TECHNISCHE UNIVERSITÄT DRESDEN

DRESDEN concept

Attacks and Protection

Thorsten Strufe, Thomas Paul, Daniel Puscher
Chair for Privacy and IT-Security

Padova, 06.09.2016

cfaed CENTER FOR ADVANCING ELECTRONICS DRESDEN

5G Lab GERMANY

HAEC

ROS

ScaDS DRESDEN LEIPZIG

Access Model

User

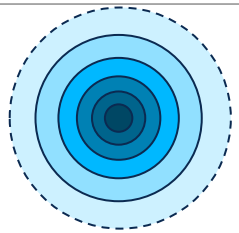

Grantable

- specific contact(s)
- contacts
- contacts of contacts
- service subscribers
- public

Implicit

- SNP

- Affiliates
  - Extenders
  - Advertisers

Everything the installing user can see

Not much (aggregates)
Unless they pay really well

- ISP

Everything their subscribers see/write *(until Nov 21st '12)*

Cloning Attacks on Social Networks

Target: get on the friend list of real users to get access to their personal information and their circle of trust

Two Cloning Attacks

- Clone the account of an existing user inside the same network and send friend requests to her contacts
- Clone the victim profile into a different social network where she is not registered and contact her friends

Aim of Attacks and Experiments

Is it possible for an attacker to launch impersonation attacks on a large scale against a number of popular social networking sites?

=>Obtain illegitimate authorizations

- Facebook (international)
- XING (international)
- LinkedIn (international)
- MeinVZ (popular in Germany, Austria, Switzerland)
- StudiVZ (popular in Germany, Austria, Switzerland)

