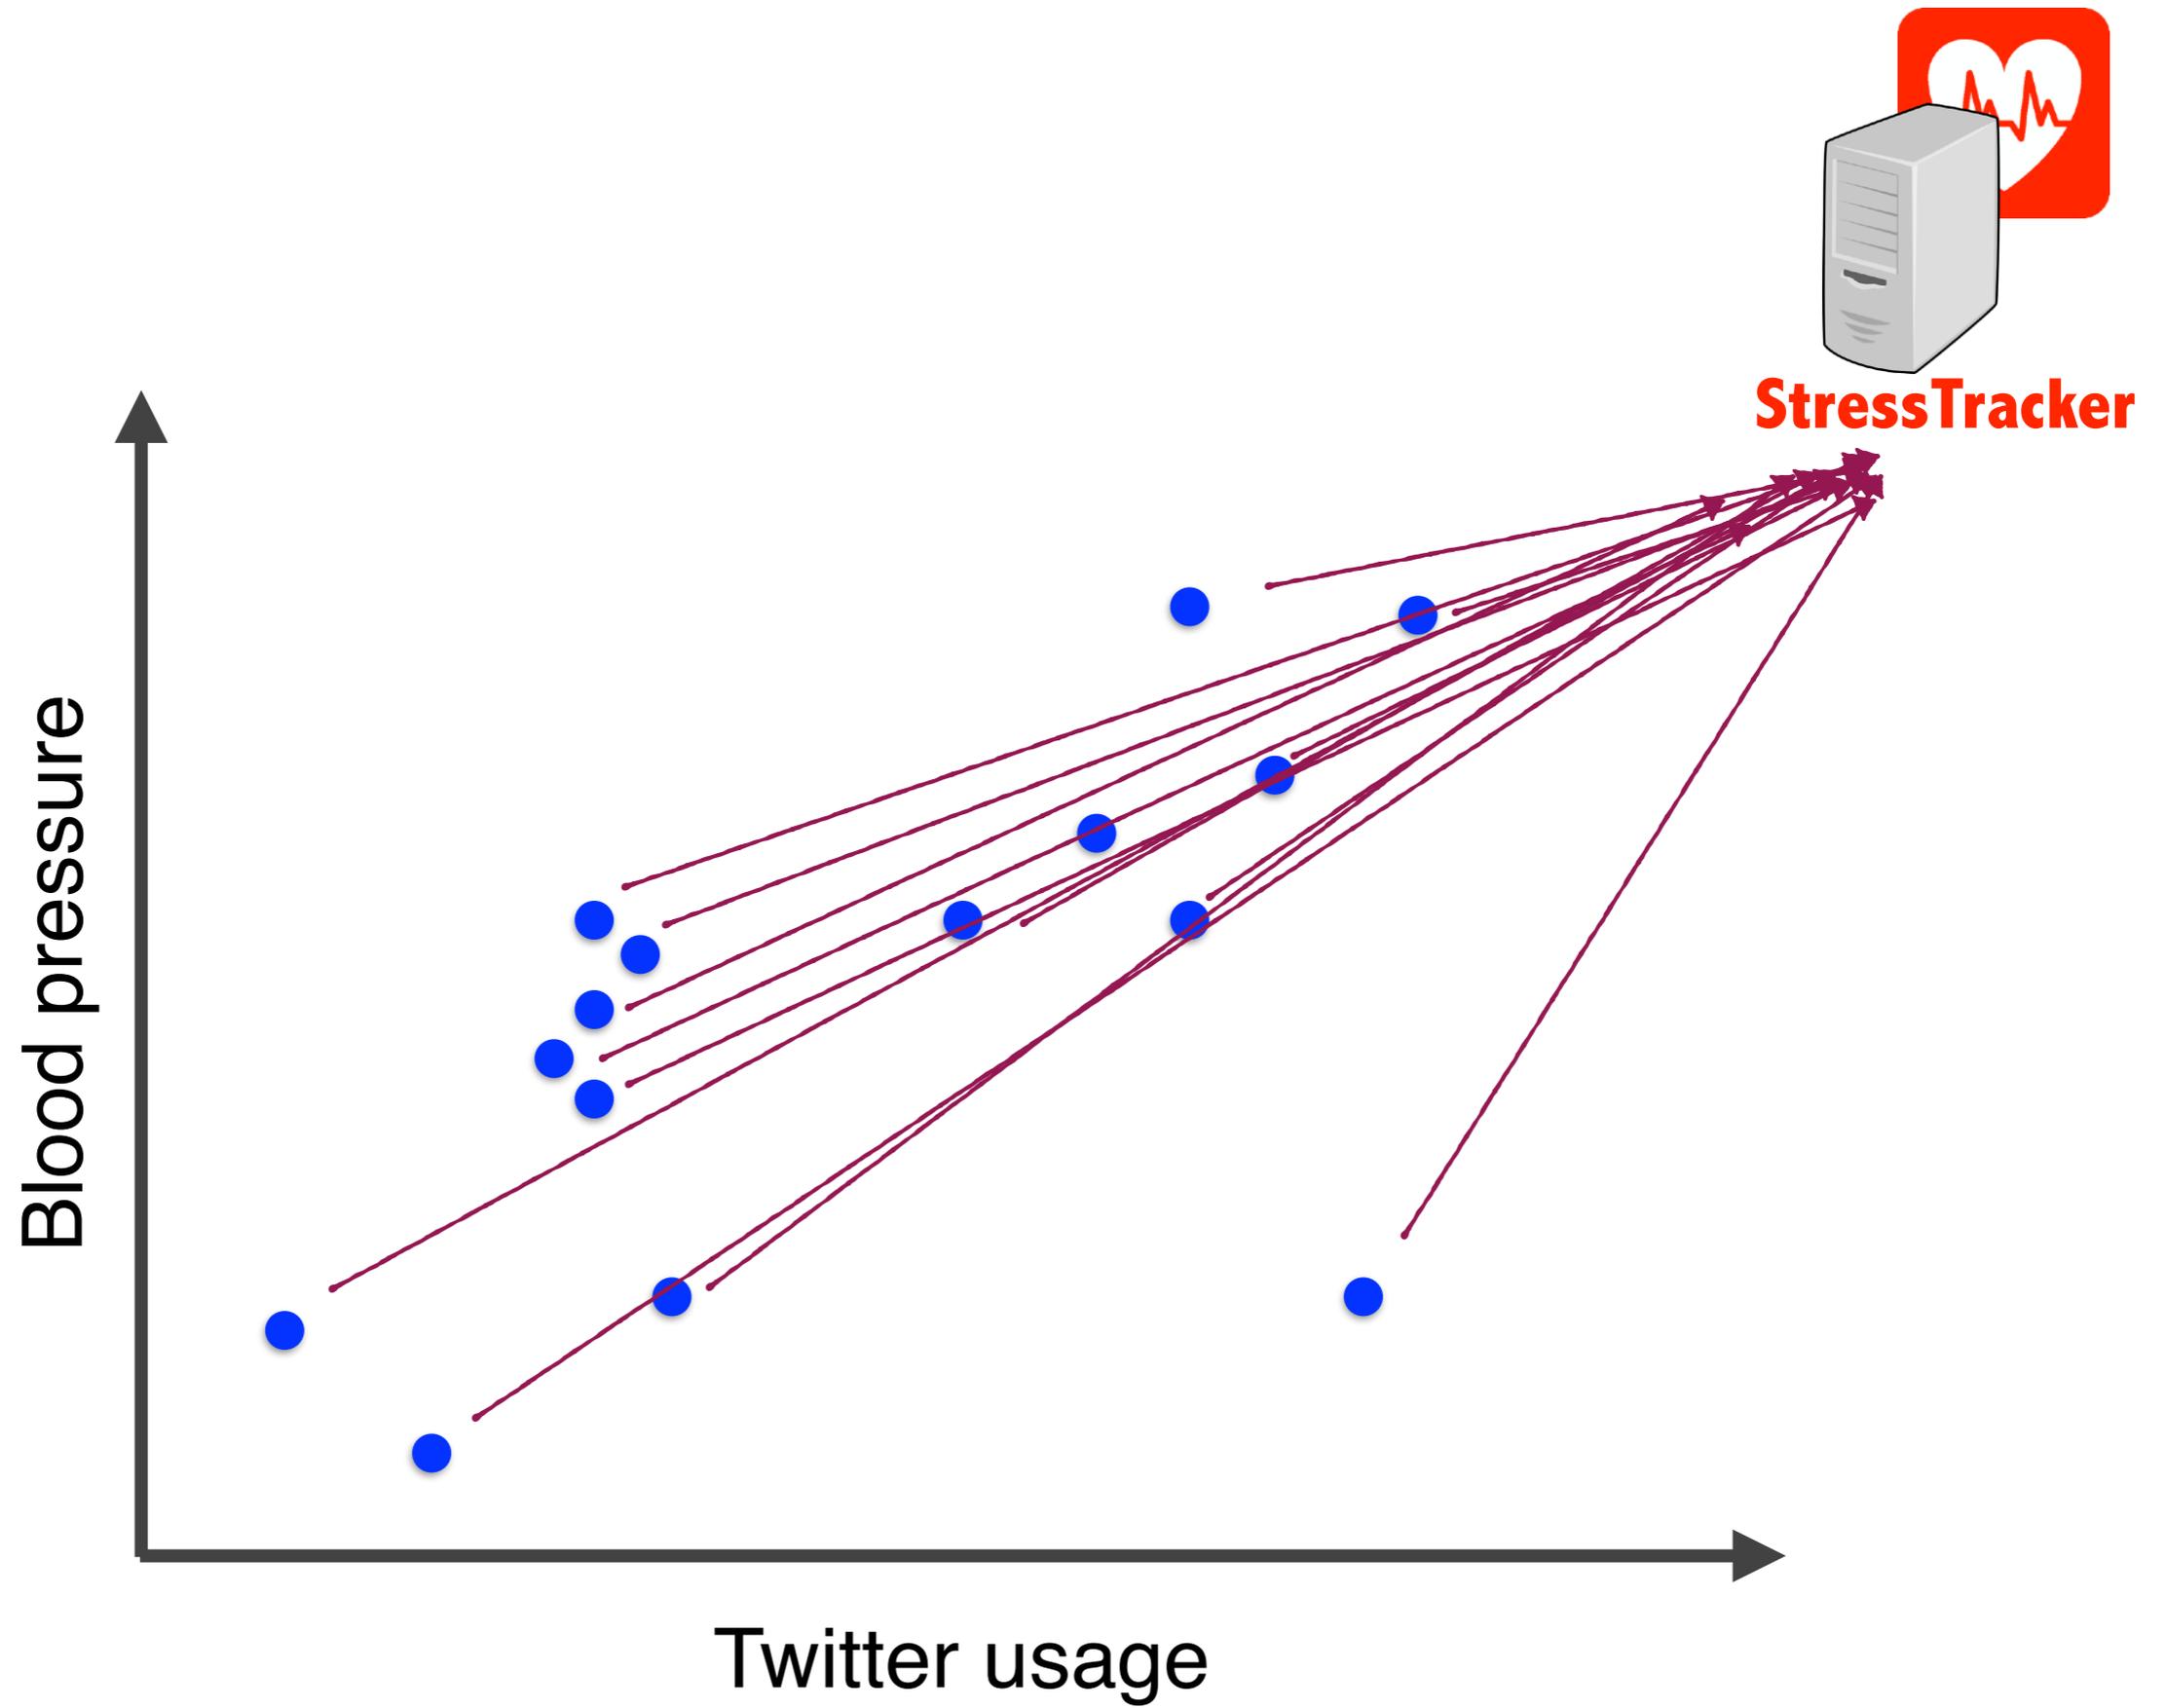
Today



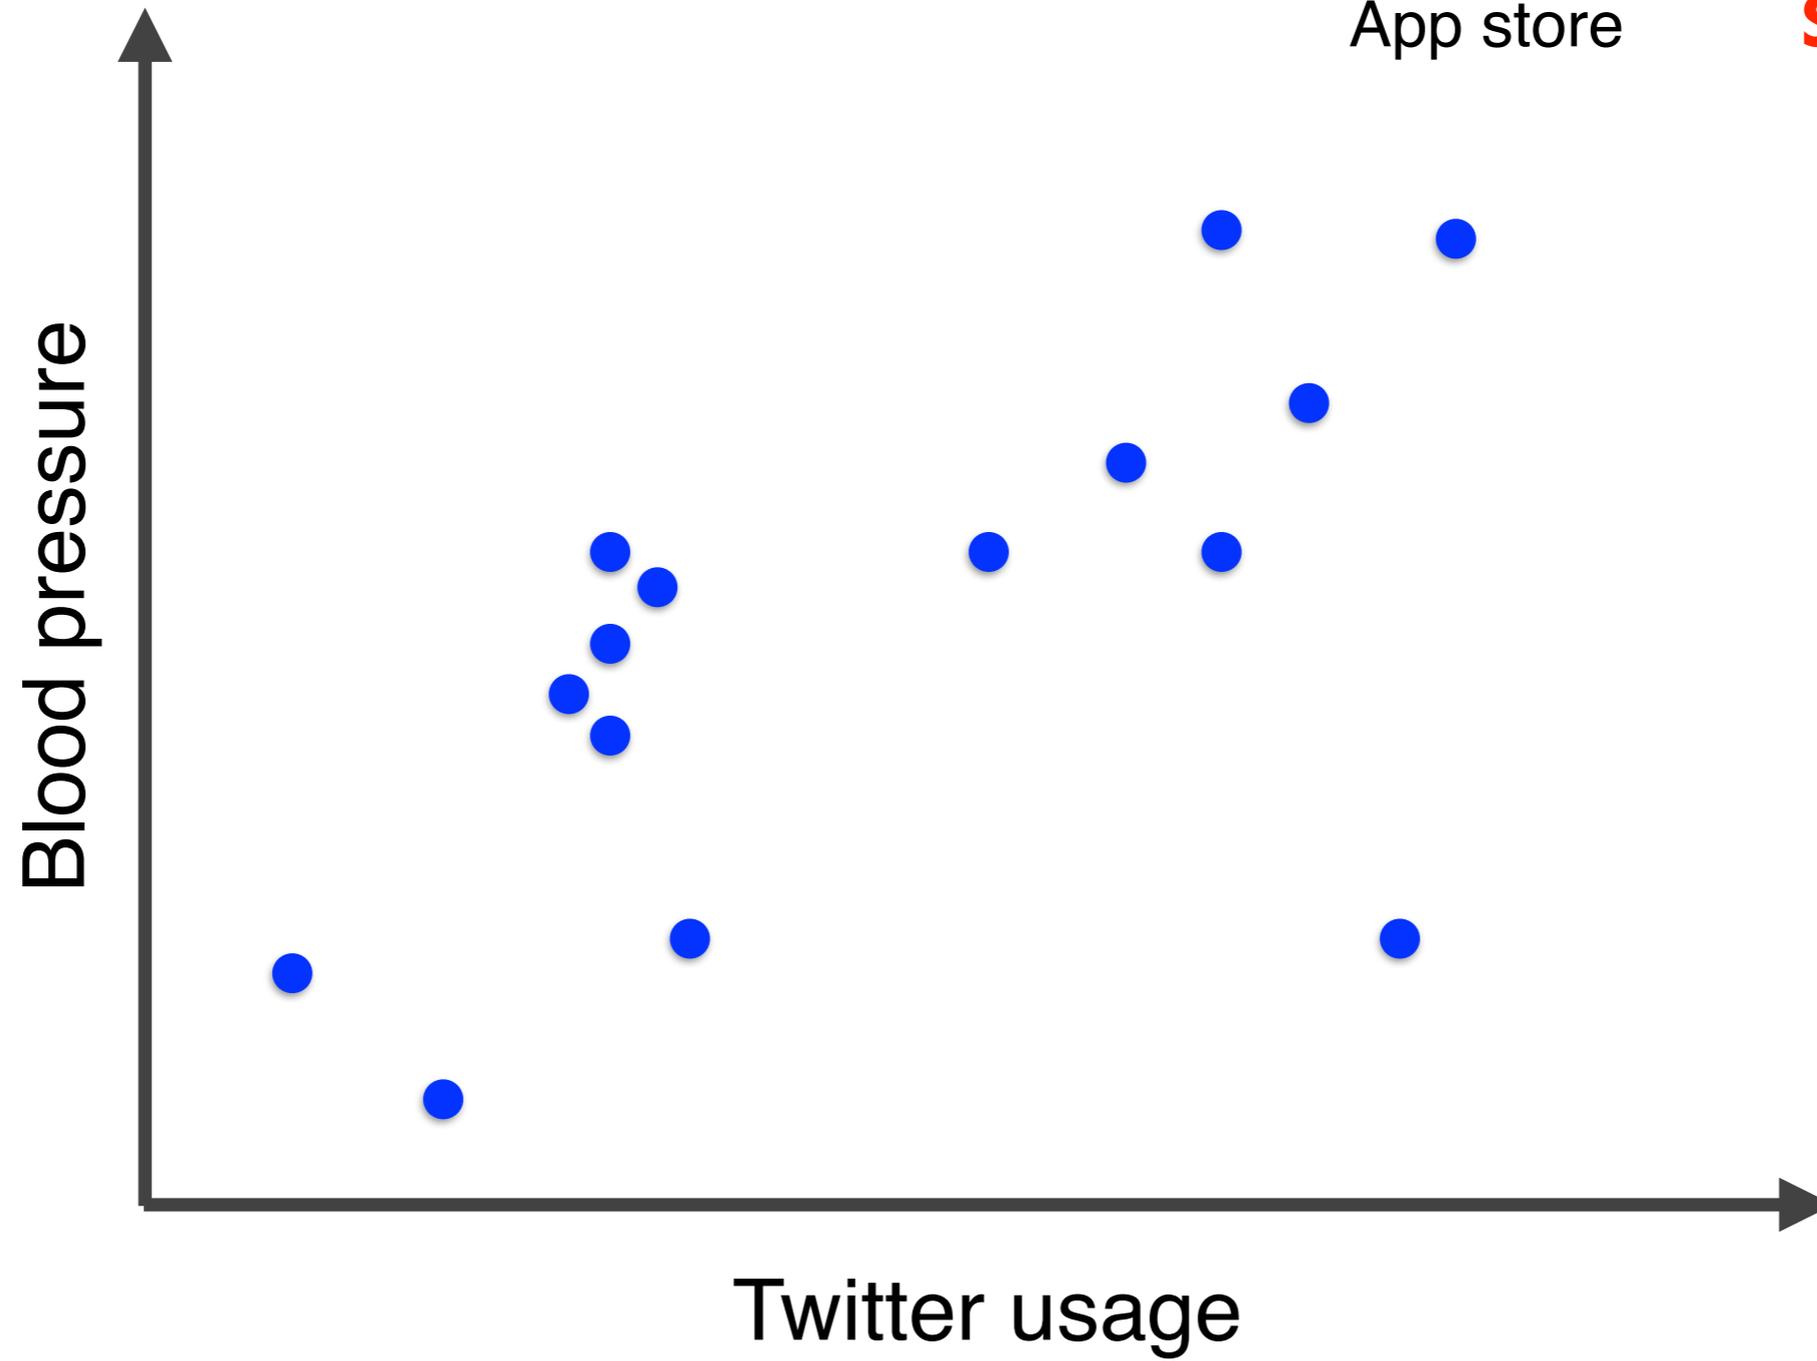
StressTracker



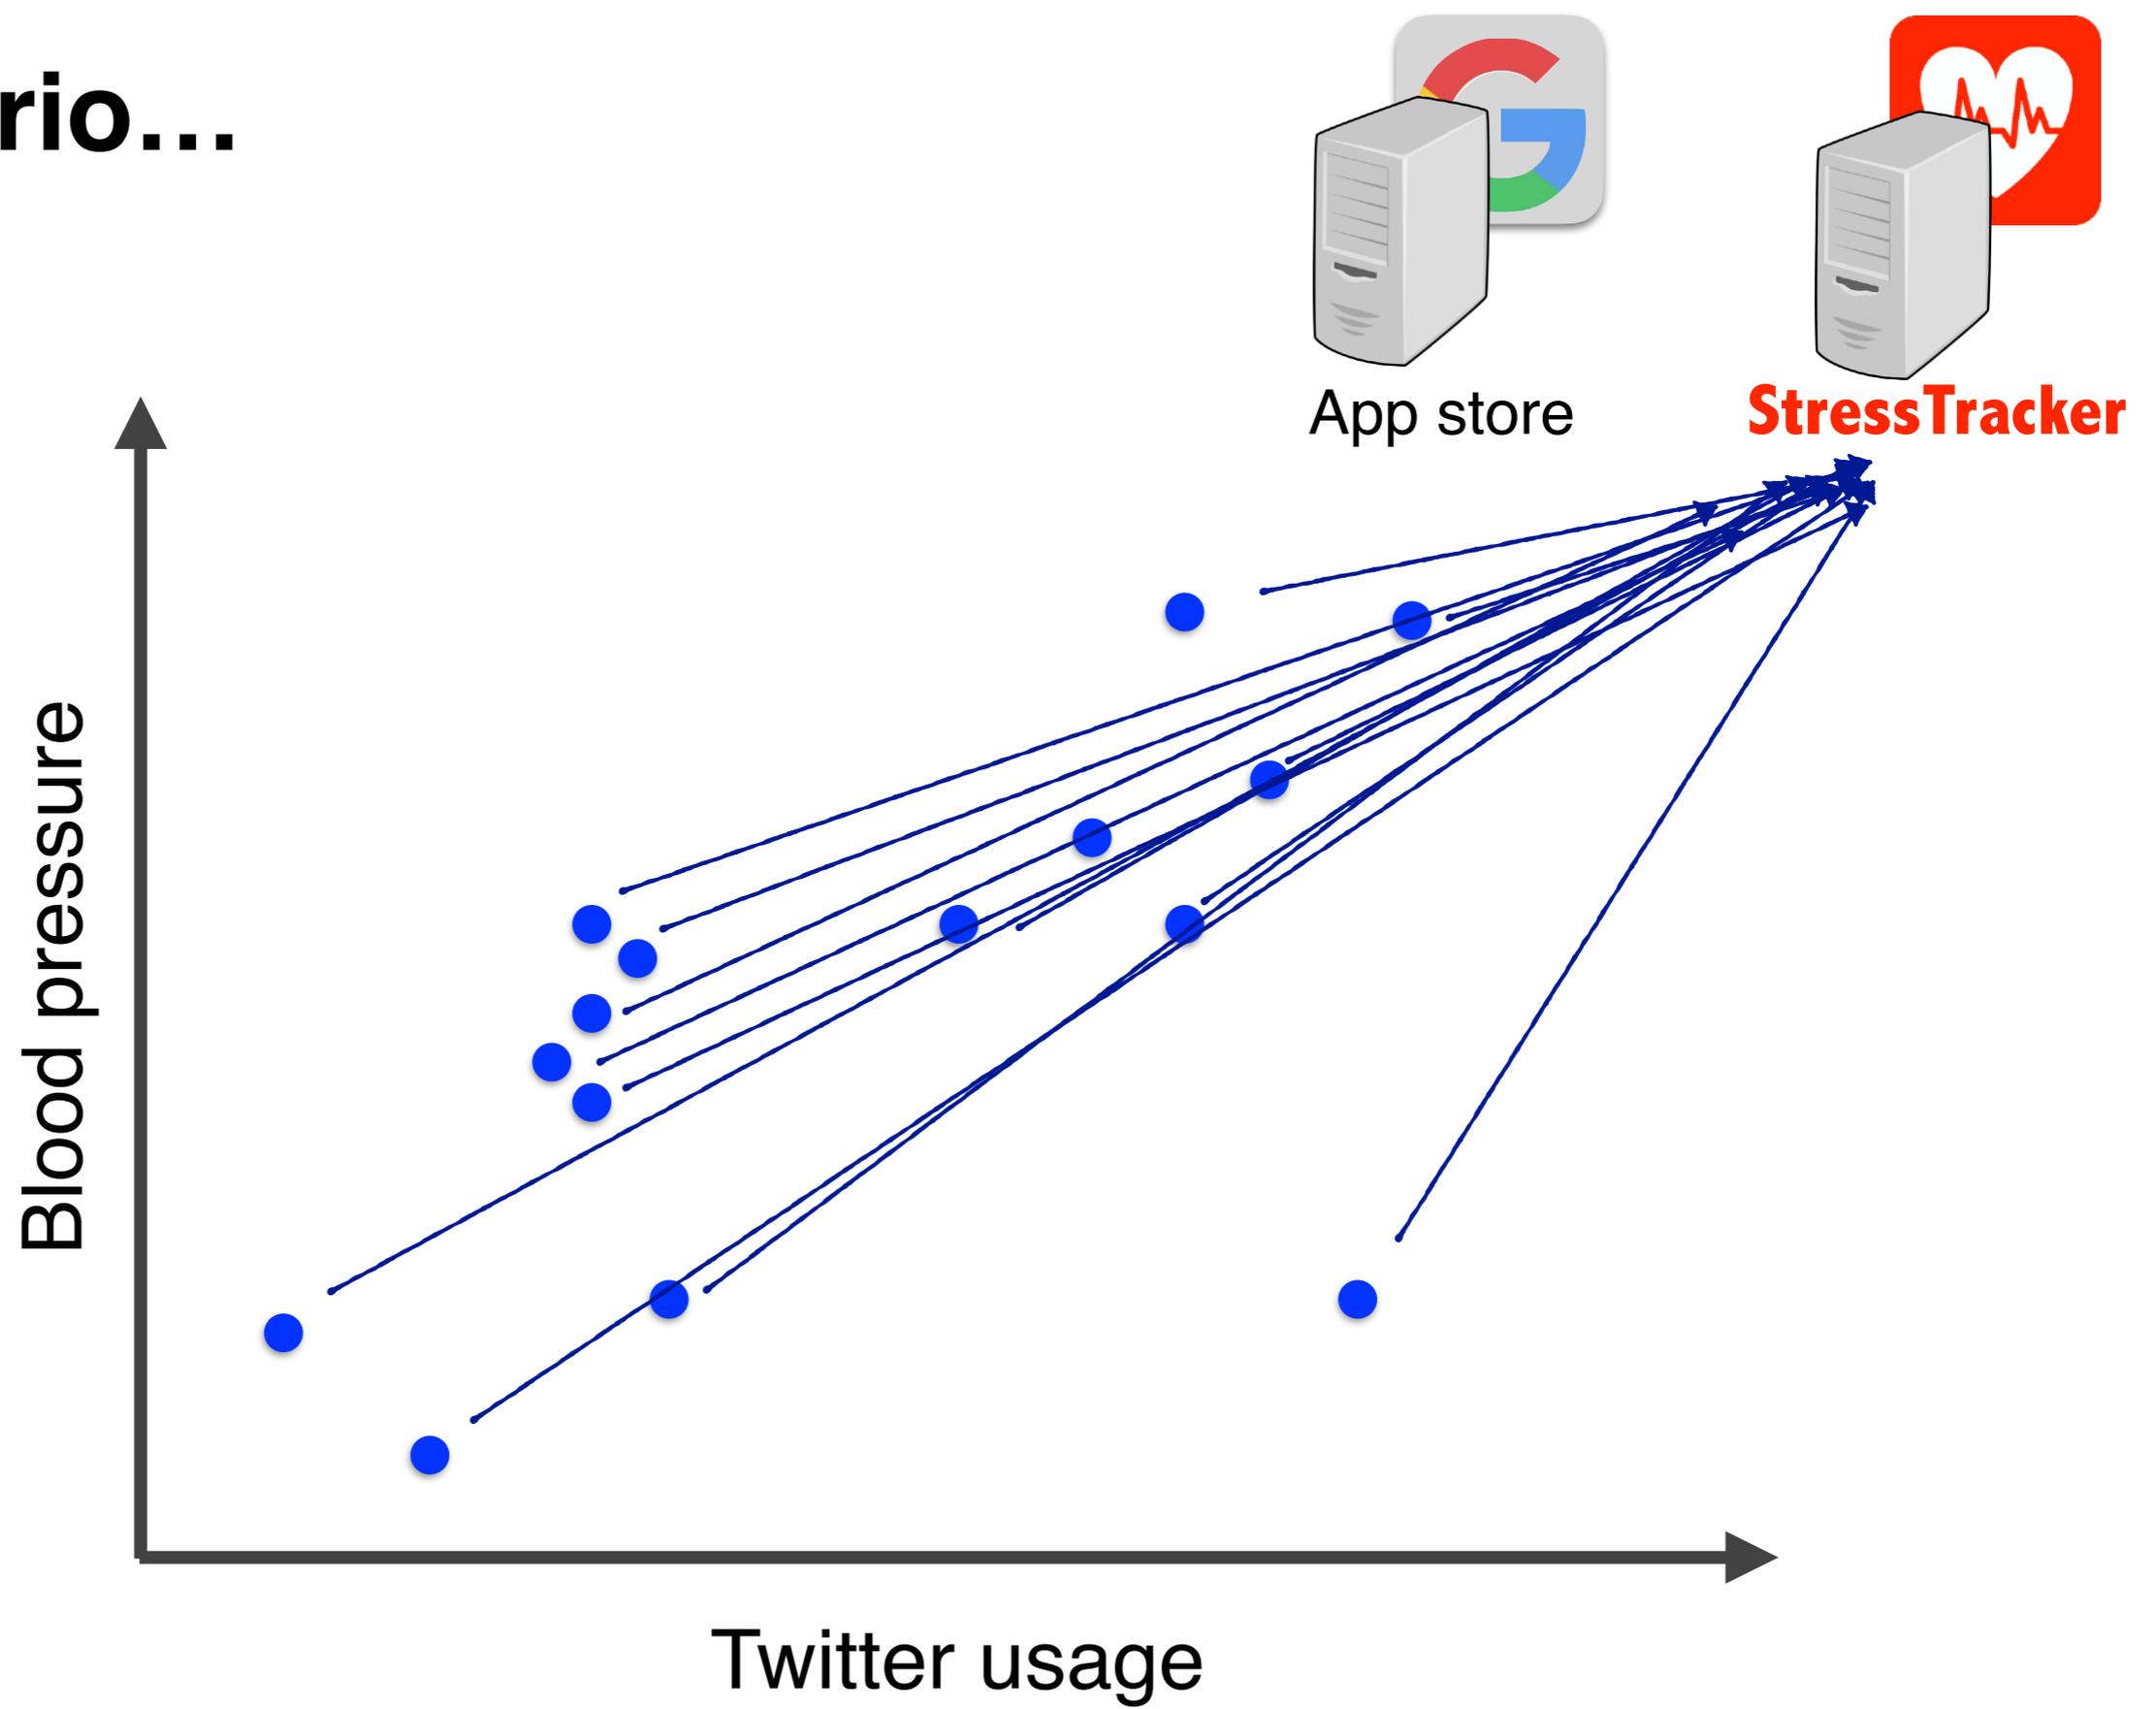
