Web3Recommend

Decentralised recommendations with trust and relevance

Rohan Madhwal
July 10th, 2023

TUDelft

Student number:
Thesis committee: Dr. ir. J.A. Pouwelse
Dr. ir. Ujwal Gadiraju

Outline

Background
Problem Statement
Web3Recommend Design & Architecture
Experiments

The BIG Data problem

- **“The information age is drowning us with an unprecedented deluge of data”**[1]
- Amplified in Social Media Platforms - anyone can be a content creator!
- 83% of TikTok’s 1 billion monthly users have published a video[2]
- 100,000 new songs on Spotify daily
- Simple search capabilities not enough!

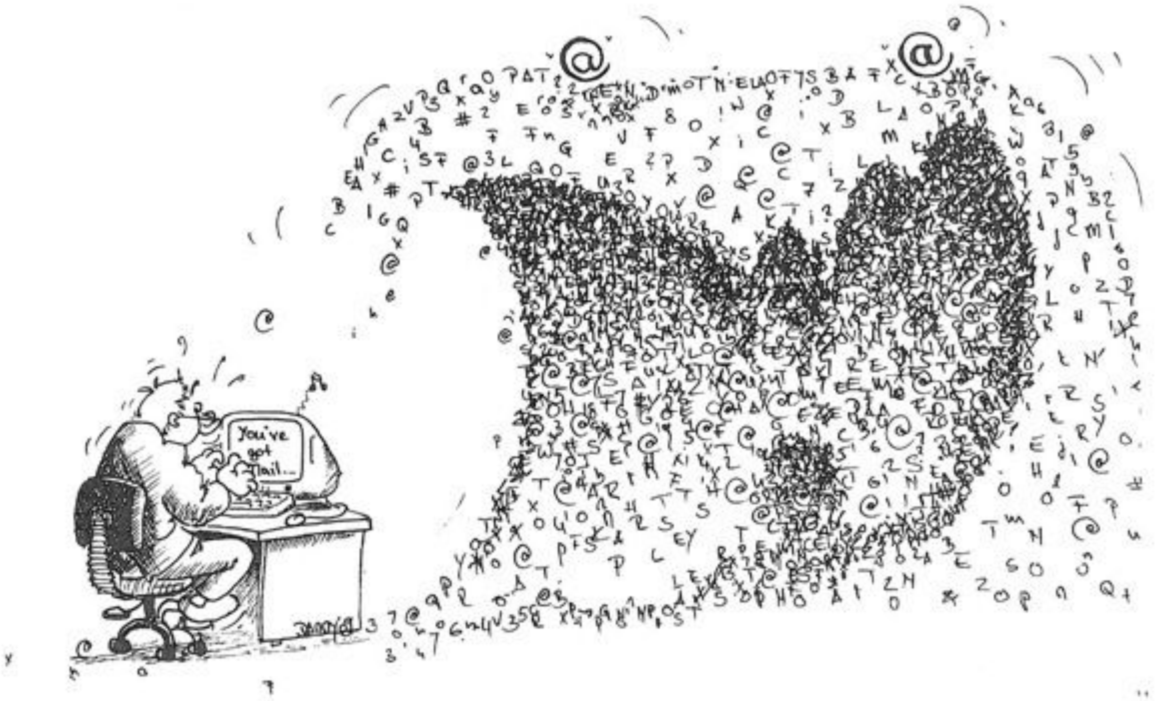

[1] Neuroscientist Daniel J. Levitin
[2] “Tiktok statistics - everything you need to know [Mar 2023 update],” Mar 2023. [Online]. Available: https://wallaroomedia.com/ blog/social-media/tiktok-statistics/

Social Recommender Systems

- Recommendations for social network users based on **personalised needs**
- Inferred through unique **implicit and explicit interactions** in the network
- Symbiotic relationship with network

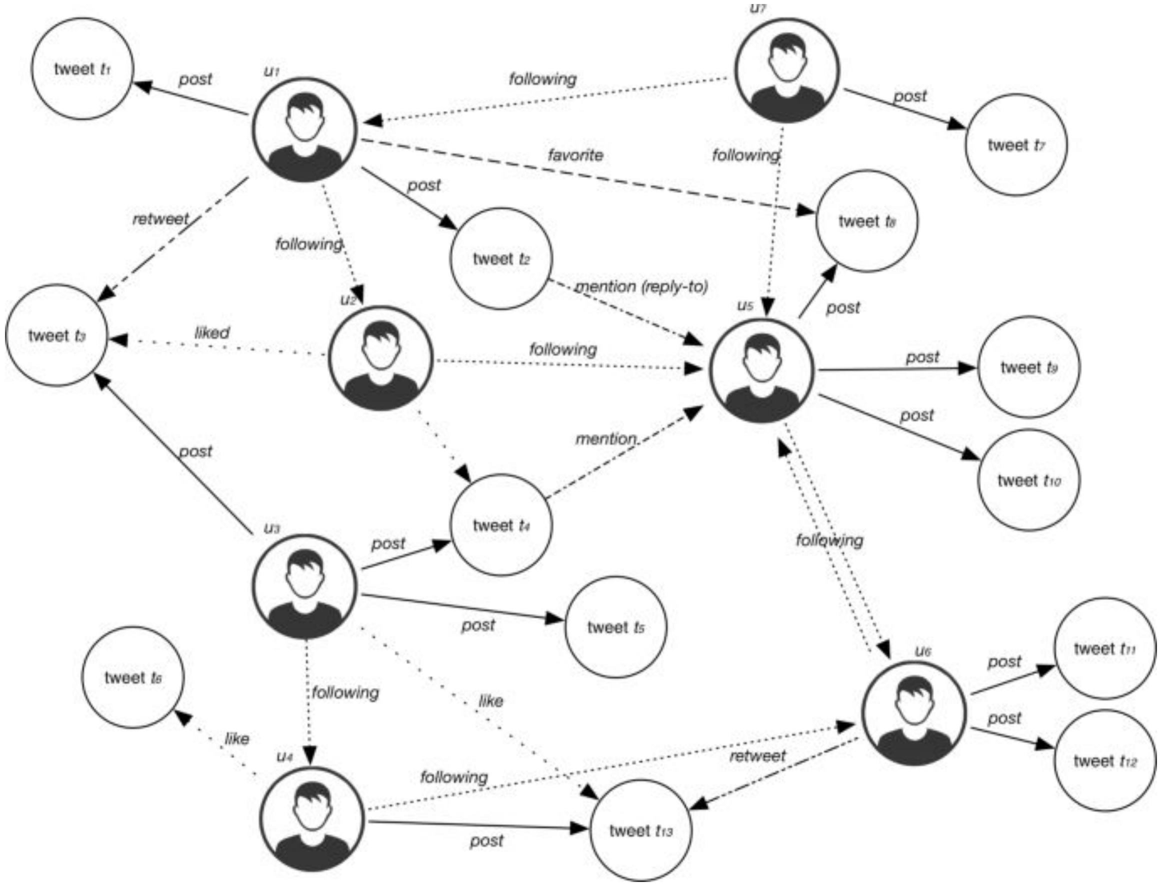

Popular Example: TikTok's "For You" Page

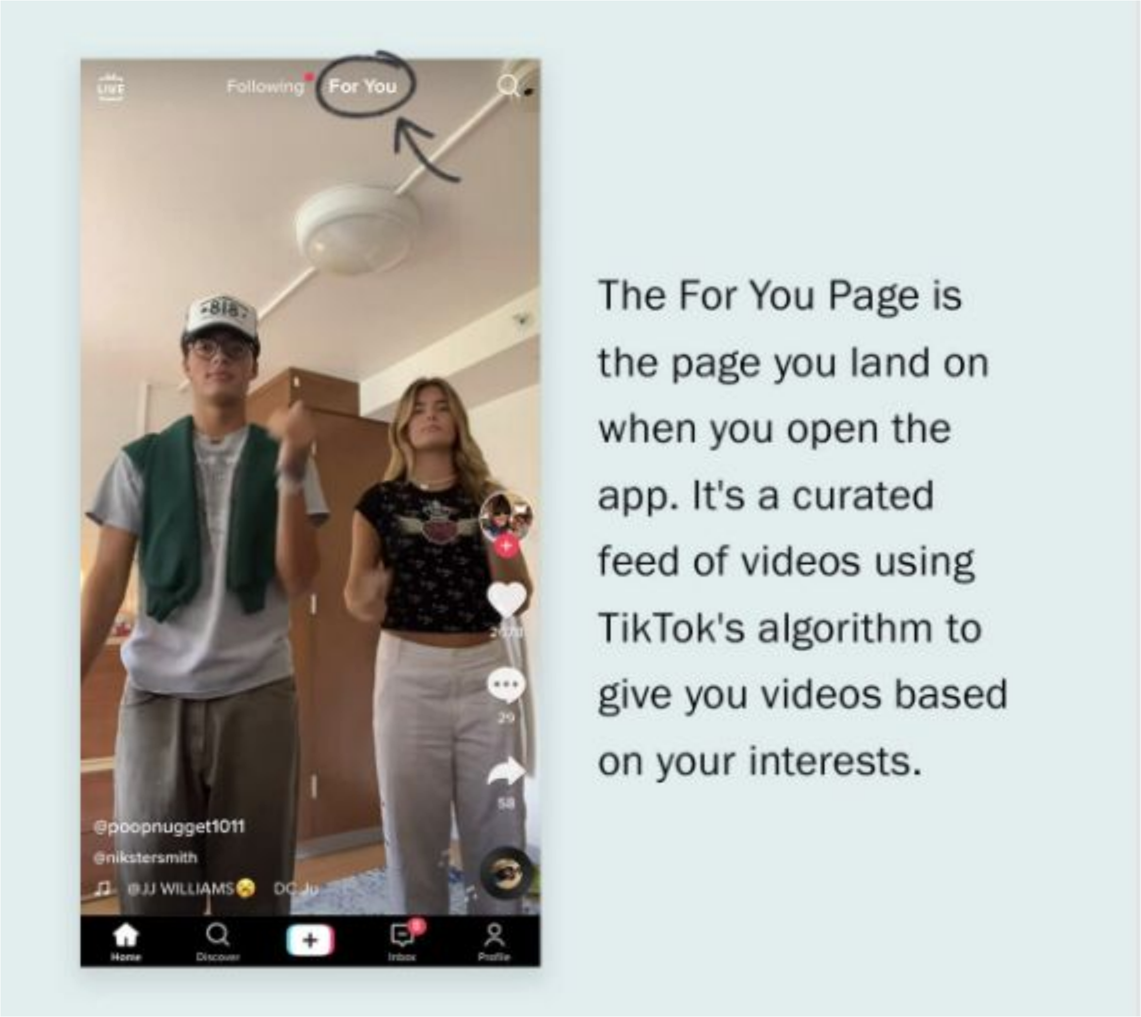

Erosion in public trust in Centralised platforms

- Conventional Social Recommender systems run on Centralised social platforms
- Facebook/Meta, TikTok, Twitter, etc.
- The last decade has witnessed a fall in public trust in these platforms

TU Delft

Intentional Violations of Trust

An official website of the United States government Here's how you know

Español | Report Fraud | Sign Up for Consumer Alerts

FEDERAL TRADE COMMISSION
PROTECTING AMERICA'S CONSUMERS

Enforcement | Policy | Advice and Guidance | News and Events | Abou

Home / Business Guidance / Business Blog

Business Blog

Twitter to pay $150 million penalty for allegedly breaking its privacy promises – again

Facebook users raise privacy complaints over tracking for marketing

By Anick Jesdanun and Rachel Metz
Associated Press

NEW YORK — Some users of the online hangout Facebook are complaining that its two-week-old marketing program is publicizing their purchases for friends to see.

Those users say they never noticed a small box that appears on a corner of their friends' activities through the feeds. About 40 Web sites have decided to embed a free tool from Facebook, known as a Beacon, to enable the marketing feeds.

The idea is that if users see a friend buy or do something, they'd take that action as an endorsement for a movie, a band or a soft drink.

But it also raises privacy concerns.

Mike Mayer, for instance, ... inquiries to Facebook, which issued a statement defending its practices. Facebook officials have also said advertising supports the free service.

"Beacon gives users an easy way to share relevant information from other sites with their friends on Facebook," the statement said. "Information is shared with a small selection of a user's trusted network of friends, not publicly on the Web or with all Facebook users.

... ting companies use names for endorsements without "explicit permission."

"We want Facebook to realize that their users are rightly concerned that private information is being made public," MoveOn spokesman Adam Green said, adding that Facebook could quell concerns by seeking "opt in" consent rather than leaving it to users to "opt out" by taking steps to decline sharing.

Australia's privacy watchdog to enter talks with Facebook owner over Cambridge Analytica lawsuit

Federal court orders commissioner and Meta to start mediation to end protracted, costly legal proceedings

- Follow our Australia news live blog for the latest updates
- Get our morning and afternoon news emails, free app or daily news podcast

Facebook's owner Meta has been ordered to enter mediation with the Office of the Australian Information Commissioner over legal proceedings related to the Cambridge Analytica scandal. Photograph: Jaap Arriens/NurPhoto/Shutterstock

The New York Times

Why Countries Are Trying to Ban TikTok

Governments have expressed concerns that TikTok, which is owned by the Chinese company ByteDance, may endanger sensitive user data.

