

Decentralized Systems Engineering

CS-438 – Fall 2025

DEDIS

Pierluca Borsò-Tan

EPFL

Credits: B. Ford, Wikimedia Commons, NSA, A. Dakhnovich

Anonymous Communication

Building privacy-preserving systems

Who has eyes on your internet usage ?

○

○

○

○

○

○

○

○

○

○

Hasn't TLS / encryption solved the problem ?

Metadata absolutely tells you everything about somebody's life.

If you have enough metadata, you don't really need content. – Stewart Baker, NSA

Ex-NSA Chief: 'We Kill People Based on Metadata'

By Lee Ferran May 12, 2014

(Gen. Michael Hayden)

- Location, device, used software, visited websites, sensor usage, etc.

Identifying HTTPS-Protected Netflix Videos in Real-Time

Andrew Reed, Michael Kranch
Dept. of Electrical Engineering and Computer Science
United States Military Academy at West Point

- Based on a fingerprint database of 42,027 videos, they identified 99.5% of 200 random 20-minutes video streams correctly, ~ 90% within 8 minutes.

Why desire anonymity online ?

- Privacy (individuals), Security (business, governments)
- Freedom of speech / journalism / activists
 - escaping censorship
 - avoid speech being linked to oneself
- Avoid ad targeting, tracking
- Bypass geo-blocking
- Helps criminals stay out of jail
- Helps cops investigate online crimes

Threat model

Is our desire to remain anonymous a secret on its own ?

Who are we keeping our identity from ?

- A website
- Advertisers
- A platform (e.g. Meta, Google)
- A well-funded government

What are their capabilities ?

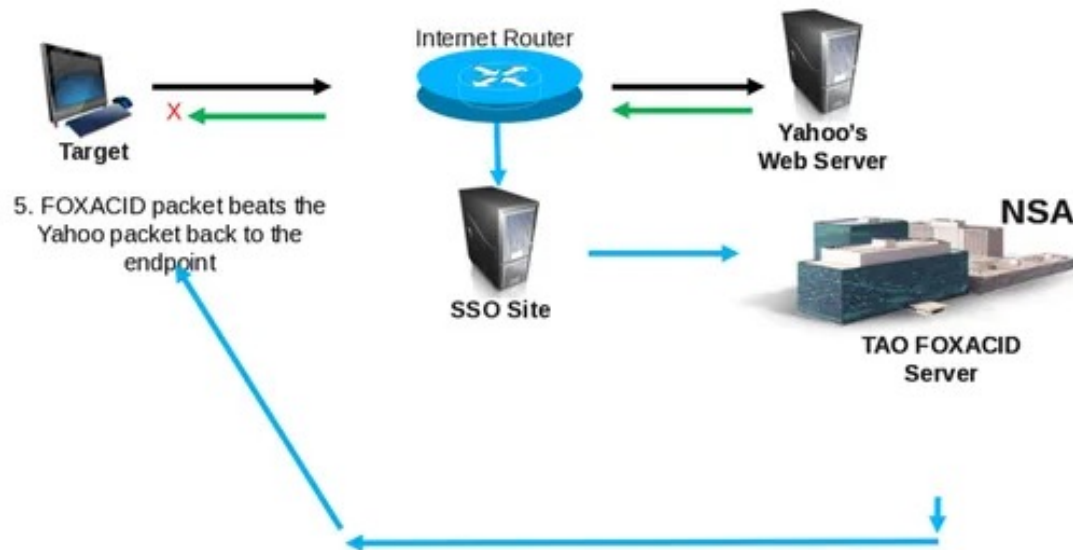
- Cookies, “Supercookies”, fingerprinting
- Semi-honest nodes (“honest but curious”)
- Malicious nodes
- NSA – Xkeyscore, Quantum & FoxAcid (MITM, MOTS)
- CAC (国家互联网信息办公室) – Censorship, MOTS, control over platforms