Today



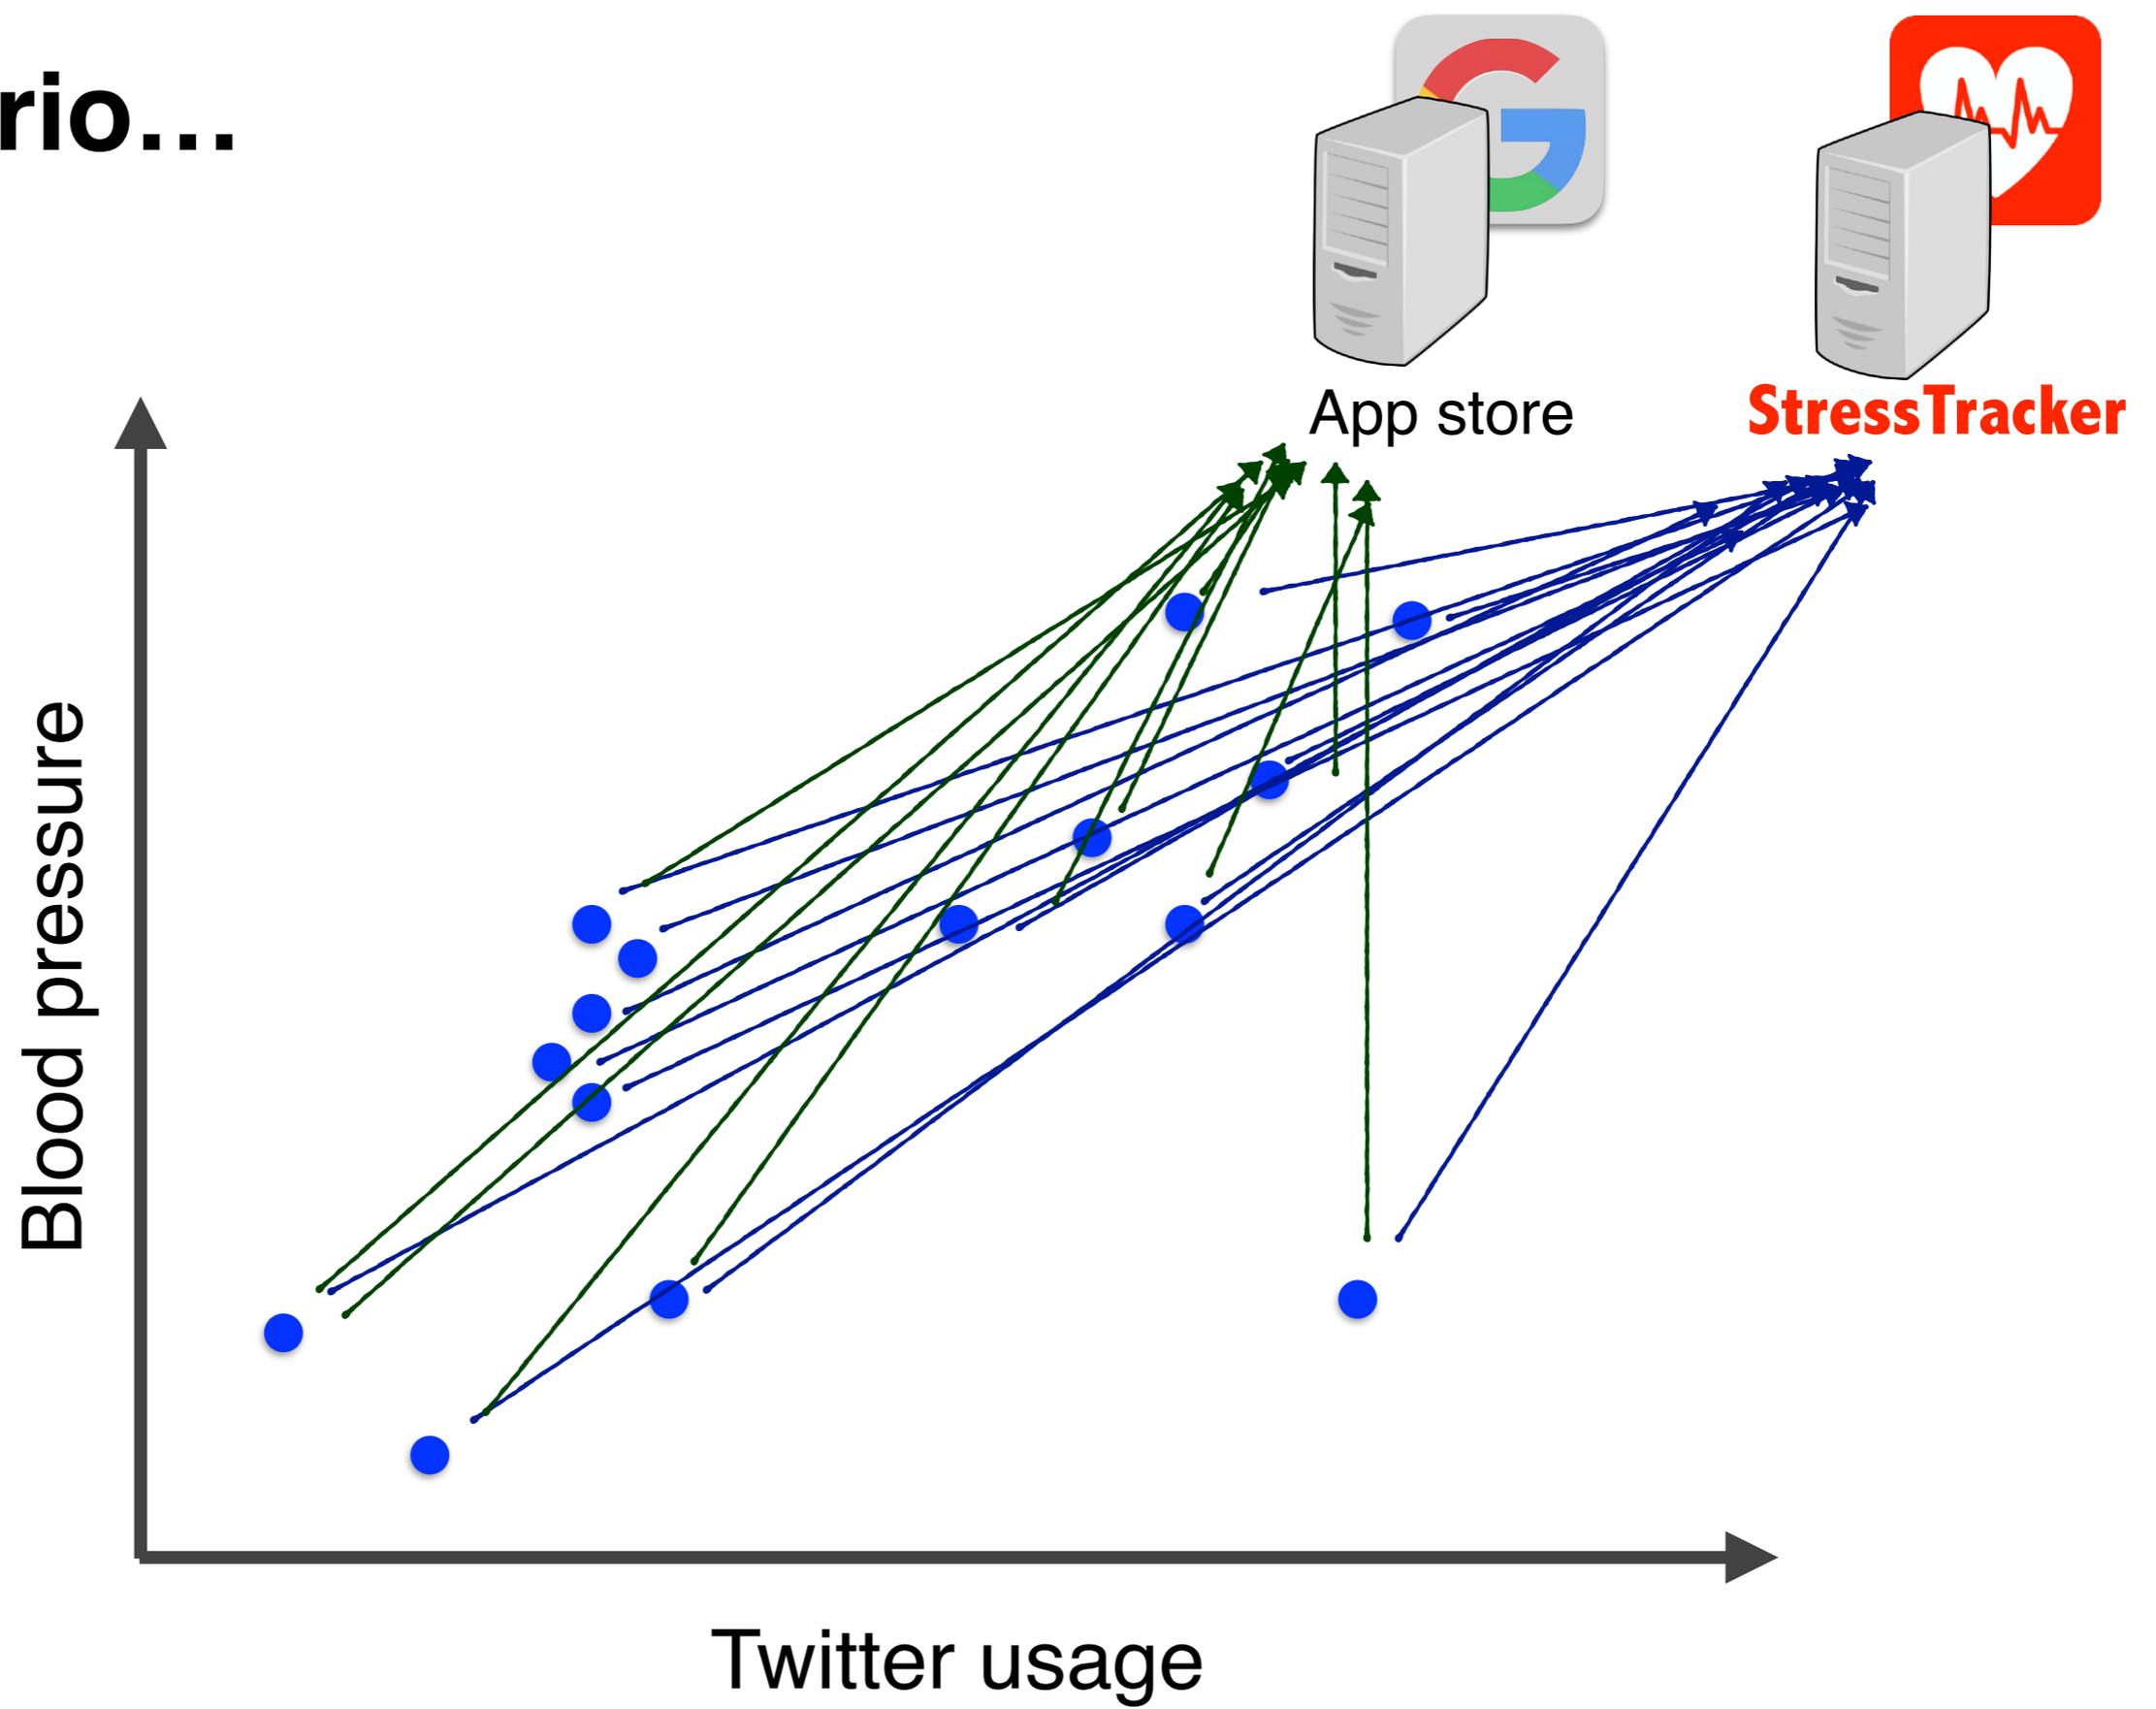
With Prio...



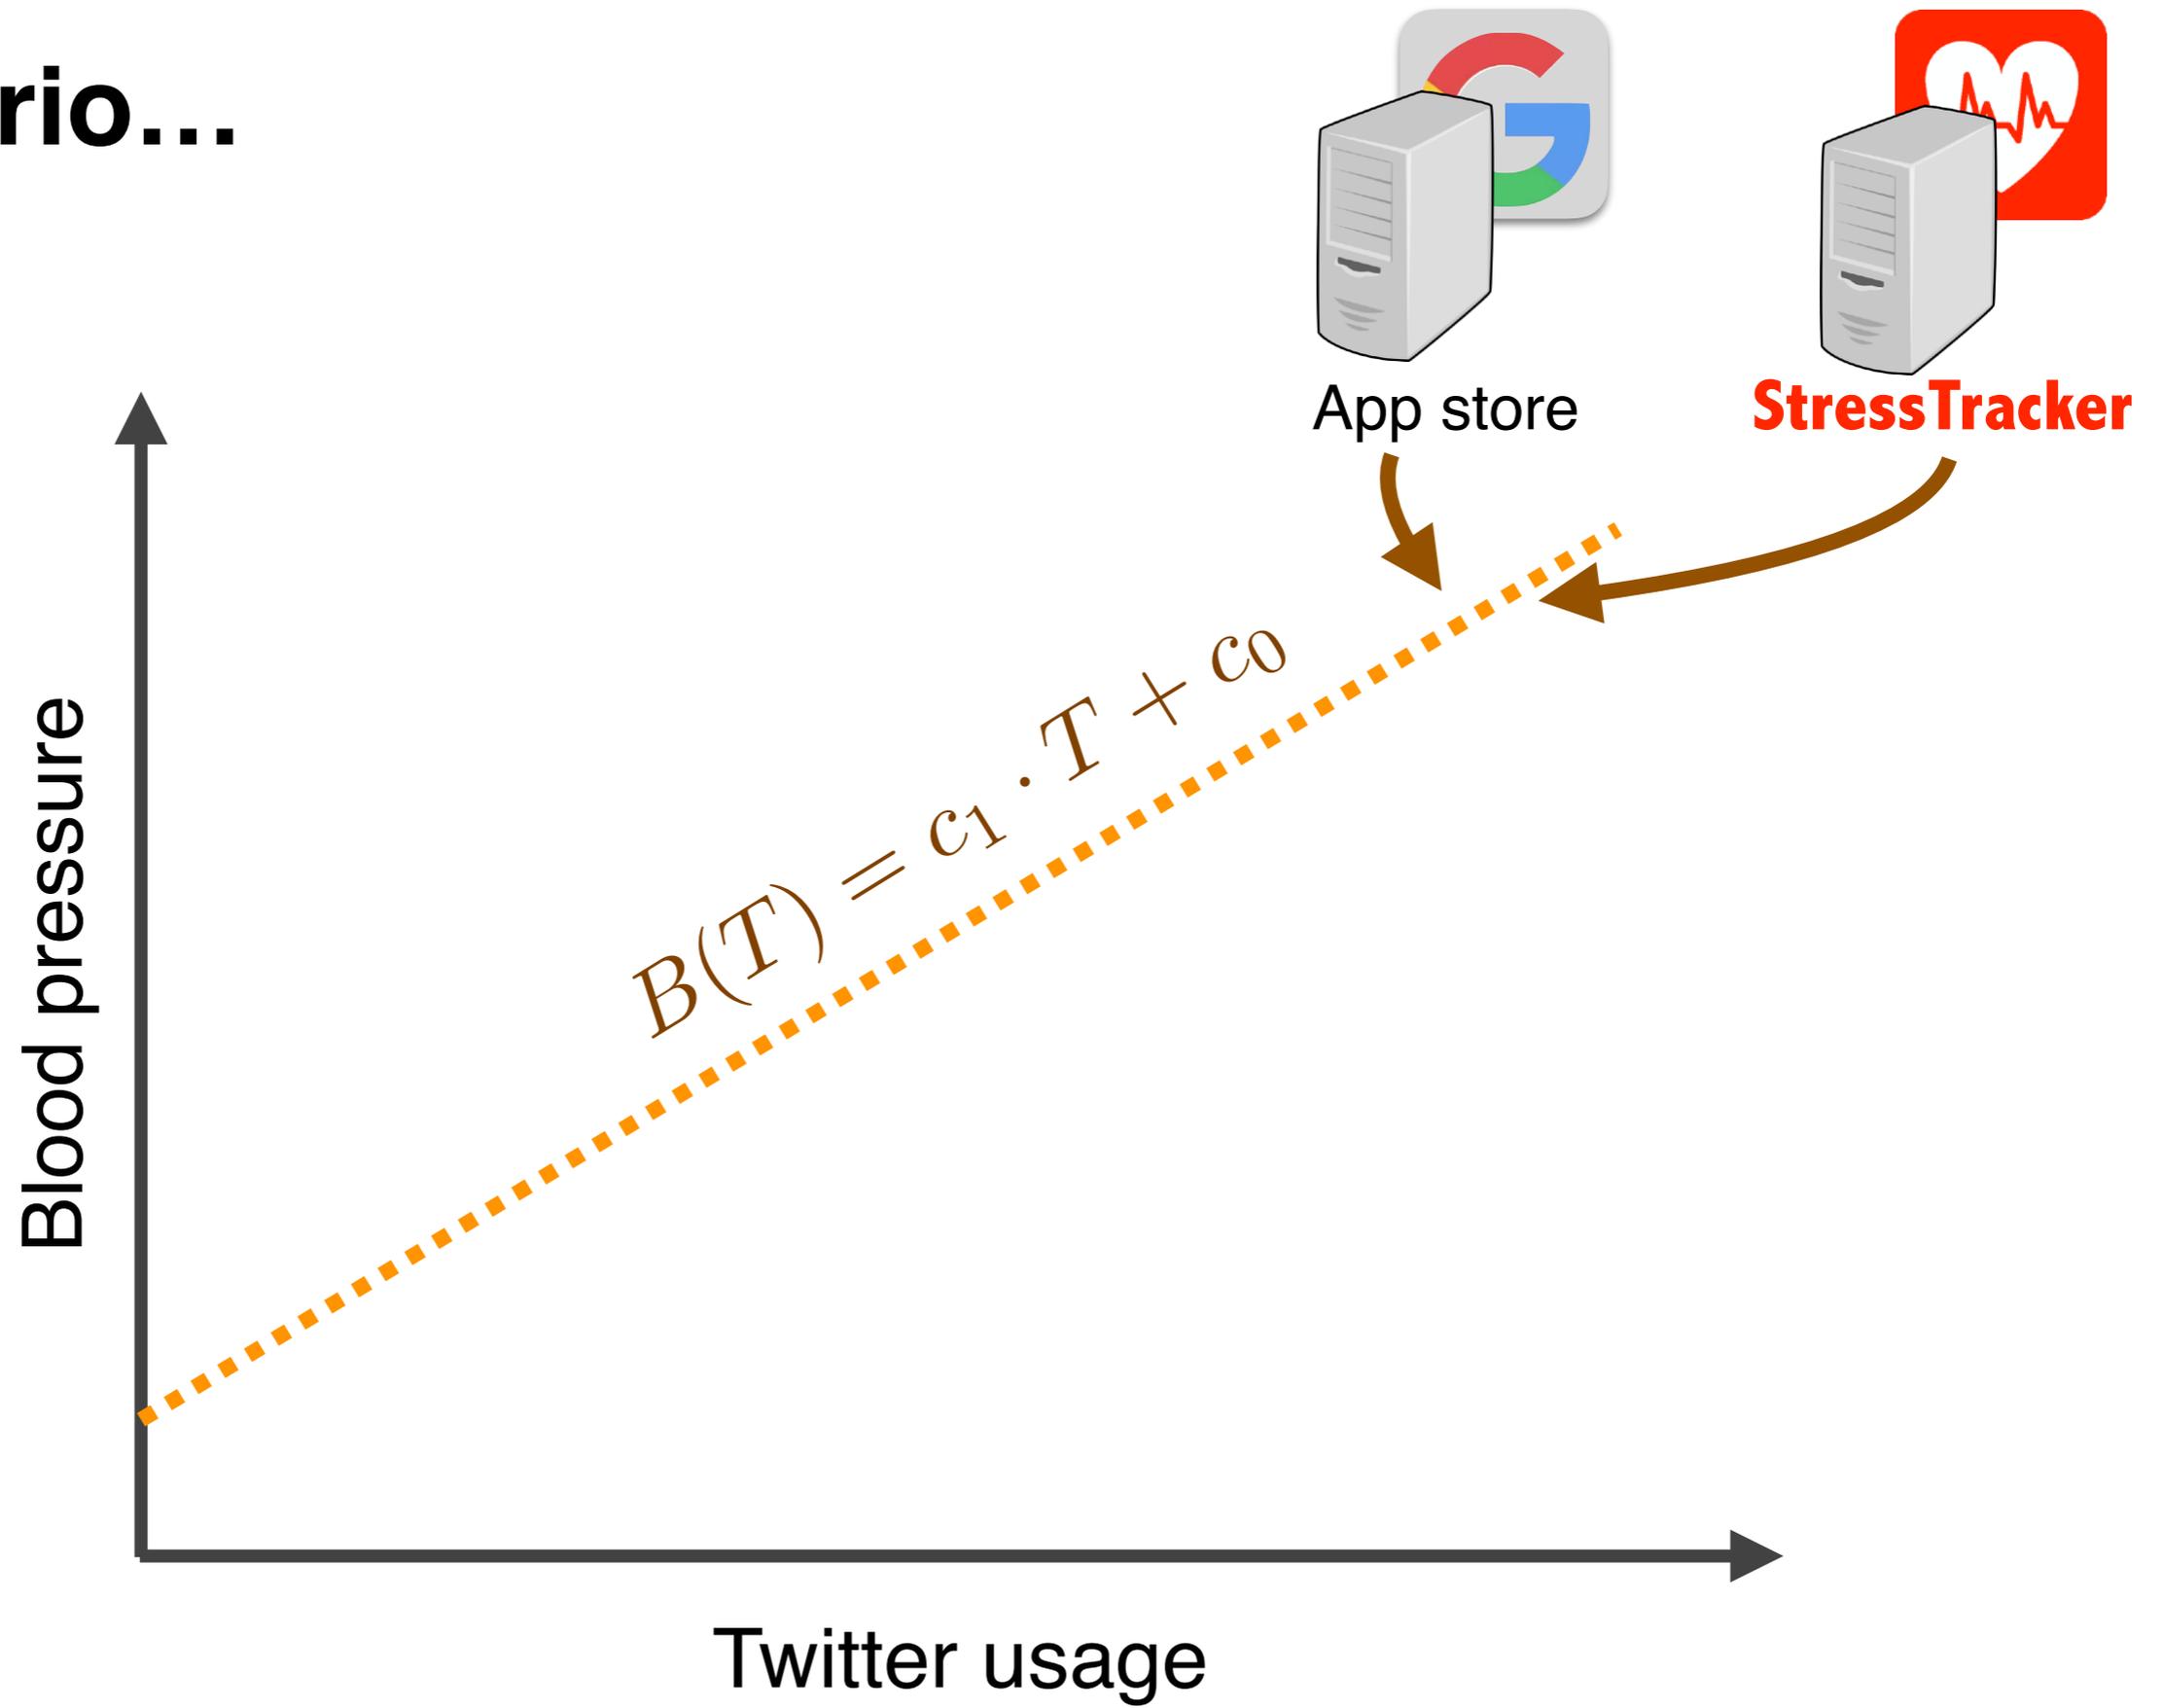
With Prio...



