Express: Lowering the Cost of Metadata-Hiding Communication with Cryptographic Privacy

Saba Eskandarian, Stanford
Henry Corrigan-Gibbs, MIT CSAIL
Matei Zaharia, Stanford
Dan Boneh, Stanford

How Can We Protect Whistleblowers?

Private Communication Tools

End to end encrypted messaging apps

E.g. Signal, WhatsApp

Problem: **metadata**

Private Communication Tools

End to end encrypted messaging apps

E.g. Signal, WhatsApp

Problem: **metadata**

Anonymizing proxy

E.g. Tor, SecureDrop

Problem: **global adversaries**

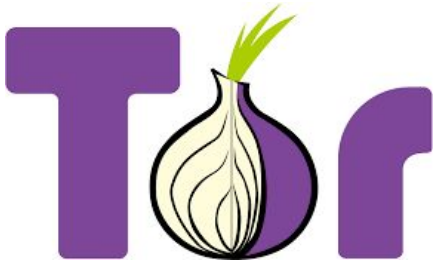

Private Communication Tools

Metadata-hiding communication systems with cryptographic privacy

Private Communication Tools

Metadata-hiding communication systems with cryptographic privacy

Drawback: **heavy requirements placed on clients**

- Requirement to run in synchronized rounds
- High communication costs

Private Communication Tools

Metadata-hiding communication systems with cryptographic privacy

Drawback: **heavy requirements placed on clients**

- Requirement to run in synchronized rounds
- High communication costs

Fundamental issue: whistleblowing tools need *cover traffic*, which must be possible for clients to generate at minimal cost

Private Communication Tools

Metadata-hiding communication systems with cryptographic privacy

Drawback: **heavy requirements placed on clients**

- Requirement to run in synchronized rounds
- High communication costs

Fundamental issue: whistleblowing tools need *cover traffic*, which must be possible for clients to generate at minimal cost

Can we get around high client costs?

Express: Practical Metadata-Hiding Whistleblowing

Qualitative improvement: users do not access the system in synchronized rounds

Asymptotic improvements:

Client computation: O(1)

Communication: O(1)

Prior work: O(√N)

Practical improvements:

6x faster server

8x faster client

>10x communication reduction

6x reduction in dollar cost to run

Express Overview

2 server system, secure against:

- Up to one corrupt server
- Arbitrarily many corrupt users

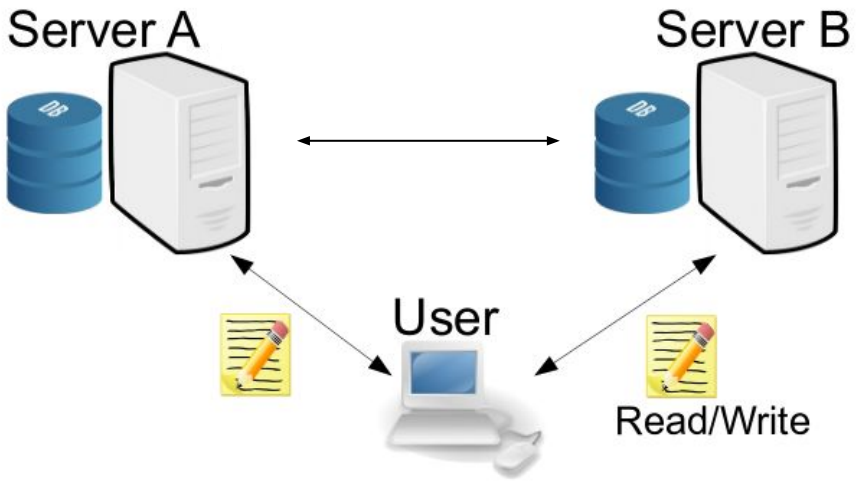

Express Overview

2 server system, secure against:

- Up to one corrupt server
- Arbitrarily many corrupt users

Supported operations:

