Decentralized Systems Engineering

CS-438 – Fall 2025

DEDIS

Bryan Ford and Pierluca Borsò-Tan

EPFL

Credits: B. Ford, Wikimedia Commons, VISA, Swiss Govt.

So far…

DHT, Consensus, etc.

Failure modes

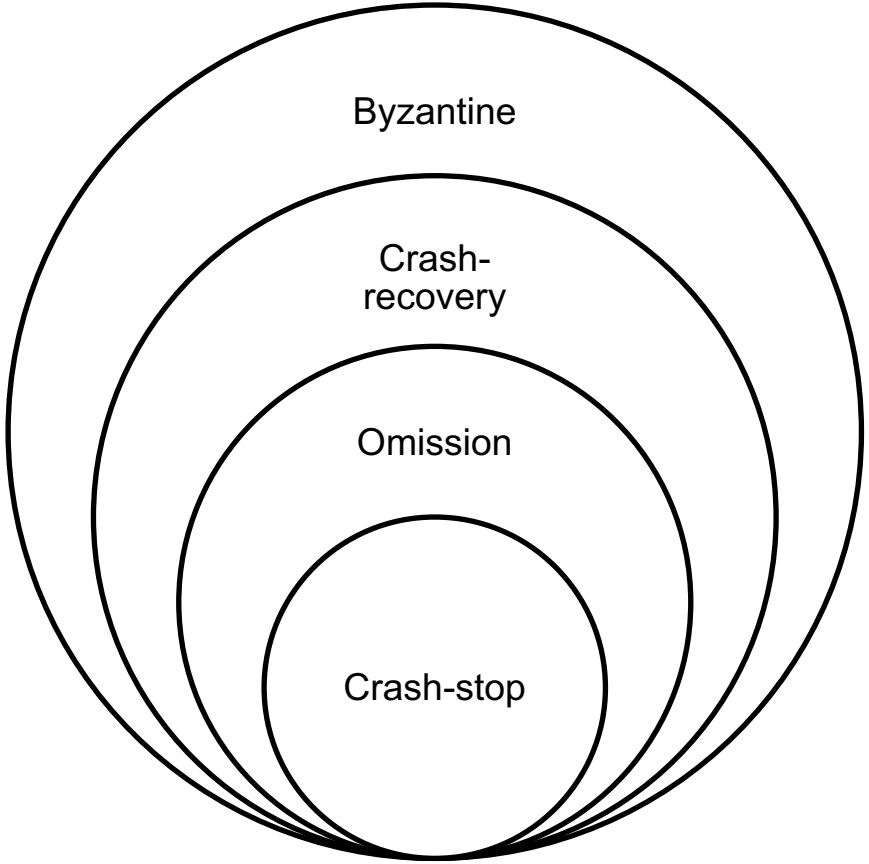

At worst: Byzantine

Nodes can misbehave *in arbitrary ways*

We still assumed a known total number of nodes (and known identities)

… not today !

Sybil Attacks and Defenses

Fake it till you break it - thwarting Sybils !

Sybil What ?

1973: “Sybil” – study of a patient diagnosed with multiple personality disorder

2002: The Sybil Attack – John R. Douceur

- *“One can have, some claim, as many electronic personas as one has time and energy to create.”* – Judith Donath

- Fake identities
  -
  -
  -
  -
  -
  -

Sybil Attack – Implications

- DHTs: eclipse attacks
  - Censor nodes
  - Censor key-value pairs

- Compromise threshold-based security (t-of-n)
  - Creeping compromise: slowly increase t, n

- Compromise consensus
  - Force particular decisions
  - Rewrite history
  - Equivocate (multiple histories)

Sybil defenses – an overview

- Permissioned systems
- Stronger identity
- Adding artificial costs
- Social network-based
- Proof of Personhood

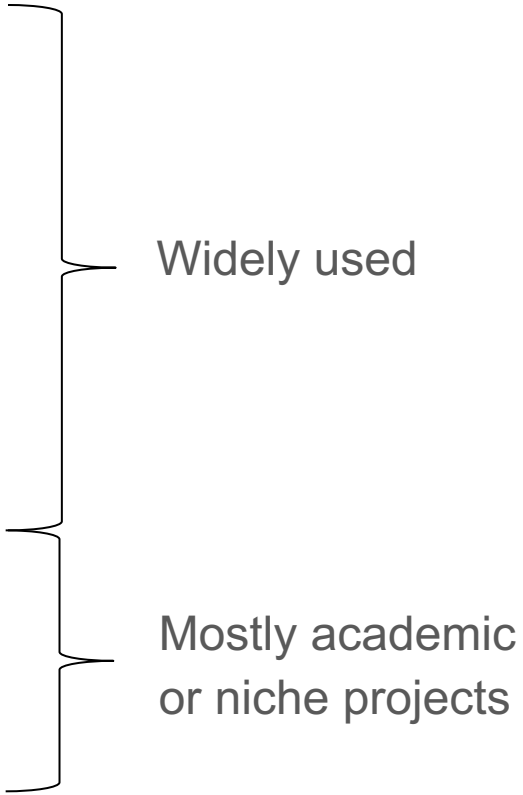

Stronger identities (1/2)