Automated Profile Cloning

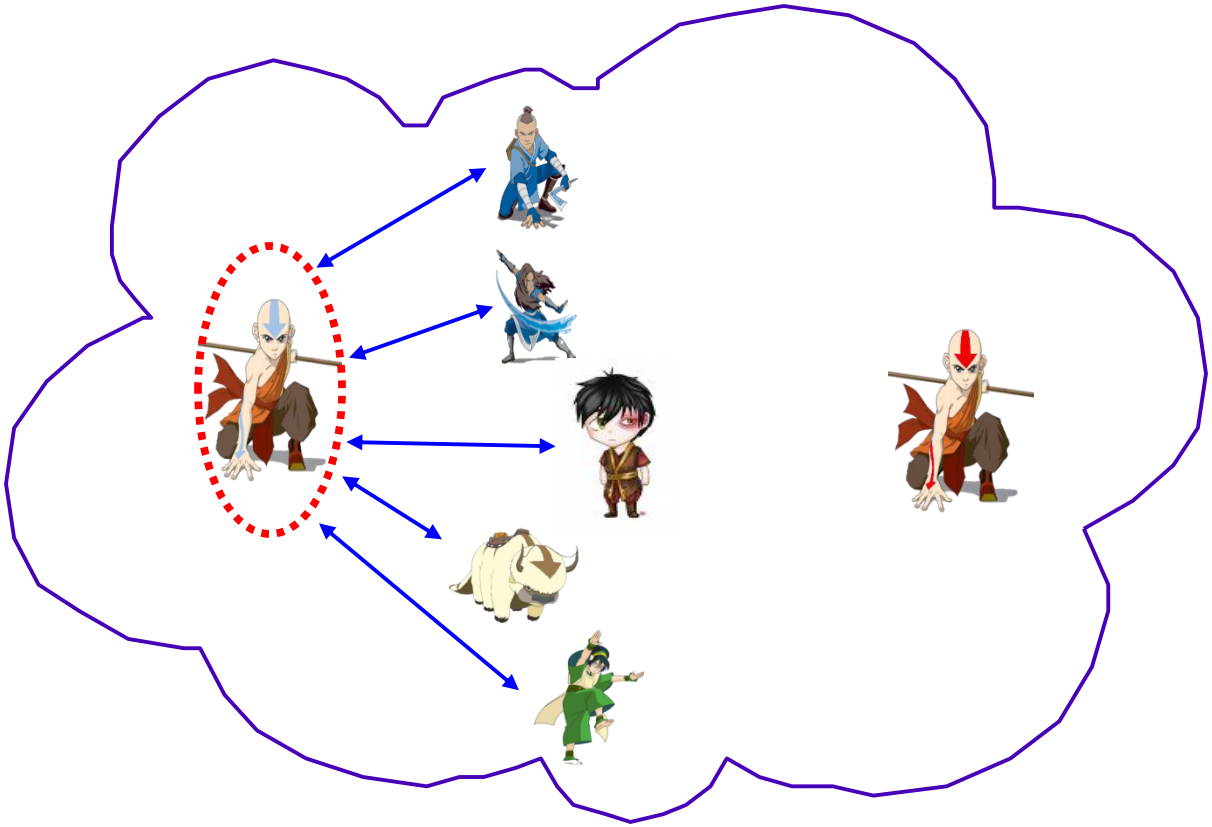

Automated Profile Cloning

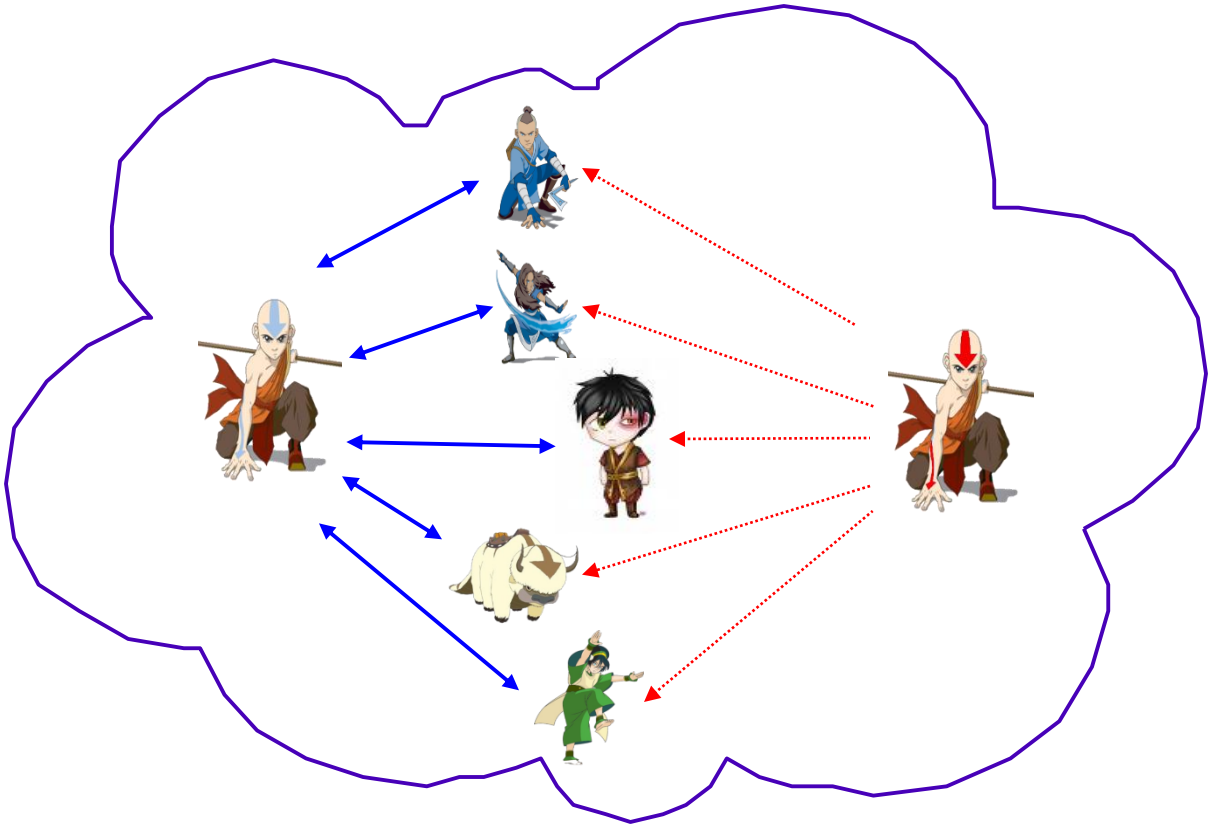

Automated Profile Cloning

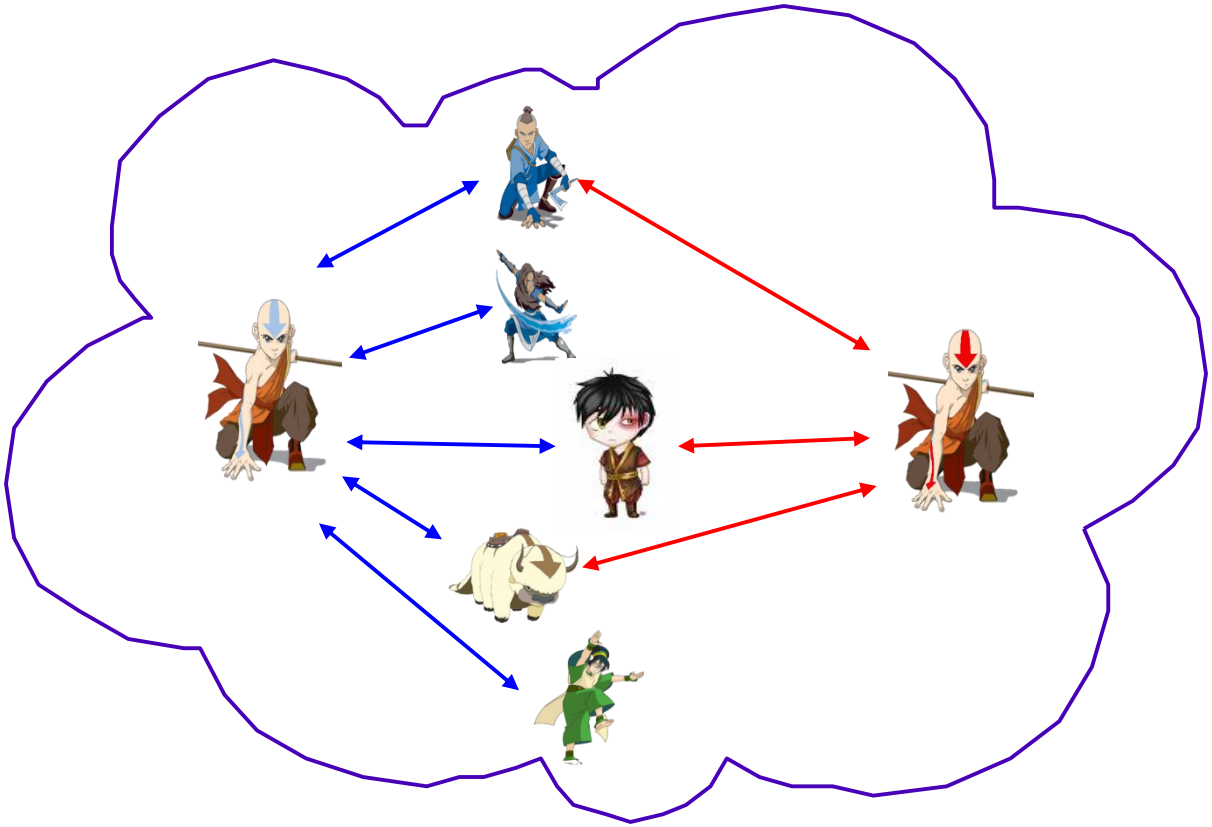

Cross-Site Profile Cloning

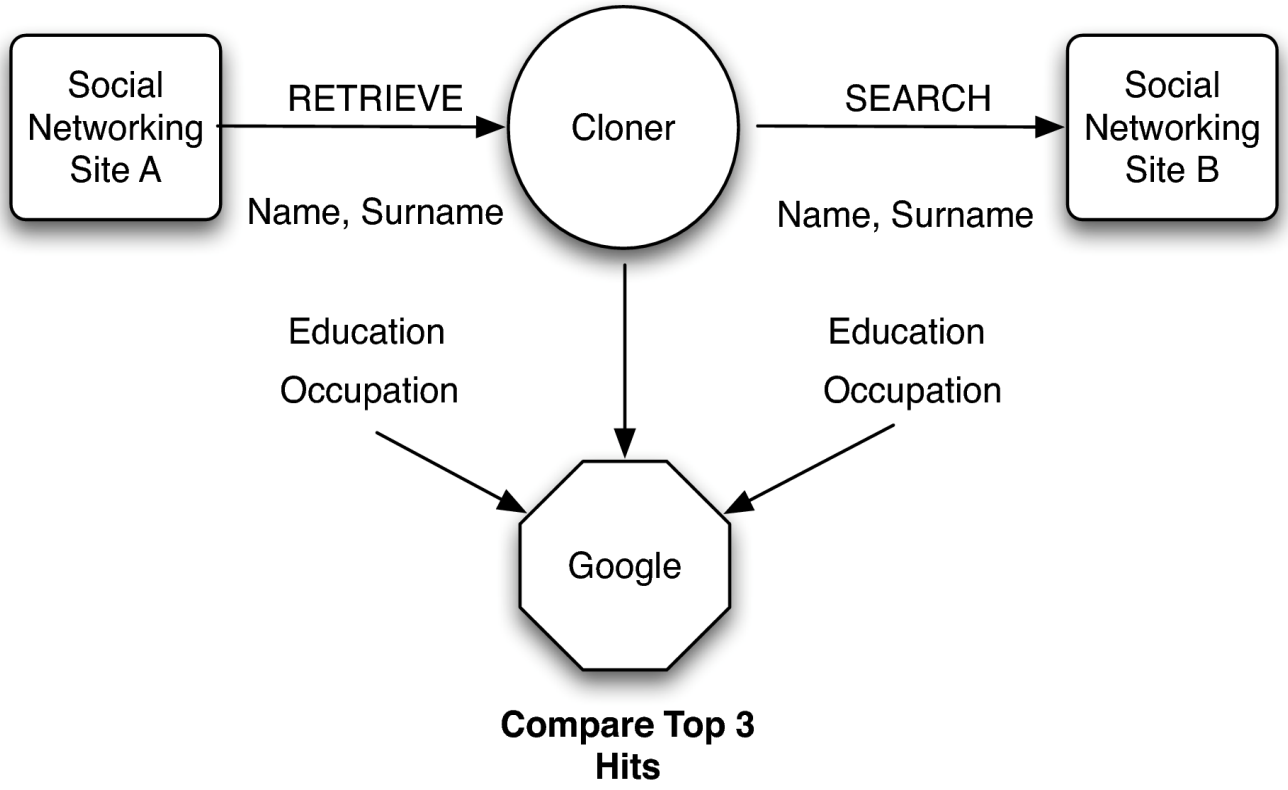

Obstacle: CAPTCHAs

CAPTCHA: Completely Automated Public Turing test to tell Computers and Humans Apart

CAPTCHAs are employed to prevent automated programs from accessing and abusing the services

In order to automate the attacks, a number of CAPTCHA breaking techniques were developed

- "Quick and dirty", techniques are not perfect
- The aim is to break the CAPTCHAs efficiently enough to make automated attacks against several social networking sites possible

MeinVZ and StudiVZ CAPTCHAs

GD Library (PHP) CAPTCHAs

CAPTCHAs always contain 5 letters

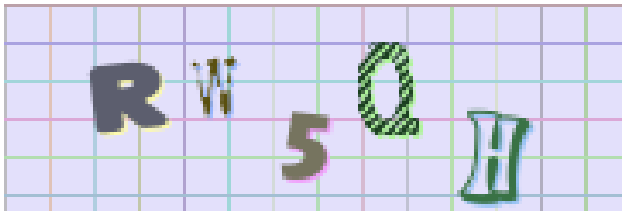

Each letter is written in

- Different font
- Different background and foreground color

Often tilted, scaled or blurred

A simple grid-base noise is added to the image

Quick script* with success rate of 88.7%

*Cracking the CAPTCHAs was done with serious amounts of help from Michael Roßberg/TU-Ilmenau

Facebook CAPTCHAs

Adopts ReCAPTCHAs

- Asks words that are encountered while digitizing books that cannot be correctly recognized by the OCR program
- By solving the CAPTCHAs, the user contributes to the effort to increase the accuracy of the text of the digitized book

ReCAPTCHA asks meaningful words. Therefore, after solution is found, the word is sought in a dictionary

- Result additionally submitted to Google as check

Script with success rate of 7%

Might seem small, but…

- If every bot is capable of solving 7 CA
- per day, a botnet that consists of 10.000 can
- send 70.000 friend requests per da
- Attack against Microsoft Live Hotma
- similar success rat

Prototype Implementation: iCloner

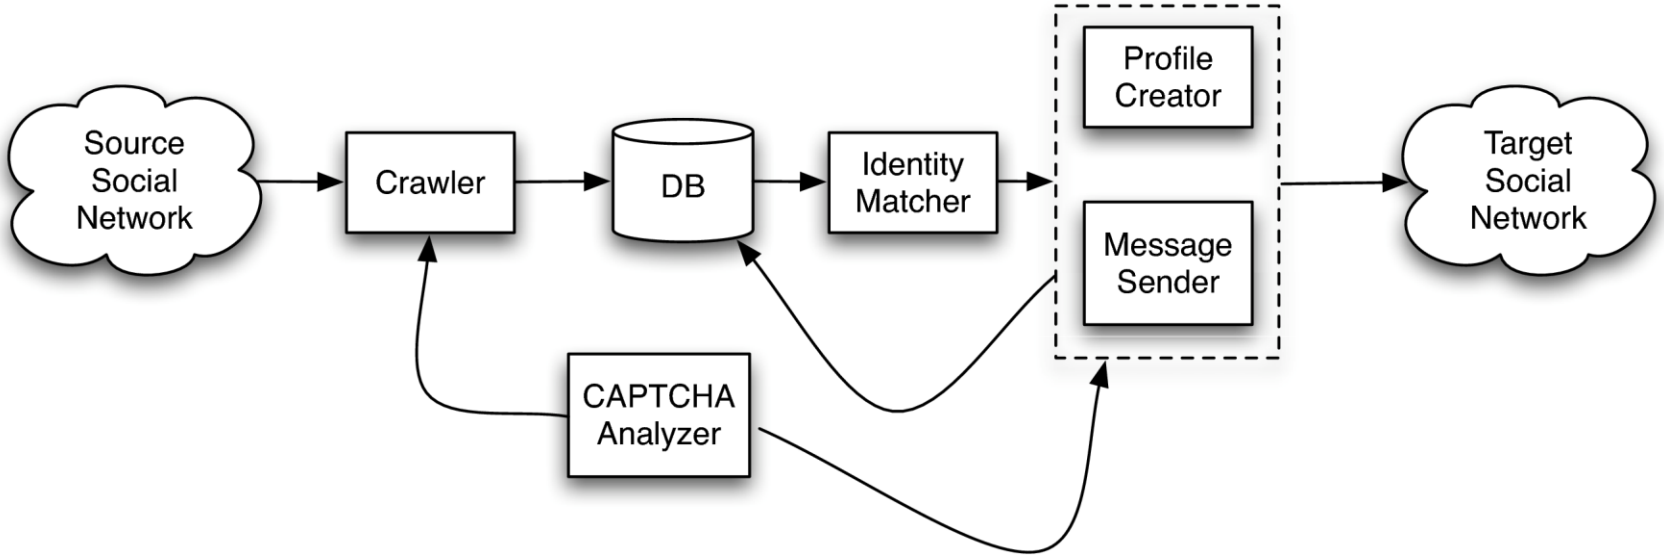

Real-World Experiments

Is it feasible to perform cloning attacks in the real-world?

