Privacy Enhancing Technologies

Carmela Troncoso, Gradiant

PRIPARE Workshop on Privacy by Design
Ulm 9th-10th March 2015

Outline

- What are privacy enhancing technologies?
- Privacy Enhancing Technologies
  - PETs for personal data management
  - PETs for data disclosure minimization
- Conclusions

What are privacy enhancing technologies?

What is privacy?

- So far in the workshop:
  - Abstract and subjective concept, hard to define
  - Popular definitions:
    - “The right to be let alone”: freedom from intrusion
    - “Informational self-determination” : focus on control
  - EU Regulation Data Protection Directive (95/46/EC)
    - What data can be collected and how should it be protected
  - Privacy controls: more detailed high level description

- And from a technical point of view?
  - Privacy properties

Privacy properties: **Anonymity**

- Hiding link between identity and action/piece of information.
  - Reader of a web page, person accessing a service
  - Sender of an email, writer of a text
  - Person to whom an entry in a database relates
  - Person present in a physical location

- Definitions:
  - Pfitzmann-Hansen (PH)[1] *"Anonymity is the state of being not identifiable within a set of subjects, the* anonymity set [...] The *anonymity set is the set of all possible subjects who might cause an* action" [pattern Anonymity set]
  - ISO 29100[2]"defines anonymity as a characteristic of information that does not permit a personally identifiable information principal to be identified directly or indirectly"

- In practice it is a Probabilistic definition

Privacy properties: **Pseudonymity**

- *PH*[1] *"Pseudonymity is the use of pseudonyms as IDs* [...] A *digital pseudonym is a bit string which is unique as ID and which can* be used to authenticate the holder" [pattern Pseudonymous identity]
- ISO15408[3] "pseudonymity ensures that a user may use a resource or service without disclosing its identity, but can still be accountable for that use. "

One time pseudonyms (Anonymity)

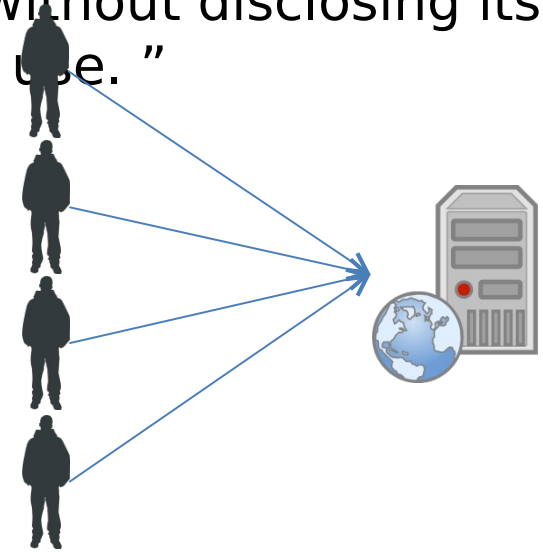

Persistent pseudonyms (Identity!)

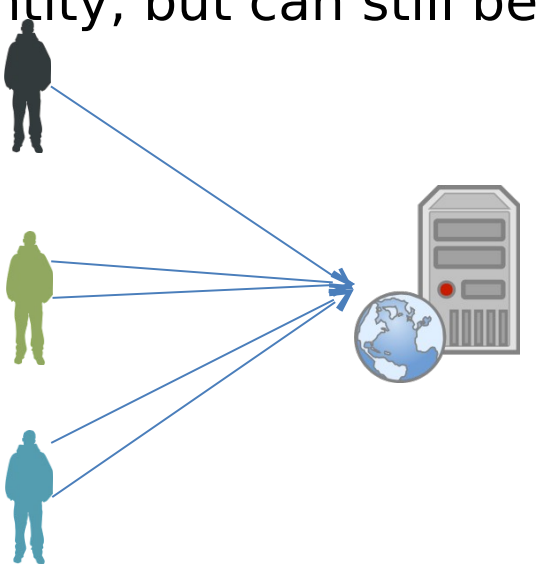

Hybrid (Multiple Identities)

Privacy properties: **Unlinkability**

- Hiding link between two or more actions/identities/info pieces
  - Two anonymous letters written by the same person
  - Two web page visits by the same user
  - Entries in two databases related to the same person
  - Two people related by a friendship link
  - Same person spotted in two locations at different points in time
- Definitions
  - PH[1] “Unlinkability of two or more items means that within a system, these items are no more and no less related than they are related concerning the a-priori knowledge”
  - ISO15408[3]“unlinkability ensures that a user may make multiple uses of resources or services without others being able to link these uses together.”

Privacy properties: **Unobservability**

- Hiding user activity.
  - whether someone is accessing a web page
  - whether an entry in a database corresponds to a real person
  - whether someone or no one is in a given location

- Definitions
  - PH[1] *"Unobservability is the state of items of interest being* indistinguishable from any item of interest at all [...] *Sender unobservability then means that it is not noticeable* whether any sender within the unobservability set sends."
  - ISO15408[3] "unobservability ensures that a user may use a resource or service without others, especially third parties, without being able to observe that the resource or service is being used."

Privacy properties: **Plausible deniability**

- Not possible to prove user knows, has done or has said something
  - Off-the-record conversations
  - Resistance to coercion:
    - Not possible to prove that a person has hidden information in a computer
    - Not possible to know that someone has the combination of a safe
  - Possibility to deny having been in a place at a certain point in time
  - Possibility to deny that a database record belongs to a person

Privacy properties

- So far it was about de-coupling identity and actions
- but we could keep identity and hide data
  - Cryptographic security properties
  - Not similar widely accepted for other means (the previous properties are building blocks)
- Differential privacy: a data base looks “almost” the same before and after an event occurs.
  - Special noise

Privacy enhancing technologies

- Technologies that enable users to preserve their privacy
  - In terms of technical properties

- From whom?
  1. Third parties = trust on data controller/processor (or must disclose data)

- PETs for data disclosure minimization (i.e., minimize trust)
- "Ultimate" Data Protection

Privacy enhancing technologies

- Technologies that enable users to preserve their privacy
  - In terms of technical properties

- From whom?
  1. Third parties = trust on data controller/processor (or must disclose data)
     - PETs for personal data management ["soft privacy"]
     - Support to Data Protection
  2. Data controller/processor = no trust
     - PETs for data disclosure minimization (i.e., minimize trust) ["hard privacy"]
     - "Ultimate" Data Protection

PETs for personal data management

PETs for decision support

- Provide insight in how user's data is being collected, stored, processed and disclosed to the data subject to enable well-informed decisions [pattern Protection against tracking]

- Transparency-Enhancing Technologies[4]
  - *Google Dashboard*: what personal data is stored and who has access
  - *Collusion (Firefox addon)*: list of entities tracking users
  - *Mozilla Privacy Icons:* simple visual language to make privacy policies more understandable
  - *Privacy Bird (IE Add-on):* shows user whether webpage complies with her preferred policy based on images

- Challenges
  - How to provide information useful to users
    - How to convey it
    - How to make users understand

Privacy as Control
Privacy as Practice

PETs for consent support

- Provide users with means to express their privacy preferences and give consent [pattern Protection against tracking]

- Privacy policies languages (P3P, S4P, SIMPL)
  - Automated processing and comparison with users' preferences
  - Difficult to make unambiguous and inform users (TETs)
  - Difficult to standardize and make them expressive

- Anti-tracking
  - Do Not Track options
    - Browser tag expressing who can collect persona
  - Track Me Not plugin
    - Renders collection useless

Privacy as Control
Privacy as Practice

PETs for enforcement support

- Provide users with means to enforce their preferences
- Locally “easy”: blockers (pop-ups, ads, cookies,...)
- Remotely
  - Sticky policies associated to data(e.g., trusted third party stores encryption keys only disclosed in certain cases)
  - Use of trusted hardware (HSMs, TPMs) to process data “out” of the server’s control

Privacy as Control Privacy as Practice

PETs for accountability support

- Data controllers should be able to demonstrate compliance with Data Protection.