Give this article

TikTok has long denied allegations that it puts sensitive user data into the hands of the Chinese government. Valerie Macon/Agence France-Presse — Getty Images

Unintentional Violations of Trust

WIRED BACKCHANNEL BUSINESS CULTURE GEAR IDEAS SCIENCE SECURITY PRIME DAY

What Twitter's 200 Million-User Email Leak Actually Means

The exposure of hundreds of millions of email addresses puts pseudonymous users of the social network at risk.

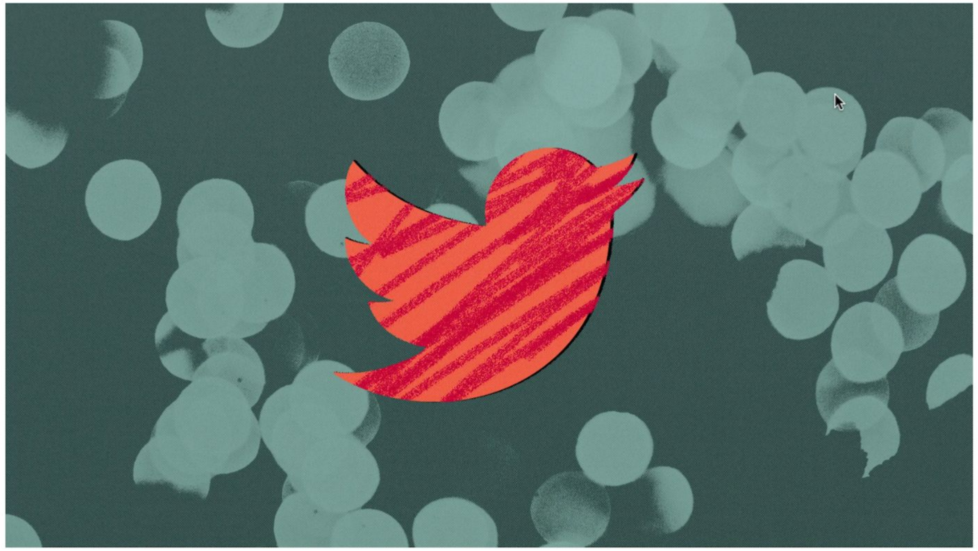

ILLUSTRATION: ROSIE STRUVE; GETTY IMAGES

META / TECH / PRIVACY

Meta fined $276 million over Facebook data leak involving more than 533 million users

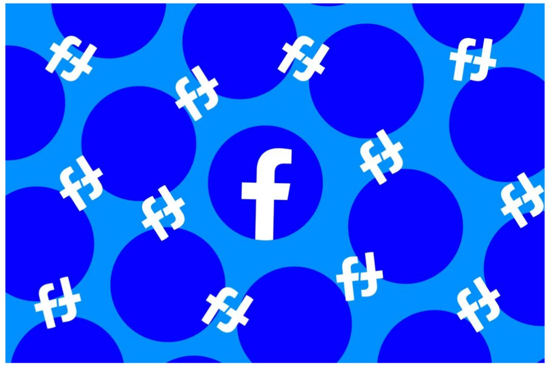

Illustration by Nick Barclay / The Verge

/ The April 2021 leak exposed the phone numbers, locations, and birthdates of Facebook users on the platform from 2018 to 2019.

By Emma Roth, a news writer who covers the streaming wars, consumer tech, crypto, social media, and much more. Previously, she was a writer and editor at MUO.

Nov 28, 2022, 4:02 PM GMT+1 | 6 Comments / 6 New

BENZINGA

Google To Cough Up $392M For User Privacy Breach

Anusuya Lahiri

November 15, 2022 · 2 min read

TU Delft

Rise of Decentralised Technologies and Web3

- Direct interactions between users without third-party intermediation
- **Web3** platforms promise trusted alternatives to profit-driven institutions
- Leverage communal infrastructure and participant resources

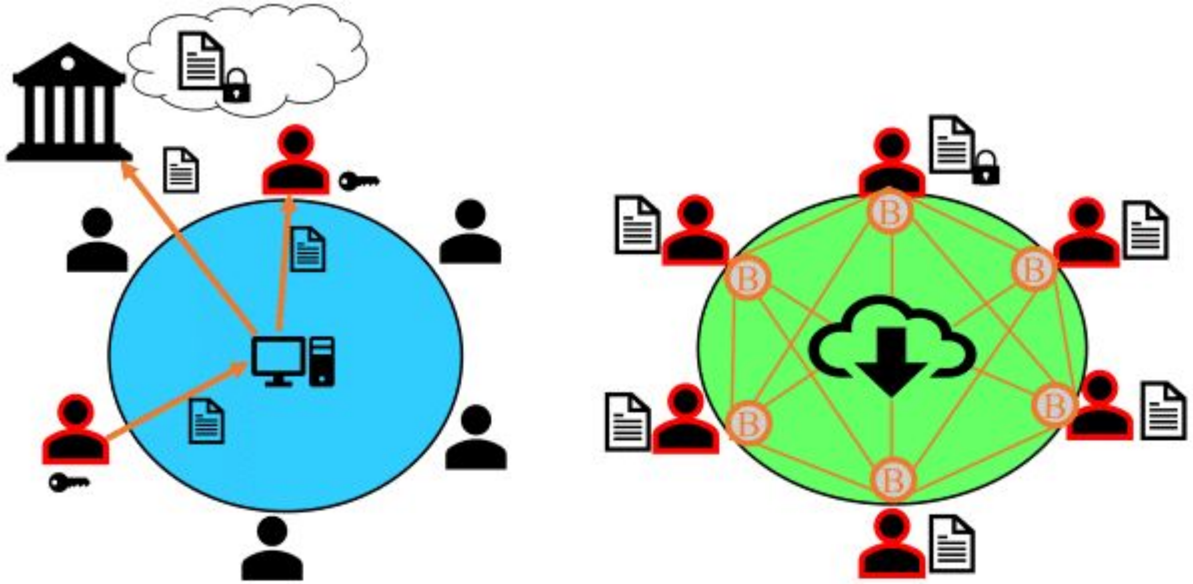

Centralized Framework   Decentralized Framework

First Major Challenge - Web3 Social Recommender System: Lack of a Global Perspective

- In any single node, difficult to have holistic view of the rest of the network
- Require more design, planning, and management
- Lack of any leader
- No *decentralized police* or *bug fix authority*
- No room for ad-hoc decisions, everything needs to be decided upfront

Second Major Challenge - Web3 Social Recommender System: Sybil Attack (1)

- No central authority to perform identity verification or monitoring
- Pseudonymity/Anonymity is a feature of many Web3 Platforms
- Attackers can effortlessly create potentially unlimited fake identities (Sybils)

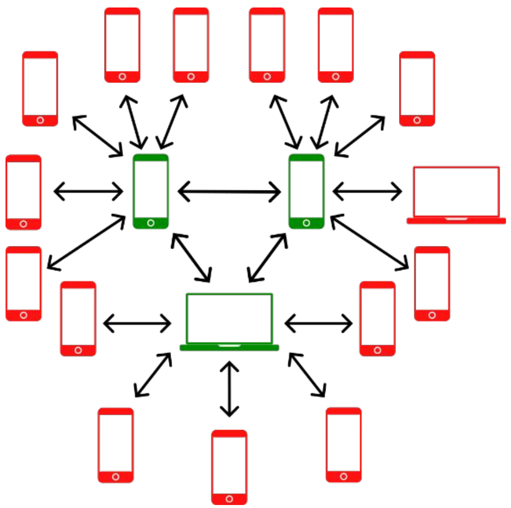

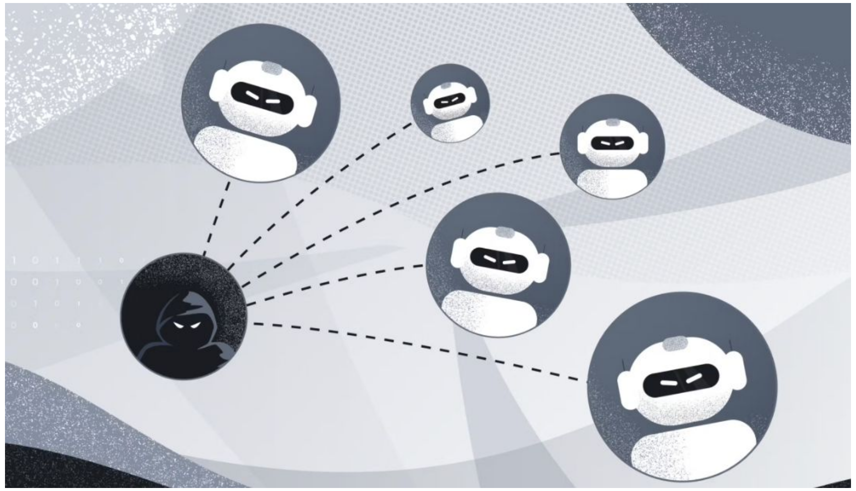

TU Delft

Second Major Challenge - Web3 Social Recommender System: Sybil Attack (2)

- In a Social Recommender System, a user's prior experience is a *vote* for item
- Thus, attacks allow manipulation of a naive Social Recommender System
- Ensure that recommender recommends their *sybil items*
- Leads to Spam/Malicious content
- Places a burden on the network's resources

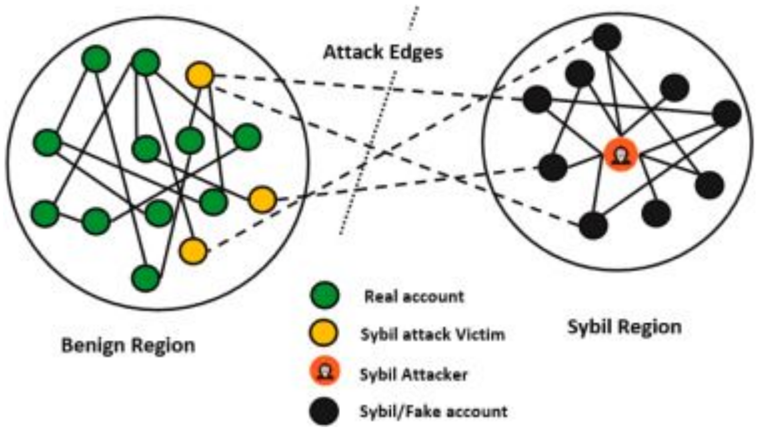

Problem Description

Generating decentralised, globally-informed social recommendations for Web3 platforms that are tolerant to adversarial Sybil attacks

Hypothesis

Generating decentralised, globally-informed social recommendations for Web3 platforms that are tolerant to adversarial Sybil attacks

⬇

A decentralised implementation of a random walk based Social Recommender System with Sybil resistance added to recommendations

Validation

Generating decentralised, globally-informed social recommendations for Web3 platforms that are tolerant to adversarial Sybil attacks

A decentralised implementation of a random walk based Social Recommender System with Sybil resistance added to recommendations

Validate relevance and sybil tolerance of recommendations through experiments on real world data

TU Delft

Deployment as Proof of Principle

Generating decentralised, globally-informed social recommendations for Web3 platforms that are tolerant to adversarial Sybil attacks

↓

A decentralised implementation of a random walk based Social Recommender System with Sybil resistance added to recommendations

↓

Validate relevance and sybil tolerance of recommendations through experiments on real world data

↓

Deploy solution by integrating it with an existing Web3 Platform

Web3Recommend Solution Overview

- **Sybil-tolerant** Decentralised Social Recommender System
- Acts on recent interactions in network
- Real time recommendations with tight resource bound
- Functionality tested with Unit/Integration tests in JUnit
- Proof of principle: MusicDAO Deployment
- All code, experiments and tests are open-source

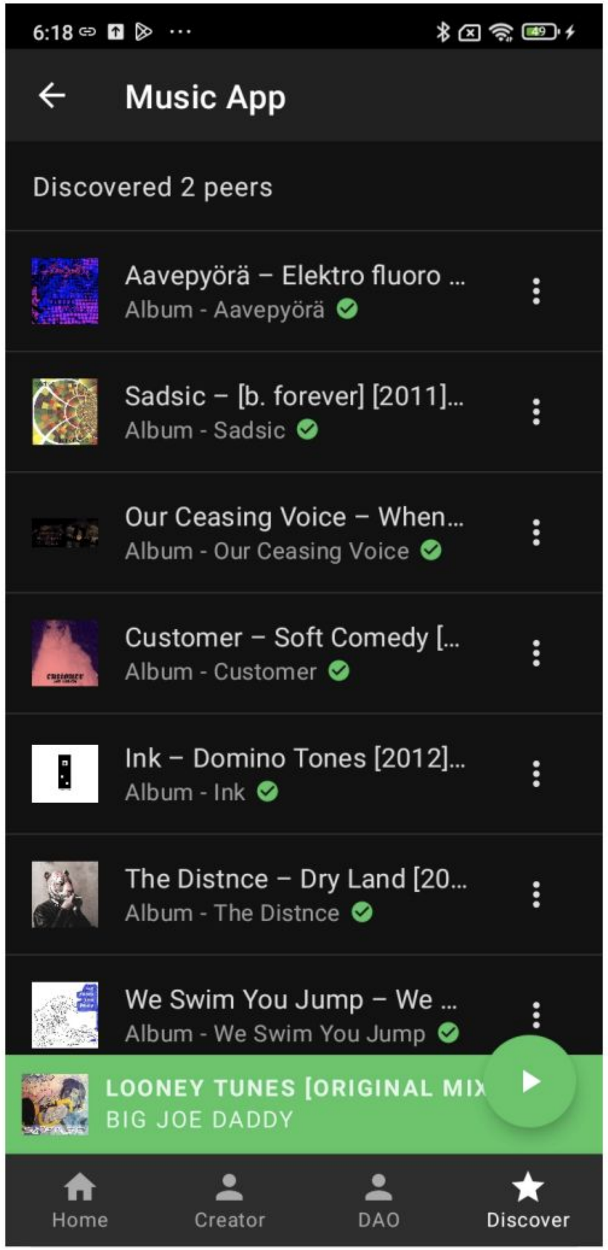

TU Delft

Background: **GraphJet**

- Graph-based system for generating real-time tweet recommendations on Twitter[1]
- Single server implementation
- Maintains and updates bipartite graph by keeping track of user-tweet interactions over the most recent *n* hours
- Based on a **personalized SALSA algorithm**, which involves random walks in a bi-partite graph of users and tweets

TU Delft

[1] Sharma, Aneesh, et al. "GraphJet: Real-time content recommendations at Twitter." *Proceedings of the VLDB Endowment* 9.13 (2016): 1281-1292.