- Register mailbox
- (Private) write to mailbox
- Read from mailbox

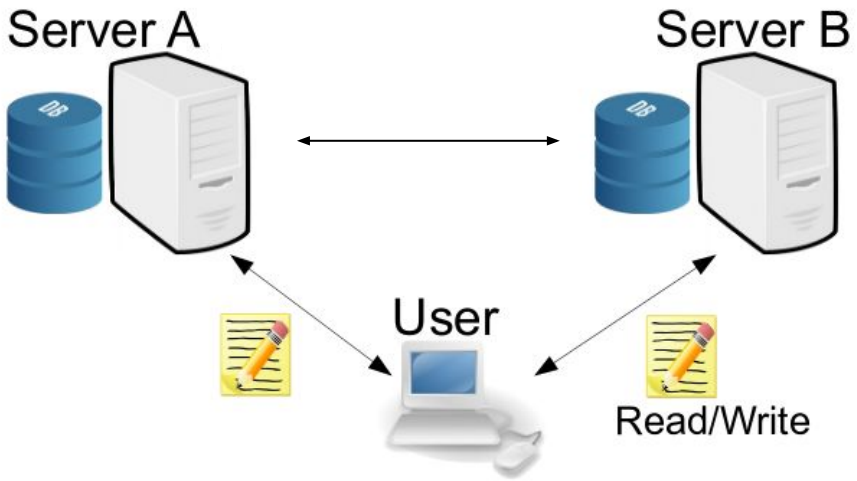

Express Overview

2 server system, secure against:

- Up to one corrupt server
- Arbitrarily many corrupt users

Supported operations:

- Register mailbox
- (Private) write to mailbox
- Read from mailbox

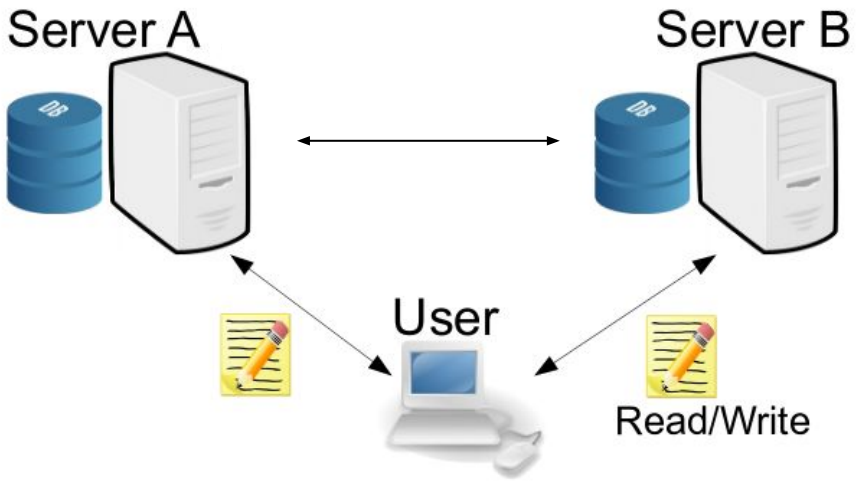

Security: can't tell who the *recipient* of a message is

Tool: Private Writing

Addr	Data
0	0
1	0
2	0
3	0
4	0

Addr	Data
0	0
1	0
2	0
3	0
4	0

Distributed Point Functions and their Applications, Niv Gilboa, Yuval Ishai, *Eurocrypt’14.*
Private Information Storage, Rafail Ostrovsky, Victor Shoup, *STOC’97*

Tool: Private Writing

x	f(x)
0	0
1	0
2	0
3	“Hi!”
4	0

Addr	Data
0	0
1	0
2	0
3	0
4	0

Addr	Data
0	0
1	0
2	0
3	0
4	0

Distributed Point Functions and their Applications, Niv Gilboa, Yuval Ishai, *Eurocrypt’14.*
Private Information Storage, Rafail Ostrovsky, Victor Shoup, *STOC’97*

Tool: Private Writing

x	$f_1(x)$
0	“abc”
1	“xf$”
2	“^tg”
3	“!7≈”
4	“jhV”

x	$f_2(x)$
0	“abc”
1	“xf$”
2	“^tg”
3	“‘2!)”
4	“jhV”