- Non repudiable logs
  - Backups, distributed logging
  - Forward integrity (hash chains)

- Verifiable Audits
  - Automated tools for log audits

Data Management vs. Minimization

- Previous techniques are applied once personal data has been disclosed
- Aim at:
  - Help the user understand and decide
  - Make data controllers more responsible
- But they cannot guarantee that privacy is not lost
- Can we reduce the amount of data disclosed?

PETs for personal data disclosure minimization

Privacy as confidentiality!

Anonymous credentials

- Authentication is the first step before any security policy can be applied

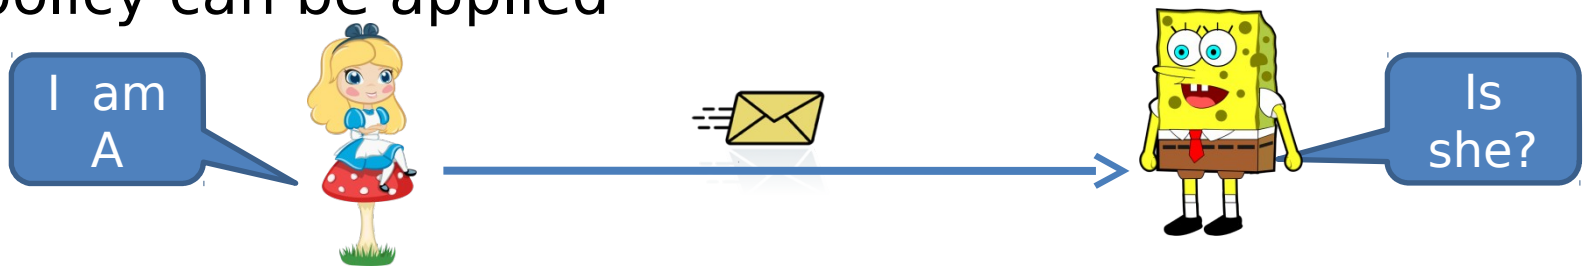

- Makes sense in government, military, even commercial
  - ...but if there is no closed group? (e.g., peer-to-peer)
  - The **Identity Management** concept

- Possible solutions:
  - Private authentication: hide against 3rd parties
  - Anonymous credentials: protect against everybody

Idea behind credentials

- Many transactions involve attribute certificates
  - ID docs: state certifies name, birth dates, address
  - Letter reference: employer certifies salary
  - Club membership: club certifies some status
  - PKI certificate: RRN in Belgian eID, NIF in Spain
- Do you want to show all of them?
- Credential: token certifying one attribute
  - e.g. ticket to the cinema (“i have paid”)
  - Digital credentials: string, boolean attributes, range

Properties

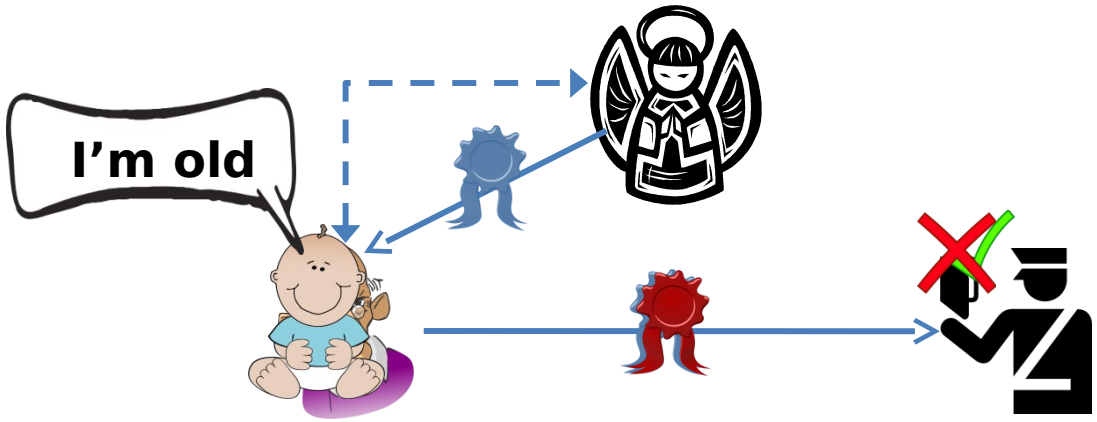

- **Completeness**: if the statement is true, the verifier will be convinced
- **Zero-knowledge**: if the statement is true no cheating verifier learns anything other than this fact
- **Soundness**: no cheating prover can convince the honest verifier

- **Unlinkability**: two requests cannot be linked to the same user

- Holds even if verifier and prover collide

Zero-knowledge proofs

- One party to prove to another that a statement is true, without revealing anything other than the veracity of the statement.

- J.J. Quisquater: "How to Explain Zero-Knowledge Protocols to Your Children"

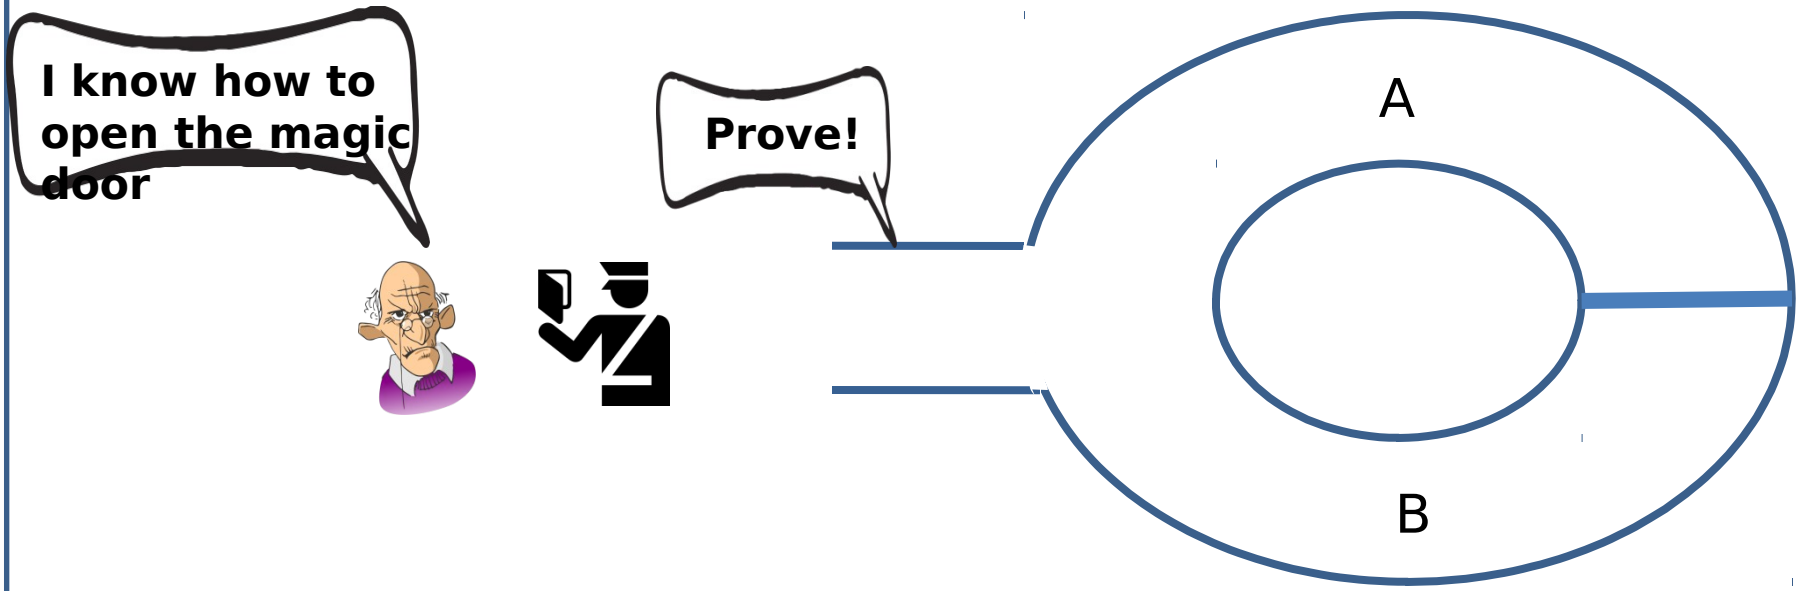

Zero-knowledge proofs

- One party to prove to another that a statement is true, without revealing anything other than the veracity of the statement.

