How to Bootstrap Anonymous Communication

Sune K. Jakobsen[1] Claudio Orlandi[2]

[1]Queen Mary, University of London

[2]Aarhus University

January 16, 2016

How can you get anonymity?

Contact a journalist or publisher, and tell them you want to be anonymous.

How can you get anonymity?

Contact a journalist or publisher, and tell them you want to be anonymous.

Use Tor network. Here your message will go through 3 different servers, before it is sent to the recipient.

How can you get anonymity?

Contact a journalist or publisher, and tell them you want to be anonymous.

Use Tor network. Here your message will go through 3 different servers, before it is sent to the recipient.

Use SecureDrop. A hidden service on Tor that media can host.

How can you get anonymity?

Contact a journalist or publisher, and tell them you want to be anonymous.

Use Tor network. Here your message will go through 3 different servers, before it is sent to the recipient.

Use SecureDrop. A hidden service on Tor that media can host.

Other suggestions: Vuvuzela, Riposte, Dissent, cMix/Privategrity.

What if no one can help you?

If anonymous communication is banned, these method are not going to work anymore.

What if no one can help you?

If anonymous communication is banned, these method are not going to work anymore.

What can you do if no one will help you?

What if no one can help you?

If anonymous communication is banned, these method are not going to work anymore.

What can you do if no one will help you?

Cryptogenography: Without assumption on the computational power of the adversary, many people can each reveal 3.1 bits while keeping 5%'s doubt about who is leaking.

What if no one can help you?

If anonymous communication is banned, these method are not going to work anymore.

What can you do if no one will help you?

Cryptogenography: Without assumption on the computational power of the adversary, many people can each reveal 3.1 bits while keeping 5%'s doubt about who is leaking.

What can we do if the adversary has bounded computational power?

Problem

One person, Lea, has some information x she wants to reveal to a journalist Joe.

Problem

One person, Lea, has some information x she wants to reveal to a journalist Joe.

She do not want Joe to learn that the information came from her.

Problem

One person, Lea, has some information x she wants to reveal to a journalist Joe.

She do not want Joe to learn that the information came from her.

She can publish files on a site where other people publish files, e.g. Instagram or YouTube.

Problem

One person, Lea, has some information x she wants to reveal to a journalist Joe.

She do not want Joe to learn that the information came from her.

She can publish files on a site where other people publish files, e.g. Instagram or YouTube.

We assume that she has access to a limited anonymous channel.

Problem

One person, Lea, has some information x she wants to reveal to a journalist Joe.

She do not want Joe to learn that the information came from her.

She can publish files on a site where other people publish files, e.g. Instagram or YouTube.

We assume that she has access to a limited anonymous channel.

Can she send x to Joe, if x has more bits than what she can send over the channel?

Steganography

Steganography means concealed writing.

Unlike cryptography, steganography hides the fact that there is a secret message.

Steganography

Steganography means concealed writing.

Unlike cryptography, steganography hides the fact that there is a secret message.

This is used by Message in a Bottle. [Invernizzi-Kruegel-Giovanni 2013]

