

SiTP Dec. 2017



Cryptography for IoT

Dan Boneh
Stanford University

... but first: Computer Security at Stanford



Alex Aiken

software analysis



Dan Boneh

applied Crypto,
crypto currencies



Matei Zaharia

security and big data



Dawson Engler

automated
bug finding

David Mazières

Op. Systems



Phil Levis

IoT Security




John Mitchell

protocol design,
online ed.



Mendel Rosenblum

VM' s in security



Courses

➤ Courses:

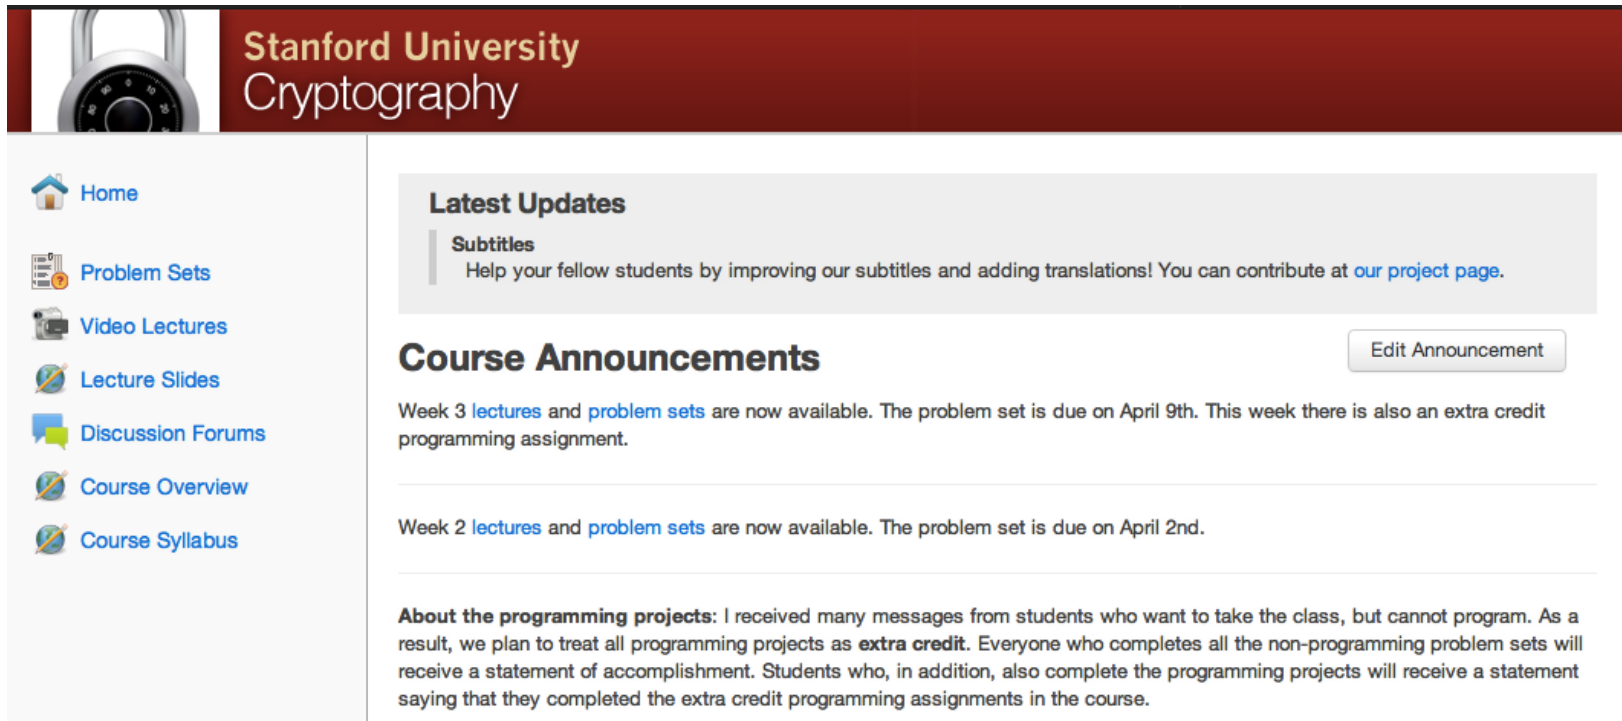
- CS55N (freshmen seminar): ten ideas in computer security
- **CS155: Computer Security**
- **CS251: Crypto currencies and blockchain technologies**
- **CS255: Intro to Crypto**
- CS259: Security analysis of network protocols
- CS355: Graduate course in cryptography

➤ Stanford Advanced Computer Security Certificate

<http://scpd.stanford.edu/computerSecurity/>

Online Courses

//www.coursera.org/learn/crypto



The image shows a screenshot of the Stanford University Cryptography course page on Coursera. The header features a padlock icon and the text "Stanford University Cryptography". A left sidebar contains navigation links: Home, Problem Sets, Video Lectures, Lecture Slides, Discussion Forums, Course Overview, and Course Syllabus. The main content area includes a "Latest Updates" section with a "Subtitles" announcement, a "Course Announcements" section with two announcements, and a section titled "About the programming projects" explaining the extra credit policy.

Stanford University
Cryptography

- Home
- Problem Sets
- Video Lectures
- Lecture Slides
- Discussion Forums
- Course Overview
- Course Syllabus

Latest Updates

Subtitles
Help your fellow students by improving our subtitles and adding translations! You can contribute at [our project page](#).

Course Announcements [Edit Announcement](#)

Week 3 [lectures](#) and [problem sets](#) are now available. The problem set is due on April 9th. This week there is also an extra credit programming assignment.

Week 2 [lectures](#) and [problem sets](#) are now available. The problem set is due on April 2nd.

About the programming projects: I received many messages from students who want to take the class, but cannot program. As a result, we plan to treat all programming projects as **extra credit**. Everyone who completes all the non-programming problem sets will receive a statement of accomplishment. Students who, in addition, also complete the programming projects will receive a statement saying that they completed the extra credit programming assignments in the course.

Course open to the public

Free Book Draft

A Graduate Course in Applied Cryptography

Dan Boneh and Victor Shoup

Free at: [//cryptobook.us](http://cryptobook.us)

Please send us comments

Multiparty computation (MPC) and SGX

MPC for genomic data analysis

[Jagadeesh, Wu, Birgmeier, Boneh, Bejerano, *Science* 2017]

What genes causes a specific disorder?

$$v_1: \begin{bmatrix} 0 & 1 & 0 & 2 & 0 & 1 \\ 1 & 0 & 1 & 2 & 0 & 1 \\ v_3: \begin{bmatrix} 2 & 0 & 0 & 2 & 1 & 1 \\ 0 & 0 & 1 & 2 & 0 & 1 \end{bmatrix}$$



People with Kabuki syndrome

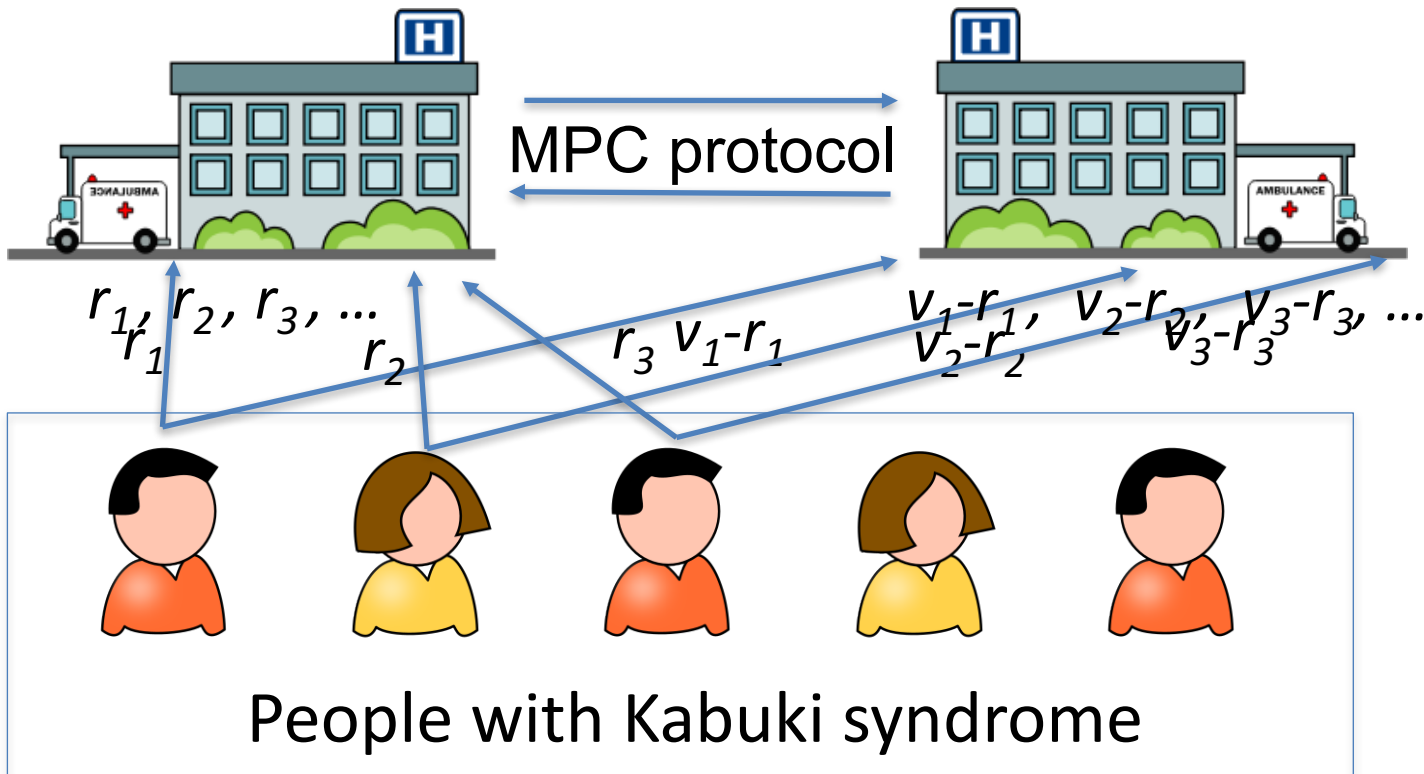


Each has 211 to 374 rare genes out of $\approx 20,000$ genes

Patient i : vector v_i of dim 20,000 that is 0 for normal genes

MPC for genomic data analysis

[Jagadeesh, Wu, Birgmeier, Boneh, Bejerano, 2017]

