Traffic Analysis

high latency anonymous communications

Carmela Troncoso*
IMDEA Software Institute

*Thanks to George Danezis for sharing slides

Privacy in electronic communications

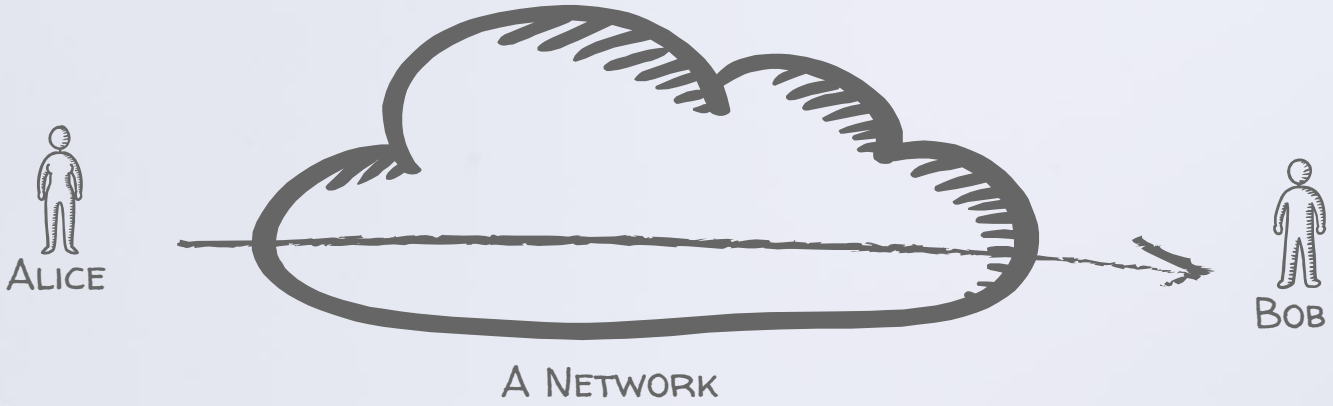

Privacy in electronic communications

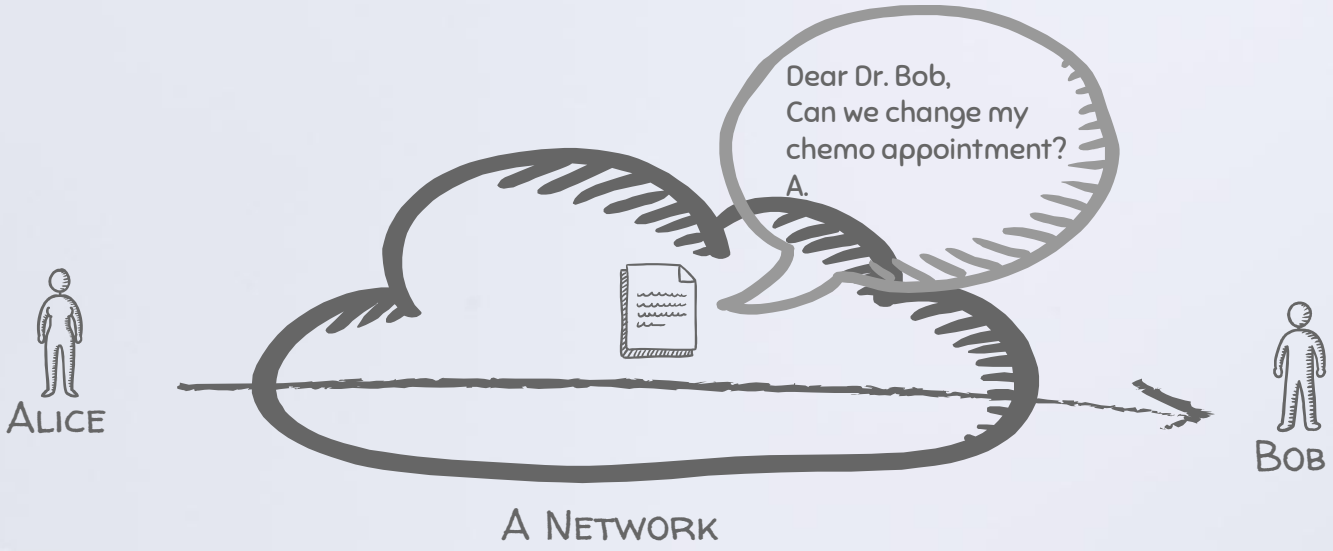

Privacy in electronic communications

Intelligence agencies

Privacy in electronic communications

Intelligence agencies

ISPs

Alice

Dear Dr. Bob,
Can we change my
chemo appointment?
A.

Bob

A Network

Privacy in electronic communications

Intelligence agencies

The Boss

ISPs

Alice

A Network

Bob

Privacy in electronic communications

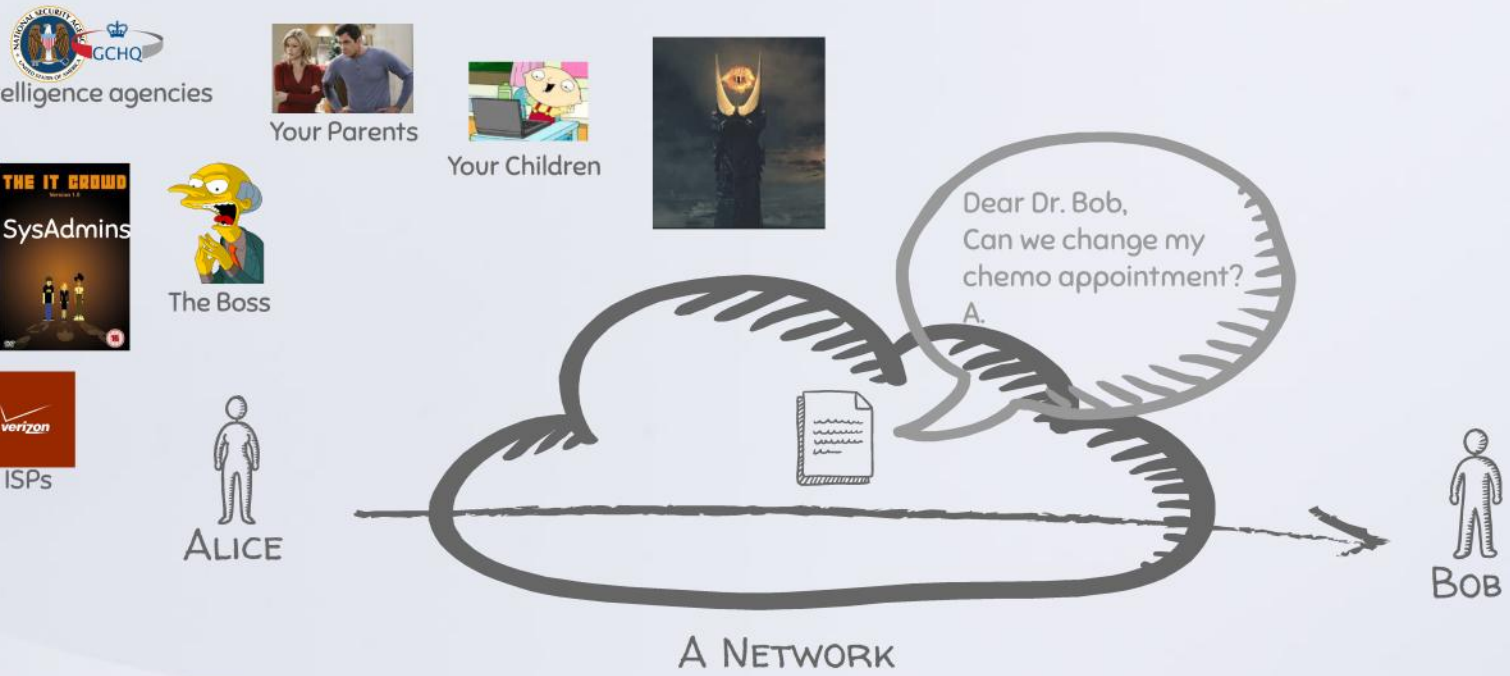

Privacy in electronic communications

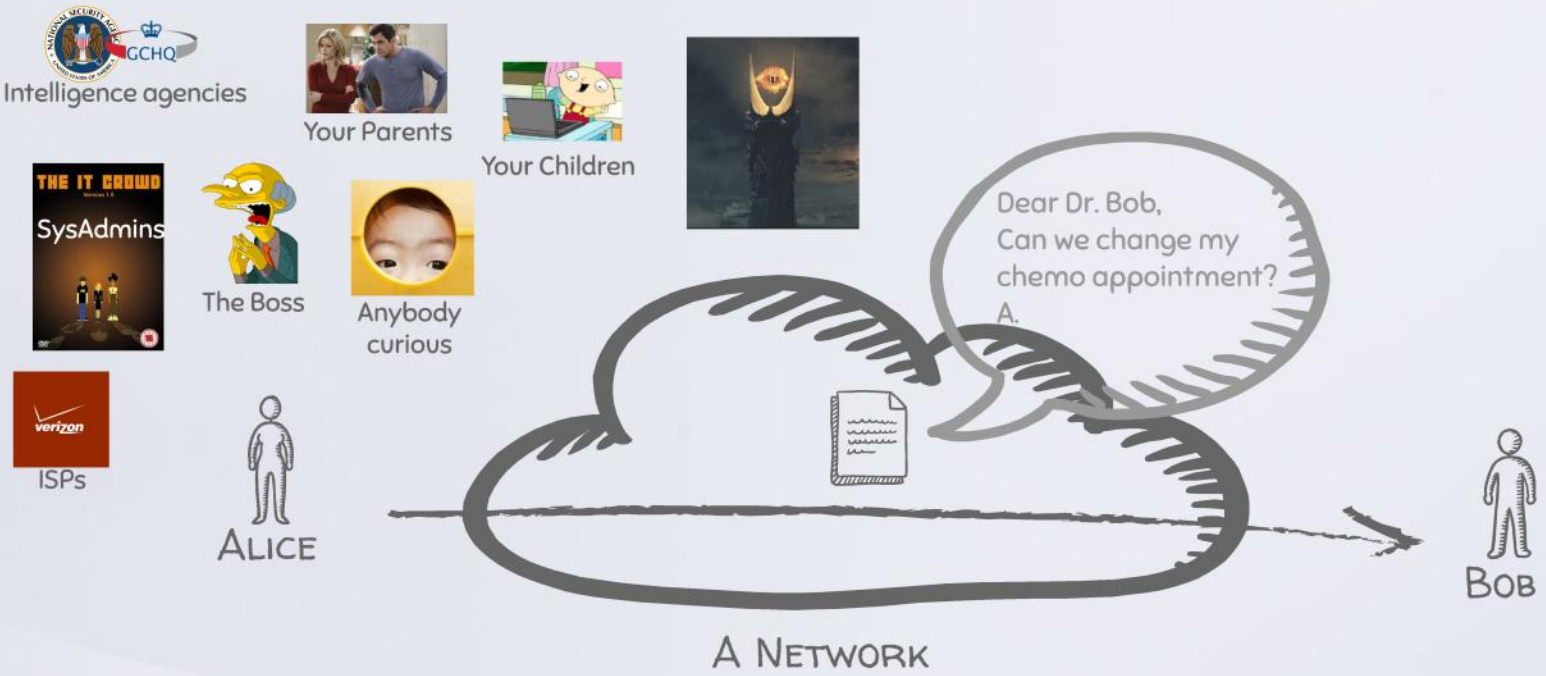

Privacy in electronic communications

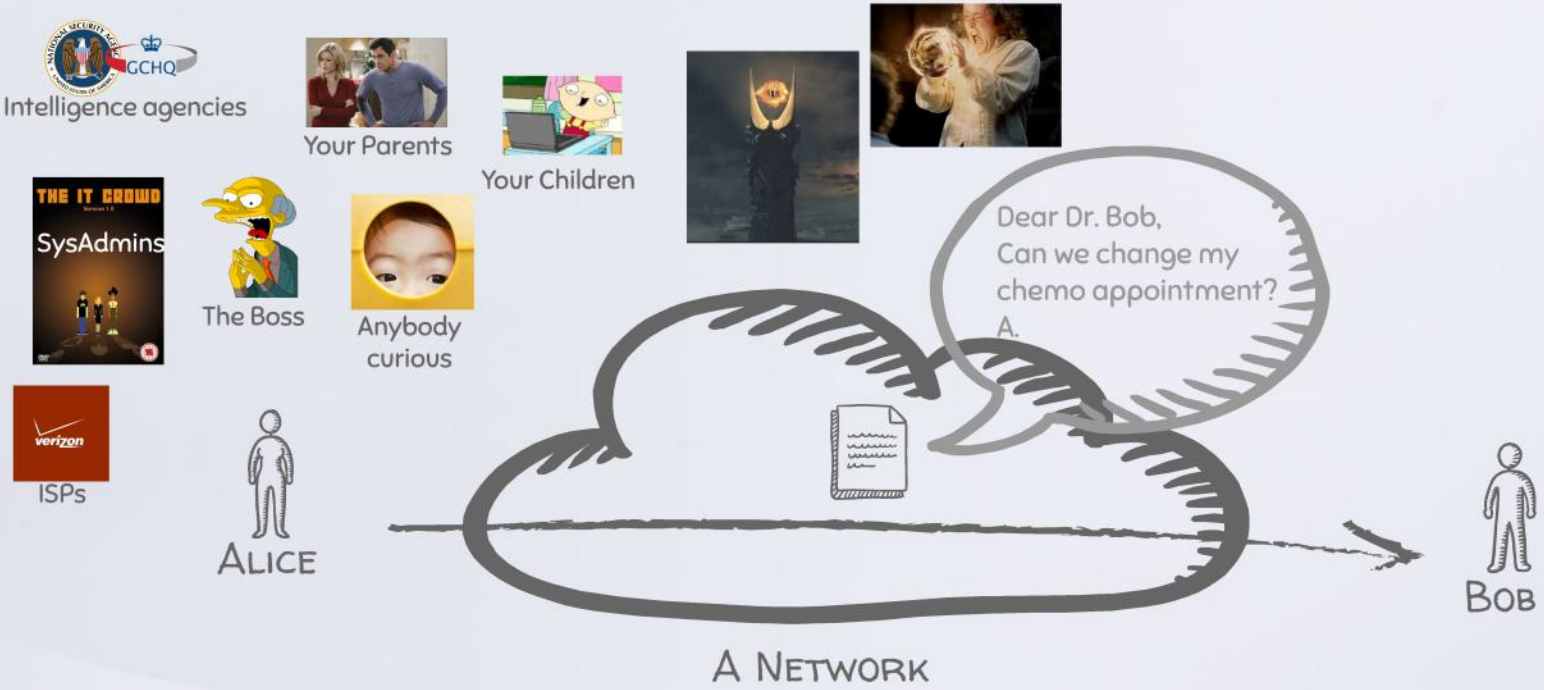

Privacy in electronic communications

Intelligence agencies

Your Parents

Your Children

The Boss

Anybody curious

ISPs

Alice

Bob

A Network

But we can encrypt! What is the problem?

