Myco: Unlocking Polylogarithmic Accesses in Metadata-Private Messaging

Darya Kaviani[1] Deevashwer Rathee[1] Bhargav Annem[2] Raluca Ada Popa[1]

[1]*UC Berkeley* [2]*CalTech*

IEEE Security & Privacy '25

Why Hide Metadata?

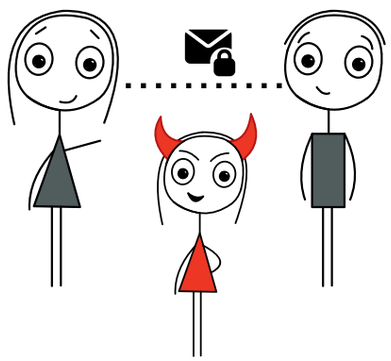

Adversaries in E2EE messaging can see **who** communicates with whom, **when**, & **how many** messages they exchange.

"We **kill** people based on **metadata**."[1]

General Michael Hayden, Former NSA & CIA Director

NSA can derive **25,000** people's networks given one person's call metadata.[2]

Tweet metadata alone identifies users with **96.7%** accuracy.[3]

The Prior Work Trilemma

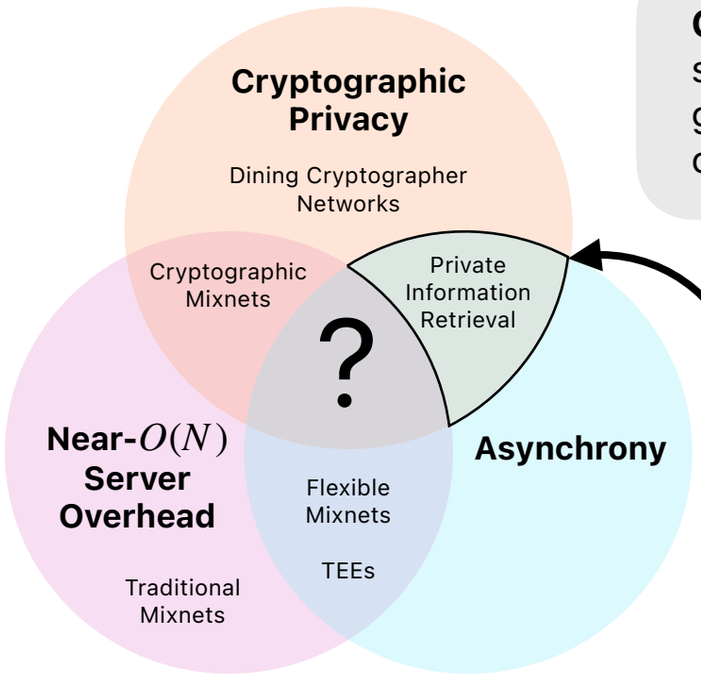

Goal: high throughput, strong cryptographic guarantees, asynchronous client support.

Server Overhead is $O(N^2)$ in the number of system users N!

[1] https://www.nybooks.com/online/2014/05/10/we-kill-people-based-metadata/
[2] https://news.stanford.edu/stories/2016/05/stanford-computer-scientists-show-telephone-metadata-can-reveal-surprisingly-sensitive-personal-information
[3] https://www.wired.com/story/twitter-metadata-user-privacy/

Key Insights

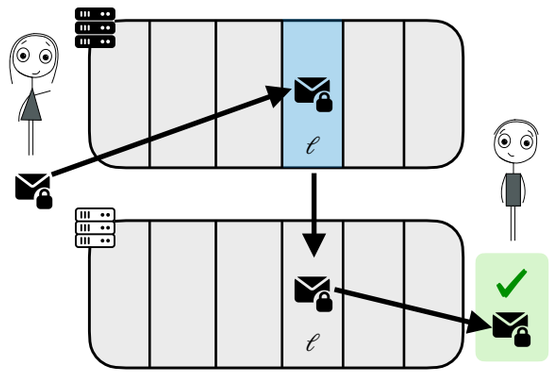

Oblivious **batch** evictions

Use **asymmetric** distributed-trust servers to unlink reads & writes.

Independent bucket assignment

Read-static messaging tree

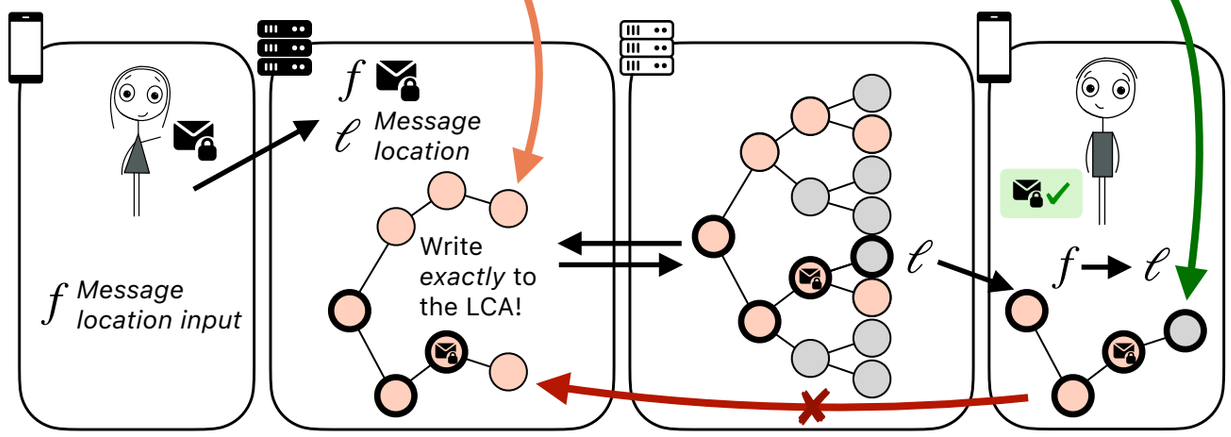

Results

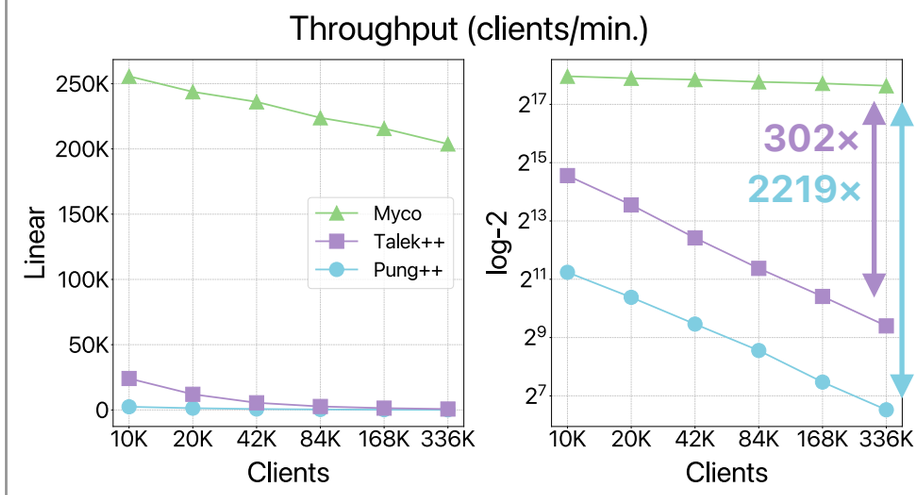

Server overhead: $O(N \log^2 N)$

Future Work

Improve **concrete** communication (<500 KB/client)