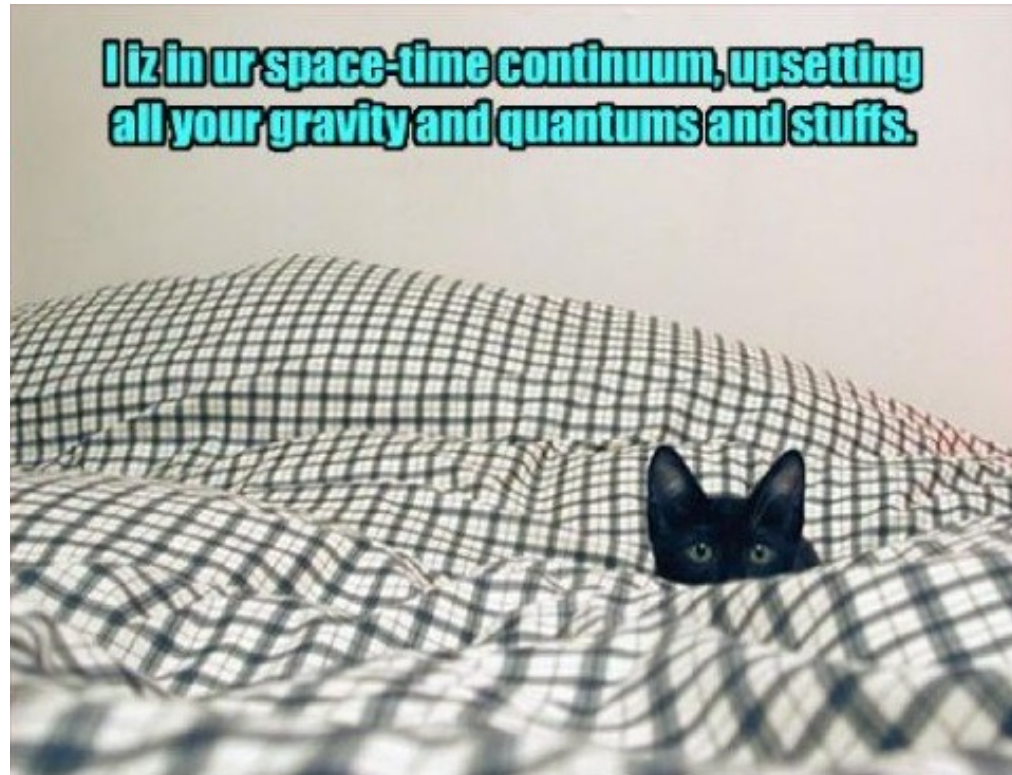
Threat model

What is QUANTUM?

QUANTUM Generic Animation – High Level of How It Works




Threat model



The goal


- Sender and receiver cannot be “linked” by a 3rd party
- Sender and receiver both remain anonymous, including to each other within an *anonymity set*
- Metadata must be unusable for traffic analysis
→ What does this entail ?
- Ideally: censorship-resistant



How to achieve anonymity

1-hop approach:

- Proxy / Commercial VPN



Advantages:

- Shields user from website IP-based tracking
- Prevents geolocation


Problems:

- VPN knows incoming \leftrightarrow outgoing mapping
- Vulnerable to traffic analysis
- Vulnerable to hacking / coercion
- Vulnerable to censorship

How to achieve anonymity

2-hops approach:

- Apple Private Relay



Advantages:

- Shields user from website IP-based tracking
- In theory, no single party sees both sender & receiver

Problems:


- Restricted to countries allowing it
- Apple + 3rd party jurisdiction
- Limited to user's geography
- Only works with some applications
- Vulnerable to traffic analysis

Mix networks (e.g. Mixminion)

