Lecture 8: Anonymity

CPSC 444/544 Real-World Cryptography
2026 Spring

Prof. Fan Zhang

Yale

Computer Science

Updates

- Video recording in Canvas
- Next section

Messaging under Strong Adversary

8. (02-04) Mixnets, Onion routing, Tor
    - Suggestion: Loopix (USENIX'17)
9. (02-09) DC nets, Dissent
10. (02-11) E2EE encrypted messaging
11. (02-16) Censorship
    - How the Great Firewall of China Detects and Blocks Fully Encrypted Traffic (USENIX'23)
12. (02-18) Presentations:
    - Attacking and Improving the Tor Directory Protocol (S&P'24, RWC'25)
    - Pudding: Private User Discovery in Anonymity Networks

Question: If your communication is encrypted, is your privacy protected?

Problem: *metadata*

- TLS conceals data content, but information could be leaked through *metadata*
- **Traffic analysis**: Making use of (merely) the *metadata* of a communication to extract information.

What are *metadata?*

- Who talked to whom?
  - PKs, IP addresses, online pseudonyms
- When and where did they talk?
- For how long?
- How many bytes where sent?
- …

“Just Metadata”

Evaluating the privacy properties of telephone metadata

Jonathan Mayer[a,b,1], Patrick Mutchler[a], and John C. Mitchell[a]

[a]Security Laboratory, Department of Computer Science, Stanford University, Stanford, CA 94305; and [b]Stanford Law School, Stanford University, Stanford, CA 94305

Edited by Cynthia Dwork, Microsoft Research Silicon Valley, Mountain View, CA, and approved March 1, 2016 (received for review April 27, 2015)

Since 2013, a stream of disclosures has prompted reconsideration of surveillance law and policy. One of the most controversial principles, both in the United States and abroad, is that communications metadata receives substantially less protection than communications content. Several nations currently collect telephone metadata in bulk, including on their own citizens. In this paper, we attempt to shed light on the privacy properties of telephone metadata. Using a crowdsourcing methodology, we demonstrate that telephone metadata is densely interconnected, can trivially be reidentified, and can be used to draw sensitive inferences.

surveillance | privacy | telephone | metadata | social network

participants of the categories of social network information that the application was requesting. Each screen of the application, until information upload was complete, provided participants with an opportunity to withdraw. Furthermore, participants were furnished contact information for research staff such that they could request deletion of their information after using the application. The university institutional review board suggested helpful methodological refinements, and we began collecting data only after receiving the board's approval.

We also took a number of security precautions to safeguard participant information. Our application transmitted information to a cloud storage service only over an encrypted and authenticated connection [transport layer security (TLS)], and we retrieved information only over TLS. Credentials for accessing the data were restricted to the research team, and once the data were retrieved, the data were stored on encrypted devices at academic facilities.

J. Mayer, P. Mutchler, & J.C. Mitchell, Evaluating the privacy properties of telephone metadata, Proc. Natl. Acad. Sci. U.S.A. 113 (20) 5536-5541, https://doi.org/10.1073/pnas.1508081113 (2016).

Examples from Mayer et al

Health conditions: A person making multiple calls to an oncologist, followed by a pharmacy, could be inferred to have cancer.

Religious affiliation: Repeated calls to a church at specific times (e.g., Friday afternoons or Sunday mornings) indicated religious practices.

Romantic relationships: Frequent, late-night calls between two individuals…

if followed by a sudden stop, suggested a breakup 💔

“Metadata absolutely tells you everything about somebody’s life. **If you have enough metadata, you don’t really need content.”**

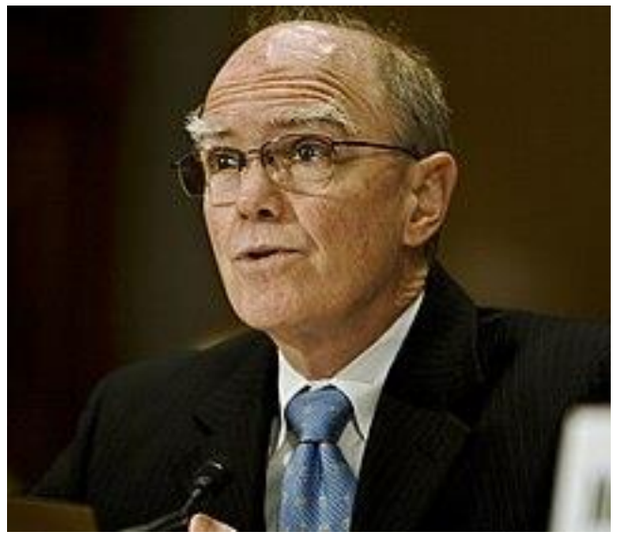

Former NSA General Counsel
Stewart Baker

“We kill people based on metadata.”

General Michael Hayden, former director of the NSA and the CIA

How easy is it to collect and exploit metadata?

- Exposed **by default** in core internet protocols:
  - TCP/IP, HTTP, UDP, FTP, TLS, DNS, …
- Available to **a large number of intermediaries**
  - Local LAN or WiFi router
  - Internet Service Provider (ISP), Mobile network operator
  - BGP routers, Autonomous Systems, Internet Exchanges
  - Internet backbone cables
- Metadata has **lower legal protection** than data content
- Metadata is **machine-readable**, lower volume than content and much **easier to interpret automatically** than content
- Metadata is **difficult** and **expensive** to protect

Metadata-Protecting Communication Systems

Powerful global network adversaries

- The adversary:
  - Can monitor all links in the network
  - Can compromise entities in the network by injecting corrupt nodes or through coercion
    - Can read, inject, delete, modify messages
    - **cannot** break crypto primitives or see inside nodes it does not control
- Example: State controlled ISP
- **Main objective**: break anonymity goals through traffic analysis

Goal: Anonymity

Anonymity is the state of being not identifiable within a set of subjects, the ***anonymity set.***

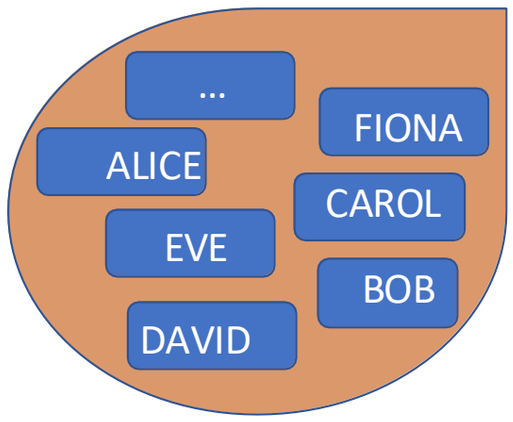

You CANNOT be anonymous on your own
You need a crowd of other (diverse) people

You are MORE anonymous when:
(1) The anonymity set contains more people
(2) You do not stand out within that set

The big picture

Classical secure communication model

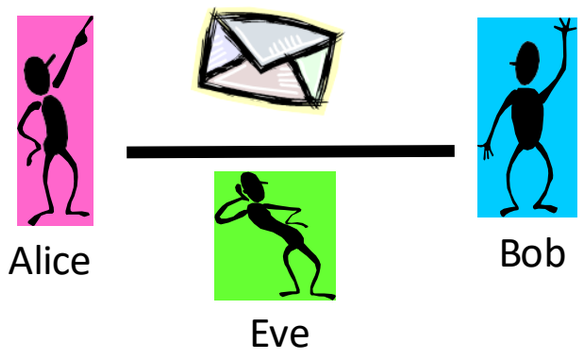

Anonymous communication model

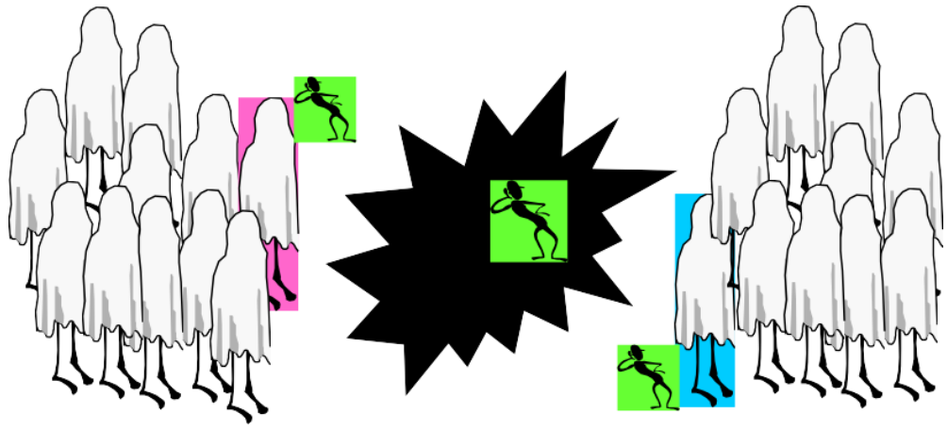

- Many potential senders and receivers
- Attacker may corrupt some of them
- Attacker may eavesdrop on all of the channels

Remark: Anonymity is fragile!