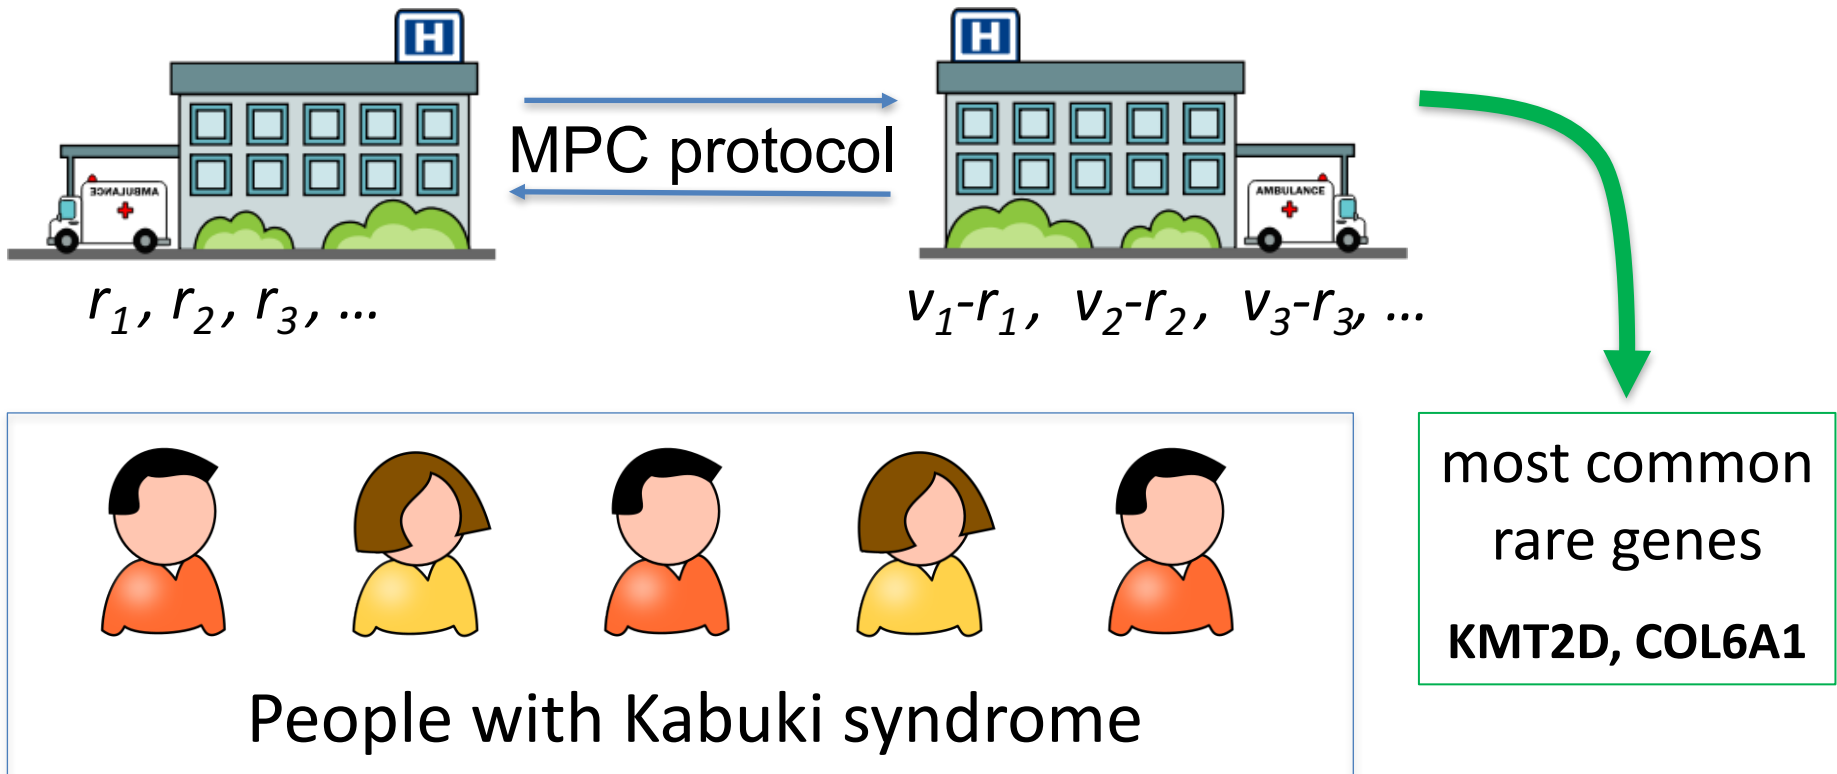


Each has 211 to 374 rare genes out of $\approx 20,000$ genes

Patient i : vector v_i of dim 20,000 that is 0 for normal genes

MPC for genomic data analysis

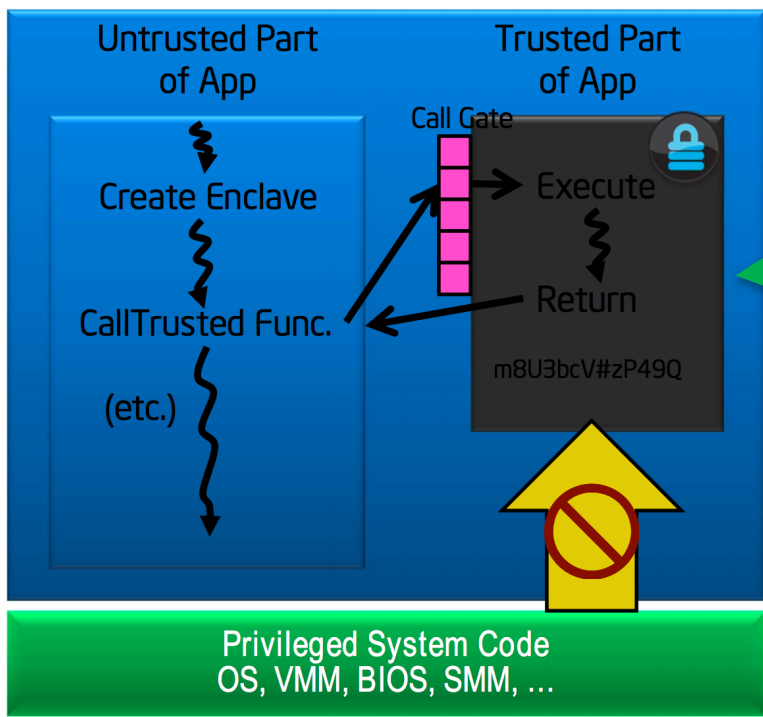
[Jagadeesh, Wu, Birgmeier, Boneh, Bejerano, 2017]



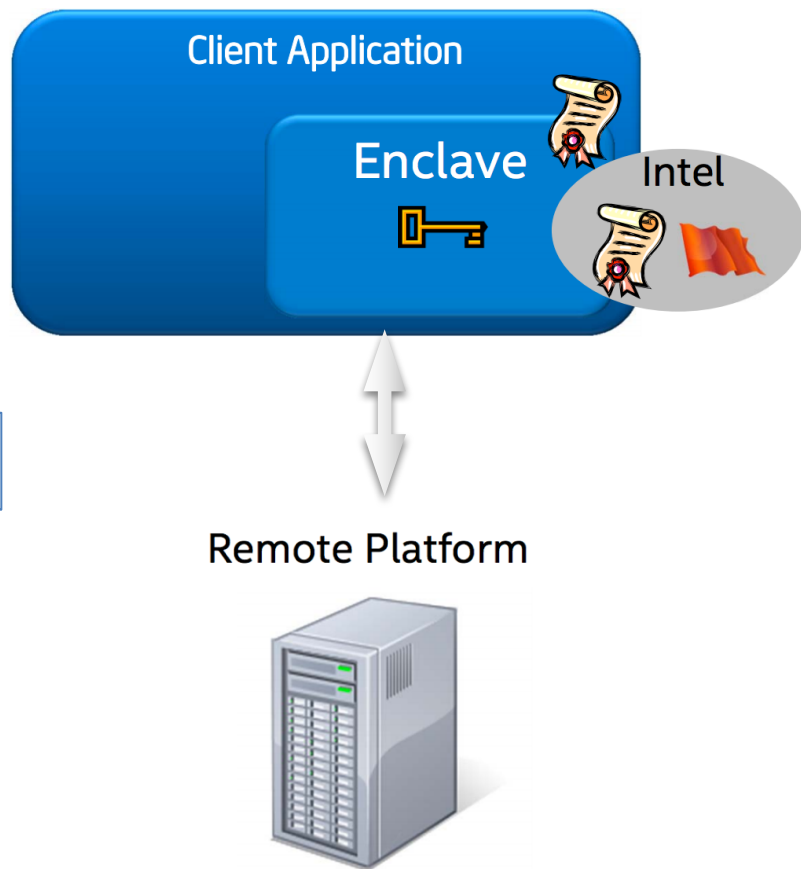
Nothing else is revealed about the individual genomes !!

Can we do this with Intel's SGX ?

Enclave Application



Remote Attestation



Enclave

Iron: Functional encryption and obfuscation using Intel SGX

Ben Fisch, Dhinakaran Vinayagamurthy,

Dan Boneh, Sergey Gorbunov

In proc. ACM CCS 2017

Functional Encryption

[Boneh-Sahai-Waters, 2011]

Master-key Authority



mpk

program P

functional key K_P

researcher



K_P

$c_1 \leftarrow E(\text{mpk}, v_1)$

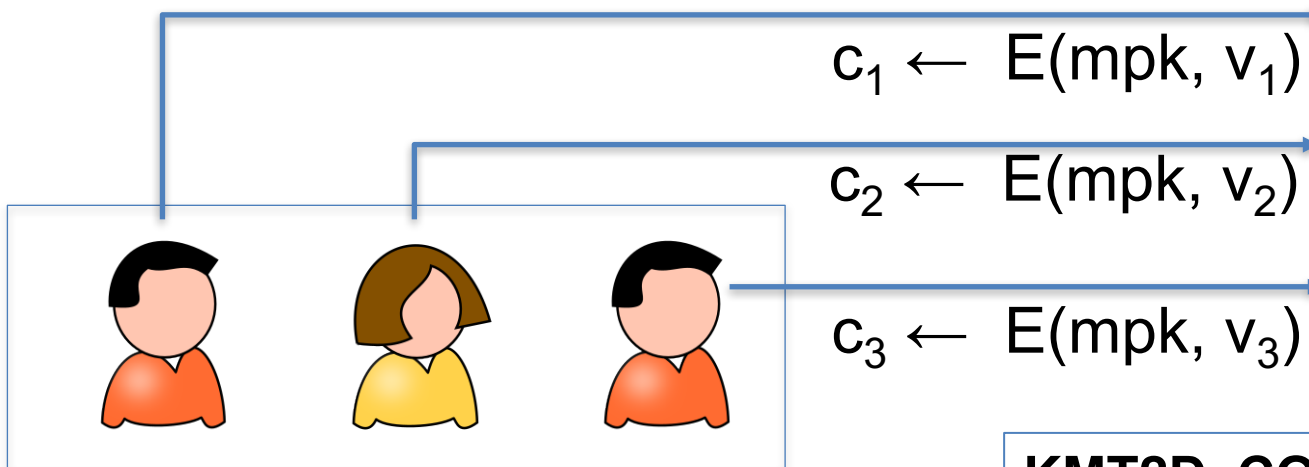
$c_2 \leftarrow E(\text{mpk}, v_2)$

$c_3 \leftarrow E(\text{mpk}, v_3)$

Decrypt

$P(v_1, v_2, v_3)$

KMT2D, COL6A1



Functional Encryption

[Boneh-Sahai-Waters, 2011]

Why is functional encryption hard?
no interaction during decryption
can't use MPC techniques

researcher



K_P

$c_1 \leftarrow E(\text{mpk}, v_1)$

$c_2 \leftarrow E(\text{mpk}, v_2)$

$c_3 \leftarrow E(\text{mpk}, v_3)$

Decrypt

Satisfy regulators?
(GDPR)

$P(v_1, v_2, v_3)$

But not so simple ...

- Enclave memory access pattern leaks and can break FE security
- How to represent the program P:
 - Cannot move code into enclave after EINIT
 - Difficult to safely implement interpreter in enclave: performance and memory access pattern leak
- Side channel attacks (timing, power)

Iron architecture

Key manager enclave:

manage master key

Decryption enclave:

initialized at startup

Function enclave: for specific program P.

if approved, signed by key manager

mpk

Key Manager Enclave

$(pk_{pke}, sk_{pke}) \leftarrow \text{PKE.Keygen}$
 $(vk_{\text{sign}}, sk_{\text{sign}}) \leftarrow \text{Sig.Keygen}$

sk_{pke} 

sk_{pke} 

Decryption Enclave

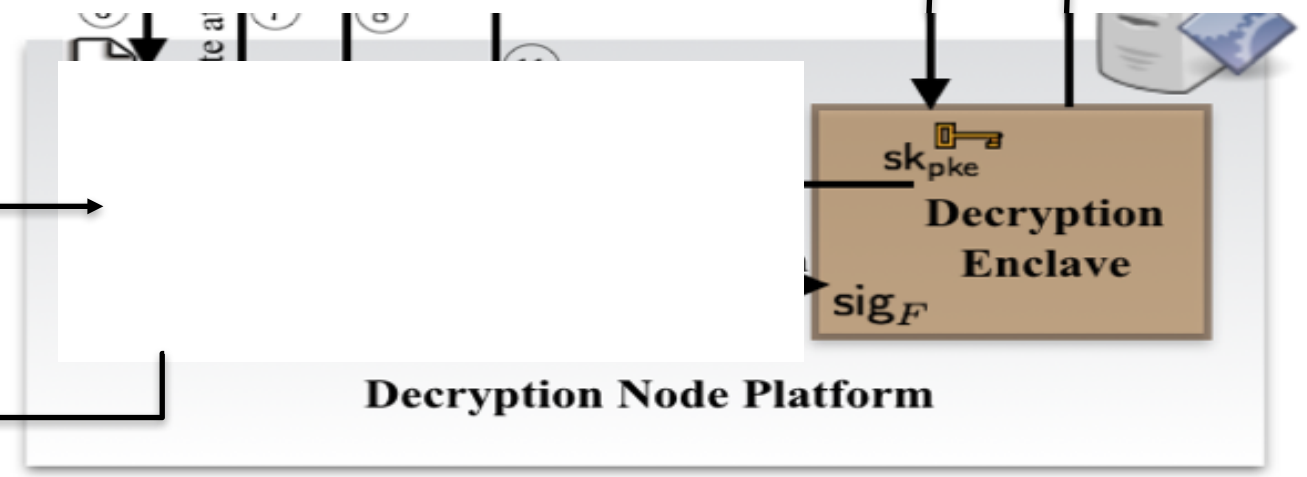
sig_F

Decryption Node Platform

$c_1 \leftarrow E(\text{mpk}, m_1)$

$c_2 \leftarrow E(\text{mpk}, m_2)$

$P(m_1, m_2)$



Security

- Formally model the SGX HW interface:
**Setup, Load, Run, Run&Report,
Run&Quote, ReportVerify, QuoteVerify**