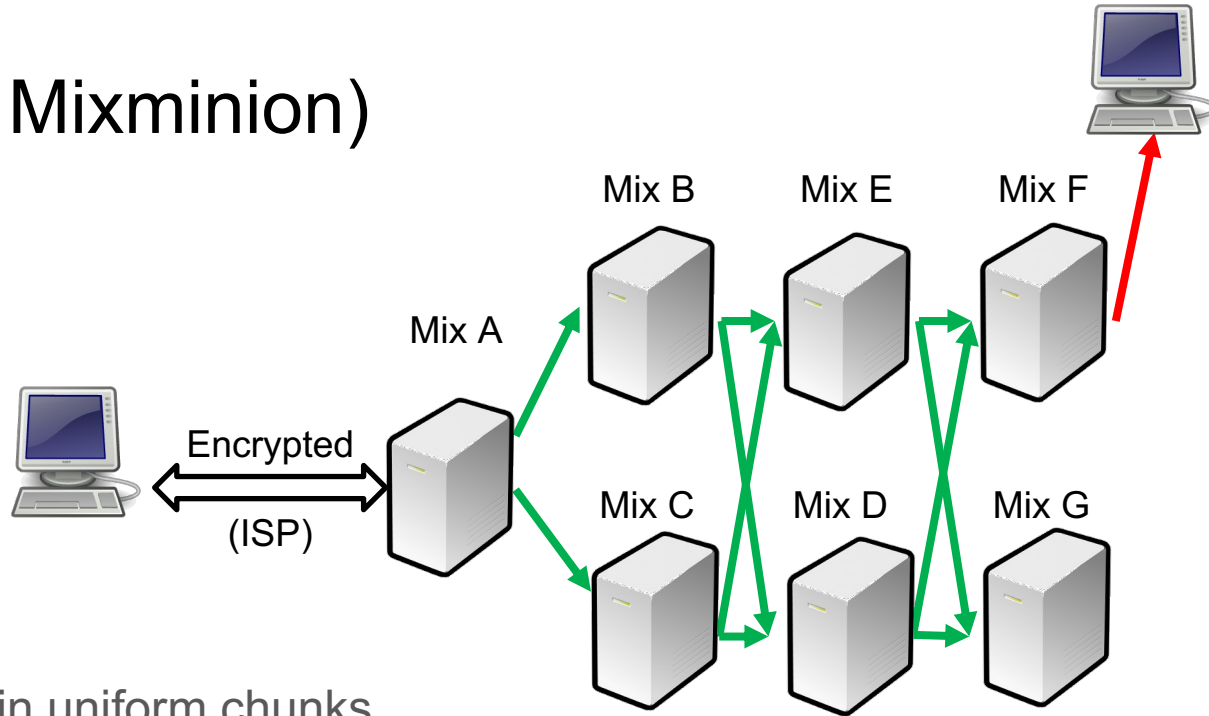
- Goal: anonymize e-mail / Usenet-like traffic

- Key intuition

- Client splits message M in uniform chunks, padded as needed, and encrypts each chunk C for a path through the mix-net



Mix networks (e.g. Mixminion)



- Client splits message M in uniform chunks, padded as needed, and encrypts each chunk C for a path through the mix-net

$$S = K_a \left(R_0, C, K_c \left(R_1, D, K_d \left(R_2, F, K_f(C) \right) \right) \right)$$

Mix networks


- Client can stay anonymous & provide encrypted return path for replies
- Works with just 1 honest mixer

Advantage

- provable (strong) anonymity
- may resist traffic analysis

Problem:

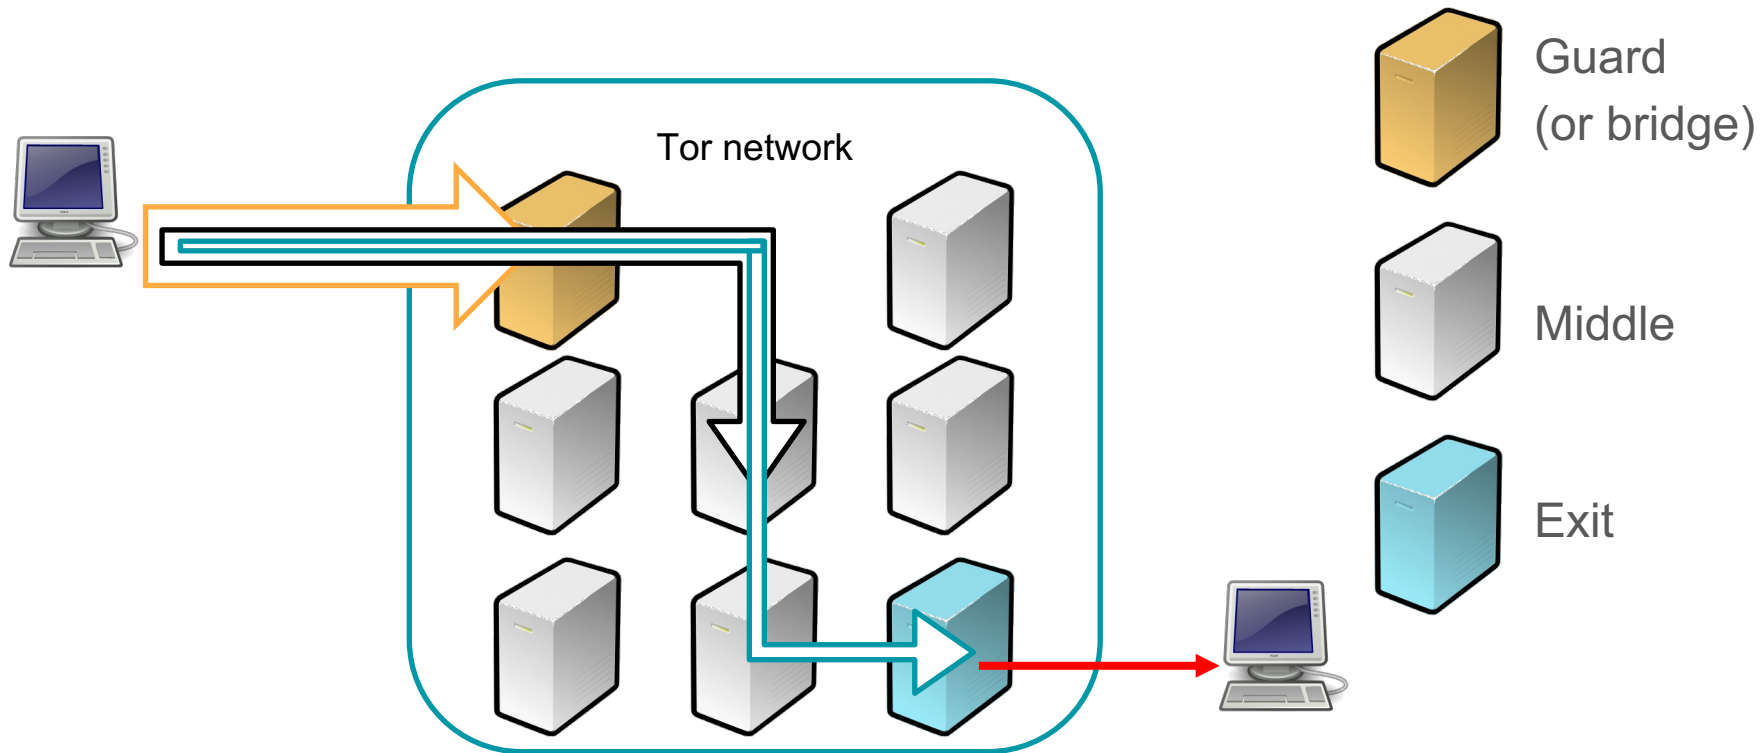
- Very slow, high latency (hours)
- Few users → small anonymity set



Tor (The Onion Router)

- Can we make mix-net work at interactive speeds ?
 - trade-off with robustness to traffic analysis
- Intuition: could we nest multiple VPN connections ?