- We will focus on the network layer, but anonymity may break at other layers.
  - E.g., a whistleblower may be partially identified through timing of post (ET or PT), writing style, or their online avatar…

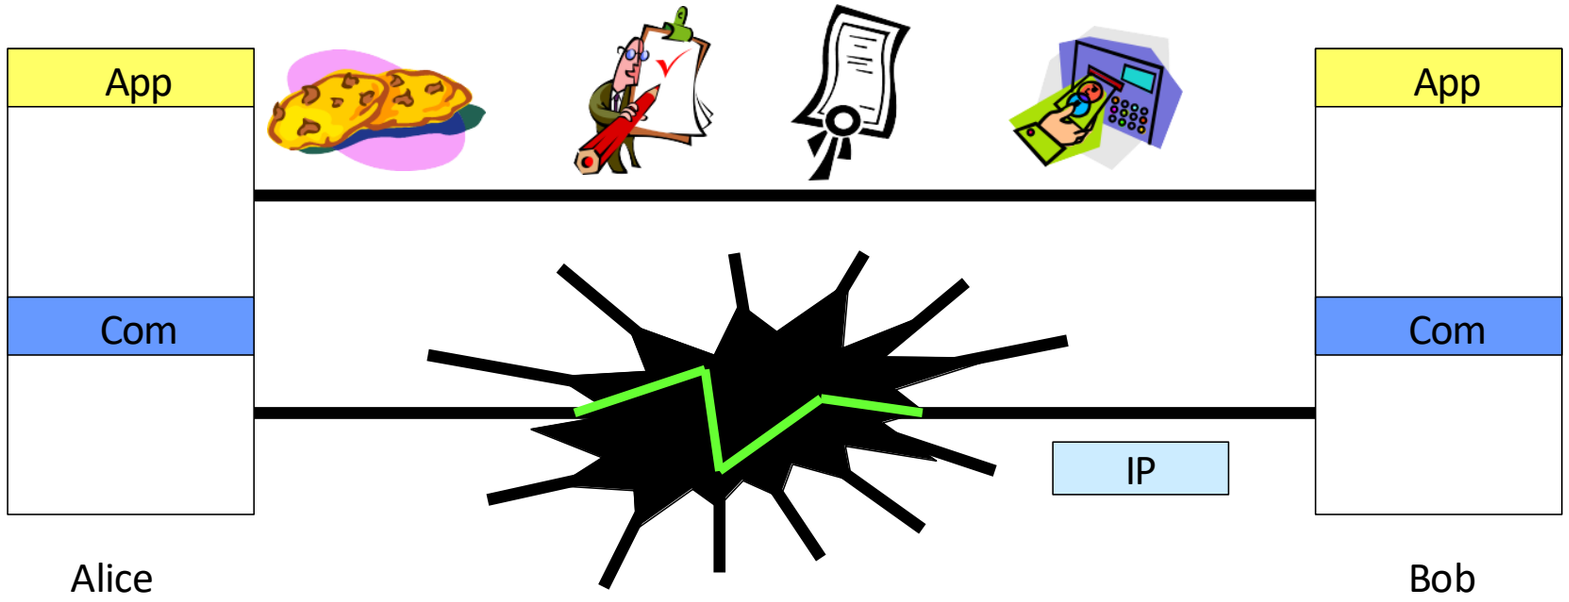

…
Writing style
Persistent behavior
Device+SW fingerprint
Network IP / metadata
Device RF fingerprints

Leakage that enables deanonymisation can occur at multiple layers!

Positive uses of anonymity

- Investigative journalism
- Whistleblowing
- Voting
- Complain while avoiding retaliation
- Personal privacy protection

Harms of anonymity

- Spamming
- Deception
- Hate mail
- Impersonation and misrepresentation
- Online financial fraud
- Other illegal activities

How to build Metadata-Protecting Communication Systems?

- Existing systems can be categorized into 6 families, based on the cryptographic tools:
  - Mix net
  - Dining Cryptographer net
  - Differential Privacy
  - Private Information Retrieval (PIR)
  - Reverse PIR
  - Secure multi-party computation (MPC)

Based on “SoK: Metadata-Protecting Communication Systems”

Our agenda

1) **Mix network (mixnet)**
2) **Tor**: weakened version of mixnet
3) **Dining Cryptographer nets (DC nets)**: anonymous broadcast

Anonymity method #1:

Mix network

Concept of “mixing”

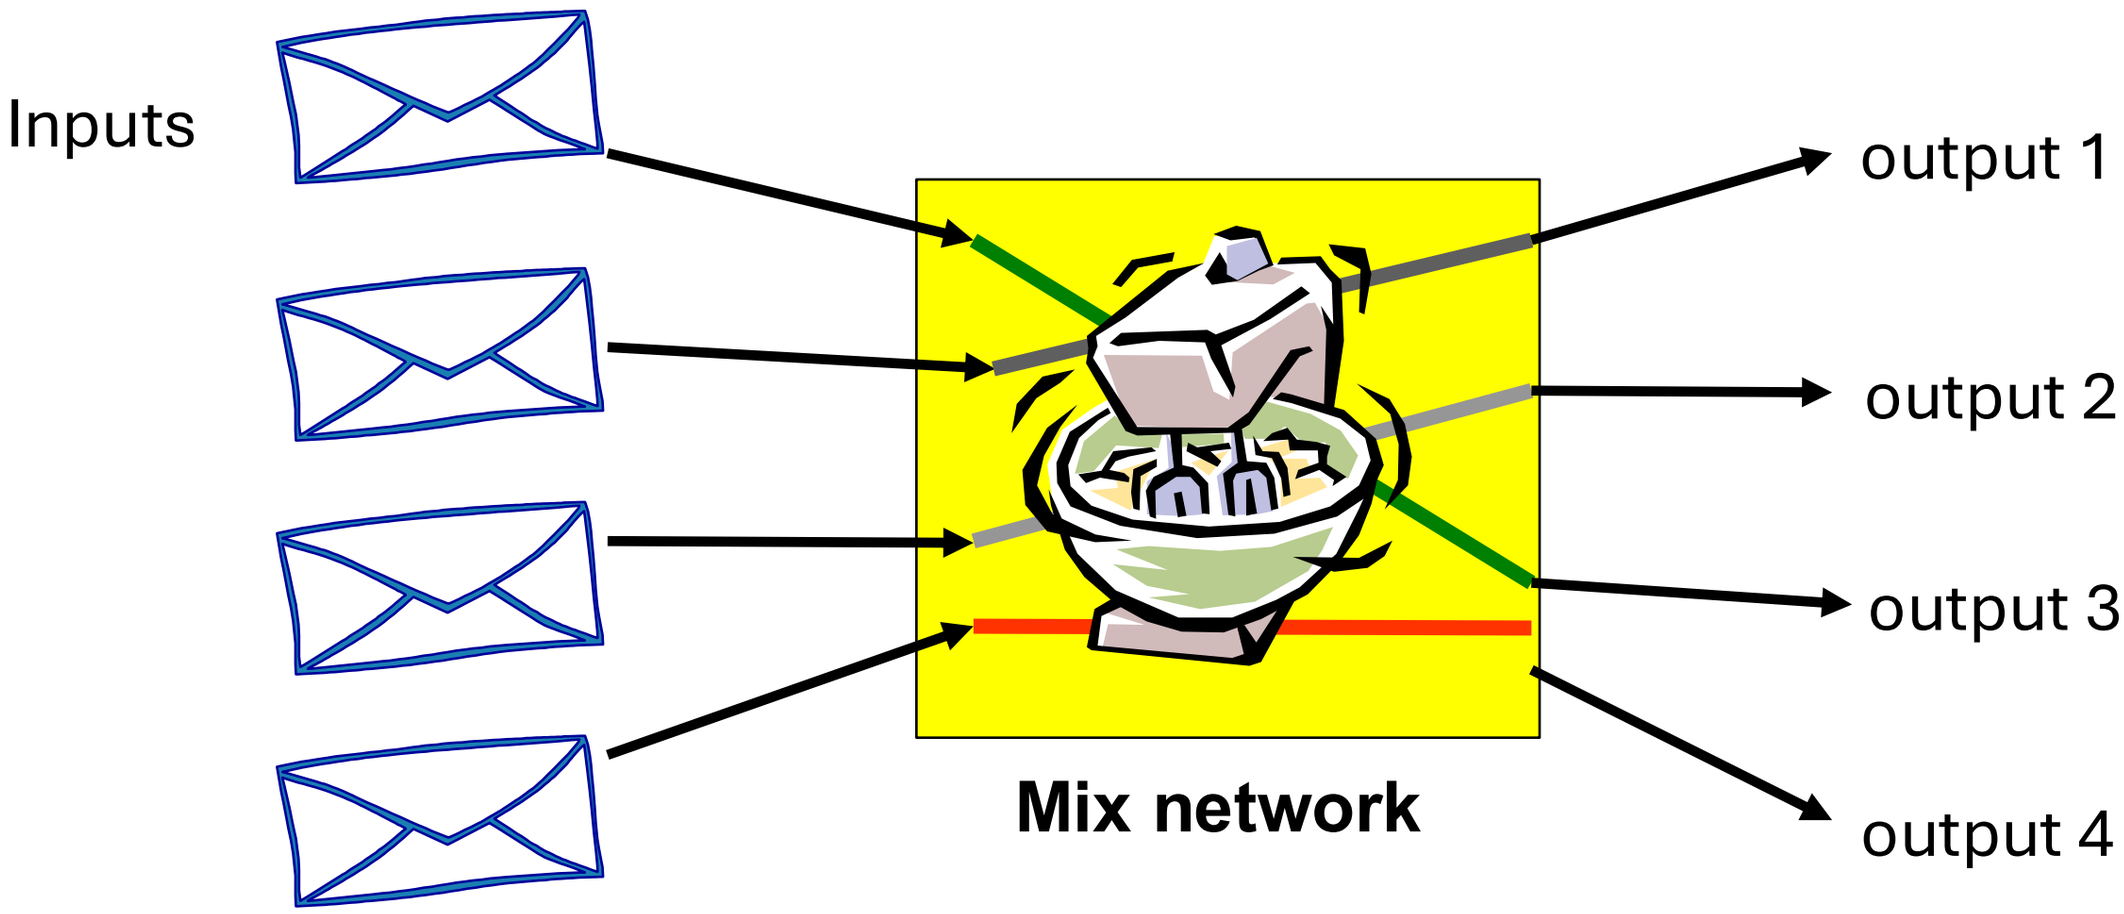

To the adversary, any input is equally likely.

Chaum's Mix

- One of many fundamental inventions by David Chaum
  - *David Chaum. "Untraceable electronic mail, return addresses, and digital pseudonyms". Communications of the ACM, February 1981.*
- Public key crypto + trusted re-mailer (Mix) for **anonymous email**

Basic Mix Design

Notation: $\{M\}_{pk}$ means PK encryption after padding M to fixed length.