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!

Example Encoding: Average and Variance

[PrivStats11]

Example Encoding: Average and Variance

[PrivStats11]

- Each of N clients holds a value x_i
- Servers want the **AVG** and **VAR** of the x_i s.

Each client i encodes her value x as the pair

$$(x, y) = (x, x^2)$$

Simple to check that the encoding is valid:

$$\text{Valid}(x, y) = (x^2 - y) \quad [\text{outputs zero if valid}]$$

Example Encoding: Average and Variance

[PrivStats11]

- Each of N clients holds a value x_i
- Servers want the **AVG** and **VAR** of the x_i s.

Each client i encodes her value x as the pair

$$(x, y) = (x, x^2)$$

Simple to check that the encoding is valid:

$$\text{Valid}(x, y) = (x^2 - y) \quad [\text{outputs zero if valid}]$$

Use Prio to compute the sum of encodings $\sum_i (x_i, y_i)$

Example Encoding: Average and Variance

[PrivStats11]

- Each of N clients holds a value x_i
- Servers want the **AVG** and **VAR** of the x_i s.

Each client i encodes her value x as the pair

$$(x, y) = (x, x^2)$$

Simple to check that the encoding is valid:

$$\text{Valid}(x, y) = (x^2 - y) \quad [\text{outputs zero if valid}]$$

Use Prio to compute the sum of encodings $\sum_i (x_i, y_i)$

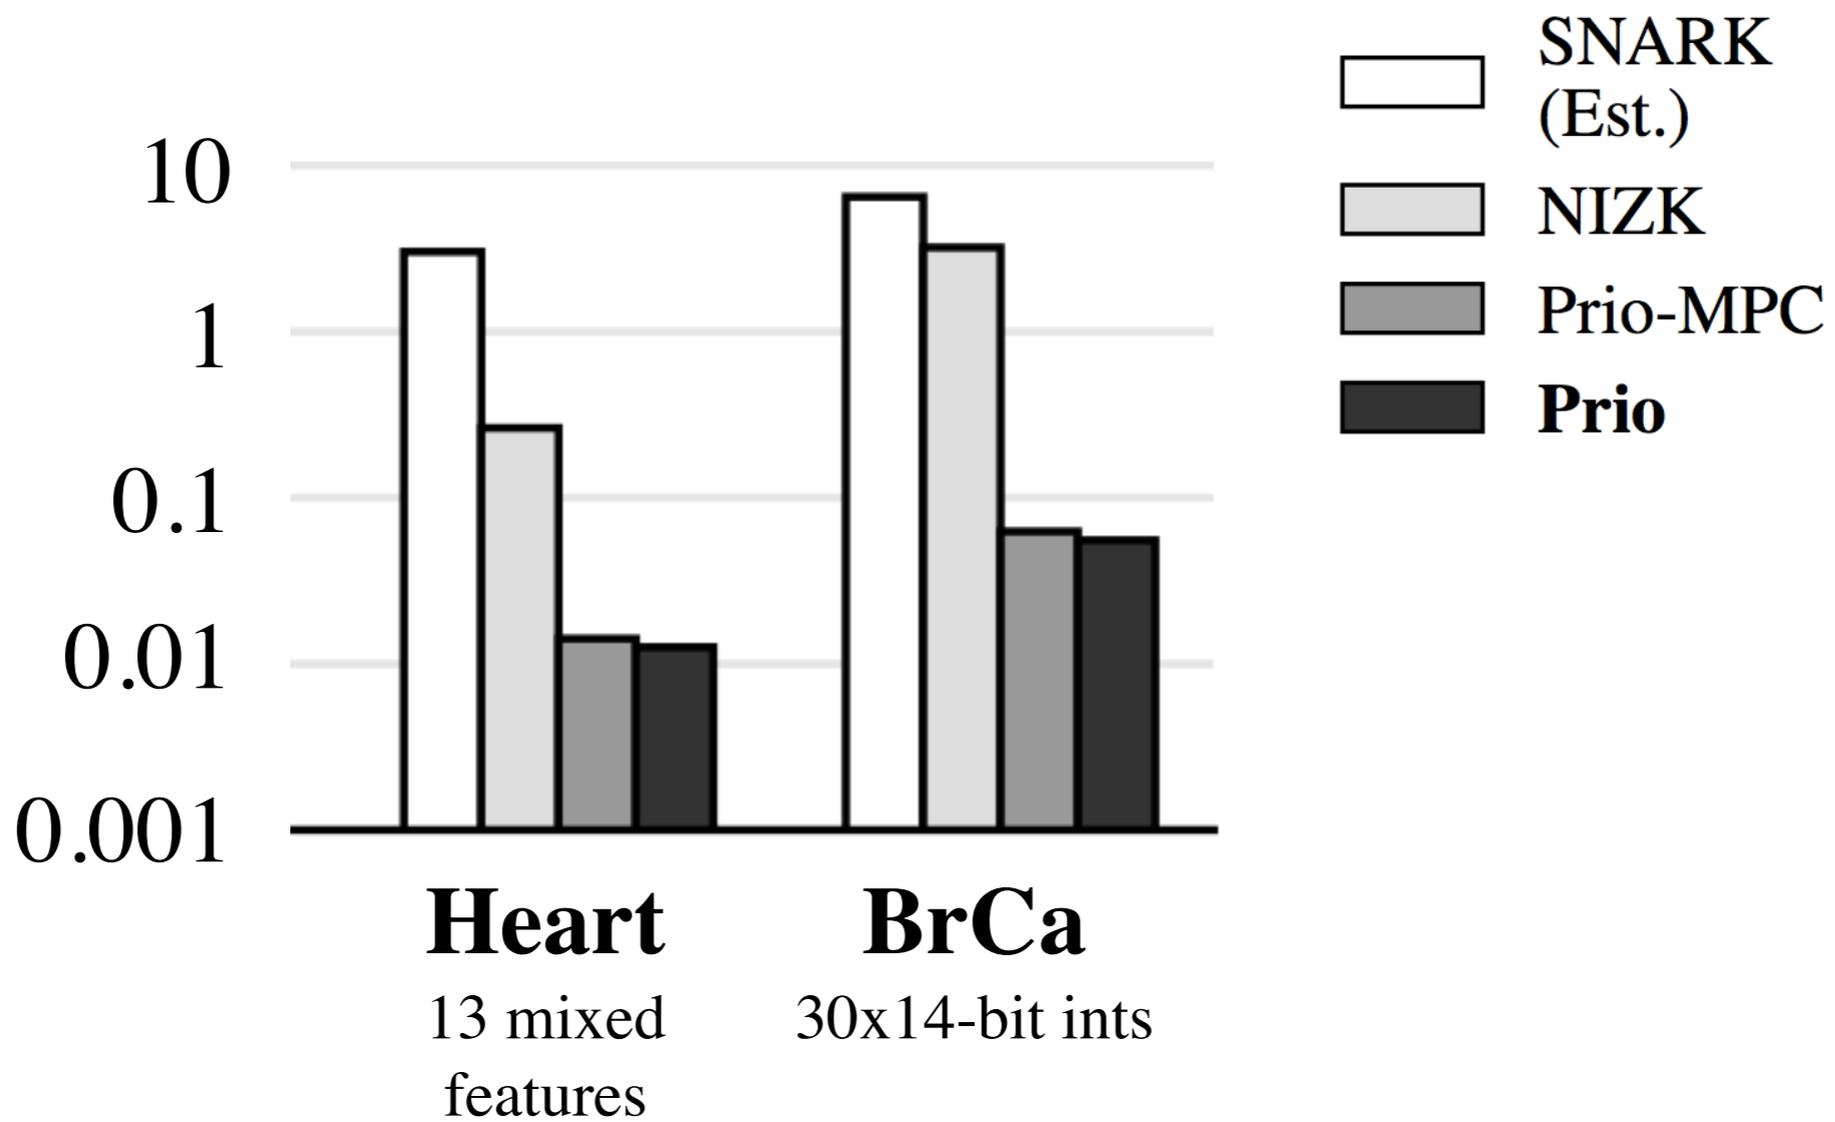
Then recover the statistics:

$$\text{AVG}(X) = (\sum_i x_i) / N$$

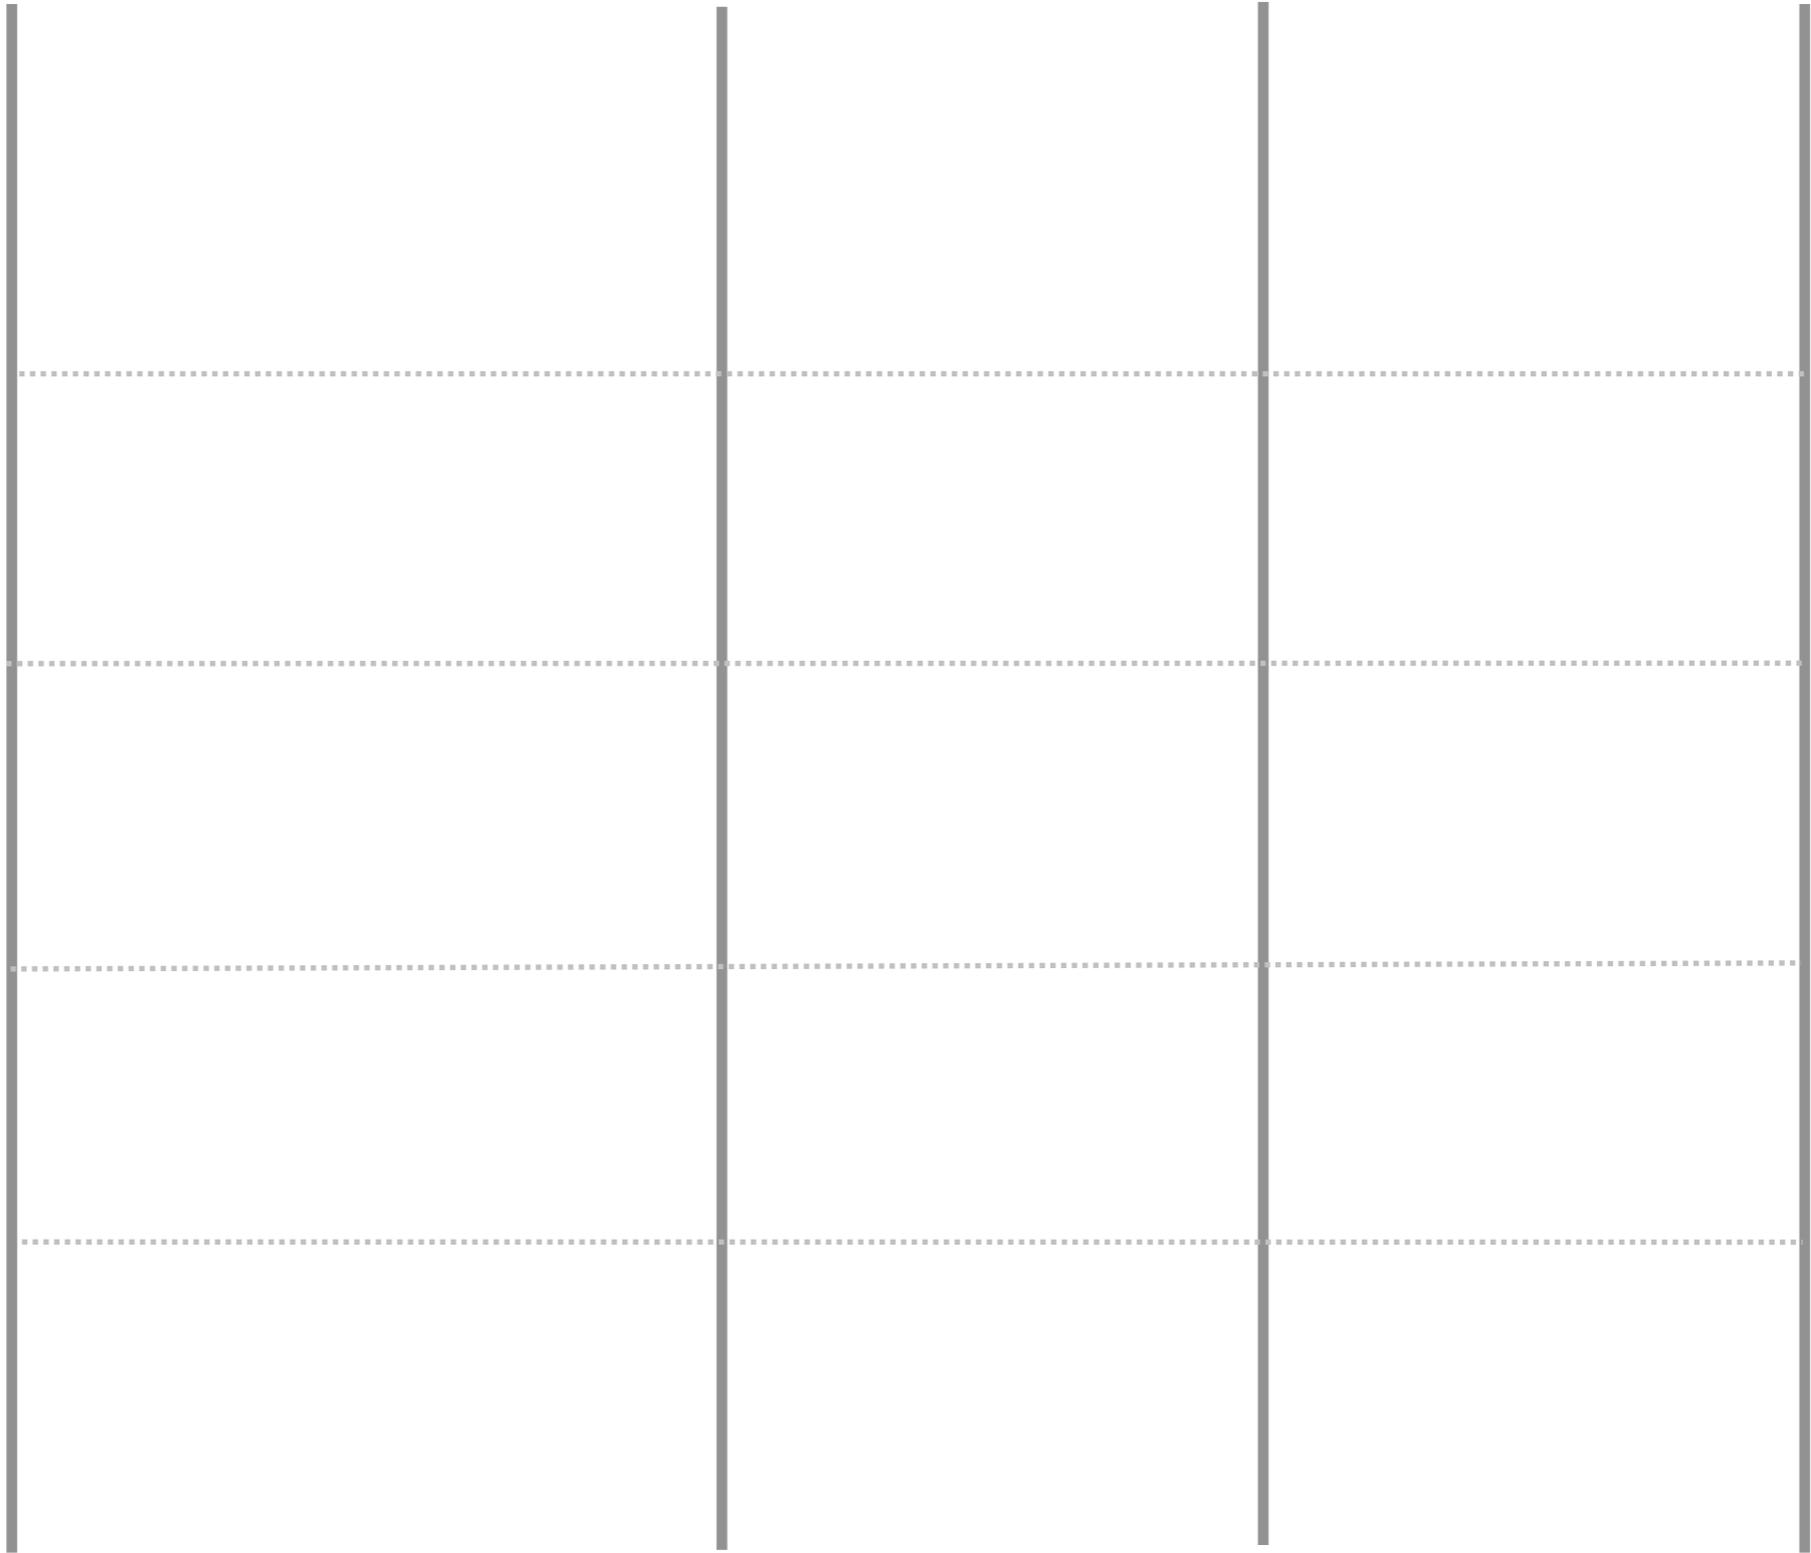
$$\text{AVG}(X^2) = (\sum_i y_i) / N = (\sum_i x_i^2) / N$$

$$\text{VAR}(X) = \text{AVG}(X^2) - \text{AVG}(X)^2$$

Client time (s)
Lower is better.