Addr	Data
0	0
1	0
2	0
3	0
4	0

Addr	Data
0	0
1	0
2	0
3	0
4	0

Distributed Point Functions and their Applications, Niv Gilboa, Yuval Ishai, *Eurocrypt’14.*
Private Information Storage, Rafail Ostrovsky, Victor Shoup, *STOC’97*

Tool: Private Writing

f_1

f_2

Addr	Data
0	0
1	0
2	0
3	0
4	0

Addr	Data
0	0
1	0
2	0
3	0
4	0

Distributed Point Functions and their Applications, Niv Gilboa, Yuval Ishai, *Eurocrypt'14.*
Private Information Storage, Rafail Ostrovsky, Victor Shoup, *STOC'97*

Tool: Private Writing

f_1

f_2

Addr	Data
0	0
1	0
2	0
3	0
4	0

Addr	Data
0	0
1	0
2	0
3	0
4	0

Distributed point function (DPF): f_1, f_2 have size O(logN), generated in time O(logN), where N = # of mailboxes

Distributed Point Functions and their Applications, Niv Gilboa, Yuval Ishai, *Eurocrypt'14.*
Private Information Storage, Rafail Ostrovsky, Victor Shoup, *STOC'97*

Tool: Private Writing

f_1

f_2

Addr	Data
0	“abc”
1	“xf$”
2	“^tg”
3	“!7≈”
4	“jhV”

Addr	Data
0	“abc”
1	“xf$”
2	“^tg”
3	“‘2!)”
4	“jhV”

Distributed Point Functions and their Applications, Niv Gilboa, Yuval Ishai, *Eurocrypt’14.*
Private Information Storage, Rafail Ostrovsky, Victor Shoup, *STOC’97*

