Systematic design of privacy-preserving systems

Carmela Troncoso

institute
imdea
software

Modern times – daily life

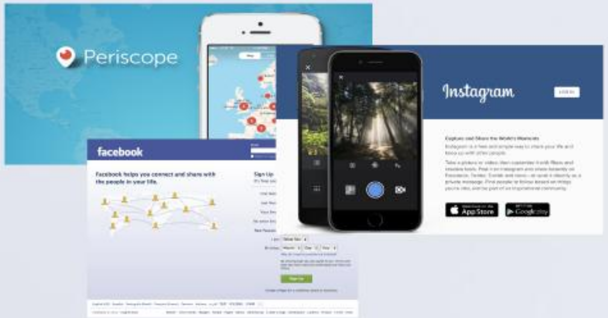

Modern times – daily life

Modern times – daily life

Modern times – daily life

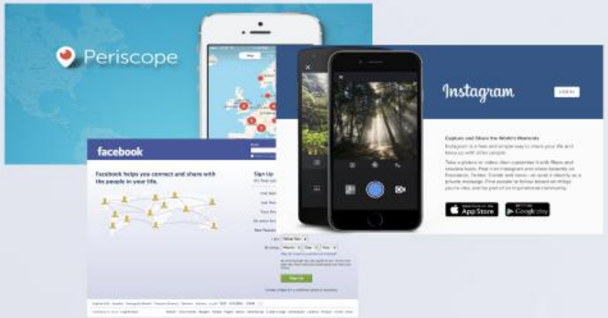

Modern times – daily life

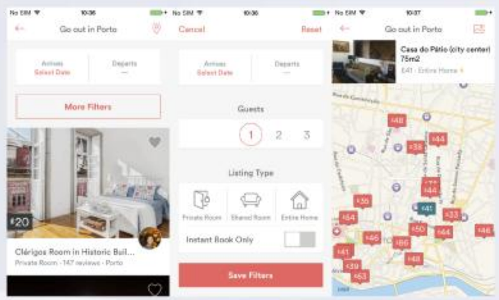

Modern times – daily life

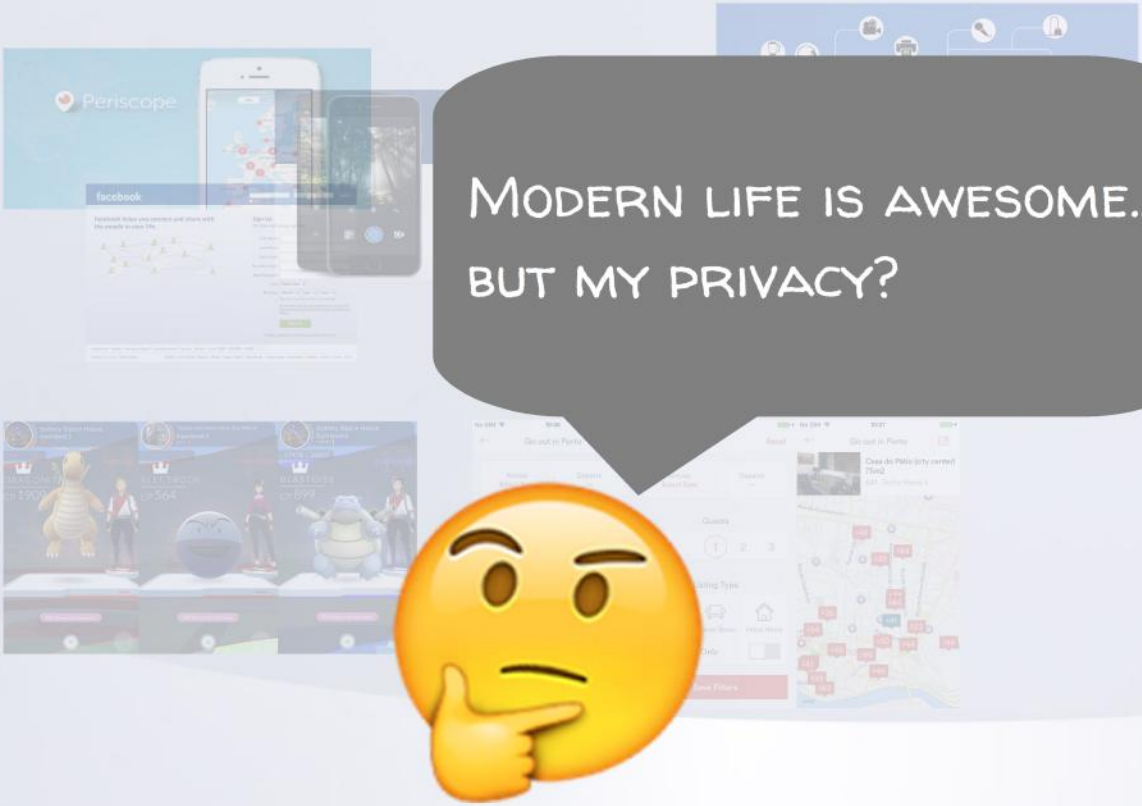

Modern life is awesome…

but my privacy?

Modern times – daily life

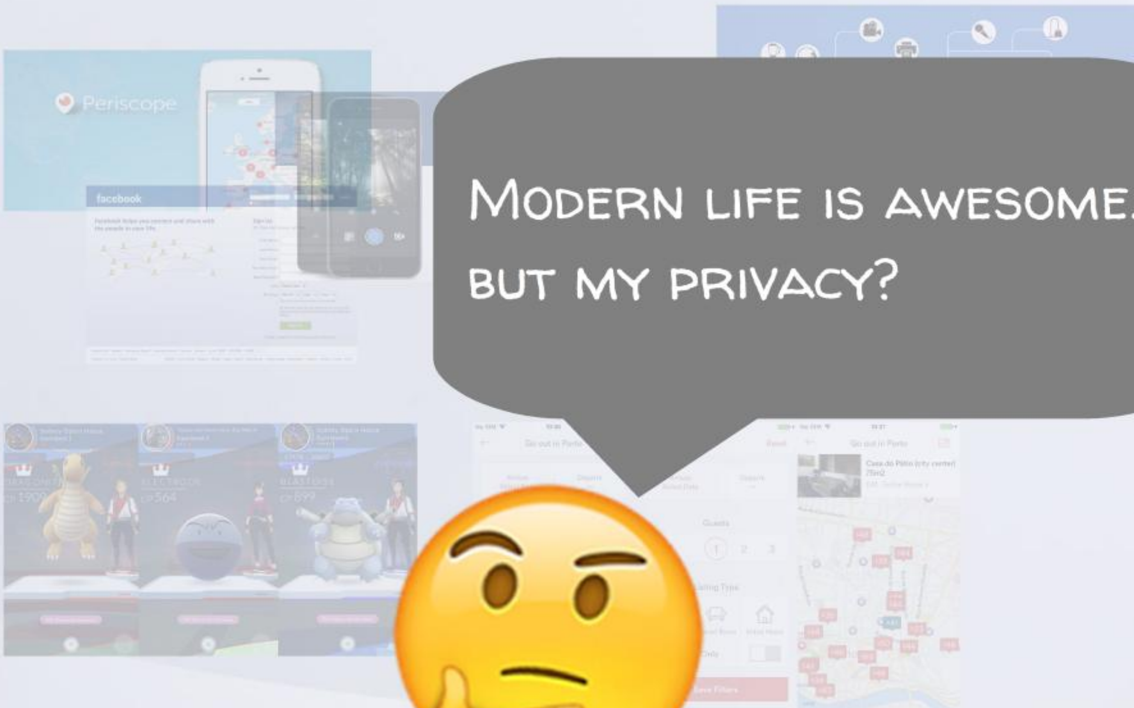

Modern life is awesome…

but my privacy?

"Personal data is the new oil of the Internet and the new currency of the digital world."

Meglena Kuneva

European Consumer Commissioner

Modern times – daily life

Modern life is awesome…

but my privacy?

"Personal data is the new oil of the Internet and the new currency of the digital world."

Meglena Kuneva

European Consumer Commissioner

Privacy by Design – Let's have it!

Information and Privacy Commissioner of Ontario

PbD

Privacy by Design

Privacy by Design principles

1. Proactive not Reactive; Preventative not Remedial
2. Privacy as the Default Setting
3. **Privacy Embedded into Design**
4. Full Functionality: Positive-Sum, not Zero-Sum
5. End-to-End Security — Full Lifecycle Protection
6. Visibility and Transparency — Keep it Open
7. Respect for User Privacy — Keep it User-Centric

Cavoukian et al. (2010)

https://www.ipc.on.ca/images/resources/7foundationalprinciples.pdf

Privacy by Design – Let's have it!

Information and Privacy Commissioner of Ontario

PbD

Privacy by Design

Privacy by Design principles

1. Proactive not Reactive; Preventative not Remedial
2. Privacy as the Default Setting
3. **Privacy Embedded into Design**
4. Full Functionality: Positive-Sum, not Zero-Sum
5. End-to-End Security — Full Lifecycle Protection
6. Visibility and Transparency — Keep it Open
7. Respect for User Privacy — Keep it User-Centric

Cavoukian et al. (2010)

Article 25 European General Data Protection Regulation

GDPR

EU General Data Protection Regulation

"the controller shall [...] implement appropriate technical and organisational measures […] which are designed to implement data-protection principles[...] in order to meet the requirements of this Regulation and protect the rights of data subjects."

https://www.ipc.on.ca/images/resources/7foundationalprinciples.pdf
http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679&from=EN

Privacy by Design – Let's have it!

Information and Privacy Commissioner of Ontario

Privacy by Design

Privacy by Design principles

1. Proactive not Reactive; Preventative not Remedial
2. Privacy as the Default Setting
3. **Privacy Embedded into Design**
4. Full Functionality: Positive-Sum, not Zero-Sum
5. End-to-End Security — Full Lifecycle Protection
6. Visibility and Transparency — Keep it Open
7. Respect for User Privacy — Keep it User-Centric

Cavoukian et al. (2010)

Article 25 European General Data Protection Regulation

GDPR

EU General Data Protection Regulation

"the controller shall [...] implement appropriate technical and organisational measures […] which are designed to implement data-protection principles[...] in order to meet the requirements of this Regulation and protect the rights of data subjects."

Actually… "Data Protection by design and by default"

https://www.ipc.on.ca/images/resources/7foundationalprinciples.pdf
http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679&from=EN

Privacy by Design – Let's have it!

Information and Privacy Commissioner of Ontario

PbD
Privacy by Design

Privacy by Design principles

1. Proactive not Reactive; Preventative not Remedial
2. Privacy as the Default Setting
3. **Privacy Embedded into Design**
4. Full Functionality: Positive-Sum, not Zero-Sum
5. End-to-End Security — Full Lifecycle Protection
6. Visibility and Transparency — Keep it Open
7. Respect for User Privacy — Keep it User-Centric

Cavoukian et al. (2010)

Article 25 European General Data Protection Regulation

GDPR
EU General Data Protection Regulation

"the controller shall [...] implement appropriate technical and organisational measures […] which are designed to implement data-protection principles[...] in order to meet the requirements of this Regulation and protect the rights of data subjects."

Actually… "Data Protection by design and by default"

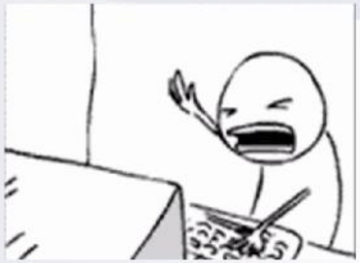

BUT HOW ??????????

https://www.ipc.on.ca/images/resources/7foundationalprinciples.pdf
http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679&from=EN

PART I:
Reasoning about Privacy when designing systems

This talk: Engineering Privacy by Design

PART II:
Designing Technologies to support Privacy-aware designs

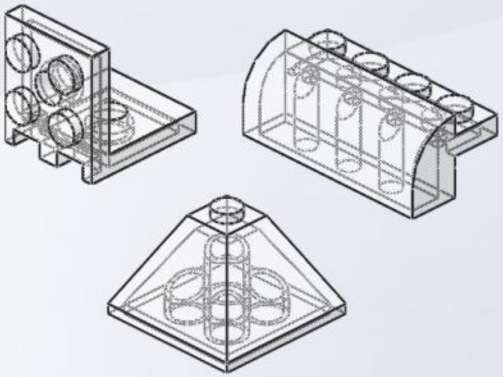

PART I:
Reasoning about Privacy when designing systems

PRIVACY


PoPETs
IEEE S
25th USENIX
SECURITY SYMPOSIUM
N D S S
SYMPOSIUM
acm
SIGSAC
SECURITY & PRIVACY


KICK ASS

PRIVACY

PRIVACY

WHY?? Not only motivation....