Builds on HW security models of:

Pass et. al. [PST'17], Bahmani et. al. [BBB+'16]

- MIFE simulation-based security, assuming:
adversary cannot distinguish
black-box HW interface and real SGX

Side-channel attacks

Security proof does not capture side channel attacks on SGX


- Cache-timing attacks [CD16] leak memory access patterns at cache-line granularity
- Page-fault attacks [XCP15] leak memory access patterns at 4KB page granularity
- Branch shadowing attacks [LSG+16] can directly view branch history (saved for pipeline branch prediction)

DEFENSE: only sign function enclaves whose memory access pattern is independent of sensitive data (e.g. ORAM based)

Implementation and Evaluation

- C++ using the Intel(R) SGX SDK 1.6 for Windows
Intel Skylake i7-6700, 3.40 GHz, 8 GiB RAM,
Windows Server 2012 R2 Standard
- Function enclave implementation is *data-oblivious* to resist side-channels

Comparing Iron to cryptographic constructions



Private data aggregation

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

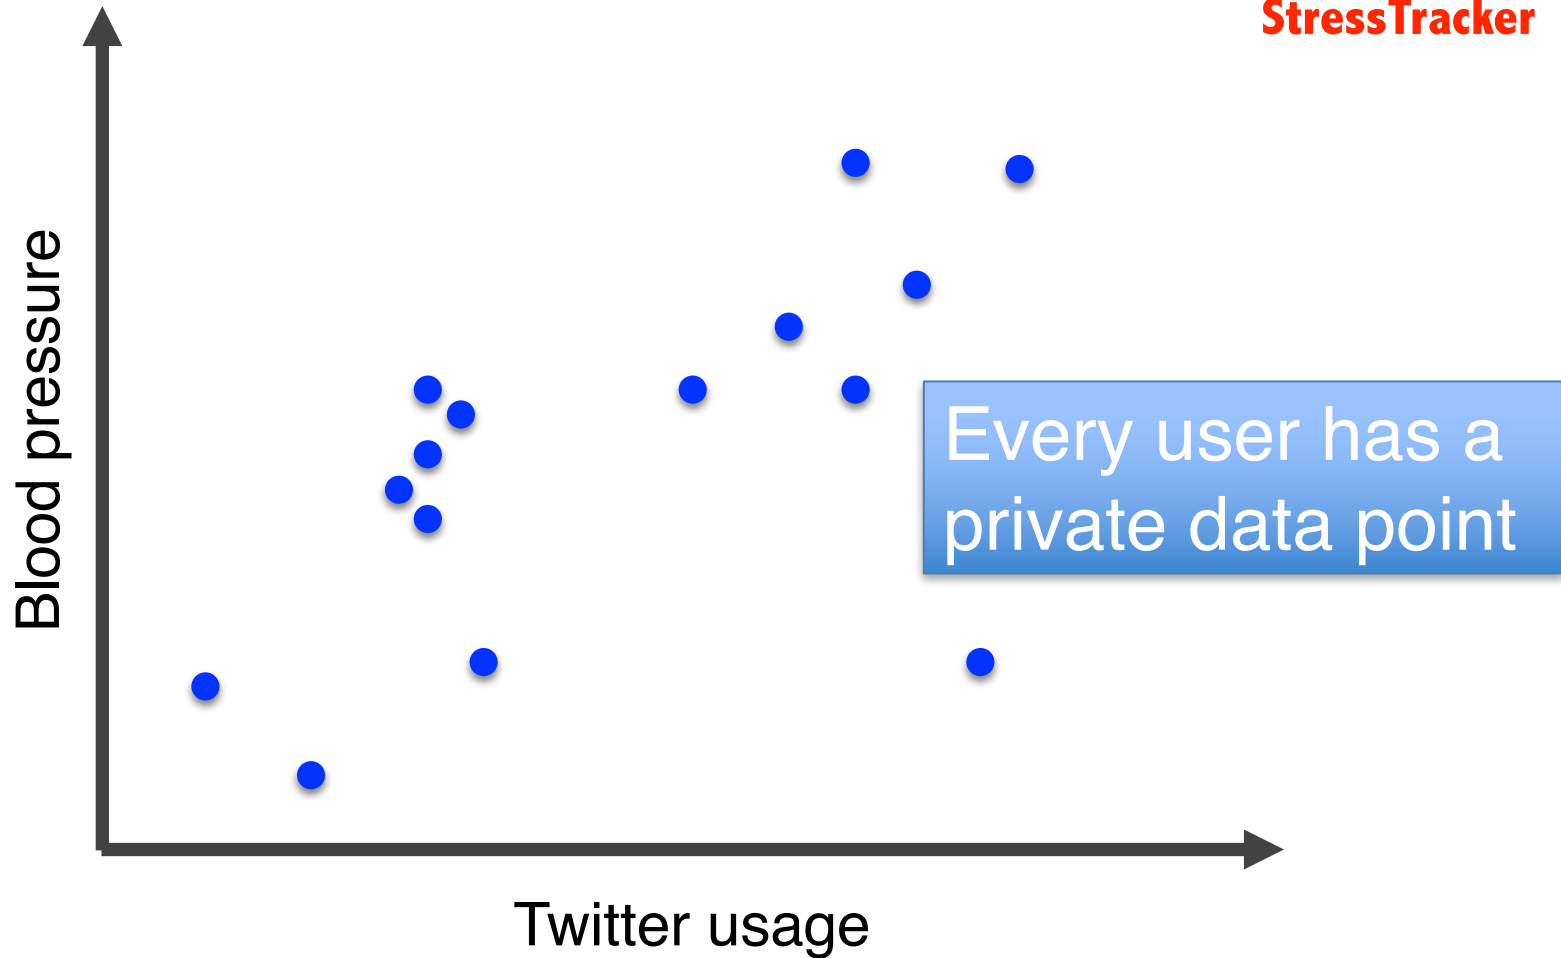
Joint work with Henry Corrigan-Gibbs

NSDI 2017

Today: Non-private aggregation



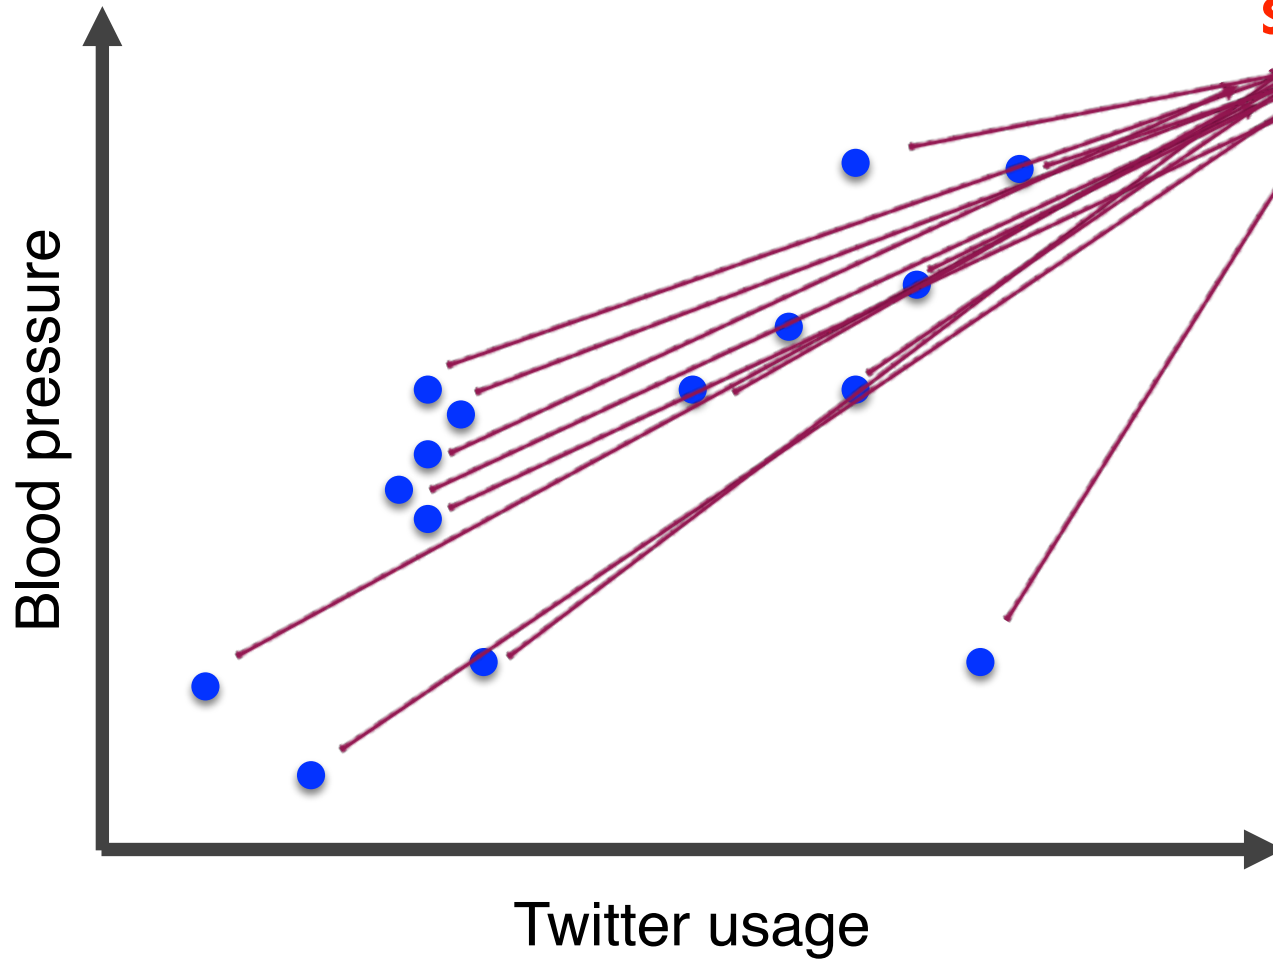
StressTracker



Today: Non-private aggregation



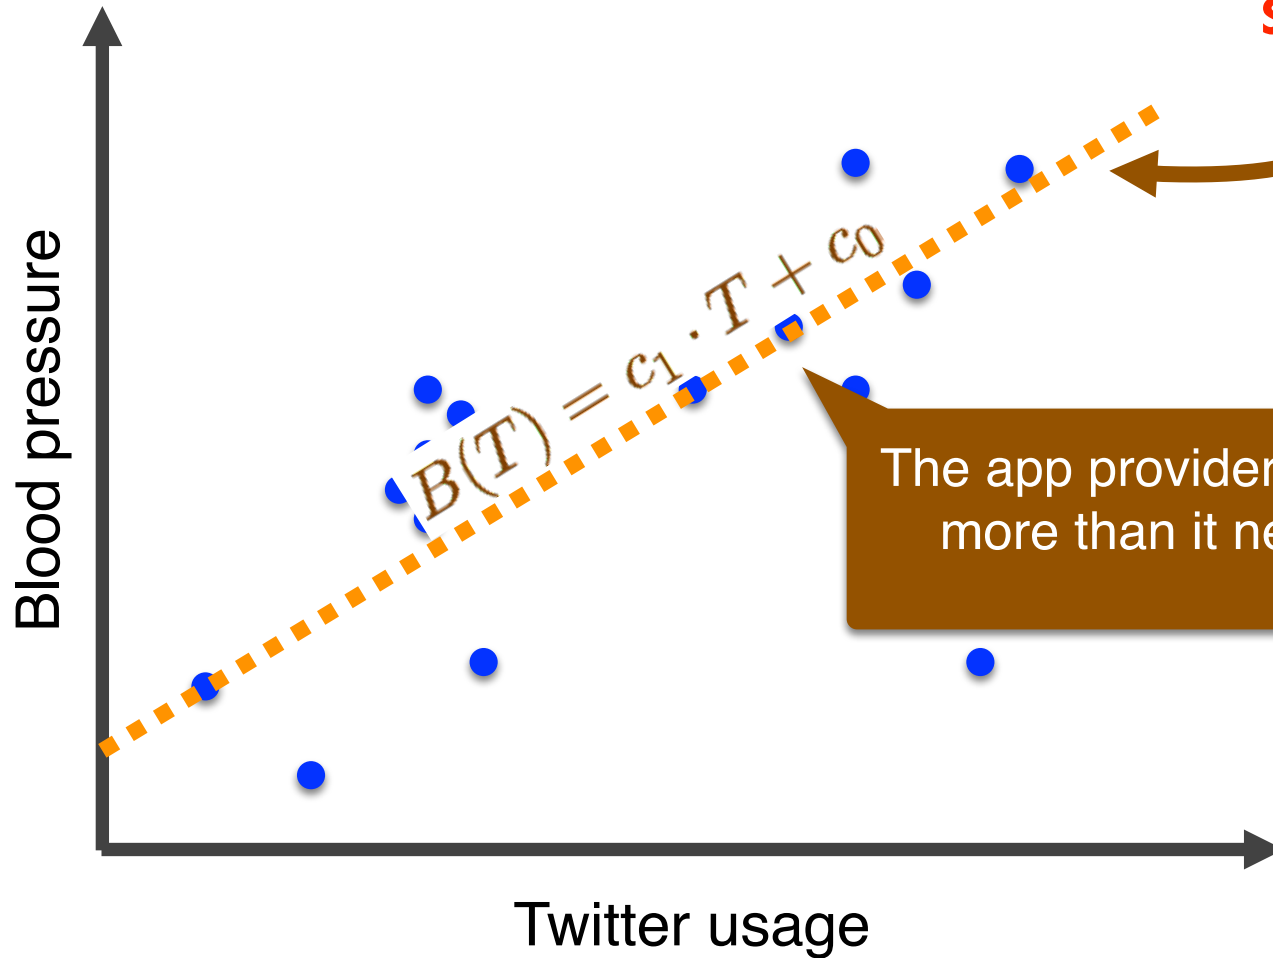
StressTracker



Today: Non-private aggregation



StressTracker



Prio: Private aggregation


Clients send one *share* of their data to each aggregator

Blood pressure

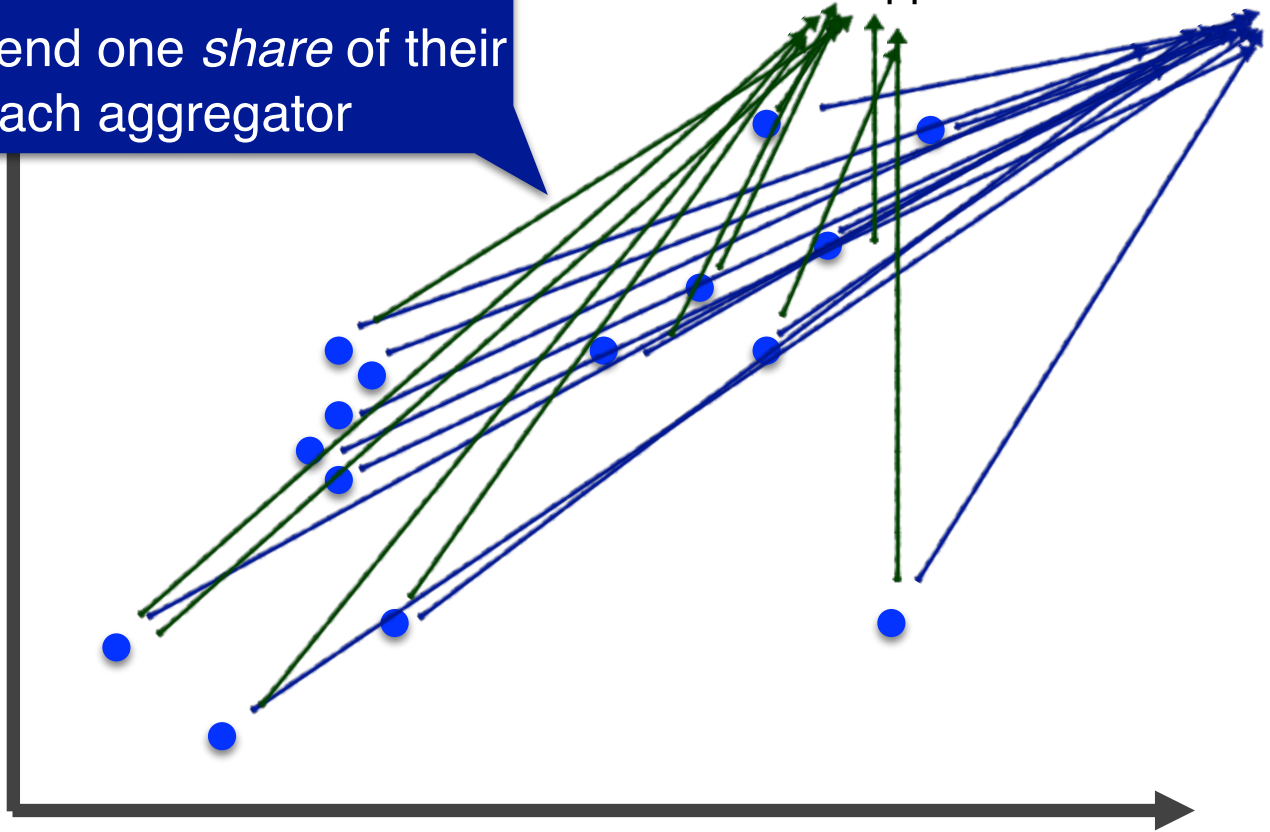
Twitter usage



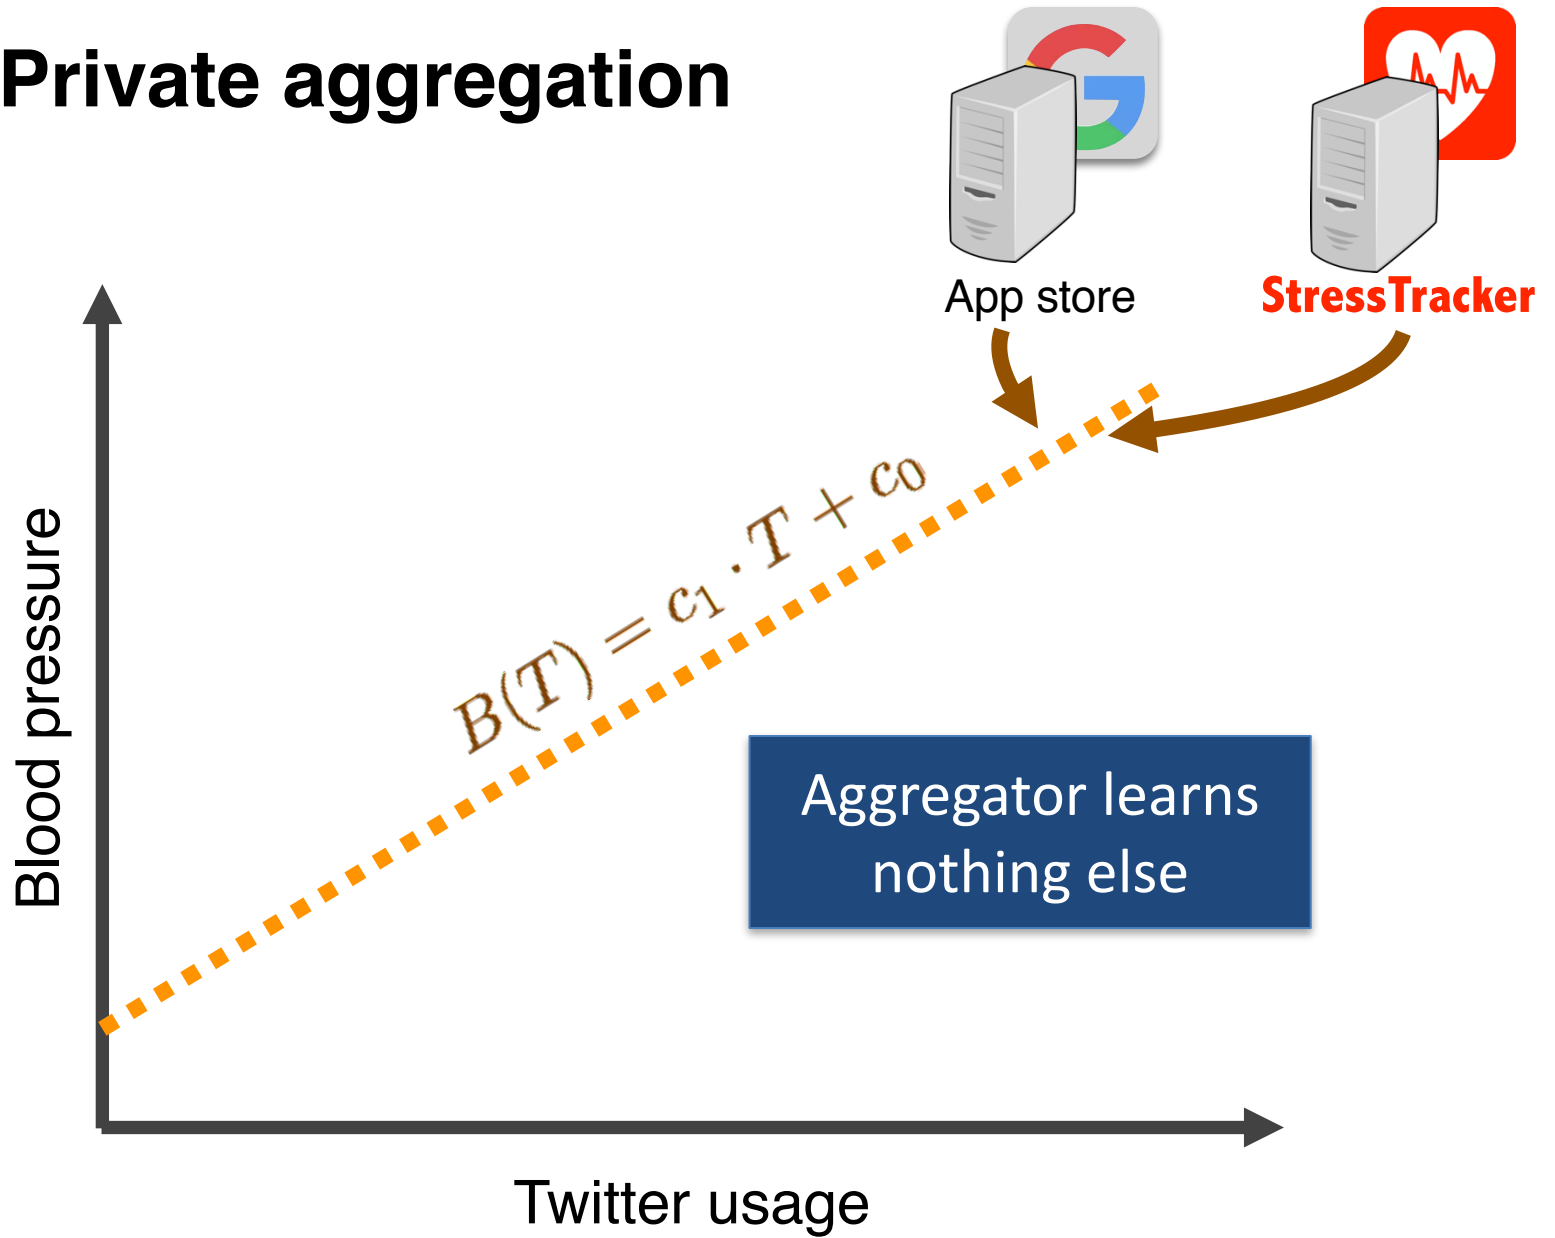