But we can encrypt! What is the problem?

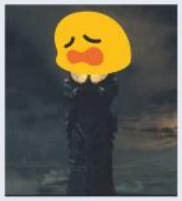

But we can encrypt! What is the problem?

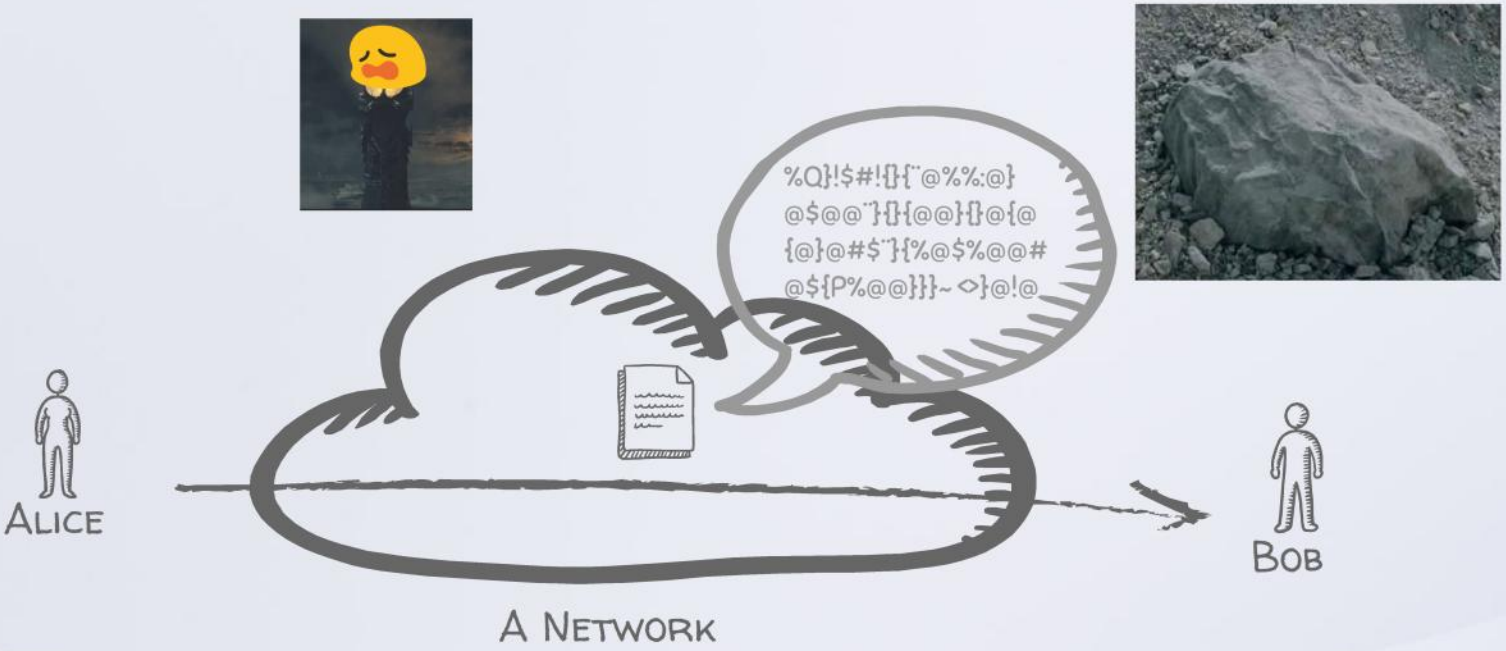

But we can encrypt! What is the problem?

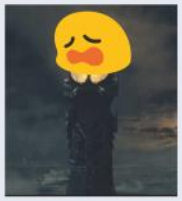

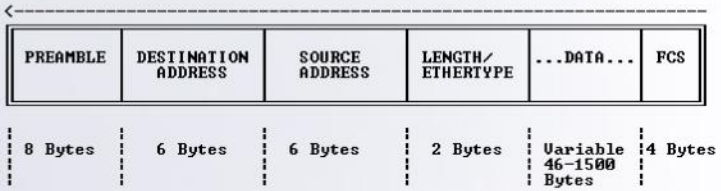

Ethernet
(IEEE 802.3, 1997)

But we can encrypt! What is the problem?

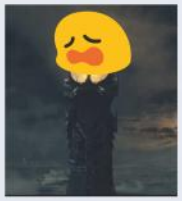

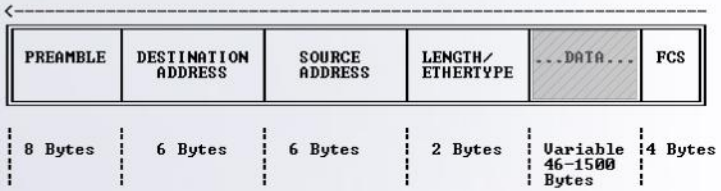

Ethernet
(IEEE 802.3, 1997)

But we can encrypt! What is the problem?

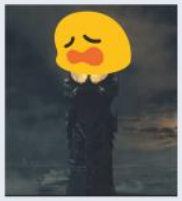

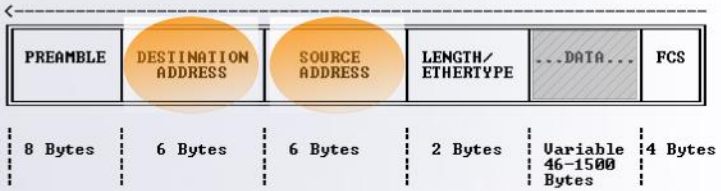

Ethernet
(IEEE 802.3, 1997)

But we can encrypt! What is the problem?

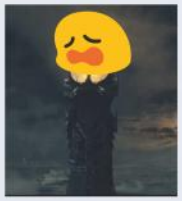

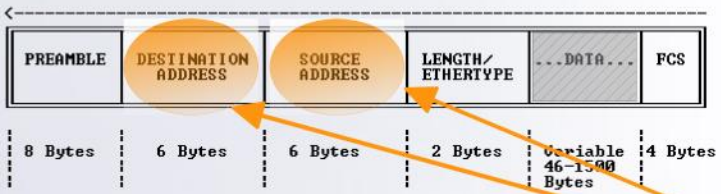

Ethernet
(IEEE 802.3, 1997)

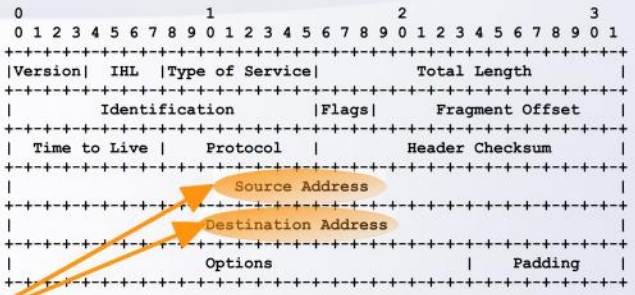

IPv4 Header
(RFC 791, 1981)

Weak identifier

Same for TCP, SMTP, IRC, HTTP, ...

But we can encrypt! What is the problem?

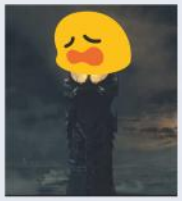

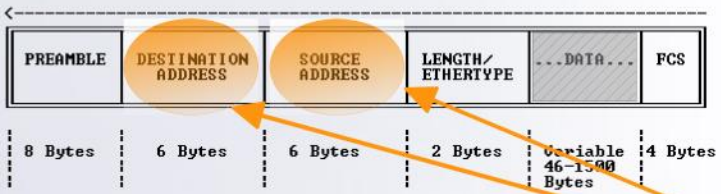

Ethernet
(IEEE 802.3, 1997)

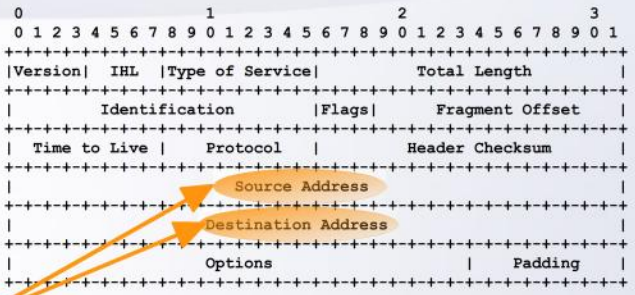

IPv4 Header
(RFC 791, 1981)

Weak identifier

Same for TCP, SMTP, IRC, HTTP, ...

OMG!! The problem is Traffic Analysis!!

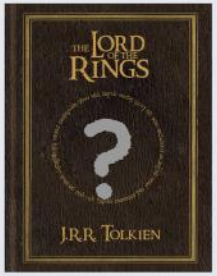

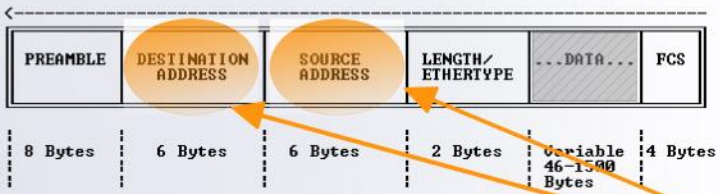

Ethernet
(IEEE 802.3, 1997)

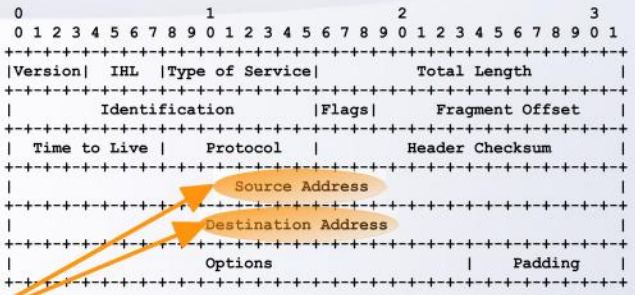

IPv4 Header
(RFC 791, 1981)

Weak identifier

Same for TCP, SMTP, IRC, HTTP, ...

Traffic WHAT?

Wikipedia: traffic analysis is the process of intercepting and examining messages in order to deduce information from patterns in communication

Making use of "just" traffic data of a communication (aka metadata) to extract information

(as opposed to analyzing content or perform cryptanalysis)

Traffic WHAT?

Wikipedia: traffic analysis is the process of intercepting and examining messages in order to deduce information from patterns in communication

Making use of "just" traffic data of a communication (aka metadata) to extract information (as opposed to analyzing content or perform cryptanalysis)

Identities of communicating parties

Timing, frequency, duration

Location

Volume

Device

Traffic WHAT?

Wikipedia: traffic analysis is the process of intercepting and examining messages in order to deduce information from patterns in communication

Making use of "just" traffic data of a communication (aka metadata) to extract information (as opposed to analyzing content or perform cryptanalysis)

Identities of communicating parties

Timing, frequency, duration

Location

Volume

Device

Military Roots

– M. Herman: "These non-textual techniques can establish **targets' locations**, order-of-battle and **movement**. Even when messages are not being deciphered, traffic analysis of the target's Command, Control, Communications and intelligence system and its patterns of behavior provides indications of his **intentions** and **states of mind**"

– **WWI**: British troops finding German boats.

– **WWII**: assessing size of German Air Force, fingerprinting of transmitters or operators (localization of troops).

Herman, Michael. Intelligence power in peace and war. Cambridge University Press, 1996.
Diffie, Whitfield, and Susan Landau. Privacy on the line: The politics of wiretapping and encryption. MIT press, 2010.
http://www.theguardian.com/world/interactive/2013/nov/01/snowden-nsa-files-surveillance-revelations-decoded

Traffic WHAT?

Wikipedia: traffic analysis is the process of intercepting and examining messages in order to deduce information from patterns in communication

Making use of "just" traffic data of a communication (aka metadata) to extract information (as opposed to analyzing content or perform cryptanalysis)


Identities of communicating parties


Timing, frequency, duration


Location


Volume


Device

Military Roots

– M. Herman: "These non-textual techniques can establish **targets' locations**, order-of-battle and **movement**. Even when messages are not being deciphered, traffic analysis of the target's Command, Control, Communications and intelligence system and its patterns of behavior provides indications of his **intentions** and **states of mind**"

– **WWI**: British troops finding German boats.

– **WWII**: assessing size of German Air Force, fingerprinting of transmitters or operators (localization of troops).

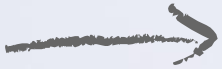

Nowadays

– Diffie&Landau: "Traffic analysis, not cryptanalysis, is the backbone of communications intelligence"

– Stewart Baker (NSA): "metadata **absolutely tells you everything about somebody's life**. If you have enough metadata, you don't really need content."

– Tempora, MUSCULAR → XkeyScore, PRISM

– Also "good" uses: recommendations, location-based services,

Herman, Michael. Intelligence power in peace and war. Cambridge University Press, 1996.
Diffie, Whitfield, and Susan Landau. Privacy on the line: The politics of wiretapping and encryption. MIT press, 2010.
http://www.theguardian.com/world/interactive/2013/nov/01/snowden-nsa-files-surveillance-revelations-decoded

Actually, any **Meta Data** is sensitive!!

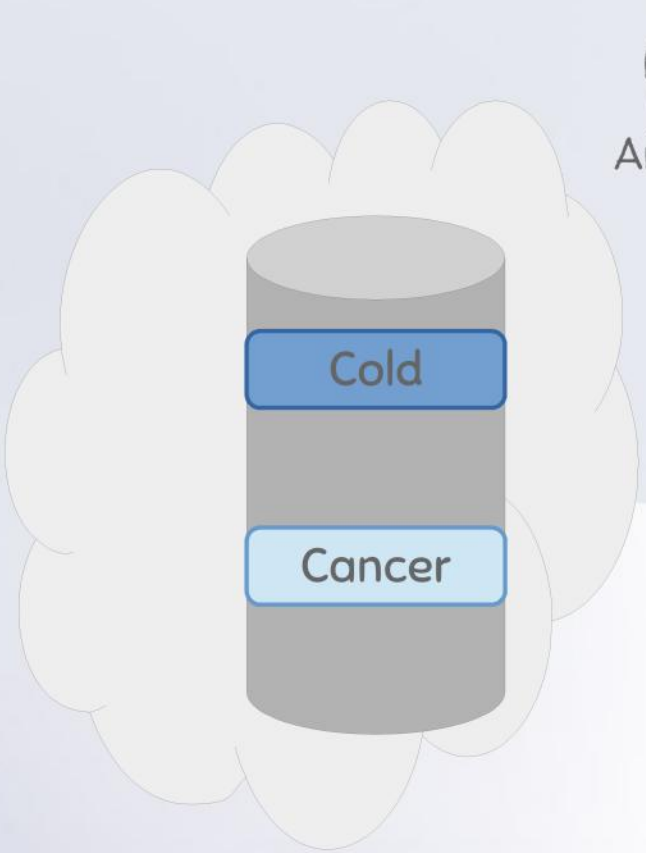

Actually, any Meta Data is sensitive!!