Questions:

- Can an attacker launch large-scale attacks?
- How willing are users to accept friendship requests from forged profiles of people who are already in their friendship lists?
- Is it possible to efficiently find two identical accounts in two different social networks?

Profile Crawling Experiments

StudiVZ and MeinVZ

- Displays CAPTCHA if large number of requests come from one account
- To collect as much information as possible, without being noticed, 16 accounts were created, and separately used for crawling
- Collected 5M profiles with contact information, and 1.2M complete user profiles

XING

- Does not display CAPTCHA, but disables the account if the account requests around 2000 pages consecutively
- 118,000 accounts were crawled

Profile Cloning Experiments

Attack: duplicate the profiles of five users (D1,...,D5) and create fictitious profiles (F1,…,F5 as control group)

Measure ratio of accepted re-friending requests

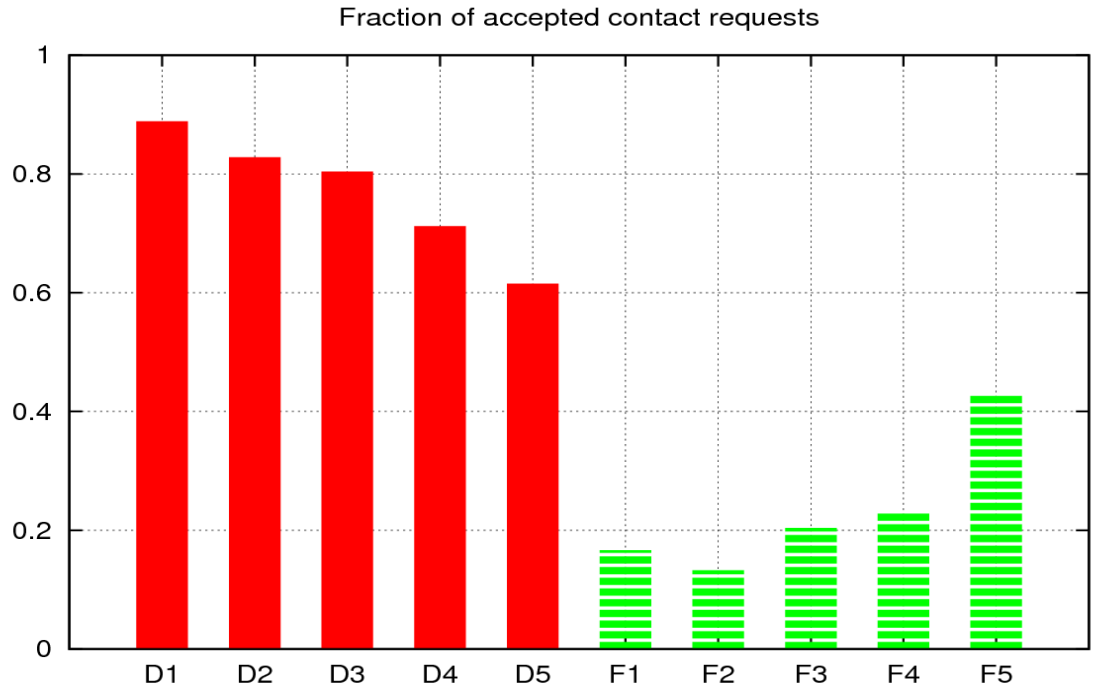

Profile Cloning Experiments

Do the users really trust their friends in their friend list?

Would they click the link seen in the message below?

```
Hey, I put some more pictures online. Check them
here!:

http://193.55.112.123/userspace/pix?user=<account>
&guest=<contact>&cred=3252kj5kj25kjk325hk}

Ciao, <account first-name>
```

Profile Cloning Experiments

Click through rate for messages from duplicate / fictitious profiles

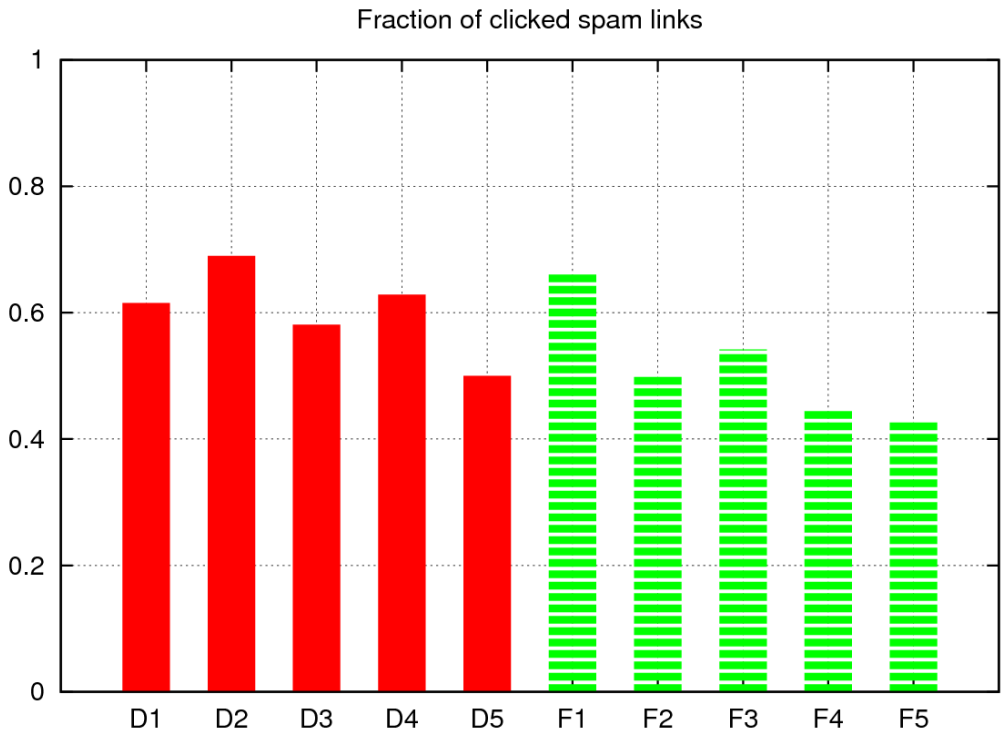

Cross-Site Profile Cloning Experiments

Cloning profiles that exist on XING, but not on LinkedIn

The success of the cross-site profile cloning depends on the number of users that have a profile in both of the networks

From around 30.000 crawled profiles in XING, 3.700 were also registered in LinkedIn

Cross-Site Profile Cloning Experiments

Clone 5 users from XING to LinkedIn

iCloner identified 78 out of 443 XING friend contacts that were also registered in LinkedIn

Fraction that has actually accepted the contact requests:

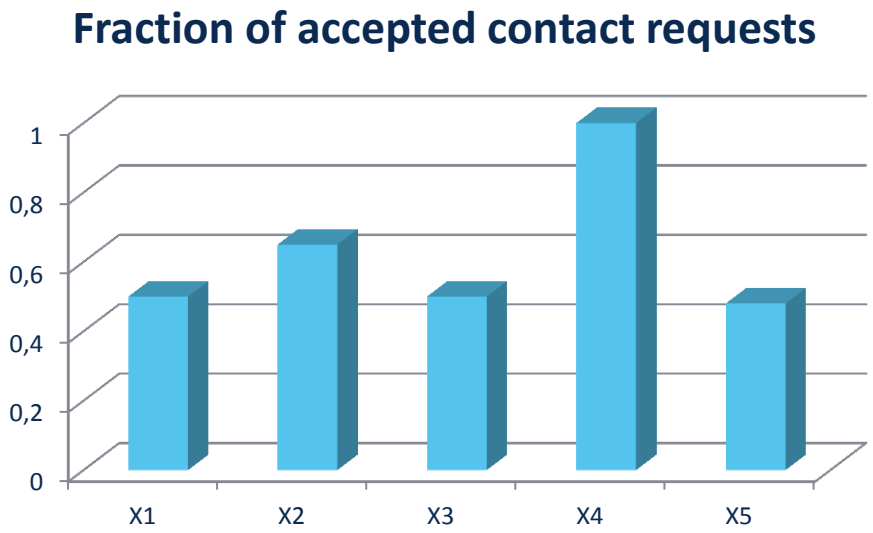

Summarizing the iCloner Experiments

Large scale profile retrieval (used to be) easy

Captchas obstacle but no deal-breaker

Cloning profiles (and obtaining „friendships“) is easy

- Sybil assumption not realistic (interaction, may be)

Online Services: Social Media and the Web

Web traffic is converging to sites of 6 corporations

- Success due to integration and strong personalization
- Data minimization conflict with business modell
- Trackers snoop on remaining pages