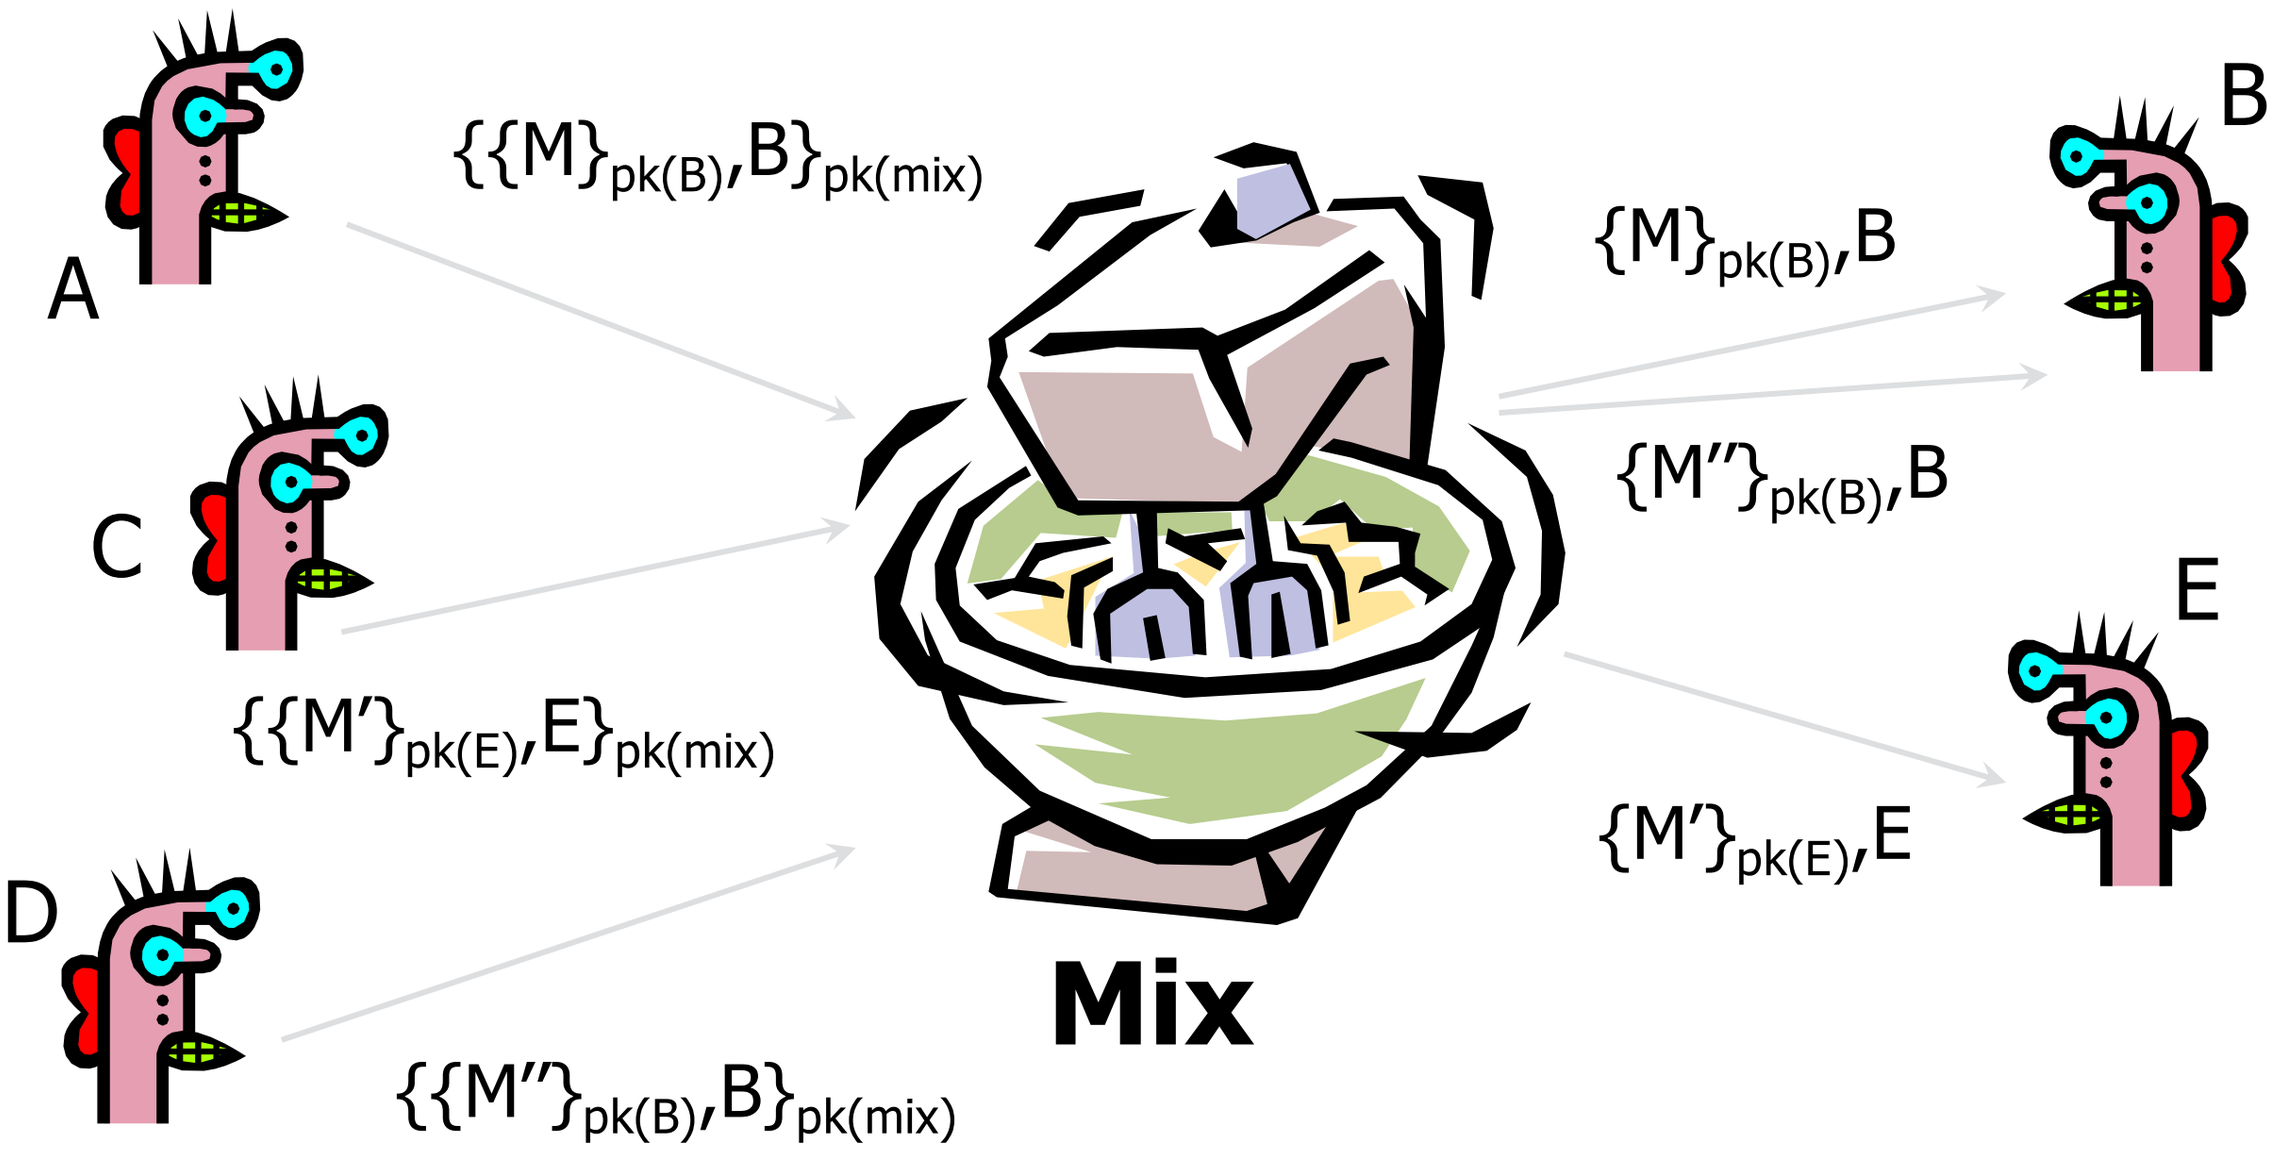

- Batch, decrypt, permute, forward

Basic Mix Design

- Assuming a trusted **mixer…**
- Adversary cannot link a sent message with a received one, due to **encryption** and **padding**
- **Batching** prevents traffic analysis based on message arrival timing
  - Downside: introduces high latency
- **Permuting** prevents linkage based on message order