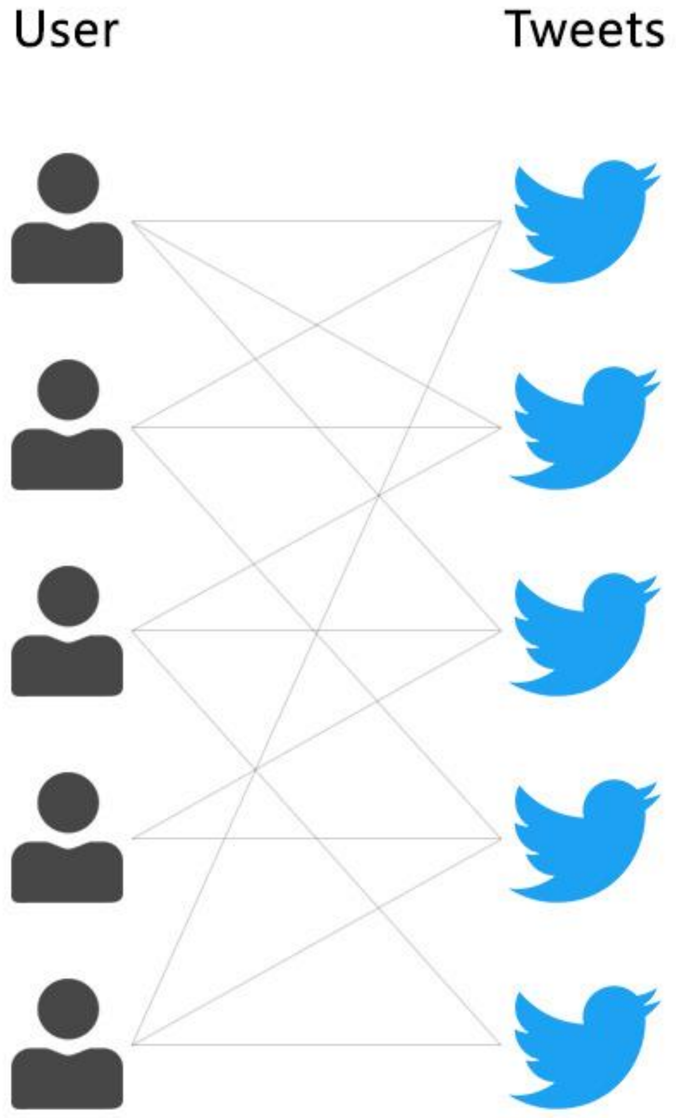
User
Tweets

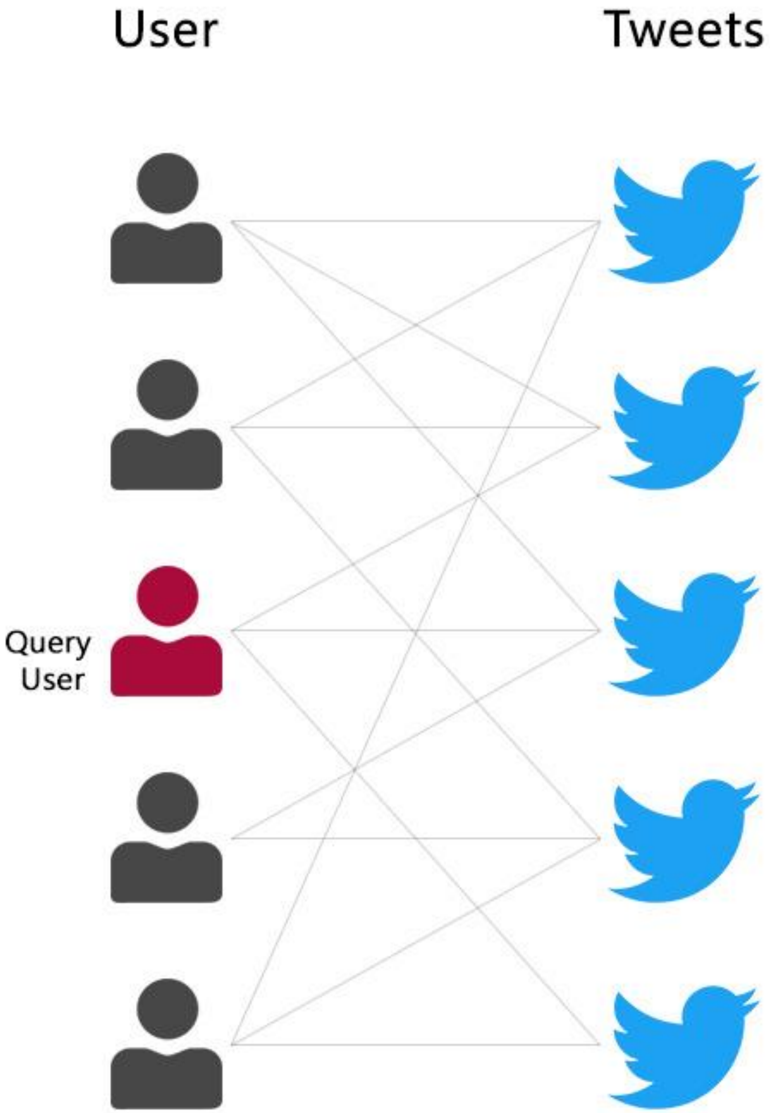
User
Tweets
Query
User


User
Tweets
Start random walk

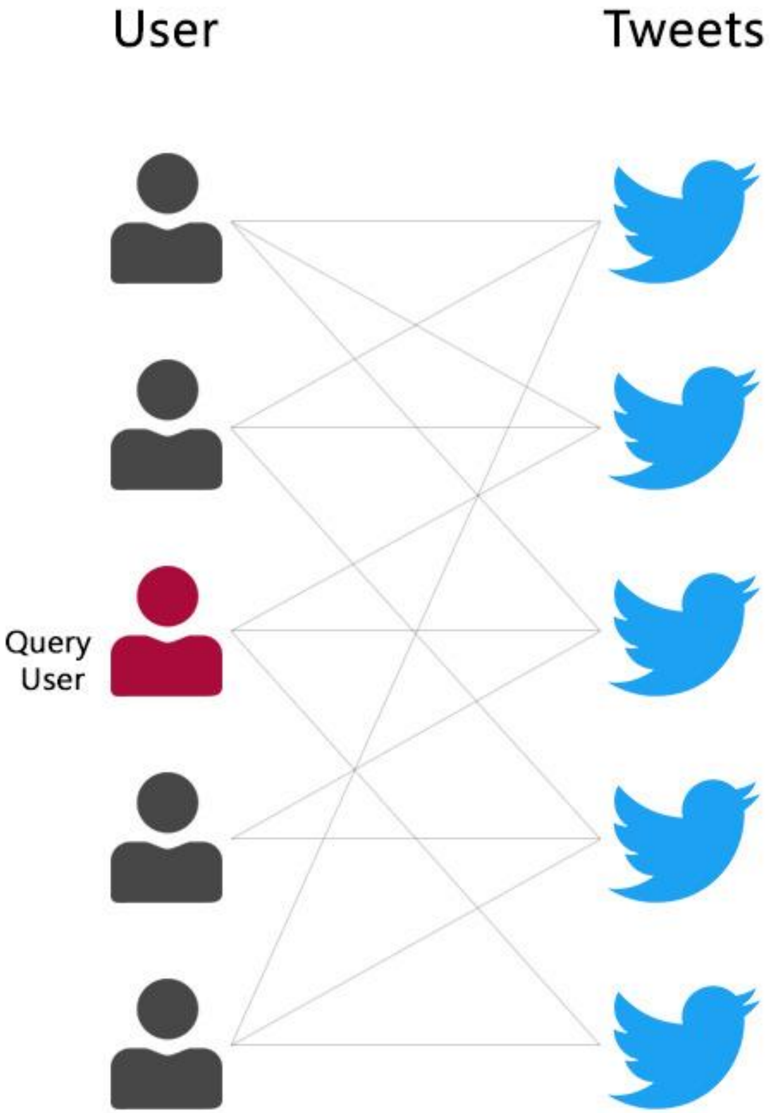
User
Tweets
Query
User

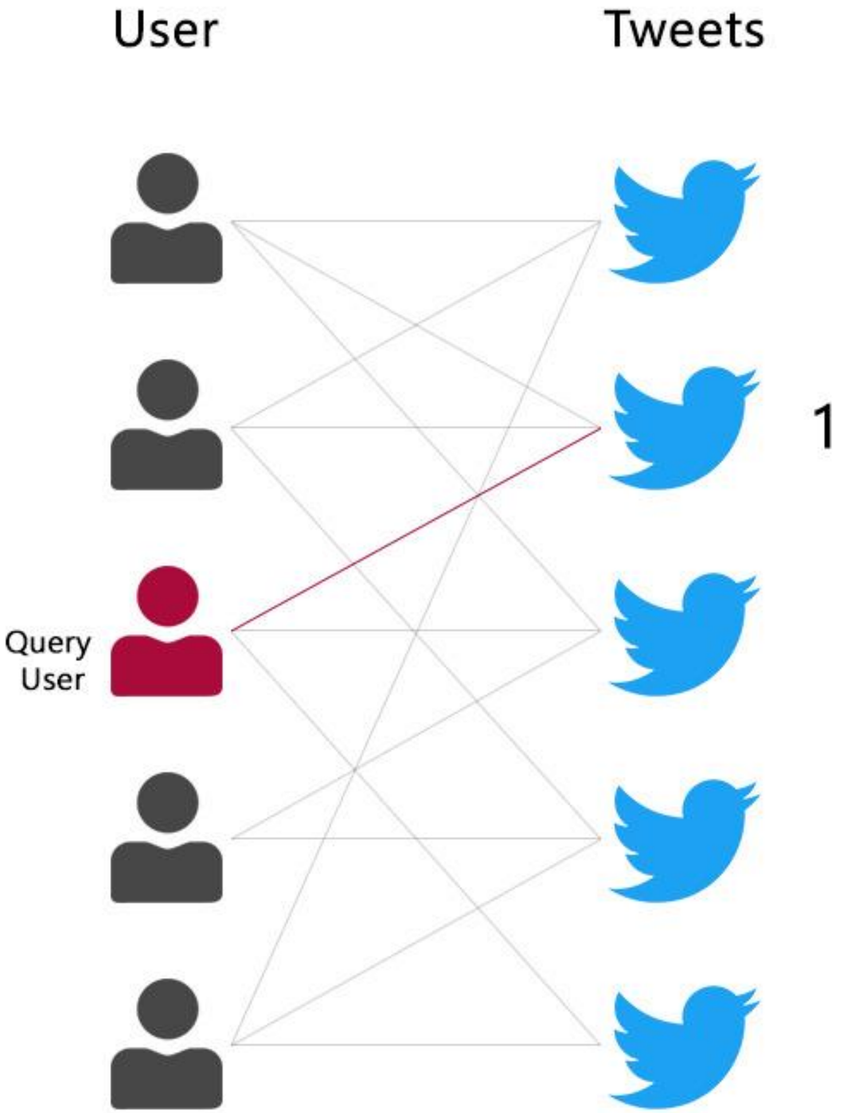
User
Tweets
Query
User
1

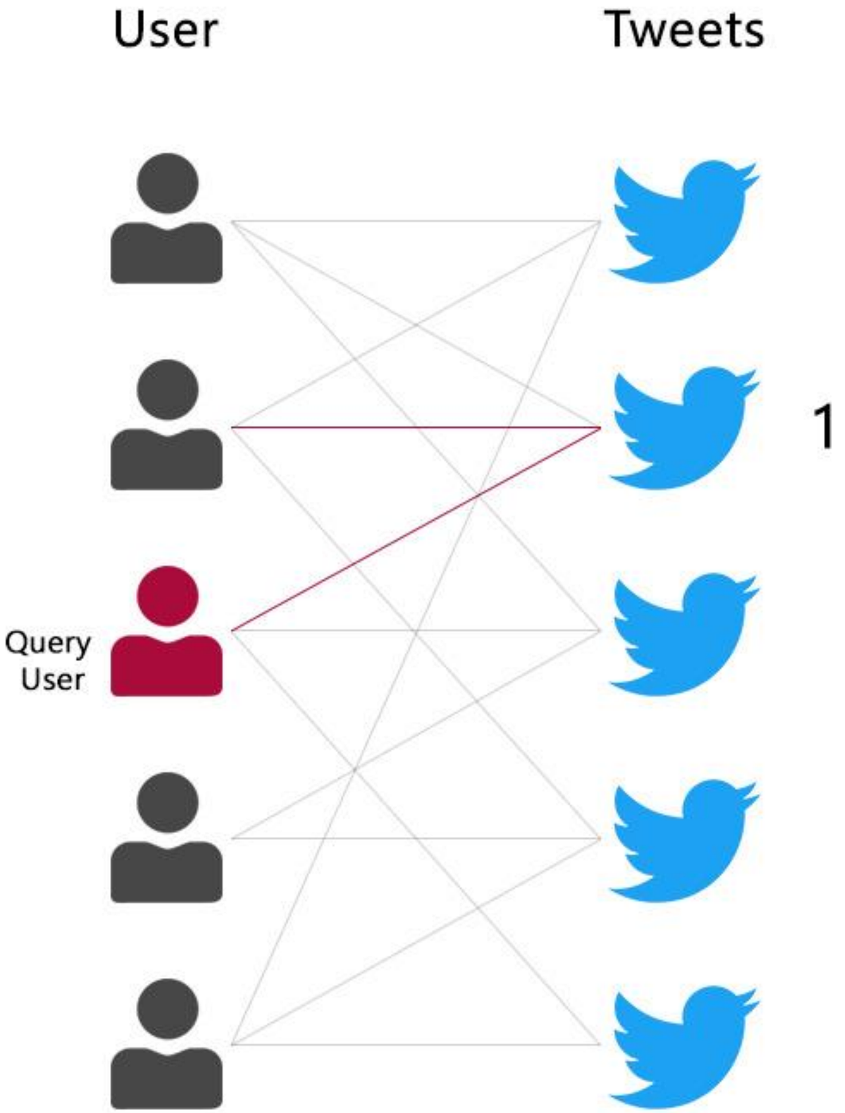
User
Tweets