Anonymous Steganography Scheme

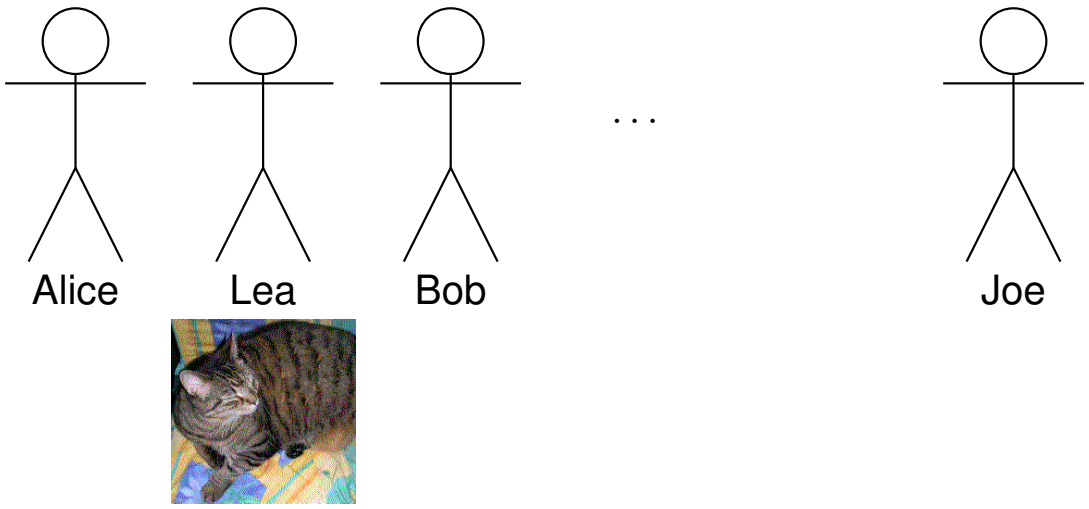

Lea uses an algorithm Gen to generate a key ek, and then use the key to generate a random looking string $c \leftarrow \mathrm{Enc}_{ek}(x)$. This string is then embedded into a picture using steganography.

Anonymous Steganography Scheme

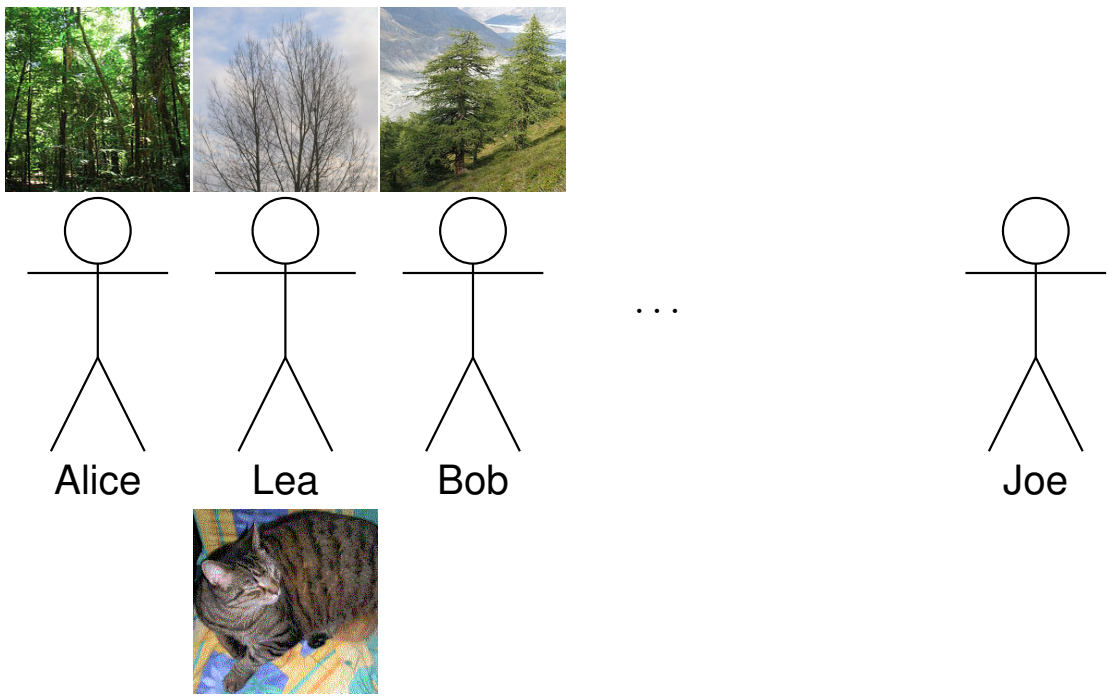

Everyone uploads a picture. Lea uploads a picture with c embedded.