- J.J. Quisquater: "How to Explain Zero-Knowledge Protocols to Your Children"

- If there are doubts repeat!
  - 50% chance
  - Likelihood decreases

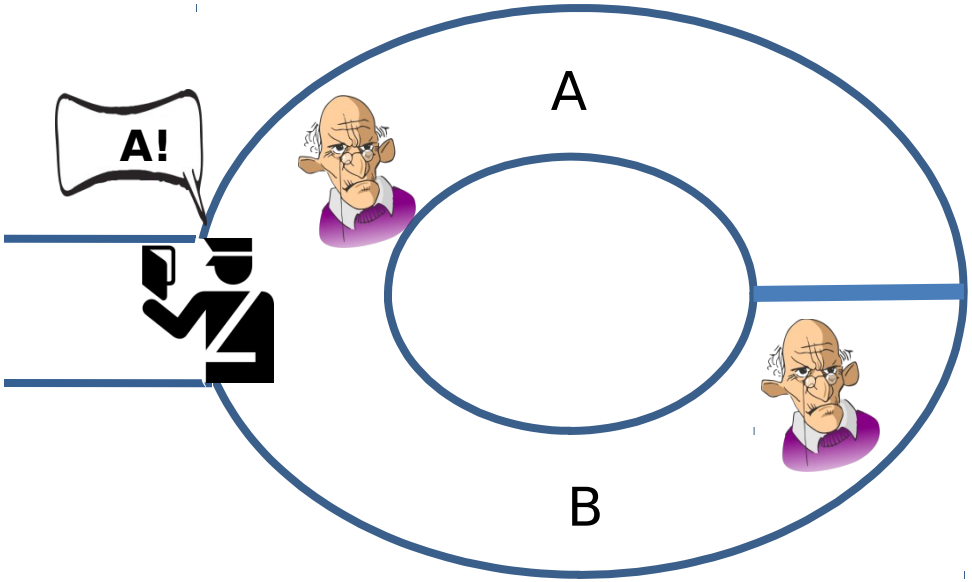

Optional properties

- **Revokation**: some schemes allow for revokation of credential
  - Total revokation
  - Blacklisting
- **Linkability**: some schemes allow to link credential shows
- **Limited shows**: some schemes allow to limit the number of shows
- **Re-identification**: some schemes allow to de-anonymize the subject

PKI vs Anonymous Credentials

PKI

Signed by a trusted issuer

Certification of attributes

Authentication (secret key)

Double-signing detection

No data minimization

Users are identifiable

Users can be tracked
(Signature linkable to other contexts where PK is used)

Anonymous credentials

Signed by a trusted issuer

Certification of attributes

Authentication (secret key)

Double-signing detection

Data minimization

Users are anonymous

Users are unlinkable in different contexts

Other privacy-preserving crypto

- Private Information Retrieval
  - Query databases without revealing query
- Multiparty computation
  - Group computation where only result is revealed
- Cryptographic commitments
  - “Vaults” that allow to commit to secret values
- eCash
  - Digital cash with anonymity and unlinkablity properties (like real cash!)
- Private set intersection
  - Find matching elements in sets without revealing further information

Anonymous communications

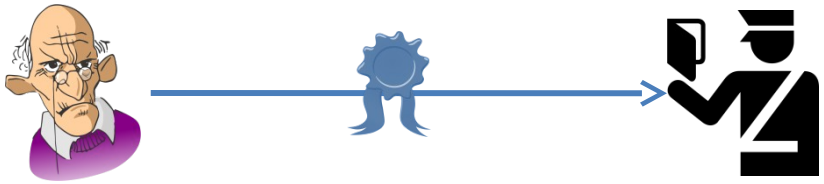

- Hidden assumptions
  - Secure channel
  - The channel does not break the privacy property
- But IP is a pseudo-identifier!
  - anonymous credentials are useless in this case...
- Need protection against **traffic analysis**
  - the military also use internet...

Traffic analysis

- Even if communication is encrypted, traffic data can reveal a lot of information: source, destination, timing, volume, etc.
- Examples from WW Ii:
  - British at Bletchley Park assesing the size of Germany's air-force
  - Discover/Uncover inminent actions
    - Japanese countermeasures key in Pearl Harbour (1941)
    - D-day decoys
  - Identifying people by their typing
- Examples from today
  - Amazon profiling based on clicks and hoovers
  - Fraud analysis in banks and Credit card companies

System model

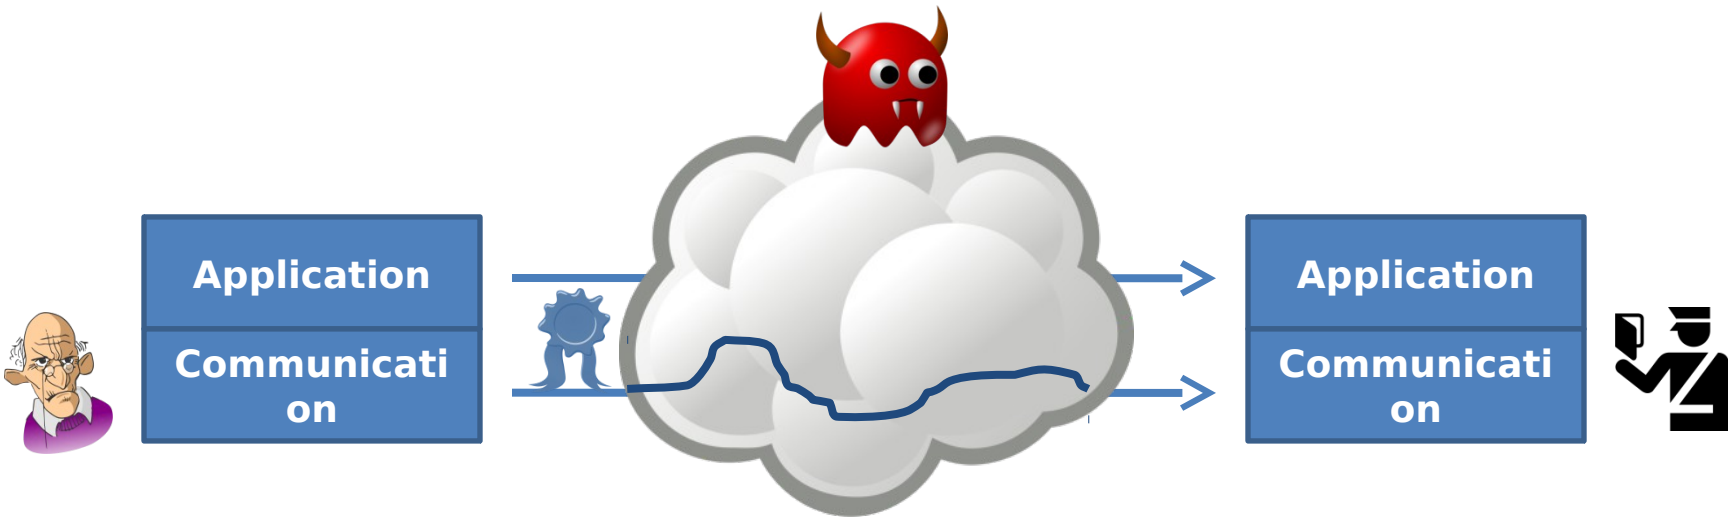

Attacker assumptions

- Attacker abilities:
  - Observe
    - All links (*Global Passive Adversary)*
    - *Some links*
  - Modify, delay, delete or inject messages.
  - Control some nodes in the network.
- Attacker limitations:
  - Cannot break cryptographic primitives.
  - Cannot see inside nodes he does not control.