Is privacy engineering a craft?

Engineering Privacy by Design 1.0

Two case studies:

- anonymous e-petitions: no identity attached to petitions
- privacy-preserving road tolling: no fine grained data sent to server

Seda Gurses, Carmela Troncoso, Claudia Diaz. Engineering Privacy by Design.Computers, Privacy & Data Protection. 2011

Engineering Privacy by Design 1.0

Two case studies:

- anonymous e-petitions: no identity attached to petitions
- privacy-preserving road tolling: no fine grained data sent to server

The Key is *"data minimization"*

Seda Gurses, Carmela Troncoso, Claudia Diaz. Engineering Privacy by Design.Computers, Privacy & Data Protection. 2011

Engineering Privacy by Design 1.0

Two case studies:

- anonymous e-petitions: no identity attached to petitions
- privacy-preserving road tolling: no fine grained data sent to server

The Key is *"data minimization"*

Seda Gurses, Carmela Troncoso, Claudia Diaz. Engineering Privacy by Design.Computers, Privacy & Data Protection. 2011

Engineering Privacy by Design 1.0

Two case studies:

- anonymous e-petitions: no identity attached to petitions
- privacy-preserving road tolling: no fine grained data sent to server

The Key is *"data minimization"*

but, it's not "data" that is minimized (in the system as a *whole*)

- kept in user devices
- sent encrypted to a server (only client has the key)
- distributed over multiple servers: only the user, or colluding servers, can recover the data

Seda Gurses, Carmela Troncoso, Claudia Diaz. Engineering Privacy by Design.Computers, Privacy & Data Protection. 2011

Engineering Privacy by Design 1.0

Two case studies:

- anonymous e-petitions: no identity attached to petitions
- privacy-preserving road tolling: no fine grained data sent to server

The Key is *"data minimization"*

but, it's not "data" that is minimized (in the system as a *whole*)

- kept in user devices
- sent encrypted to a server (only client has the key)
- distributed over multiple servers: only the user, or colluding servers, can recover the data

"data minimization" is a bad metaphor!!!

Seda Gurses, Carmela Troncoso, Claudia Diaz. Engineering Privacy by Design.Computers, Privacy & Data Protection. 2011

Unpacking "Data Minimization": Privacy by Design Strategies

Overarching goal

Minimizing privacy **risks** and
trust assumptions placed on other entities

Seda Gurses, Carmela Troncoso, Claudia Diaz. Engineering Privacy by Design Reloaded. Amsterdam Privacy Conference. 2015

Unpacking "Data Minimization": Privacy by Design Strategies

Overarching goal

Minimizing privacy **risks** and **trust assumptions** placed on other entities

The Adversary

Seda Gurses, Carmela Troncoso, Claudia Diaz. Engineering Privacy by Design Reloaded. Amsterdam Privacy Conference. 2015
Seda Gurses and Claudia Diaz. "Two tales of privacy in online social networks." IEEE Security & Privacy Magazine. 2013

Unpacking "Data Minimization":

Privacy by Design Strategies

Overarching goal

Minimizing privacy **risks** and **trust assumptions** placed on other entities

Social Privacy

Other users
3rd parties

The Adversary

Seda Gurses, Carmela Troncoso, Claudia Diaz. Engineering Privacy by Design Reloaded. Amsterdam Privacy Conference. 2015
Seda Gurses and Claudia Diaz. "Two tales of privacy in online social networks." IEEE Security & Privacy Magazine. 2013

Unpacking "Data Minimization": Privacy by Design Strategies

Overarching goal

Minimizing privacy **risks** and **trust assumptions** placed on other entities

Social Privacy

Institutional Privacy (data protection)

Other users
3rd parties

+

semi-trusted
service provider

The Adversary

Seda Gurses, Carmela Troncoso, Claudia Diaz. Engineering Privacy by Design Reloaded. Amsterdam Privacy Conference. 2015

Seda Gurses and Claudia Diaz. "Two tales of privacy in online social networks." IEEE Security & Privacy Magazine. 2013

Unpacking "Data Minimization": Privacy by Design Strategies

Overarching goal

Minimizing privacy **risks** and **trust assumptions** placed on other entities

Social Privacy

Institutional Privacy (data protection)

Other users 3rd parties + semi-trusted service provider

The Adversary

Anti-surveillance Privacy (PETS)

EVERYONE

Seda Gurses, Carmela Troncoso, Claudia Diaz. Engineering Privacy by Design Reloaded. Amsterdam Privacy Conference. 2015

Seda Gurses and Claudia Diaz. "Two tales of privacy in online social networks." IEEE Security & Privacy Magazine. 2013

Unpacking "Data Minimization": Privacy by Design Strategies

Overarching goal

Minimizing privacy **risks** and **trust assumptions** placed on other entities

Strategies

Minimize Collection	Minimize Disclosure	Minimize Linkability
Minimize Centralization	Minimize Replication	Minimize Retention

Seda Gurses, Carmela Troncoso, Claudia Diaz. Engineering Privacy by Design Reloaded. Amsterdam Privacy Conference. 2015

Unpacking "Data Minimization":
Privacy by Design Strategies

Overarching goal

Minimizing privacy **risks** and **trust assumptions** placed on other entities

Strategies

Minimize Collection	Minimize Disclosure	Minimize Linkability
Minimize Centralization	Minimize Replication	Minimize Retention

Great! but... how do we use these strategies?

We make explicit the activities and reasoning in **privacy engineering design** process

Seda Gurses, Carmela Troncoso, Claudia Diaz. Engineering Privacy by Design Reloaded. Amsterdam Privacy Conference. 2015

Case study: Electronic Toll Pricing

Starting assumptions

1) Well defined functionality
   - Charge depending on driving

2) Security, privacy & service integrity requirements
   - User location should be private
   - No cheating clients

3) Initial reference system

Case study: Electronic Toll Pricing

Starting assumptions

1) Well defined functionality
   - Charge depending on driving

2) Security, privacy & service integrity requirements
   - User location should be private
   - No cheating clients

3) Initial reference system

Case study: Electronic Toll Pricing

Activity 1: Classify Entities in domains

User domain: components under the control of the user, eg, user devices

Service domain: components outside the control of the user, eg, backend system at provider

Case study: Electronic Toll Pricing

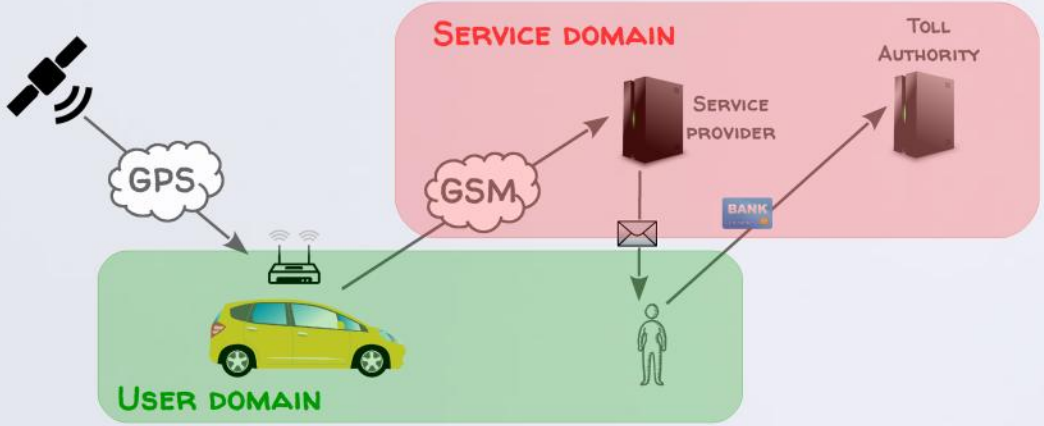

Activity 1: Classify Entities in domains

- **User domain**: components under the control of the user, eg, user devices
- **Service domain**: components outside the control of the user, eg, backend system at provider

Case study: Electronic Toll Pricing

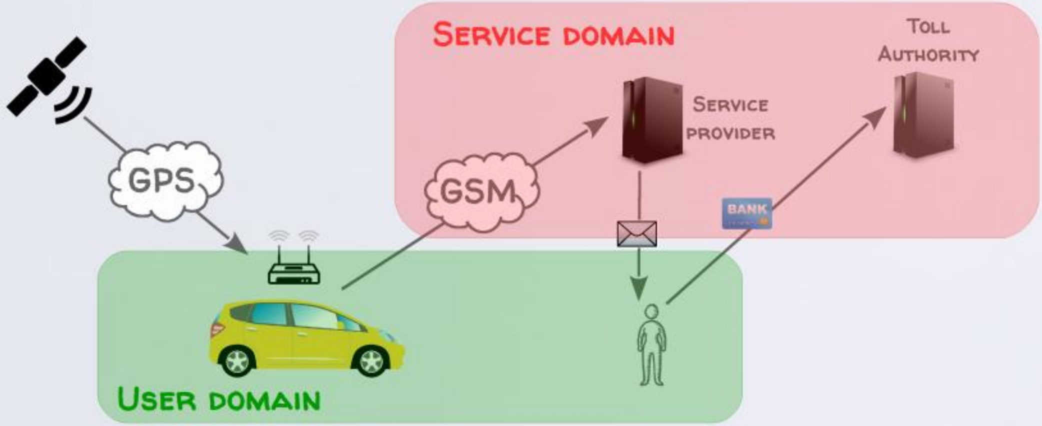

Activity 1: Classify Entities in domains

- **User domain**: components under the control of the user, eg, user devices
- **Service domain**: components outside the control of the user, eg, backend system at provider

Activity 2: Identify necessary data for providing the service

- Location data – compute bill
- Billing data – charge user
- Personal data – send bill
- Payment data – perform payment

Case study: Electronic Toll Pricing

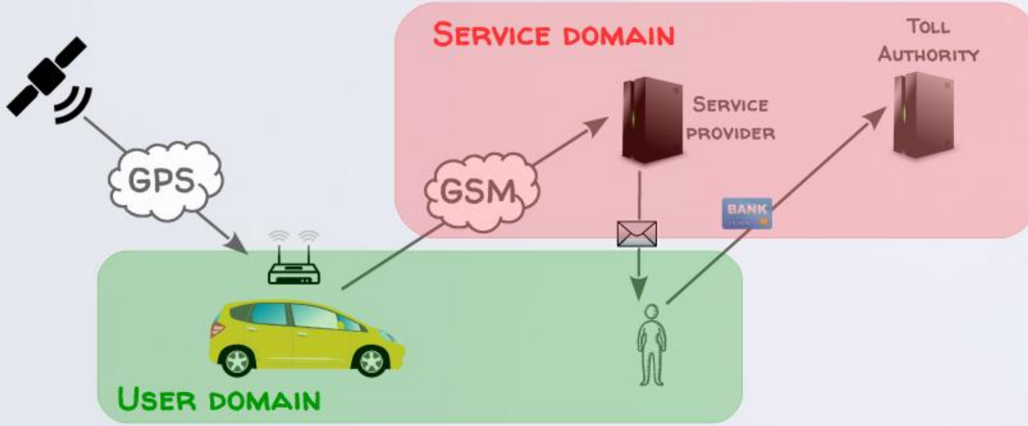

Activity 1: Classify Entities in domains

User domain: components under the control of the user, eg, user devices

Service domain: components outside the control of the user, eg, backend system at provider