Tor (The Onion Router)

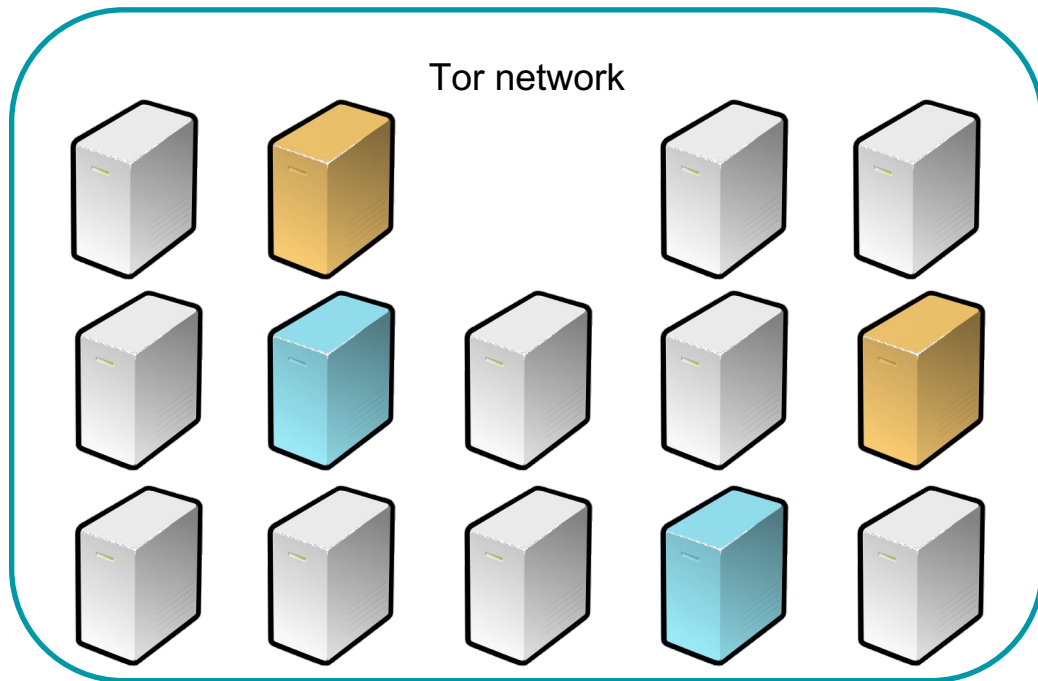


Tor (The Onion Router)

- Making traffic look uniform: each packet is 514 bytes
- How do we find Tor relays ?
 - Hardcoded (10) directory servers !
- New list of all known relays every hour
 - How do they agree on the list ?

Tor (The Onion Router)

- How can this system be attacked ?



Relays



Guard
(or bridge)



Middle



Exit



Tor (The Onion Router)