Query
User
1

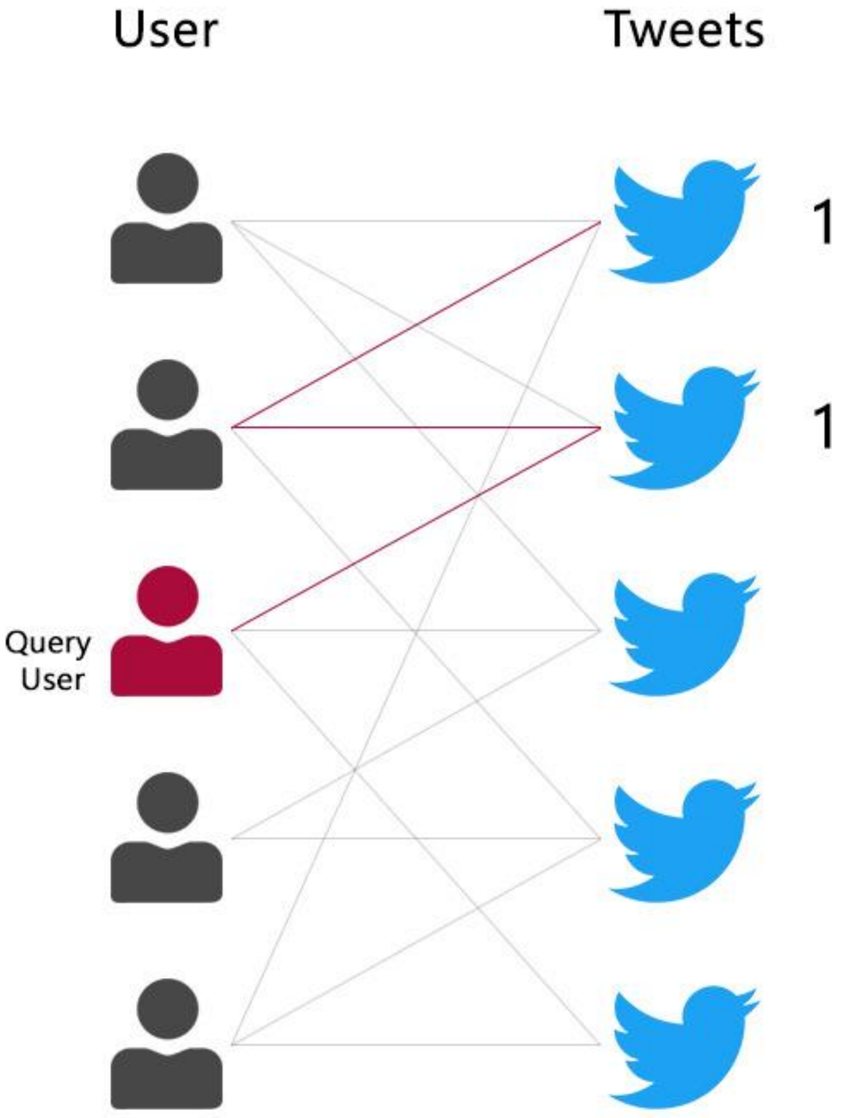
User
Tweets
Query
User
1
1

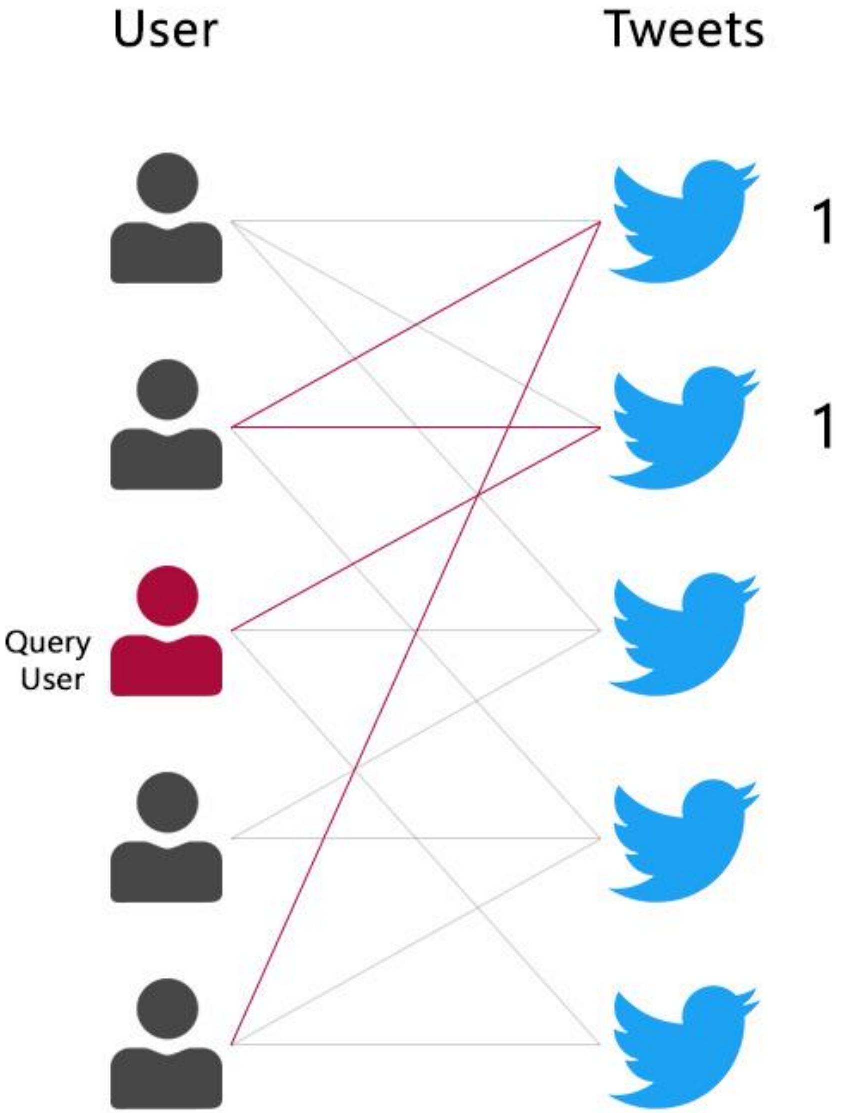
User
Tweets
1
1
Query User

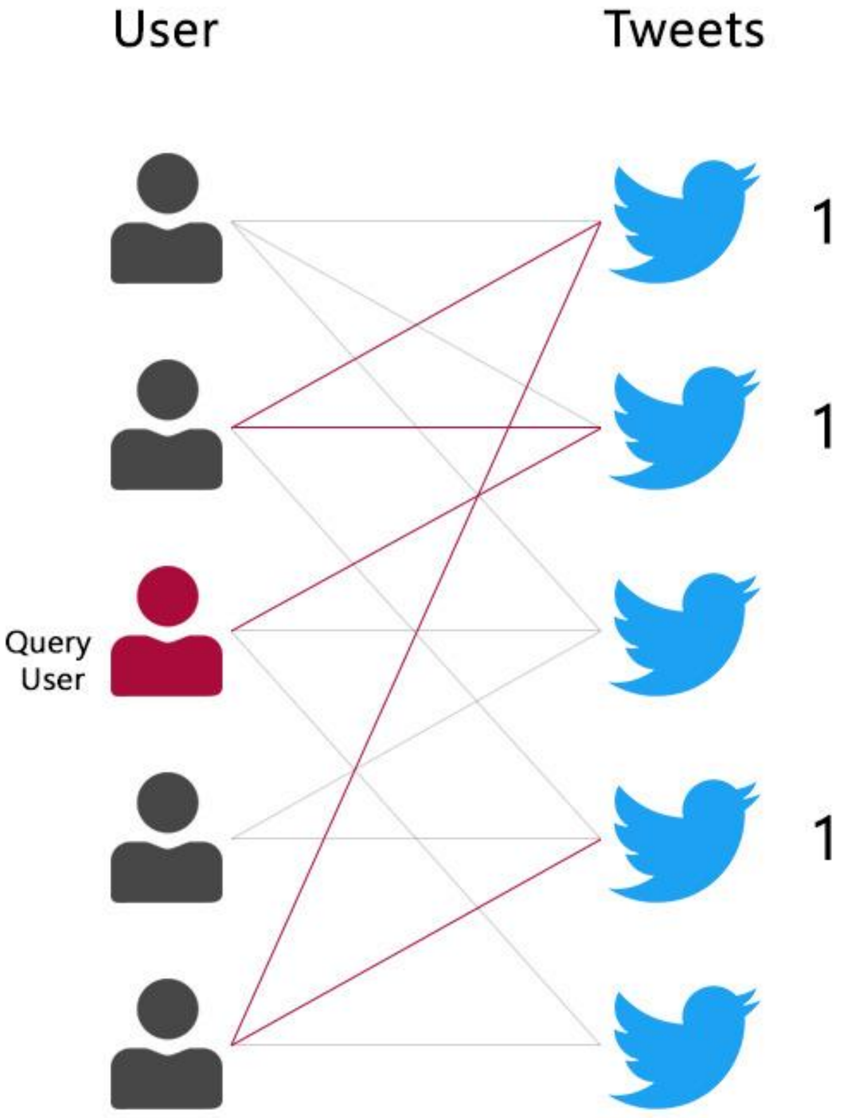
User
Tweets
1
1
Query
User
1

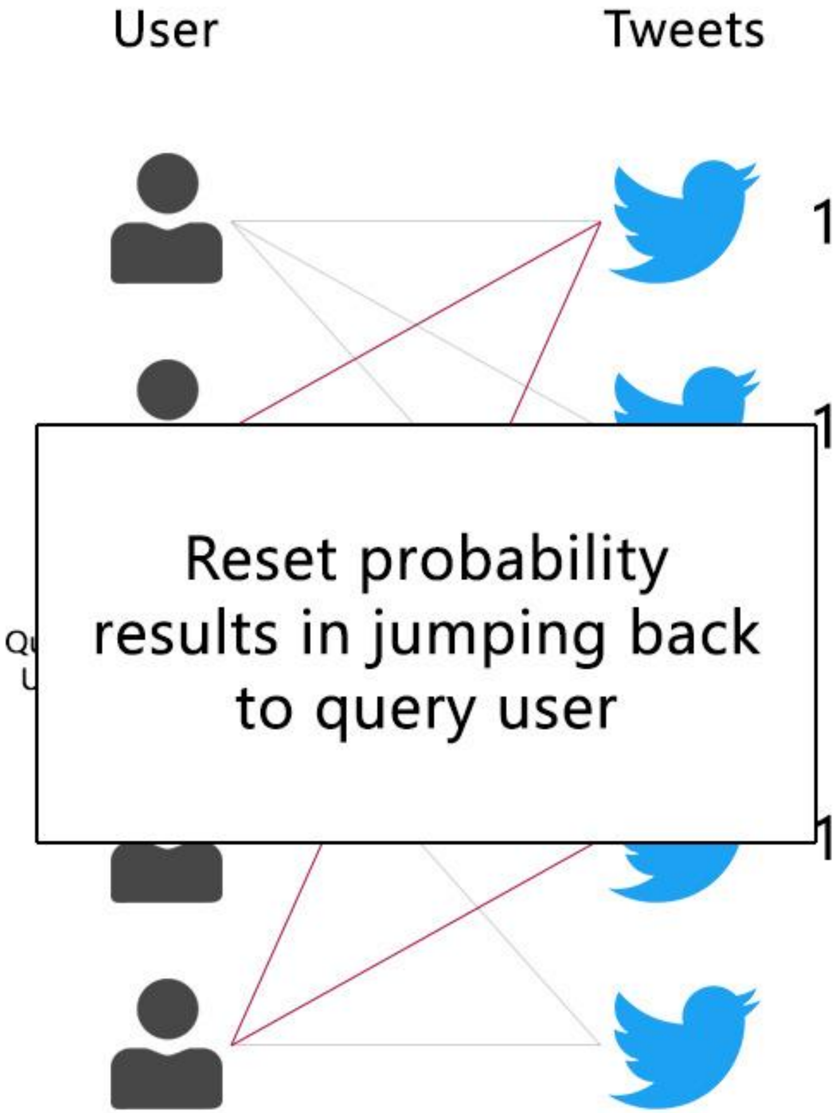
User
Tweets
1
1
1
Reset probability results in jumping back to query user