App store



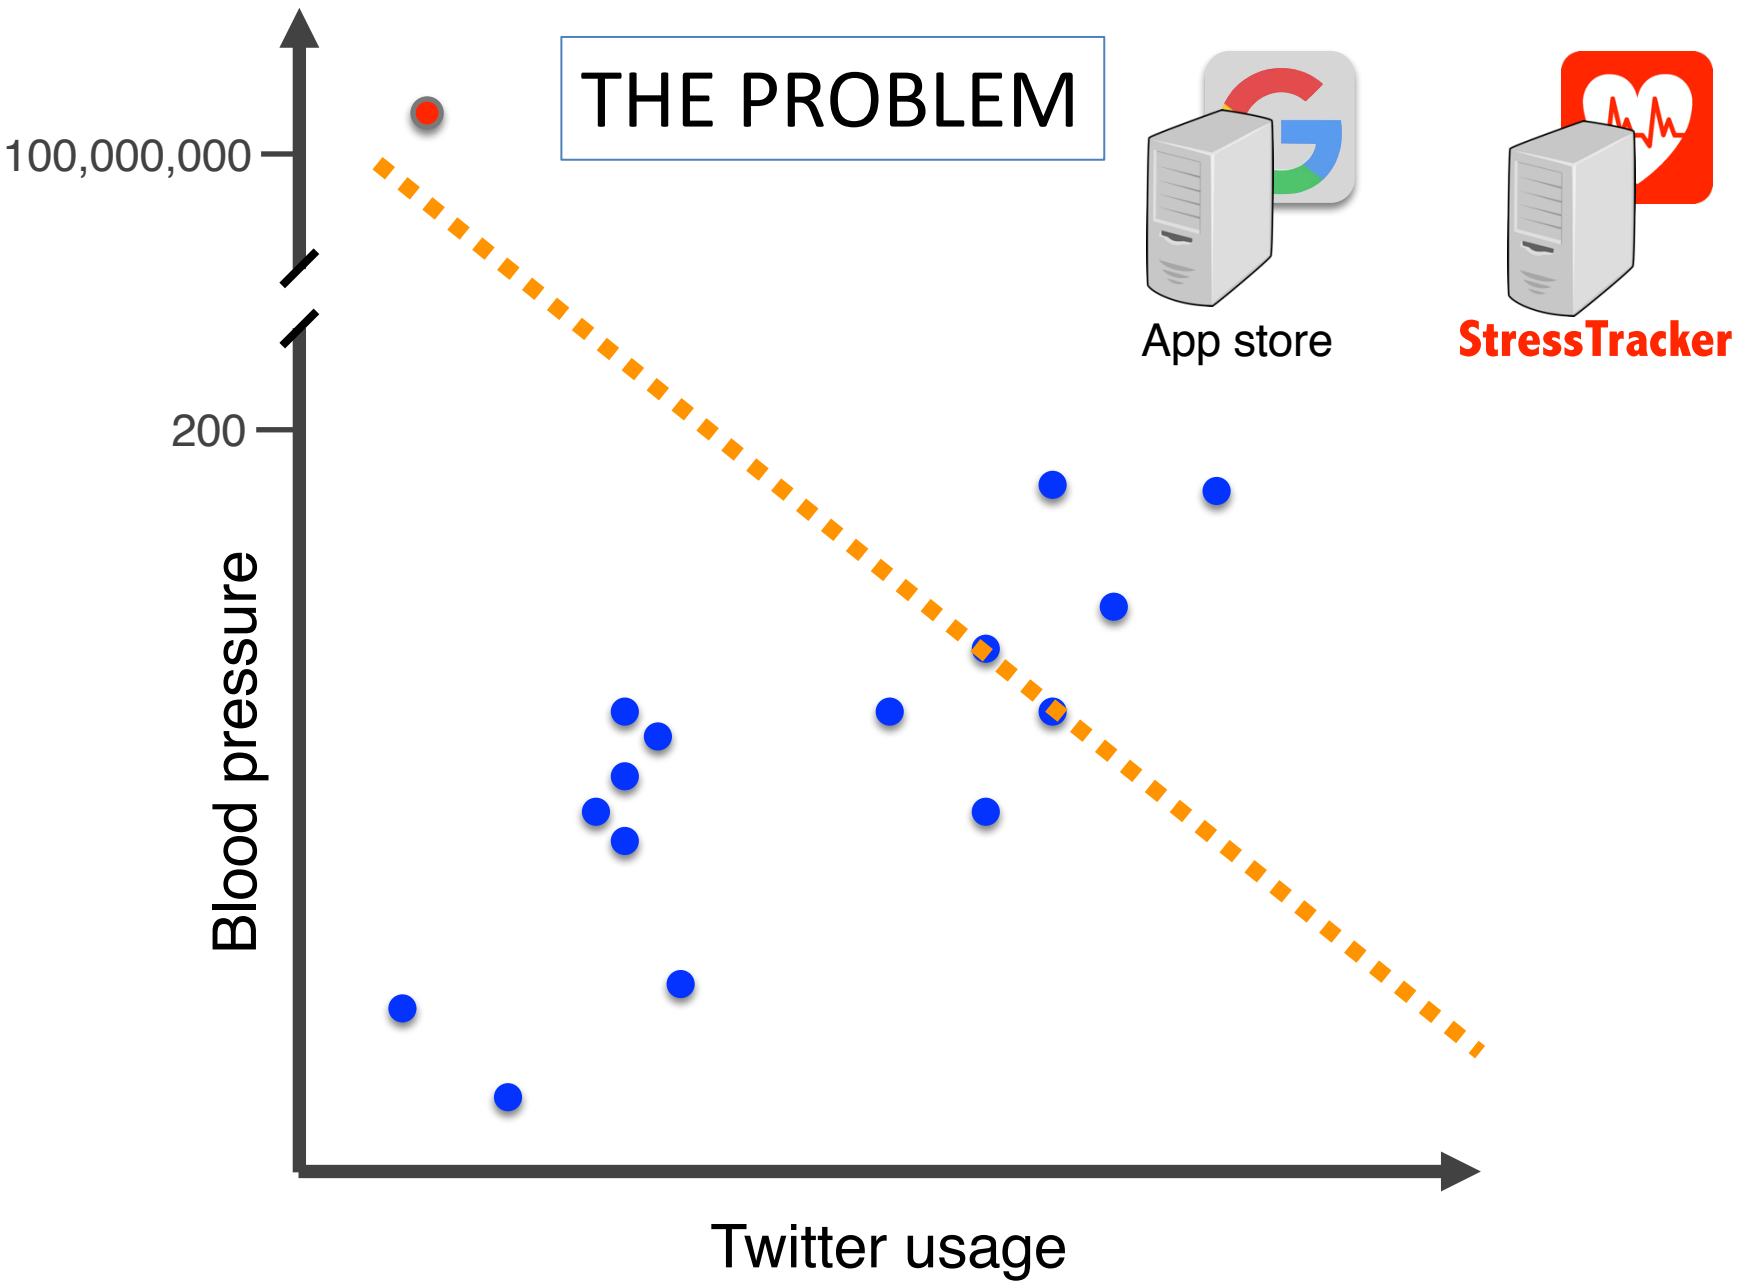
StressTracker



Prio: Private aggregation



THE PROBLEM



Private aggregation



Exact correctness: if all servers are honest they learn $f(x_1, \dots, x_n)$

Privacy: if one server is honest they learn only $f(x_1, \dots, x_n)$

Robustness: malicious clients have bounded influence

Scalable: no public-key crypto (other than TLS)

Prio contributions

Achieves all four goals

1. Robustness using secret-shared non-interactive proofs (SNIPs)

- Every client efficiently proves to servers that its submission is well formed
- Takes advantage of non-colluding servers (verifiers)

2. Aggregatable encodings

Compute sums privately \implies

compute $f(\cdot)$ privately for many f 's of interest

Existing approaches

- **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)

Private aggregation needed in many settings

Private client value (X_i)	Aggregate $f(X_1, \dots, X_N)$
Location data (phones/cars)	<ul style="list-style-type: none">• Number of devices in location L• Ten most popular locations• Locations with weakest signal strength
Web browsing history	<ul style="list-style-type: none">• Most common bug-triggering websites• Websites with TLS certificate errors
Health information	<ul style="list-style-type: none">• Min, max, avg, stddev heart rate• ML model relating BP to Twitter usage
Text messages	<ul style="list-style-type: none">• Min, max, average number per day• ML model relating time of day to emotion

Warm-up: Computing private sums