Onion encryption

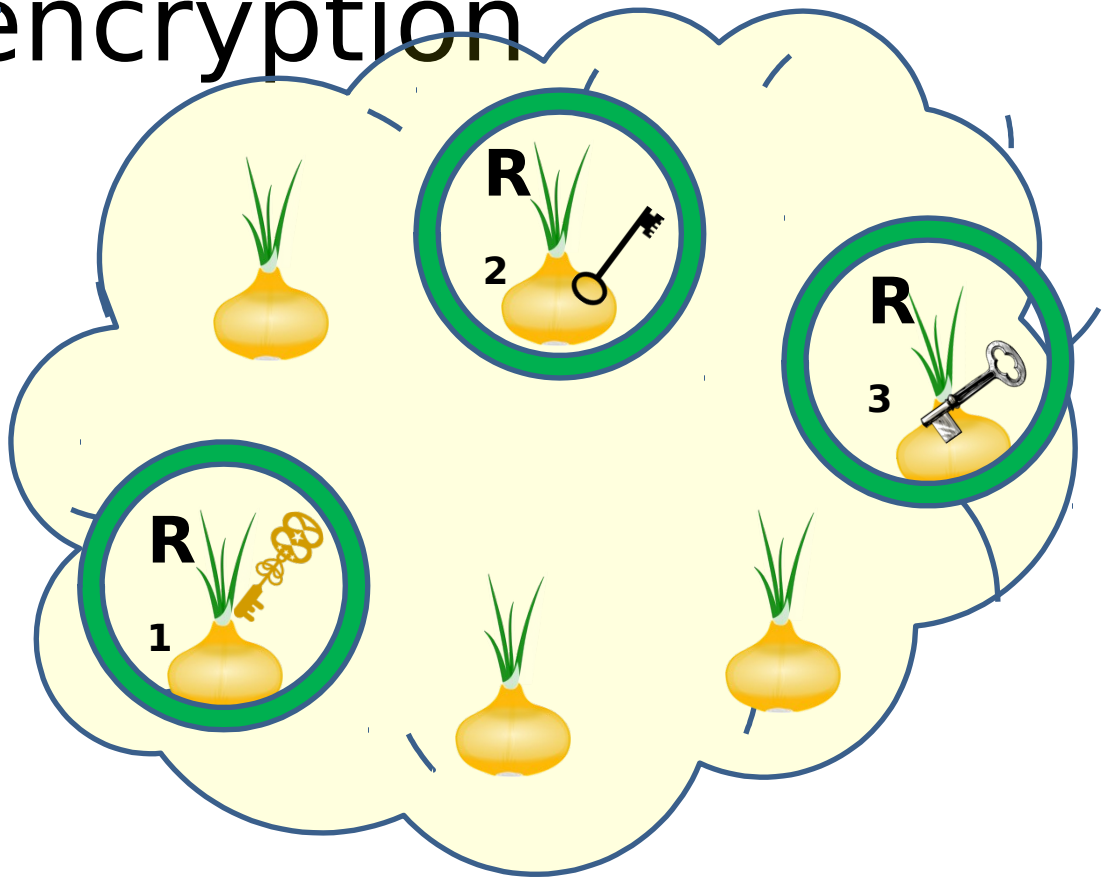

D

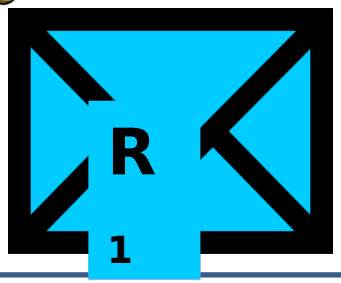

Onion Routing

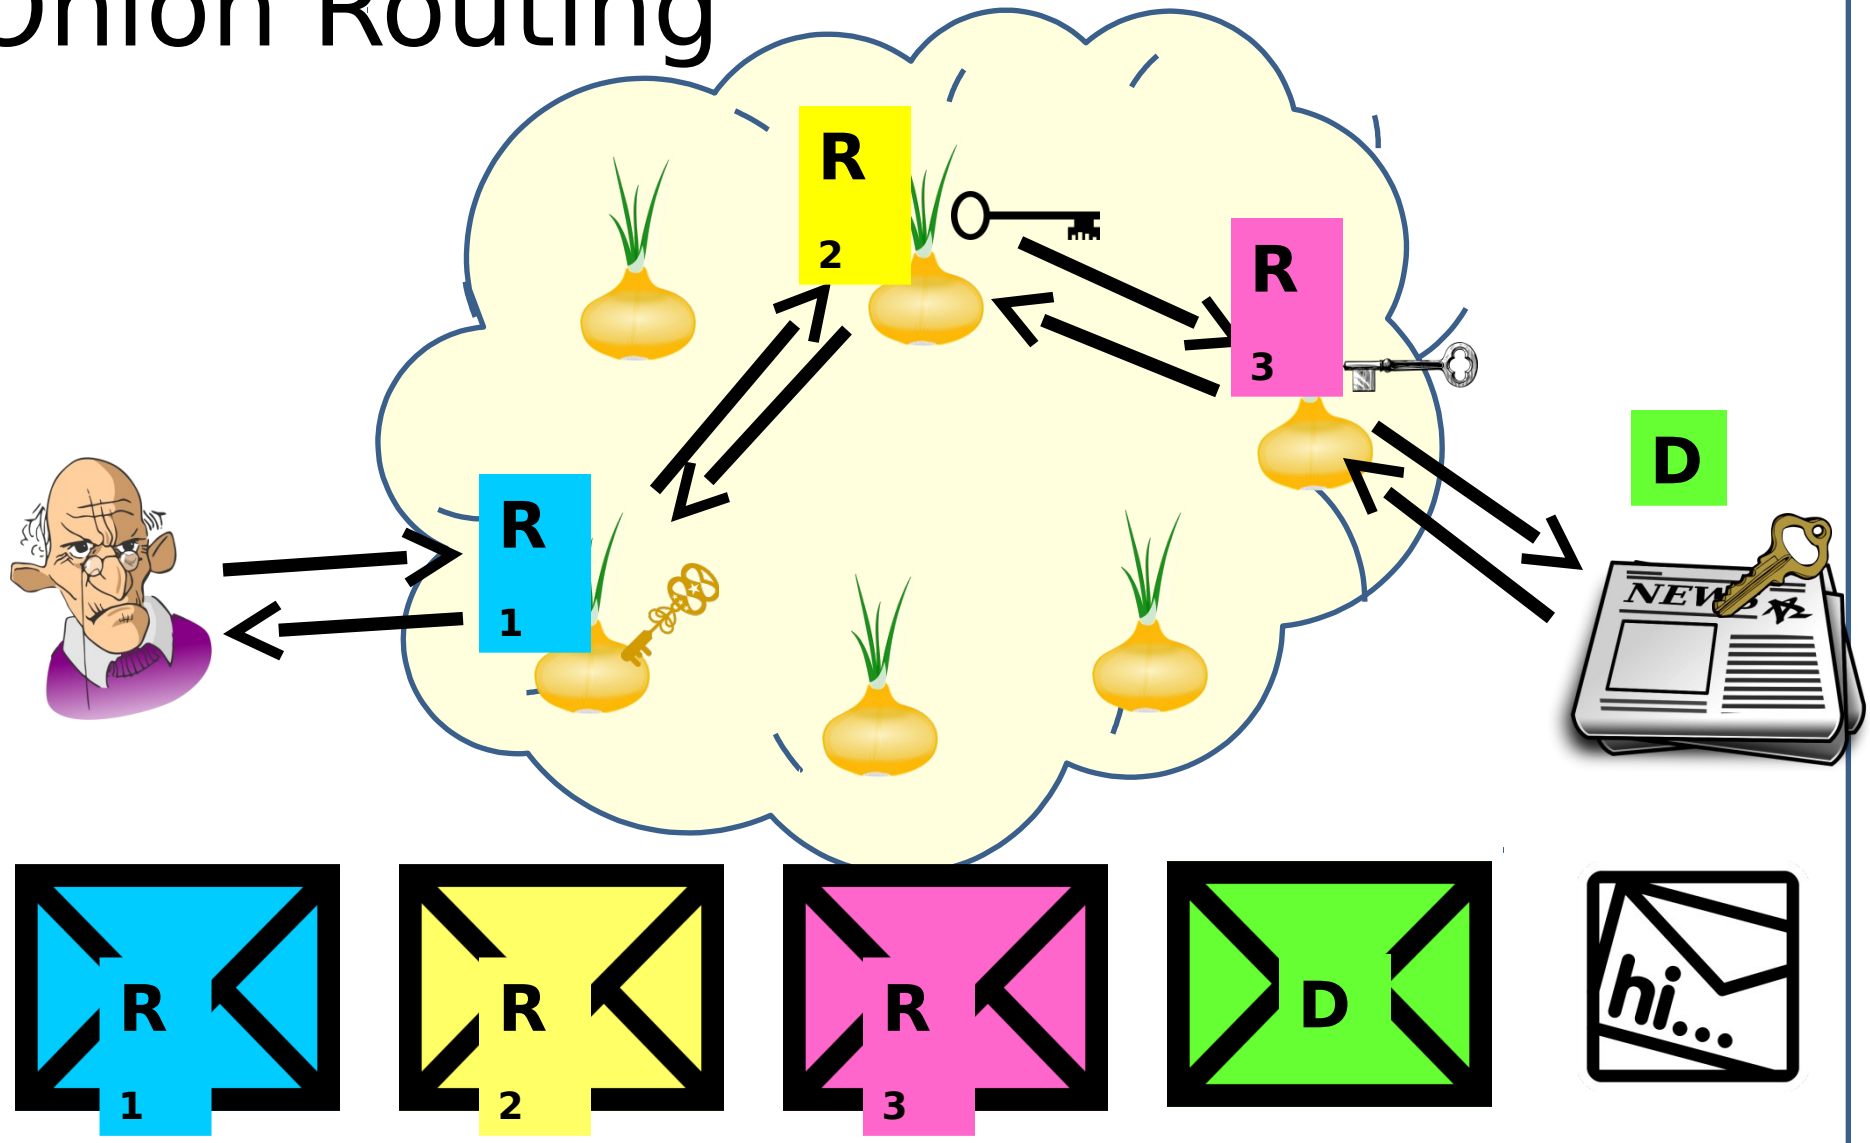

TOR – adversary model

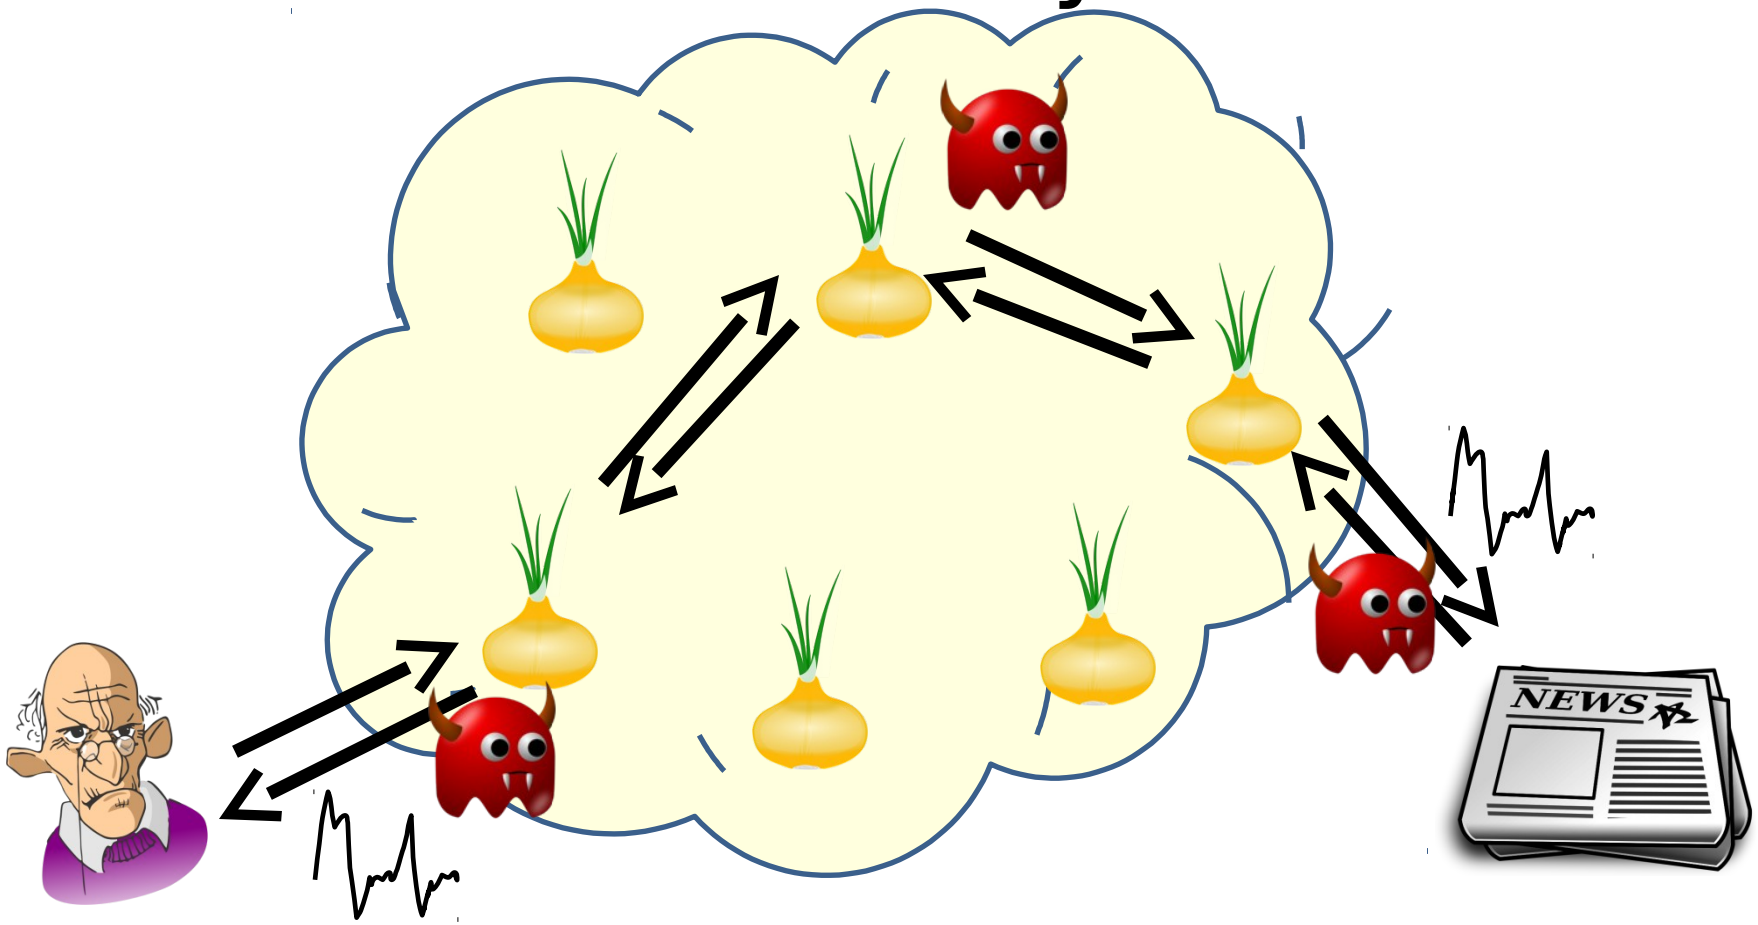

Data Anonymization

- Gzillion anonymization techniques
  - Remove identifier (removing, hashing, encrypting)
  - Add noise
  - Modify graph information
  - Generalise (k-anonymity, cloaking, …)

Art. 29 WP's opinion on anonymization techniques

3 criteria to decide a dataset is non-anonymous (pseudonymous):

- is it still possible to single out an individual,
- is it still possible to link two records within a dataset (or between two datasets)
- can information be inferred concerning an individual?