- Note: still leaks **traffic volume (B got two messages!)**

Cascade

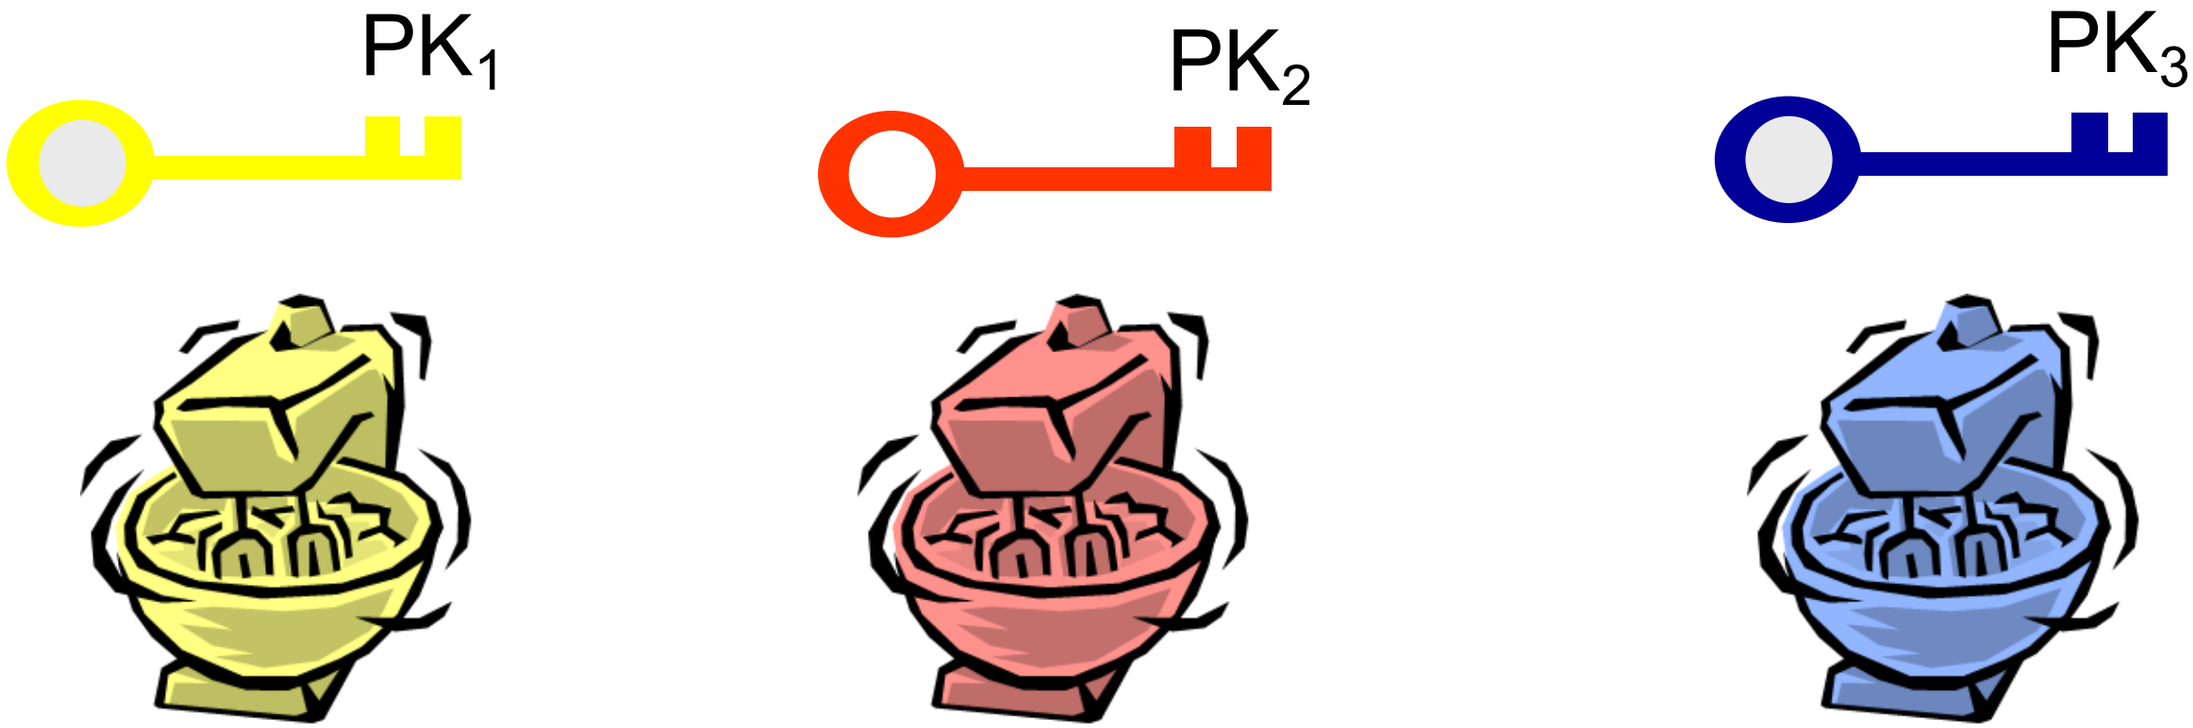

Mix servers can form a network topology.
Most example: a chain

Cascade (chain of three)

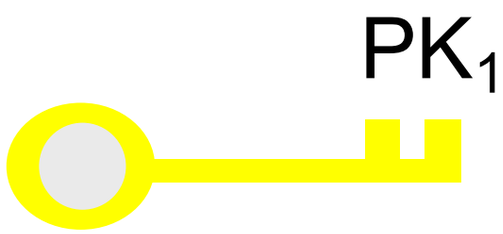

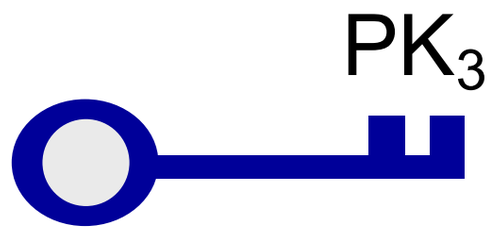

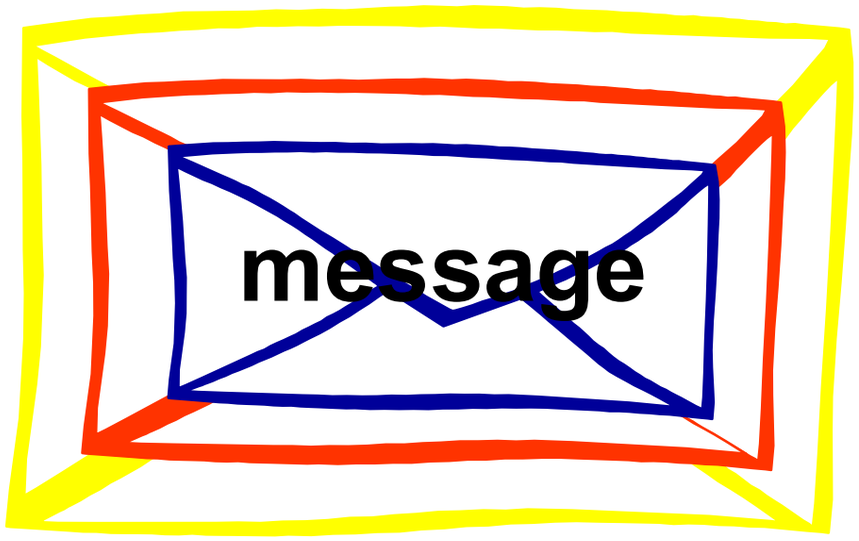

Ciphertext = $E_{PK1}[E_{PK2}[E_{PK3}[message]]]$

Cascade (chain of three)

Server 1

Server 2

Server 3

Generally, n users and m servers.