Anonymous Steganography Scheme

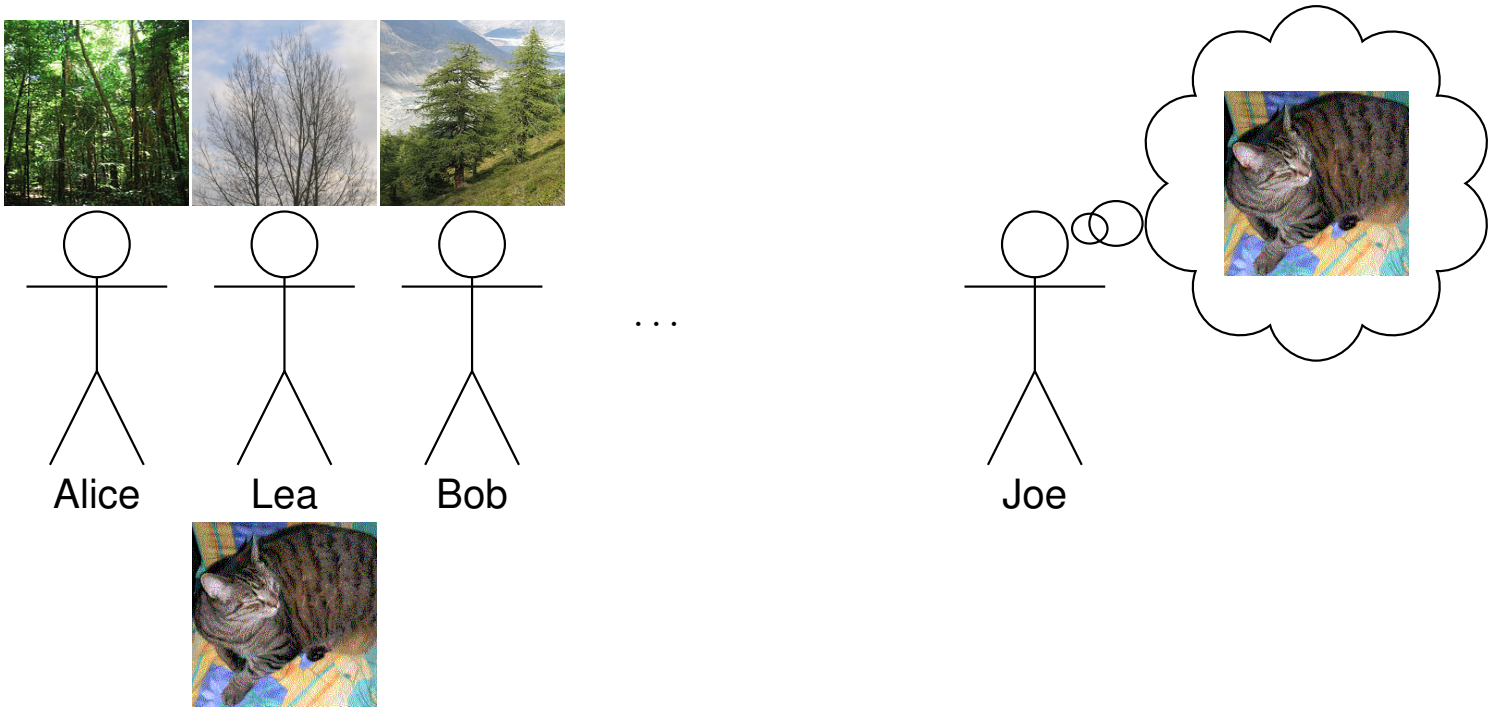

We want Joe to be able to extract x using an algorithm Dec(t). However, if he could do this independently of the other pictures, he could figure out who sent x.