Singling out - metadata tends to be unique

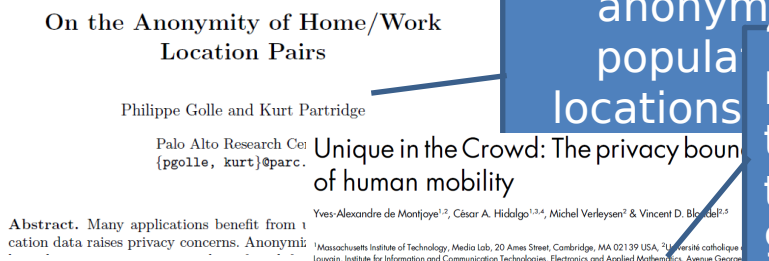
On the Anonymity of Home/Work Location Pairs

Philippe Golle and Kurt Partridge

Palo Alto Research Ce
{pgolle, kurt}@parc.

Abstract. Many applications benefit from u
cation data raises privacy concerns. Anonymi

the median size of the individual's anonymity set in the U.S. working popula
locations

Unique in the Crowd: The privacy boun
of human mobility

Yves-Alexandre de Montjoye[1,2], César A. Hidalgo[1,3,4], Michel Verleysen[2] & Vincent D. Bl

[1]Massachusetts Institute of Technology, Media Lab, 20 Ames Street, Cambridge, MA 02139 USA, [2]Université catholique
Louvain, Institute for Information and Communication Technologies, Electronics and Applied Mathematics, Avenue George
Lemaître 4, B-1348 Louvain-la-Neuve, Belgium, [3]Harvard University, Center for International Development, 79 JFK Street,
Cambridge, MA 02138, USA, [4]Instituto de Sistemas Complejos de Valparaíso, Paseo 21 de Mayo, Valparaíso, Chile,
[5]Massachusetts Institute of Technology, Laboratory for Information and Decision Systems, 77 Massachusetts Avenue, Cam
MA 02139, USA.

We study fifteen months of human mobility data for one and a half million individuals and find that human mobility traces are highly unique. In fact, in a dataset where the location of an individual is specified hourly, and with a spatial resolution equal to that given by the carrier's antennas, four spatio-temporal points are enough to uniquely identify 95% of the individuals. We coarsen the data spatially and temporally to find a

"if the location of an individual is specified hourly, and with a spatial resolution equal to that given by the carrier's antennas, four spatio-temporal points are enough to uniquely identify 95% of the individuals." [15 montsh, 1.5M people]"

Location

Privacy Working Paper 3. Pittsburgh 2000.

Simple Demographics Often Identify People Uniquely

Demographics

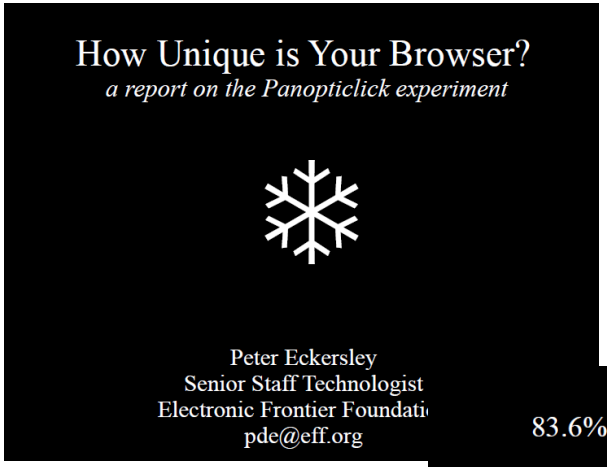
How Unique is Your Browser?
a report on the Panopticlick experiment

Peter Eckersley
Senior Staff Technologist
Electronic Frontier Foundati
pde@eff.org

83.6% had completely unique fingerprints
(entropy: 18.1 bits, or more)

94.2% of "typical desktop browsers" were unique
(entropy: 18.8 bits, or more)

Web browser


Latanya Sweeney
Carnegie Mellon University
latanya@andrew.cmu.edu

"It was found that 87% (216 million of 248 million) of the population in the United States had reported characteristics that likely made them unique based only on {5-digit ZIP, gender, date of birth}"

Link records relating to an individual

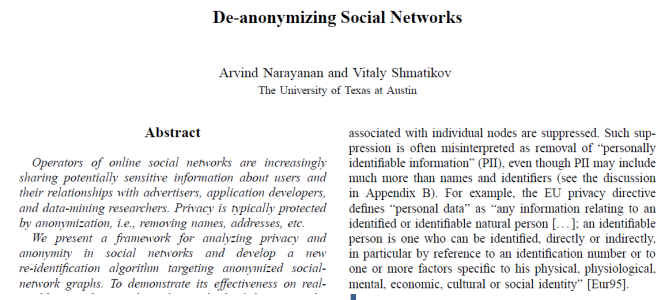

De-anonymizing Social Networks

Arvind Narayanan and Vitaly Shmatikov
The University of Texas at Austin

Abstract

Operators of online social networks are increasingly sharing potentially sensitive information about users and their relationships with advertisers, application developers, and data-mining researchers. Privacy is typically protected by anonymization, i.e., removing names, addresses, etc.

We present a framework for analyzing privacy and anonymity in social networks and develop a new re-identification algorithm targeting anonymized social-network graphs. To demonstrate its effectiveness on real-

associated with individual nodes are suppressed. Such suppression is often misinterpreted as removal of "personally identifiable information" (PII), even though PII may include much more than names and identifiers (see the discussion in Appendix B). For example, the EU privacy directive defines "personal data" as "any information relating to an identified or identifiable natural person [. . .]; an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity" [Eur95].

take two graphs representing social networks and map the nodes to each other based on the *graph structure alone*—no usernames, no nothing

Netflix Prize, Kaggle contest

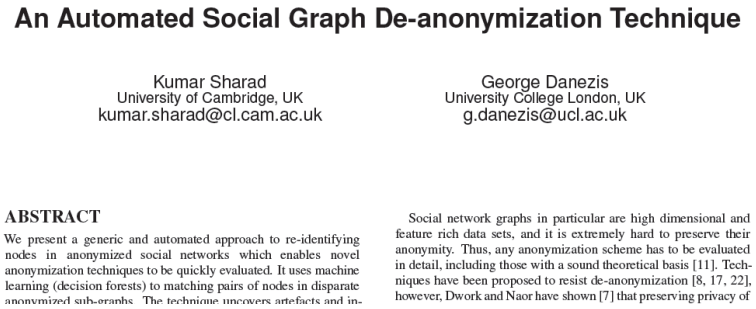

An Automated Social Graph De-anonymization Technique

Kumar Sharad
University of Cambridge, UK
kumar.sharad@cl.cam.ac.uk

George Danezis
University College London, UK
g.danezis@ucl.ac.uk

ABSTRACT

We present a generic and automated approach to re-identifying nodes in anonymized social networks which enables novel anonymization techniques to be quickly evaluated. It uses machine learning (decision forests) to matching pairs of nodes in disparate anonymized sub-graphs. The technique uncovers artefacts and in-

Social network graphs in particular are high dimensional and feature rich data sets, and it is extremely hard to preserve their anonymity. Thus, any anonymization scheme has to be evaluated in detail, including those with a sound theoretical basis [11]. Techniques have been proposed to resist de-anonymization [8, 17, 22], however, Dwork and Naor have shown [7] that preserving privacy of

Technique to automate graph de-anonymization based on machine learning. Does not need to know the algorithm!

Inferring information about an individual

the latitude and longitude of their homes. From these locations, we used a free Web service to do a reverse "white pages" lookup, which takes a latitude and longitude coordinate as input and gives an address and name. [172

Inference Attacks on Location Tracks

John Krumm

Microsoft Research
One Microsoft Way
Redmond, WA, USA
jckrumm@microsoft.com

Abstract. Although the privacy threats and countermeasures associated with location data are well known, there has not been a thorough experiment to assess the effectiveness of either. We examine location data gathered from volunteer subjects to quantify how well four different algorithms can identify

"We investigate the subtle cues to user identity that may be exploited in attacks on the privacy of users in web search query logs. We study the application of simple classifiers to map a sequence of queries into the gender, age, and location of the user issuing the queries."