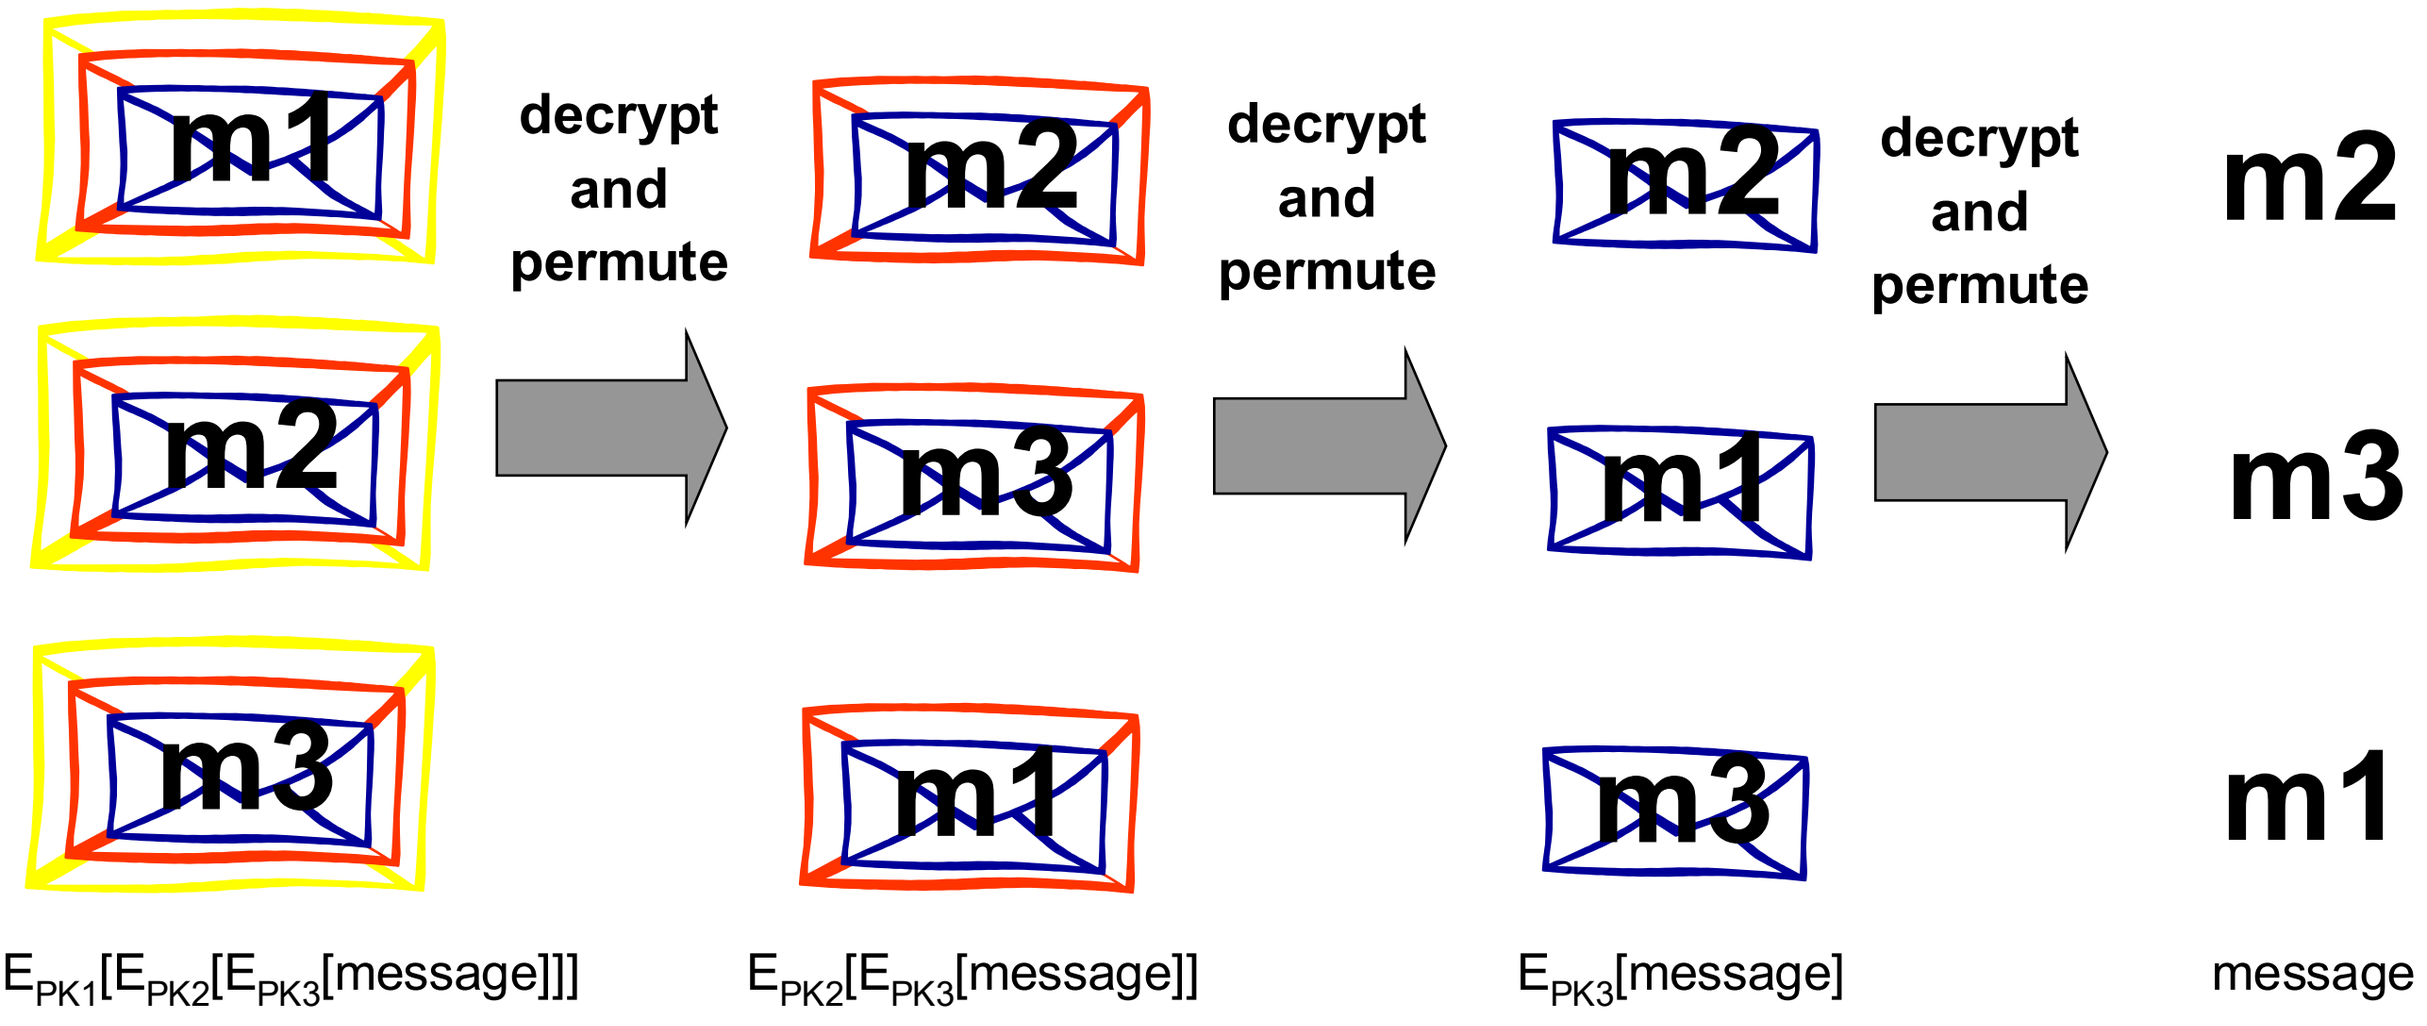

$E_{PK1}[E_{PK2}[E_{PK3}[message]]]$

$E_{PK2}[E_{PK3}[message]]$

$E_{PK3}[message]$

message

Cascade (chain of three)

Server 1

Server 2

Server 3

- In a passive adversary model, a single honest mix can guarantee anonymity! “Any trust”
- Easily broken by active attack, examples:
- **Tagging attack**: duplicate the msg from the target by 1st server and observe which recipient get duplicate messages.
  - One fix: CCA secure encryption + deduplicate per hop
- **Deletion attack**: drop messages to decrease anonymity.

Re-encryption mix

- Decryption mixnet: decrypt then permute
  - With RSA encryption in Chaum's paper
- Re-encryption mixnet: re-encrypt then shuffle
  - Easier to be made verifiable
- Also called *verifiable shuffle*

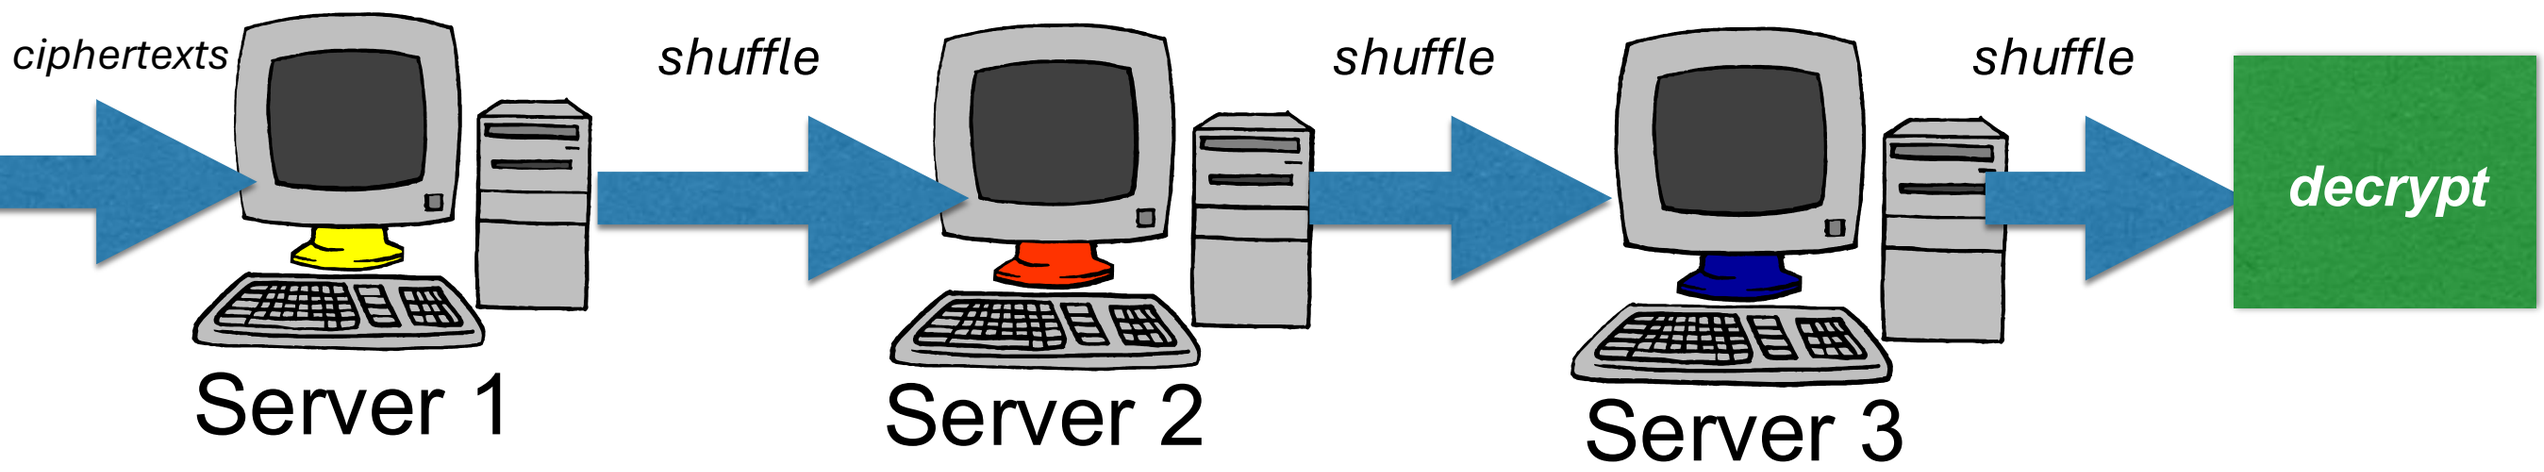

El Gamal re-encryption

- Recall El Gamal
  - Fix (G, q, g) where DDH is hard
  - Encryption of **m** under public key y is $(g^r, m \cdot y^r)$
- To re-encrypt (c_1, c_2), choose a random $s \in \mathbb{Z}_q$, output $(c_1 g^s, c_2 y^s)$
- Observe:
  - 1) Re-encryption preserves plaintext
  - 2) To an observer, they are unlinkable to (c_1, c_2)