Tool: Private Writing

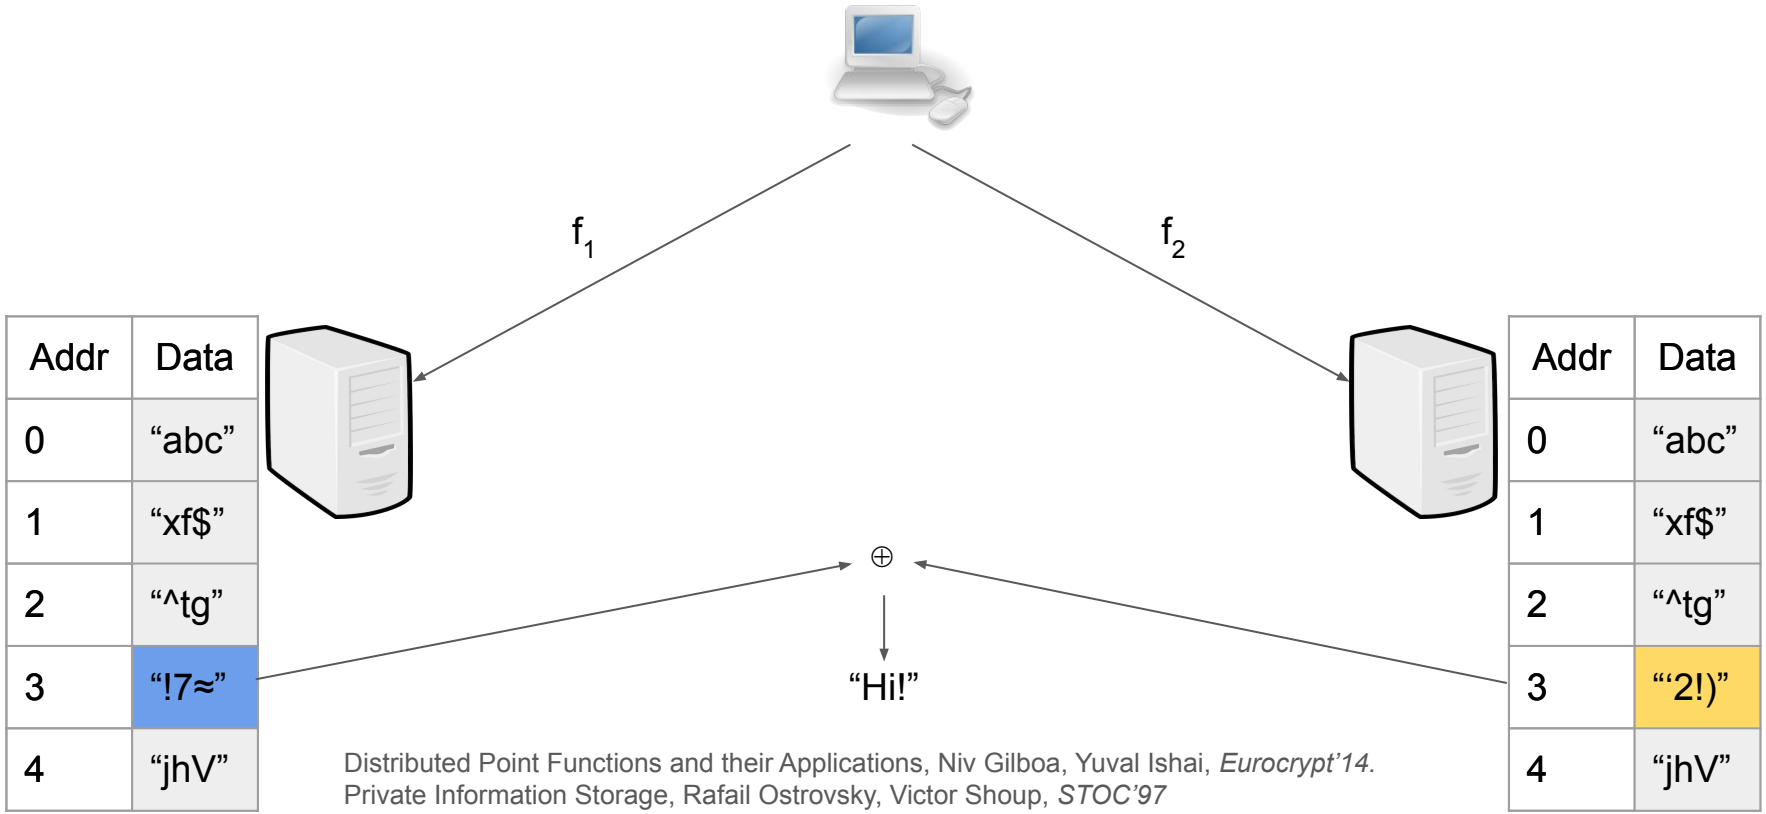

Distributed Point Functions and their Applications, Niv Gilboa, Yuval Ishai, *Eurocrypt'14.*
Private Information Storage, Rafail Ostrovsky, Victor Shoup, *STOC'97*

Filtering out Malformed DPFs

Problem: disruptive user sends malformed message to corrupt mailboxes

x	f(x)
0	989f4
1	dDf73
2	08dji3
...	...
N	89hfif

Filtering out Malformed DPFs

Problem: disruptive user sends malformed message to corrupt mailboxes

Solution: servers blindly *audit* all incoming write requests

Filtering out Malformed DPFs

Problem: disruptive user sends malformed message to corrupt mailboxes

Solution: servers blindly *audit* all incoming write requests

Prior work: third server audits requests

- O(√N) communication
- O(√N) client/auditor computation

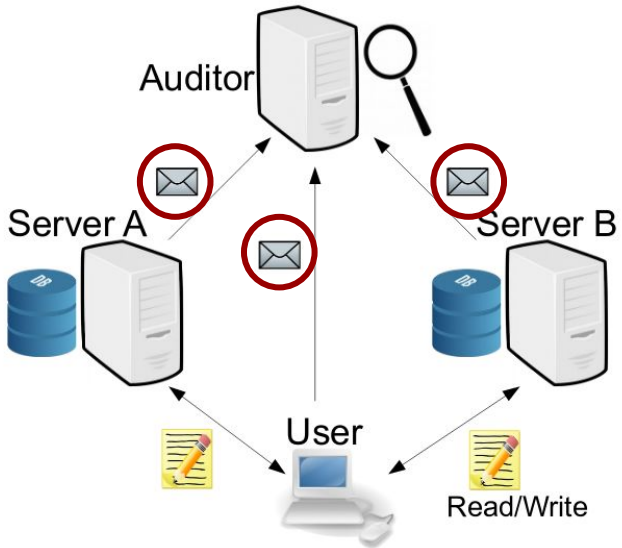

Riposte: An Anonymous Messaging System Handling Millions of Users, Henry Corrigan-Gibbs, Dan Boneh, David Mazieres, *Oakland'15.*

Filtering out Malformed DPFs

Problem: disruptive user sends malformed message to corrupt mailboxes

Solution: servers blindly *audit* all incoming write requests

<u>New auditing protocol:</u>

- O(1) communication
- O(1) client computation
- No additional server!