Activity 2: Identify necessary data for providing the service

Location data – compute bill

Billing data – charge user

Personal data – send bill

Payment data – perform payment

Activity 3: Distribute data in architecture

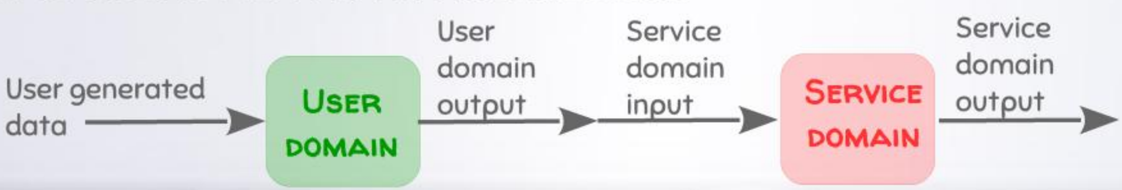

Case study: Electronic Toll Pricing

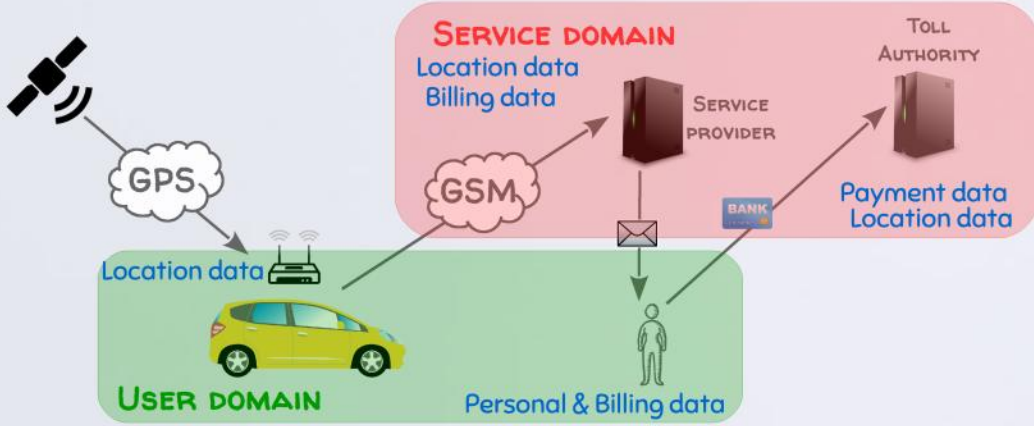

Activity 1: Classify Entities in domains

User domain: components under the control of the user, eg, user devices

Service domain: components outside the control of the user, eg, backend system at provider

Activity 2: Identify necessary data for providing the service

Location data – compute bill

Billing data – charge user

Personal data – send bill

Payment data – perform payment

Activity 3: Distribute data in architecture

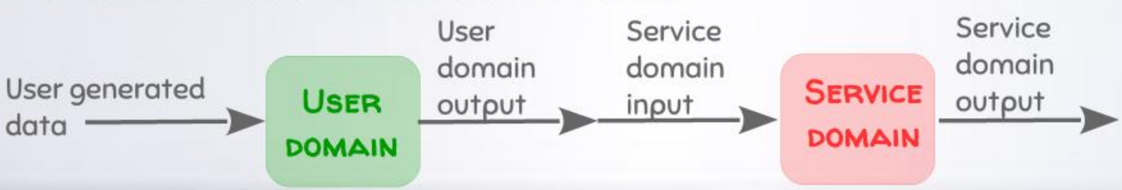

Case study: Electronic Toll Pricing

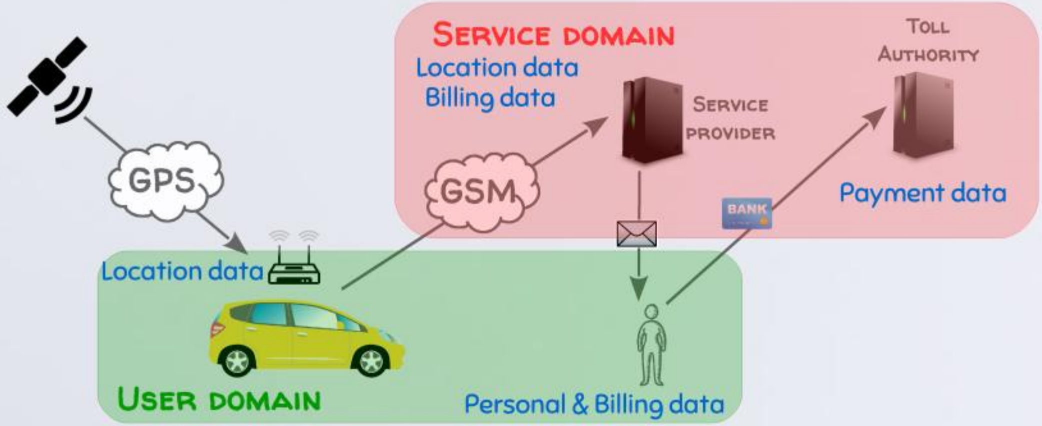

Activity 4: Select technological solutions following →

not sending the data (local computations)
encrypting the data
advanced privacy-preserving protocols
obfuscate the data
anonymize the data

Minimizing privacy **risks** and **trust assumptions** placed on other entities

Minimize Collection
Minimize Disclosure
Minimize Linkability
Minimize Centralization
Minimize Replication
Minimize Retention

Case study: Electronic Toll Pricing

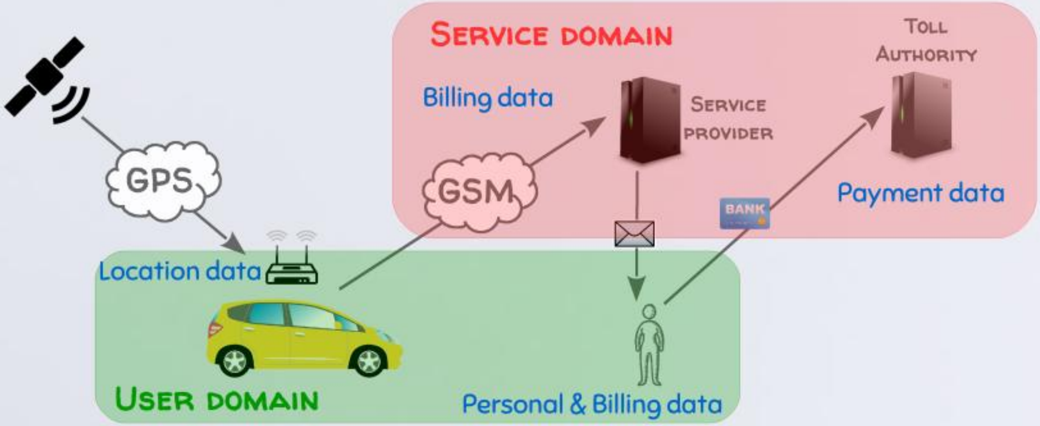

Location is not needed,
only the amount to bill!

Activity 4: Select technological solutions following →

- not sending the data (local computations)
- encrypting the data
- advanced privacy-preserving protocols
- obfuscate the data
- anonymize the data

Minimizing privacy **risks** and **trust assumptions** placed on other entities

Minimize Collection	Minimize Disclosure	Minimize Linkability
Minimize Centralization	Minimize Replication	Minimize Retention

J. Balasch, A. Rial, C. Troncoso, B. Preneel, I. Verbauwhede, C. Geuens. PrETP "Privacy-Preserving Electronic Toll Pricing" USENIX Security Symposium 2010
C. Troncoso, G. Danezis, E. Kosta, J. Balasch, B. Preneel. "PriPAYD. Privacy-Friendly Pay-As-You-Drive Insurance" IEEE TDSC 2011

Case study: Electronic Toll Pricing

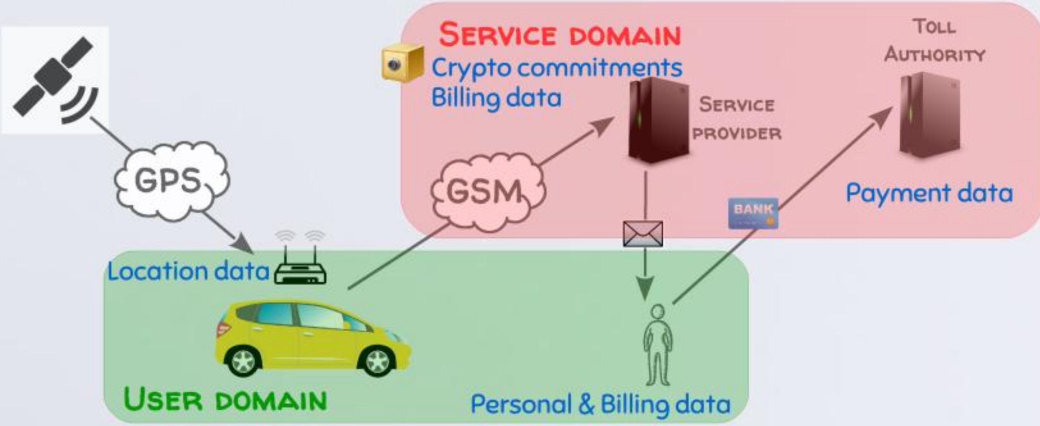

Location is not needed, only the amount to bill!

Service integrity?

Activity 4: Select technological solutions following →

- not sending the data (local computations)
- encrypting the data
- advanced privacy-preserving protocols
- obfuscate the data
- anonymize the data

Minimizing privacy **risks** and **trust assumptions** placed on other entities

Minimize Collection	Minimize Disclosure	Minimize Linkability
Minimize Centralization	Minimize Replication	Minimize Retention

J. Balasch, A. Rial, C. Troncoso, B. Preneel, I. Verbauwhede, C. Geuens. PrETP "Privacy-Preserving Electronic Toll Pricing" USENIX Security Symposium 2010
C. Troncoso, G. Danezis, E. Kosta, J. Balasch, B. Preneel. "PriPAYD. Privacy-Friendly Pay-As-You-Drive Insurance" IEEE TDSC 2011

Case study: Electronic Toll Pricing

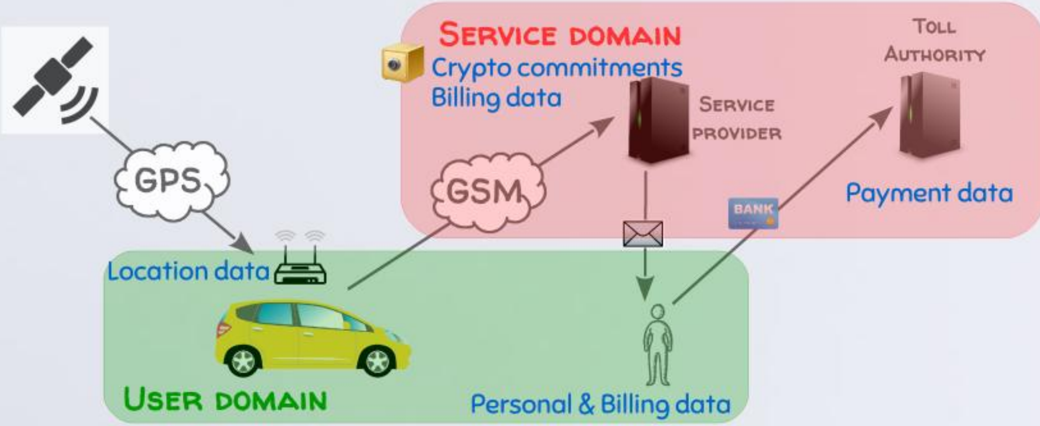

Location is not needed, only the amount to bill!

Service integrity?

Activity 4: Select technological solutions following →

- not sending the data (local computations)
- encrypting the data
- advanced privacy-preserving protocols
- obfuscate the data
- anonymize the data

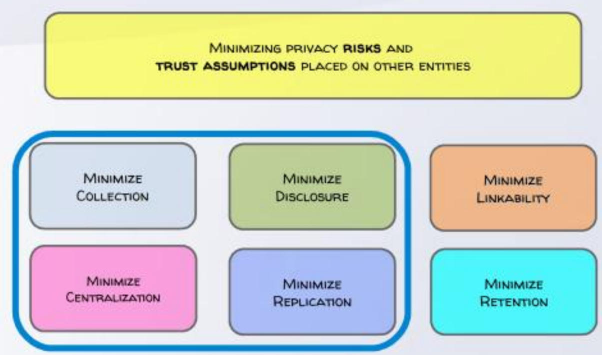

J. Balasch, A. Rial, C. Troncoso, B. Preneel, I. Verbauwhede, C. Geuens. PrETP "Privacy-Preserving Electronic Toll Pricing" USENIX Security Symposium 2010

C. Troncoso, G. Danezis, E. Kosta, J. Balasch, B. Preneel. "PriPAYD. Privacy-Friendly Pay-As-You-Drive Insurance" IEEE TDSC 2011

Case study: Electronic Toll Pricing

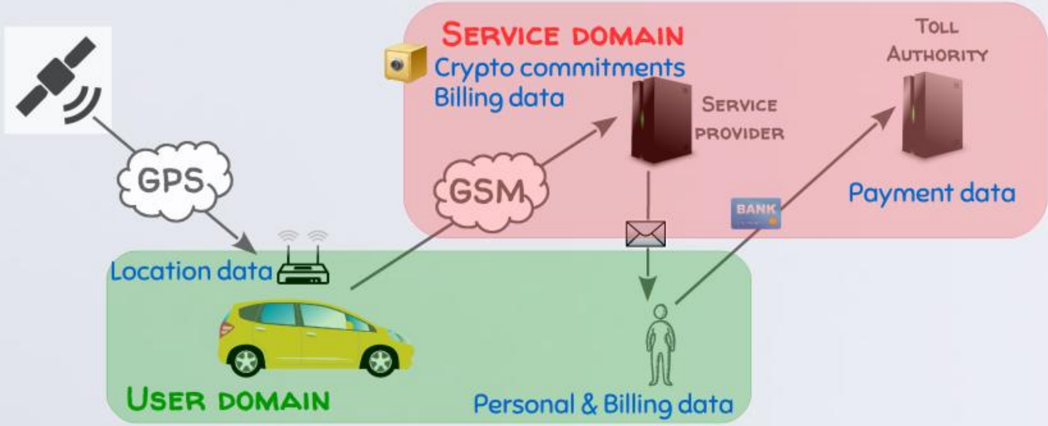

Location is not needed, only the amount to bill!