Convergence of communication and expression

- Facebook evolves to integrated communication platform with 1.6 Bn users
- Google, g+: 350 Mio user
- Clear name: perfectly identifiable

Increasingly mobile utilization

- Perfect location, easy tracking
- Configuration more tedious

TOP 10 WEB BRANDS BY UNIQUE AUDIENCE (U.S. TOTAL)

Rank	Brand	Unique Audience	Time Per Person (h
1	Google	170,629,000	2:05:30
2	Facebook	145,297,000	6:41:44
3	Yahoo!	135,100,000	2:32:52
4	YouTube	124,073,000	1:57:28
5	MSN/WindowsLive/Bing	123,133,000	1:15:40
6	Microsoft	86,986,000	0:47:26
7	Amazon	84,735,000	0:38:14
8	AOL Media Network	83,826,000	2:09:36
9	Wikipedia	76,310,000	0:24:25

[Nielsen]

Model and Adversaries

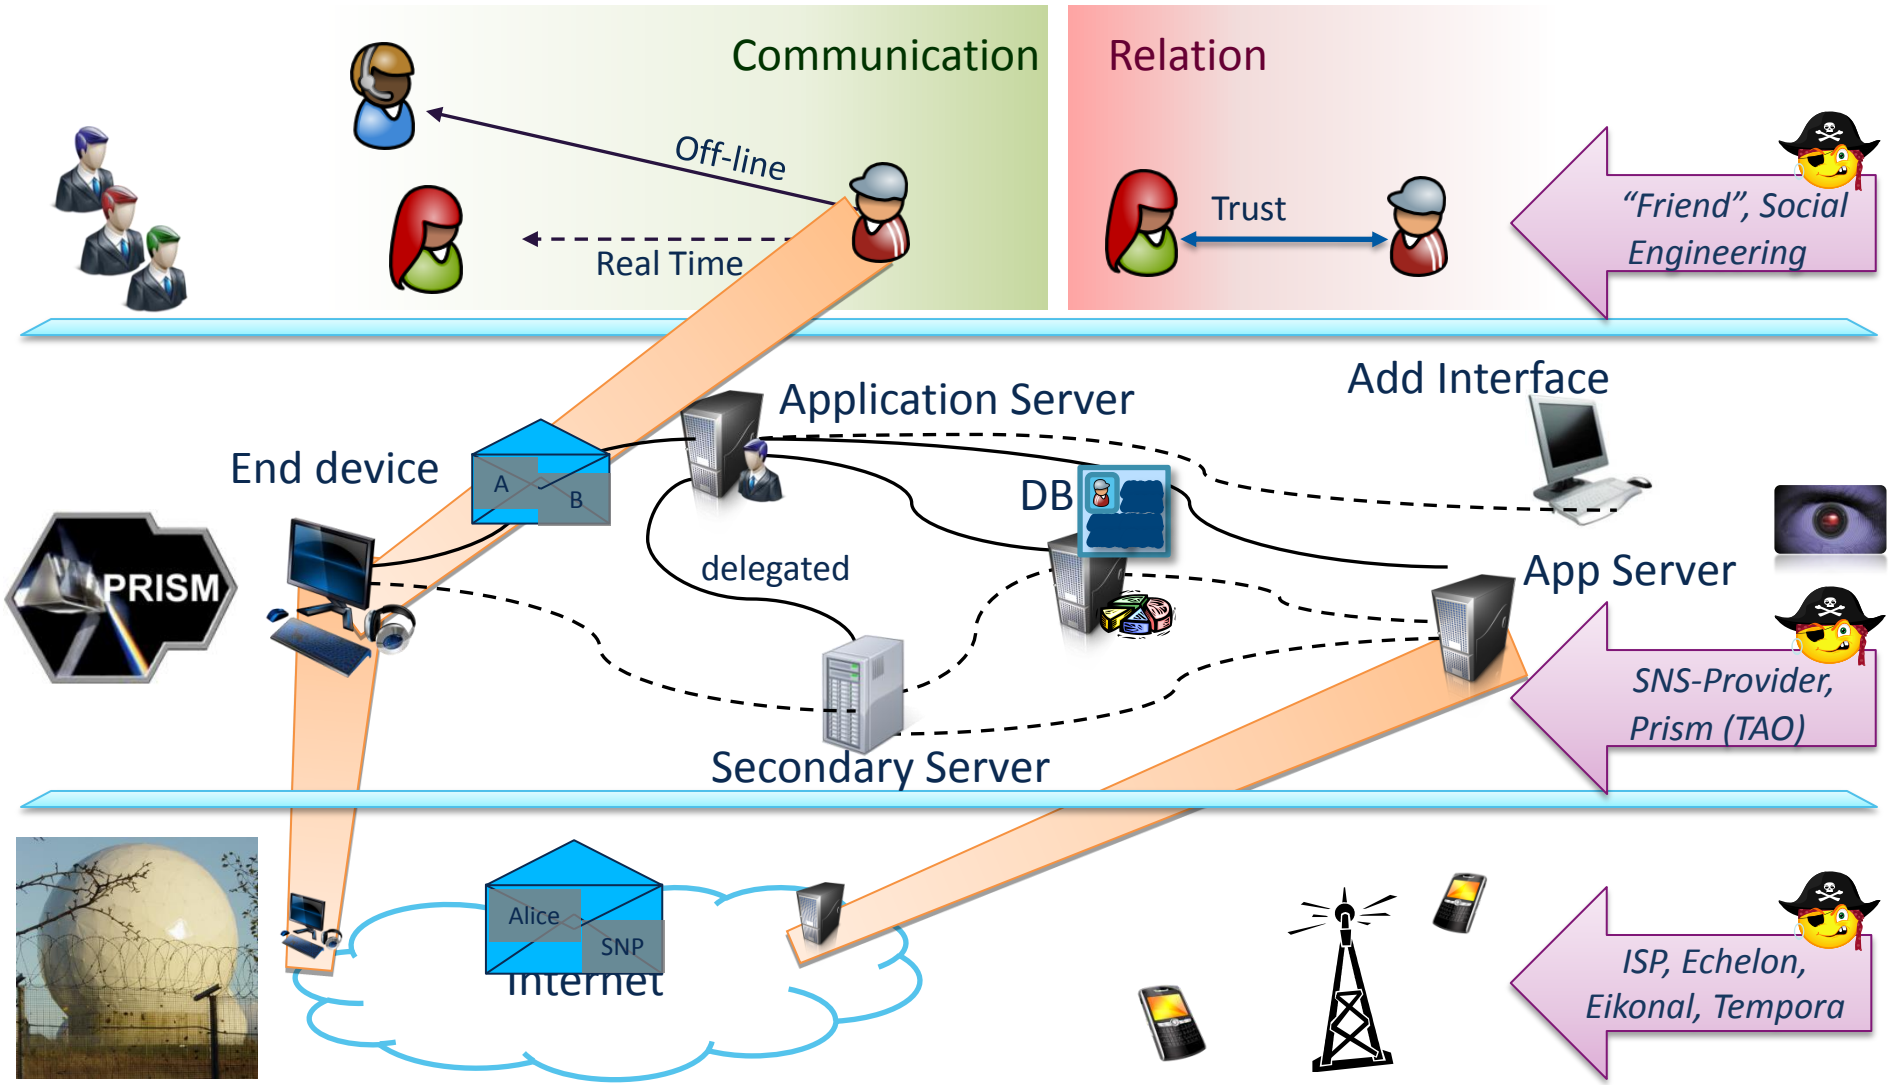

What You can do –

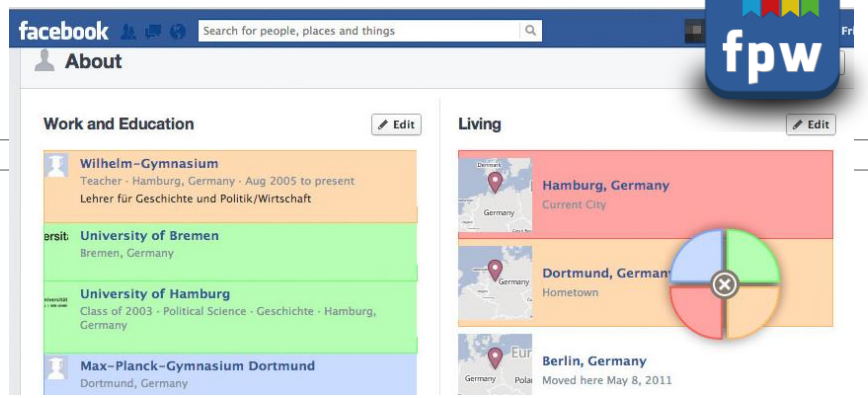

- *Authorize actively! (Privacy Controls)*
- *Communicate confidential (Encrypt your traffic)*
- *Lock out the mediator (E2E encryption)*

