Webs of Trust: Choosing Who to Trust on the Internet

Matteo Dell'Amico

Annual Privacy Forum
October 23, 2020

Outline

1. A Reputation System for the Internet?
2. Privacy and Robustness
3. Decentralization and Scalability
4. Conclusions

Sharing Economy

- Users engage in **“peer-to-peer” transactions** with each other
- Unlocks value that **would otherwise be wasted**

Reputation Systems

- Predict who will behave correctly
- Give an **incentive** to behave well

Reputation Systems Shortcomings

Proprietary

- **Privacy**: users' data is kept and controlled by a third entity
- **Economics**: tendence towards a monopoly of those who have more data

Siloed

- **Cold start for users** who don't have reputation in a given system
- **Cold start for new applications** which don't have reputation info

(star rating icon by Magicon from the Noun Project)

Respecting Privacy

Black Mirror is coming true in China, where your 'rating' affects your home, transport and social circle

- Can we avoid a **big brother** through **decentralization**?
- Can we let users **control and minimize what they share**?

Our Problem

Goal

- Study the feasibility of a **decentralized reputation system for the Internet**

Constraints

- Should be effective in **incentivizing cooperation**
- Should give users GDPR-style control on **which data to include**, and **for what**

This Work

- An exploration of
  - **design issues**
  - related **state of the art**

Direct Reciprocation

- **Tit for tat**: I will treat you as you treated me
- Makes selfish agents behave cooperatively
- Very effective when people have **repeated interactions**
- Key to the success of the BitTorrent protocol

Reputation

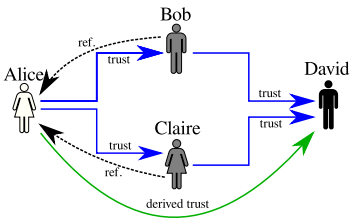

Subjective Reputation

- You trust the “friends of your friends”
  - and friends of the friends of your friends...
- Reputation travels along paths in a **web of trust**: a graph whose edges are **trust relationships**

Indirect Reciprocation

- I will behave well with those that treated my friends well
  - and the friends of my friends...
- Solves the cooperation problem when **interactions aren’t repeated**

Reputation Function

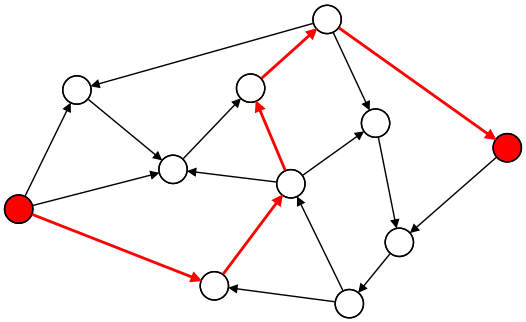

- When Alice evaluates the reputation of Bob, it's a **function of paths from Alice to Bob** in the web of trust
- Some natural choices: min path length, max flow, ...

Cheap Pseudonyms

Disposable Identities

- It's good for privacy if people can freely create different *personas* for different activities
- **Whitewashing**: people can erase a bad reputation by just using new pseudonyms

Consequences

- Introduces a **cold start** situation
  - because newcomers are indistinguishable from misbehavers
- **Cooperation can still emerge** (Friedman and Resnick, 2001)
- We should probably enable **both** persistent and disposable identities

Sybil Attack

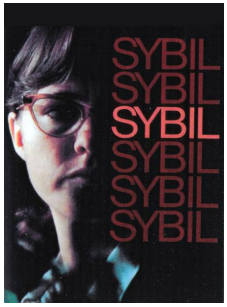

Douceur 2002

- Attack to peer-to-peer systems
- A very large number of **fake user profiles** is created to subvert the system behavior

Cheng and Friedman 2005, Dell'Amico and Capra 2010

- Several **subjective** reputation metrics are **resilient** against this attack
  - Max flow, Personalized PageRank, ...
- Indeed, reputation is a **defense** against Sybil attack

Decentralized Reputation

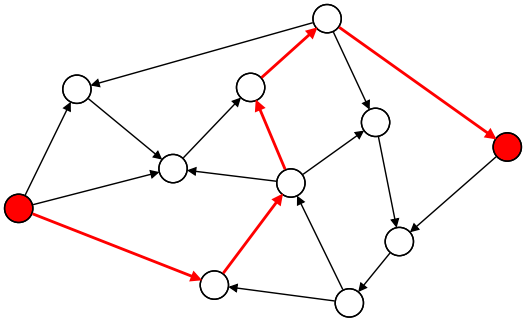

- We want to find **paths** on a web of trust—i.e., **navigate** the network
- But we want to find them **without** storing the full graph anywhere
- Can we do it on nodes that **only know their neighbors** in the network?
- Surprisingly, **several works target this problem** without acknowledging each other

Decentralized Network Embeddings

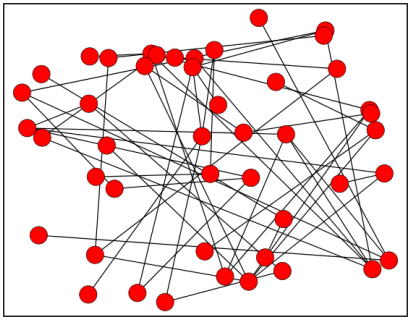

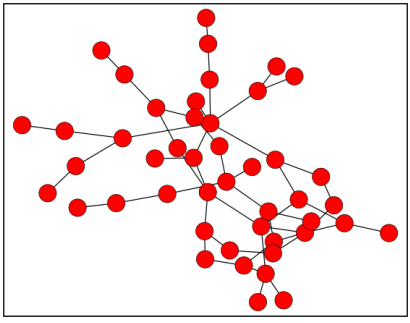

- Obtain **“coordinates”** for each node in the graph
- Navigate it with a strategy that takes the steps that bring closer to the destination
- Designed for several different use cases—implementing & comparing them is under way

Consistency

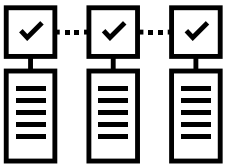

Distributed Ledgers

- Blockchains are based on **distributed ledgers** (DLs)—append-only, unmodifiable data structures that are readable and writeable by anybody
- We can use them for **non-repudiable records** of what happened, that can be used to prove misbehavior afterwards
- Ironically, the fact that writing on DLs is slow and expensive can help us against some types of Sybil attacks

(logo by James Fok from the Noun Project, CC BY 3.0)

Thank You!

- We discussed the problem of creating a **privacy-preserving reputation system for everybody on the Internet**

- While this looks like a **huge task**, plenty of work handles many related problems
- This work's contribution is to **highlight the main problems** and **point to existing solutions** in the state of the art