Anonymous Steganography Scheme

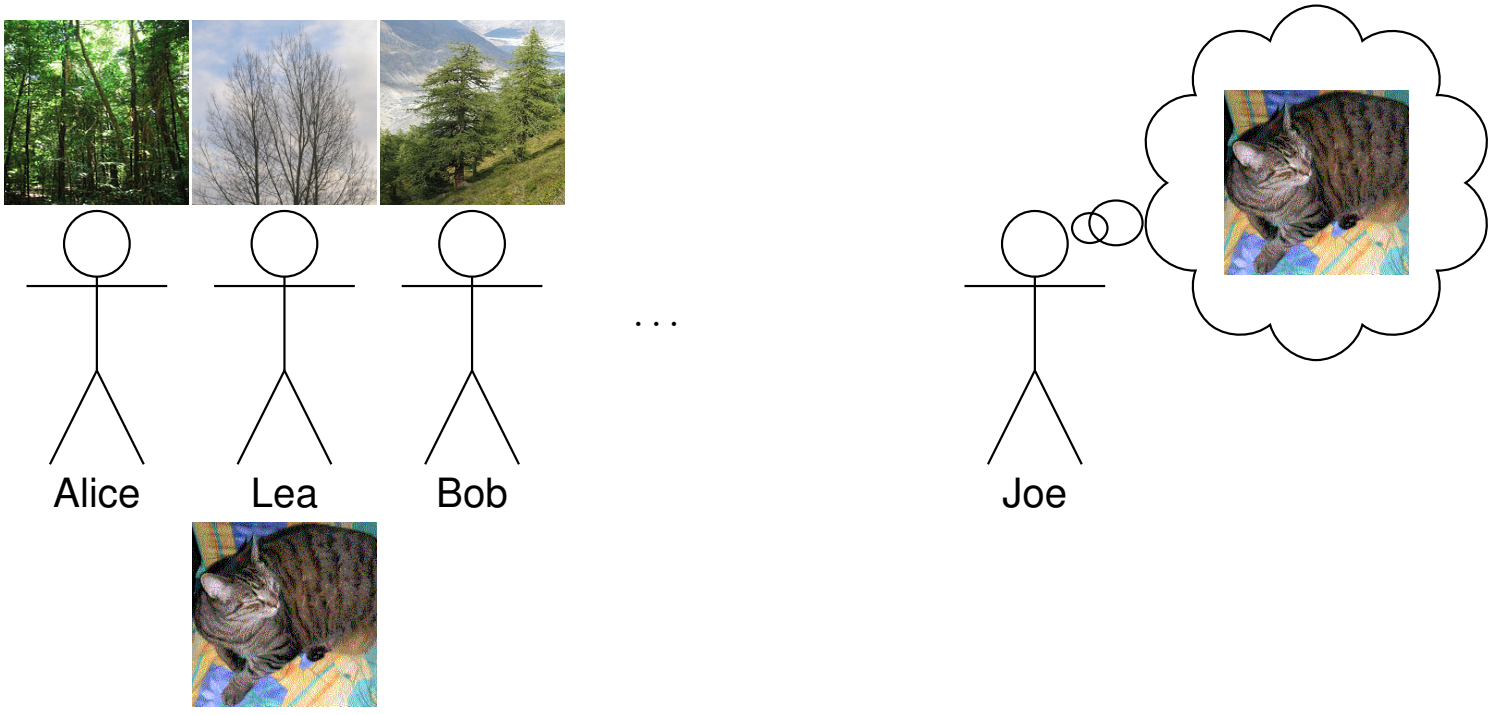

To avoid this, we have to ensure that Joe can only use Dec on the entire transcript t. We let Lea generate a key $dk \leftarrow \mathrm{KeyEx}_{ek}(t, i)$. Now Lea sends dk over the anonymous channel. Joe computes $x' \leftarrow \mathrm{Dec}_{dk}(t)$.

Anonymous Steganogarphy Scheme

An *anonymous steganography scheme* it a tuple $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{KeyEx}, \mathsf{Dec})$ with

$$
\begin{aligned}
ek &\leftarrow \mathsf{Gen}(1^\lambda) \\
c &\leftarrow \mathsf{Enc}_{ek}(x) \\
dk &\leftarrow \mathsf{KeyEx}_{ek}(t, i) \\
x' &= \mathsf{Dec}_{dk}(t)
\end{aligned}
$$

which achieves *correctness*, *compactness* ($|dk| < |x|$) and is *anonymous* (next slide).