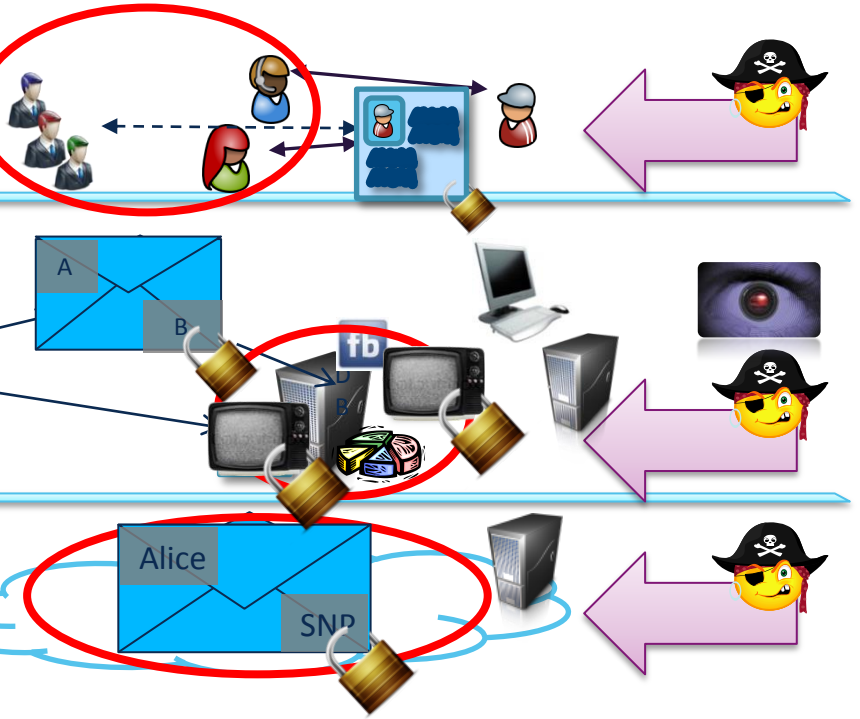

Stakeholders and the Data at their Hands

Explicit

- Created content
  - Profile, posts
- Annotations/comments
- Preferences/structural interaction (contacts, +1, etc)

Extracted

- Profiling aggregates
- Preference models
- Image recognition models

Incidental / „metadata“

- Observed:
  - ***Session artifacts*** (time of actions), ***interest*** (retrieved profiles; membership in groups/ participation in discussions), ***influence*** (users)
  - Clickstreams, ad preferences, ***communication*** (end points, type, intensity, frequency, extent), ***location*** (IP; shared; gps coordinates), udid
- Inferred
  - ..derived from observations
  - ..from homophily

Externally correlated

- Interest/preferences (clickstreams through ad networks, fb-connect)

Content Protected, so: „It's only Meta Data"

Metadata privacy

In controlled (opt-in!) study [1], participants

Called their family,…

- *… adult establishments,*
- *… firearms dealer,*
- *… headshop, hydroponics- and hardware store,*
- *…different groups of medical specialists,*
- *…family and planned parenthood offices*

„Facebook Mining" attacks

single-term lecture (students without any prior knowledge on ML)

Information (ab)used:

- Partial profiles
- Neighborhood (homophily)

Inferred (with high accuracy):

- Gender
- Age
- Education level
- Expected tenure with employer
- Sexual preferences
- Political preferences

[1] https://cyberlaw.stanford.edu/blog/2013/11/what%27s-in-your-metadata

PETs: Entire Distribution of Data and Control

Decentralize the services

System classes

- Federated SNS
- P2P- / DOSN
- Social overlays and darknets

diaspora*

PeerSoN

Frankfurter Allgemeine
ZEITUNG FÜR DEUTSCHLAND

Deutscher Bundestag

May be better…

Safebook – Privacy through Decentralization

Centralized service identified as vulnerability

Safebook: Secure Social Networking through decentralization

- Remove centralized instance
- Distribute storage and control

- Decentralization requires: discovery, trust, *controlled access, availability*

- Friends in social networking services trust each other in the real world
  - Leverage existing „social trust" to encourage ***cooperation***
  - ***Data replication*** at trusted nodes to facilitate availability
  - Suspect all other service providers: encrypt everything

This is difficult! Or is it?