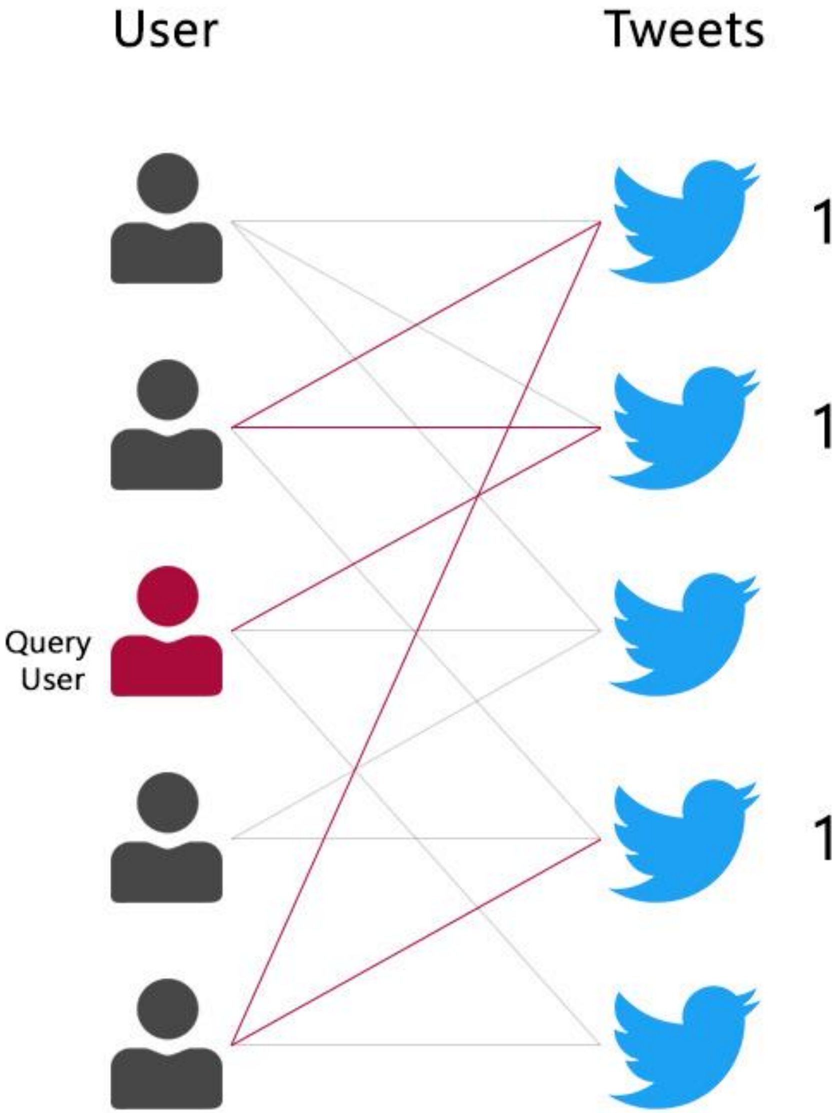
User
Tweets
1
1
Query
User
1

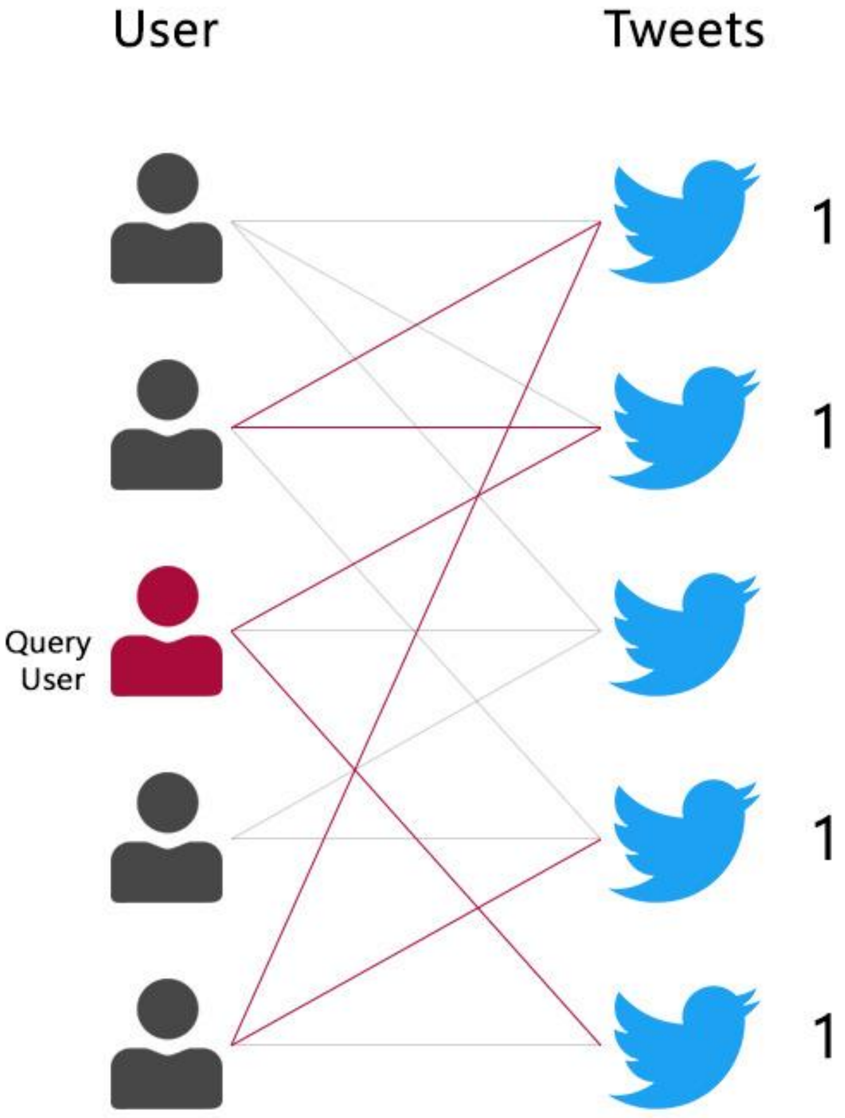
User
Tweets
1
1
Query
User
1
1

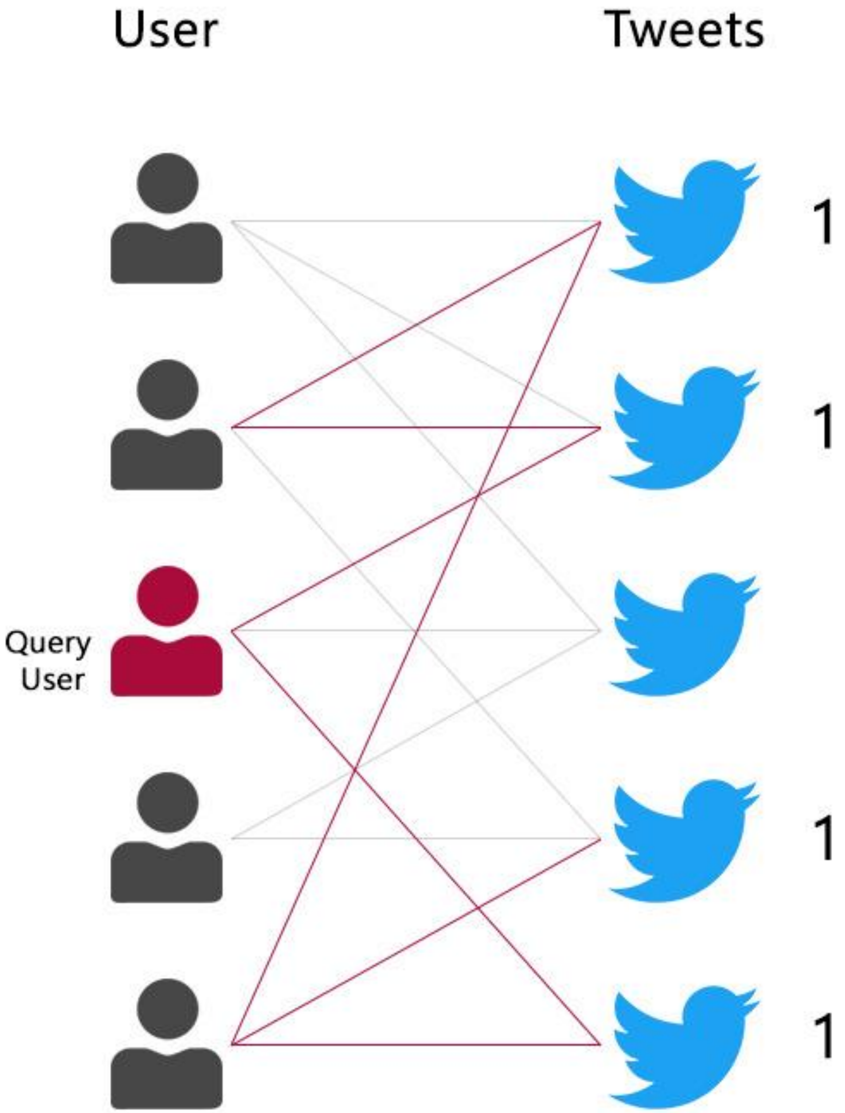
User
Tweets
1
1
Query User
1
1

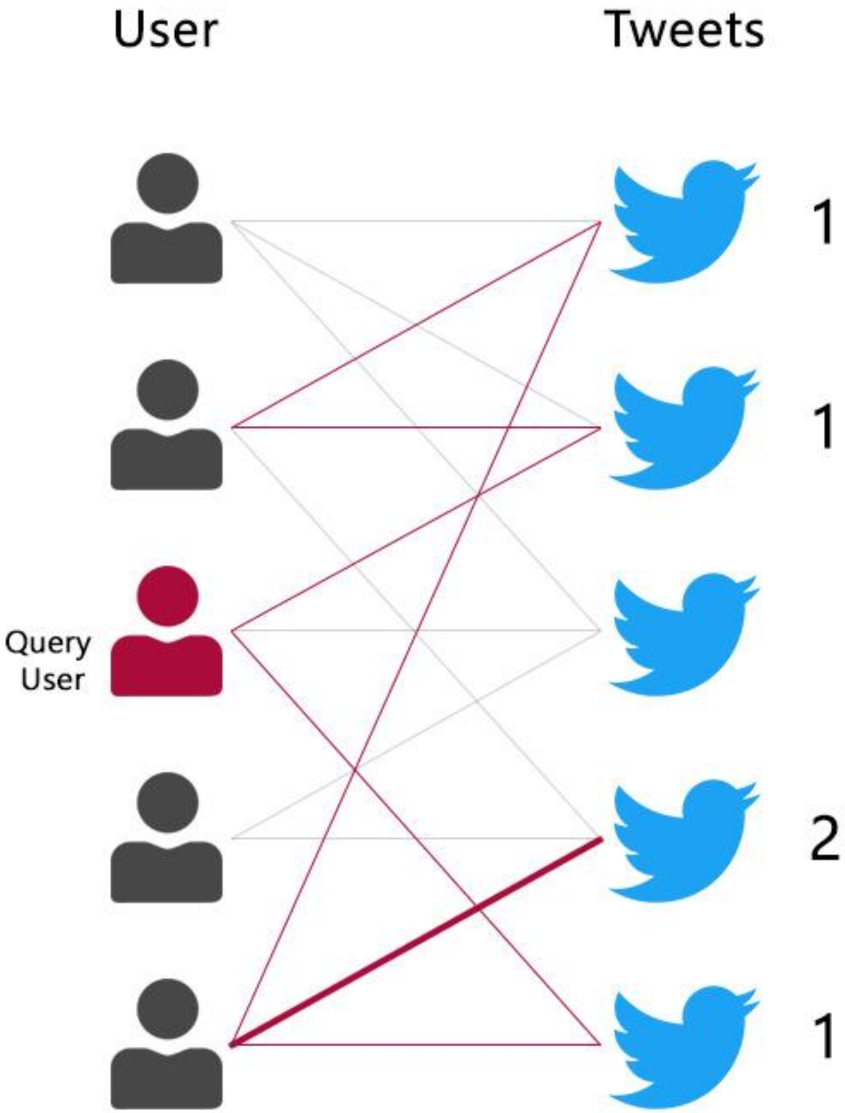
User
Tweets
Query User
1
1
2
1

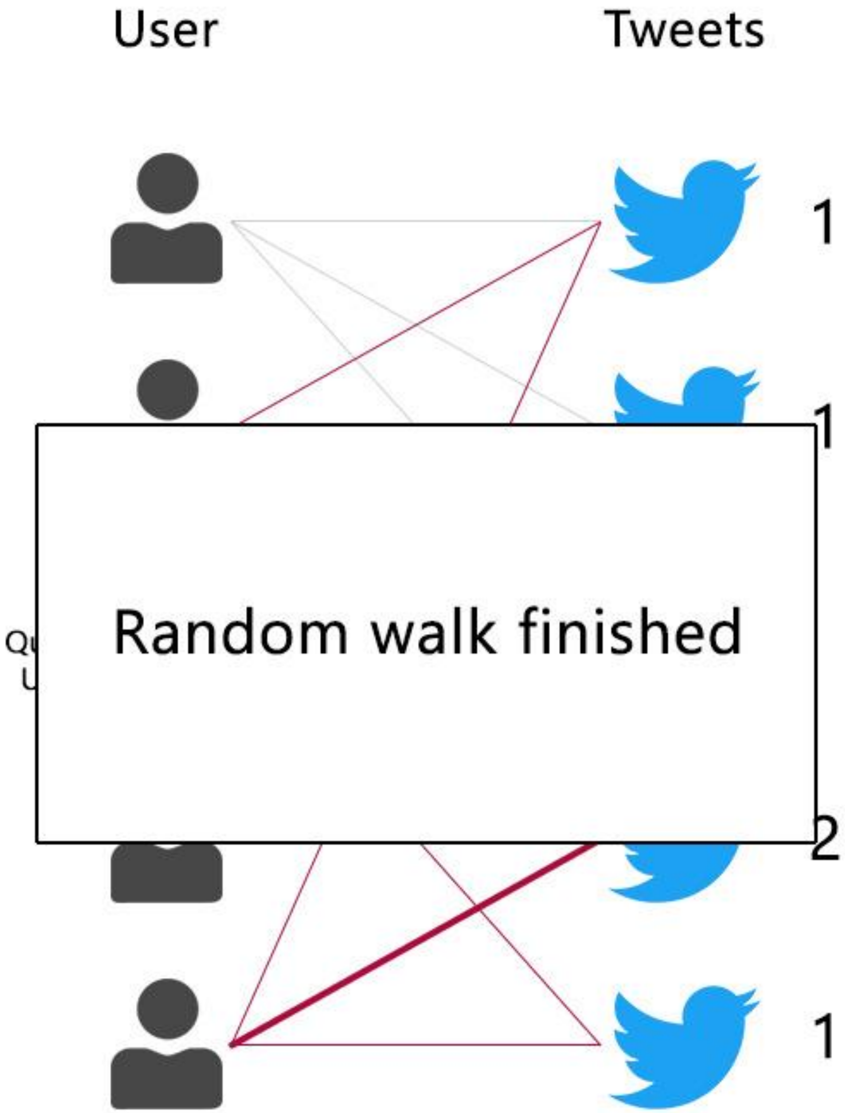
User
Tweets
1
1
Random walk finished
2
1

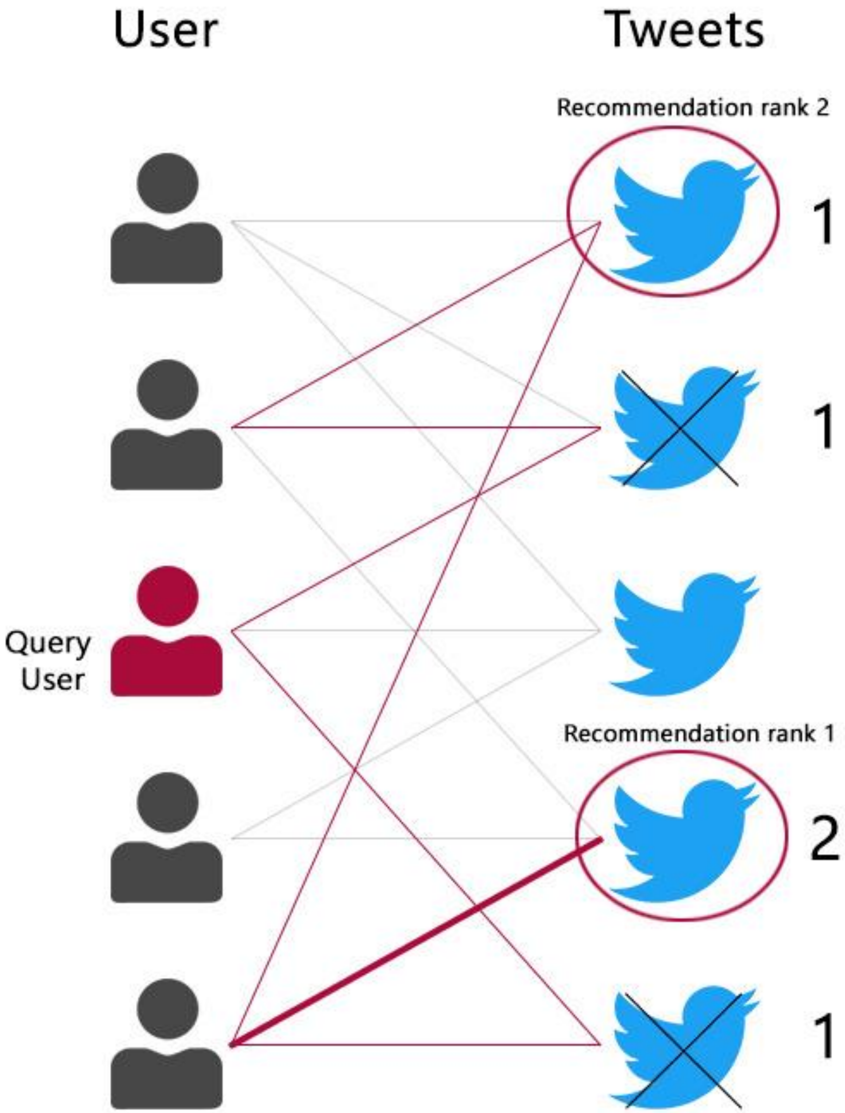
User
Tweets
Recommendation rank 2
1
1
Query User
Recommendation rank 1
2
1

Background: **MeritRank**

- In the context of open, permissionless systems, complete elimination of Sybil Attacks is not possible[1]
- Further, emulating Sybil resistant properties of closed systems undermines privacy and other desirable properties of decentralized solutions
- Instead, MeritRank provides guidelines for Sybil tolerance in feedback aggregation through random walk **decays**

TU Delft

[1] Nasrulin, Bulat, Georgy Ishmaev, and Johan Pouwelse. "Meritrank: Sybil tolerant reputation for merit-based tokenomics." *2022 4th Conference on Blockchain Research & Applications for Innovative Networks and Services (BRAINS)*. IEEE, 2022.

Background: MeritRank **Sybil Attack Strategies** in Random Walks

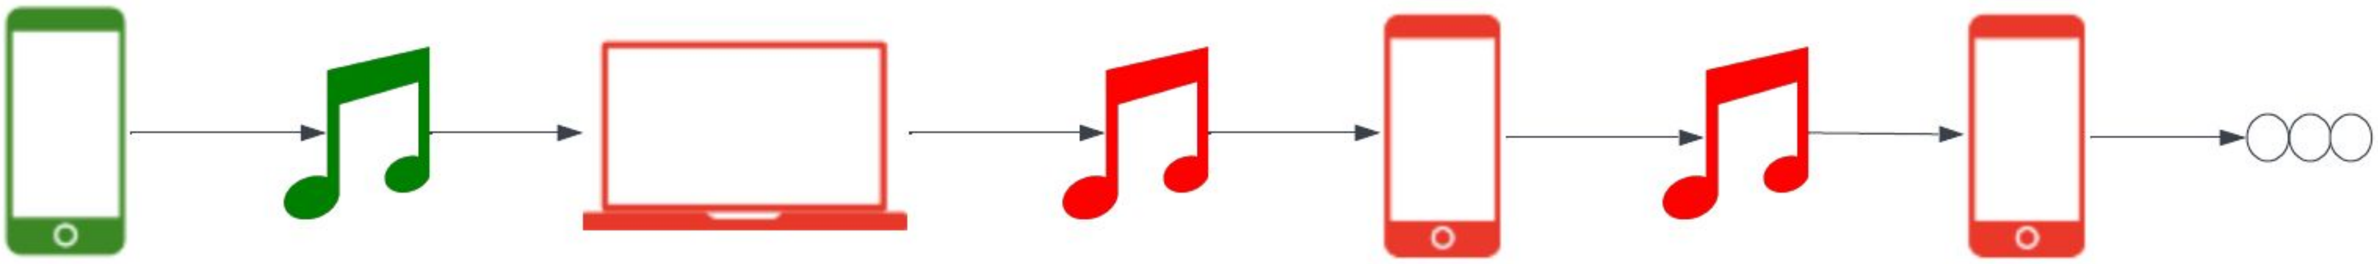