Anonymity

Challenger

Adversary

x, i_0, i_1 ←

$b \leftarrow \{0,1\}$

$ek \leftarrow \mathsf{Gen}(\lambda)$

$t_{i_b} \leftarrow \mathsf{Enc}_{ek}(x)$

$t_{i_{1-b}} \leftarrow \{0,1\}^l$ → t_{i_0}, t_{i_1}

← $t_1, t_2, \ldots, t_n$

$dk \leftarrow \mathsf{KeyEx}_{ek}(i_b, t)$

→ dk

Guess b

Results

Theorem

Assuming the existence of homomorphic encryption and indistinguishability obfuscators for all polynomially sized circuits, there exist an anonymous steganography scheme.

Results

Theorem

Assuming the existence of homomorphic encryption and indistinguishability obfuscators for all polynomially sized circuits, there exist an anonymous steganography scheme.

Theorem

Any anonymous steganography scheme must have dk of length more than $O(\log(\lambda))$

Results

Theorem

Assuming the existence of homomorphic encryption and indistinguishability obfuscators for all polynomially sized circuits, there exist an anonymous steganography scheme.

Theorem

Any anonymous steganography scheme must have dk of length more than $O(\log(\lambda))$

The lower bound holds even if we only require polynomially small probability of success, and allow the leaker to send multiple messages.

Construction, sketch

Each $c^j = t_i^j$ is an encryption of x^j.

Construction, sketch

Each $c^j = t_i^j$ is an encryption of x^j.

dk contains a homomorphic encryption of i.

Construction, sketch

Each $c^j = t_i^j$ is an encryption of x^j.

dk contains a homomorphic encryption of i.

For each j Joe can compute an encryption of t_i^j, without knowing i.

Construction, sketch

Each $c^j = t_i^j$ is an encryption of x^j.

dk contains a homomorphic encryption of i.

For each j Joe can compute an encryption of t_i^j, without knowing i.

If Joe only got this information he could use a vector commitment scheme to commit to these encryptions.

Construction, sketch

Each $c^j = t_i^j$ is an encryption of x^j.

dk contains a homomorphic encryption of i.

For each j Joe can compute an encryption of t_i^j, without knowing i.

If Joe only got this information he could use a vector commitment scheme to commit to these encryptions.

Lea can also make these computations, and build a circuit that takes as input j, an encryption of t_i^j and a correct opening, and decrypts to x^j.

Construction, sketch

Each $c^j = t_i^j$ is an encryption of x^j.

dk contains a homomorphic encryption of i.

For each j Joe can compute an encryption of t_i^j, without knowing i.

If Joe only got this information he could use a vector commitment scheme to commit to these encryptions.

Lea can also make these computations, and build a circuit that takes as input j, an encryption of t_i^j and a correct opening, and decrypts to x^j.

Lea includes an obfuscation of this circuit in dk and send it all to Joe at the same time.

Construction, sketch

Each $c^j = t_i^j$ is an encryption of x^j.

dk contains a homomorphic encryption of i.

For each j Joe can compute an encryption of t_i^j, without knowing i.

If Joe only got this information he could use a vector commitment scheme to commit to these encryptions.

Lea can also make these computations, and build a circuit that takes as input j, an encryption of t_i^j and a correct opening, and decrypts to x^j.

Lea includes an obfuscation of this circuit in dk and send it all to Joe at the same time.

To make the proof work, you need to have two independent encryptions of i and use a somewhere statistically binding vector commitment scheme [Hubáček-Wichs 2015].

Open problems

Can we make an anonymous steganography scheme without use indistinguishability obfuscation?

Can the leaker avoid downloading all the uploaded files, and instead use a hash of the files?

Questions?