Architecture

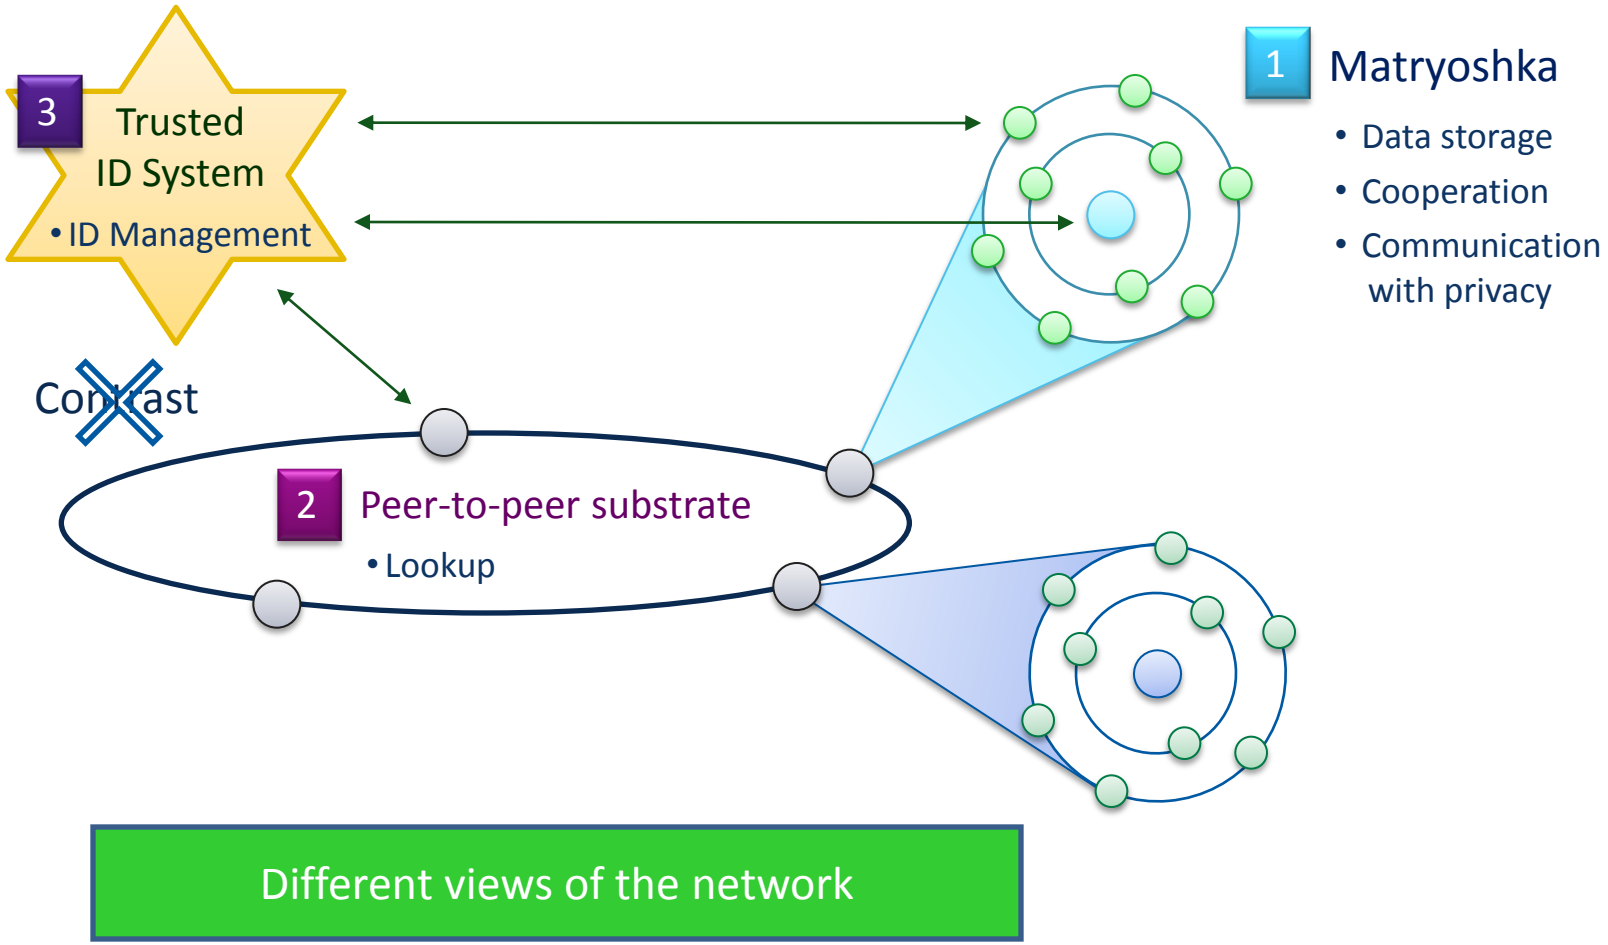

The following slides cf.: Cutillo, “Safebook”, 2009

User i's Matryoshka

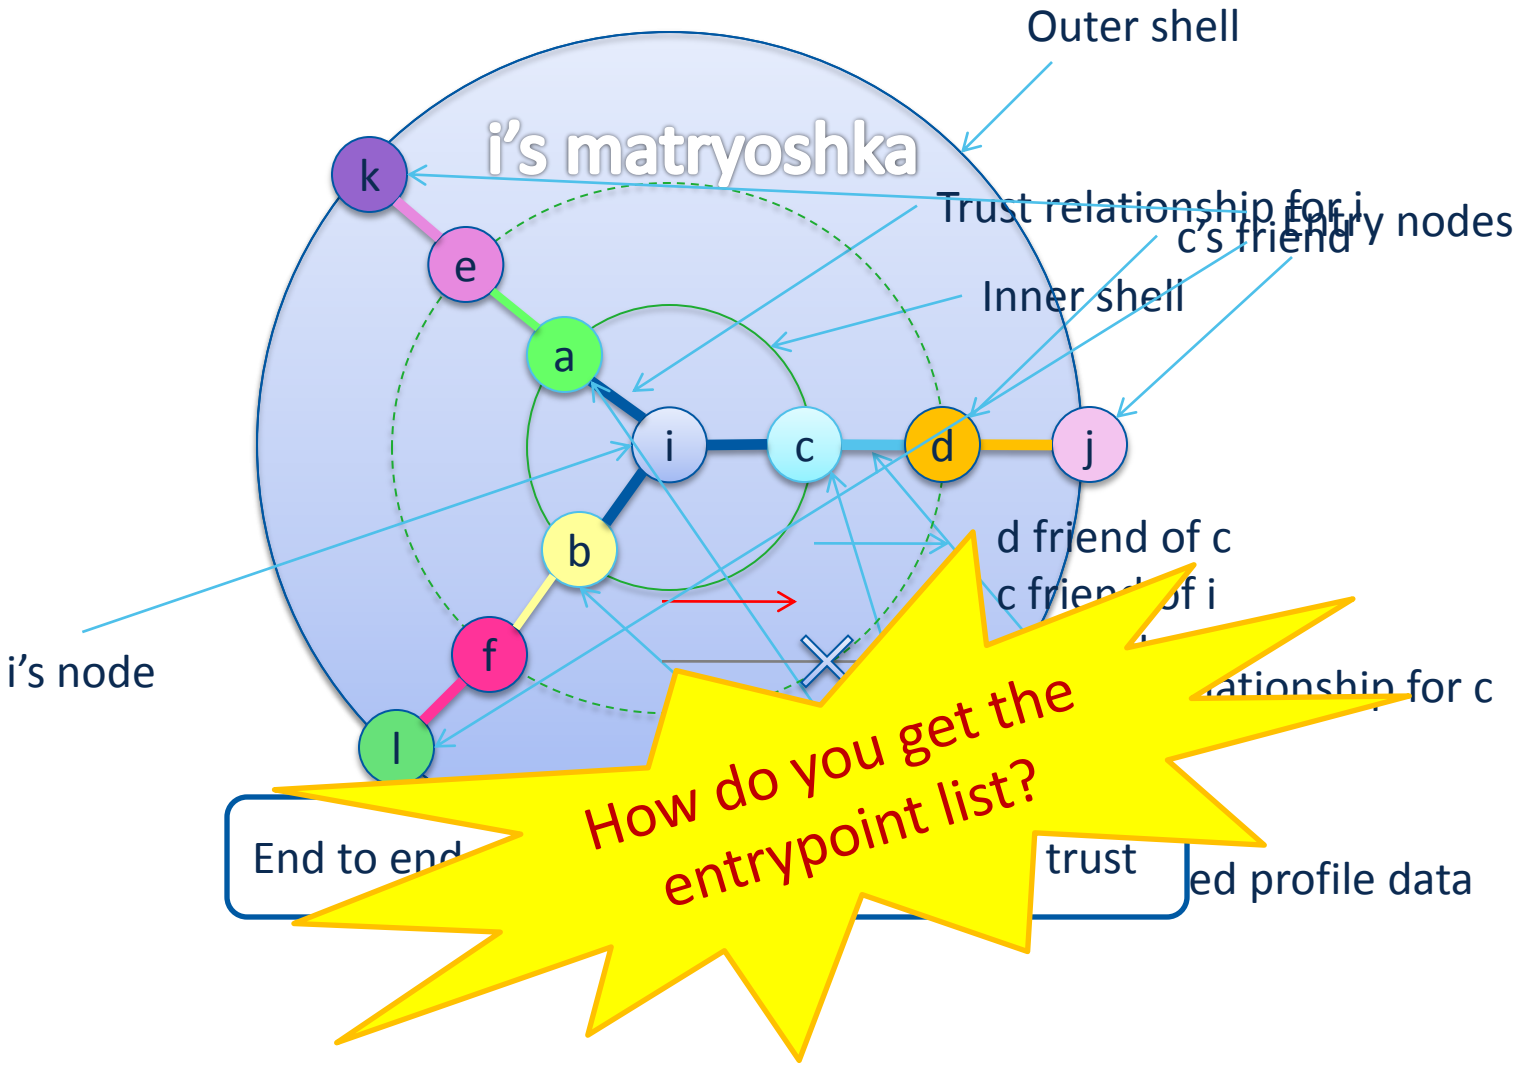

Finding it, using P2P: *a* looks for *b*

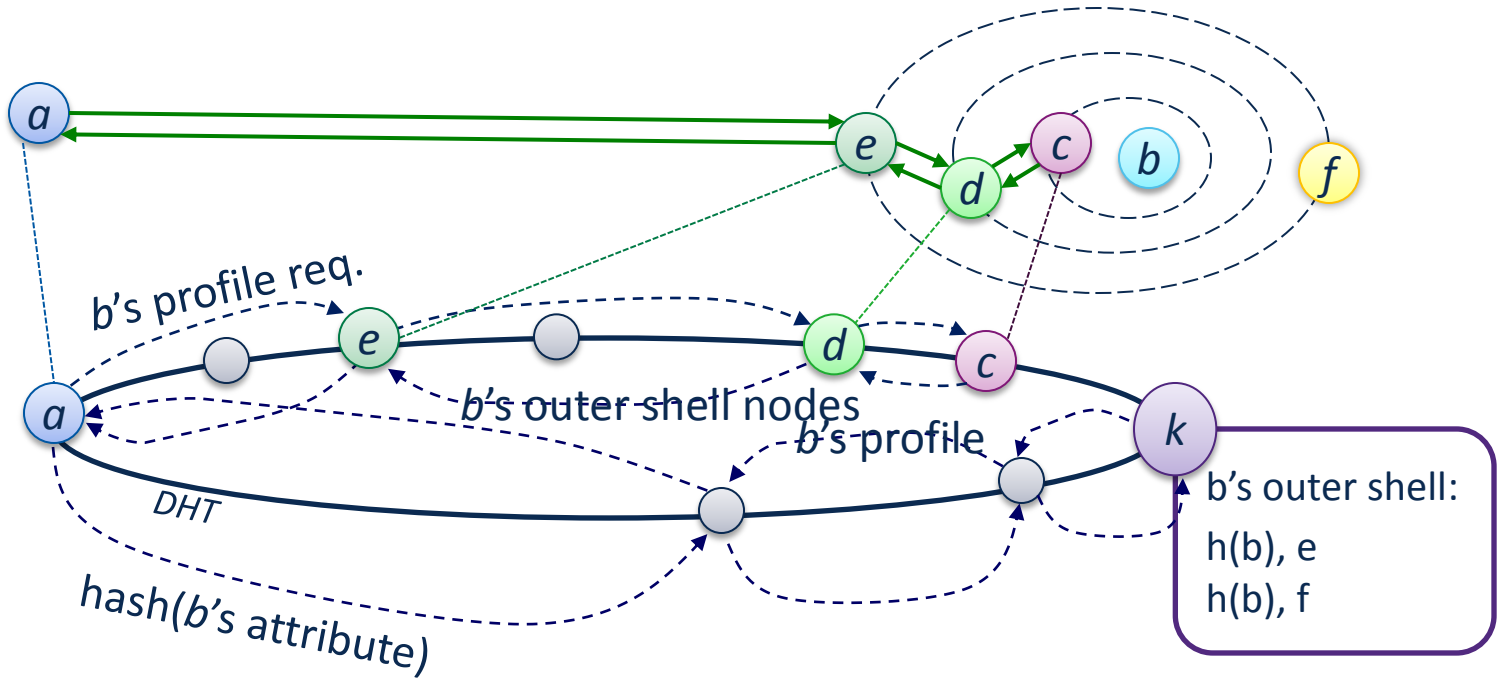

lookup
- *a* looks for *b*'s entry nodes
- *k* provides *b*'s outer shell nodes

data request
- *a* sends profile data request to a *b*'s entry node

Data reply
- One of b's inner shell nodes answers

Safebook

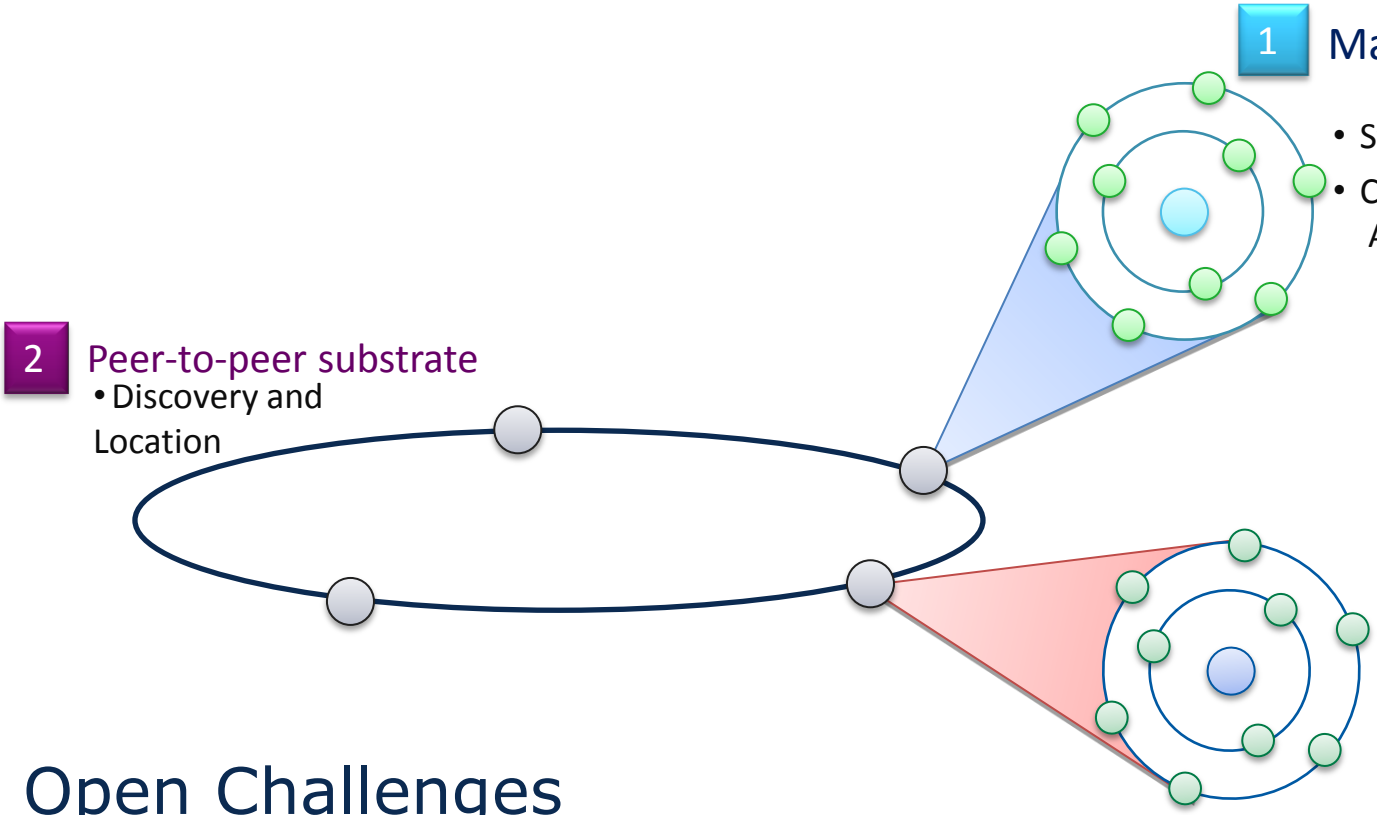

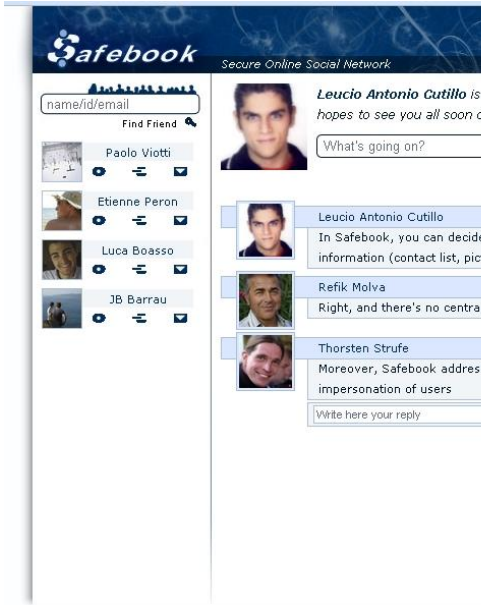

Open Challenges

- ***Performance*** insufficient
- ***Availability*** questionable (correlated churn)
- ***Concealed participation*** impossible

Social Overlays ("Darknets")

Decentralized OSN don't achieve what we want…

Stricter requirements

- Anonymity/ Pseudonymity (sender and receiver)
- Hidden participation (no 3rd party disclosure: hidden „friendships")
- Efficient discovery and interactive communication

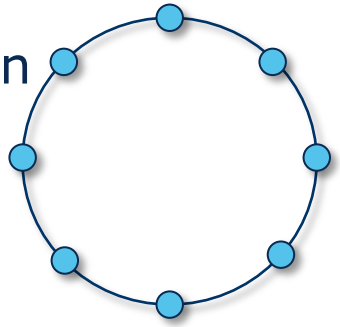

Concepts

- Connectivity constraints: mutual trust in RL
  - Overlay reflects social trust graph, topology is fixed
- Information containment: source rewriting, mixing
- Addressing and routing
  - log / polylog expected routing length required