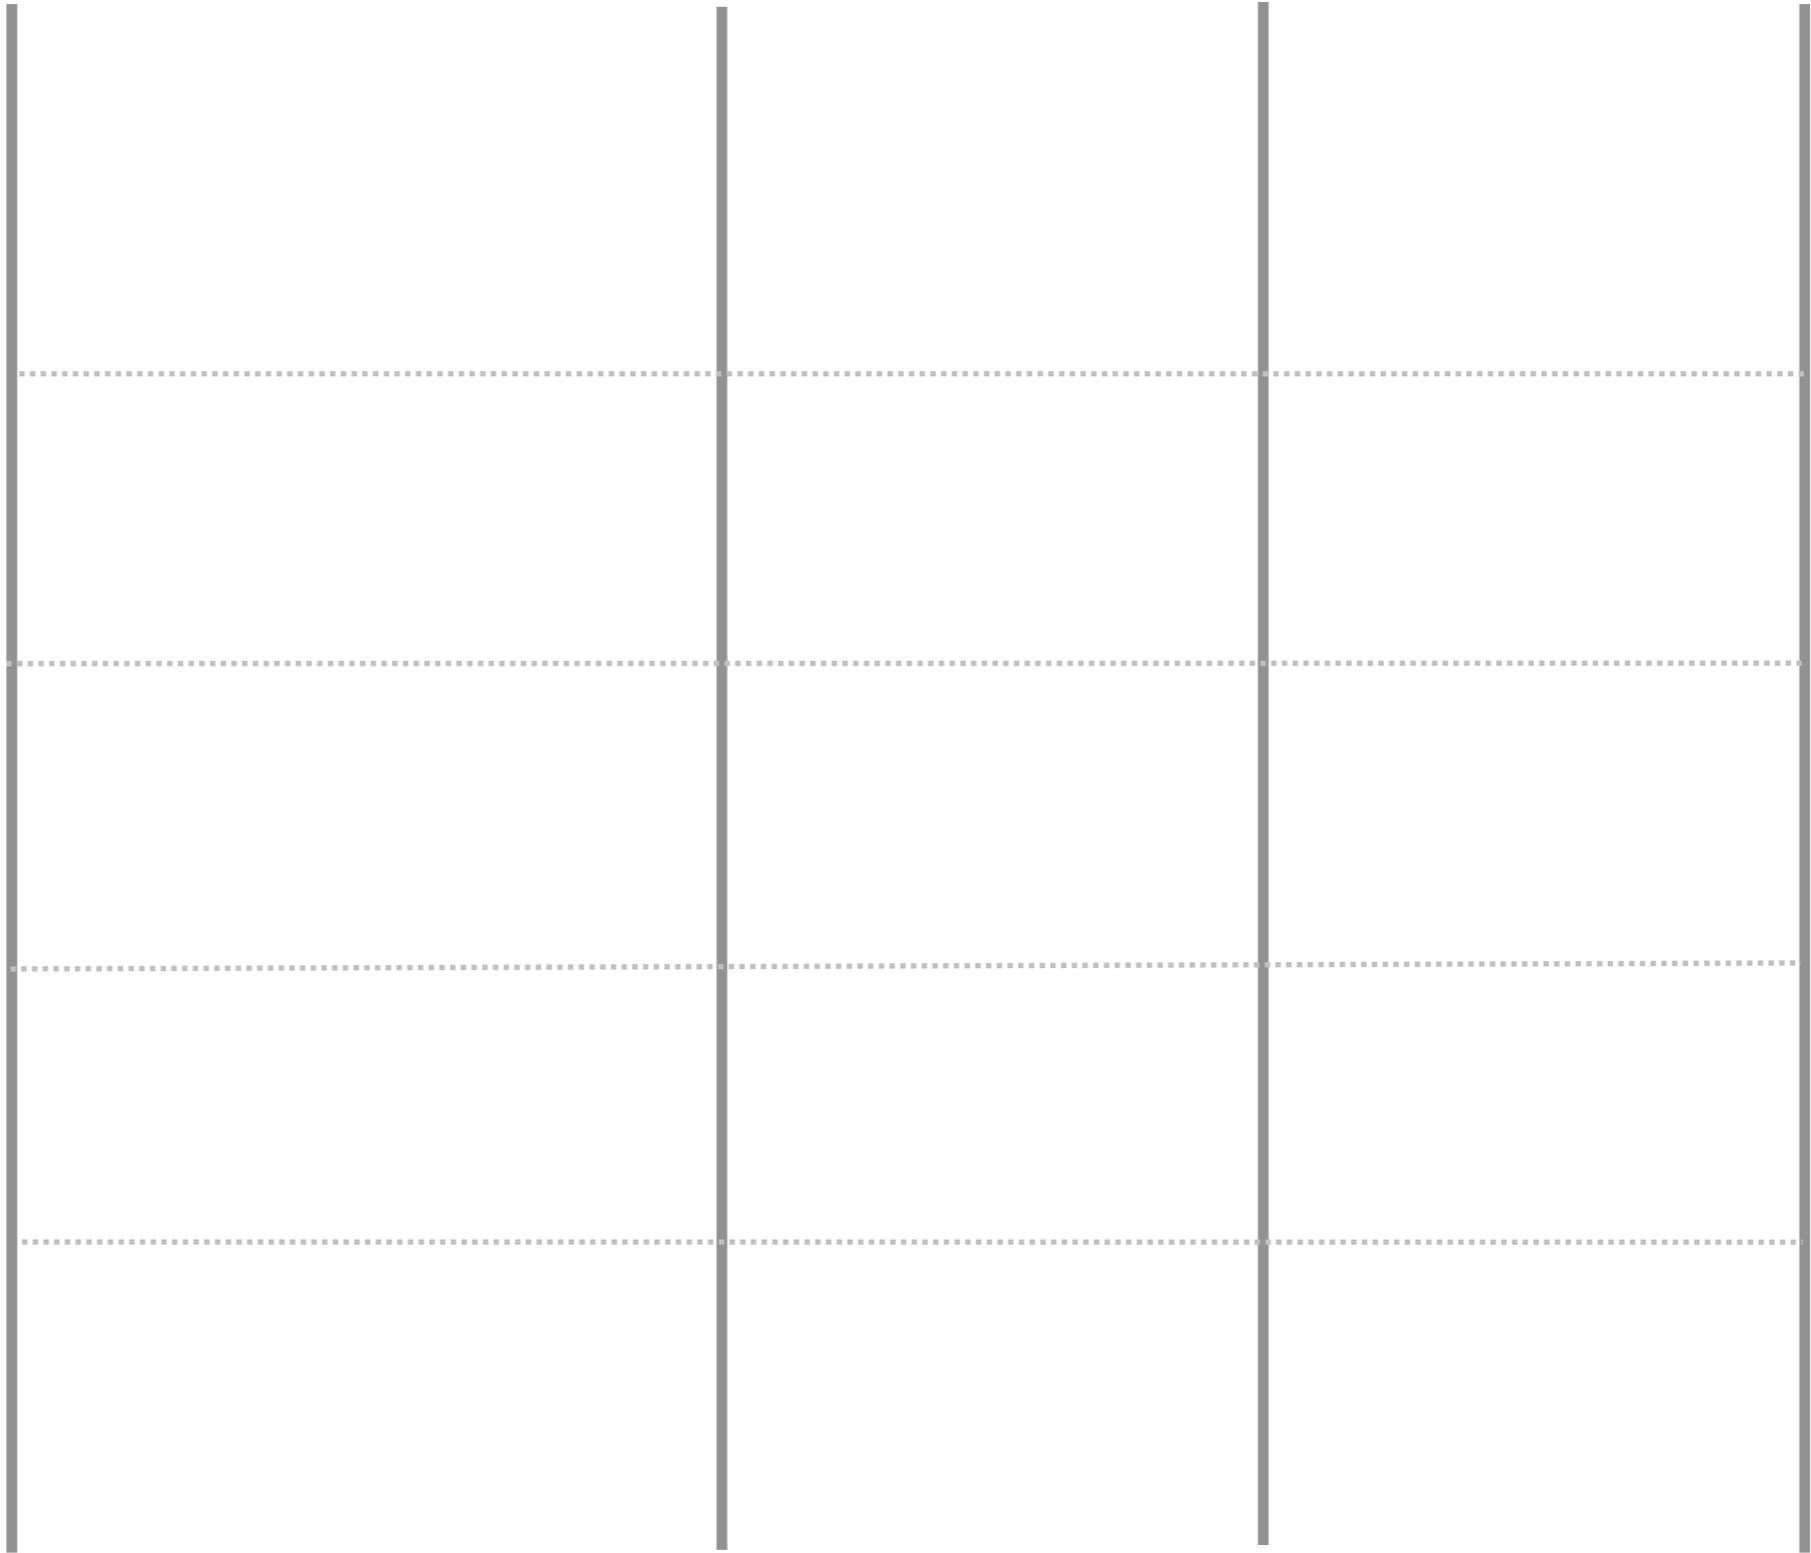
Using 128-bit integers



Using 128-bit integers



Submit data

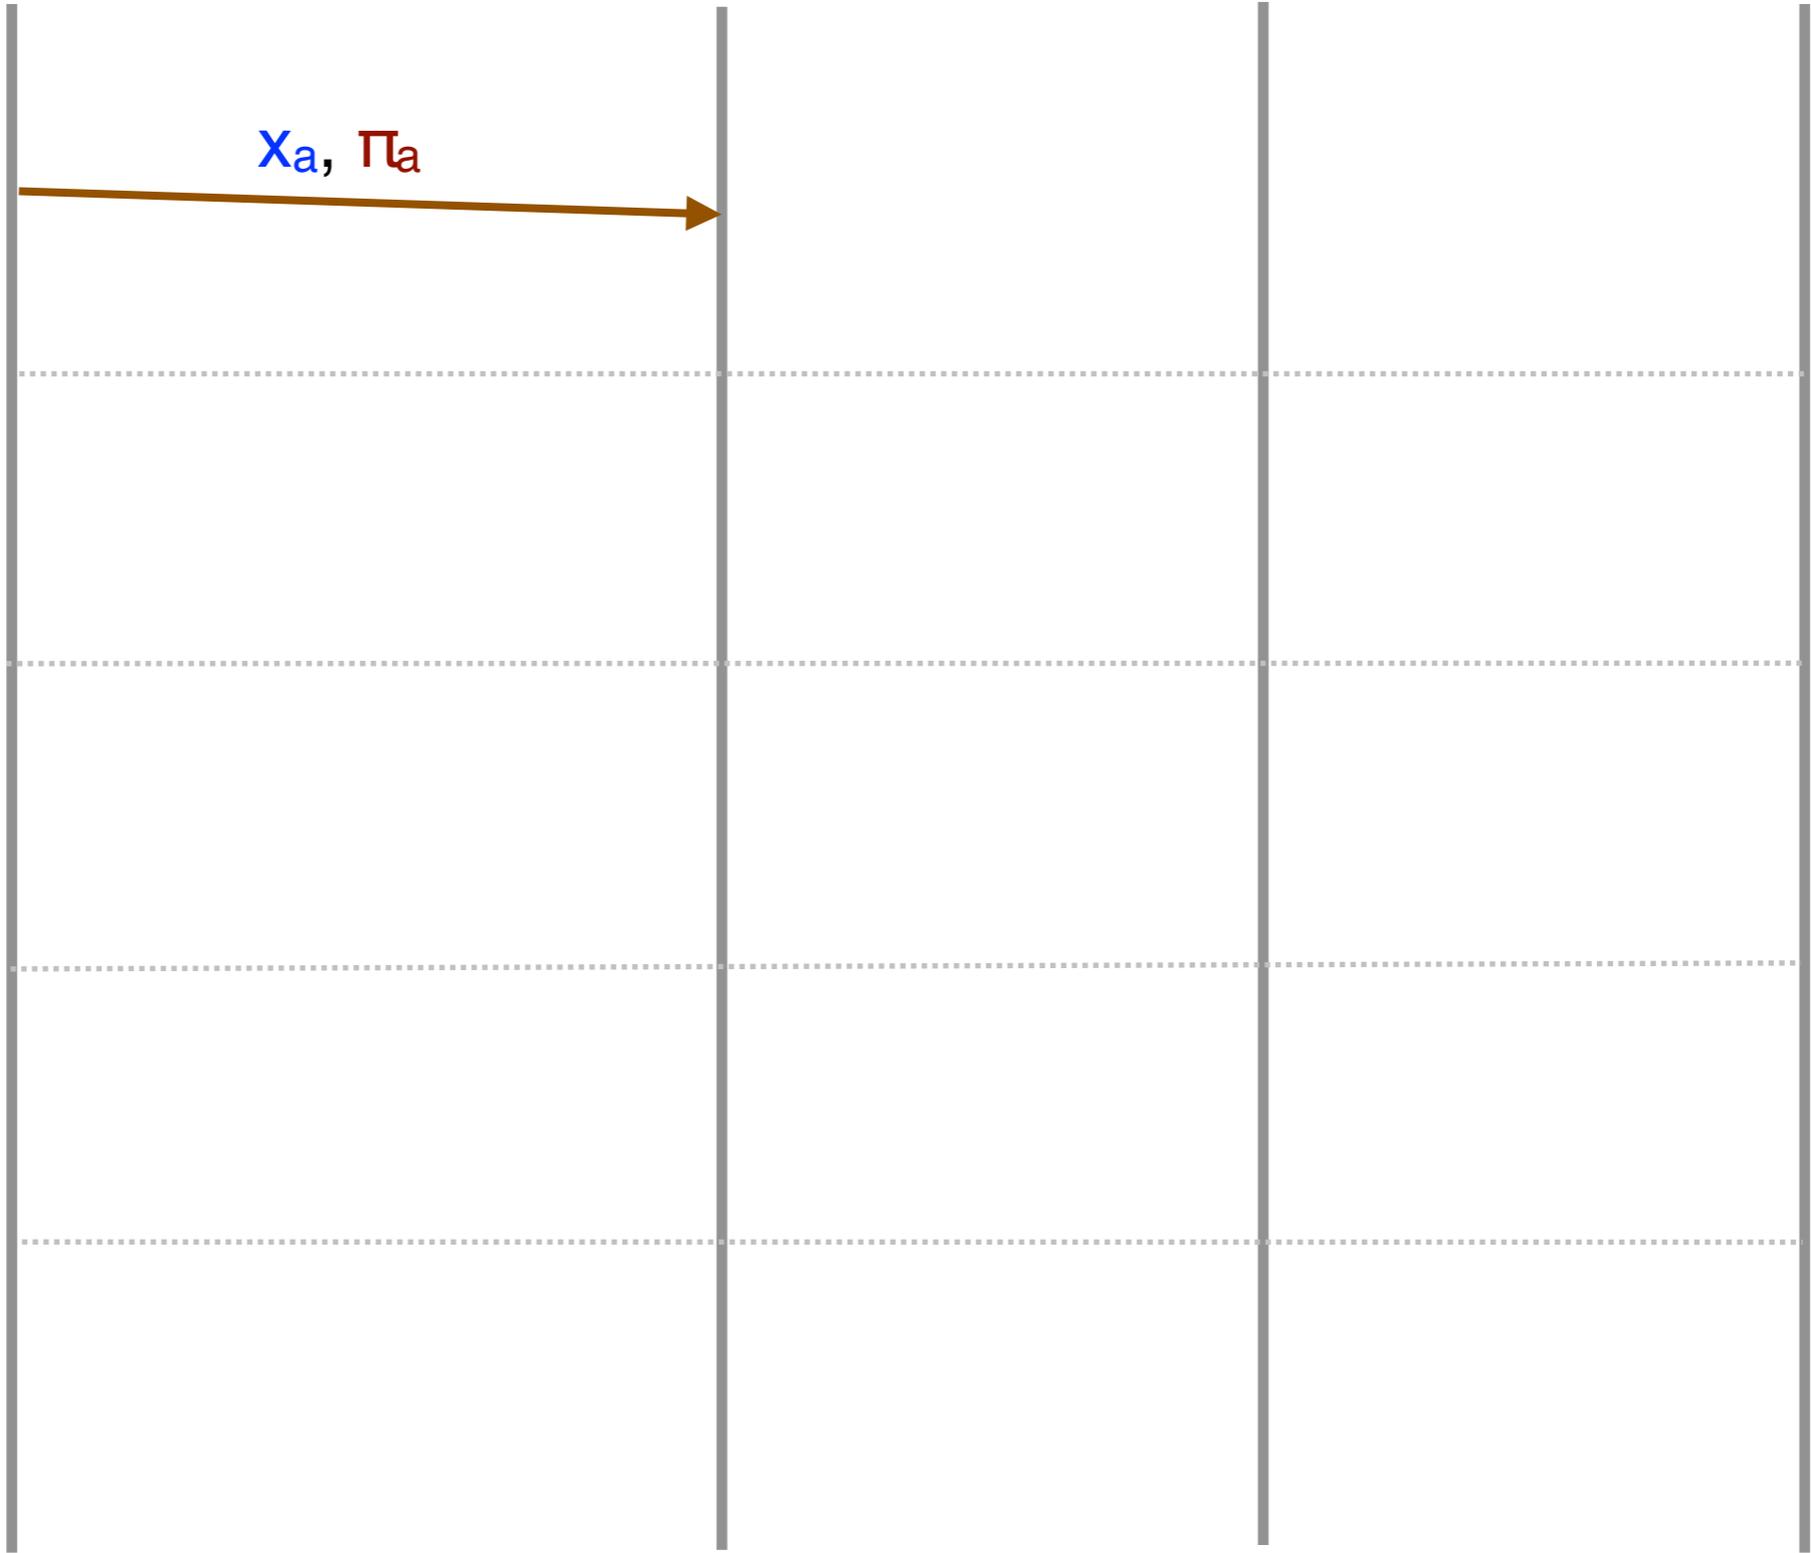


Using 128-bit integers

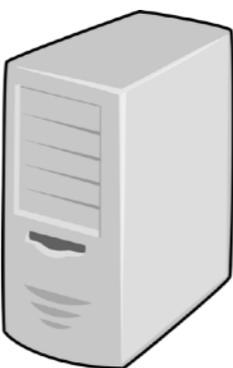


Submit data

X_a, Π_a



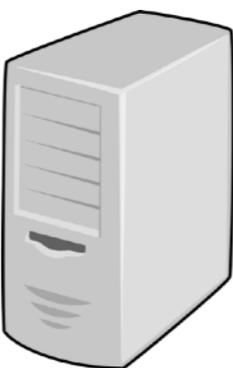
Using 128-bit integers