- Sign up with phone number (e.g., WhatsApp)
- Sign up with credit card
- Sign up with e-mail

  me@gmail.com vs. me+cs438@gmail.com

- ID verification
  - Regulatory requirement

    e.g. “Know your customer” (KYC)

  - Deterrents: cost, jail, paper trail

Stronger identities (2/2)

- Biometrics
  - Face
  - Fingerprints
  - Iris

Biggest biometrics databases ?

- Aadhaar (India) – 1.38B
- China – ?
- Common Identity Repository (EU) – 350M
- Dpt. of Homeland Security (US) – 270M

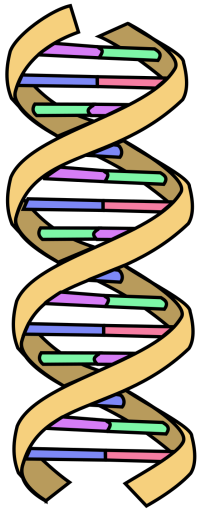

Stronger identity – weaknesses ?

- Privacy
  - Needs centralized database
  - DB encoding ?
  - DB usage for authentication
  - DB usage for Sybil resistance

- Forgeability
  - Fake “fingerprints”
  - Fake “iris”
  - Biometrics synthesis

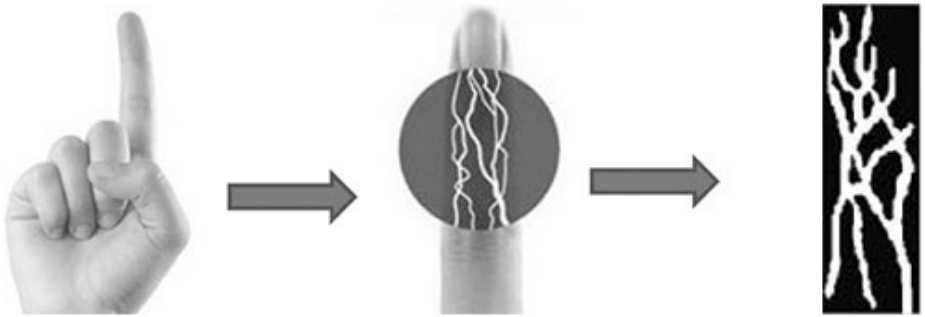

Artificial Costs

- Key idea: increase the cost to Sybil identities
  -
  -
  -
  -
  -
  -

Sybil defenses – artificial costs

- Proof-of-work
  - First proposed for E-mail anti-spam
  - Popularized by Bitcoin

- Crypto puzzle
  - $H(data, nonce) = \underbrace{000 \ldots 000}_{\text{Proof-of-work threshold}} xxxxxxxxx$ → find the $nonce$
  - Doesn't prevent an attack, just increases its costs
  - Not efficient, not environmentally friendly !

Sybil defenses – artificial costs

- Proof-of-stake
  - Nodes must stake money to participate in consensus
  - Randomized validators, likelihood based on stake
  - Misbehaviour punished by loss of stake
  - Risks: hostile takeover, devolution to plutocracy

Social / Trust Network Defenses (1/2)

- PGP “Web of Trust” model
  - Alternative to PKI
  - “Key signing” parties
  - “Alice” → Key A “Bob” → Key B

- PKI / Client-side TLS certificates
  - Company-managed ?
  - Email-challenge ?
  - Not Sybil-resistant !

Social / Trust Network Defenses (2/2)

- Algorithms: generic
  - SybilGuard
  - SybilLimit
  - SybilRank

- Algorithms: application-specific
  - SumUp (recommendations / vote aggregation)
  - Whānau (DHT)
  - dSybil

Social Network Defenses – Assumptions

- Social Graph
- Edges denote “trust”
- Honest region is well-connected

- “Sybil region” scenario
- Attack edges are expensive
- Attack edges are rare/few