Shuffle = re-encrypt then permute

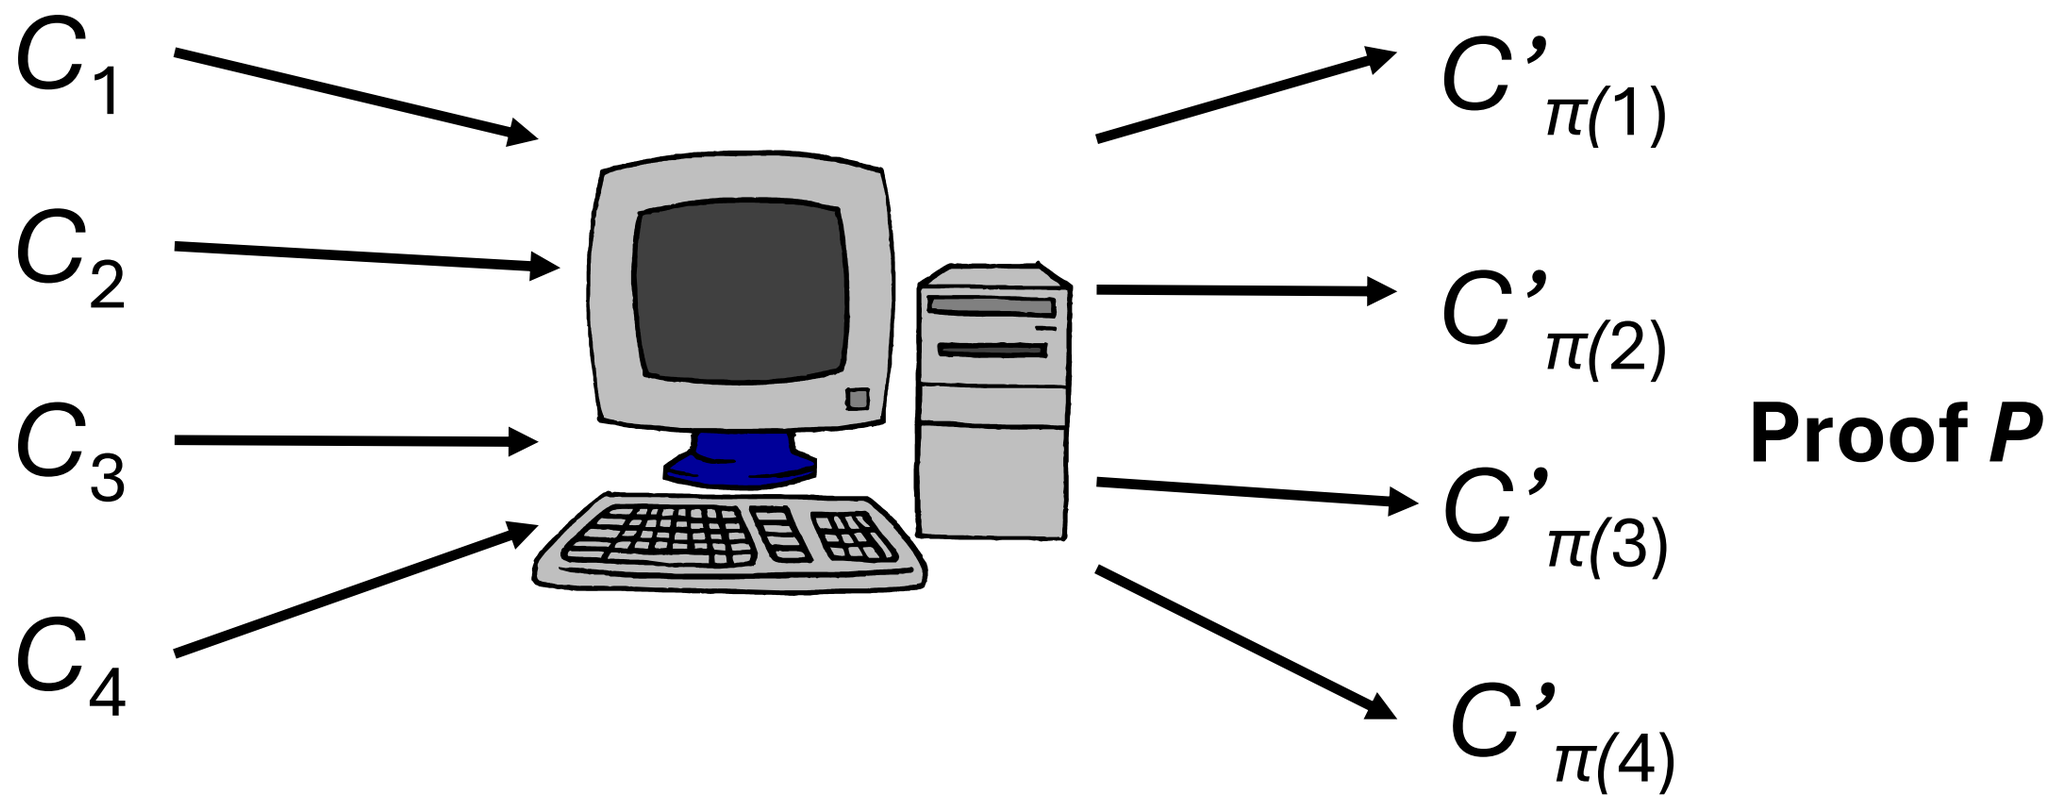

- π is a permutation on {1,2,3,4}
- $C'_{\pi(i)}$ is a re-encryption of C_i
- Verifiable shuffle: *P* proves correctness of above two claims.
- Assumption: the very first input set is fixed.

Approaches to verifiable shuffle

- Standard idea: Each mix proves that $C_1, C_2, \cdots, C_N$ is a re-encryption of $(C_{\pi(1)}, \cdots, C_{\pi(N)})$ while hiding the permutation π.
  - Neff's [CCS'01]
    - Plaintext equivalence is an ingredient (c.f. HW3)
- Smarter ideas:
  - Boneh and Golle [BG02]: the product of a random subset of its inputs is computed, and the mix server is required to produce a subset of outputs of equal products
  - Jacobsson, Juels and Rivest [JJR02]: revealing a pseudo-randomly selected subset of its input/output relations.