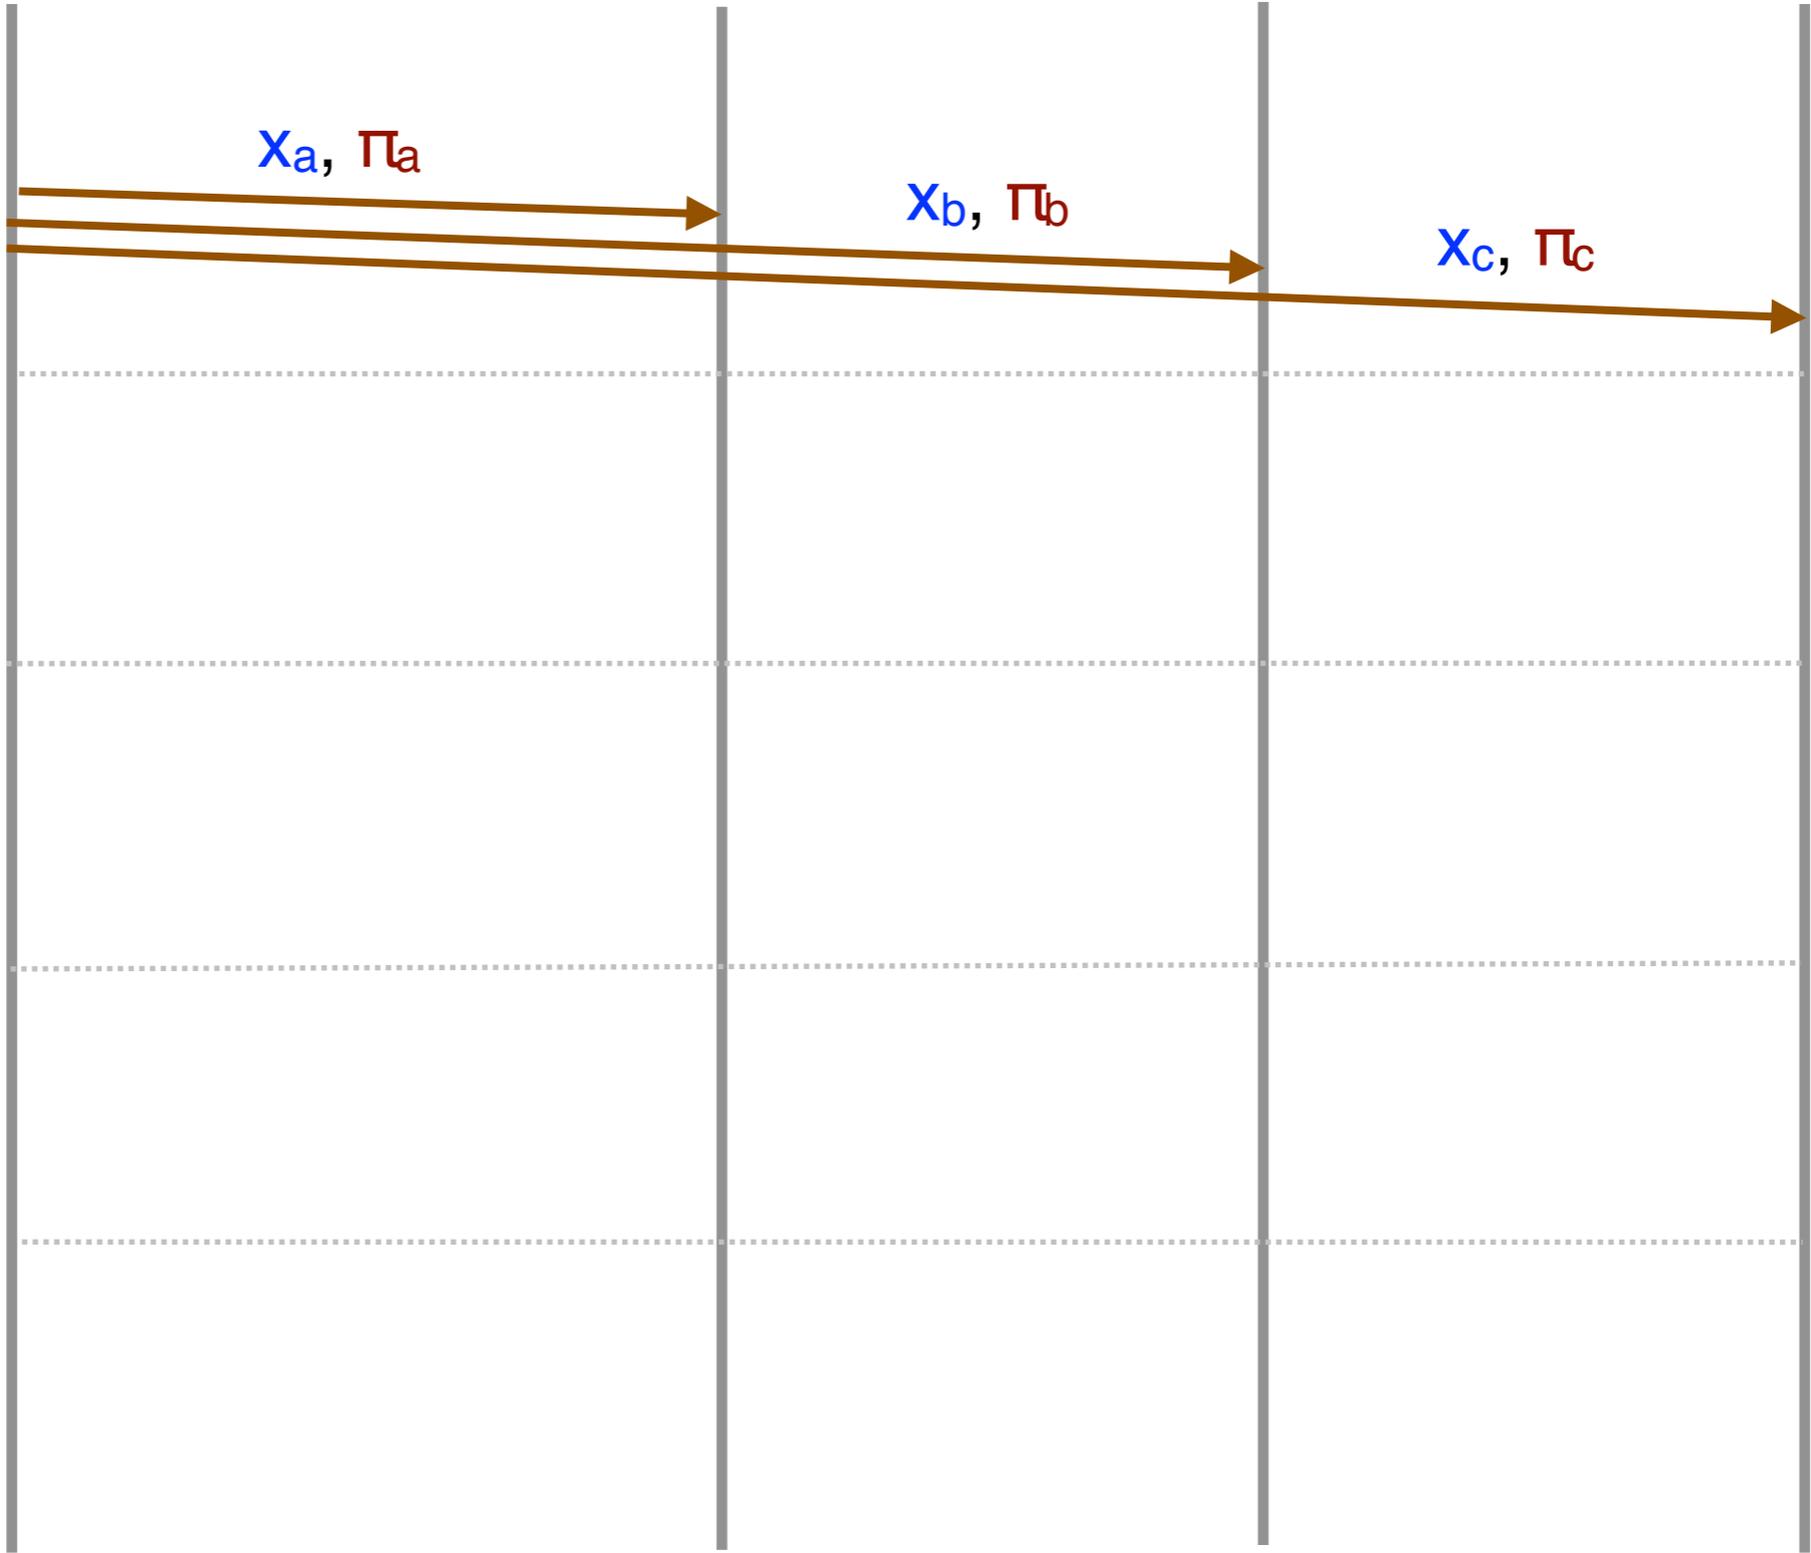
Submit data



Using 128-bit integers



Submit data



Using 128-bit integers



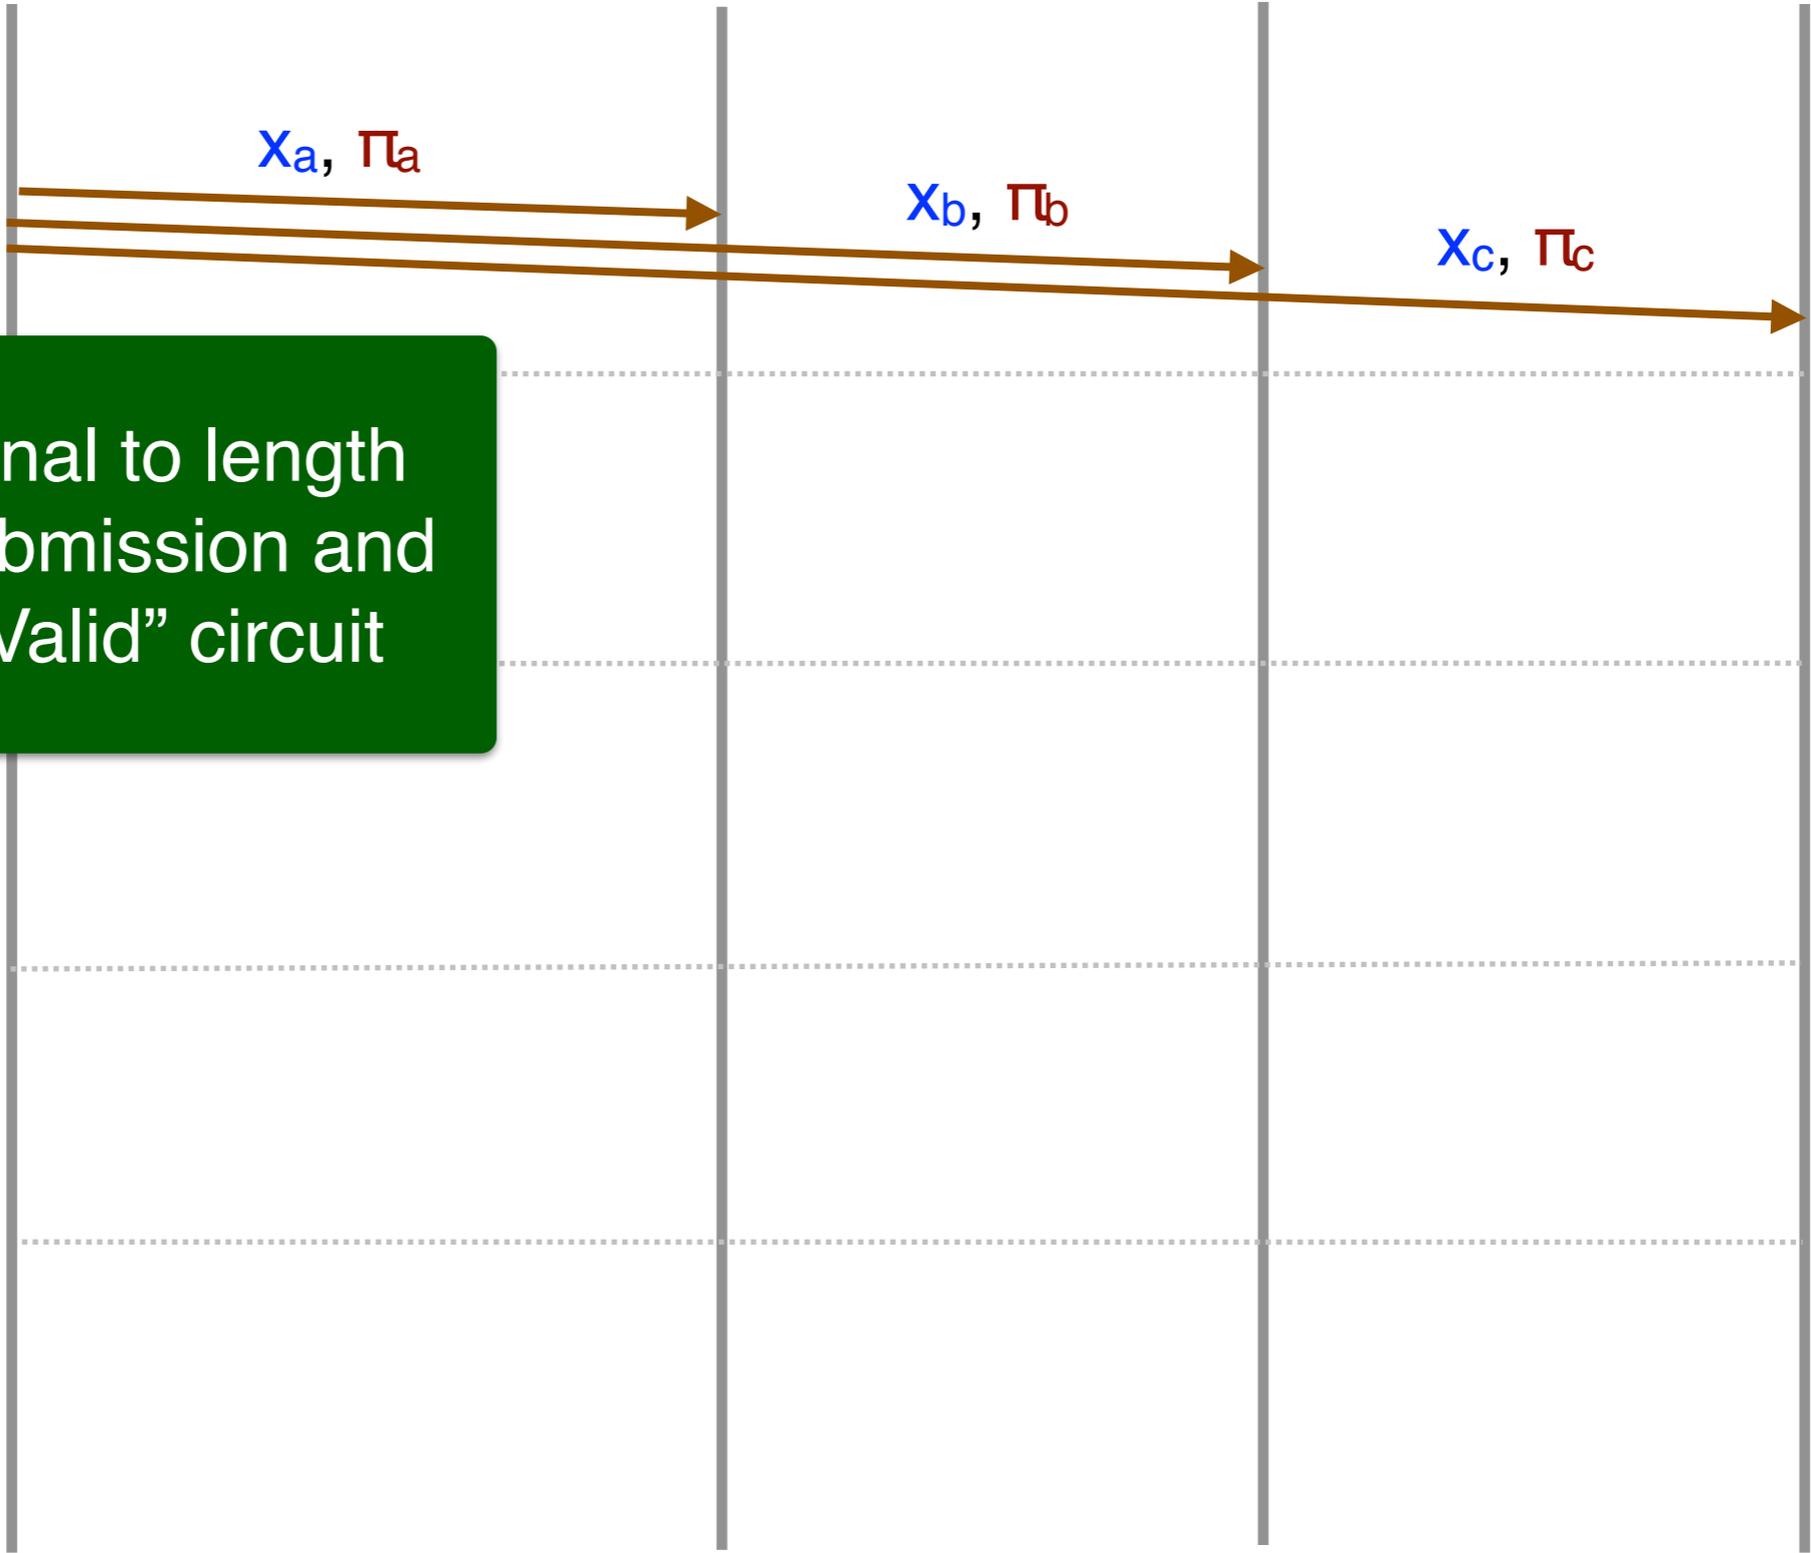
Submit data

X_a, Π_a

X_b, Π_b

X_c, Π_c

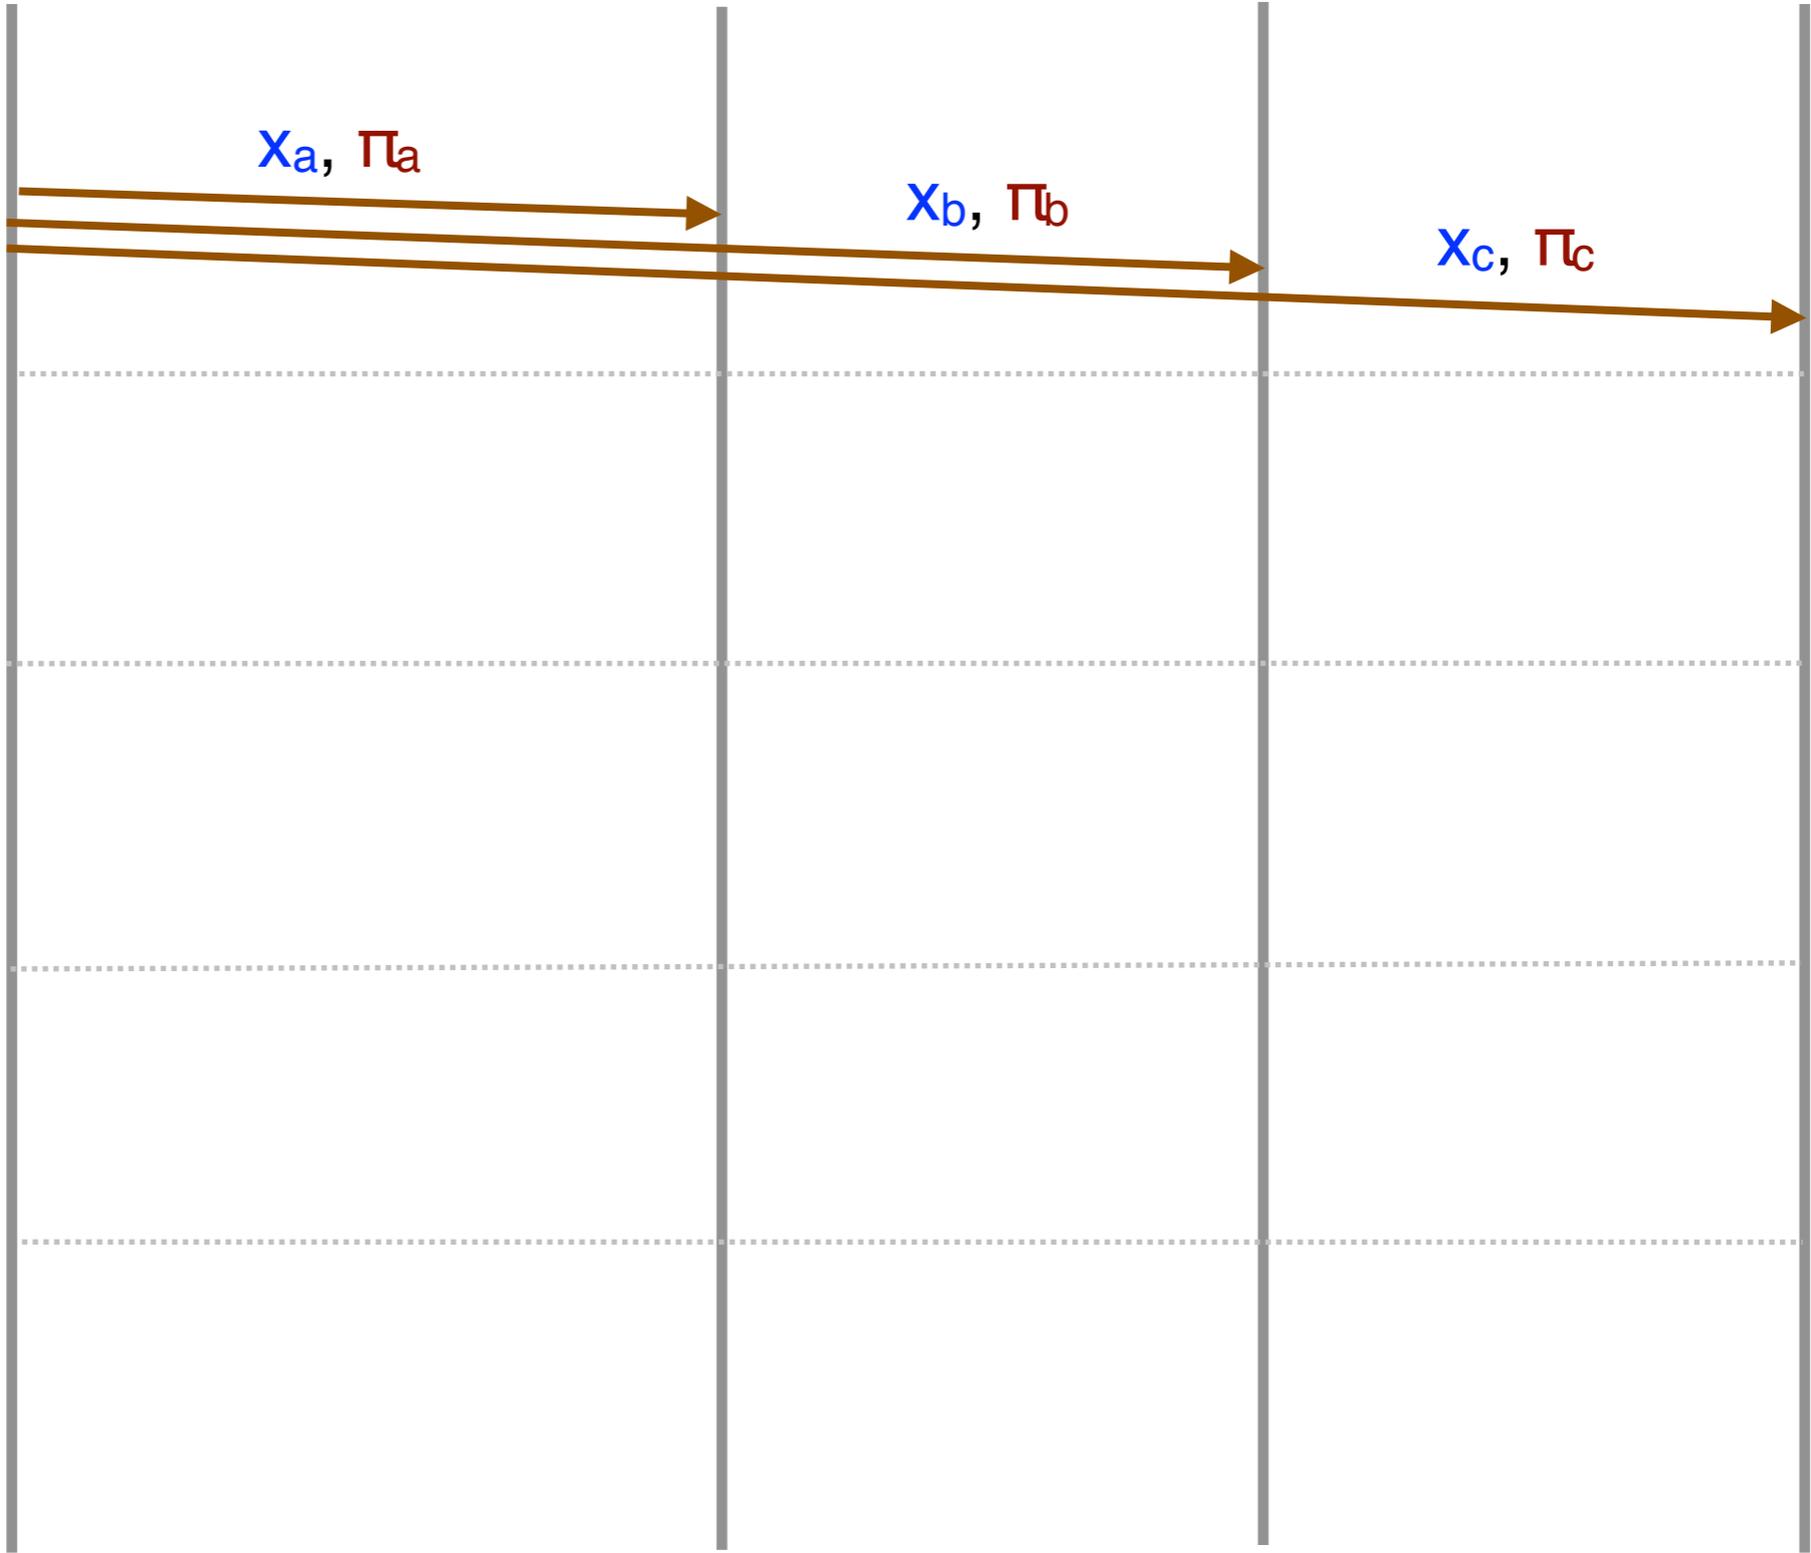
Proportional to length of data submission and size of "Valid" circuit