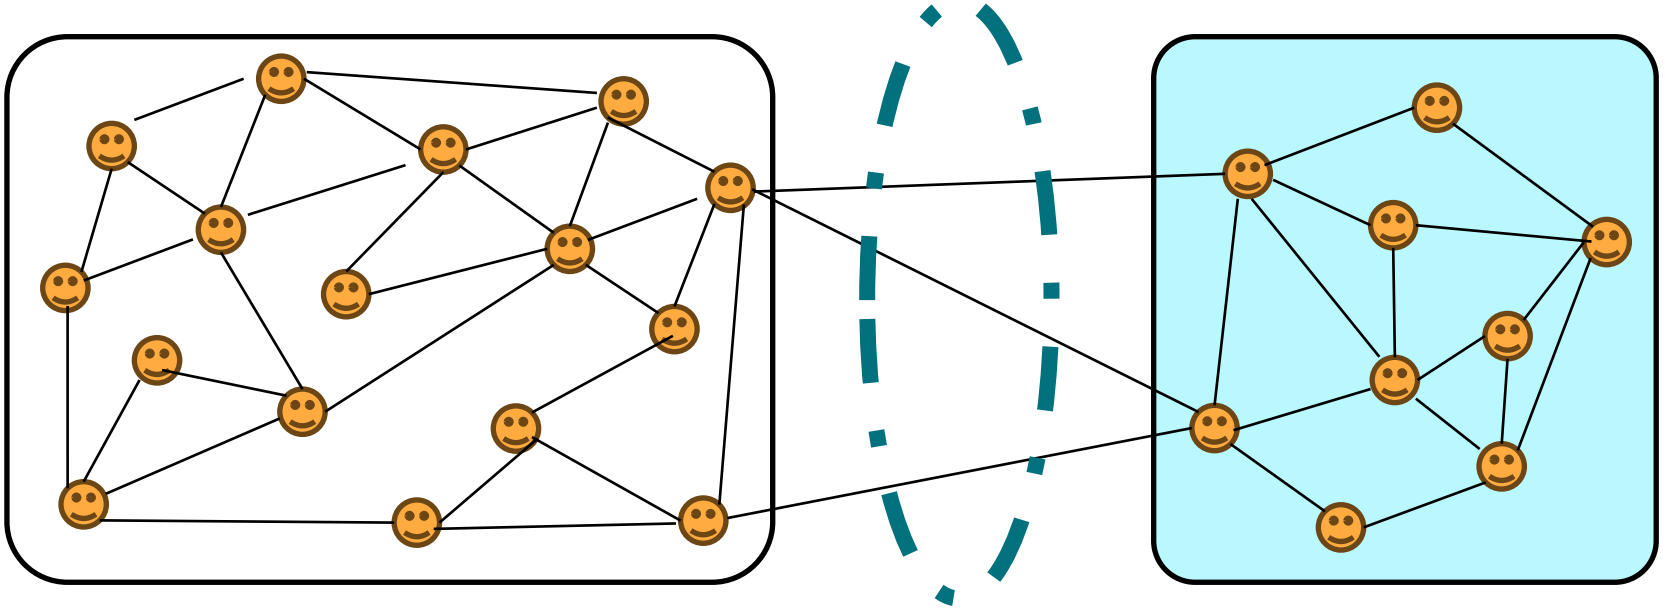

SumUp

- Random walk in the graph
- Assign voting rights to end node
- Repeat

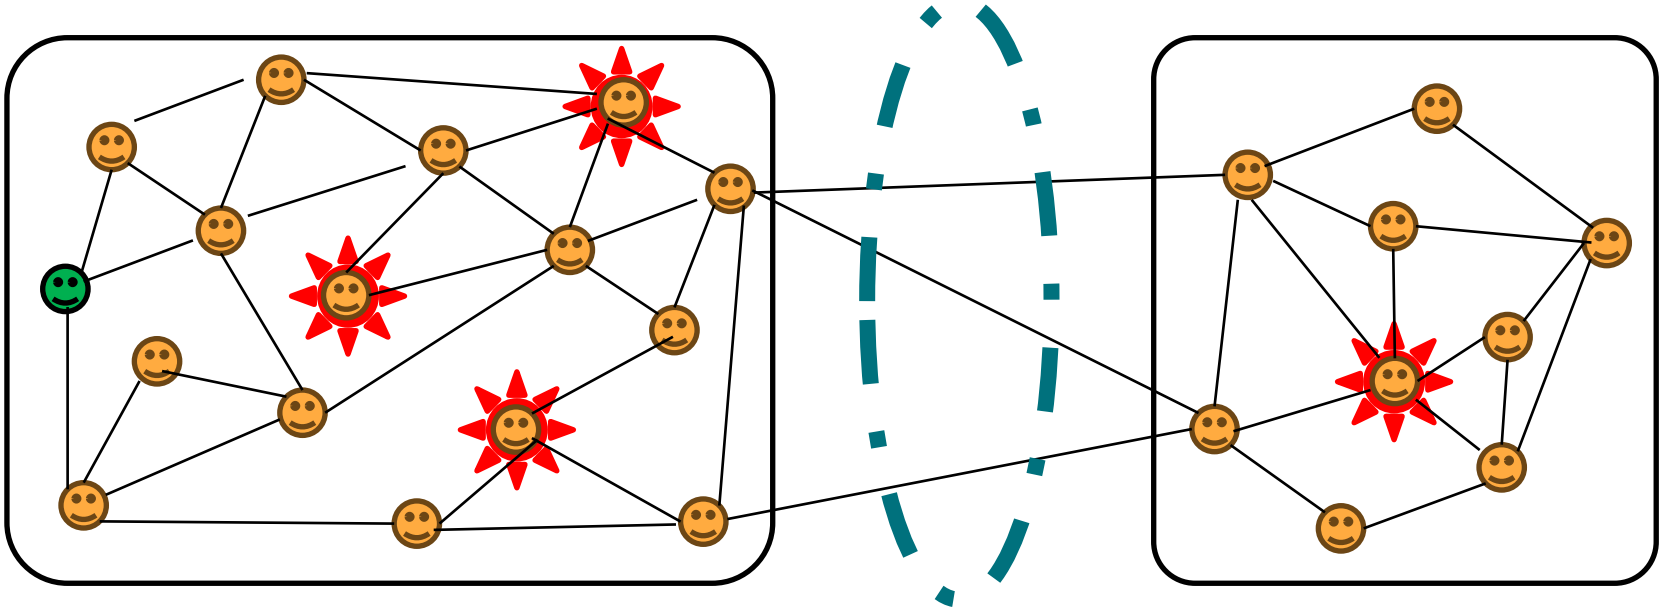

Social Network Defenses – Weaknesses

Basics:

- Privacy
- Performance

Re-thinking the “movie plot threat”

- Crowd-sourcing
- Sparse infiltration
- Small-scale attacks

Sybils on Facebook

Let’s do a thought experiment !

We’re Facebook and trying to detect fake accounts

How?

Proof of Personhood

Key intuition: can we link identity *only* to “being a physical person” ?

Goals:

- Inclusion
  low cost to participation (permissionless)
- Equality
  one person, one vote (strictly)
- Security
  against identity theft/loss and Sybils