Every device i holds a value $x_i \in \mathbb{F}$

Think: integers modulo a prime p

Cloud wants to compute

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

without learning any users' private value $x_i \in \mathbb{F}$

Example: Privately measuring traffic congestion

$$x_i = \begin{cases} 1 & \text{if user } i \text{ is on Golden Gate Bridge} \\ 0 & \text{otherwise} \end{cases}$$



$x_1 + \dots + x_N$ gives number of users on bridge

Private sums: A “straw-man” scheme

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

Server A



Server B





Server C



Assume that at least one
server is honest.

Private sums: A “straw-man” scheme




Split into shares s.t.

$$x_1 = [x_1]_a + [x_1]_b + [x_1]_c \in \mathbb{F}$$

$[x]$ means
“additive share of x ”


Private sums: A “straw-man” scheme



$$x_1 \in \mathbb{F}$$




Private sums: A “straw-man” scheme



$$x_1 \in \mathbb{F}$$



Private sums: A “straw-man” scheme

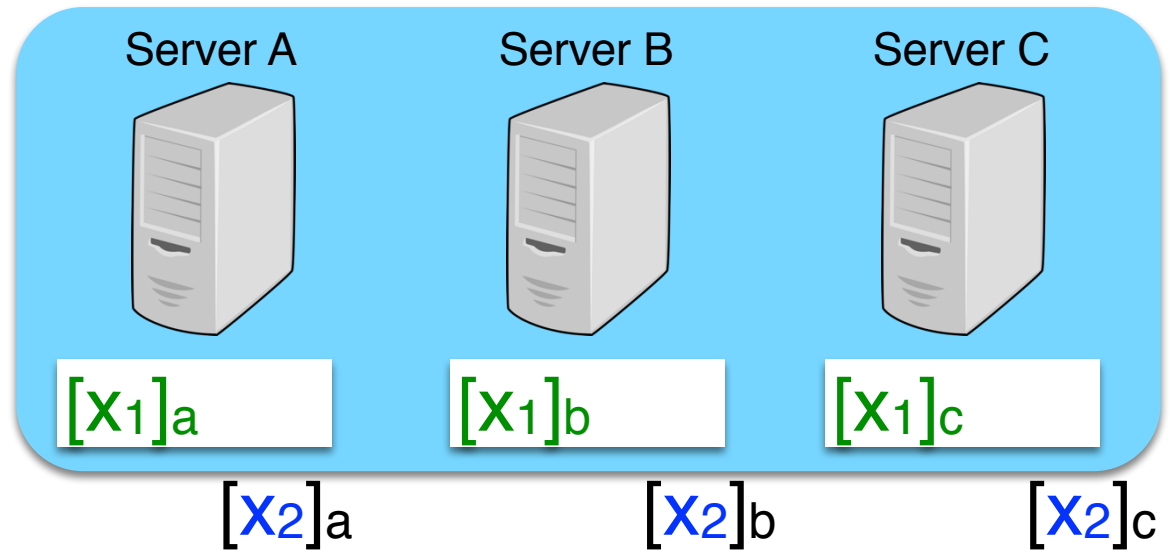


$[X_2]_a$ $[X_2]_b$ $[X_2]_c$

$X_2 \in \mathbb{F}$



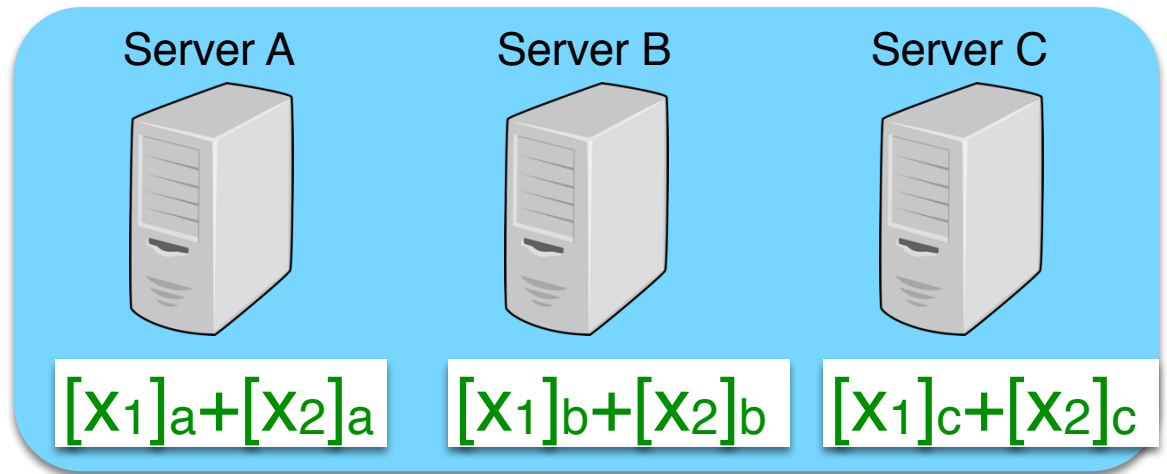
Private sums: A “straw-man” scheme



$X_2 \in \mathbb{F}$




Private sums: A “straw-man” scheme




$X_2 \in \mathbb{F}$



Private sums: A “straw-man” scheme



Private sums: A “straw-man” scheme



$$S_A + S_B + S_C = [X_1]_a + [X_1]_b + [X_1]_c + \dots \in \mathbb{F}$$
$$= X_1 + X_2 + \dots + X_N$$



Learn that three phones are on the
Bridge—but not which three

Strawman computing private sums


Correctness: if everyone follows the protocol, servers compute the sum of all x_i s.