Serial Attack

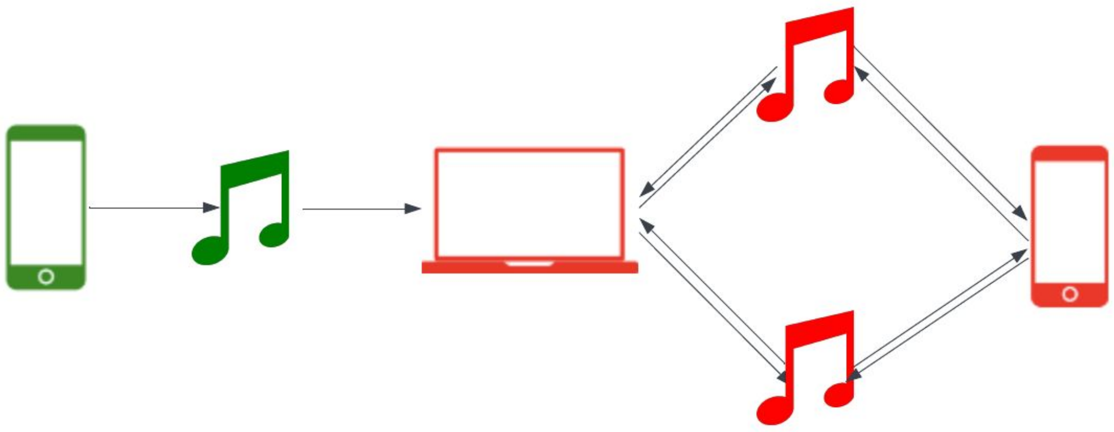

Parallel/Cyclic Attack

Background: **MeritRank Decays**

1. **Transitivity α decay**

   Limits random walks length

2. **Connectivity β decay**

   Introduces punishment for a node for being in a separate component/island

3. **Epoch γ decay**

   Prevent exploitation of old trust edges

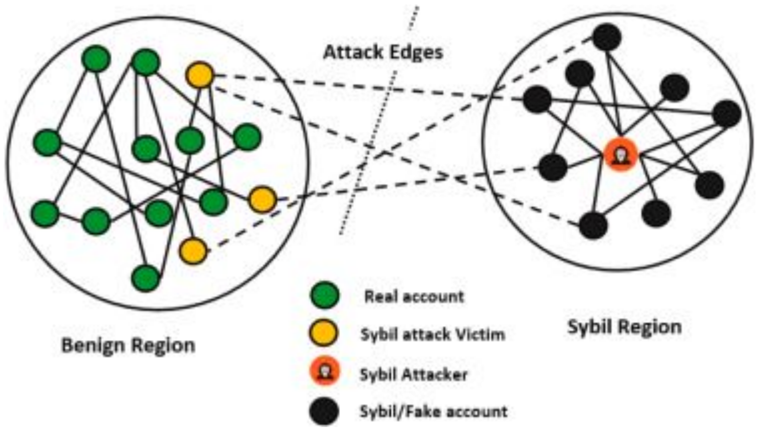

TU Delft

Web3Recommend Architectural Components

- Central data Structure: **TrustNetwork**
- **Time-biased edge gossiping mechanism**
- Modified **Personalized SALSA for** recommendations
- Modified **Personalized PageRank** for global user trust estimation
- **Compact Serialization Techniques**
- **Bootstrap mechanisms**

TrustNetwork

Implemented using a combination of two data structures:

1. **User To User Network**

   Weighted directed acyclic graph of all user

   Edges represent *trust relationships*

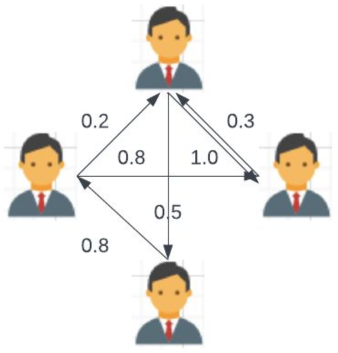

2. **User To Item Network**

   Weighted undirected acyclic bi-partite graph of users & items

   Edges represent *affinity relationships*

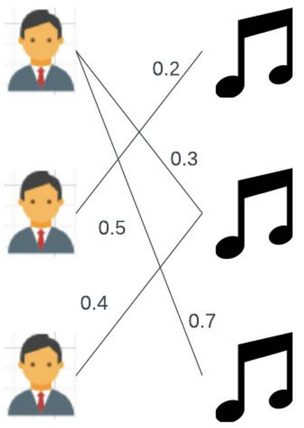

Recommendation Algorithm

Four modifications on the GraphJet Personalized SALSA algorithm:

1. **Weighted Random Walks**
2. **Incremental Random Walks[1]**
3. **MeritRank Decays to limit influence of Sybil Attacks**
4. **“Trusted” SALSA Random Walks**

TU Delft

[1] B. Bahmani, A. Chowdhury, and A. Goel, “Fast incremental and personalized pagerank,” arXiv preprint arXiv:1006.2880, 2010

Beta Decay Calculation

- Measure the diversity of users voting for an item
- Punish items whose random walks always include the same users
- Punishment for item i, is beta item decay b[i]:

$$b[i] = \begin{cases} 1-\beta & \text{if } \exists u \in U : div(u,i) > \tau \\ 1 & \text{else} \end{cases}$$

Where div(u,i) is a measure of “sybilness” of user u on item i:

$$div(u,i) = \frac{\sum_{r \in R(i)} \begin{cases} 1 & \text{if } (u \in r) \cap (r[u] < r[i])) \\ 0 & \text{else} \end{cases}}{|R(i)|}$$

TU Delft

Personalized Ranking Score Calculation

The score s[i] for item i is calculated as:

$$s[i] = \frac{|R(i)|}{\sum_{x \in I} |R(x)|} b[i]$$

Items are then ranked by their score and presented as **recommendations**

Bootstrap Mechanisms

1. **Circle of trust**

   Start random walks from “seed set” instead of the source node

2. **New User**

   Improved User Collaborative Filtering based on [1]

TU Delft

[1] B. Zhang and B. Yuan, “Improved collaborative filtering recommendation algorithm of similarity measure,” in AIP Conference Proceedings

Compact Serialization

- Optimisation for devices with limited resources
- Based on the format used in the 2nd DIMACS challenge

```
c
c SOURCE: Generated using a Custom Graph Exporter
c
p nodeToSong 4 3
n 1 someNode 0.0
n 2 randomNode 0.3
s 3 someTorrentHash 0.5
s 4 anotherTorrentHash 0.8
e 1 3 0.5 1585451228000
e 2 3 0.3 0
e 1 4 0.4 1
```

Experimental Setup

- To evaluate Web3Recommend, a network of non-Sybil users was needed
- Dataset from the taste profile subset of the Million Song Dataset was chosen
- Sourced from The Echo Nest, an online resource that provides music applications on the web, smartphones, etc
- Provides us with real world users, songs and user-song-playcount triplets

Total Users	50000
Total Edges	201114
Max UtU Edges	5
Min UtU Edges	0
Avg UtU Edges	4.02228

a) User to User Graph

Total Users	50000
Total Songs	386213
Total Edges	659962
Max UtS Edges	53
Min UtS Edges	0
Avg UtS Edges	13.19924

b) User to Song Graph

Leave out one Cross Validation

Results

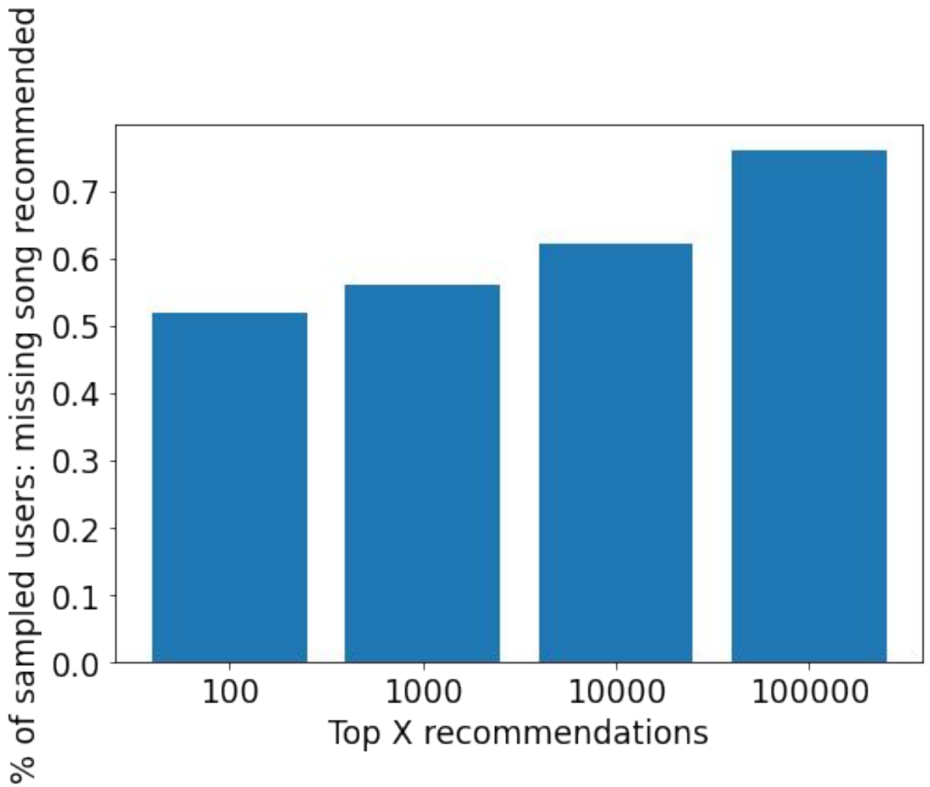

Leave out one experiment with α = 0.1 and β = 0.0

TU Delft

Leave out one Cross Validation : Comparison to Personalized SALSA

Results

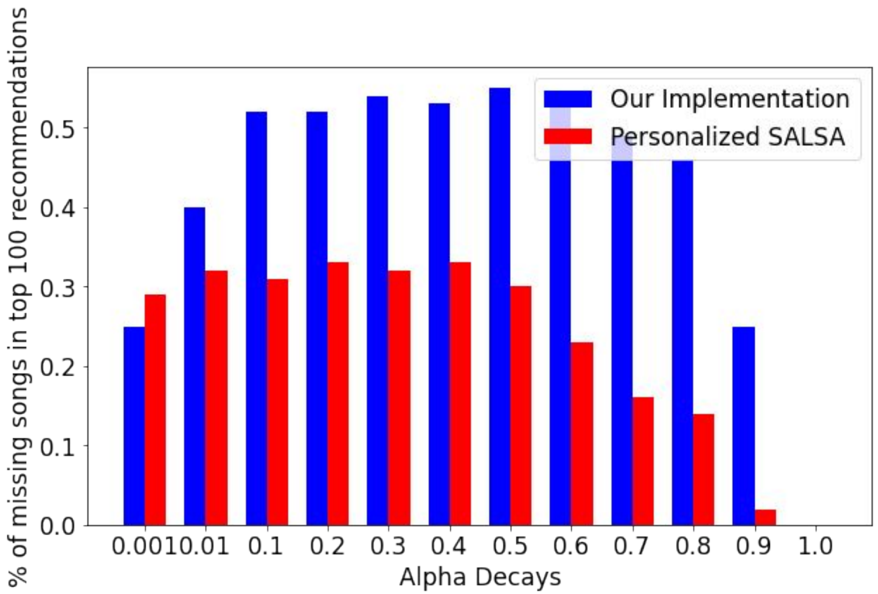

Top 100 Recommendations

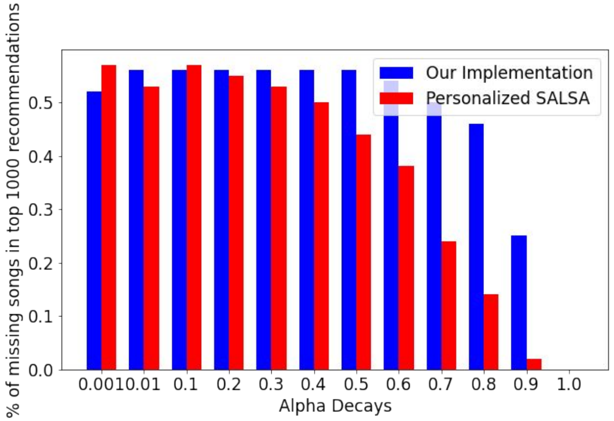

Top 1000 Recommendations

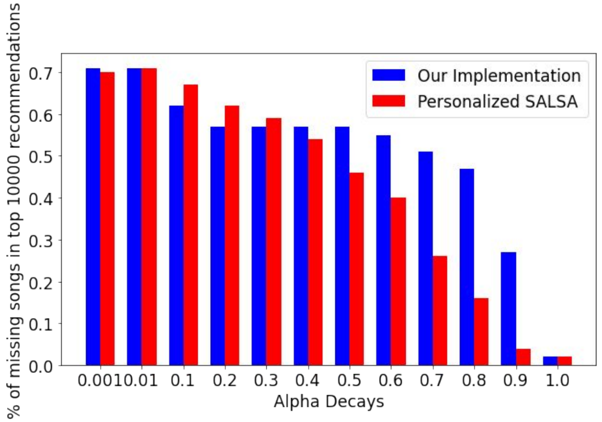

Top 10000 Recommendations

TU Delft

Effect on Ranking (RBO[1]) of Increasing Decays

Results

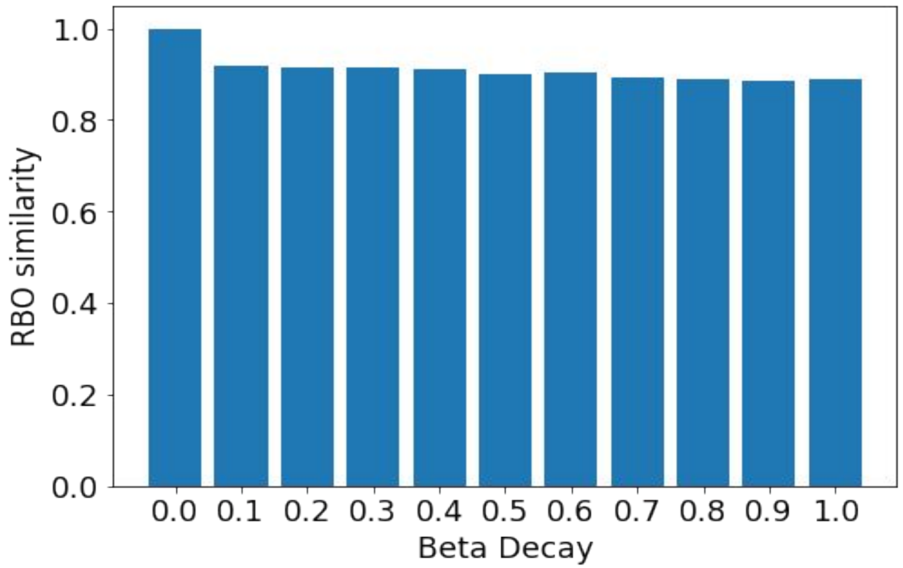

RBO vs increasing Beta Decay

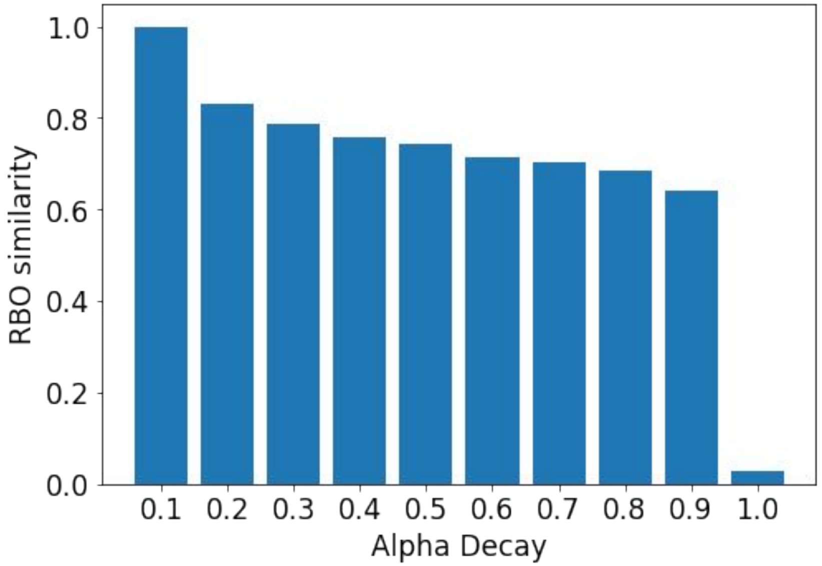

RBO vs increasing Alpha Decay

TU Delft

[1] W. Webber, A. Moffat, and J. Zobel, “A similarity measure for indefinite rankings,” ACM Transactions on Information Systems (TOIS), vol. 28, no. 4, pp. 1–38, 2010.