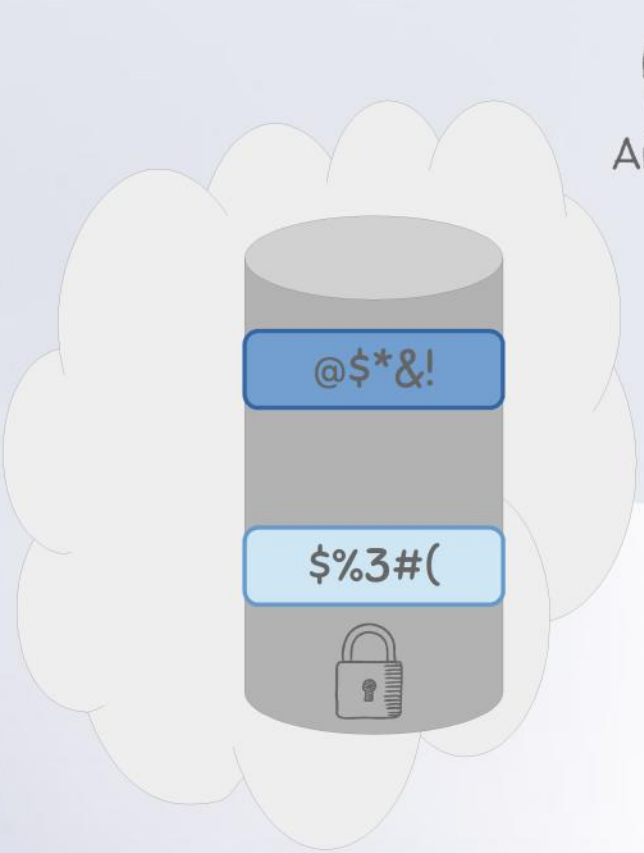

Actually, any Meta Data is sensitive!!

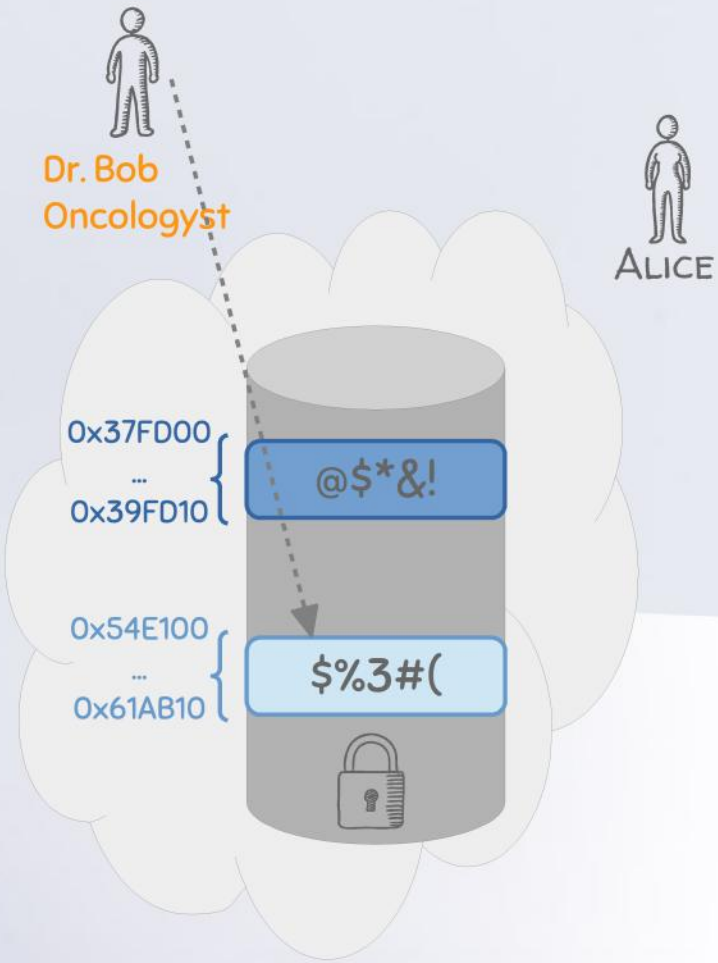

Actually, any Meta Data is sensitive!!

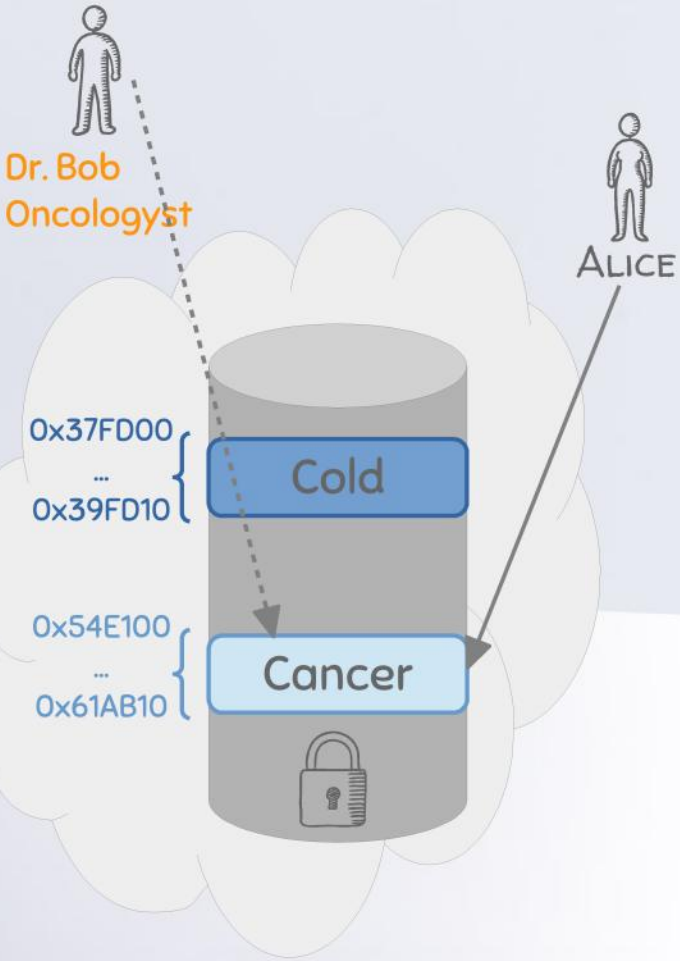

Actually, any Meta Data is sensitive!!

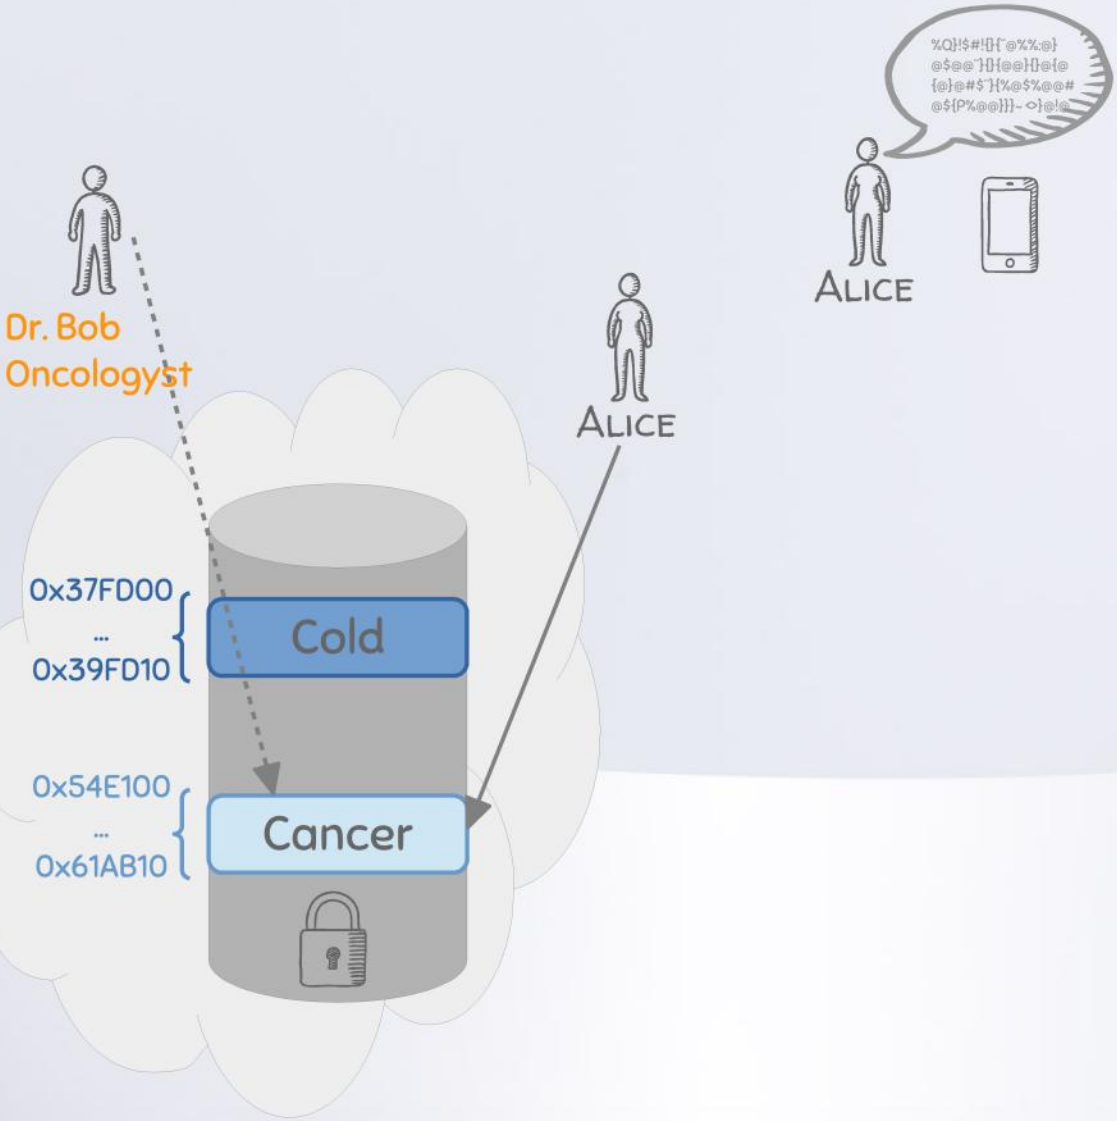

Actually, any Meta Data is sensitive!!

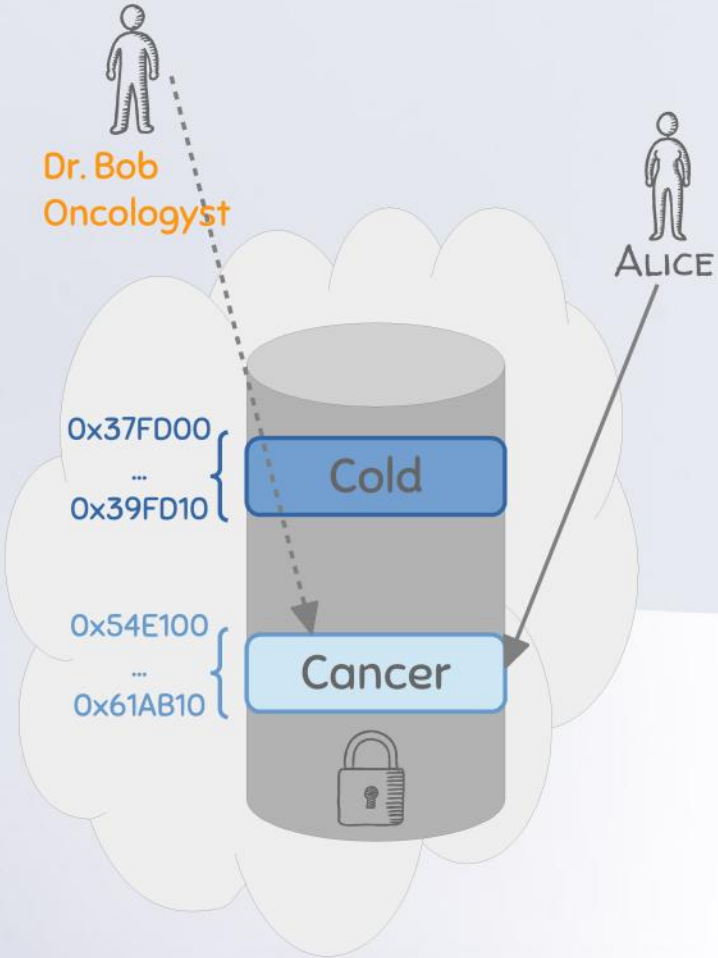

Actually, any Meta Data is sensitive!!

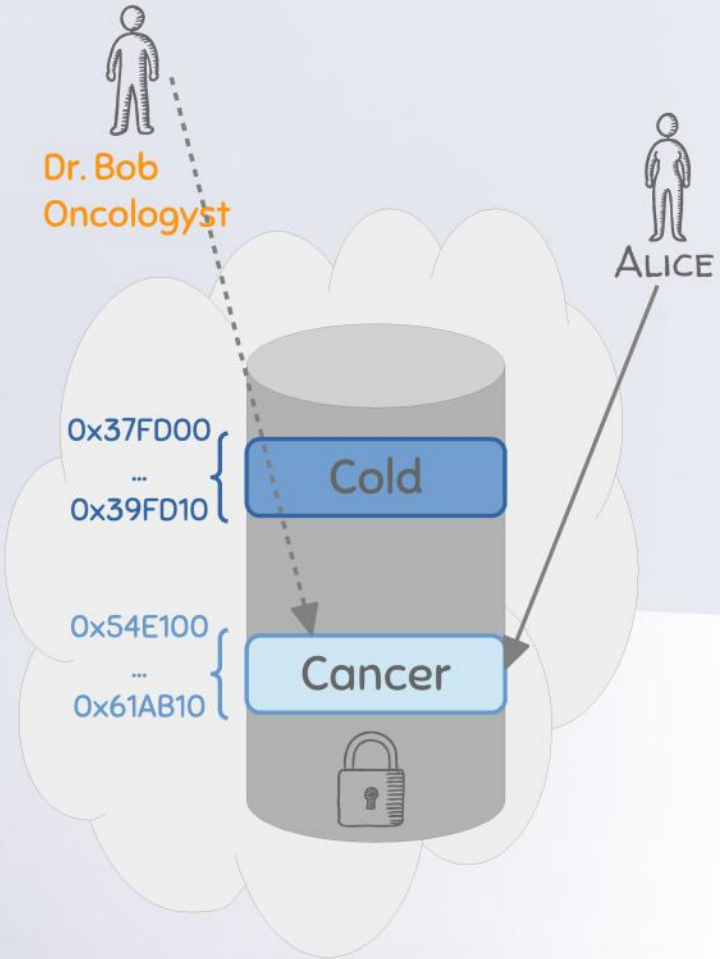

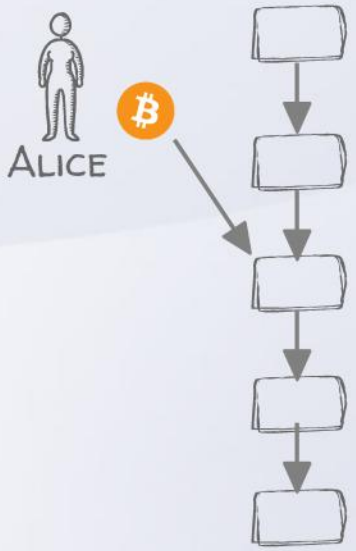

Actually, any Meta Data is sensitive!!