Privacy: any proper subset of the servers can simulate everything given
(a) the public parameters, and
(b) the sum of the x_i s.

Scalability: by inspection.

Robustness: ???

Private sums: A “straw-man” scheme




x_3 is supposed to
be a 0/1 value

\mathbb{F}

$x_3 \in \mathbb{F}$



Private sums: A “straw-man” scheme




An evil client needn't
follow the rules!


$$[r]_a \quad [r]_b \quad [r]_c \quad \leftarrow \mathbb{F}$$



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



Solution: SNIP Proofs



$[x]_a$ $[x]_b$ $[x]_c$

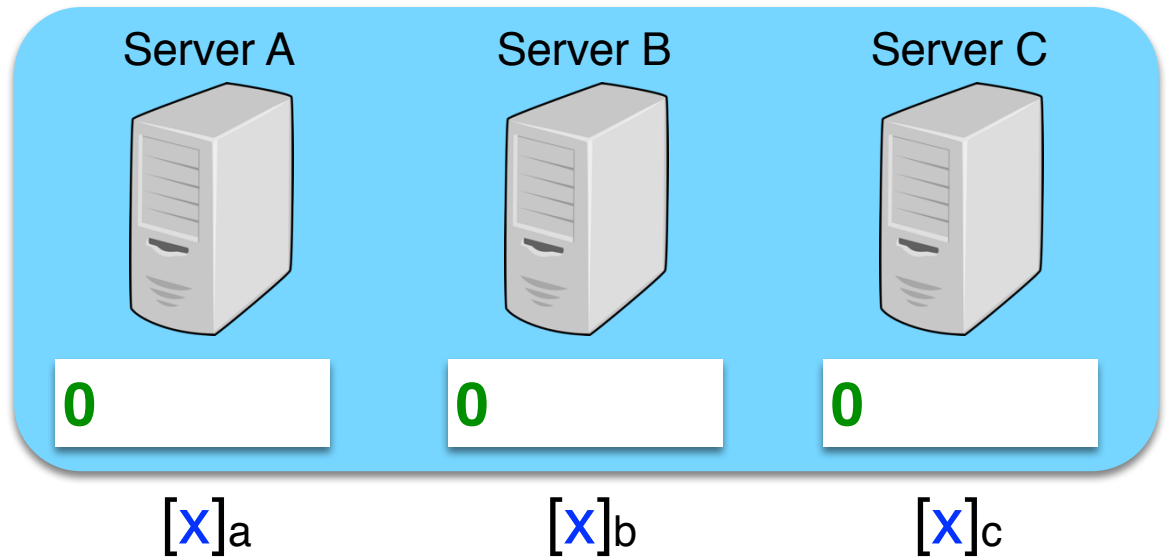
$x \in \mathbb{F}$

An arrow points from the text $x \in \mathbb{F}$ towards the three server icons above.



x is supposed to
be a 0/1 value

Solution: SNIP Proofs



$x \in \mathbb{F}$

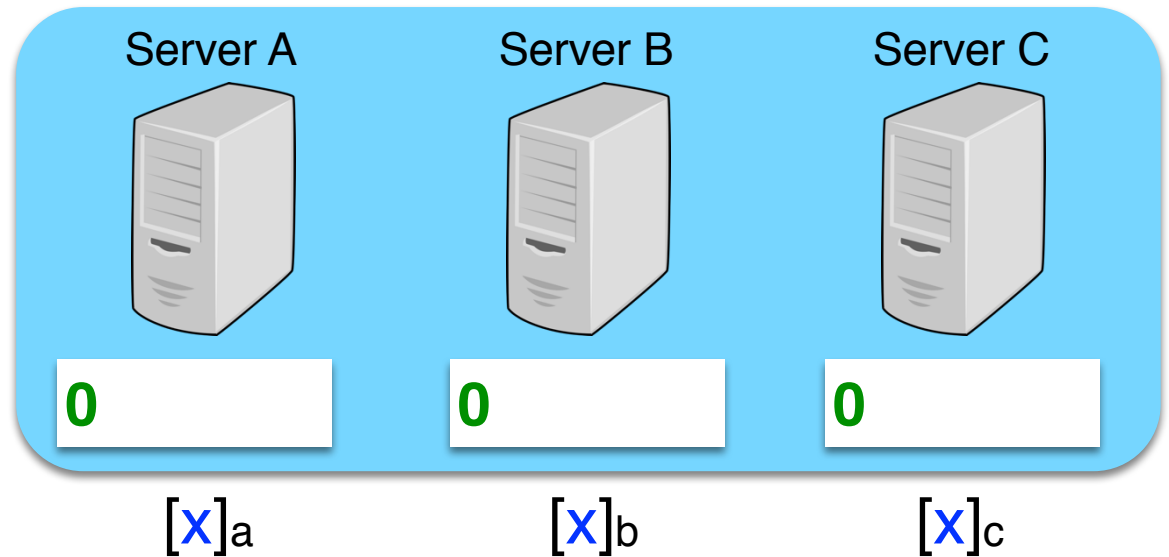


Without learning x ,
the servers want to ensure that:

$$[x]_a + [x]_b + [x]_c \in \{0, 1\}$$

Remember: these are
big integers mod p

Solution: SNIP Proofs



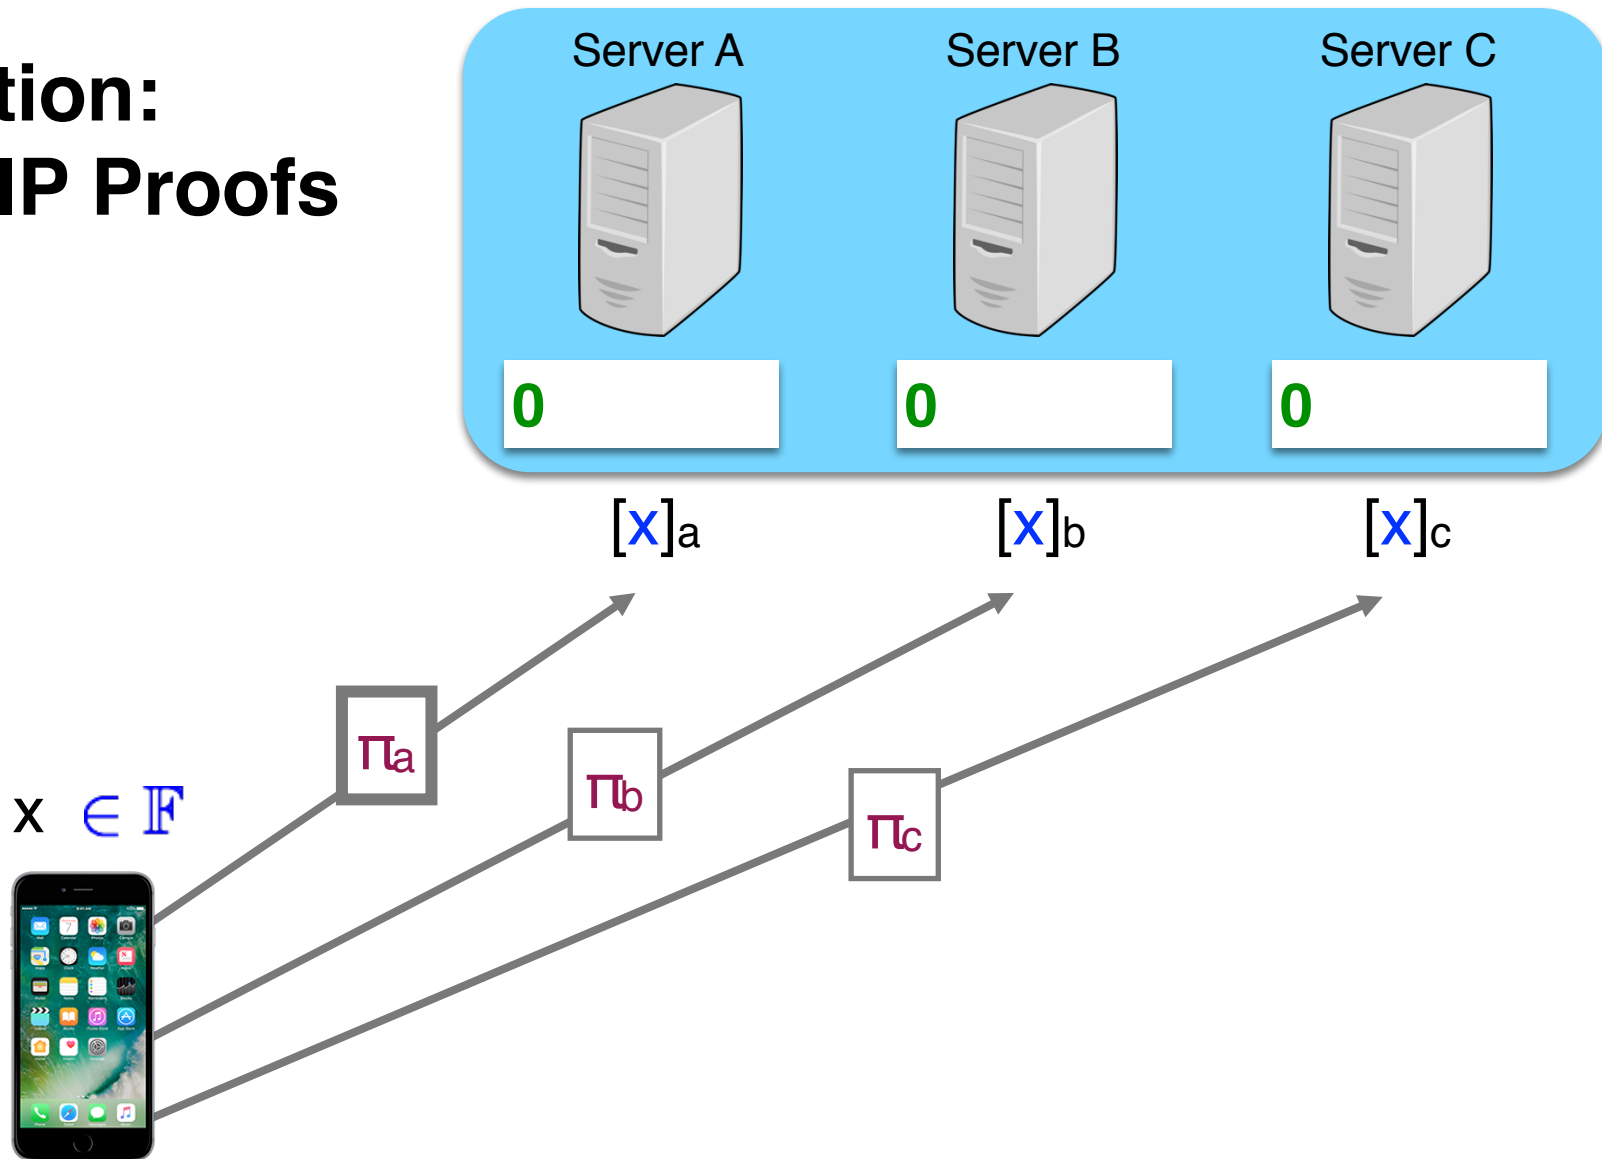
$x \in \mathbb{F}$




- Servers hold shares of x and a public predicate $\text{Valid}(\cdot)$
- Servers want to test if “ $\text{Valid}(x) = 0$ ” without leaking anything else about x
- The Valid predicate can be an arbitrary circuit:

$$\text{Valid}(x_1, x_2) = “3 < x_1 < 19 \text{ and } x_2 \in \{0, 1, 2\}”$$

Solution: SNIP Proofs



Solution: SNIP Proofs




$x \in \mathbb{F}$



Prio servers detect and reject malformed client submissions

\Rightarrow a client can influence aggregates by at most ± 1



Security goals for SNIPs

Completeness: Honest client convinces honest servers

Soundness: Dishonest client almost never convinces honest servers

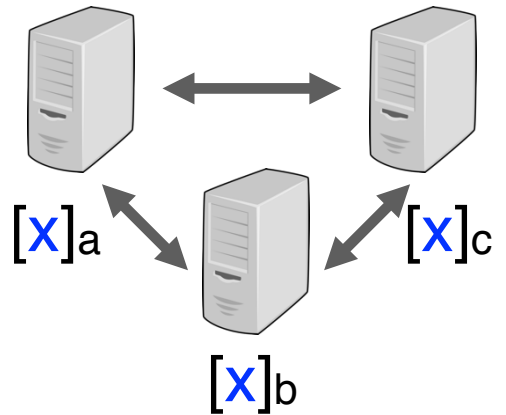
Zero-knowledge: Any proper subset of malicious servers learns nothing about x , except that x is valid

Client



A "valid" x

Servers



π_a, π_b, π_c

Existing techniques

Full blown MPC

Commitments + NIZKs

Commitments + SNARKs

Func. secret sharing [BGJ'1

SNIP

Limitations

Heavy setup and comm.