Service integrity?

Activity 4: Select technological solutions following →

not sending the data (local computations)
encrypting the data
advanced privacy-preserving protocols
obfuscate the data
anonymize the data

J. Balasch, A. Rial, C. Troncoso, B. Preneel, I. Verbauwhede, C. Geuens. PrETP "Privacy-Preserving Electronic Toll Pricing" USENIX Security Symposium 2010
C. Troncoso, G. Danezis, E. Kosta, J. Balasch, B. Preneel. "PriPAYD. Privacy-Friendly Pay-As-You-Drive Insurance" IEEE TDSC 2011

Case study: Electronic Toll Pricing

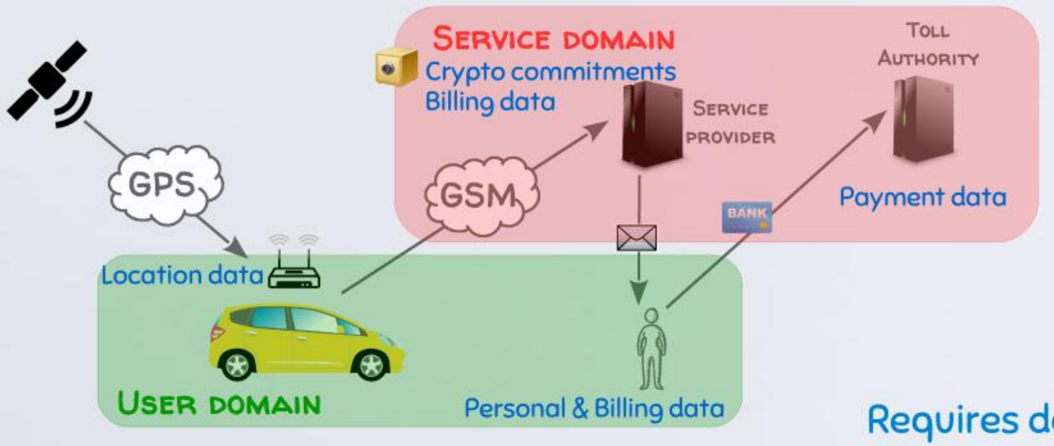

Location is not needed, only the amount to bill!

Requires deep knowledge of PETs

Service integrity?

Activity 4: Select technological solutions following →

- not sending the data (local computations)
- encrypting the data
- advanced privacy-preserving protocols
- obfuscate the data
- anonymize the data

J. Balasch, A. Rial, C. Troncoso, B. Preneel, I. Verbauwhede, C. Geuens. PrETP "Privacy-Preserving Electronic Toll Pricing" USENIX Security Symposium 2010

C. Troncoso, G. Danezis, E. Kosta, J. Balasch, B. Preneel. "PriPAYD. Privacy-Friendly Pay-As-You-Drive Insurance" IEEE TDSC 2011

Case study: Electronic Toll Pricing

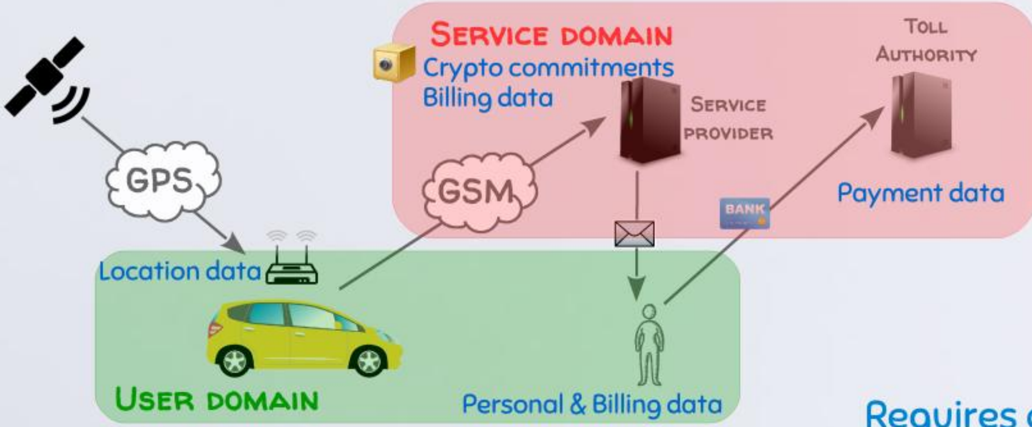

Location is not needed, only the amount to bill!

Service integrity?

Requires deep knowledge of PETs

Privacy ENABLING Technologies

Activity 4: Select technological solutions following →

- not sending the data (local computations)
- encrypting the data
- advanced privacy-preserving protocols
- obfuscate the data
- anonymize the data

Minimizing privacy **risks** and **trust assumptions** placed on other entities

Minimize Collection

Minimize Disclosure

Minimize Linkability

Minimize Centralization

Minimize Replication

Minimize Retention

J. Balasch, A. Rial, C. Troncoso, B. Preneel, I. Verbauwhede, C. Geuens. PrETP "Privacy-Preserving Electronic Toll Pricing" USENIX Security Symposium 2010
C. Troncoso, G. Danezis, E. Kosta, J. Balasch, B. Preneel. "PriPAYD. Privacy-Friendly Pay-As-You-Drive Insurance" IEEE TDSC 2011

PART I:
Reasoning about Privacy when designing systems

PART II:
Designing Technologies to support Privacy-aware designs

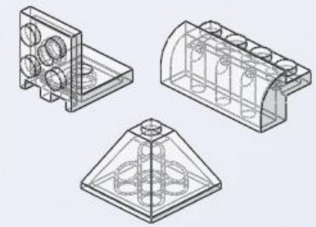

PART I:
Reasoning about Privacy when designing systems

PART II:
Designing Technologies to support Privacy-aware designs

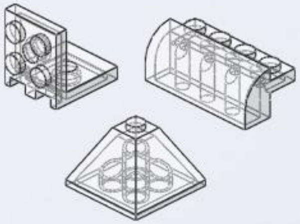

Activity 4: Select technological solutions to follow

- not sending the data (local computations)
- encrypting the data
- advanced privacy-preserving protocols
- obfuscate the data
- anonymize the data

PART I:
Reasoning about Privacy when designing systems

PART II:
Designing Technologies to support Privacy-aware designs

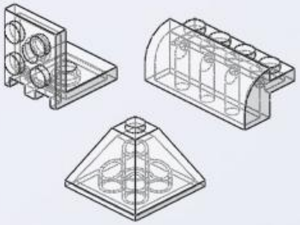

Activity 4: Select technological solutions to follow

- not sending the data (local computations)
- encrypting the data
- advanced privacy-preserving protocols
- obfuscate the data
- anonymize the data

"EASY" DESIGN but expensive

PART I:
Reasoning about
Privacy when designing
systems

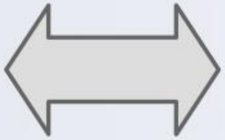

PART II:
Designing Technologies
to support Privacy-
aware designs

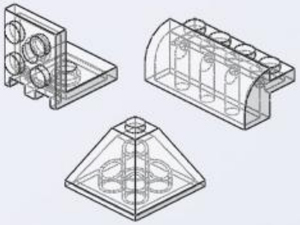

Activity 4: Select technological solutions to follow

- not sending the data (local computations)
- encrypting the data
- advanced privacy-preserving protocols
- obfuscate the data
- anonymize the data

cheap but... DIFFICULT TO DESIGN

PART II:
Designing Technologies to support Privacy-aware designs

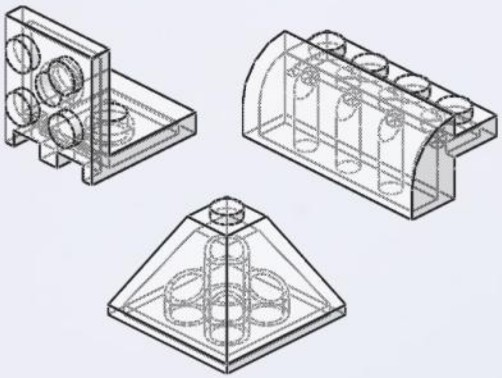

Activity 4: Select technological solutions to follow

- not sending the data (local computations)
- encrypting the data
- advanced privacy-preserving protocols
- obfuscate the data
- anonymize the data

cheap but... DIFFICULT TO DESIGN

PART II:
Designing Technologies to support Privacy-aware designs

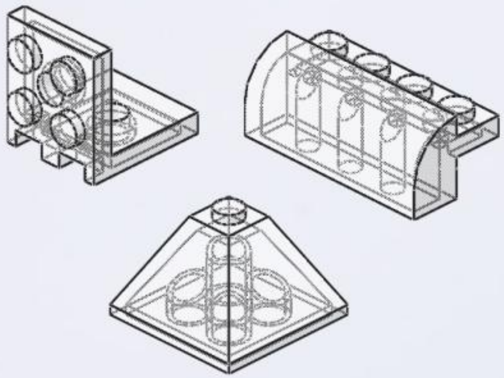

Activity 4: Select technological solutions to follow

- not sending the data (local computations)
- encrypting the data
- advanced privacy-preserving protocols
- obfuscate the data
- anonymize the data

cheap but... DIFFICULT TO DESIGN

PART II:
Designing Technologies to support Privacy-aware designs

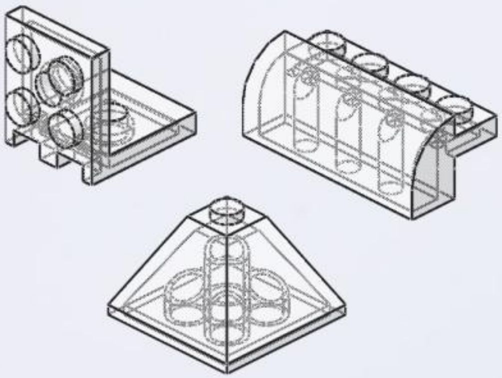

Activity 4: Select technological solutions to follow

not sending the data (local computations)
encrypting the data
advanced privacy-preserving protocols
obfuscate the data
anonymize the data

cheap but... DIFFICULT TO DESIGN

The adversary knows!

PART II:
Designing Technologies to support Privacy-aware designs

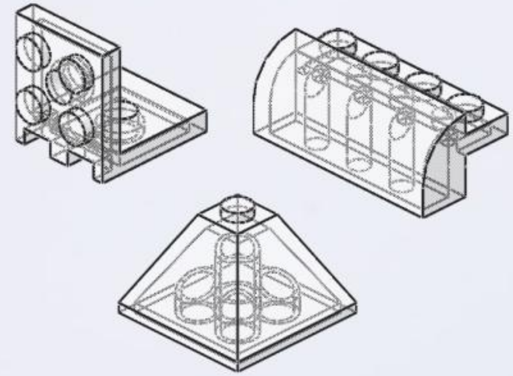

Activity 4: Select technological solutions to follow

- not sending the data (local computations)
- encrypting the data
- advanced privacy-preserving protocols
- obfuscate the data
- anonymize the data

cheap but... DIFFICULT TO DESIGN

The adversary knows!

PART II:
Designing Technologies to support Privacy-aware designs

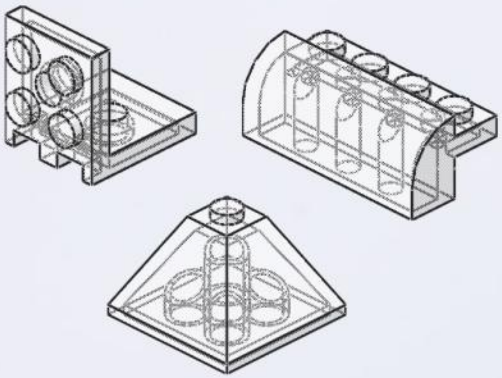

Activity 4: Select technological solutions to follow

not sending the data (local computations)
encrypting the data
advanced privacy-preserving protocols
obfuscate the data
anonymize the data

cheap but... DIFFICULT TO DESIGN

The adversary knows!