…w What You Did Last Summer" — Query Logs and User Privacy

Rosie Jones Ravi Kumar Bo Pang Andrew Tomkins
Yahoo! Research, 701 First Ave, Sunnyvale, CA 94089.
{jonesr,ravikumar,bopang,atomkins}@yahoo-inc.com

…e subtle cues to user identity that may be exploited …privacy of users in web search query logs. We …tion of simple classifiers to map a sequence of queries into the gender, age, and location of the user issuing the queries. We then show how these classifiers may be carefully combined at multiple granularities to map a sequence of queries into a

bilities; this is the goal of this paper. We initiate the study of subtle cues to user identity that exist as vulnerabilities in web search query logs, which may be exploited in attacks on the privacy of users.

Privacy attack models. We begin with a characterization of two key forms of attack against which a query log privacy scheme must be resilient. The first is a *trace attack*, in which an attacker studies a privacy-enhanced version of a sequence of searches (*trace*) made

Anonymization bottom line

- There is no known best method to anonymize and release data
  - Probably there is no way to anonymize… [Dwork et al]
- Need to quantify the information that may leak
  - Probabilistic analysis
  - Most often need for case by case analysis

Summary

- Privacy from a technical perspective: privacy properties
- Privacy Enhancing technologies
  - Enable protection of privacy
- PETs for personal data management
  - Require trust in service provider
  - State of the art in development
    - Hidden costs of securing the data silos
    - Hidden costs of public image when things go wrong
- PETs for data disclosure minimization
  - Limit trust in providers and other users (**Adversarial models!**)
  - Anonymous Credentials
  - Anonymous communications
  - Data anonymization

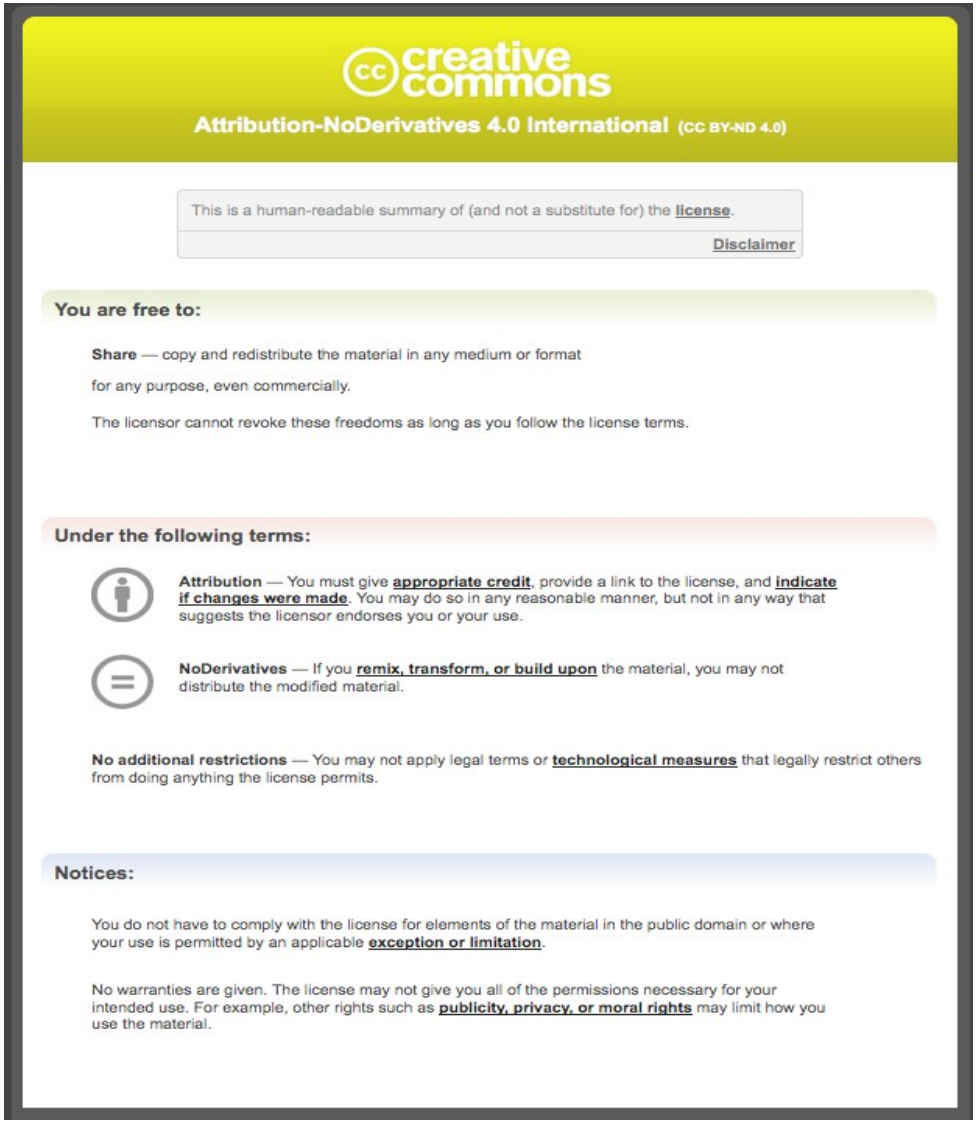

creative commons

Attribution-NoDerivatives 4.0 International (CC BY-ND 4.0)

This is a human-readable summary of (and not a substitute for) the license.

Disclaimer

You are free to:

Share — copy and redistribute the material in any medium or format

for any purpose, even commercially.

The licensor cannot revoke these freedoms as long as you follow the license terms.

Under the following terms:

Attribution — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

NoDerivatives — If you remix, transform, or build upon the material, you may not distribute the modified material.

No additional restrictions — You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Notices:

You do not have to comply with the license for elements of the material in the public domain or where your use is permitted by an applicable exception or limitation.

No warranties are given. The license may not give you all of the permissions necessary for your intended use. For example, other rights such as publicity, privacy, or moral rights may limit how you use the material.

References

1. Pfitzmann, Andreas and Hansen, Marit. A terminology for talking about privacy by data minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management. 2010.
2. International Organization for Standardization (ISO), Information technology – Security techniques – Privacy framework, ISO/IEC 29100:2011, First edition, Geneva, 15 Dec 2011.
3. International Organization for Standardization (ISO), Information technology – Security techniques – Evaluation criteria for IT security, ISO/IEC 15408-1, Third edition, Geneva, 2009.
4. Milena Janic, Jan Pieter Wijbenga, Thijs Veugen: Transparency Enhancing Tools (TETs): An Overview. STAST 2013: 18-25

Location Privacy

- Emerging Location Based Services:
  - e-Call, VII, traffic congestion control
  - Nearby...
  - Variable pricing applications (congestion pricing, pay-as-you-drive)
  - Social applications

- What can be automatically inferred about a person based on location?
  - Any important location…
    - Desk in a building [BeresfordStajano03]
    - Home location [Krumm07, Hoh et al06]
    - Future locations [Krumm06]
  - Do you want to be seen at certain locations? AIDS clinic, business competitor, or political headquarters (Google Street View)

- One pseudonym per location exposure is not enough
  - Real time
  - Space-Time relation
  - Dummy traffic?