- Privacy
  no ID, no biometrics, no databases, etc.

Pseudonym parties

Principle: real people have only one body each

- Attendees gather in “lobby” area by a deadline
- At deadline entrances close, *no one else gets in*
- Each attendee gets one token while leaving

entrances closed

Proof of Personhood – Approaches

- Pseudonym parties
- Encointer

Co-located physical bodies

- Idena

“Flip” tests (Turing tests)

- Humanity DAO

DAO / curated list

- Many others: Upala, BrightID, GoodDollar, etc.

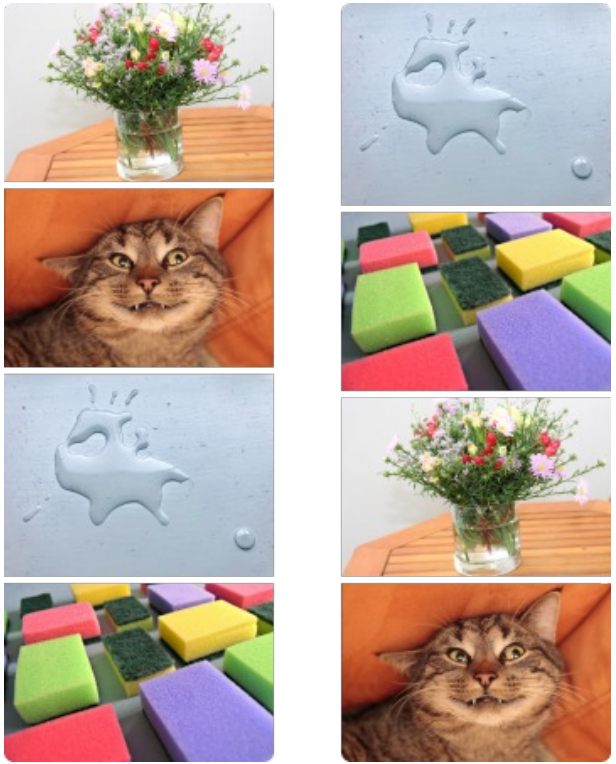

Next steps

→ Review Paxos, try to implement ←

Mandatory reading:

- “The Sybil Attack”

Optional readings:

- Plenty of papers on sybil detection and resistance