Can we design good  systematically?

Designing anonymous communications Systems

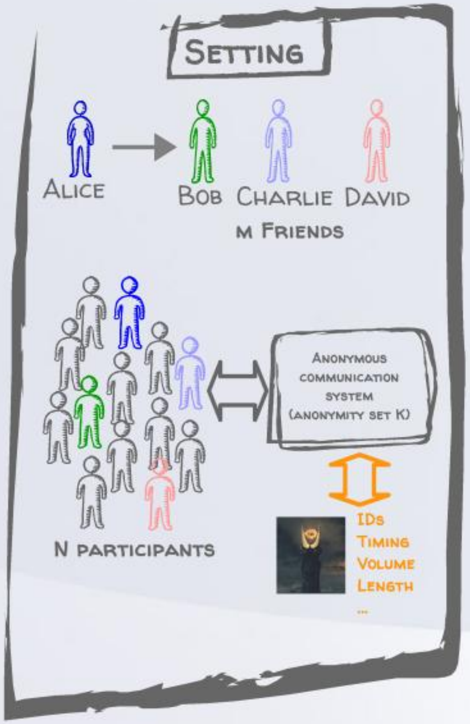

Designing anonymous communications systems

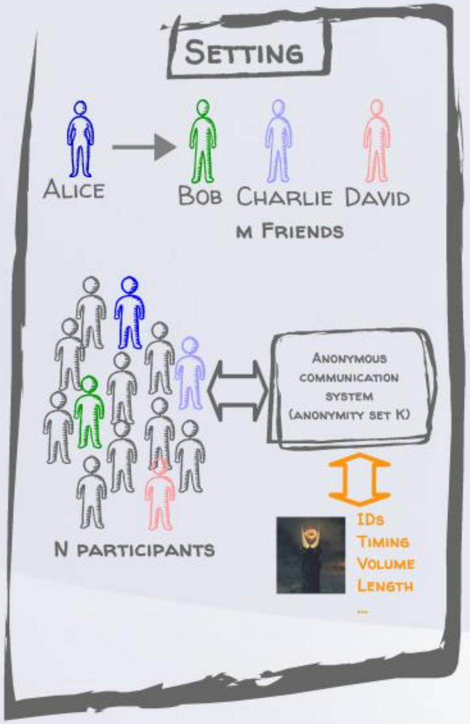

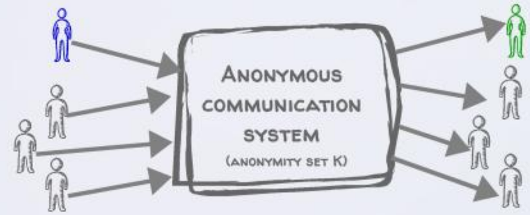

1-  sees Alice sending a single message to the system

2- Anonymity set size = K

3- Perfect! 

As time goes by and Alice sends more messages...

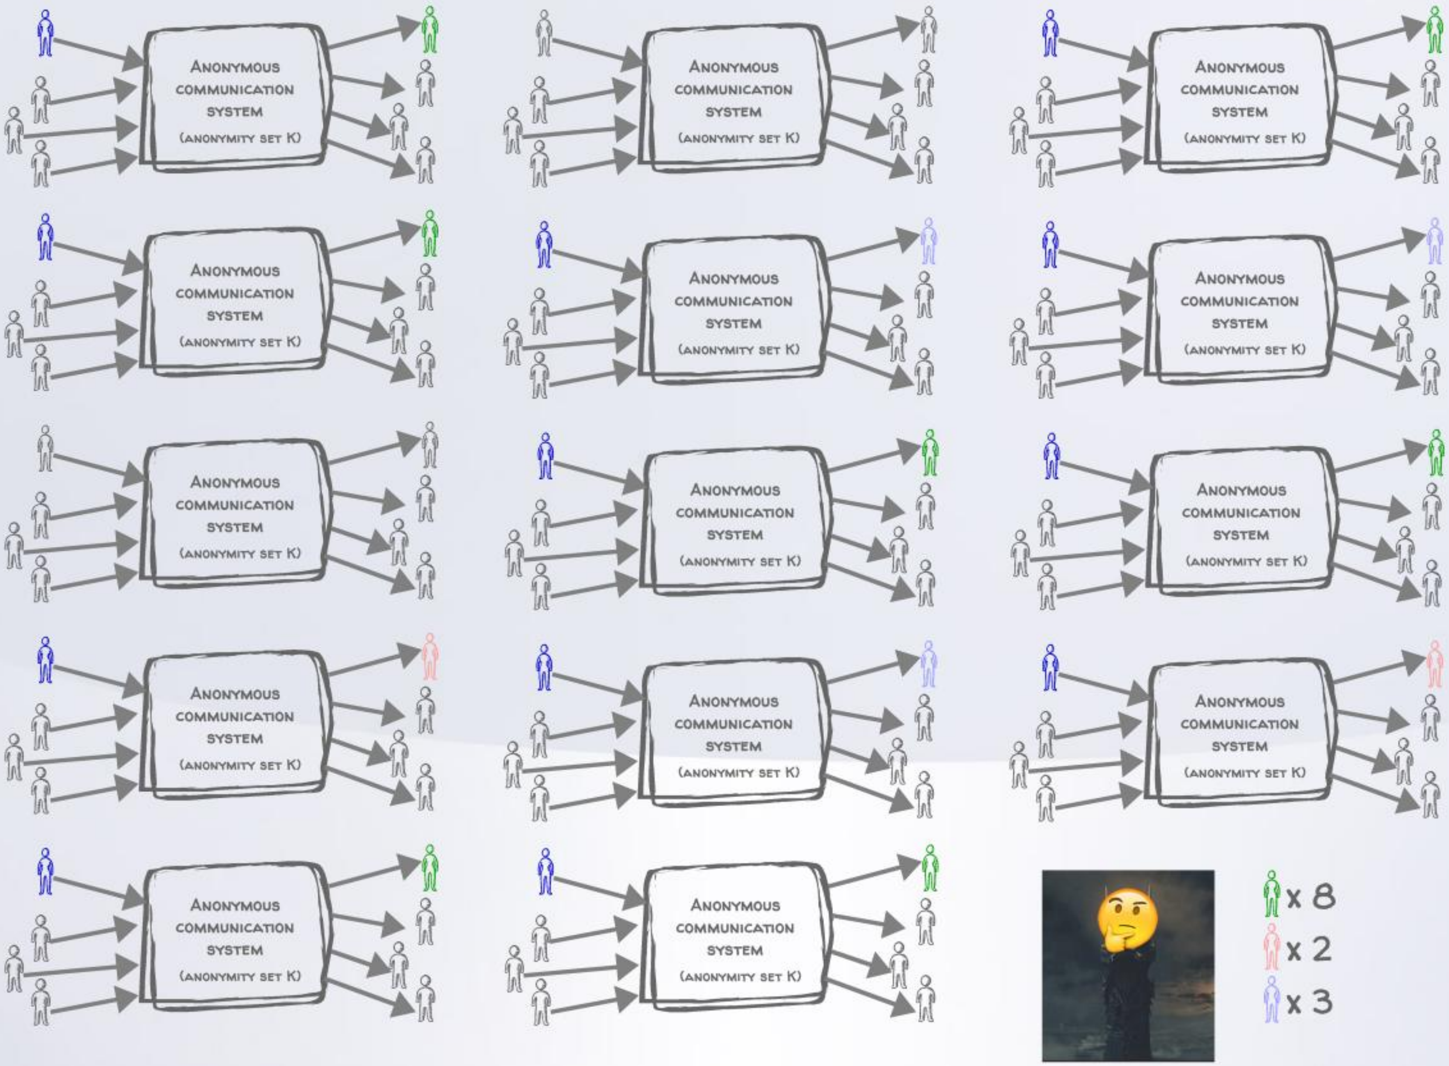

Pool-Based Anonymous Communications

λ = rate of messages

P = probability that sends a message to

α = probability of message leaving the pool

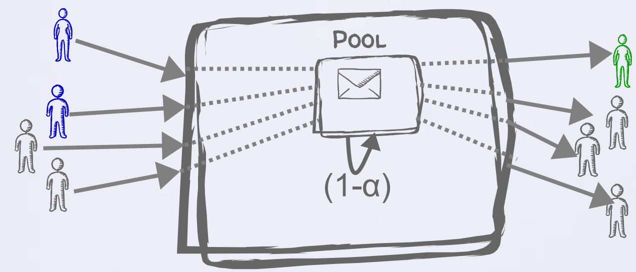

Fernando Pérez-González, , Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Simon Oya, Carmela Troncoso, Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014

Pool-Based Anonymous Communications

λ = rate of messages

P = probability that sends a message to

α = probability of message leaving the pool

δ = rate of dummy messages

δ = rate of dummy messages

P = probability that sends a dummy message to

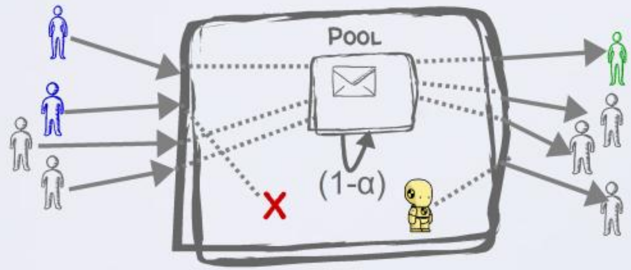

Fernando Pérez-González, , Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Simon Oya, Carmela Troncoso, Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014

Pool-Based Anonymous Communications

λ = rate of messages
P = probability that sends a message to
α = probability of message leaving the pool
δ = rate of dummy messages
δ = rate of dummy messages
P = probability that sends a dummy message to

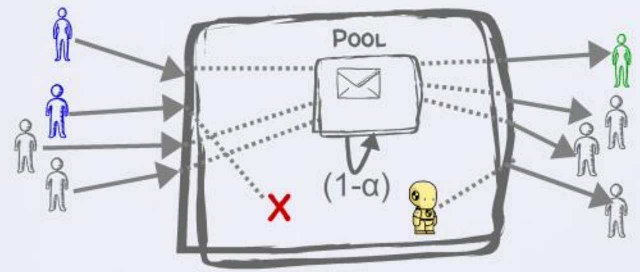

Least Squares Disclosure Attack (optimal for Mean Square Error)

x^r = vector of n# of messages sent round r (x^r =2)
y^r = vector of n# of messages received round r (y^r = 1)
$H = [x^1, x^2, x^3, \dots,]$
$Y = [y^1, y^2, y^3, \dots,]^T$

$$\hat{p} = \arg\min_p \| y - Hp \|$$
$$p_{i,j} \leqslant 1$$
$$\sum_i p_{i,j} = 1$$

Fernando Pérez-González, , Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Simon Oya, Carmela Troncoso, Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014

Pool-Based Anonymous Communications

λ = rate of messages
P = probability that sends a message to
α = probability of message leaving the pool
δ = rate of dummy messages
δ = rate of dummy messages
P = probability that sends a dummy message to

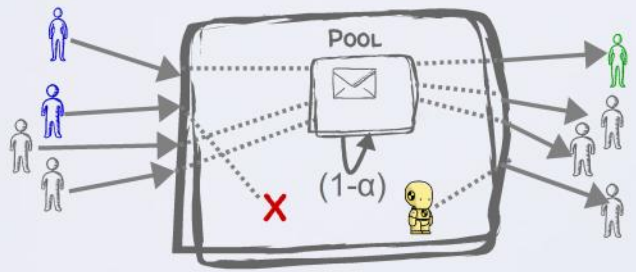

x^r = vector of n# of messages sent round r (x^r =2)
y^r = vector of n# of messages received round r (y^r = 1)
H = [x^1,x^2,x^3, ... ,]
Y = [y^1,y^2,y^3, ... ,]^T

Least Squares Disclosure Attack (optimal for Mean Square Error)

$$\hat{p} = \arg\min_{p} \| y - Hp \| \quad p_{i,j} \leq 1 \quad \sum_i p_{i,j} = 1$$

$$\hat{p} = (H^T H)^{-1} H^T y$$

Fernando Pérez-González, , Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Simon Oya, Carmela Troncoso, Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014

Pool-Based Anonymous Communications

λ = rate of messages
P = probability that sends a message to
α = probability of message leaving the pool
δ = rate of dummy messages
δ = rate of dummy messages
P = probability that sends a dummy message to