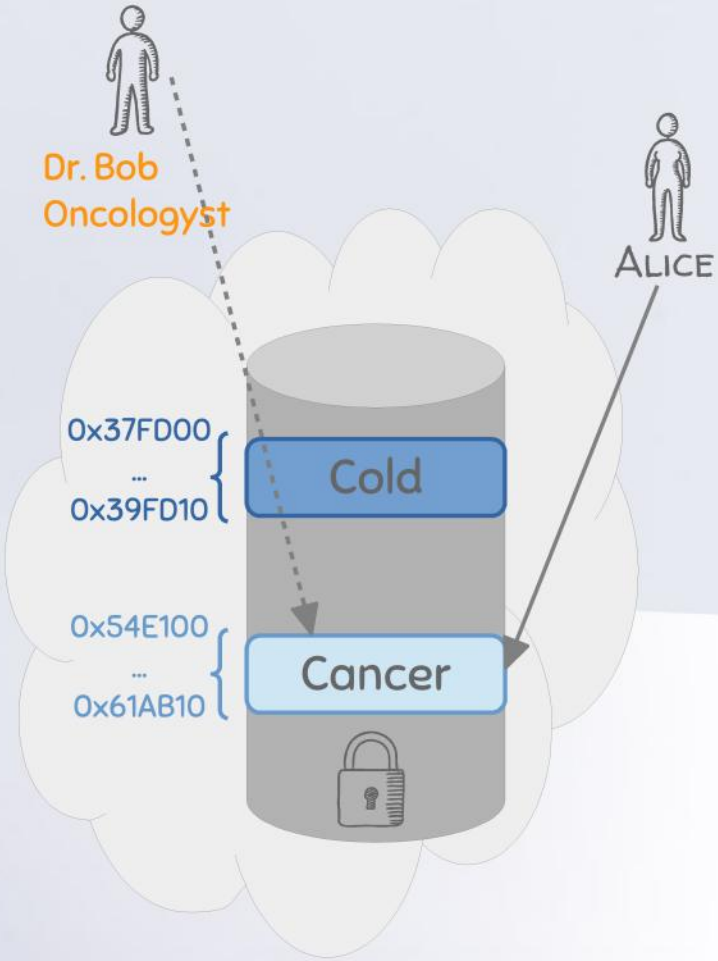

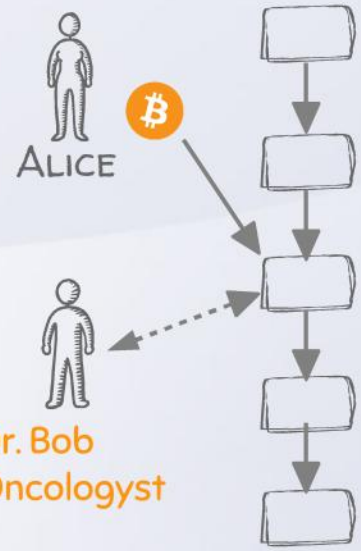

We need to protect the communication layer!

Anonymous communications

- **General applications**
  - Freedom of speech
  - Profiling / price discrimination
  - Spam avoidance
  - Investigation / market research
  - Censorship resistance
- **Specialized applications**
  - Electronic voting
  - Auctions / bidding / stock market
  - Incident reporting
  - Witness protection / whistle blowing
  - Showing anonymous credentials!

Anonymity is important to:

- the people who run some of the funniest parody Twitter accounts, such as @FeministHulk (SMASH THE PATRIARCHY!) or @BPGlobalPr during the Deepwater Horizon aftermath. San Francisco would not be better off if we knew who was behind @KarltheFog, the most charming personification of a major city's climate phenomenon.
- the young LGBTQ youth seeking advice online about coming out to their parents.
- the marijuana grower who needs to ask questions on an online message board about lamps and fertilizer or complying with state law, without publicly admitting to committing a federal offense.
- the medical patient seeking advice from other patients in coping with a chronic disease, whether it's alopecia, irritable bowel syndrome, cancer or a sexually transmitted infection.
- the online dater, who wants to meet new people but only reveal her identities after she's determined that potential dates are not creeps.
- the business that wants no-pulled-punches feedback from its customers.
- the World of Warcraft player, or any other MMOG gamer, who only wants to engage with other players in character.
- artists. Anonymity is integral to the work of The Yes Men, Banksy and Keizer.
- the low-income neighborhood resident who wants to comment on an article about gang violence in her community, without incurring retribution in the form of spray paint and broken windows.
- the boyfriend who doesn't want his girlfriend to know he's posing questions on a forum about how to pick out a wedding ring and propose. On the other end: Anonymity is important to anyone seeking advice about divorce attorneys online.
- the youth from an orthodox religion who secretly posts reviews on hip hop albums or R-rated movies.
- the young, pregnant woman who is seeking out advice on reproductive health services.
- the person seeking mental health support from an online community. There's a reason that support groups so often end their names with "Anonymous."
- the job seeker, in pursuit of cover letter and resume advice in a business blogger's comments, who doesn't want his current employer to know he is looking for work.
- many people's sexual lives, whether they're discussing online erotica or arranging kink meet-ups.
- Political Gabfest listeners. Each week, the hosts encourage listeners to post comments. Of the 262 largely positive customer reviews on iTunes, only a handful see value in using their real names.

https://www.eff.org/deeplinks/2013/10/online-anonymity-not-only-trolls-and-political-dissidents
http://geekfeminism.wikia.com/wiki/Who_is_harmed_by_a_%22Real_Names%22_policy%3F

Anonymous communications: abstract model

IDs
Timing
Volume
Length
...

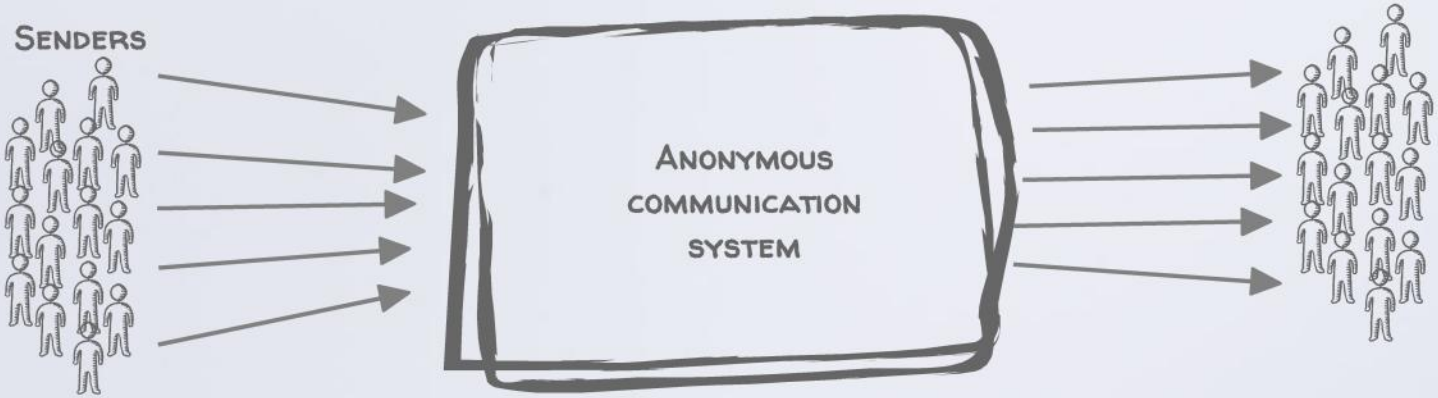

- Bitwise unlinkability
  - Crypto to make inputs and outputs bit patterns different
- (re)packetizing + (re)schedule
  - Destroy patterns (traffic analysis resistance)

Anonymous communications: abstract model

IDs
Timing
Volume
Length
...

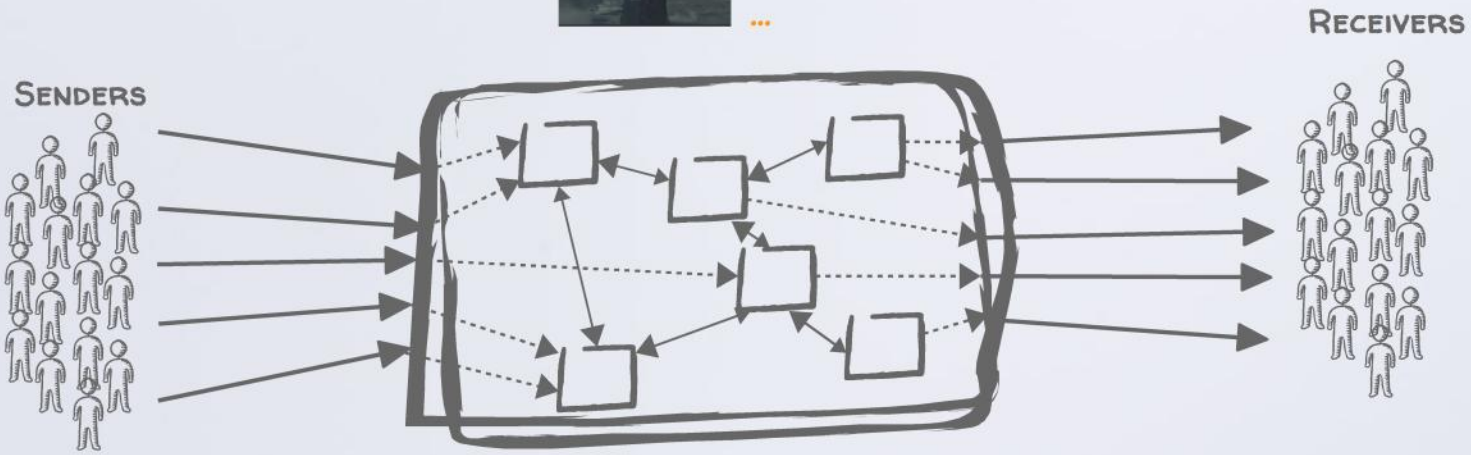

- Bitwise unlinkability
  - Crypto to make inputs and outputs bit patterns different
- (re)packetizing + (re)schedule + (re)routing,
  - Destroy patterns (traffic analysis resistance)
  - Load balancing
  - Distribute trust

Anonymous communications: abstract model

IDs
Timing
Volume
Length
...

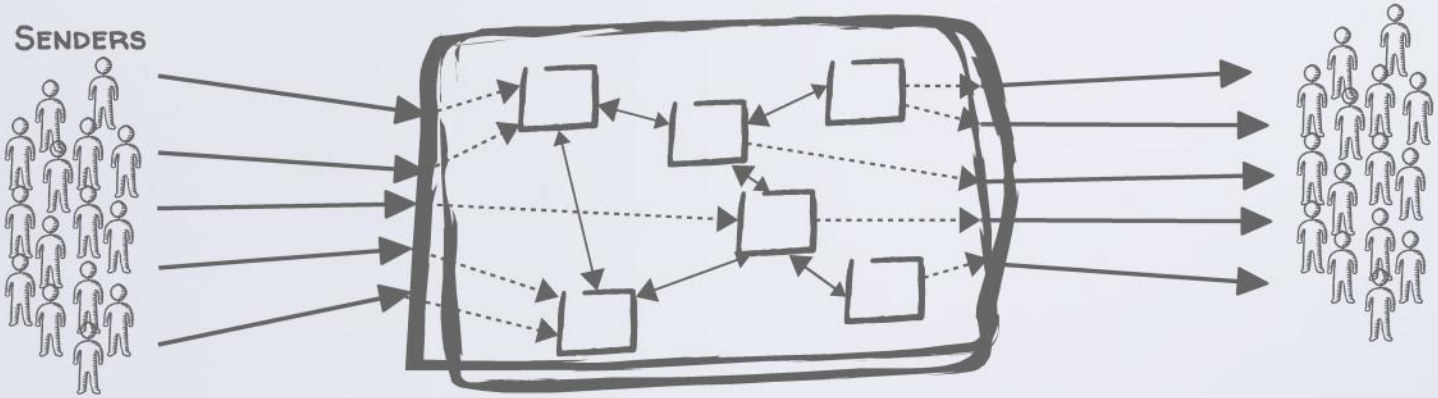

- Bitwise unlinkability
  - Crypto to make inputs and outputs bit patterns different
- (re)packetizing + (re)schedule + (re)routing,
  - Destroy patterns (traffic analysis resistance)
  - Load balancing
  - Distribute trust

In theory should work, but in practice...

IDs
Timing
Volume
Length
...