High server work


Info. theoretic techniques
 \Rightarrow little comp. overhead
 $O(1)$ server-to-server comm.

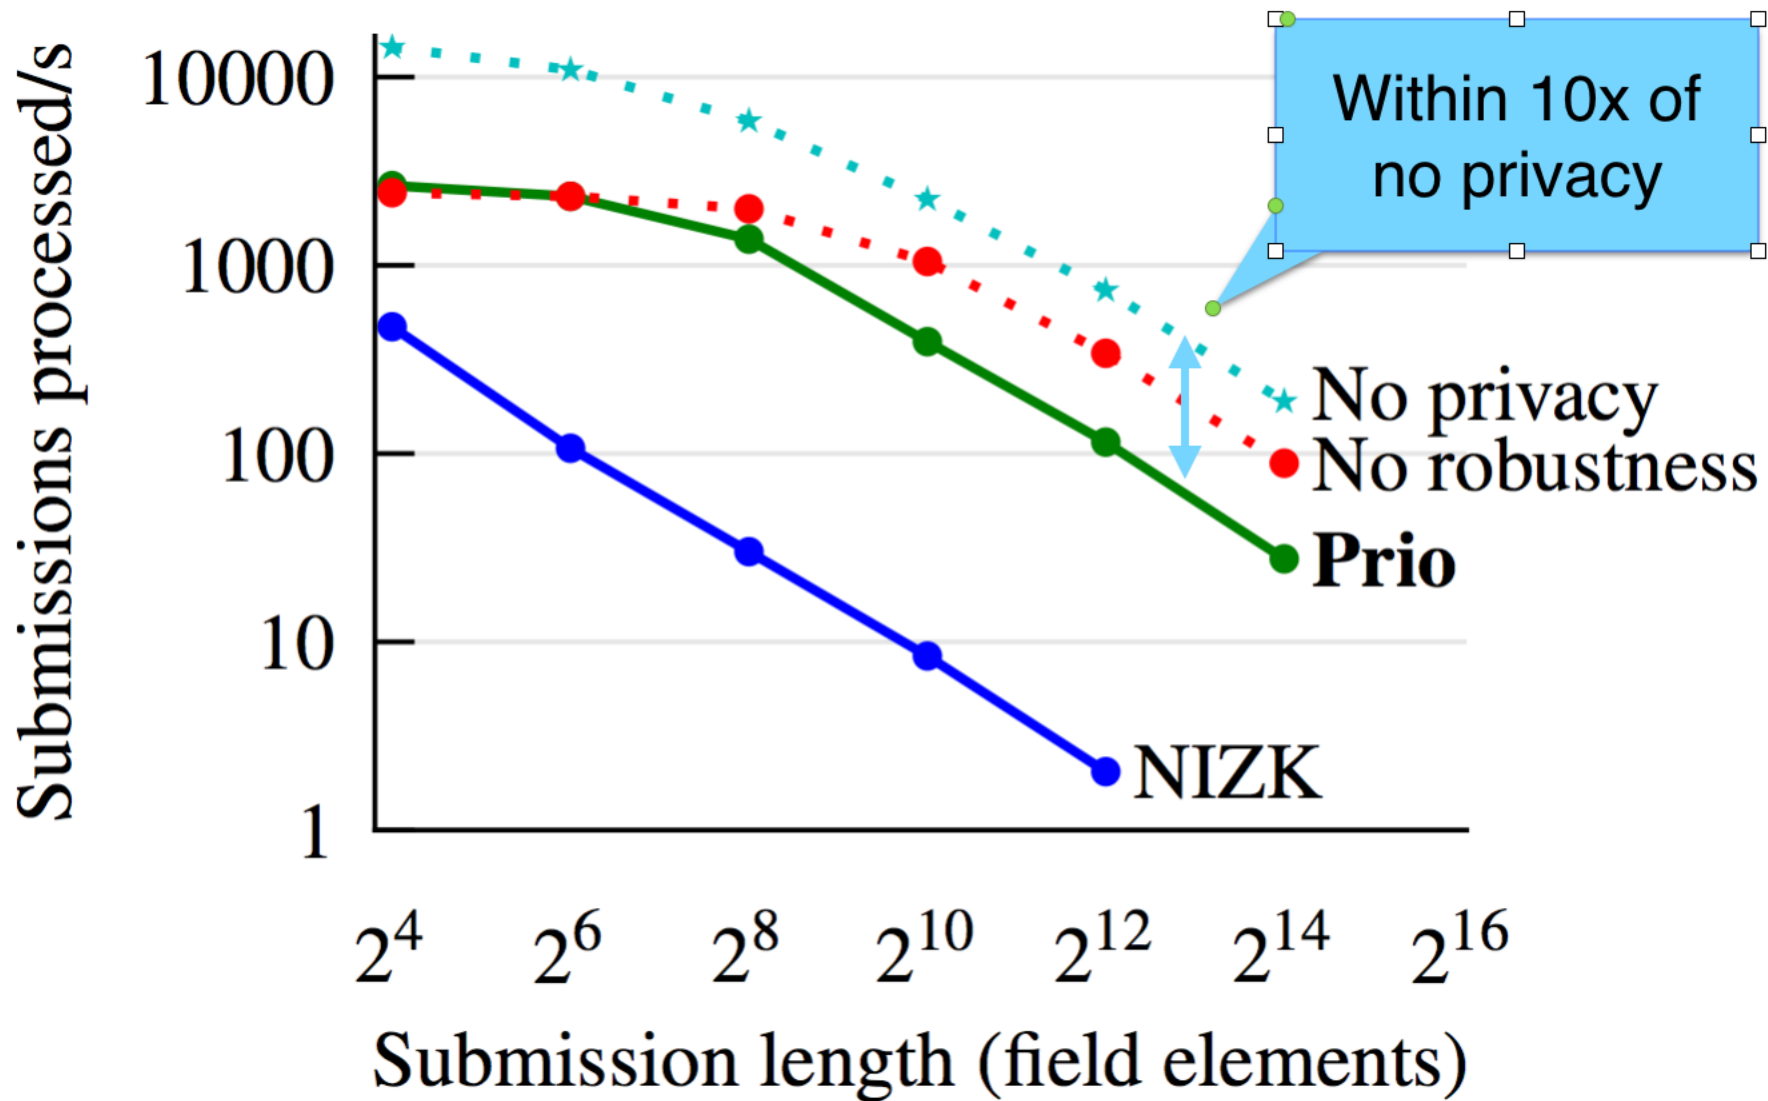
$|\pi_a|$ is linear in circuit size

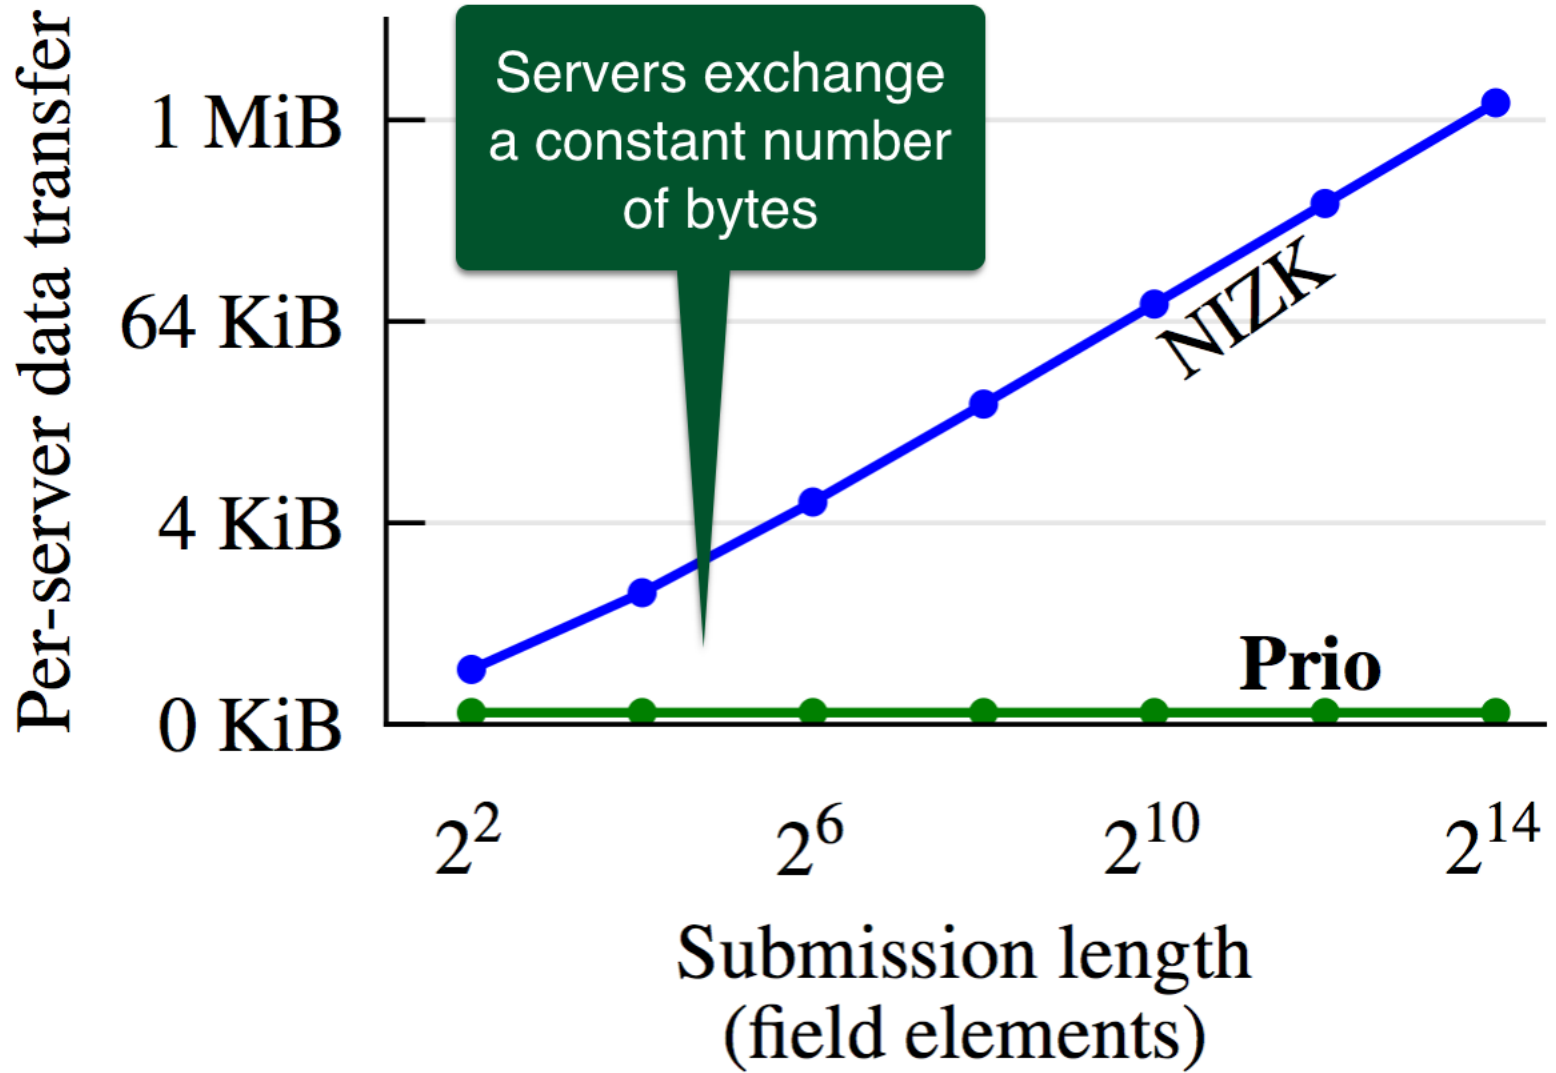
SNIPs: How?