Other applications:
Anonymizing bulletin board

Other applications:
Anonymizing bulletin board

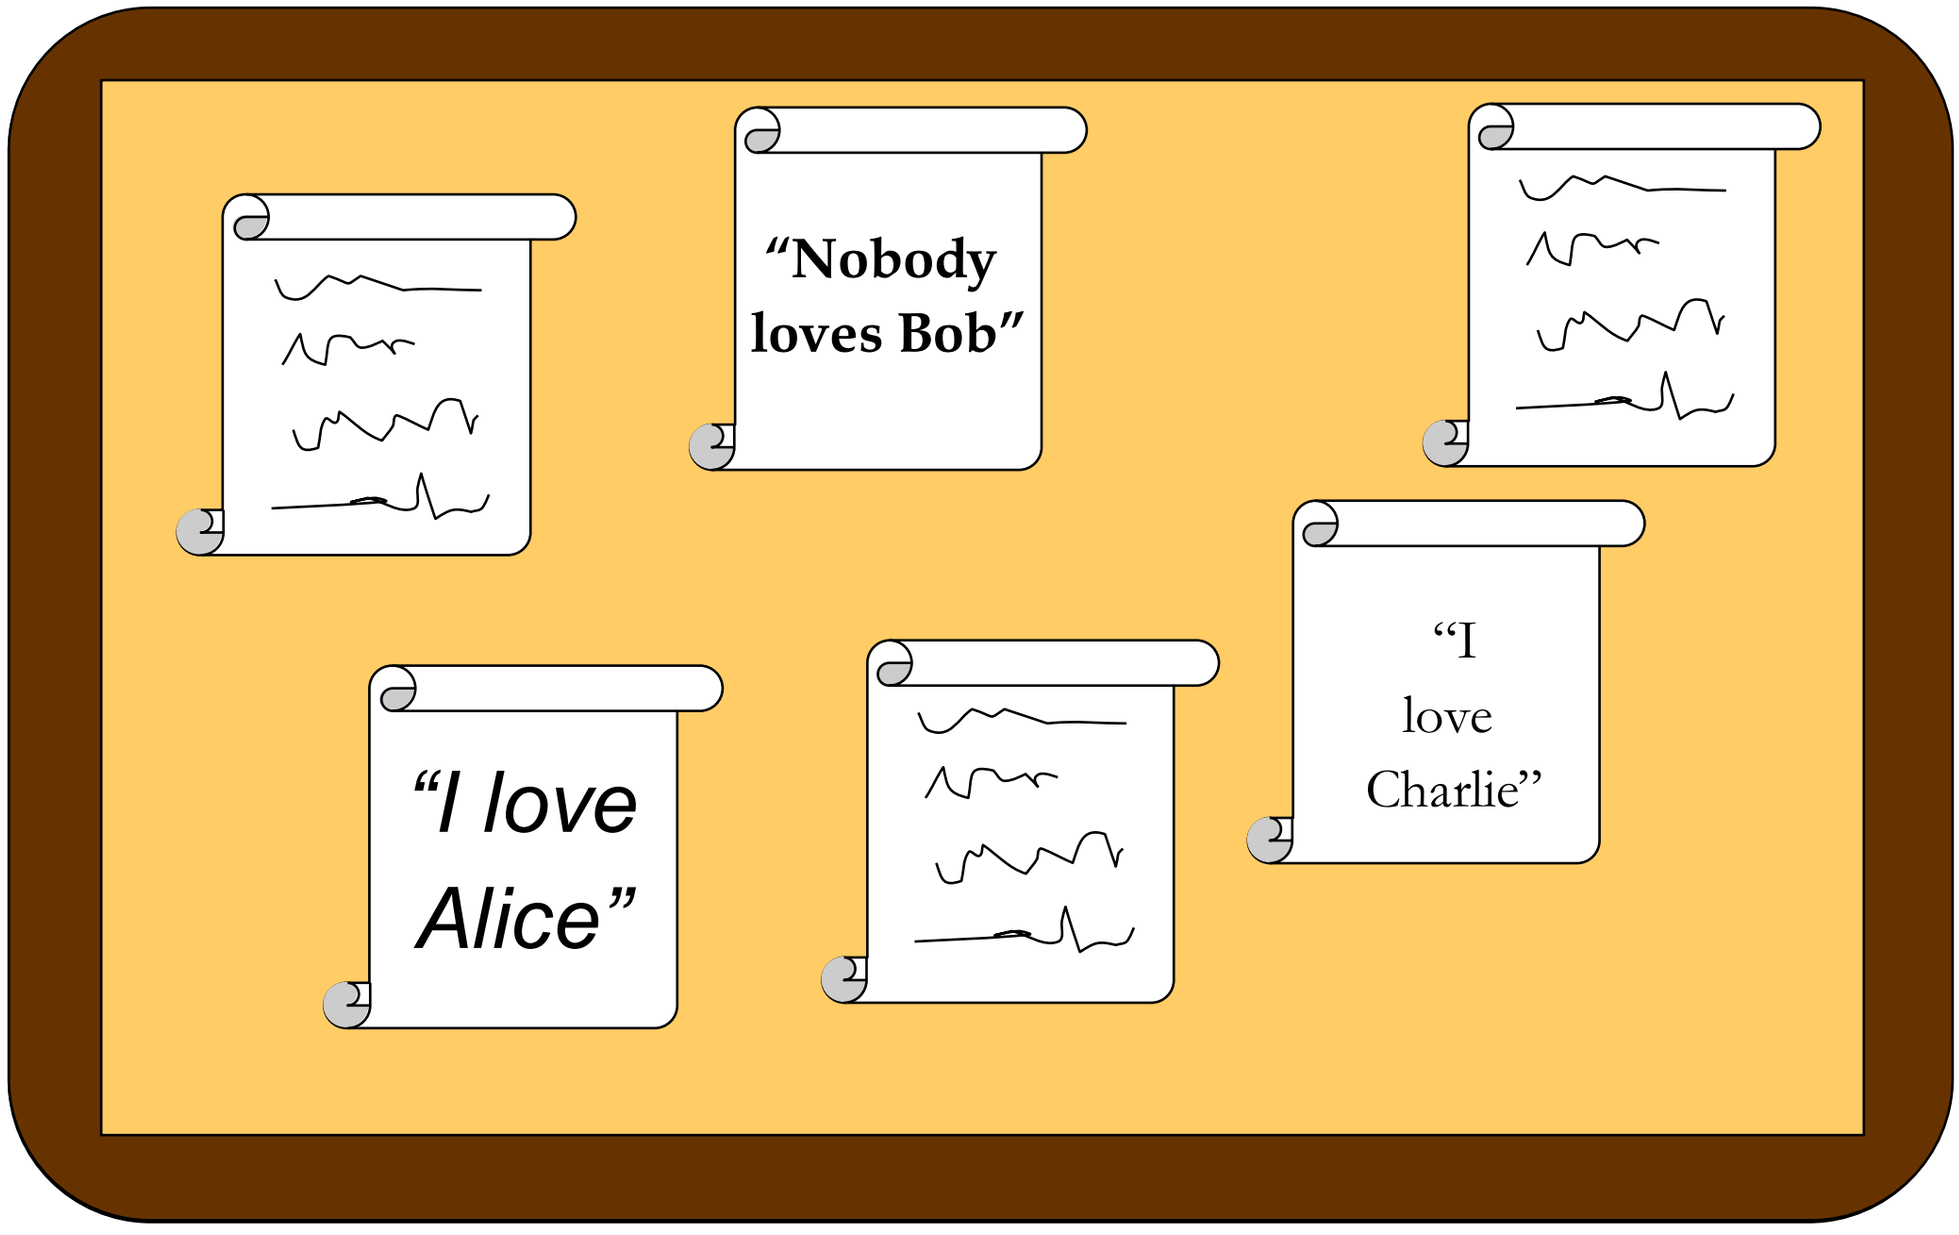

Achieves sender anonymity.