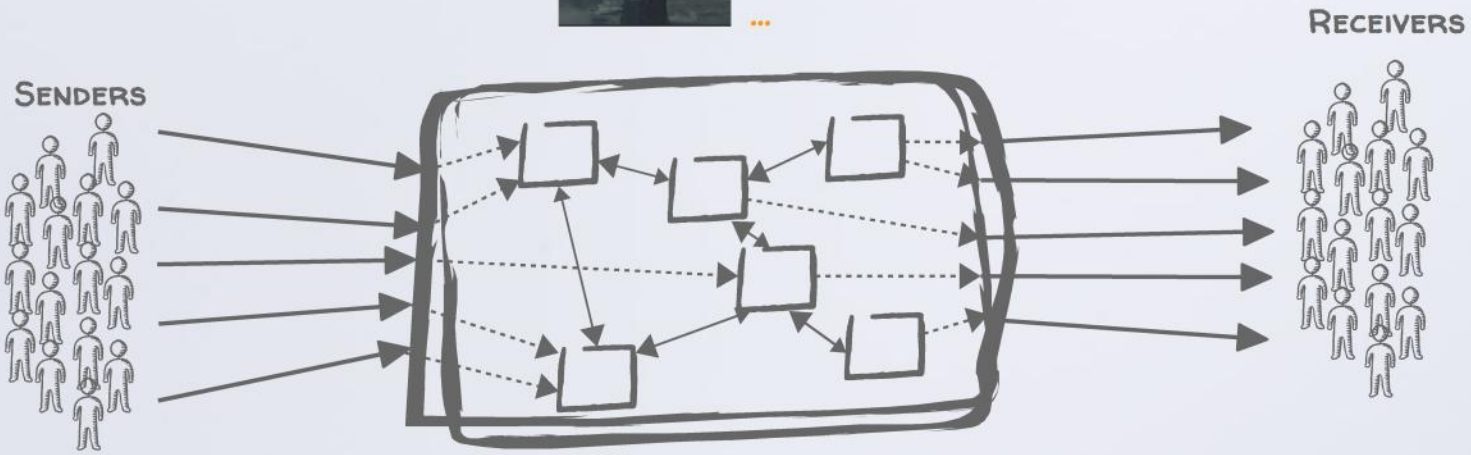

- Bitwise unlinkability
  - Crypto to make inputs and outputs bit patterns different

- (re)packetizing + (re)schedule + (re)routing,
  - Destroy patterns (traffic analysis resistance)
  - Load balancing
  - Distribute trust

In theory should work, but in practice...

IDs
Timing
Volume
Length
...

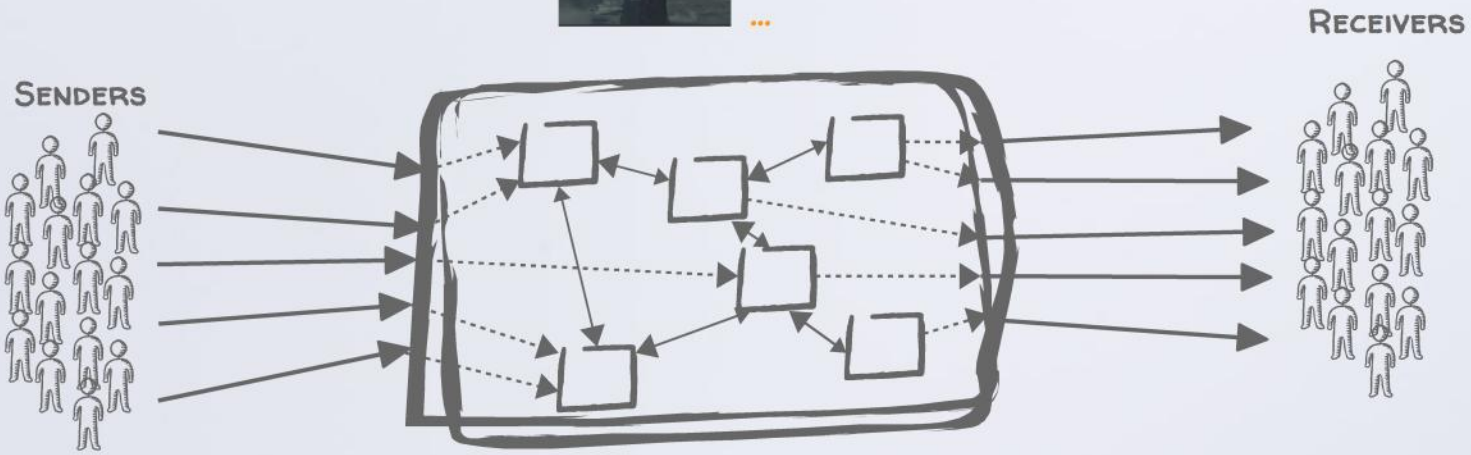

- Bitwise unlinkability
  - Crypto to make inputs and outputs bit patterns different

- (re)packetizing + (re)schedule + (re)routing,
  - Destroy patterns (traffic analysis resistance)
  - Load balancing
  - Distribute trust

Bandwidth
Delay
Churn

Intrinsic network differences
Trust?

... still vulnerable to traffic analysis

Find profiles and communication patterns
persistent relationships show up

Device identification / location
hosts' hardware particular characteristics

Identify users based on choices
not everybody can choose everything

Trace traffic based on patterns
number of packets, delays, ... differ per flow

Recover content
timing and length of packets

Identify traffic based on their patterns (e.g., website fingerprinting)
same traffic always looks similar

Trace packets based on routing algorithms
not all routes are possible

Users' past history
timing correlated to caches

Many, many, many, many, many more....

Pérez-González, Fernando, and Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Danezis, George, and Paul Syverson. "Bridging and fingerprinting: Epistemic attacks on route selection." PETS, 2008.
Houmansadr, Amir, and Nikita Borisov. "The need for flow fingerprints to link correlated network flows." PETS, 2013.
Troncoso, Carmela, and George Danezis. "The bayesian traffic analysis of mix networks."CCS, 2009.
Juarez, Marc, Sadia Afroz, Gunes Acar, Claudia Diaz, and Rachel Greenstadt. "A critical evaluation of website fingerprinting attacks." CCS, 2014.
Felten, Edward W., and Michael A. Schneider. "Timing attacks on web privacy." CCS, 2000.
Murdoch, Steven J. "Hot or not: Revealing hidden services by their clock skew." CCS, 2006.
White, A. M., Matthews, A. R., Snow, K. Z., & Monrose, F. "Phonotactic reconstruction of encrypted VoIP conversations: Hookt on fon-iks." IEEE S&P, 2011.

... STILL VULNERABLE TO TRAFFIC ANALYSIS

Find profiles and communication patterns
persistent relationships show up

Device identification / location
hosts' hardware particular characteristics

Identify users based on choices
not everybody can choose everything

Trace traffic based on patterns
number of packets, delays, ... differ per flow

Recover content
timing and length of packets

Identify traffic based on their patterns (e.g., website fingerprinting)
same traffic always looks similar

Trace packets based on routing algorithms
not all routes are possible

Users' past history
timing correlated to caches

Many, many, many, many, many more....

Pérez-González, Fernando, and Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Danezis, George, and Paul Syverson. "Bridging and fingerprinting: Epistemic attacks on route selection." PETS, 2008.
Houmansadr, Amir, and Nikita Borisov. "The need for flow fingerprints to link correlated network flows." PETS, 2013.
Troncoso, Carmela, and George Danezis. "The bayesian traffic analysis of mix networks."CCS, 2009.
Juarez, Marc, Sadia Afroz, Gunes Acar, Claudia Diaz, and Rachel Greenstadt. "A critical evaluation of website fingerprinting attacks." CCS, 2014.
Felten, Edward W., and Michael A. Schneider. "Timing attacks on web privacy." CCS, 2000.
Murdoch, Steven J. "Hot or not: Revealing hidden services by their clock skew." CCS, 2006.
White, A. M., Matthews, A. R., Snow, K. Z., & Monrose, F. "Phonotactic reconstruction of encrypted VoIP conversations: Hookt on fon-iks." IEEE S&P, 2011.

... STILL VULNERABLE TO TRAFFIC ANALYSIS

Find profiles and communication patterns
persistent relationships show up

Device identification / location
hosts' hardware particular characteristics

Identify users based on choices
not everybody can choose everything

Trace traffic based on patterns
number of packets, delays, ... differ per flow

Identify traffic based on their patterns (e.g., website fingerprinting)
same traffic always looks similar

Recover content
timing and length of packets

Trace packets based on routing algorithms
not all routes are possible

Users' past history
timing correlated to caches

Many, many, many, many, many more....

Pérez-González, Fernando, and Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Danezis, George, and Paul Syverson. "Bridging and fingerprinting: Epistemic attacks on route selection." PETS, 2008.
Houmansadr, Amir, and Nikita Borisov. "The need for flow fingerprints to link correlated network flows." PETS, 2013.
Troncoso, Carmela, and George Danezis. "The bayesian traffic analysis of mix networks."CCS, 2009.
Juarez, Marc, Sadia Afroz, Gunes Acar, Claudia Diaz, and Rachel Greenstadt. "A critical evaluation of website fingerprinting attacks." CCS, 2014.
Felten, Edward W., and Michael A. Schneider. "Timing attacks on web privacy." CCS, 2000.
Murdoch, Steven J. "Hot or not: Revealing hidden services by their clock skew." CCS, 2006.
White, A. M., Matthews, A. R., Snow, K. Z., & Monrose, F. "Phonotactic reconstruction of encrypted VoIP conversations: Hookt on fon-iks." IEEE S&P, 2011.

Where do messages go?

- **Threshold mix**: collects t messages, and outputs them changing their appearance and in a random order

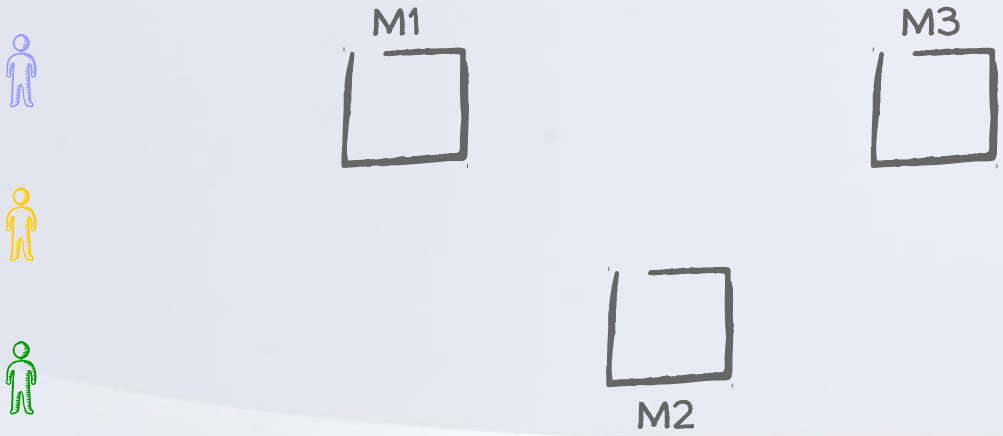

Where do messages go?

- **Threshold mix**: collects t messages, and outputs them changing their appearance and in a random order

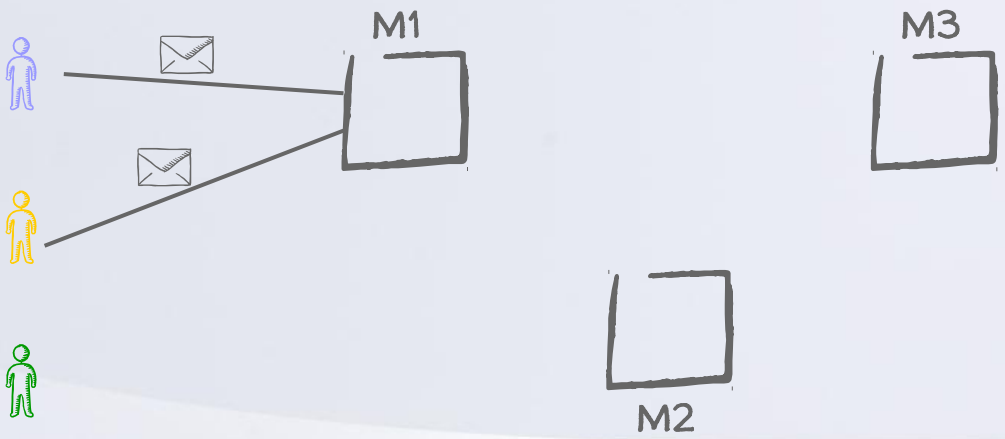

Where do messages go?

- **Threshold mix**: collects t messages, and outputs them changing their appearance and in a random order

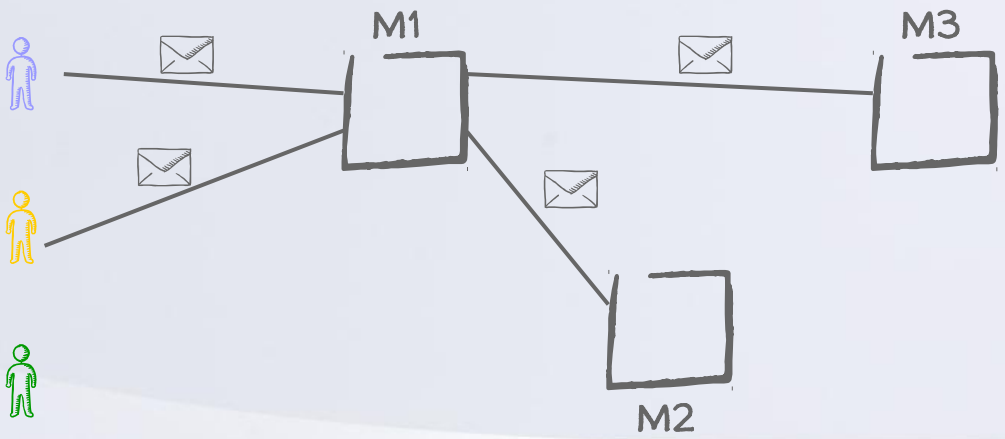

Where do messages go?

- **Threshold mix**: collects t messages, and outputs them changing their appearance and in a random order

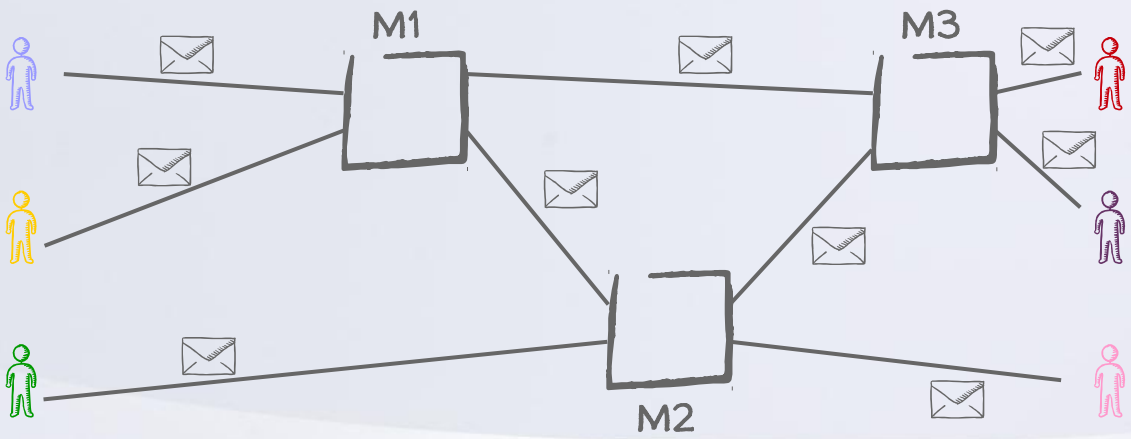

Where do messages go?