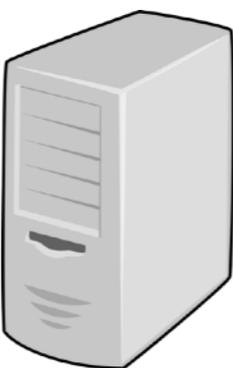
Using 128-bit integers



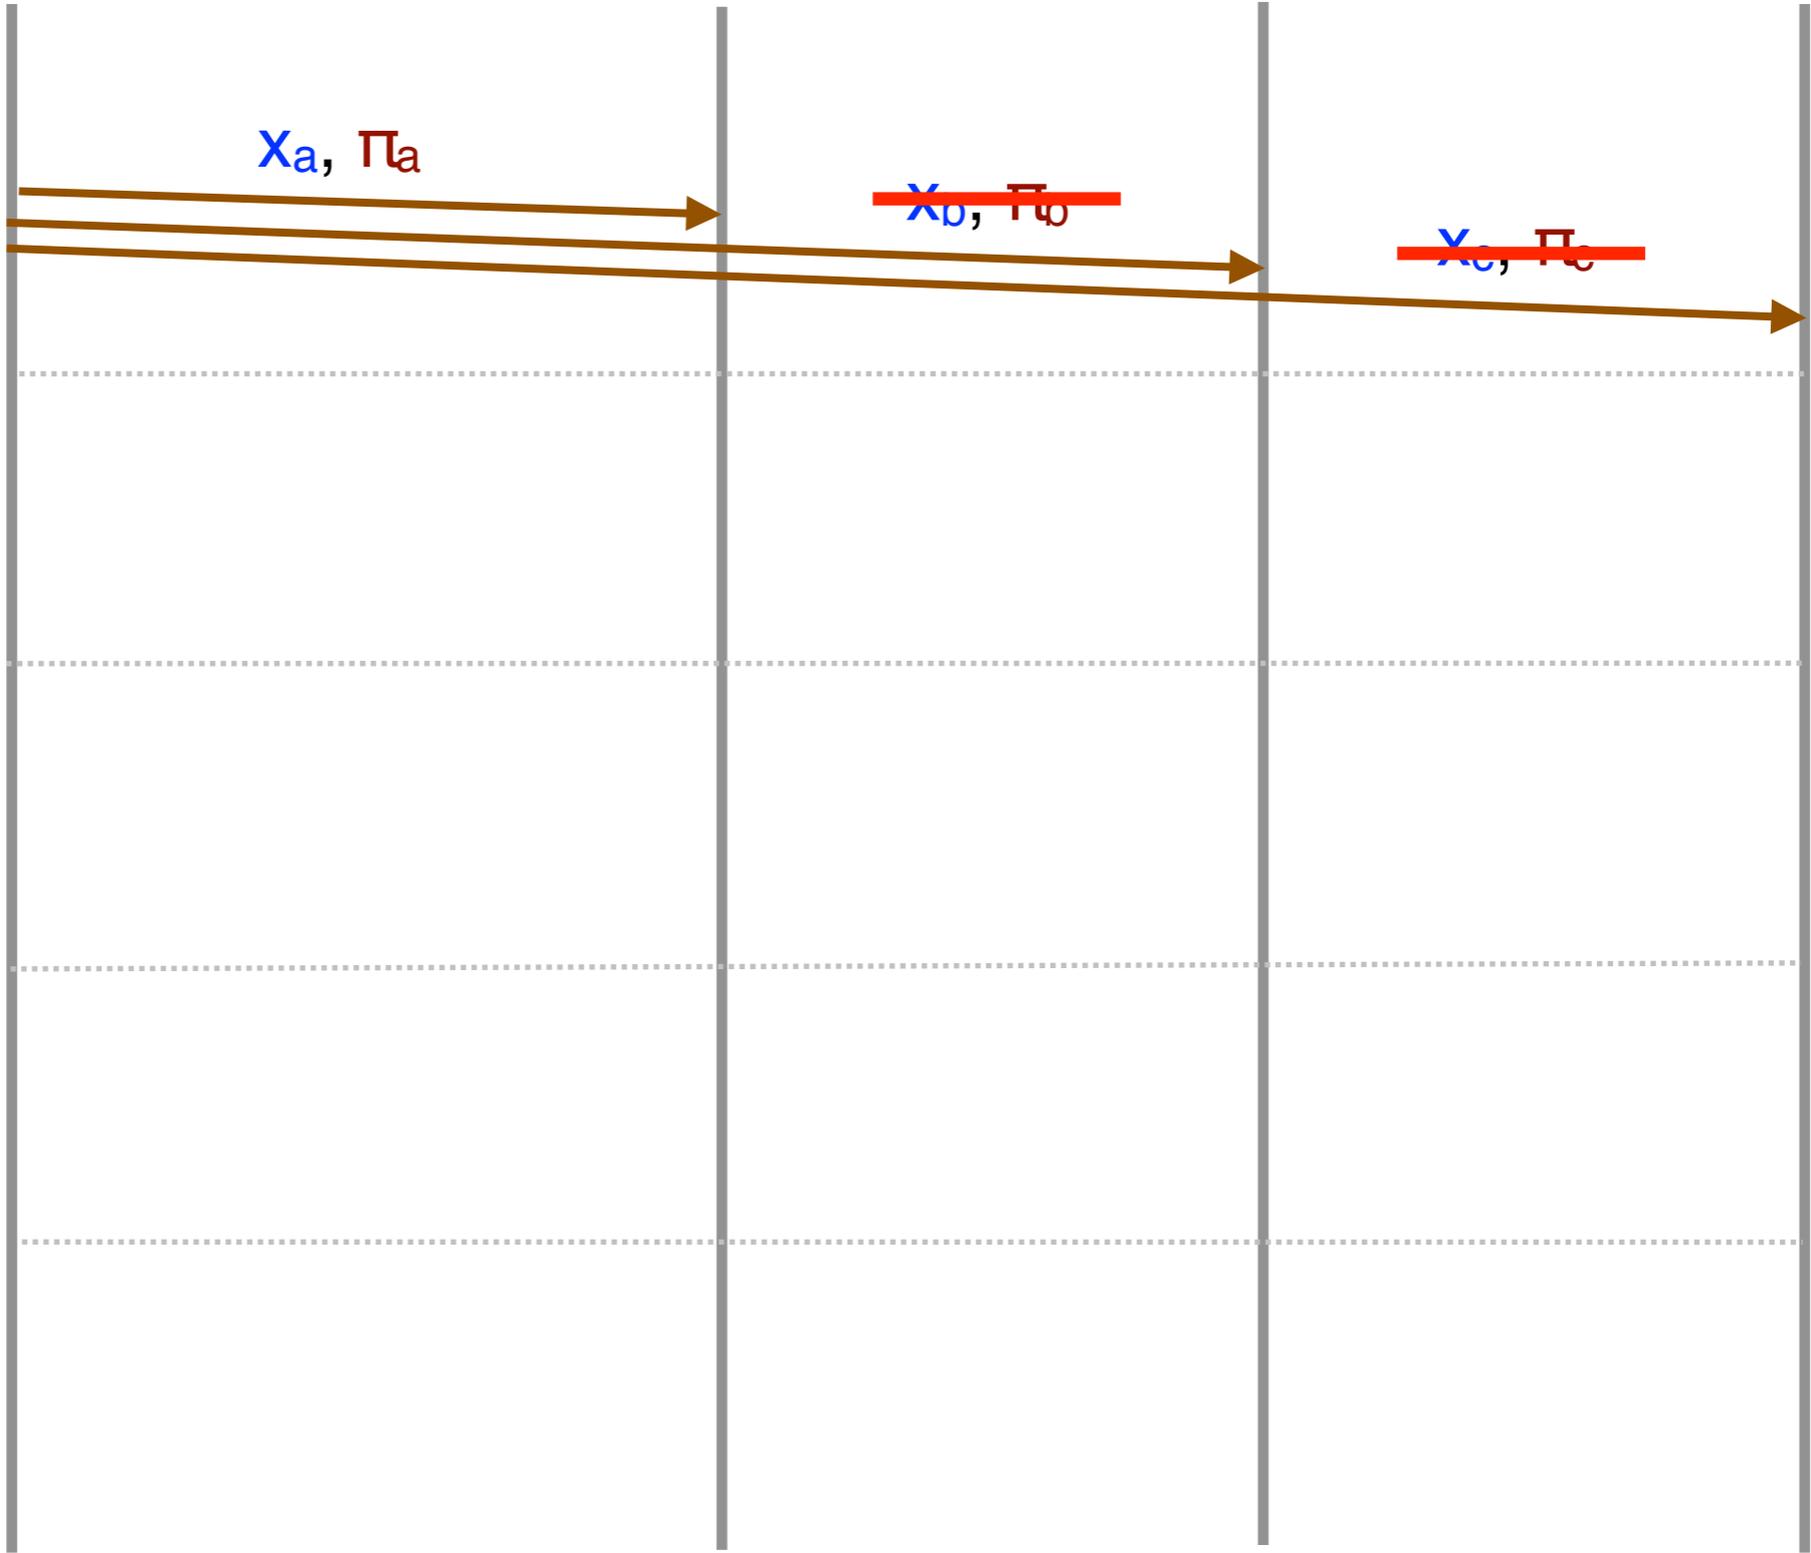
Submit data



Using 128-bit integers



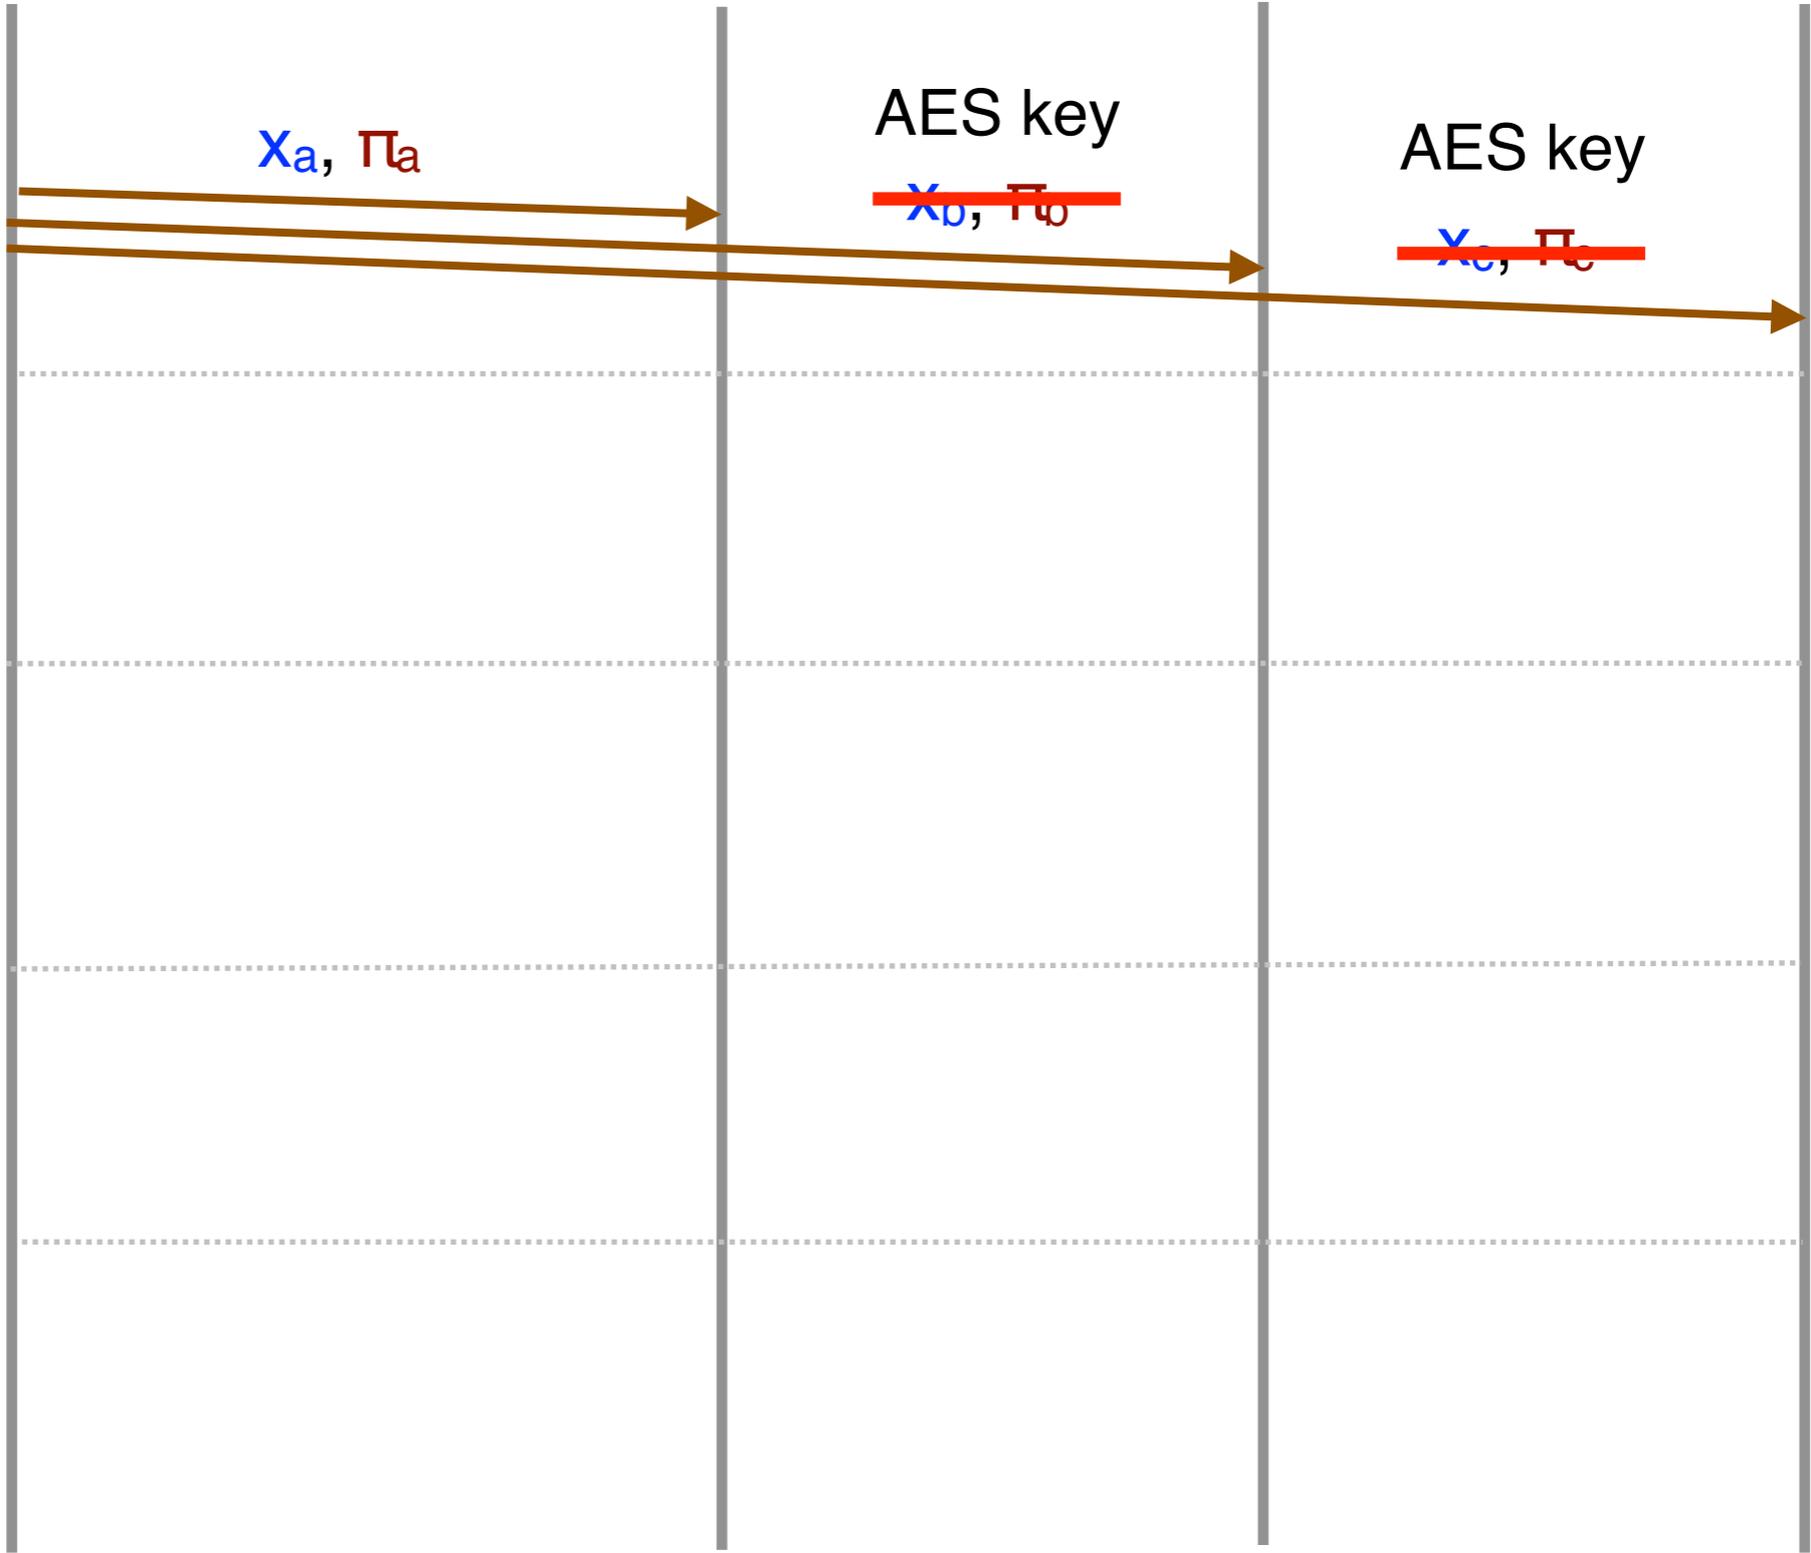
Submit data



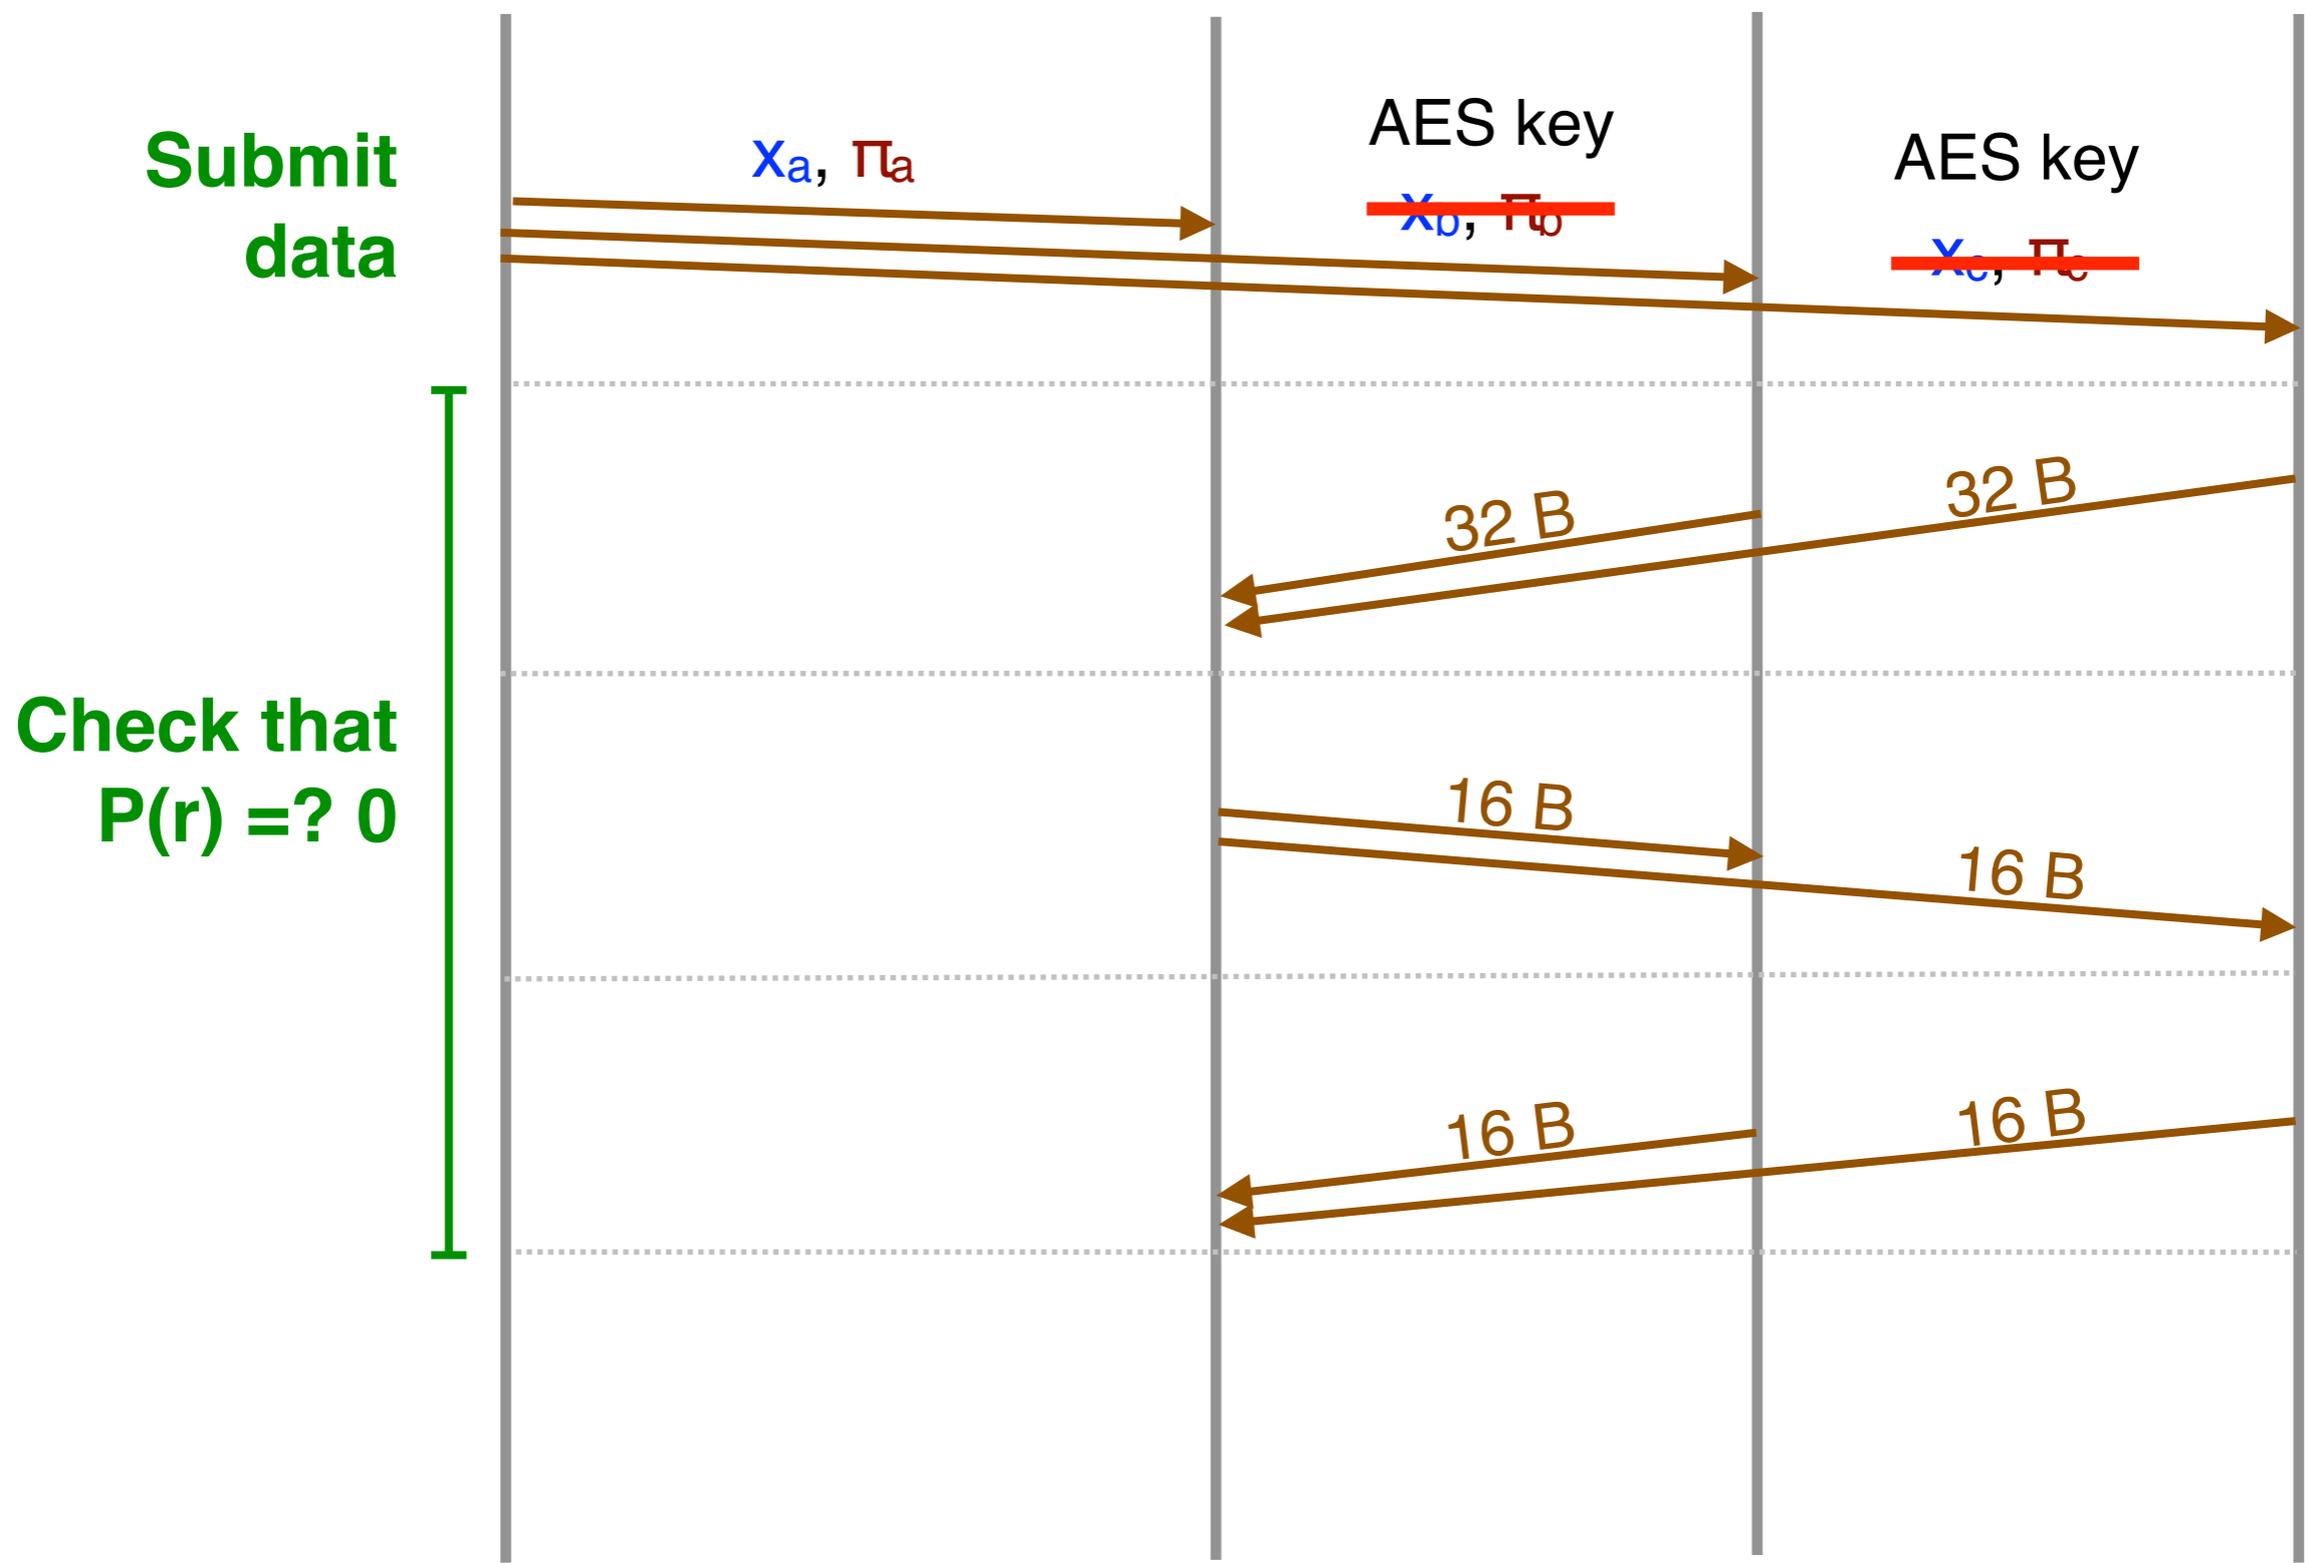