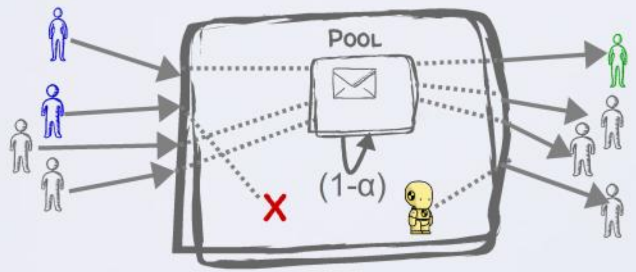

Least Squares Disclolsure Attack (optimal for Mean Square Error)

x^r = vector of n# of messages sent round r (x^r =2)
y^r = vector of n# of messages received round r (y^r = 1)
$H = [x^1, x^2, x^3, \dots,]$
$Y = [y^1, y^2, y^3, \dots,]^T$

$$\hat{p} = \arg\min_p \| y - Hp \| \quad p_{i,j} \leq 1 \quad \sum_i p_{i,j} = 1$$

$$\Rightarrow \hat{p} = (H^T H)^{-1} H^T y$$

$$MSE = \| p - \hat{p} \| = \frac{1}{t} \cdot \frac{2-\alpha}{\alpha} \cdot \frac{1}{\lambda} \cdot \left(1 + \frac{\delta}{\lambda}\right) \cdot \left(\lambda' + \delta\, p\right)$$

Fernando Pérez-González, , Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Simon Oya, Carmela Troncoso, Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014

Pool-Based Anonymous Communications

λ = rate of messages
P = probability that sends a message to
α = probability of message leaving the pool
δ = rate of dummy messages
δ = rate of dummy messages
P = probability that sends a dummy message to

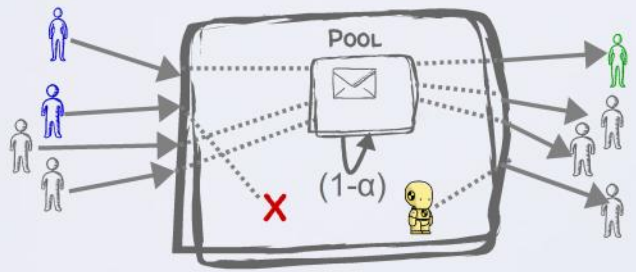

Least Squares Disclosure Attack (optimal for Mean Square Error)

x^r = vector of n# of messages sent round r (x^r =2)
y^r = vector of n# of messages received round r (y^r = 1)
H = [x^1,x^2,x^3, ... ,]
Y = [y^1,y^2,y^3, ... ,]^T

$$\hat{p} = \arg\min_{p} \| y - Hp \| \quad p_{i,j} \leq 1, \quad \sum_i p_{i,j} = 1$$

$$\Rightarrow \hat{p} = (H^T H)^{-1} H^T y$$

$$MSE = \| p - \hat{p} \| = \frac{1}{t} \cdot \frac{2-\alpha}{\alpha} \cdot \frac{1}{\lambda} \cdot \left(1 + \frac{\delta}{\lambda}\right) \cdot \left(\lambda' + \delta p\right)$$

Observations — Pool — Sender — Receiver

Fernando Pérez-González, , Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Simon Oya, Carmela Troncoso, Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014

Systematic dummy strategy design

$$MSE = \frac{1}{t} \cdot \frac{2-\alpha}{\alpha} \cdot \frac{1}{\lambda} \cdot \left(1+\frac{\delta}{\lambda}\right) \cdot \left(\lambda' + \delta\, p\right)$$

Observations   Sender   Receiver

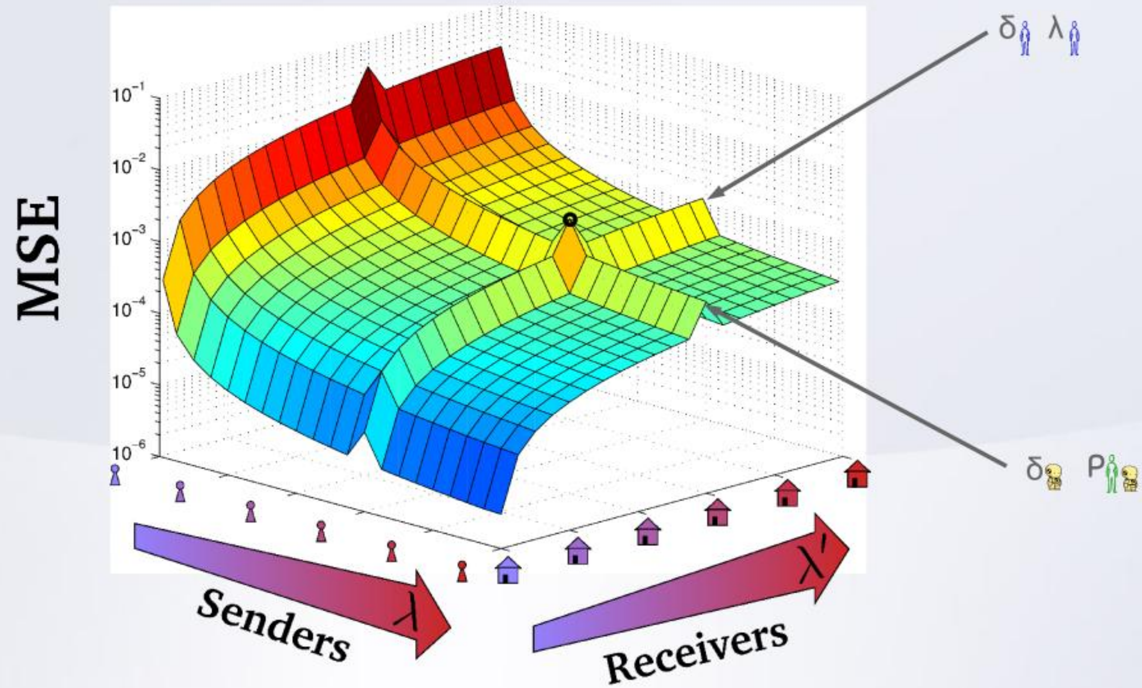

Fernando Pérez-González, , Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Simon Oya, Carmela Troncoso, Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014

Systematic dummy strategy design

$$MSE = \underbrace{\frac{1}{t}}_{\text{Observations}} \cdot \frac{2-\alpha}{\alpha} \cdot \underbrace{\frac{1}{\lambda} \cdot \left(1+\frac{\delta}{\lambda}\right)}_{\text{Sender}} \cdot \underbrace{\left(\lambda' + \delta\, p\right)}_{\text{Receiver}}$$

Given a dummy budget

Pick your favourite privacy objective and design dummy strategy!

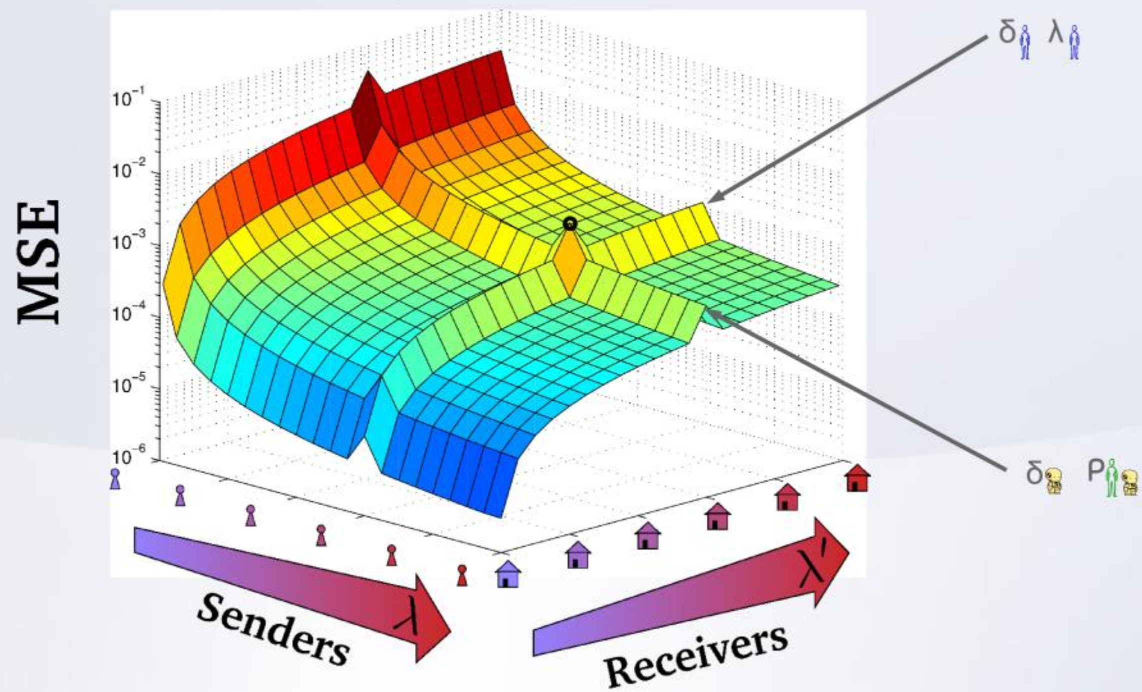

Fernando Pérez-González, , Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Simon Oya, Carmela Troncoso, Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014

Systematic dummy strategy design

$$MSE = \frac{1}{t}\cdot\frac{2-\alpha}{\alpha}\cdot\frac{1}{\lambda}\cdot\left(1+\frac{\delta}{\lambda}\right)\cdot\left(\lambda' + \delta\, p\right)$$

Observations — Sender — Receiver

Given a dummy budget

Pick your favourite privacy objective and design dummy strategy!

$$\underset{\delta,\delta,p}{\text{maximize}} \quad MSE \quad \forall$$
$$s.t. \quad MSE = \beta \cdot MSE \quad \forall$$

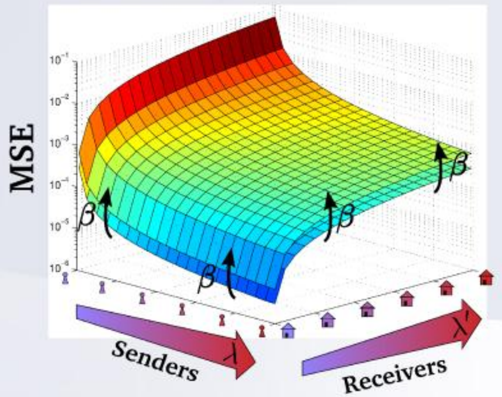

Fernando Pérez-González, , Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Simon Oya, Carmela Troncoso, Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014

Systematic dummy strategy design

$$MSE = \frac{1}{t} \cdot \frac{2-\alpha}{\alpha} \cdot \frac{1}{\lambda} \cdot \left(1+\frac{\delta}{\lambda}\right) \cdot \left(\lambda' + \delta\, p\right)$$

Observations — Sender — Receiver

Given a dummy budget

Pick your favourite privacy objective and design dummy strategy!

$$\underset{\delta, \delta, p}{maximize} \quad MSE \quad \forall$$
$$s.t. \quad MSE = \beta \cdot MSE \quad \forall$$

$$\underset{\delta, \delta, p}{maximize} \quad \min MSE \quad \forall$$

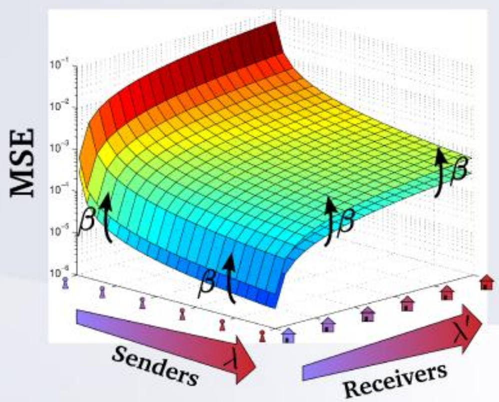

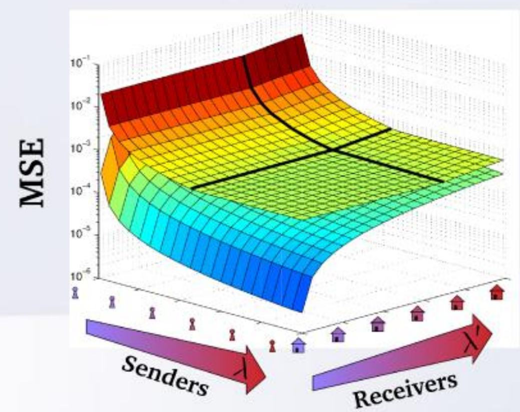

Fernando Pérez-González, , Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Simon Oya, Carmela Troncoso, Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014
Simon Oya, Fernando Pérez-González, Carmela Troncoso. Design of Pool Mixes Against Profiling Attacks in Real Conditions. IEEE/ACM Transactions on Networking, 2016