Single Sybil Attack: Serial Attack Increasing Alpha Decay

Results

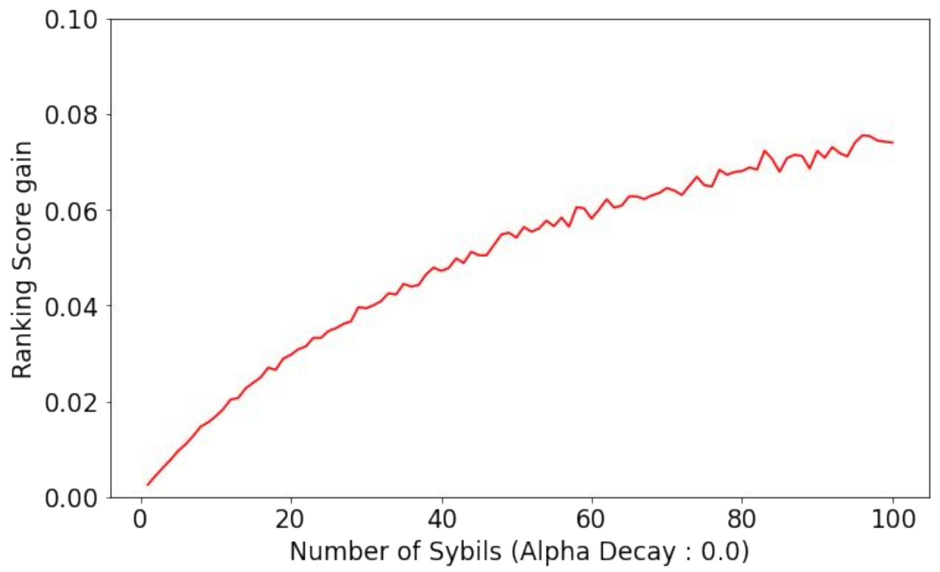

β = 0.0

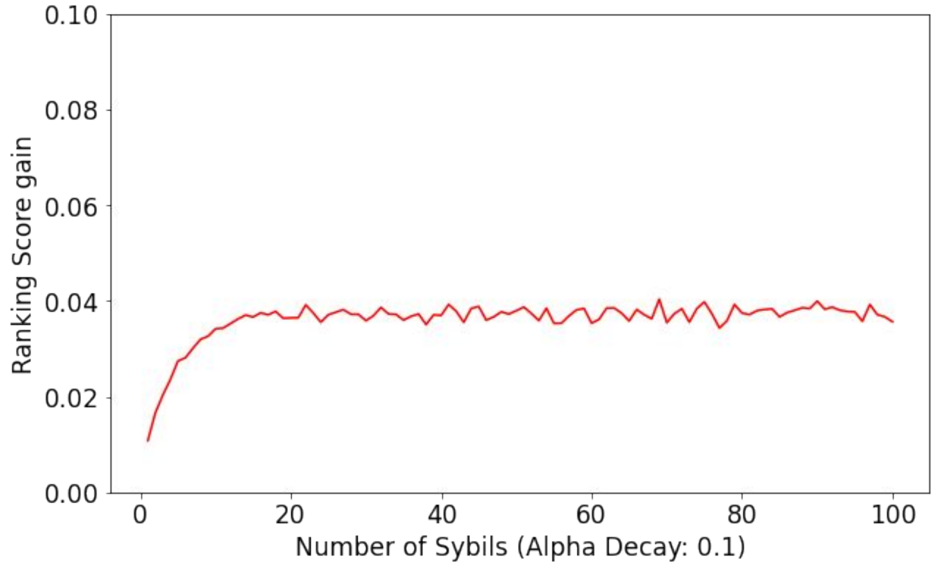

β = 0.0

Single Sybil Attack: Parallel/Cycle Attack Increasing Beta Decay

TU Delft

Results

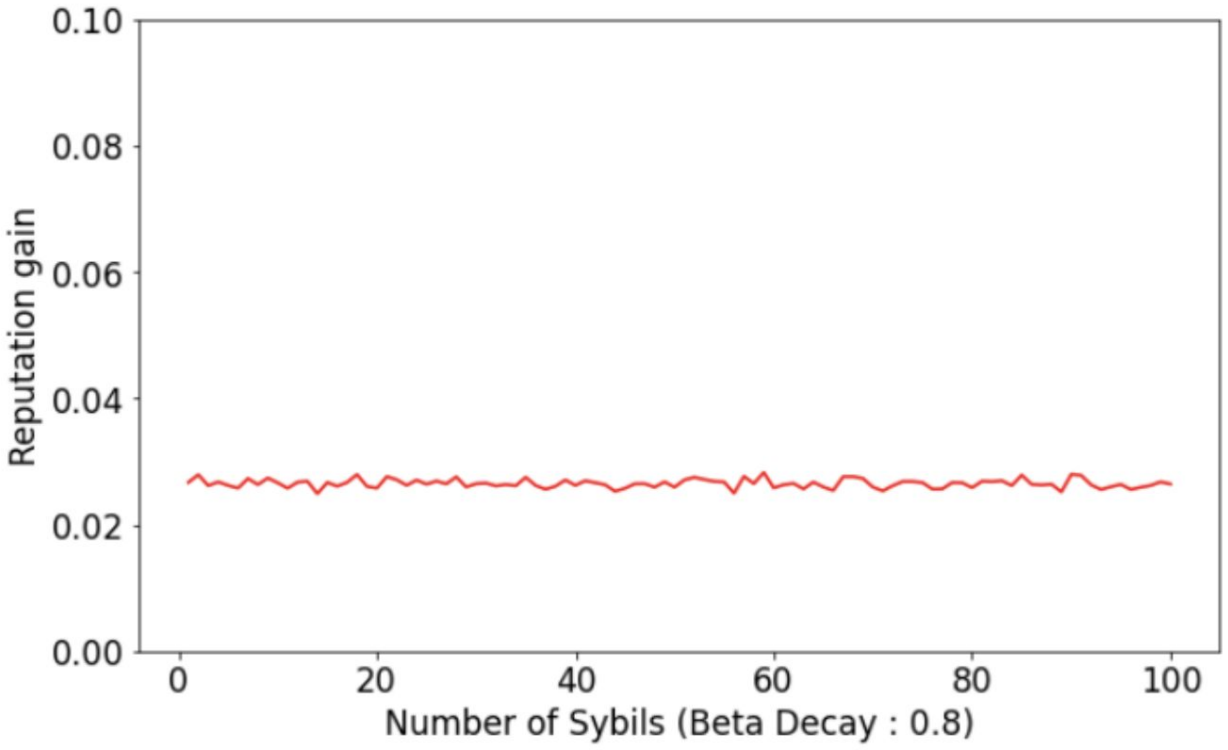

α = 0.1

Giga Sybil Attack

Attack Edges: 50% of the network

Results

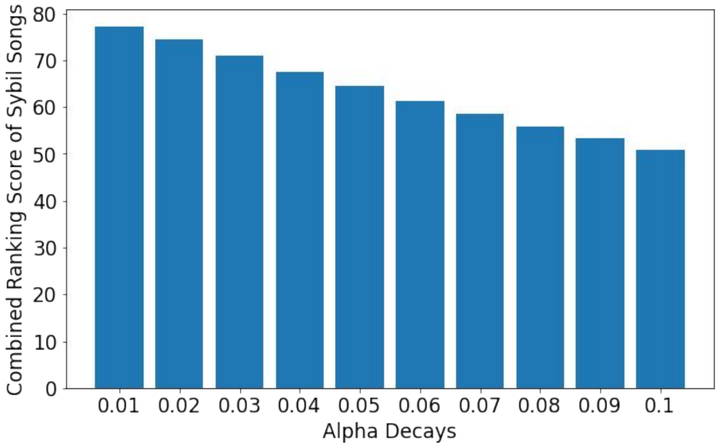

Score gained by Sybils vs Alpha Decay (β = 0.0)

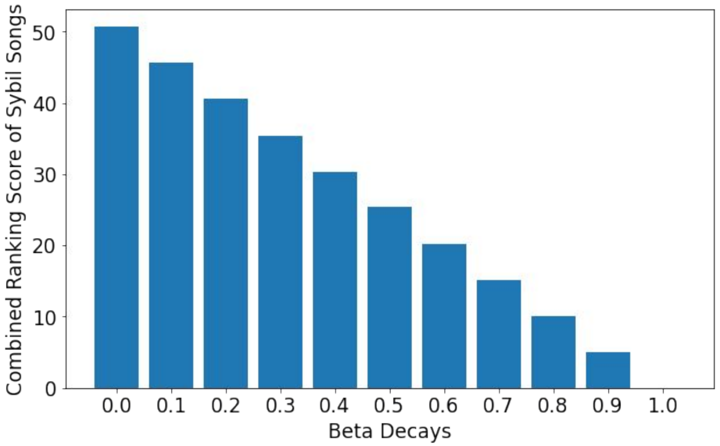

Score gained by Sybils vs Beta Decay (α = 0.1)

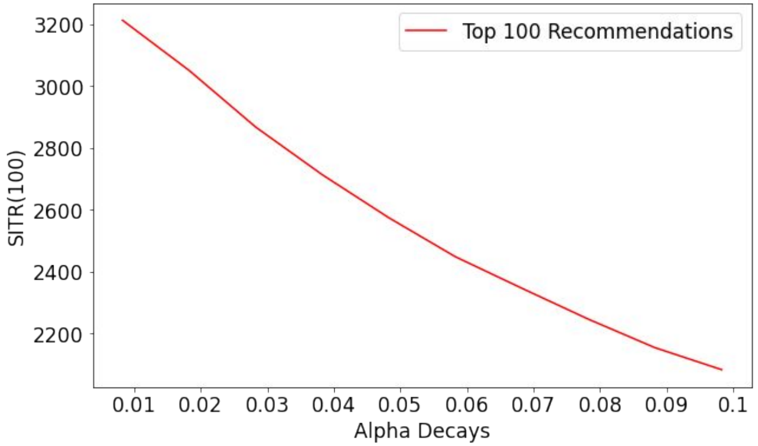

SITR(100) vs Alpha Decay (β = 0.0)

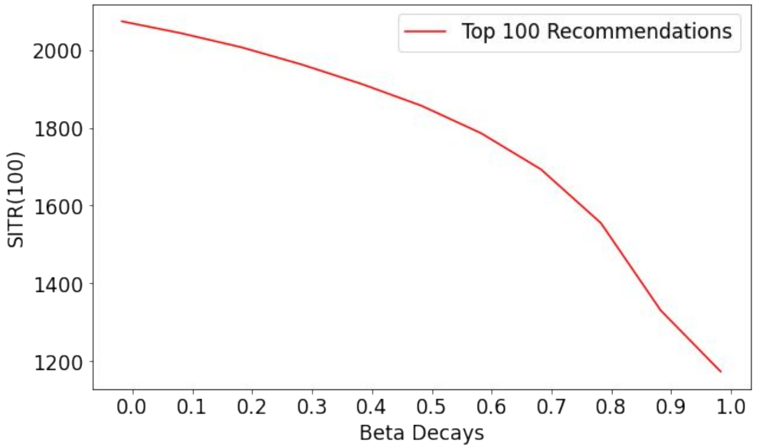

SITR(100) vs Beta Decay (α = 0.1)

Future Work/Limitations

- Notion of Trust is limited to “Sybil resistance”, explore other trust mechanisms
- Makes assumptions which might not suffice in real world scenarios (privacy guarantees, spoofing protection, timestamp synchronization)
- Evaluate performance and usability in other (non-music) scenarios
- Increasing quality of generated recommendations by feeding the output from our random walks as input to Machine Learning models

Thank you for your attention

$$RBO(S, T, p) = (1 - p) \sum p^{d-1} \cdot A_d$$

where,

$d = 1$ to ∞ (depth of the ranking being examined)

$X_d = |S_{:d} \cap T_{:d}|$ (Size of the overlap of S & T upto depth 'd')

$A_d = X_d/d$ (Agreement between S & T given by the proportion of the size of the overlap upto depth 'd')

$$SITR(x) = \sum_{i=0}^{i=x} \begin{cases} x - i & \text{if } DRankedI(i) \in \mathbb{S} \\ 0 & \text{else} \end{cases}$$

Implementation Challenges

The DIMACS Implementation Challenges help understand and improve the practical performance of algorithms for important problems, particularly those that are hard in the theoretical sense. The Challenges aid in determining realistic algorithm performance where worst-case analysis is overly pessimistic and probabilistic models are too unrealistic. Experimentation can provide insight into realistic algorithm performance when purely analytical methods fail, and it can provide new perspective that motivates deeper analytical results. Experimentation tests assumptions about implementation methods and data structures and provides an opportunity to develop and test problem instances instance generators to facilitate future comparisons.

The Implementation Challenges were inspired by David S. Johnson and date back to the early years of DIMACS. Each challenge addresses a particular problem or group of related problems and focuses the attention of many people on that problem. The challenges involve setting up a common infrastructure and library of test problems to allow researchers to evaluate their own implementations and compare them with those of others. The idea is to establish a common "playing field" in order to compare results and establish a common vision of the "state of the art."

The overarching goal of a challenge is to encourage interaction among the participants. Through these challenges, researchers exchange ideas, share test problems, combine methods, and focus on the most promising aspects of different methods. Though staged as a "competition," there are no real prizes. Implementation Challenges are about collaboration.

We model the Sybil-Tolerance as a bound on the benefit that the attacker can gain through a Sybil attack σ_S on feedback graph G. A reputation mechanism R is *Sybil tolerant* if the gain after performing a Sybil attack is limited by some constant $c \geq 0$:

$$\lim_{|S| \to \infty} \frac{\omega^+(\sigma_S)}{\omega^-(\sigma_S)} \leq c$$

Bootstrap Mechanisms

1. **Circle of trust**

   Start random walks from “seed set” instead of the source node

2. **New User**

   Improved User Collaborative Filtering based on [1] where:

$$sim(a,b) = Nsim(a,b) \times \tau + (1-\tau) \times Dsim(a,b)$$

Where Nsim is the traditional measure of similarity based on pearson correlation coefficient and Dsim is a measure of similarity calculated using the rating difference of users on their common items

More fine-tuned similarity metric than coarsely comparing items both users have interacted with