Using 128-bit integers



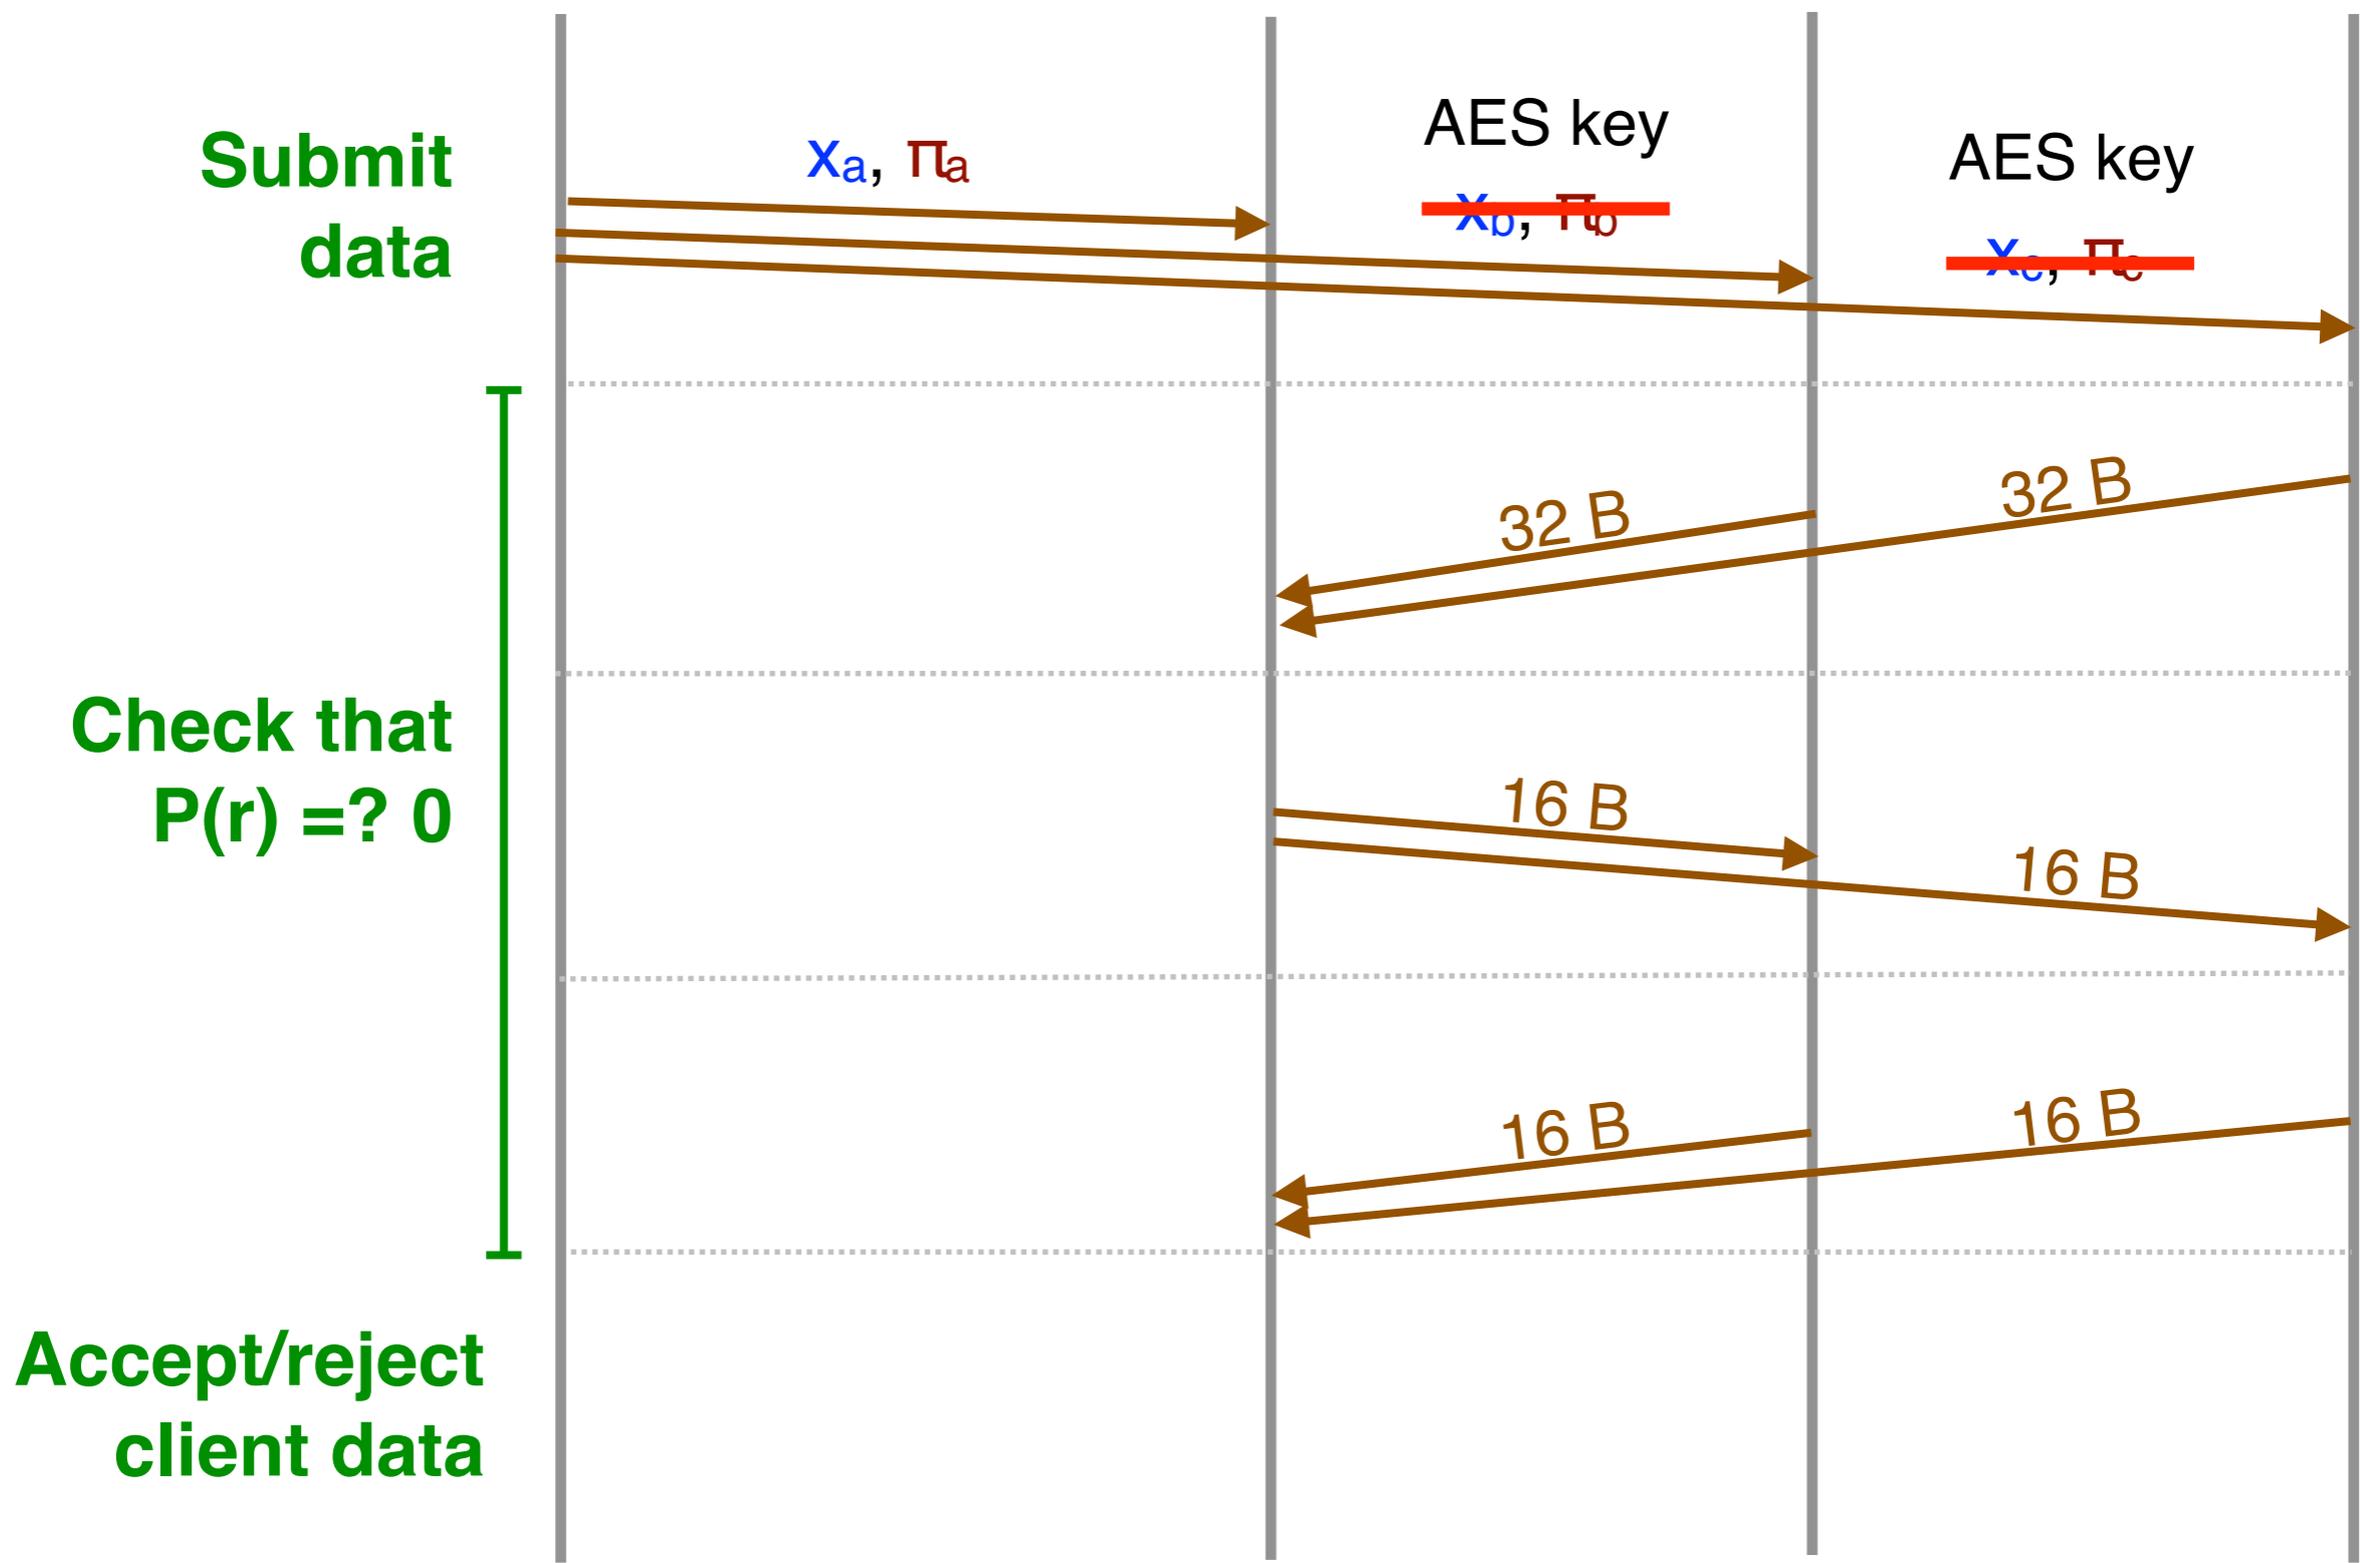
Submit data



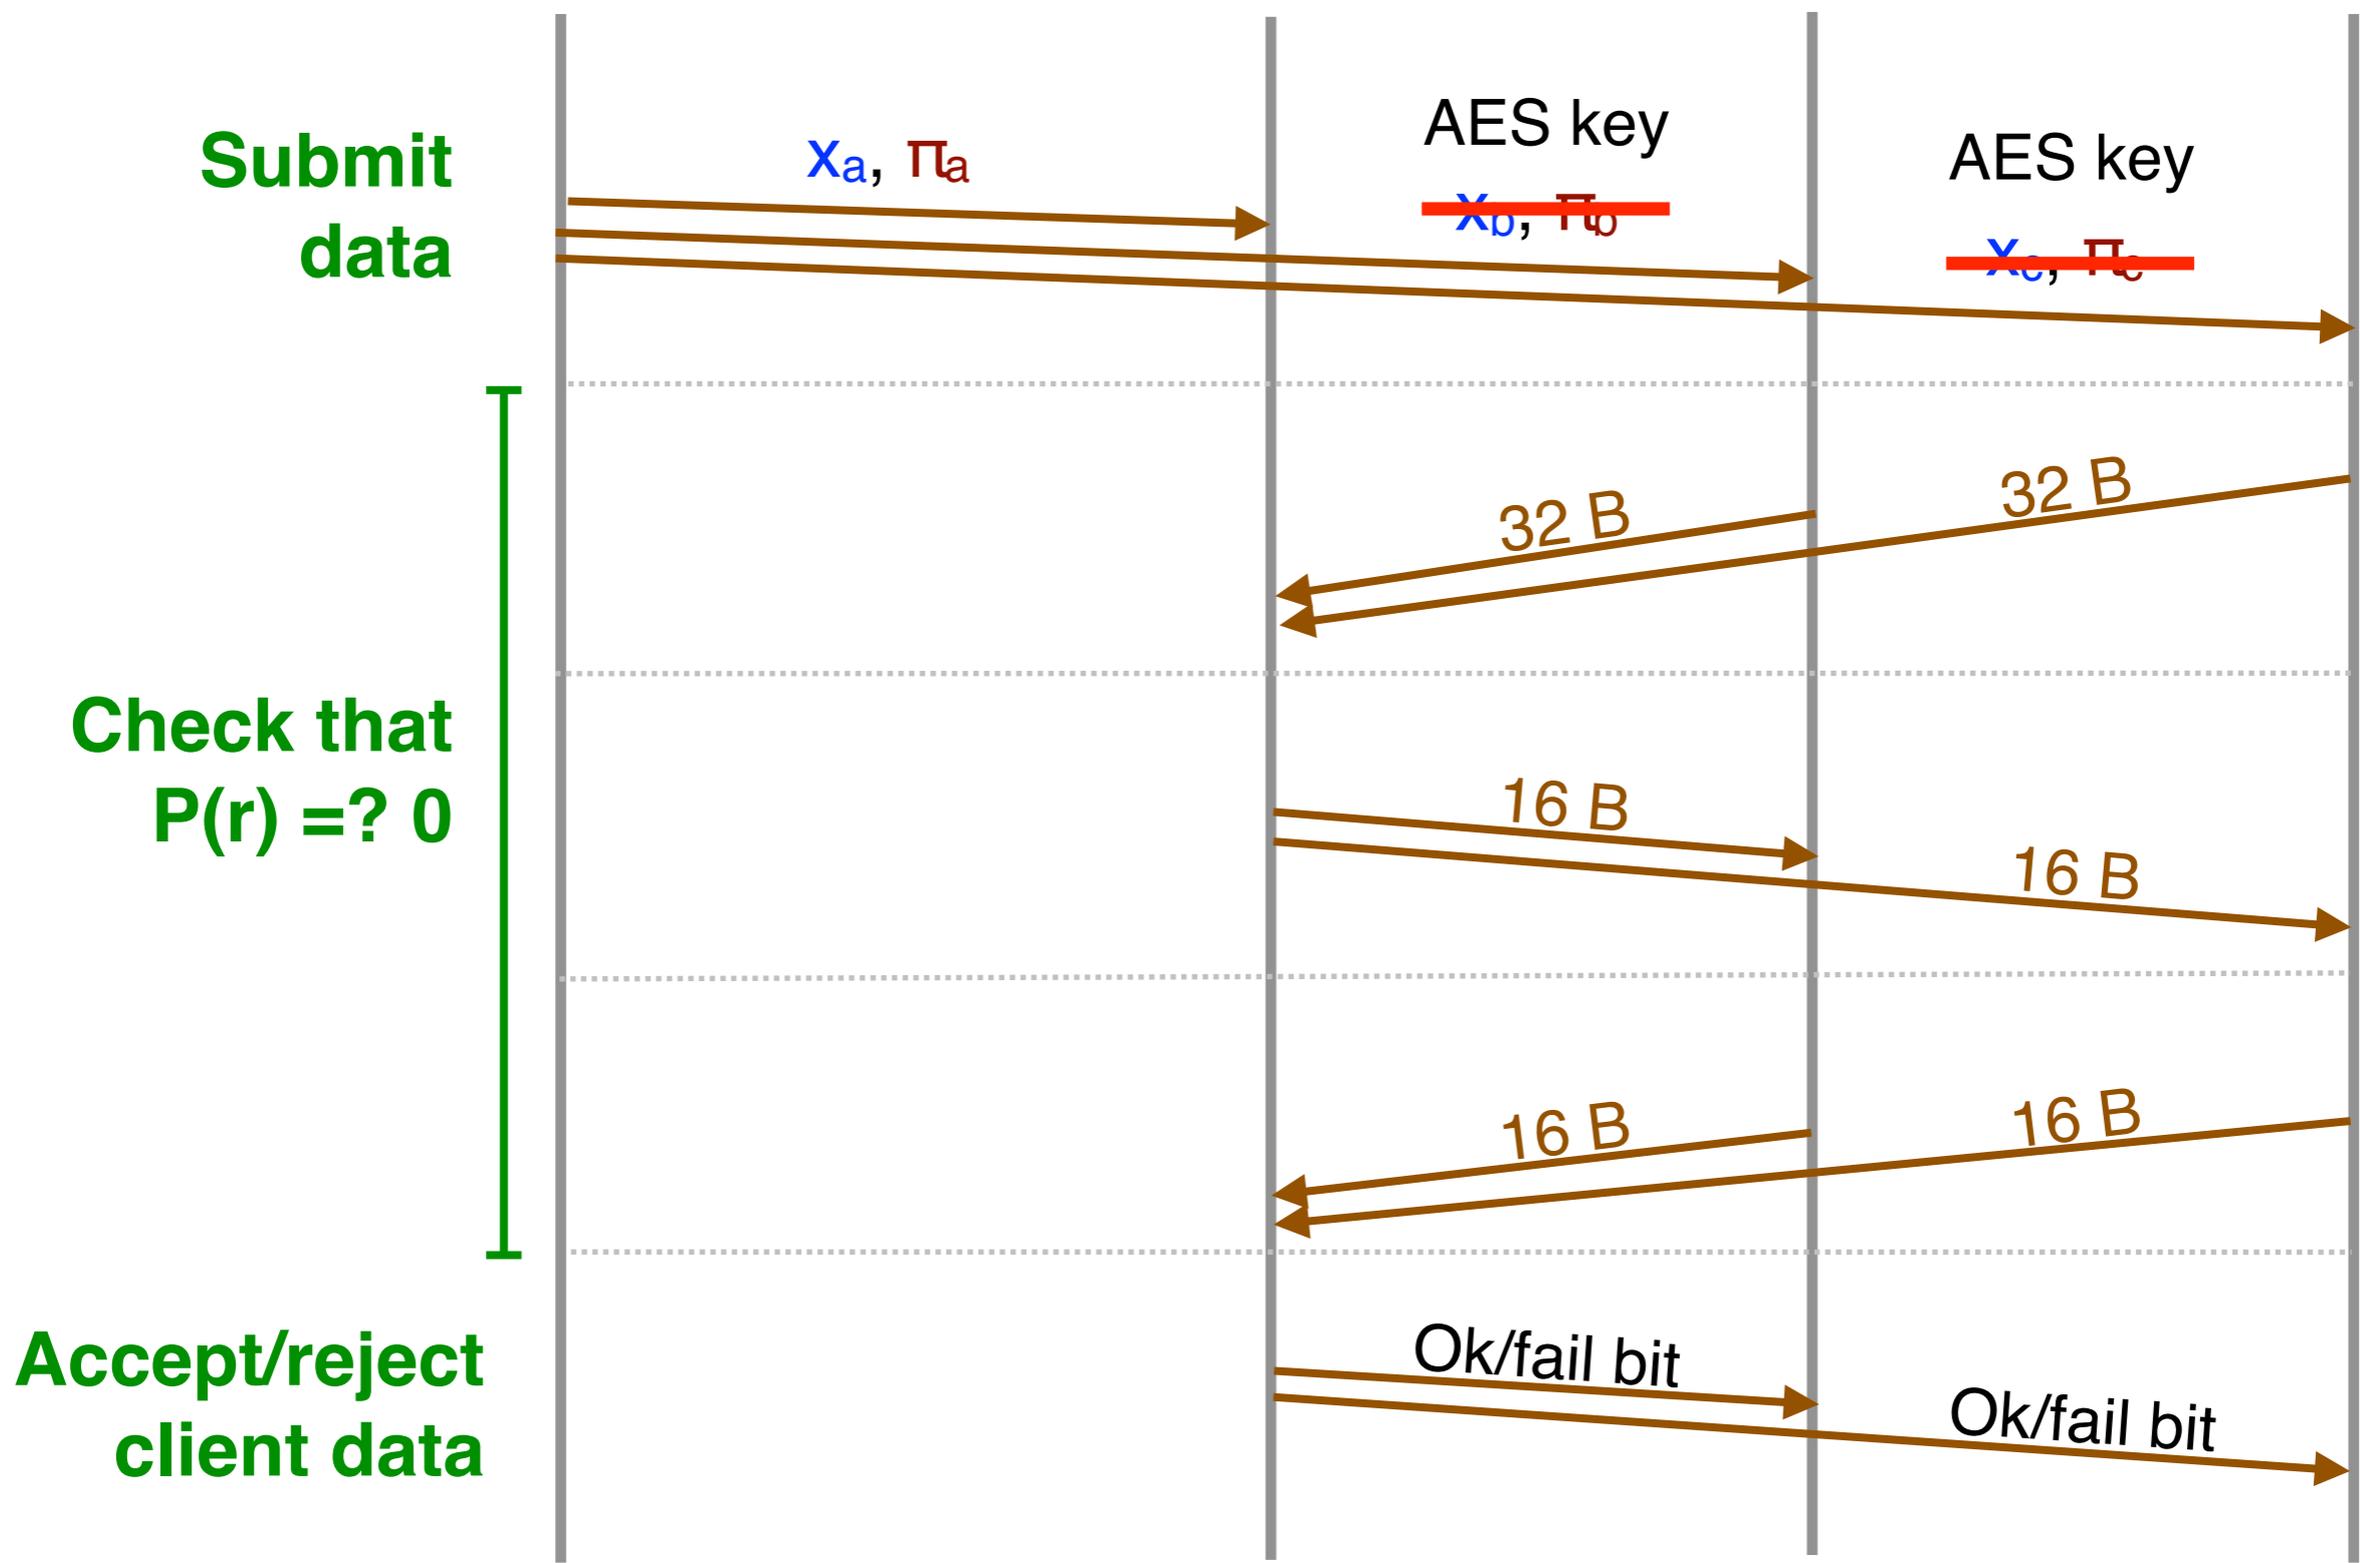
Using 128-bit integers



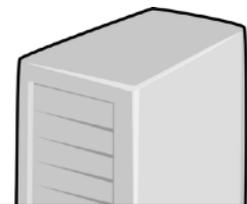
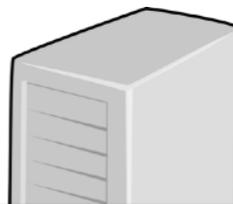
Using 128-bit integers