Defenses: Location Perturbation

- Policy-based location privacy protection requires trust
- Main ideas
  - Applications can tolerate *inaccurate location data to a certain* degree
  - Location perturbation hinders inferences on exact location
- Approaches:
  - Simple perturbation
    - Discretization
    - Random noise
  - Spatial Cloaking
  - Spatio-temporal Cloaking
  - Many more…

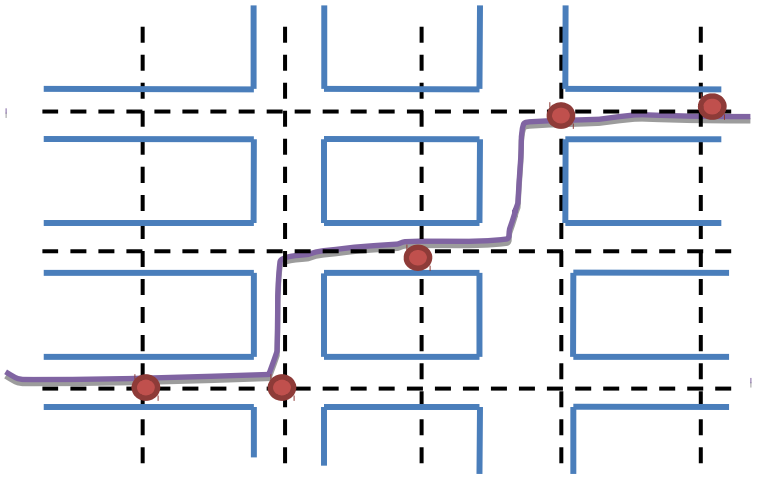

Defenses: Location Perturbation

- Policy-based location privacy protection requires trust
- Main ideas
  - Applications can tolerate *inaccurate location data to a certain* degree
  - Location perturbation hinders inferences on exact location
- Approaches:
  - Simple perturbation
    - Discretization
    - Random noise
  - Spatial Cloaking
  - Spatio-temporal Cloaking
  - Many more…

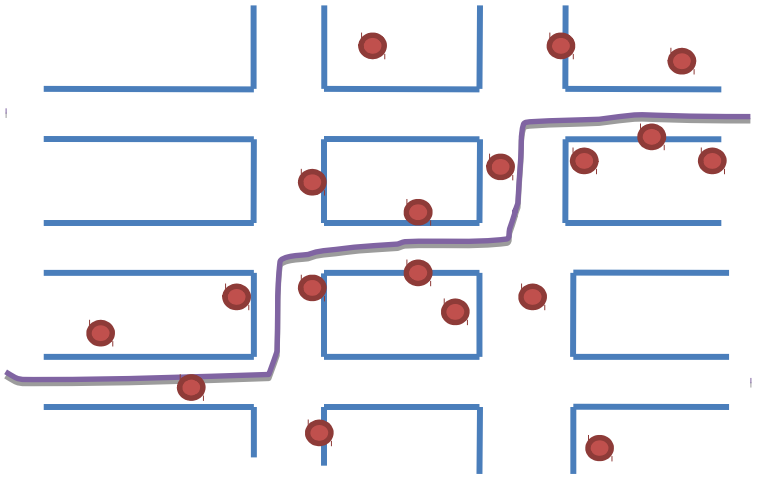

Defenses Cloaking

- Reveal a region instead of a particular place.
  - Many ways to define the region
  [pattern Location granularity]
  - Implementations

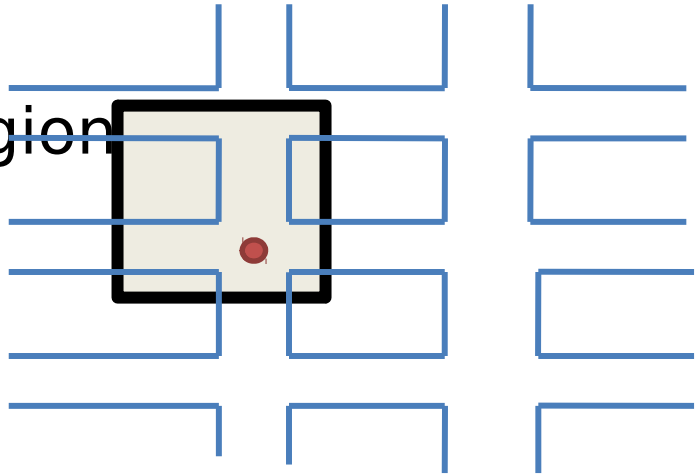

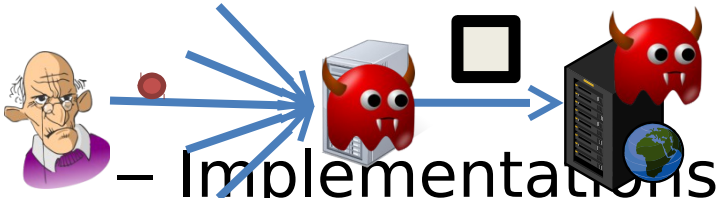

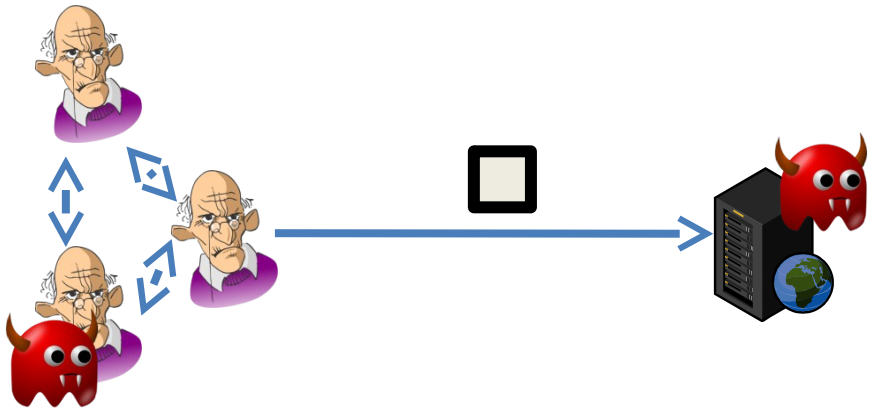

Concept of Mix (Chaum 1982)

Router that hides correspondence between inputs and outputs

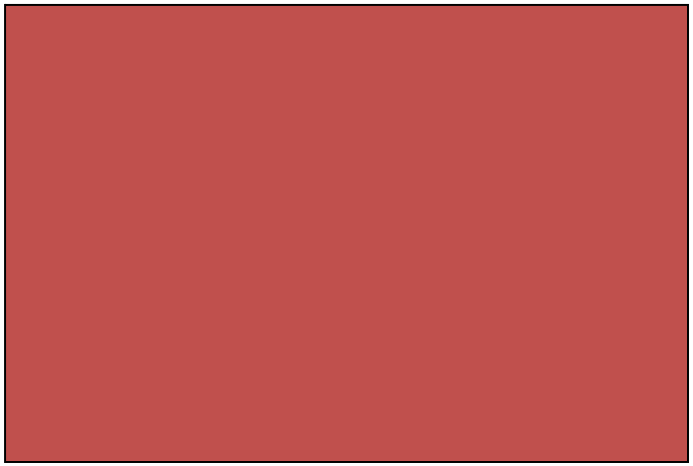

Concept of Mix: mix and flush

Router that hides correspondence between inputs and outputs

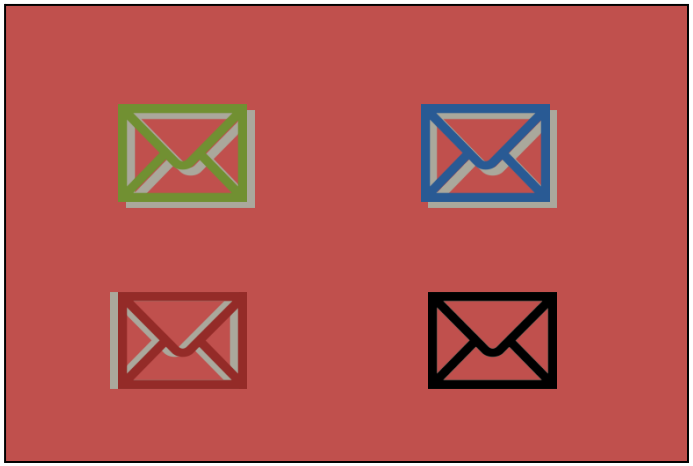

Deployed mix systems

- Mixmaster
- Mixminion