Systematic dummy strategy design

$$MSE = \frac{1}{t} \cdot \frac{2-\alpha}{\alpha} \cdot \frac{1}{\lambda} \cdot \left(1 + \frac{\delta}{\lambda}\right) \cdot \left(\lambda' + \delta\, p\right)$$

Observations — Sender — Receiver

Given a dummy budget

Pick your favourite privacy objective and design dummy strategy!

$$\underset{\delta, \delta, p}{maximize} \quad MSE \quad \forall$$
$$s.t. \quad MSE = \beta \cdot MSE \quad \forall$$

$$\underset{\delta, \delta, p}{maximize} \quad \min MSE \quad \forall$$

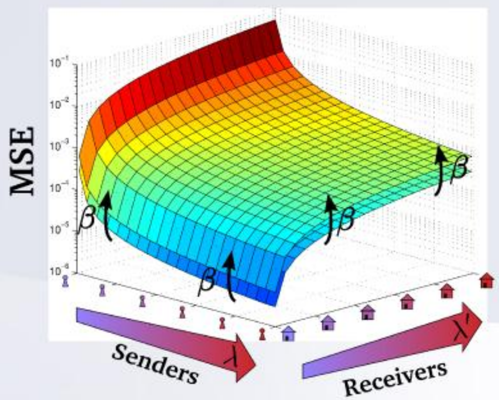

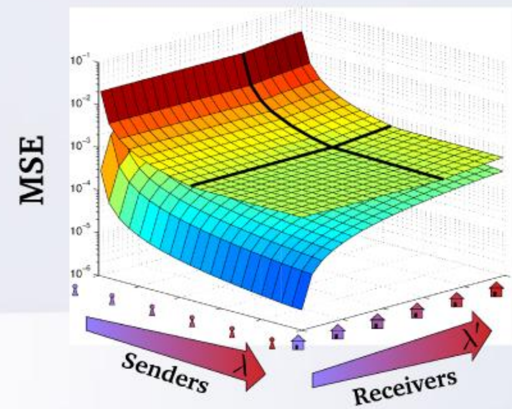

We can also use the MSE to design optimal pools

Fernando Pérez-González, , Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Simon Oya, Carmela Troncoso, Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014
Simon Oya, Fernando Pérez-González, Carmela Troncoso. Design of Pool Mixes Against Profiling Attacks in Real Conditions. IEEE/ACM Transactions on Networking, 2016

Designing Location Privacy-Preserving Mechanisms

What if the optimal attack is not known? 🤔

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Designing Location Privacy-Preserving Mechanisms

What if the optimal attack is not known? 🤔

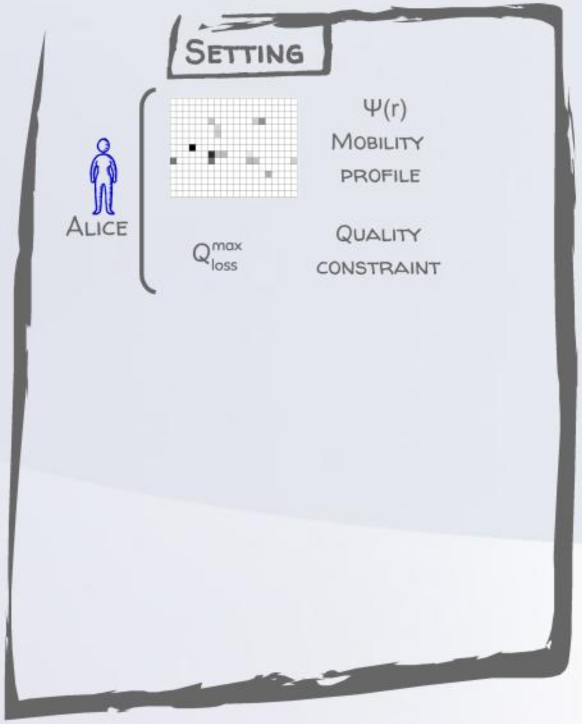

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Designing Location Privacy-Preserving Mechanisms

What if the optimal attack is not known? 🤔

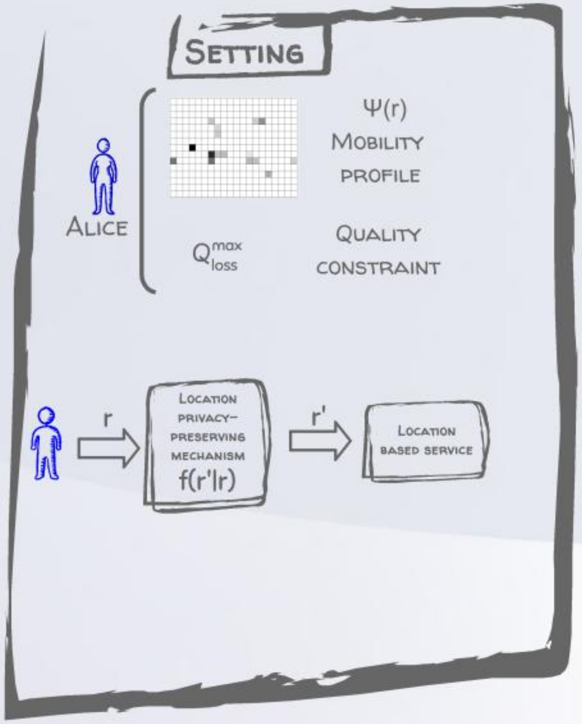

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Designing Location Privacy-Preserving Mechanisms

What if the optimal attack is not known? 🤔

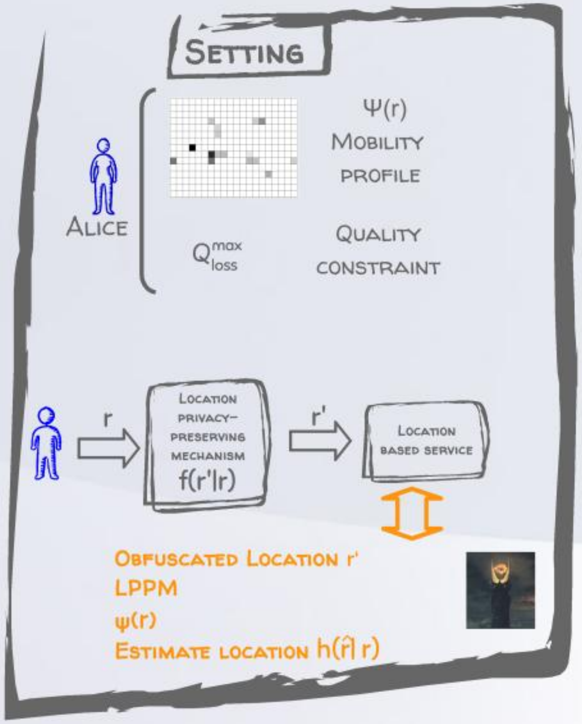

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Designing Location Privacy-Preserving Mechanisms

What if the optimal attack is not known?

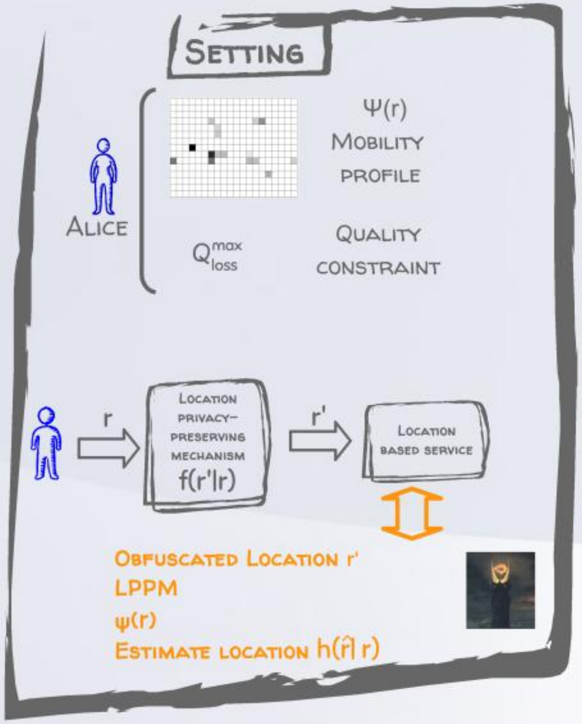

Design f(r'|r) to maximize privacy

$$Privacy(\psi, f, h, d_p) = \sum_{\hat{r}, r', r} \psi(r) f(r'|r) h(\hat{r}|r') d_p(r', \hat{r})$$

Respecting Q_{loss}^{max}

$$Q_{loss}(\psi, f, d_q) = \sum_{r', r} \psi(r) f(r'|r) d_q(r, r')$$

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Designing Location Privacy-Preserving Mechanisms

What if the optimal attack is not known?

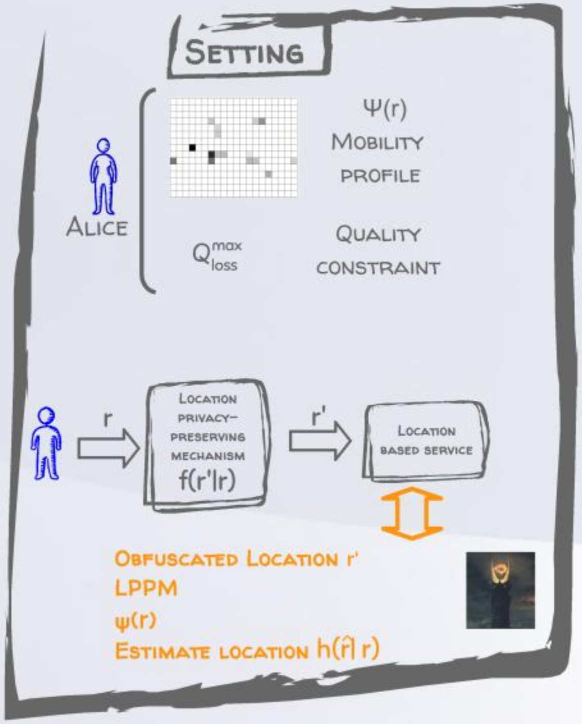

Design f(r'|r) to maximize privacy

$$Privacy(\psi, f, h, d_p) = \sum_{\hat{r}, r', r} \psi(r) f(r'|r) h(\hat{r}|r') d_p(r', \hat{r})$$

Respecting Q_{loss}^{max}

$$Q_{loss}(\psi, f, d_q) = \sum_{r', r} \psi(r) f(r'|r) d_q(r, r')$$

Traditional: arms race-based design

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Designing Location Privacy-Preserving Mechanisms

What if the optimal attack is not known? 🤔

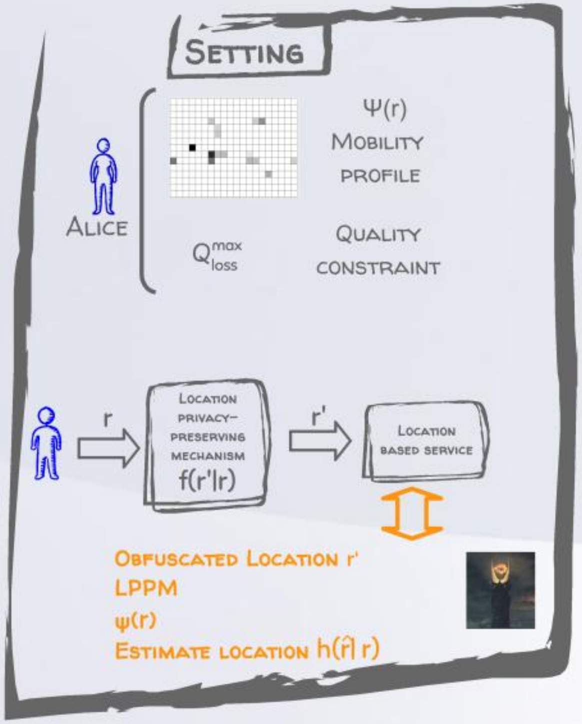

Design f(r'|r) to maximize privacy

$$Privacy(\psi, f, h, d_p) = \sum_{\hat{r}, r', r} \psi(r) f(r'|r) h(\hat{r}|r') d_p(r', \hat{r})$$

Respecting Q_{loss}^{max}

$$Q_{loss}(\psi, f, d_q) = \sum_{r', r} \psi(r) f(r'|r) d_q(r, r')$$

Traditional: arms race-based design

The race may never end...