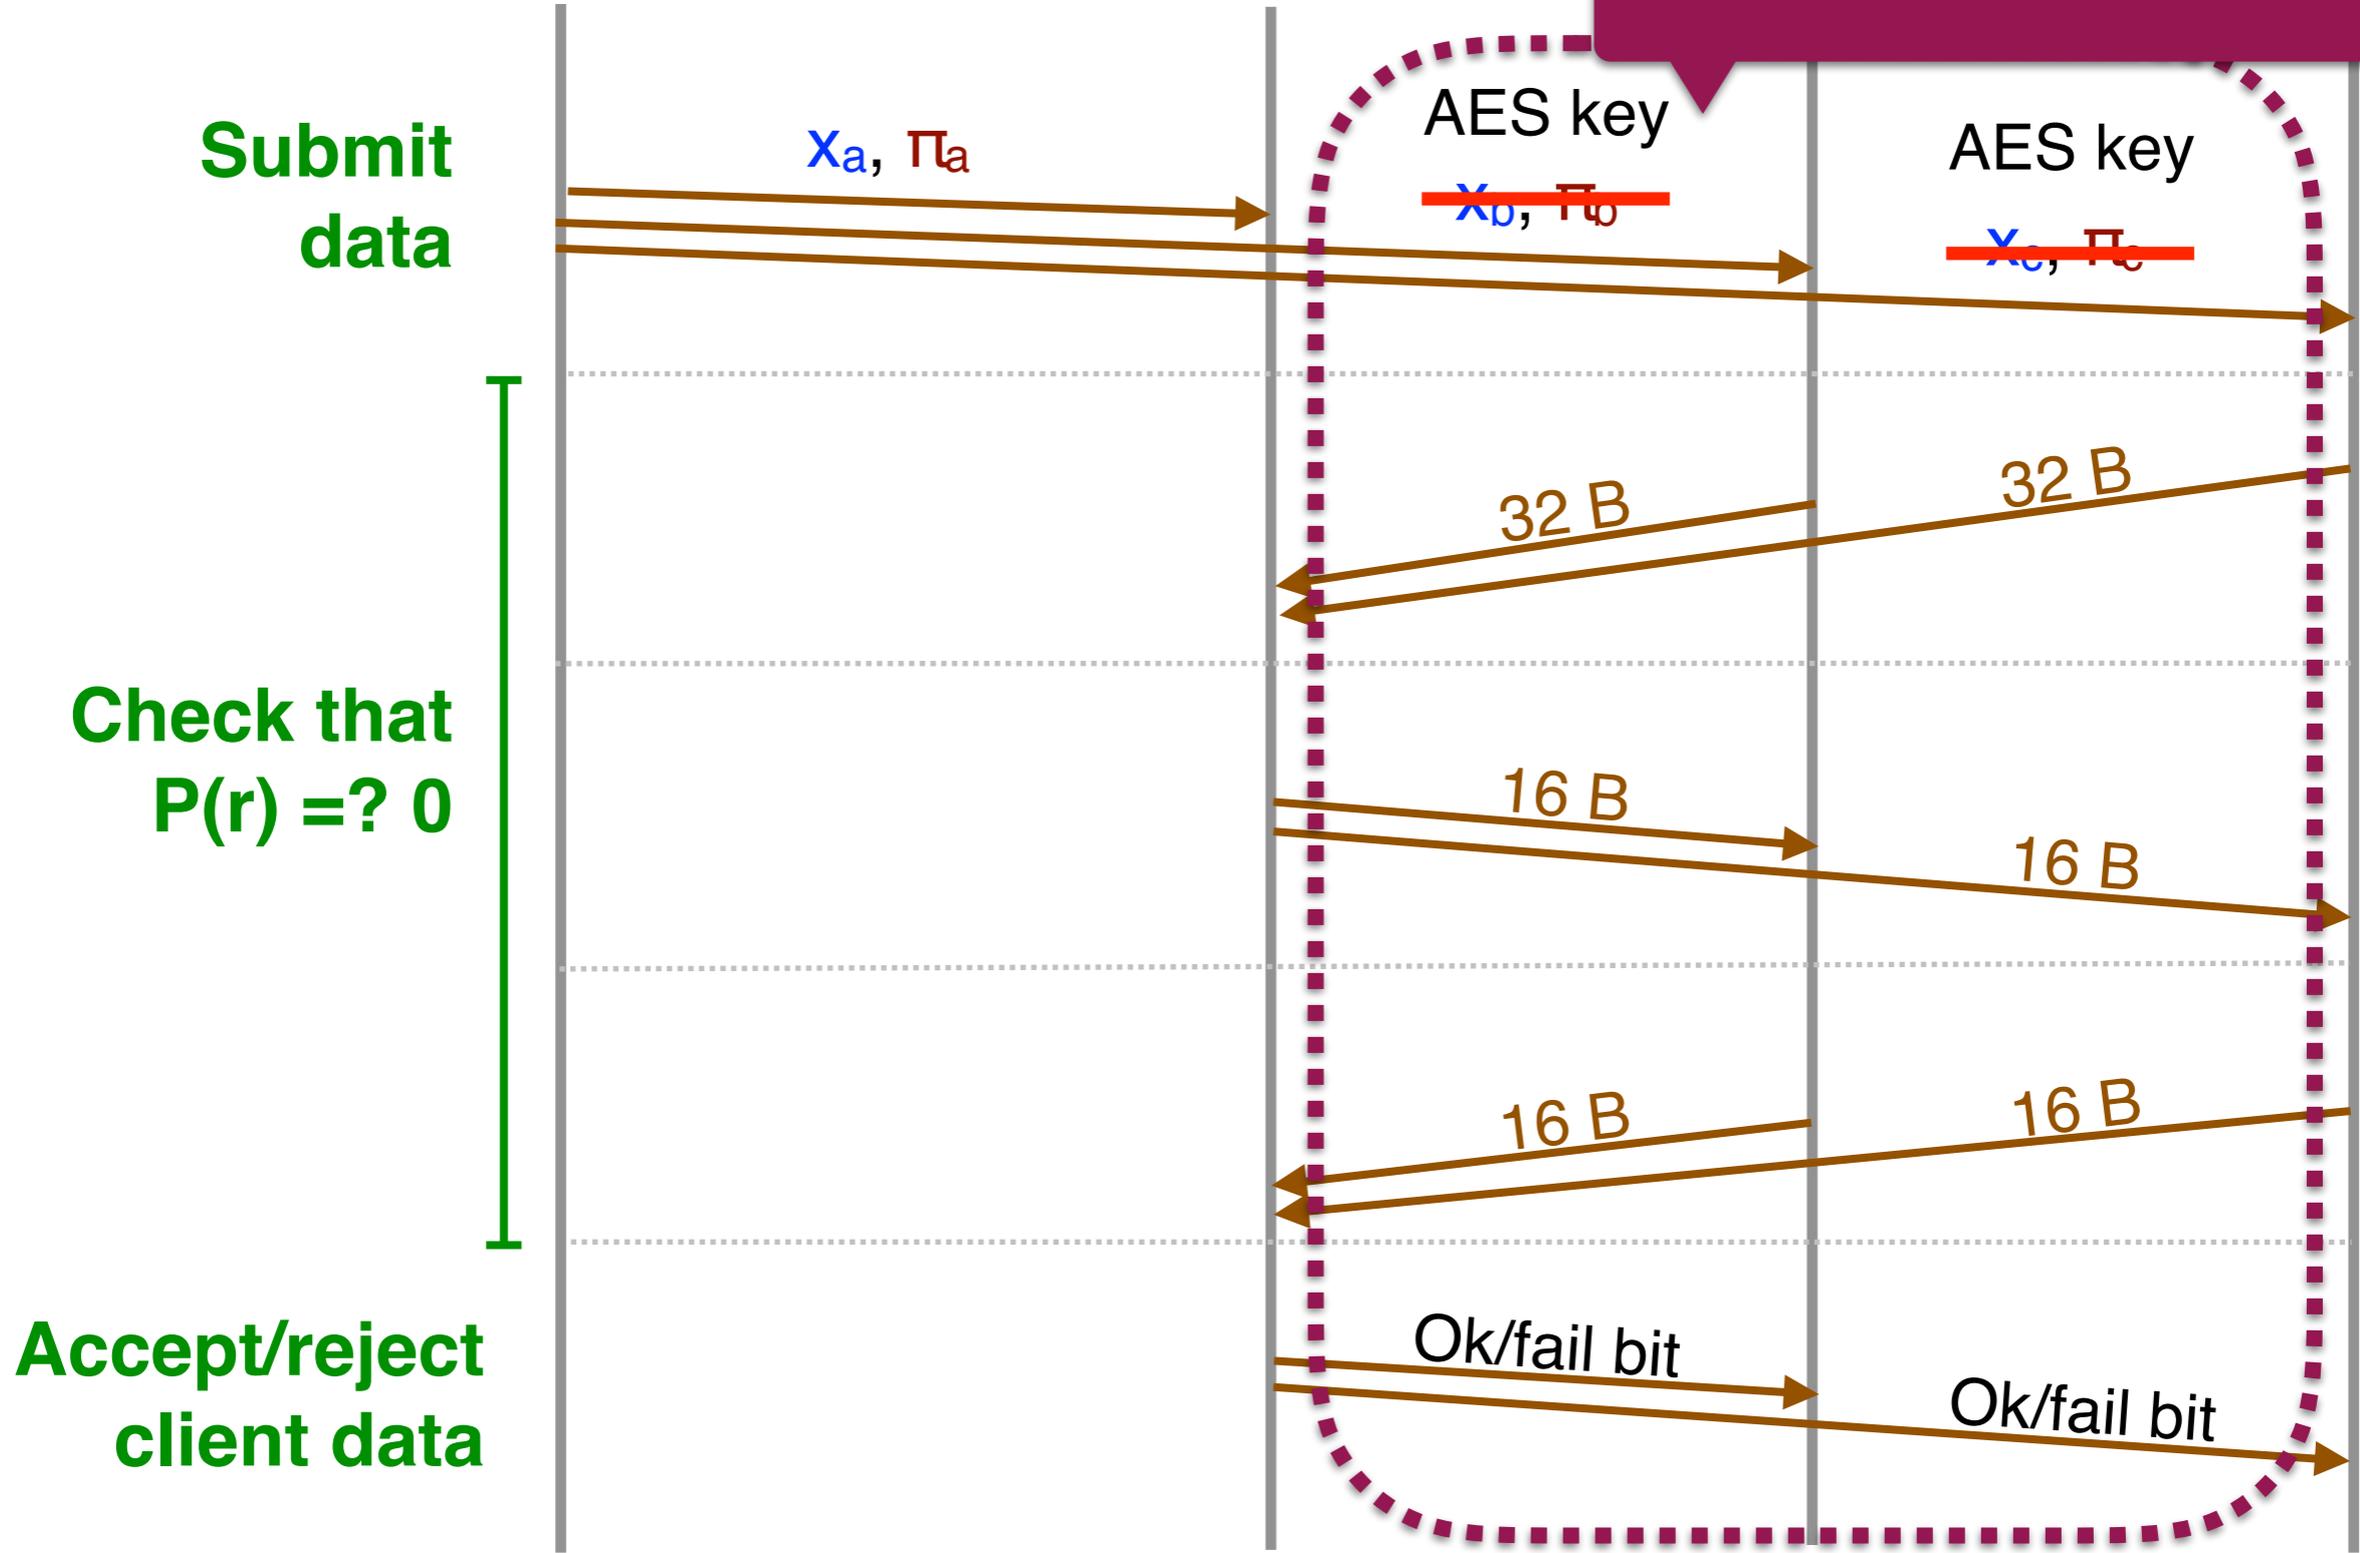
Using 128-bit integers



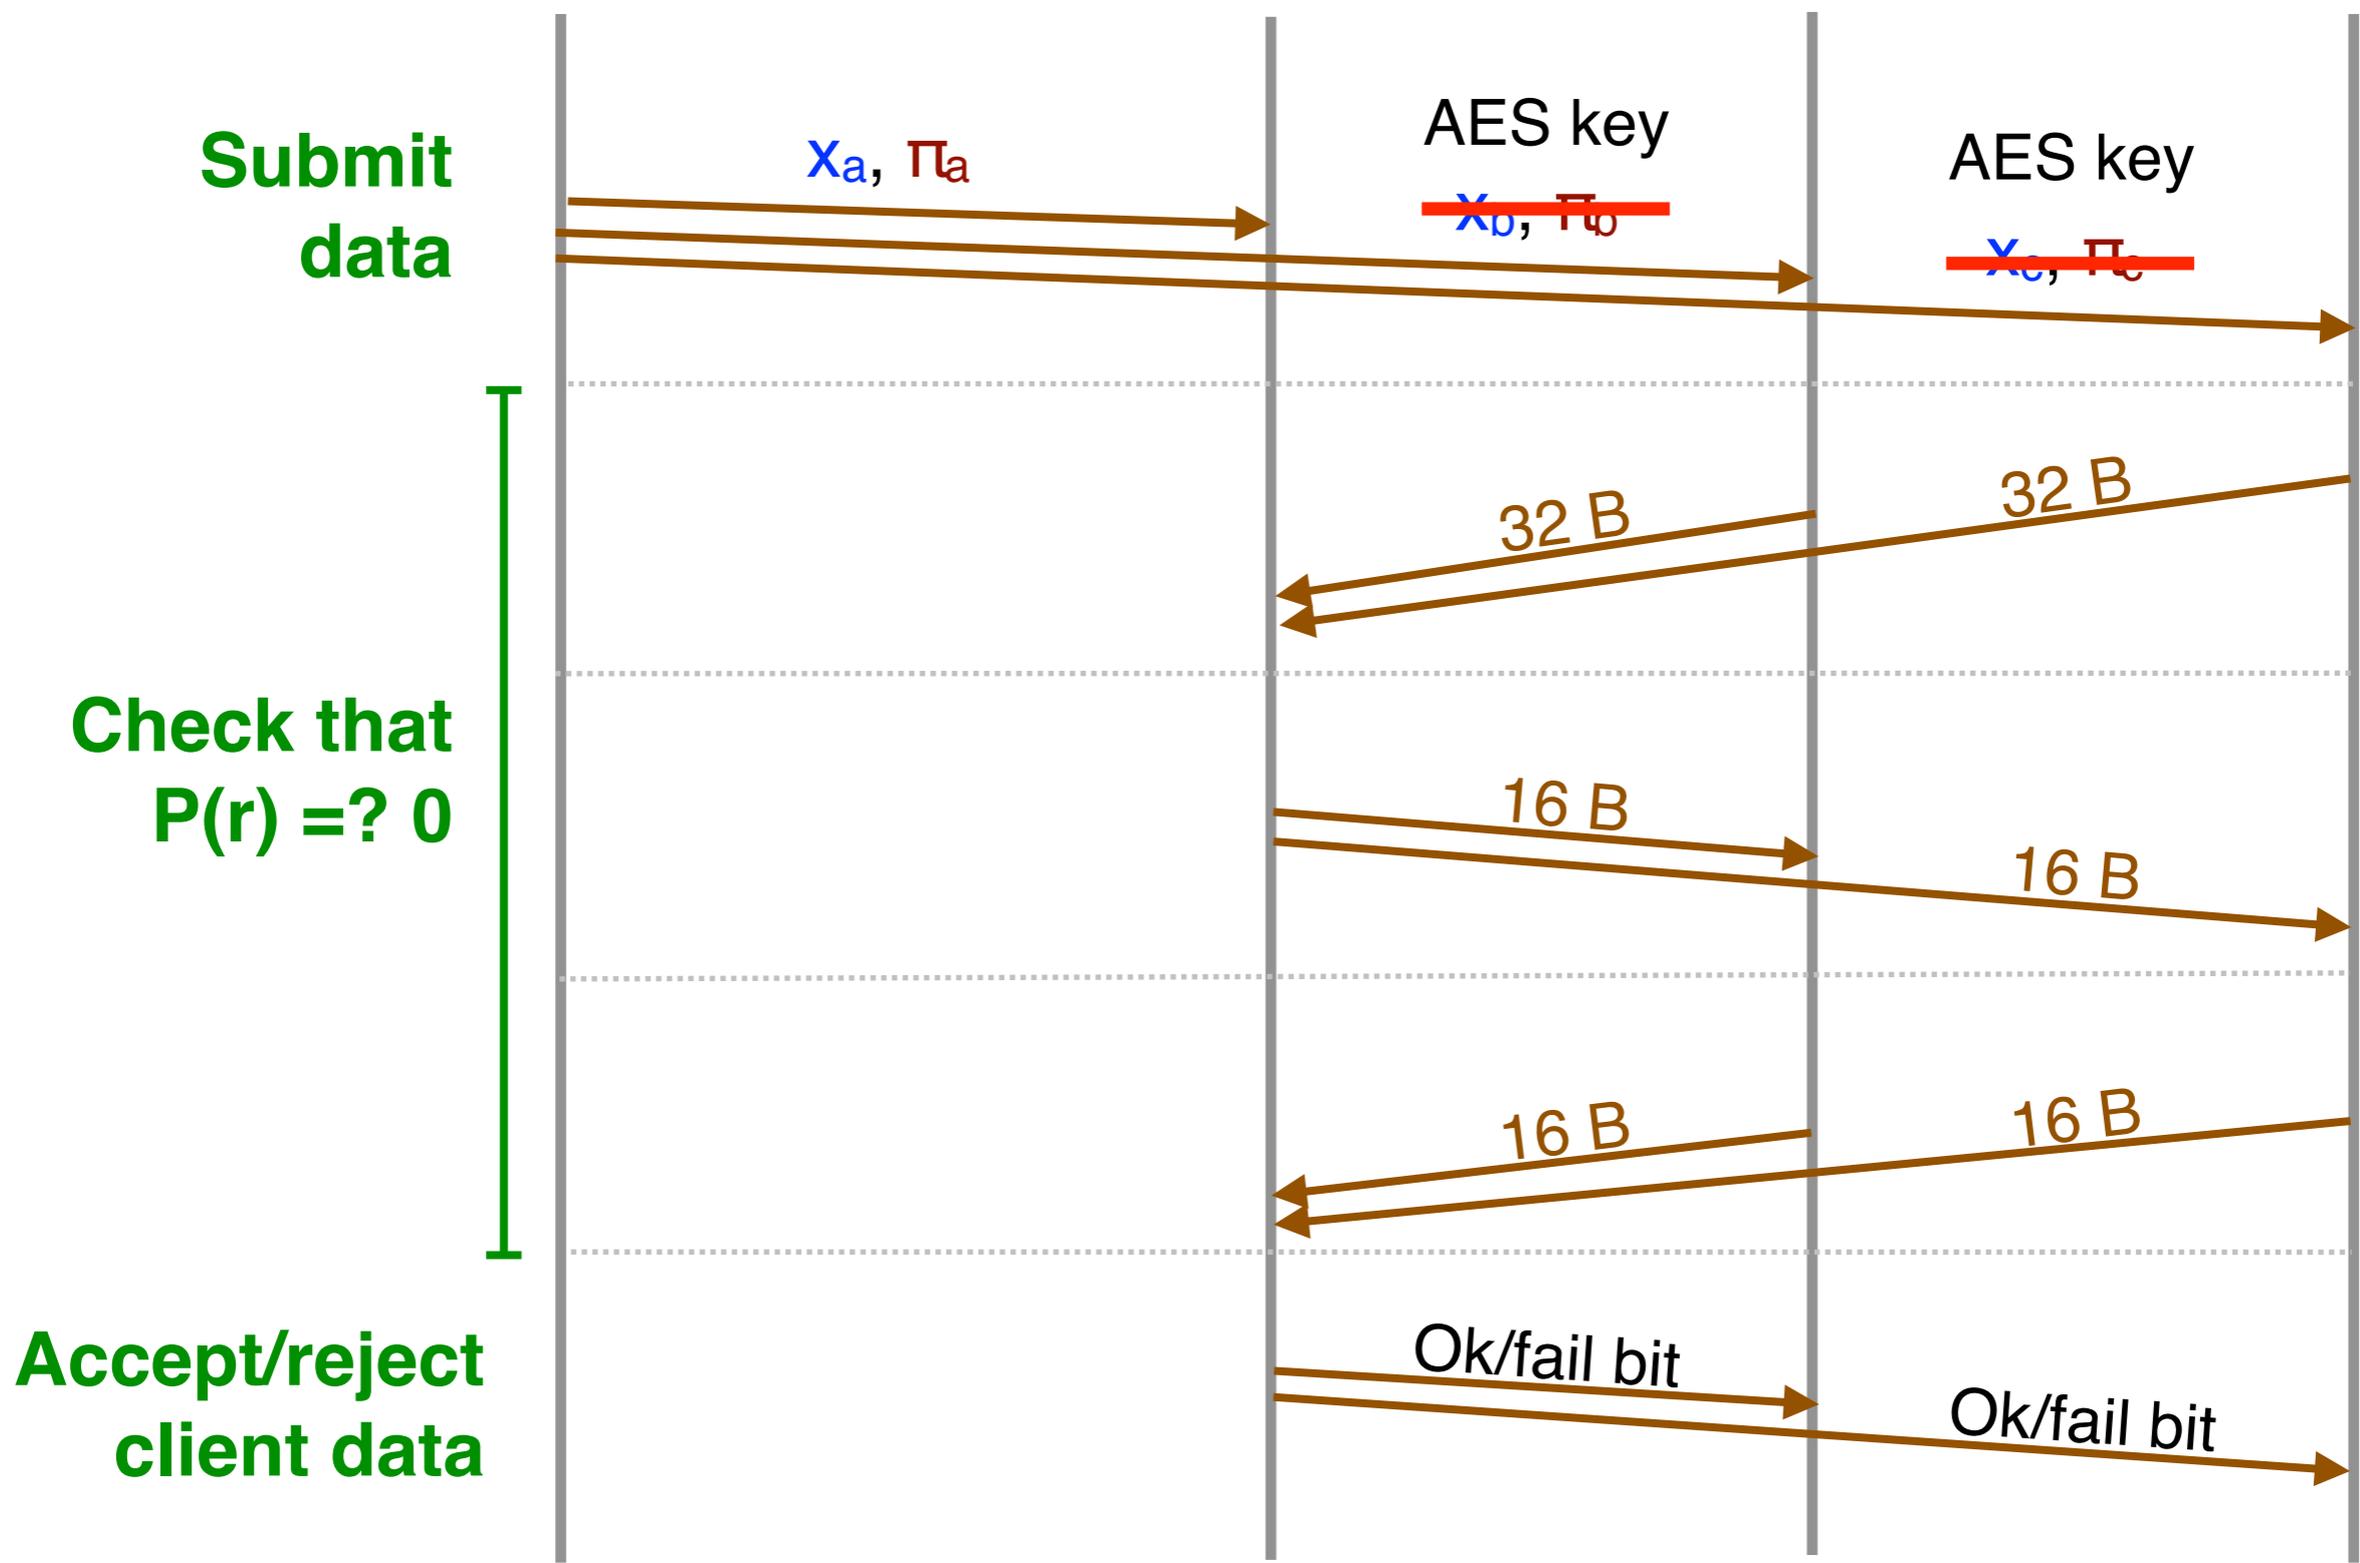
Using 128-bit integers



Does not grow with size of data or "Valid" circuit



Using 128-bit integers



Example Encoding: Average and Variance

[PrivStats11]

Example Encoding: Average and Variance

[PrivStats11]

- Each of N clients holds a 4-bit value x_i
- Servers want the **AVG** and **VAR** of the x_i s.

Each client encodes her value $x = b_3b_2b_1b_0$ as the tuple

$$(x, y) = (x, x^2, b_3, b_2, b_1, b_0)$$

Example Encoding: Average and Variance

[PrivStats11]

- Each of N clients holds a 4-bit value x_i
- Servers want the **AVG** and **VAR** of the x_i s.

Each client encodes her value $x = b_3b_2b_1b_0$ as the tuple

$$(x, y) = (x, x^2, b_3, b_2, b_1, b_0)$$

To test validity of the encoding, check that:

$$\text{Valid}(x, y) = \begin{cases} (x^2 - y) = 0 & \text{— } y \text{ is } x^2 \\ x - \sum_j 2^j b_j = 0 & \text{— } b\text{'s are the bits of } x \\ b_j \cdot (b_j - 1) = 0 & \text{— } b\text{'s are 0/1 values} \end{cases}$$