  - Structured overlays: ***(1) choose ID, (2) choose neighbors***
  - ***(2) is restricted .. adapt (1)***

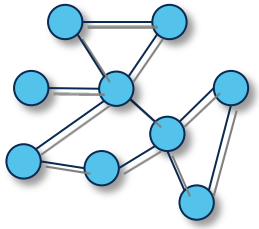

Resilient Social Overlays

Prevent identification, censorship and retribution.

From DOSN to darknets: Tightening requirements

- Concealed participation
- Unobserveability
- Metadata privacy (sender-, receiver-, relationship anonymity)

So where's the problem?

Classic overlays:

- Disclosure of IP address
- Eclipse, X-hole attacks

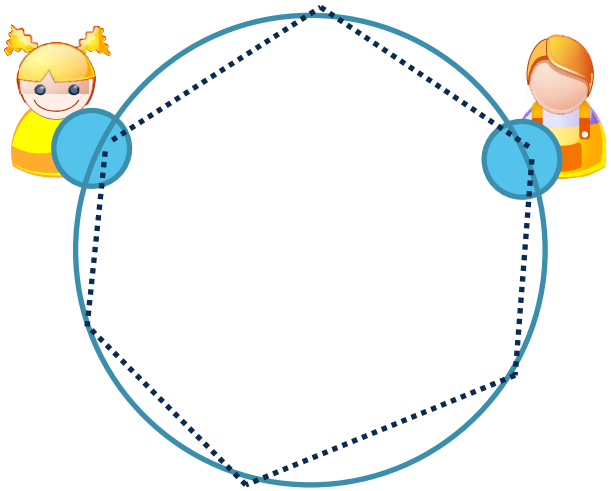

Social Overlays: Embedding/Virtual Overlays

Concepts of social overlays:

- Constrain connectivity to social links
- Contain information
- Attempt to route messages

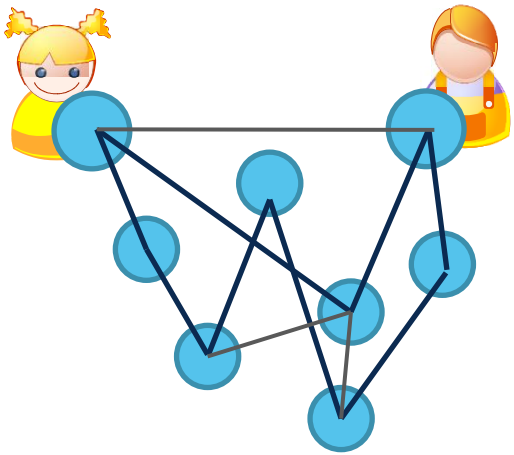

Embeddings

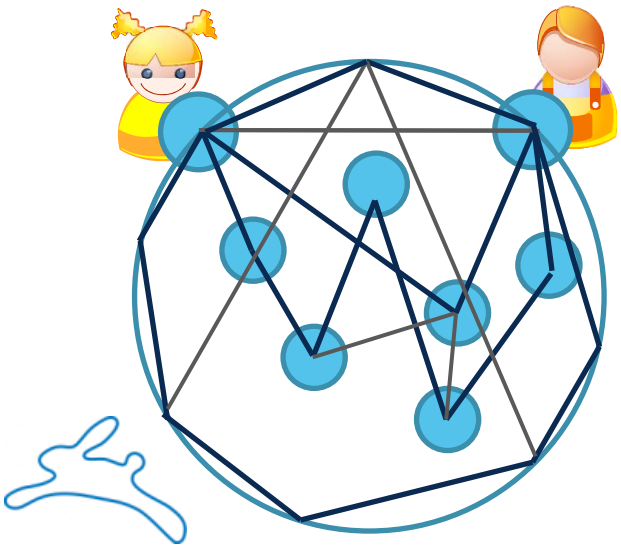

Network Embedding

A ***network embedding*** on an undirected graph $G = (V, E)$ is a function

$$ID : V \to M$$

to a metric space M equipped with a distance

$$d : M \times M \to \mathbb{R}+ .$$

For a node u ∈ V, *ID(u)* is the identifier of *u*.

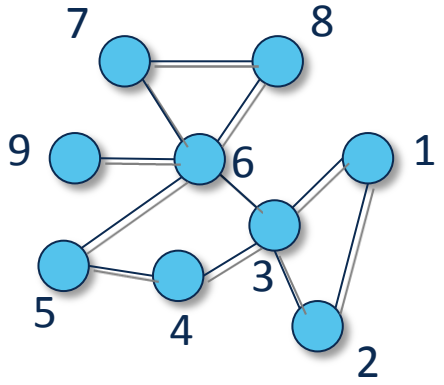

Greedy embeddings

guarantee greedy routing success (for every distinct node pair *s,t*: *s* is connected to or has a neighbor that is closer to *t*).