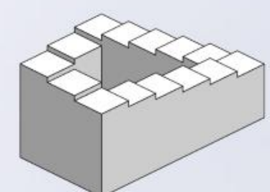

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Cutting the race short: A Game theoretic approach

Zero –sum Bayesian Stackelberg game

 Leader – chooses defense f()

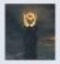 Follower – chooses attack h()

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Cutting the race short: A Game theoretic approach

Zero –sum Bayesian Stackelberg game

-  Leader – chooses defense f()
- 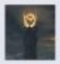 Follower – chooses attack h()

Bayesian:  incomplete information

Zero-sum:  gain is  loss (and vice versa)

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Cutting the race short: A Game theoretic approach

Zero –sum Bayesian Stackelberg game

-  Leader – chooses defense f()
- 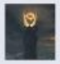 Follower – chooses attack h()

Bayesian:  incomplete information

Zero-sum:  gain is  loss (and vice versa)

Optimal Strategy for the user

Choose $f(r'|r),\ x_{r'} = \min_{\hat{r}} \sum_{r} \psi(r) f(r'|r) d_p(\hat{r}, r)$

maximize $\sum_{r'} x_{r'}$

s.t. $x_{r'} \leq \sum_{r} \psi(r) f(r'|r) d_p(\hat{r}, r), \forall \hat{r}, r'$

$\sum_{r} \psi(r) \sum_{r'} f(r'|r) d_q(r', r), \leq Q_{loss}^{max}$

$\sum_{r'} f(r'|r) = 1,\ f(r'|r) \geq 0, \forall r, r'$

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Cutting the race short: A Game theoretic approach

Zero –sum Bayesian Stackelberg game

- Leader – chooses defense f()
- Follower – chooses attack h()

Bayesian: incomplete information

Zero–sum: gain is loss (and vice versa)

Optimal Strategy for the user

Choose $f(r'|r)$ $\quad x_{r'} = \min_{\hat{r}} \sum_{r} \psi(r) f(r'|r) d_p(\hat{r}, r)$

maximize $\quad \sum_{r'} x_{r'}$

s.t. $\quad x_{r'} \leq \sum_{r} \psi(r) f(r'|r) d_p(\hat{r}, r), \forall \hat{r}, r'$

$$\sum_{r} \psi(r) \sum_{r'} f(r'|r) d_q(r', r) \leq Q_{loss}^{max}$$

$$\sum_{r'} f(r'|r) = 1, \quad f(r'|r) \geq 0, \forall r, r'$$

maximize privacy

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Cutting the race short: A Game theoretic approach

Zero –sum Bayesian Stackelberg game

-  Leader – chooses defense f()
- 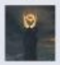 Follower – chooses attack h()

Bayesian:  incomplete information

Zero-sum:  gain is  loss (and vice versa)

Optimal Strategy for the user

Choose $f(r'|r)$, $x_{r'} = \min_{\hat{r}} \sum_{r} \psi(r) f(r'|r) d_p(\hat{r}, r)$ — maximize privacy

maximize $\sum_{r'} x_{r'}$

s.t. $x_{r'} \leq \sum_{r} \psi(r) f(r'|r) d_p(\hat{r}, r), \forall \hat{r}, r'$ — maximize privacy

$\sum_{r} \psi(r) \sum_{r'} f(r'|r) d_q(r', r) \leq Q_{loss}^{max}$ — quality constraint

$\sum_{r'} f(r'|r) = 1, \; f(r'|r) \geq 0, \forall r, r'$

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Cutting the race short: A Game theoretic approach

Zero -sum Bayesian Stackelberg game

- Leader – chooses defense f()
- Follower – chooses attack h()

Bayesian: incomplete information

Zero-sum: gain is loss (and vice versa)

Optimal Strategy for the user

Choose $f(r'|r)$, $x_{r'} = \min_{\hat{r}} \sum_{r} \psi(r) f(r'|r) d_p(\hat{r}, r)$ — maximize privacy

maximize $\sum_{r'} x_{r'}$

s.t.

$x_{r'} \leq \sum_{r} \psi(r) f(r'|r) d_p(\hat{r}, r), \forall \hat{r}, r'$ — maximize privacy

$\sum_{r} \psi(r) \sum_{r'} f(r'|r) d_q(r', r) \leq Q_{loss}^{max}$ — quality constraint

$\sum_{r'} f(r'|r) = 1, \; f(r'|r) \geq 0, \forall r, r'$ — f() is a probability distribution

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Cutting the race short: A Game theoretic approach

Zero –sum Bayesian Stackelberg game

- Leader – chooses defense f()
- Follower – chooses attack h()

Bayesian: incomplete information

Zero-sum: gain is loss (and vice versa)

Optimal Strategy for the user

Choose $f(r'|r),\ x_{r'} = \min_{\hat{r}} \sum_{r} \psi(r) f(r'|r) d_p(\hat{r}, r)$

maximize $\sum_{r'} x_{r'}$

s.t. $x_r' \le \sum_{r} \psi(r) f(r'|r) d_p(\hat{r}, r), \forall \hat{r}, r'$

$\sum_{r} \psi(r) \sum_{r'} f(r'|r) d_q(r', r), \le Q_{loss}^{max}$

$\sum_{r'} f(r'|r) = 1,\ f(r'|r) \ge 0, \forall r, r'$

Optimal Strategy for the ADV

Choose $h(\hat{r}|r'),\ y_{r'} = \max_{r'} \sum_{\hat{r}} h(\hat{r}|r) d_p(\hat{r}, r)$

minimize $\sum_{r} \psi(r) y_r + z Q_{loss}^{max}$

s.t. $y_r \ge \sum_{\hat{r}} h(\hat{r}|r') d_p(\hat{r}, r) + z d_q(r', r), \forall r, r'$

$\sum_{\hat{r}} h(\hat{r}|r') = 1,\ h(\hat{r}|r') \ge 0, \forall r', \hat{r}$

$z \ge 0$

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Cutting the race short: A Game theoretic approach

Zero –sum Bayesian Stackelberg game

- Leader – chooses defense f()
- Follower – chooses attack h()

Bayesian: incomplete information

Zero-sum: gain is loss (and vice versa)

Optimal Strategy for the user

Choose $f(r'|r),\ x_{r'} = \min_{\hat{r}} \sum_r \psi(r) f(r'|r) dp(\hat{r}, r)$

maximize $\sum_{r'} x_{r'}$

s.t. $x_r' \le \sum_r \psi(r) f(r'|r) d_p(\hat{r}, r), \forall \hat{r}, r'$

$\sum_r \psi(r) \sum_{r'} f(r'|r) d_q(r', r), \le Q_{loss}^{max}$

$\sum_{r'} f(r'|r) = 1,\ f(r'|r) \ge 0, \forall r, r'$

Optimal Strategy for the ADV

Choose $h(\hat{r}|r'),\ y_{r'} = \max_{r'} \sum_r h(\hat{r}|r) d_p(\hat{r}, r)$

minimize $\sum_r \psi(r) y_r + zQ_{loss}^{max}$

s.t. $y_r \ge \sum_{\hat{r}} h(\hat{r}|r') d_p(\hat{r}, r) + zd_q(r', r), \forall r, r'$

$\sum_{\hat{r}} h(\hat{r}|r') = 1,\ h(\hat{r}|r') \ge 0, \forall r', \hat{r}$

$z \ge 0$

minimize privacy

quality constraint

h() is a probability distribution

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Cutting the race short: A Game theoretic approach

Zero –sum Bayesian Stackelberg game

- Leader – chooses defense f()
- Follower – chooses attack h()

Bayesian: incomplete information

Zero-sum: gain is loss (and vice versa)

Optimal Strategy for the user

Choose $f(r'|r),\ x_{r'} = \min_{\hat{r}} \sum_r \psi(r) f(r'|r) dp(\hat{r}, r)$

maximize $\sum_{r'} x_{r'}$

s.t. $x_{r'} \leq \sum_r \psi(r) f(r'|r) d_p(\hat{r}, r), \forall \hat{r}, r'$

$\sum_r \psi(r) \sum_{r'} f(r'|r) d_q(r', r), \leq Q_{loss}^{max}$

$\sum_{r'} f(r'|r) = 1,\ f(r'|r) \geq 0, \forall r, r'$

Optimal Strategy for the ADV

Choose $h(\hat{r}|r'),\ y_{r'} = \max_{r'} \sum_r h(\hat{r}|r) d_p(\hat{r}, r)$

minimize $\sum_r \psi(r) y_r + z Q_{loss}^{max}$

s.t. $y_r \geq \sum_{\hat{r}} h(\hat{r}|r') d_p(\hat{r}, r) + z d_q(r', r), \forall r, r'$

$\sum_{\hat{r}} h(\hat{r}|r') = 1,\ h(\hat{r}|r') \geq 0, \forall r', \hat{r}$

$z \geq 0$

maximize privacy

quality constraint

h() is a probability distribution

shadow price

R. Shokri, G. Theodorakopoulos, C. Troncoso, J. Hubaux, J. Le Boudec, "Protecting Location Privacy: Optimal Strategy against Localization Attacks" CCS 2012
G. Theodorakopoulos, R. Shokri, C. Troncoso, J-P. Hubaux, J-Y Le Boudec "Prolonging the Hide-and-Seek Game: Optimal Trajectory Privacy for Location-Based Services" WPES 2014

Are we there yet?

Privacy by design rocks!  but realizing it is non-trivial

Are we there yet?

Privacy by design rocks!  but realizing it is non-trivial

PART I:
Reasoning about Privacy when designing systems

⇕

Explicit privacy engineering activities

PART II:
Designing Technologies to support Privacy-aware designs

⇕

Systematic design methods for obfuscation mechanisms

Are we there yet?

Privacy by design rocks!

but realizing it is non-trivial

PART I:
Reasoning about Privacy when designing systems

⇕

Explicit privacy engineering activities

Fully fledged methodology?

Requirements? Evaluation?

PART II:
Designing Technologies to support Privacy-aware designs

⇕

Systematic design methods
for obfuscation mechanisms

Are we there yet?

Privacy by design rocks!  but realizing it is non-trivial

PART I:
Reasoning about Privacy when designing systems

Explicit privacy engineering activities

PART II:
Designing Technologies to support Privacy-aware designs

Systematic design methods for obfuscation mechanisms

Fully fledged methodology?

Requirements? Evaluation?

Accessible PETS

Understanding? Implementation?

ARE WE THERE YET?

PRIVACY BY DESIGN ROCKS!  BUT REALIZING IT IS NON-TRIVIAL

PART I:
REASONING ABOUT PRIVACY WHEN DESIGNING SYSTEMS

⇕

Explicit privacy engineering activities

Fully fledged methodology

- Requirements? Evaluation?

Accessible PETS

- Understanding? Implementation?

PART II:
DESIGNING TECHNOLOGIES TO SUPPORT PRIVACY-AWARE DESIGNS

⇕

Systematic design methods for obfuscation mechanisms

Strong assumption's dependency

Are we there yet?

Privacy by design rocks!

but realizing it is non-trivial

PART I:
Reasoning about Privacy when designing systems

Explicit privacy engineering activities

- Fully fledged methodology
  - Requirements? Evaluation?
- Accessible PETS
  - Understanding? Implementation?

PART II:
Designing Technologies to support Privacy-aware designs

Systematic design methods for obfuscation mechanisms

- Strong assumption's dependency
- High computational cost

Are we there yet?

Privacy by design rocks!

but realizing it is non-trivial

PART I:
Reasoning about Privacy when designing systems

⇕

Explicit privacy engineering activities

Fully fledged methodology

- Requirements? Evaluation?

Accessible PETS

- Understanding? Implementation?

PART II:
Designing Technologies to support Privacy-aware designs

⇕

Systematic design methods for obfuscation mechanisms

Strong assumption's dependency

High computational cost

Lack of standard metrics

Are we there yet?

Privacy by design rocks!

but realizing it is non-trivial

PART I:
Reasoning about Privacy when designing systems

⇕

Explicit privacy engineering activities

Fully fledged methodology

- Requirements? Evaluation?

Accessible PETS

- Understanding? Implementation?

PART II:
Designing Technologies to support Privacy-aware designs

⇕

Systematic design methods for obfuscation mechanisms

Strong assumption's dependency

High computational cost

Lack of standard metrics

Universal?

THANKS!

ANY QUESTIONS?

More about privacy:
https://www.petsymposium.org/
http://www.degruyter.com/view/j/popets

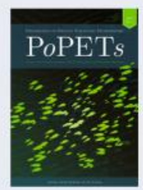

carmela.troncoso@imdea.org
https://software.imdea.org/~carmela.troncoso/
(these slides will be there soon)

Template: http://www.brainybetty.com/
Figures: SlidesCarnival