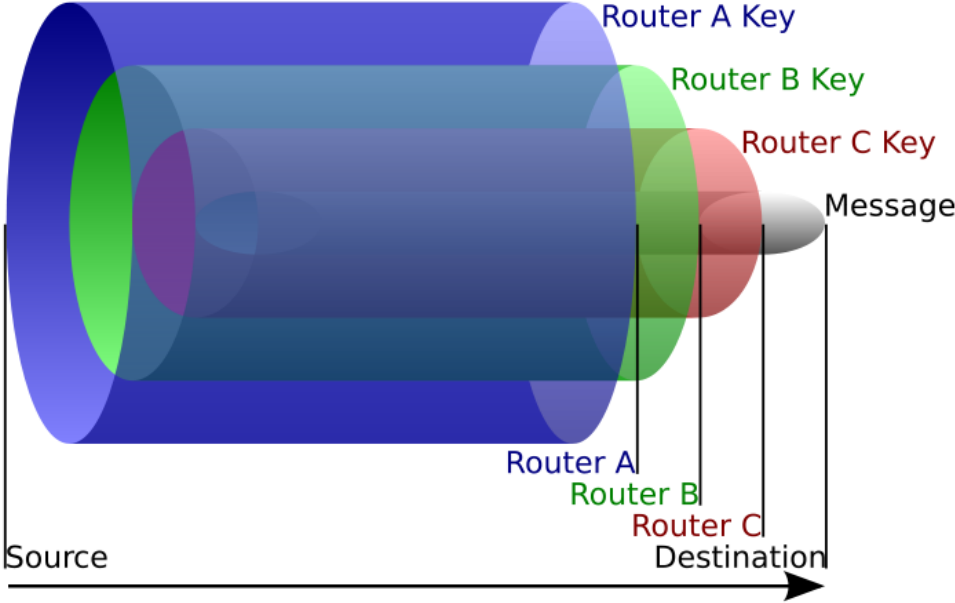
Advantage

- Larger anonymity set
- Low-latency
- Usability, interactive web
- Highly effective against weak adversaries

Problems:

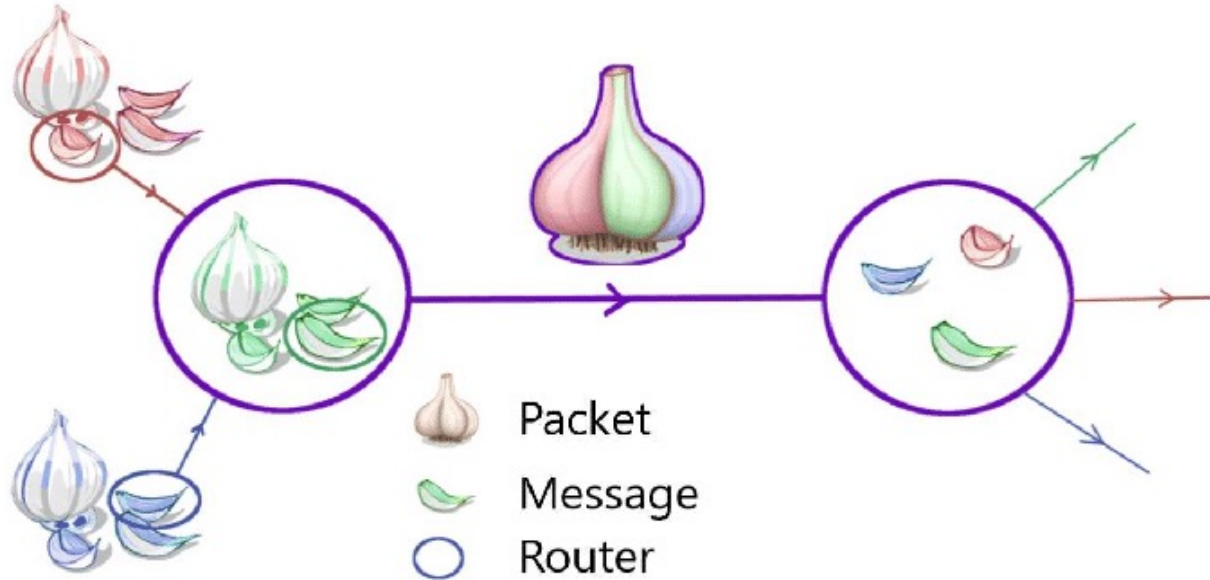
- Weak to traffic analysis attacks
- Web services may block Tor
- Adversary may become global passive adversaries

From Onion ...