- Step 1: reduce verifying circuit to verifying a single multiplication
- Step 2: Use “Beaver triple” supplied by client to verify the multiplication
- Step 3: Inject additional entropy to defend against malicious servers (similar to AMD codes)

Five-server cluster in five Amazon data centers







Complex statistics

Computing private sums \Rightarrow

can compute many other interesting aggregates

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

- Average
- Variance
- Standard deviation
- Most popular value (approx) – small universe
- “Heavy hitters” (approx)

... and even more statistics

Prio can aggregate a richer class of statistics:

- Approximate **min** and **max**
- Most popular value in a large universe
- Quality of arbitrary machine learning model (R^2)
- **Least-squares regression**


Prio supports a rich set of aggregation functions

Some limitations: cannot compute exact max

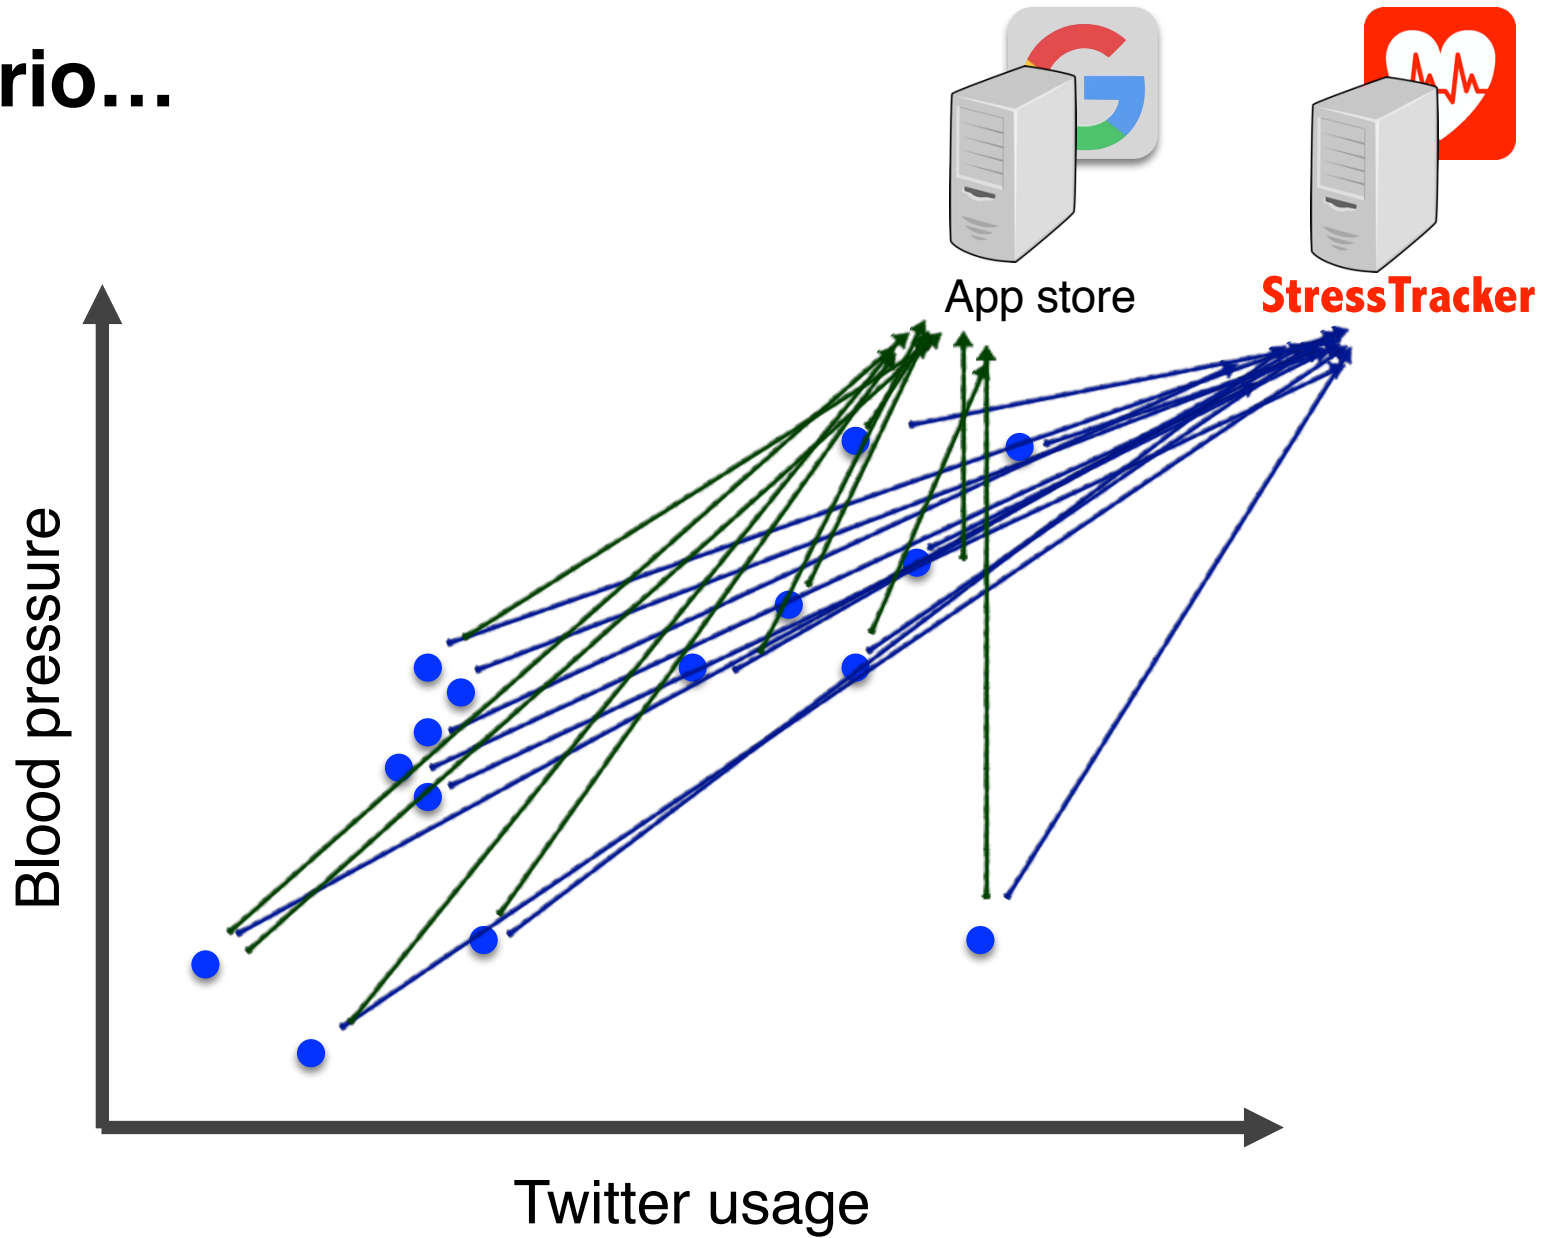
Putting it all together: Today



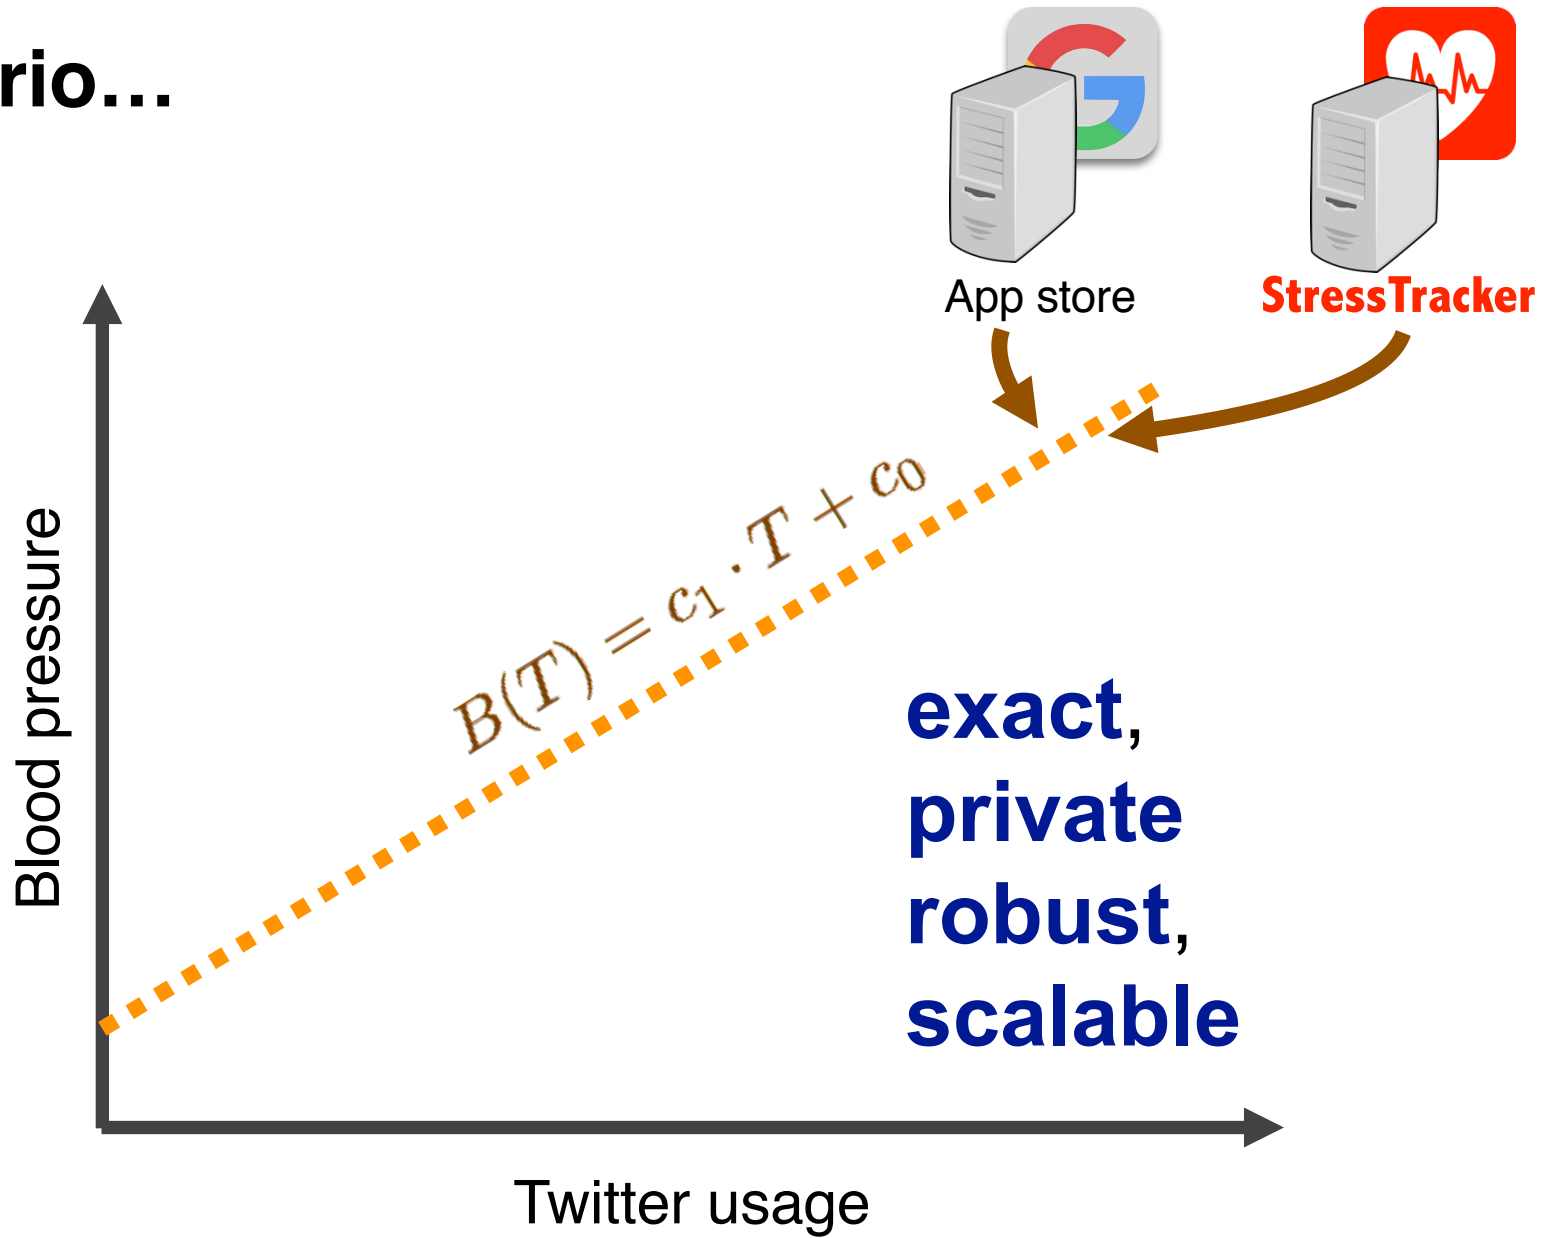
StressTracker



With Prio...



With Prio...



THE END