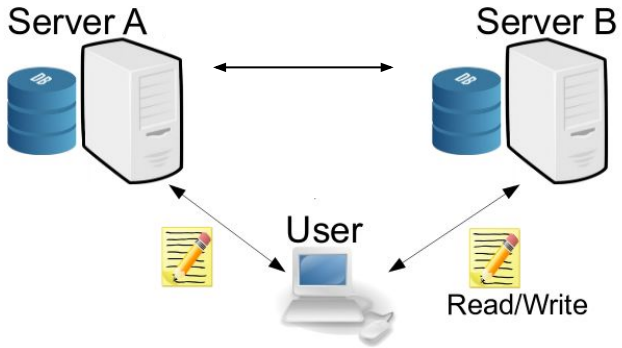

Auditing

Goal: check that values held by servers only differ at one point

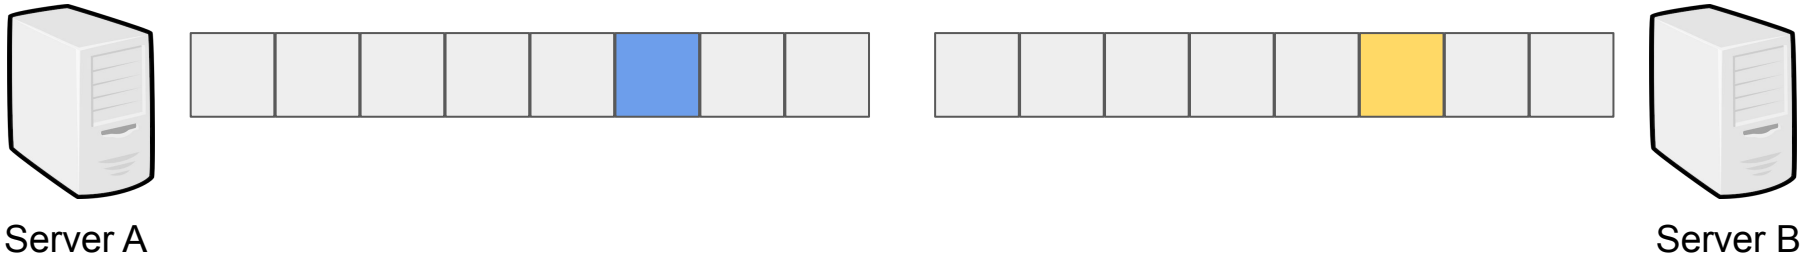

Auditing

Goal: check that values held by servers only differ at one point

Prior work has a semihonest solution where servers use a cheap MPC (only 2 multiplications) to verify this property.

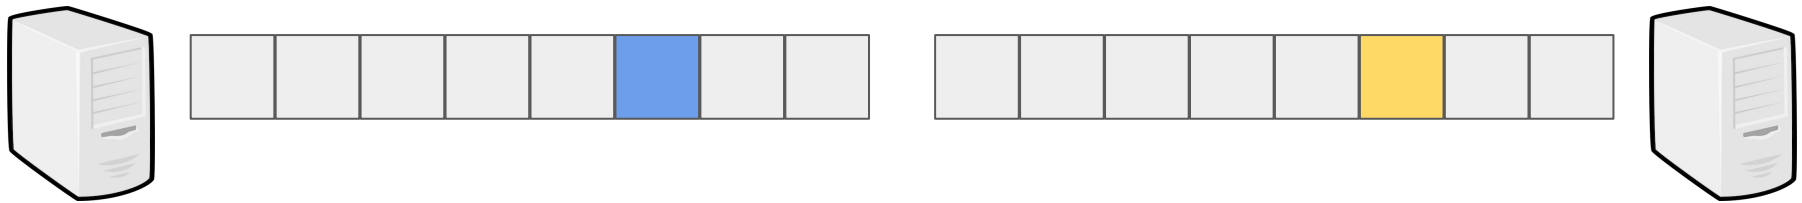

Function Secret Sharing: Improvements and Extensions, Elette Boyle, Niv Gilboa, Yuval Ishai, *CCS'16.*