Example application: Censorship-resistant publication

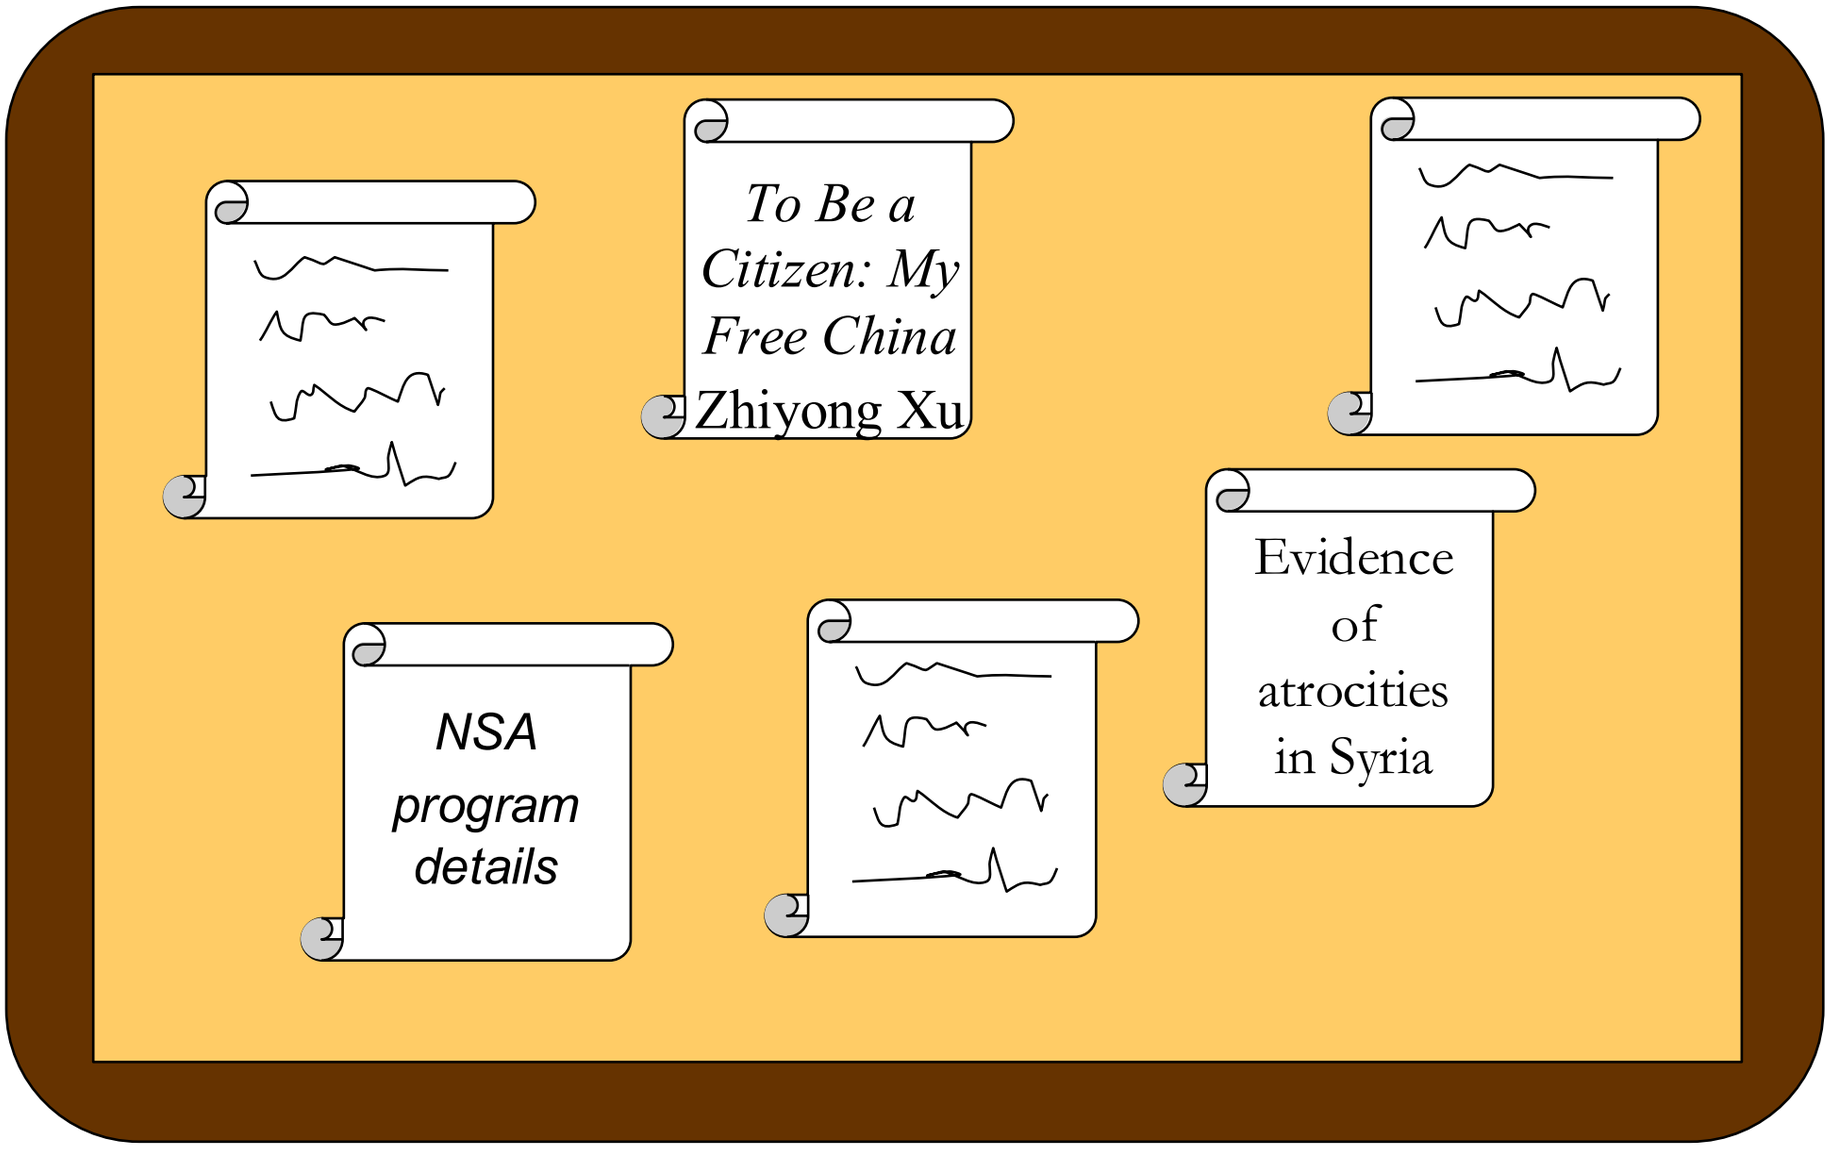

Perfect for (but not actually used by) Wikileaks

Another application: E-voting

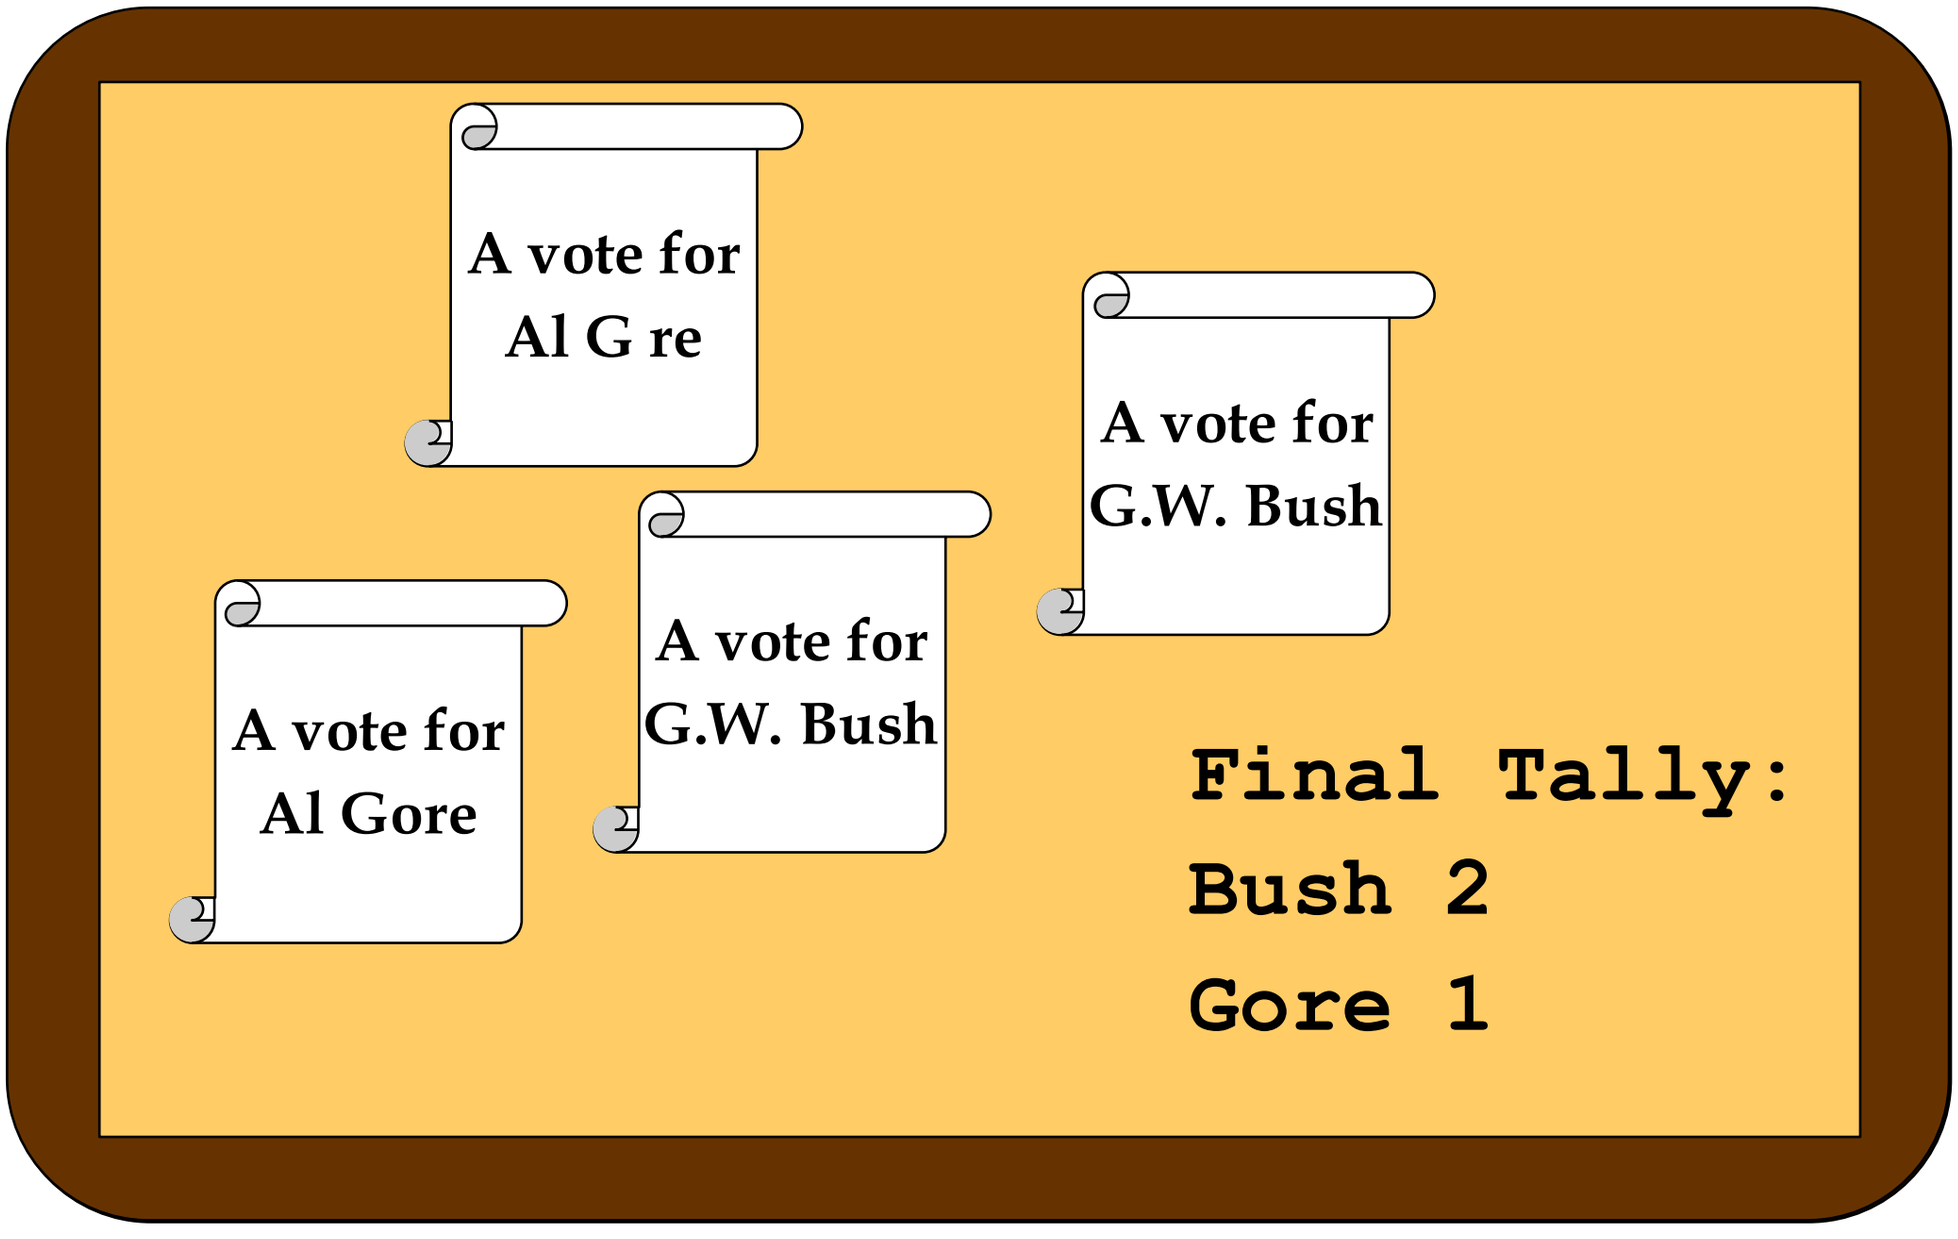

Limitations of synchronous mixnets

- All messages must be processed in a batch, resulting in high latency
  - OK for elections, anonymous e-mail, etc.
  - Bad for web browsing, anonymous voice communications, etc.
- Limited robustness: Any mix going down will stall the entire route

Generic attacks

- Anonymity relies on **batching**
  - Timeout v.s. batch sized
- **Imperfect Batching**: the mix releases messages before the batch is full (due to timeouts)
- **N-1 attacks**: Attacker drop all but the target's message (and inject N-1 adversarial messages)
- **Long-Term Statistical Analysis**: Over many rounds, an attacker can observe correlations in who is frequently active at similar times and make probabilistic inferences.

- Death Note is a mysterious notebook that kills anyone whose name is written in its pages.
- Light used it to carry out a worldwide massacre to create crime-free society.
- Death Note is completely anonymous but Light is eventually deanonymized by
- **Long-Term Statistical Analysis…**

Practical examples

- Re-mailer (Mixmaster): last release in 2008
- Tor: based on the mixing idea but weaker (next lecture)
- NymVPN: uses a 5-hop mixnet

Summary

- **Anonymity** is about hiding links between actions and identities
- **You can’t be anonymous on your own: a crowd to blend in is needed**
- Anonymity needs to be protected at **all layers**
- We saw our first anonymity approach: **Mixnet**
  - decryption-based and re-encryption based
- **Tension** between **anonymity** and **latency**
- Next up: Tor, which is low-latency, but does not withstand a global network adversary.