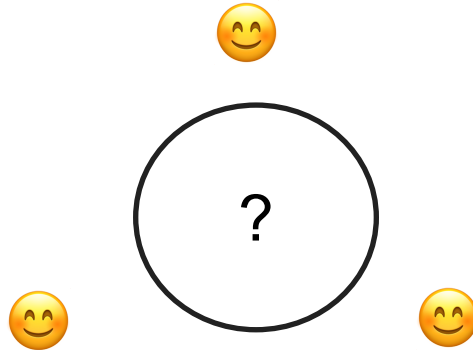
... to Garlic Routing

The I2P approach to traffic analysis resistance

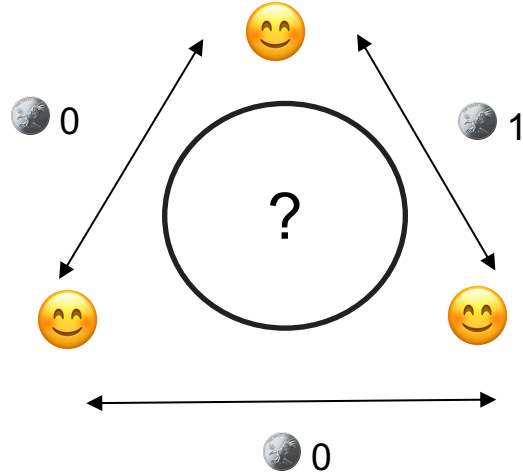


Dining Cryptographers (DC-nets)

- Fundamentally different: information coding, not relay-based
- The classic problem:
Cryptographers are having dinner & a waiter tells them the bill has been paid
They want to find out if one of them paid OR if someone else (the NSA) did
without revealing who paid




Dining Cryptographers (DC-nets)



My value = Left \oplus Right \oplus (I Paid)

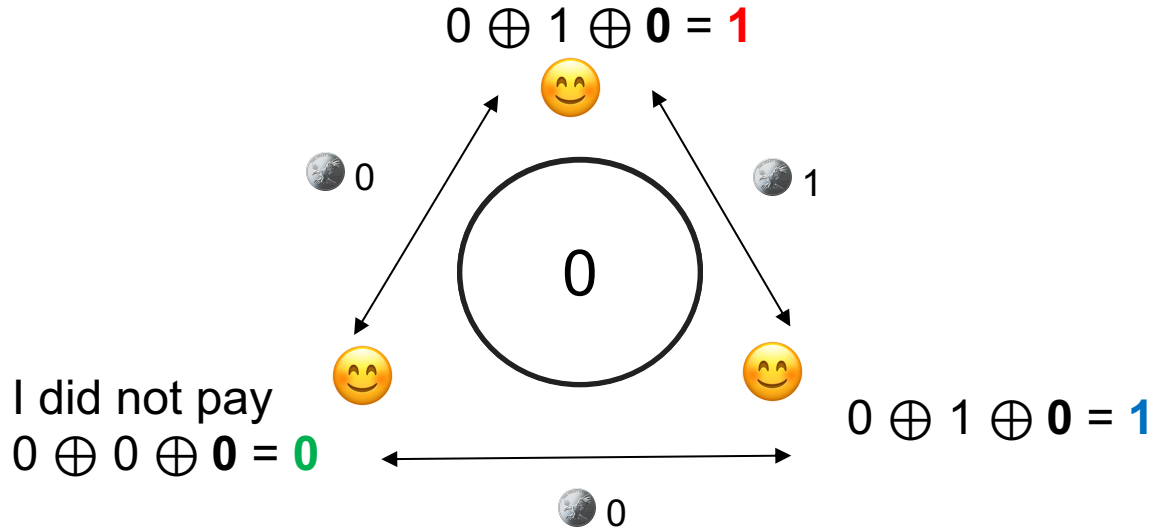
Dining Cryptographers (DC-nets)



My value = Left \oplus Right \oplus (I Paid)

Have we paid ? $1 \oplus 1 \oplus 1 = 1$

Dining Cryptographers (DC-nets)



My value = Left \oplus Right \oplus (I Paid)

Have we paid ? $1 \oplus 1 \oplus 0 = 0$

Dining Cryptographers (DC-nets)

Advantage:

- Provable, information theoretic anonymity
- Security independent of relays

Disadvantages:

- Naive implementation is inefficient $O(n^2)$, easily to disrupt internally
- Many optimizations and strengthening techniques exist and are needed

e.g. Scaling by avoiding all-to-all communication (past research at DEDIS)

- Few servers (m)
- Many clients (n)

Next steps

Reading on Moodle:

Mandatory:

- Tor: The Second-Generation Onion Router
- The Dining Cryptographers Problem

Optional:

- Plenty of papers on anonymous communication systems