Goal:

find a decentralized algorithm that approximates a greedy network embedding

Enhancing Freenet's Embedding

Distortion extends paths

Aim: *greedy embedding*
Trees can be embedded

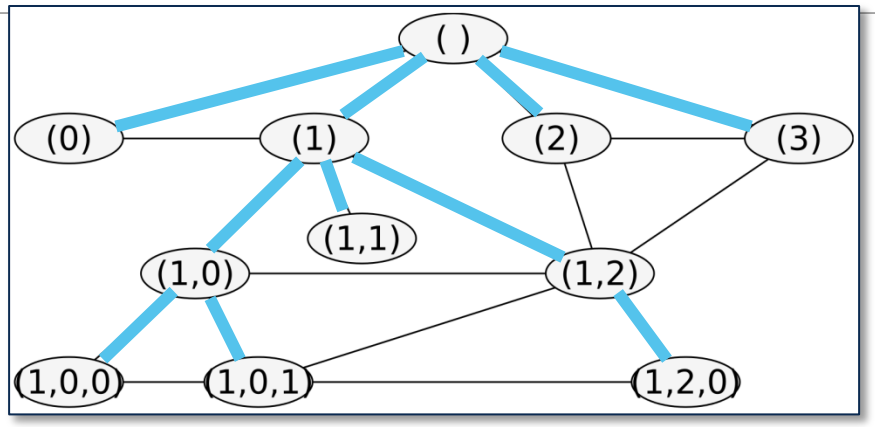

PIE tree embedding
1. Find spanning tree
2. Enumerate children

Distance metric:

$$d(s,t) := |s| + |t| - 2cpl(s,t)$$

Challenges:
- Tree addresses
  - Leak neighborhood
  - Addresses leak receiver
- Attacks on tree construction

Recovering Sender and Receiver Anonymity

Receiver anonymity

- (Return) address needed
- Distance: longest prefix match

- Blinded addresses:

1. Randomize:
   $[1,2,0] \rightarrow [r_1,r_2,r_3]$
2. Padding
   $[r_1,r_2,r_3] \rightarrow [r_1,r_2,r_3, r_{k+1}, \dots, r_L]$
3. Blinding
   $k, [r_1,..,r_L] \rightarrow (k,[h(r_1 \oplus k),h(r_2 \oplus h(r_1 \oplus k)\dots])$

- Distance metrics:

$$d_1(s,t) := |s|+|t|-2cpl(s,t)$$

$$d_2(s,t) := L - cpl(s,t) - \delta$$

Theoretical analysis

Performance Bounds

- Tree routing O(log n)
- Tree maintenance O(log n) per join/leave

Security Analysis

- Plausible deniability: Receiver cannot uniquely be identified
- Minimal information loss to allow for routing

[1] Roos, Beck, and Strufe: INFOCOM 2016

Performance Evaluation

TE is a greedy embedding

Simulation Experiment

- Topology: PGP Web of Trust
- Embeddings: Freenet/RW
- Routing: DDFS/Greedy

Is it robust?

Summary:

- *It's robust and fast!*

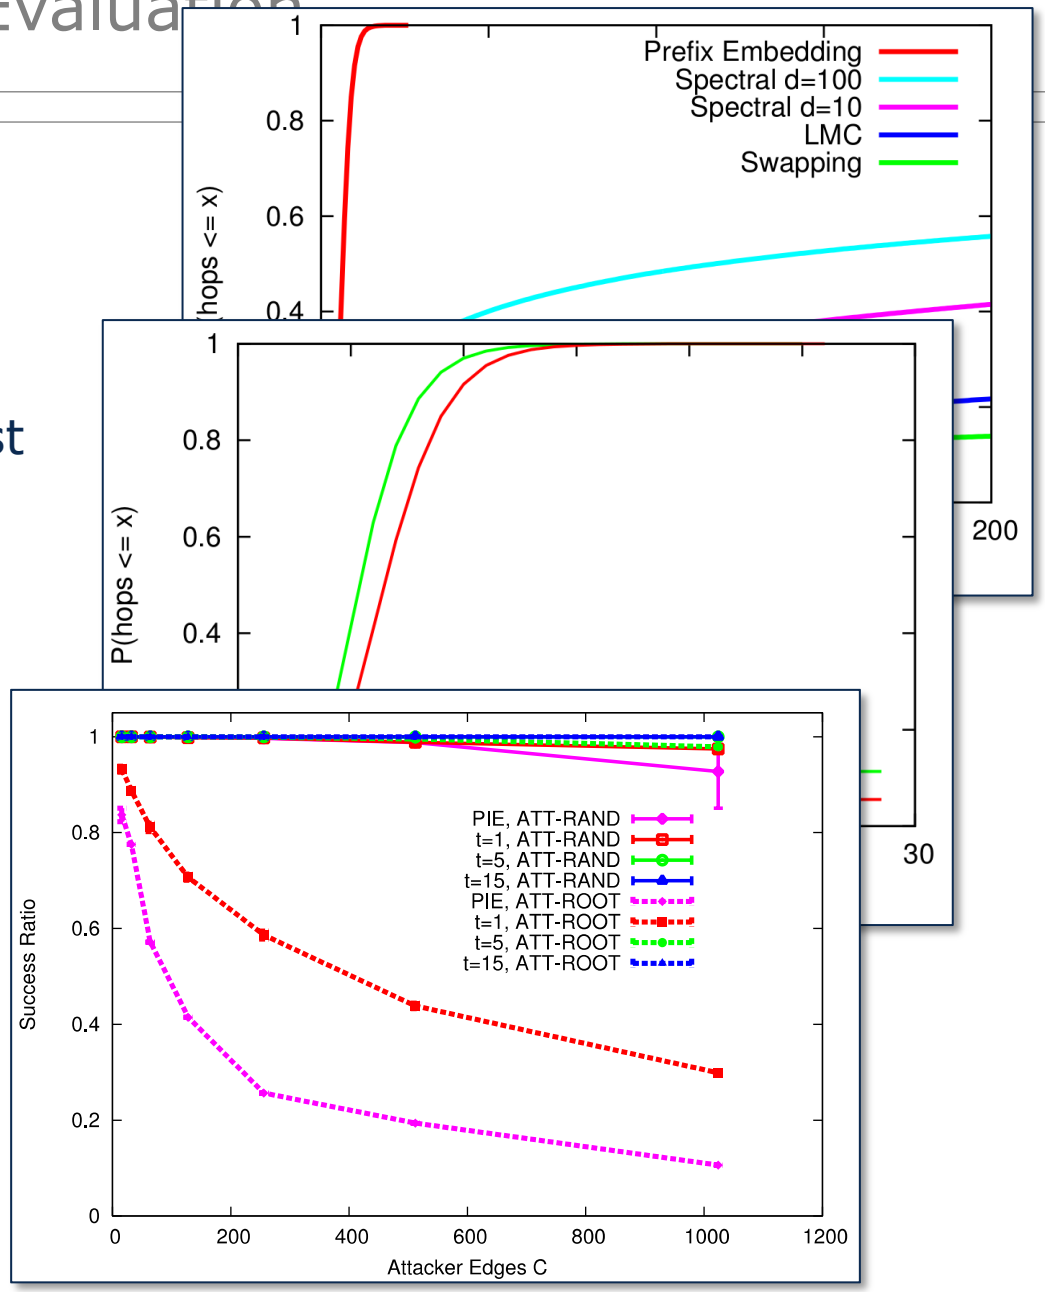

…so, are we nearly there?

Ask Martin! ;-)

My answer is complex:

- We have come far
  - embedding/routing works in simulations
  - We can build a virtual overlay on top (DHT works)
  - Reasonably stable to attacks
  - Reasonably good protection against leaks

- There‘s a lot left to be done
  - Availability (churn)
  - Performance and fairness (transfer over friends‘ links)
  - Friend „attacks“ (who‘s nosy, concern of users)
  - Extension to mobile devices
  - Get everything to run… ☺

TECHNISCHE UNIVERSITÄT DRESDEN

References

Leyla Bilge, Thorsten Strufe, Davide Balzarotti, and Engin Kirda. “All Your Contacts Are Belong to Us: Automated Identity Theft Attacks on Social Networks.”, In WWW, 2009

Cutillo, Leucio Antonio, et al. "Security and privacy in online social networks." Social Network Technologies and Applications. Springer US, 2010.

Paul, Thomas, et al. "Improving the usability of privacy settings in facebook." CoRR abs/1109.6046 (2011).

Paul, Thomas, et al. “C4PS - Colors for Privacy Settings”. In: WWW, 2012

Paul, Thomas, et al. "C4PS-helping Facebookers manage their privacy settings." International Conference on Social Informatics. Springer Berlin Heidelberg, 2012.

Paul, Thomas, Antonino Famulari, and Thorsten Strufe. "A survey on decentralized online social networks." Elsevier Computer Networks, 75, 2014

Paul, Thomas, et al. “Systematic, Large-scale Analysis on the Feasibility of Media Prefetching in Online Social Networks.” In: IEEE CCNC, 2015.

Paul, Thomas, et al. “The User Behavior in Facebook and its Development from 2009 until 2014.”, CoRR abs/1505.04943, 2015

Roos, Stefanie, et al. “Anonymous Addresses for Efficient and Resilient Routing in F2F Overlays.” In IEEE INFOCOM, 2016

Roos, Stefanie and Strufe, Thorsten. “On the impossibility of efficient self-stabilization in virtual overlays with churn.” In IEEE INFOCOM, 2015

Schulz, Stephan, and Thorsten Strufe. "d² Deleting Diaspora: Practical attacks for profile discovery and deletion." 2013 IEEE International Conference on Communications (ICC). IEEE, 2013.

All pictures credit wikimedia, unless stated differently