- **Threshold mix**: collects t messages, and outputs them changing their appearance and in a random order

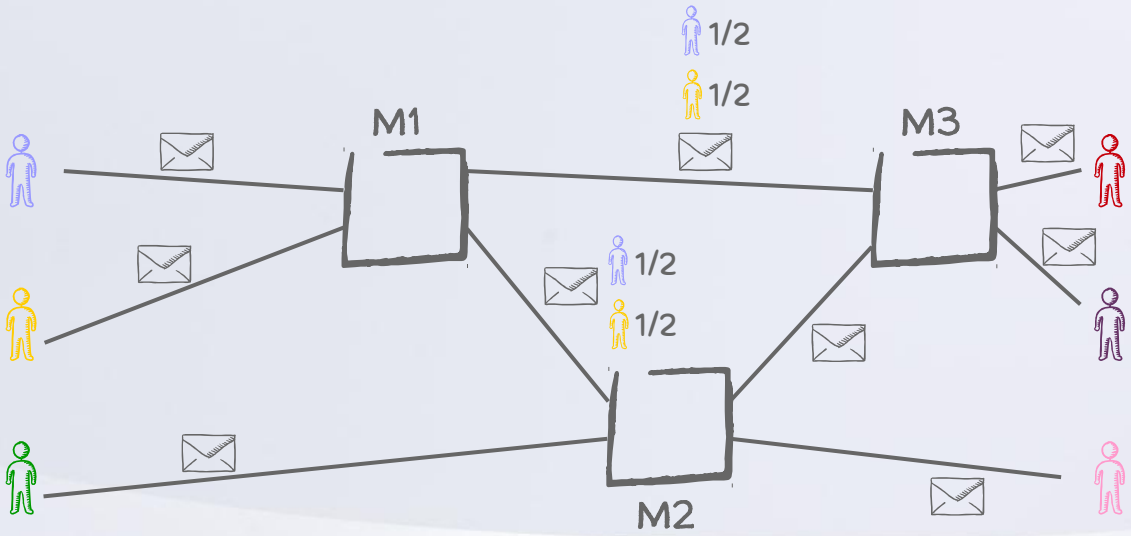

Where do messages go?

- **Threshold mix:** collects t messages, and outputs them changing their appearance and in a random order

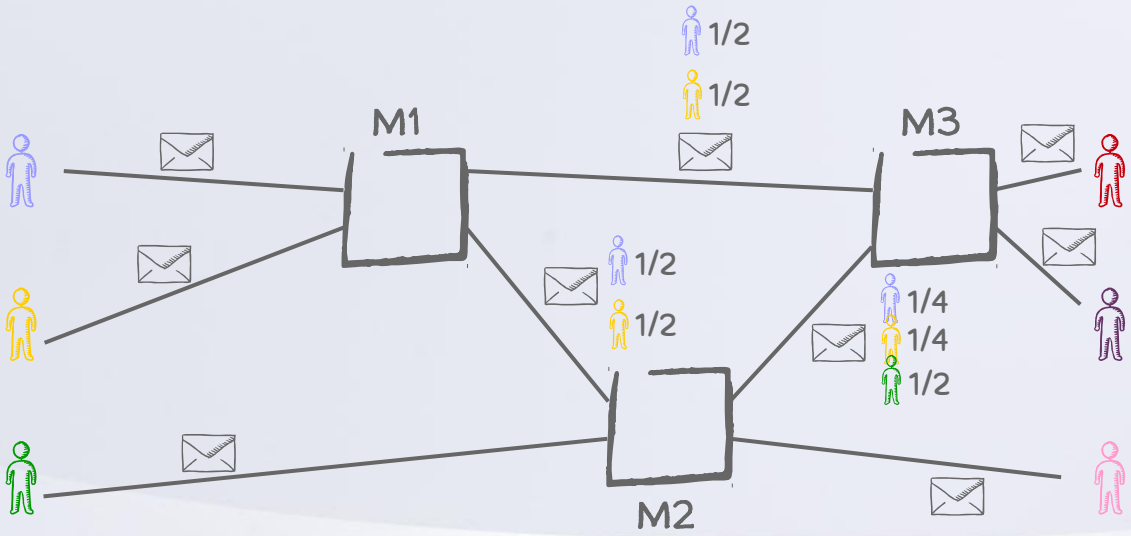

Where do messages go?

- **Threshold mix:** collects t messages, and outputs them changing their appearance and in a random order

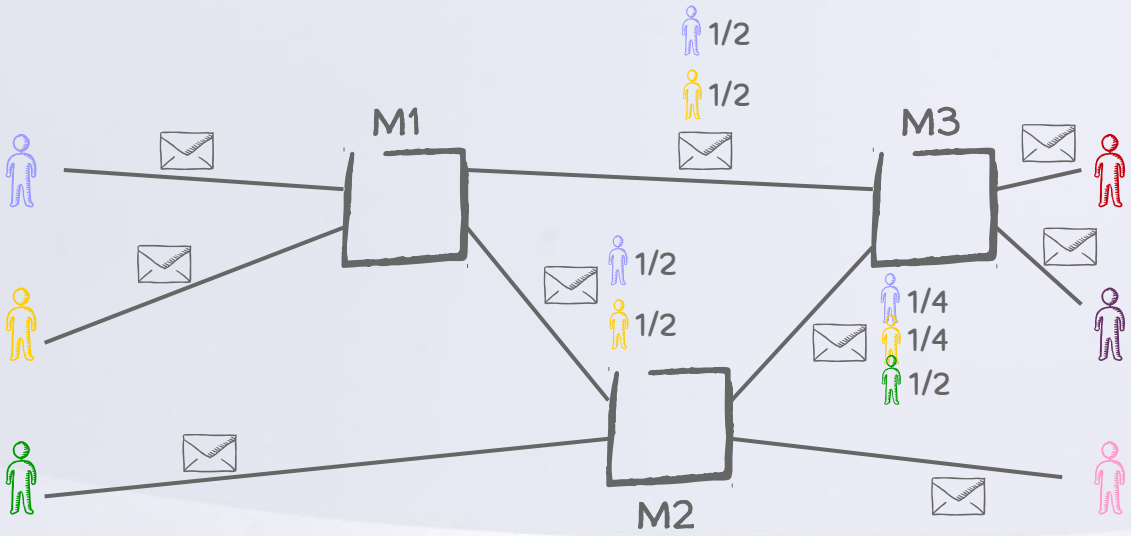

3/8	3/8	1/4
3/8	3/8	1/4
1/4	1/4	1/2

Where do messages go?

not everything is possible (e.g., max 2 hops)

- **Threshold mix:** collects t messages, and outputs them changing their appearance and in a random order

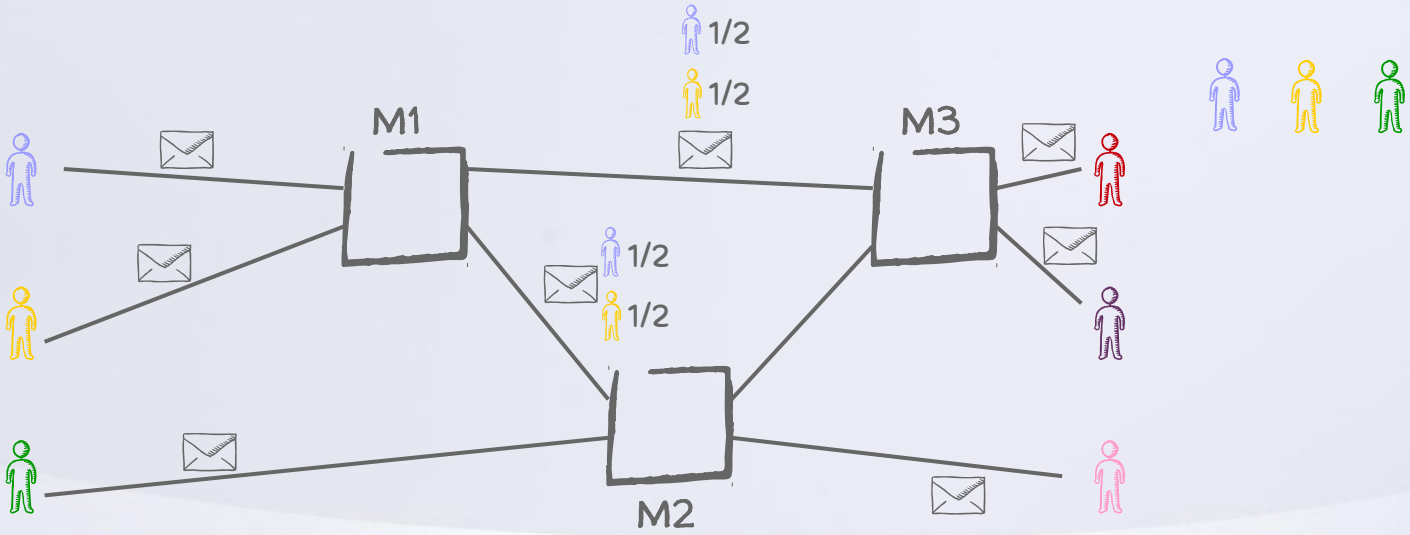

Danezis, George. "Mix-Networks with Restricted Routes". PETS 2003

Where do messages go?

not everything is possible (e.g., max 2 hops)

- **Threshold mix:** collects t messages, and outputs them changing their appearance and in a random order

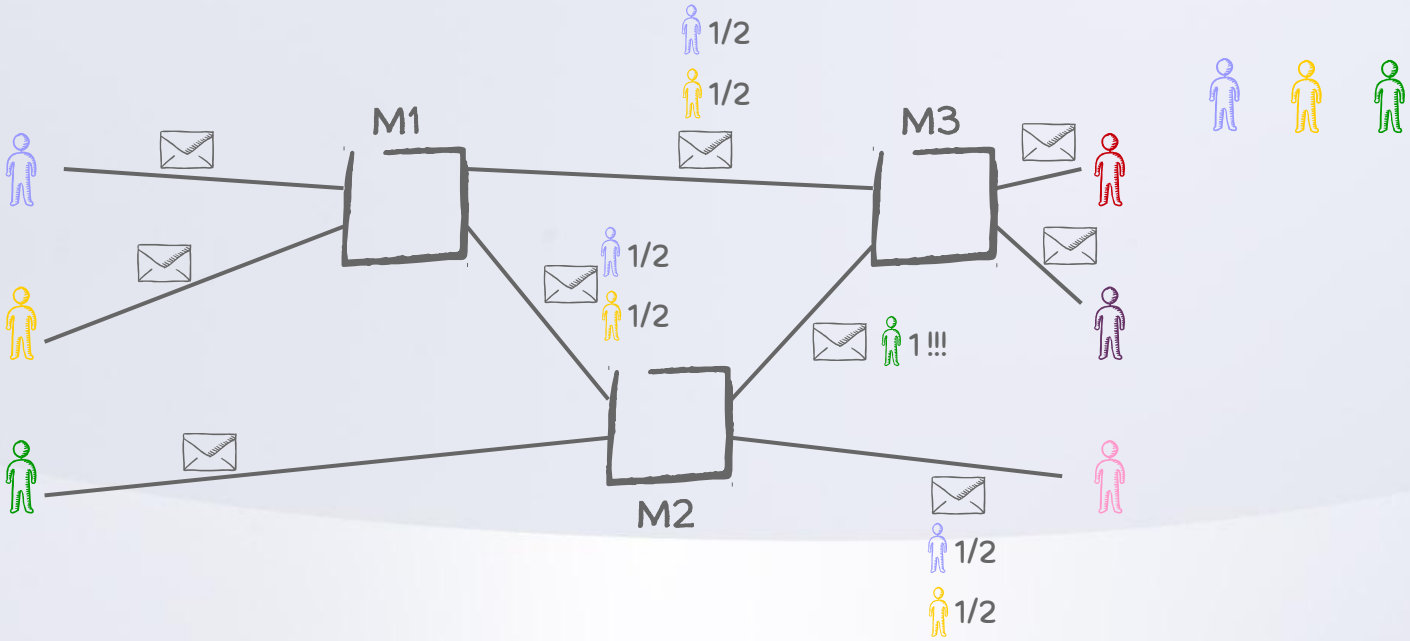

Danezis, George. "Mix-Networks with Restricted Routes". PETS 2003

Where do messages go?

not everything is possible (e.g., max 2 hops)

- **Threshold mix:** collects t messages, and outputs them changing their appearance and in a random order

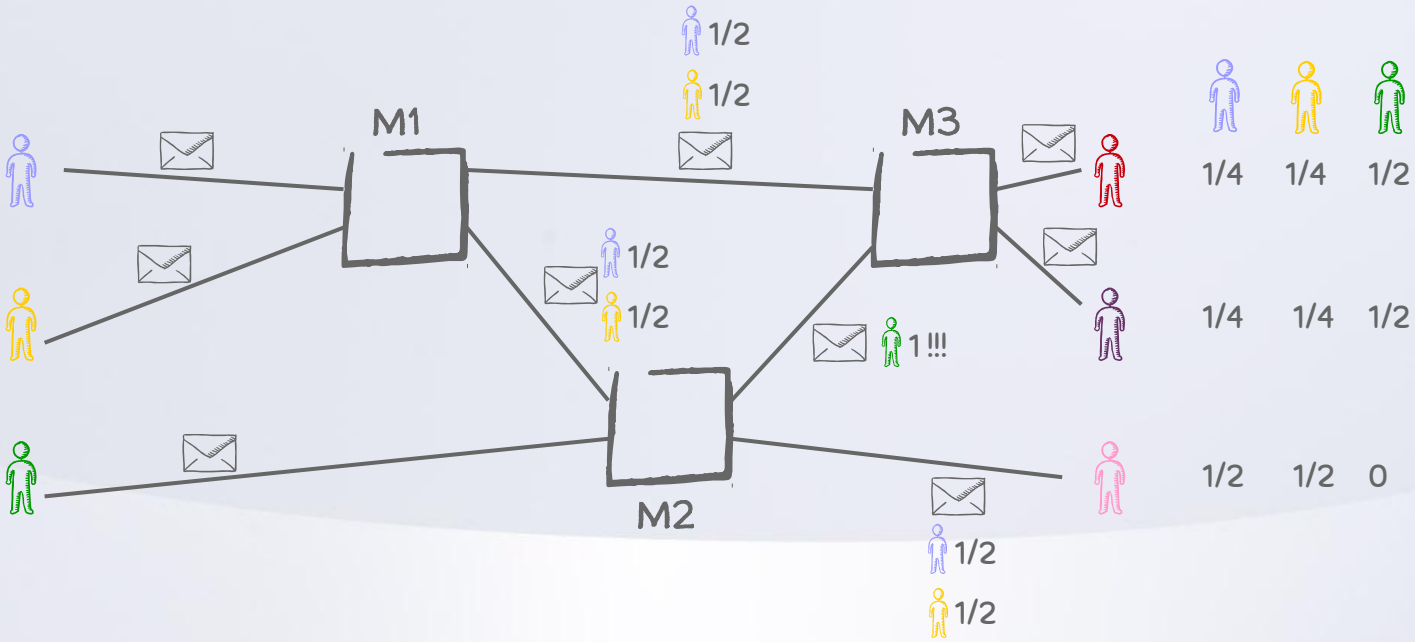

Danezis, George. "Mix-Networks with Restricted Routes". PETS 2003

Where do messages go?