Auditing

Goal: check that values held by servers only differ at one point

Prior work has a semihonest solution where servers use a cheap MPC (only 2 multiplications) to verify this property.

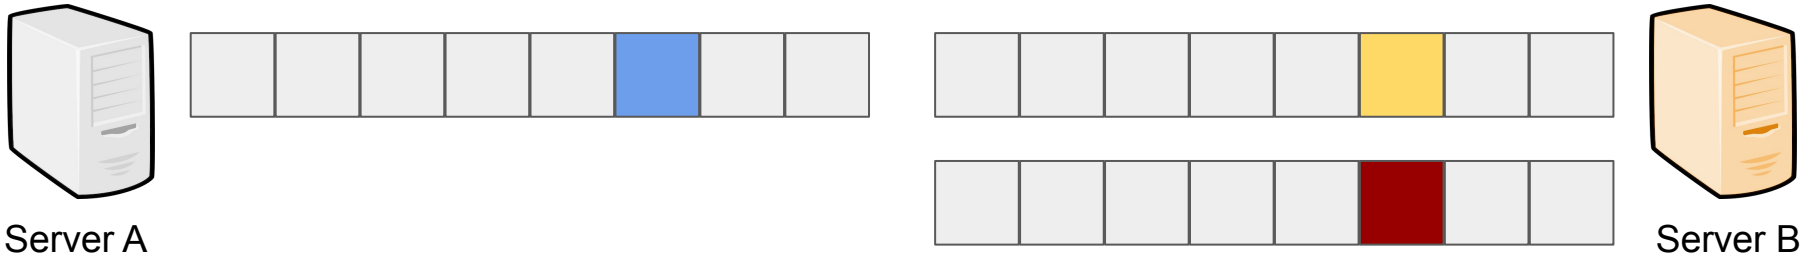

Issue: malicious server can guess and check the differing entry

Auditing

Tool: secret-shared non-interactive proofs (SNIPs)

Prio: Private, Robust, and Scalable Computation of Aggregate Statistics, Henry Corrigan-Gibbs, Dan Boneh, *NSDI'17*

Auditing

Tool: secret-shared non-interactive proofs (SNIPs)

Idea: client sends SNIP proof to servers that honest evaluation of the semihonest protocol accepts

Auditing

Tool: secret-shared non-interactive proofs (SNIPs)

Idea: client sends SNIP proof to servers that honest evaluation of the semihonest protocol accepts

Key Insight: client knows the message index, only needs O(1) work to prove facts about computation that would take servers O(N) work

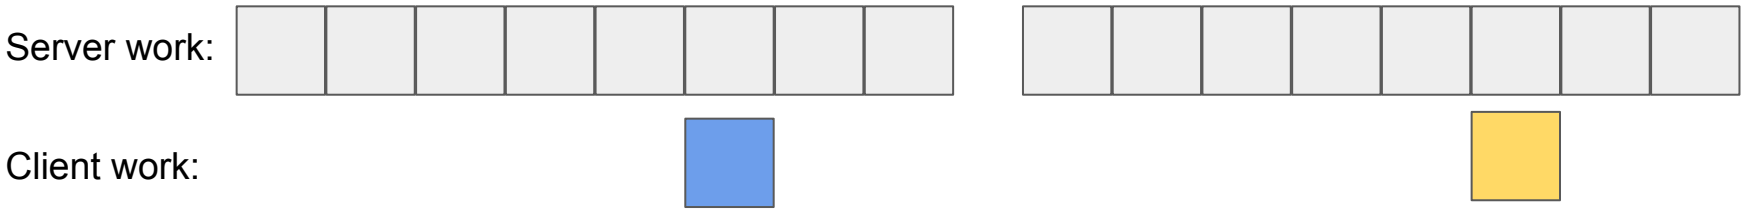

Auditing

Tool: secret-shared non-interactive proofs (SNIPs)