not everything is possible (e.g., does not know M2)

- **Threshold mix:** collects t messages, and outputs them changing their appearance and in a random order

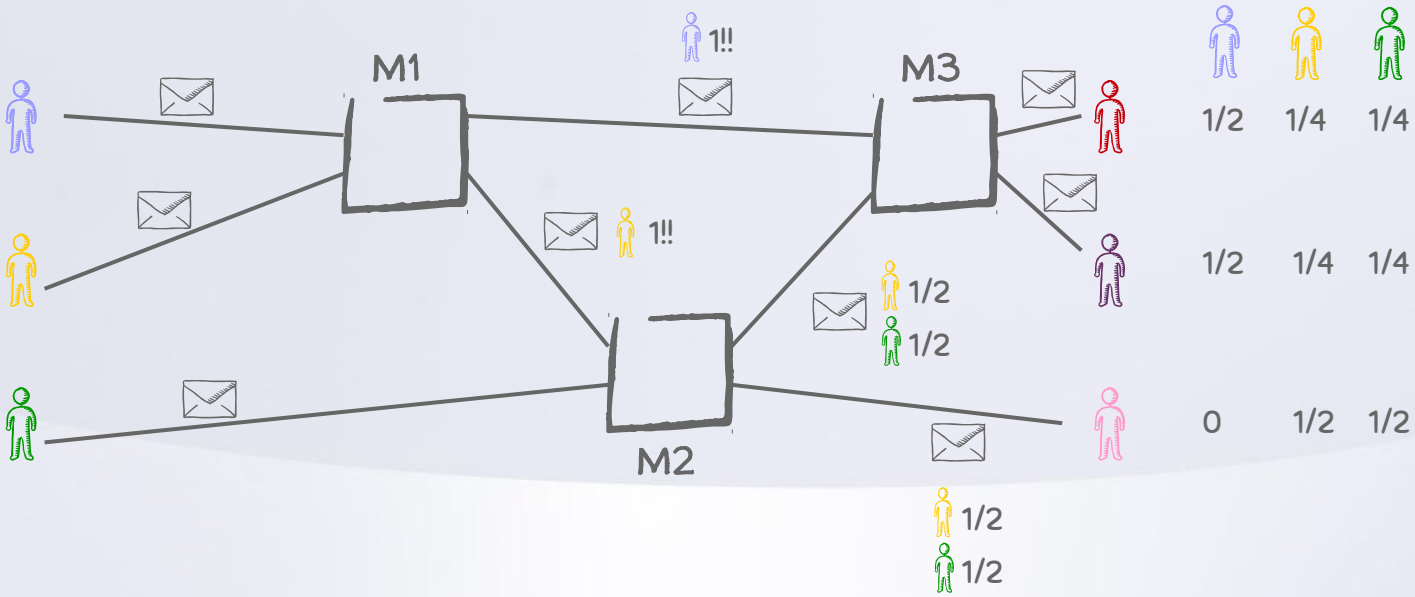

Danezis, George, and Paul Syverson. "Bridging and fingerprinting: Epistemic attacks on route selection." PETS, 2008.

Where do messages go?

not everything is possible (e.g., does not know M2)

Threshold mix: collects t messages, and outputs them changing their appearance and in a random order

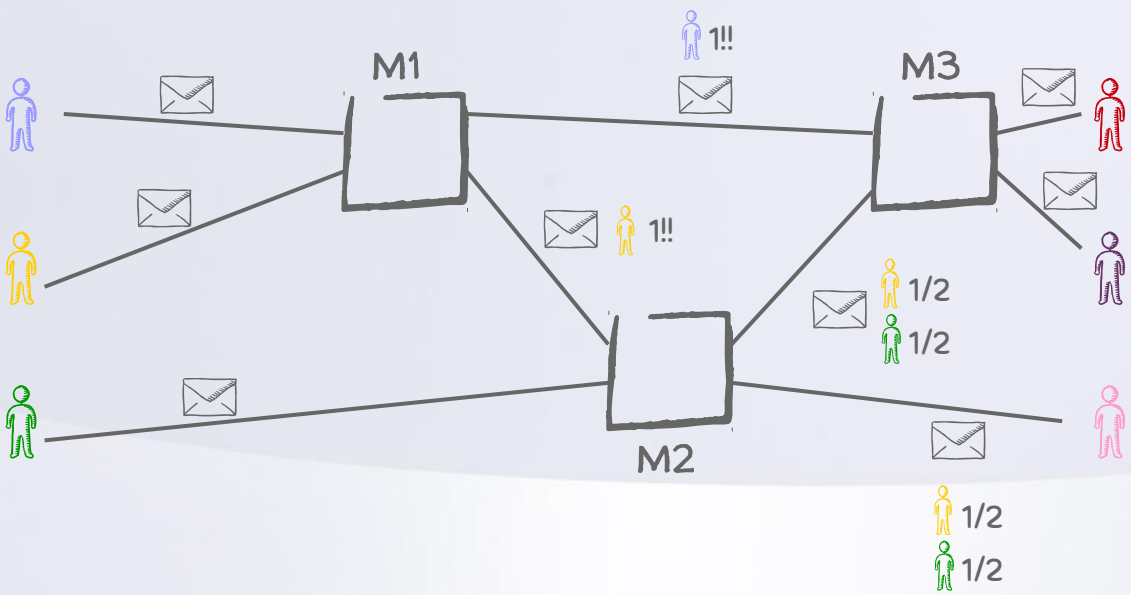

1/2	1/4	1/4
1/2	1/4	1/4
0	1/2	1/2

Non trivial given observation!!

Danezis, George, and Paul Syverson. "Bridging and fingerprinting: Epistemic attacks on route selection." PETS, 2008.

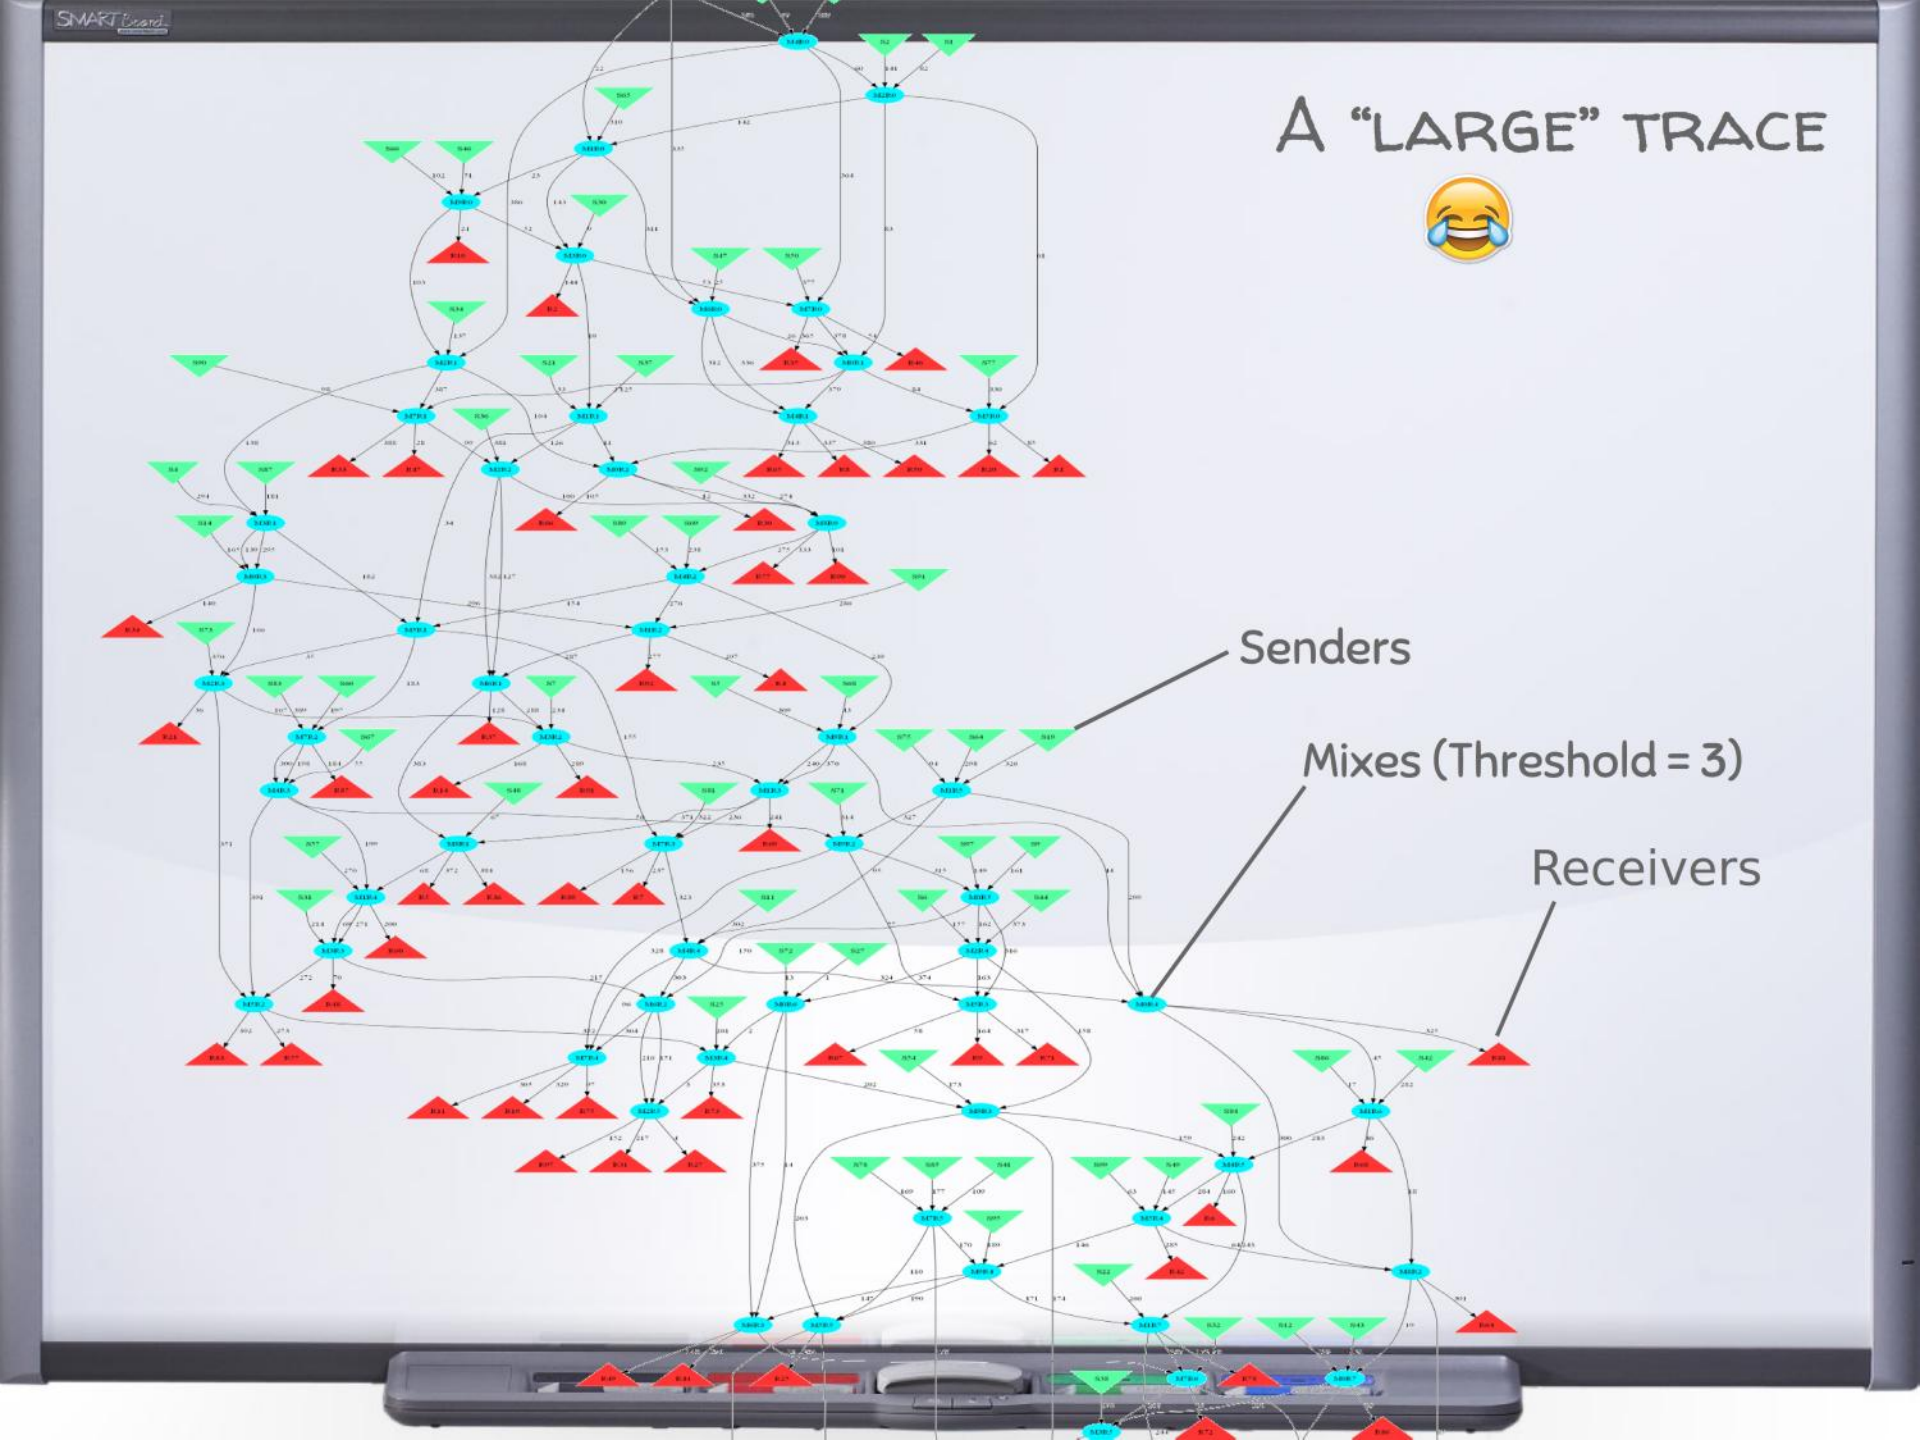
A "LARGE" TRACE
Senders
Mixes (Threshold = 3)
Receivers

Redefining the problem

Given what we see (Observation) and the system operation (Constraints)
Probability of mixes "Hidden State"?
(or Probability of each possible path?)