Idea: client sends SNIP proof to servers that honest evaluation of the semihonest protocol accepts

Key Insight: client knows the message index, only needs O(1) work to prove facts about computation that would take servers O(N) work

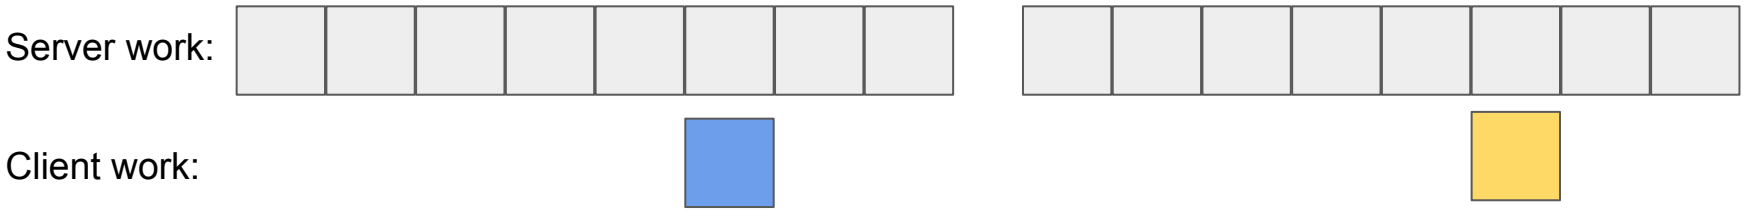

See paper for details

Evaluation: Auditing Protocol

Client runs in under 5 *micro*seconds

55,000x faster than Riposte for 1M mailboxes

Enables 8x reduction in overall client computation (now 20ms)

Comparable on server, where auditing is not the bottleneck

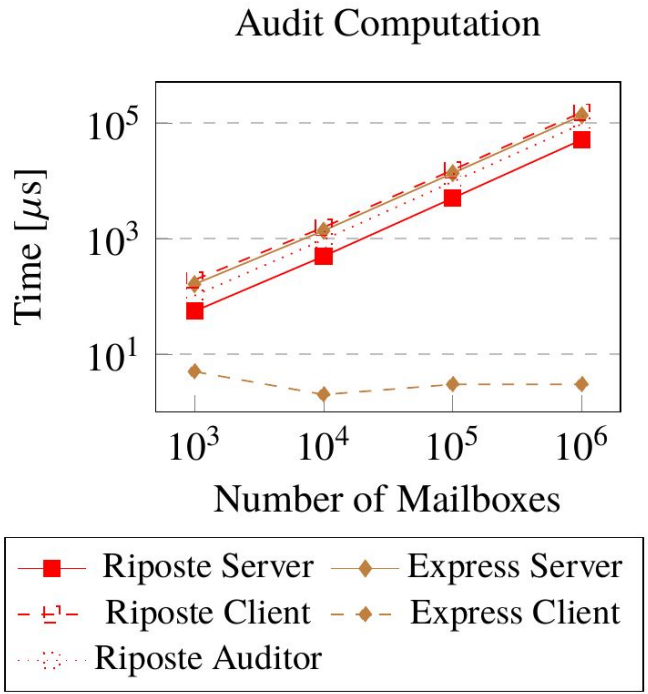

Riposte: An Anonymous Messaging System Handling Millions of Users, Henry Corrigan-Gibbs, Dan Boneh, David Mazieres, *Oakland'15.*
Unobservable Communication over Fully Untrusted Infrastructure, Sebastian Angel, Srinath Setty, *OSDI'16.*

Express

Metadata-hiding communication system with application to private whistleblowing

- Asymptotic speedup from O(√N) to O(1) for auditing
- Speedup of 8x on client, up to 6x on server (compared to Riposte)
- 6x lower dollar cost to operate system
- 13-7,000x or more reduction in communication costs

Code: https://github.com/SabaEskandarian/Express

Contact: saba@cs.stanford.edu