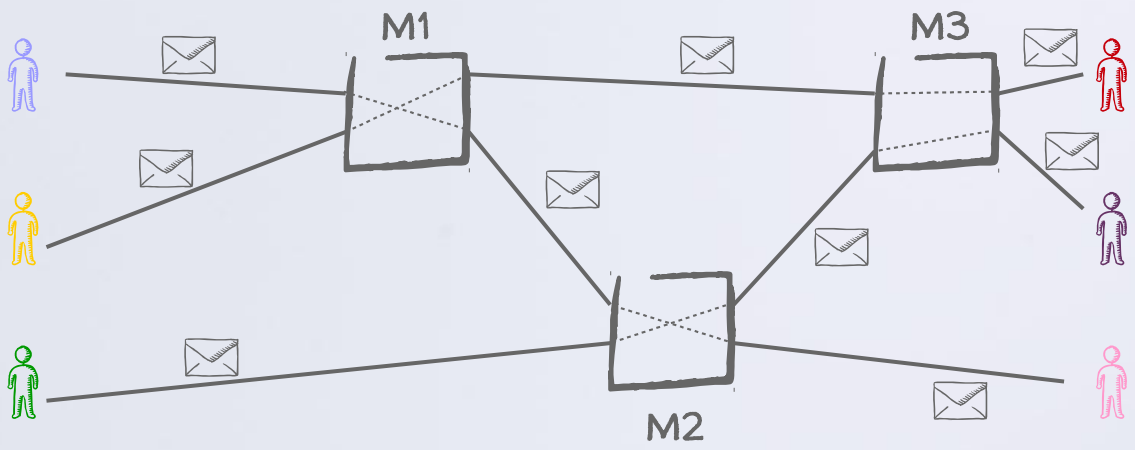

Redefining the problem

Given what we see (Observation) and the system operation (Constraints)
Probability of mixes "Hidden State"?
(or Probability of each possible path?)

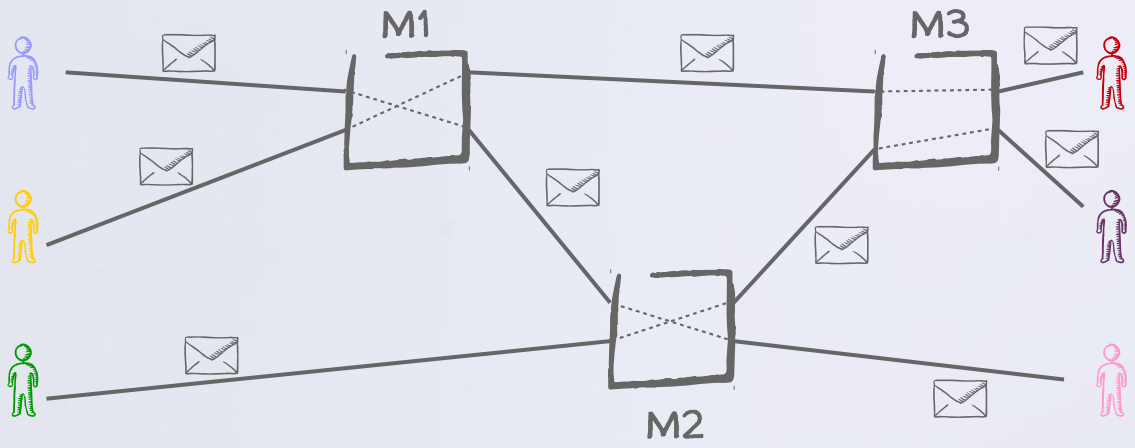

$$Pr[HS|O,C] = \frac{Pr[O|HS,C] \cdot Pr[HS|C]}{\sum_{HS} Pr[HS,O|C]}$$

Redefining the problem

Given what we see (Observation) and the system operation (Constraints)
Probability of mixes "Hidden State"?
(or Probability of each possible path?)

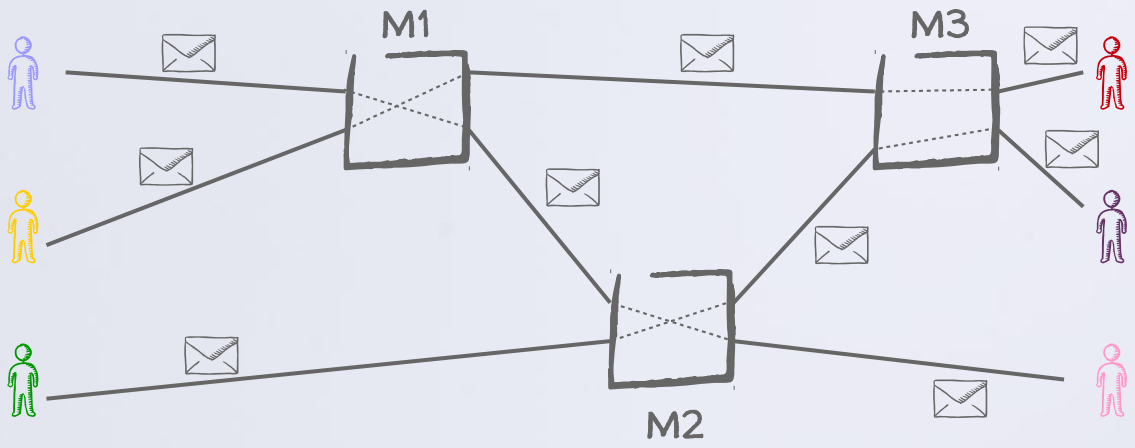

$$Pr[HS|O,C] = \frac{Pr[O|HS,C] \cdot Pr[HS|C]}{\sum_{HS} Pr[HS,O|C]}$$

Redefining the problem

Given what we see (Observation) and the system operation (Constraints)
Probability of mixes "Hidden State"?
(or Probability of each possible path?)

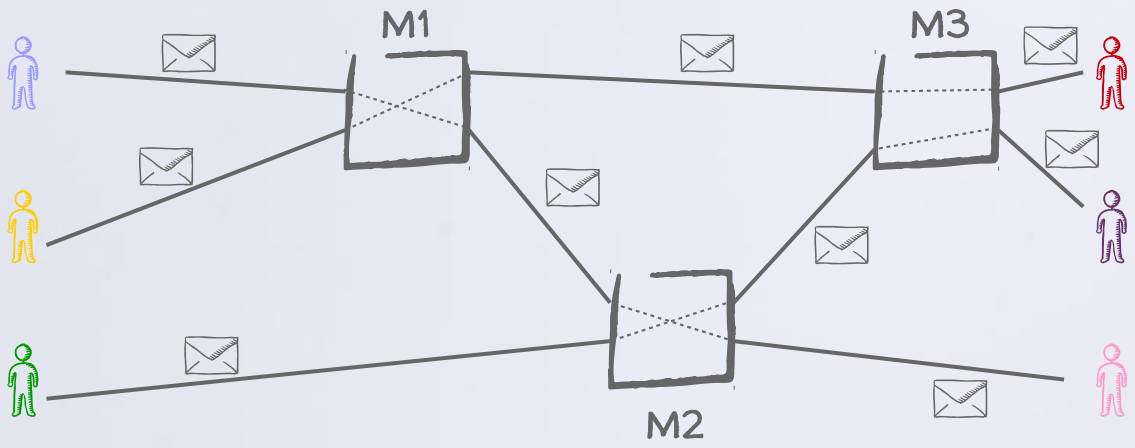

$$Pr[HS|O,C] = \frac{Pr[O|HS,C]\cdot Pr[HS|C]}{\sum_{HS} Pr[HS,O|C]} = \frac{Pr[O|HS,C]\cdot K}{Z} = \frac{Pr[Paths|C]\cdot K}{Z}$$

Redefining the problem

Given what we see (Observation) and the system operation (Constraints)
Probability of mixes "Hidden State"?
(or Probability of each possible path?)

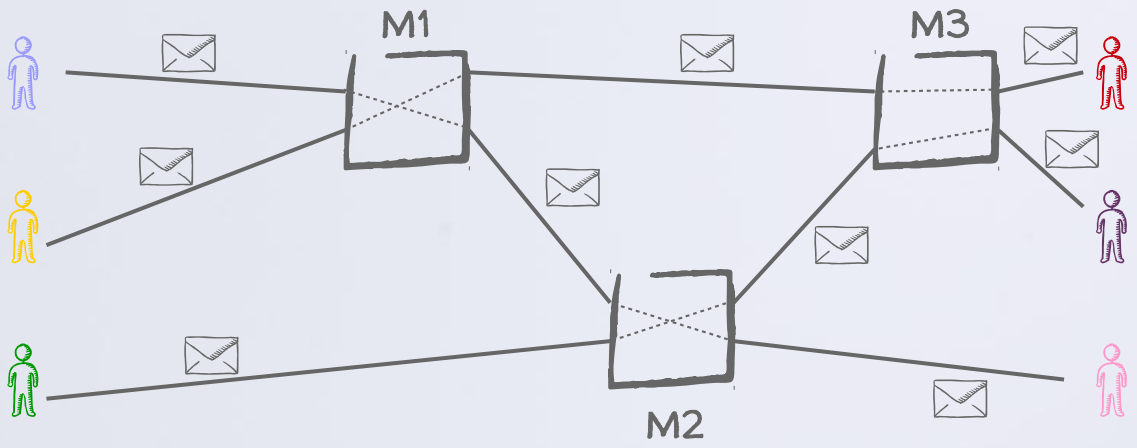

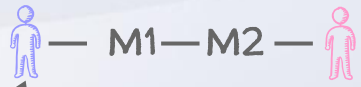

$$Pr[HS|O,C] = \frac{Pr[O|HS,C]\cdot Pr[HS|C]}{\sum_{HS} Pr[HS,O|C]} = \frac{Pr[O|HS,C]\cdot K}{Z} = \frac{Pr[Paths|C]\cdot K}{Z}$$

Redefining the problem

Given what we see (Observation) and the system operation (Constraints)
Probability of mixes "Hidden State"?
(or Probability of each possible path?)

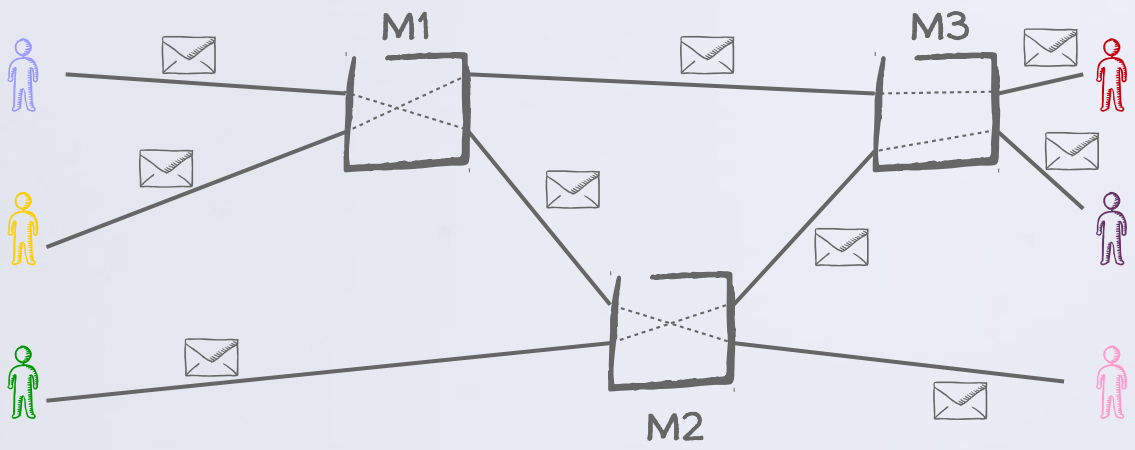

Software!! we can compute :)

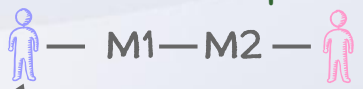

$$Pr[HS|O,C] = \frac{Pr[O|HS,C]\cdot Pr[HS|C]}{\sum_{HS} Pr[HS,O|C]} = \frac{Pr[O|HS,C]\cdot K}{Z} = \frac{Pr[Paths|C]\cdot K}{Z}$$

We usually care about marginal probabilities, not all (Pr[→ |O,C]) ← SAMPLING!!

Troncoso, Carmela, and George Danezis. "The bayesian traffic analysis of mix networks."CCS, 2009.

Takeaways attacks on routes

- Traffic analysis is non trivial when there are constraints
- Traffic analysis as inference problem: systematic!
  - Probabilistic model: can incorporate most attacks
    - Can integrate knowledge on path probability computation
      - More constraints → less anonymity but more complexity
    - Combines well with other inferences: e.g., long-term attacks (in a minute)
- Sampling methods to extract marginal probabilities

Finding persistent communications

Disclosure Attacks

In reality...

- Alice has few friends with whom she communicates often
- Alice is not always online (at least not active)

Can Sauron learn Alice's friends?

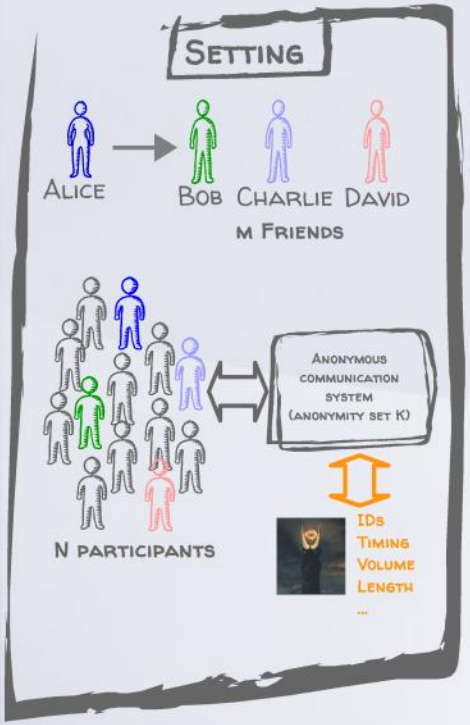

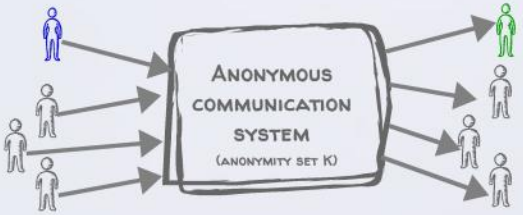

Finding persistent communications

Disclosure Attacks

In reality...

- Alice has few friends with whom she communicates often
- Alice is not always online (at least not active)

Can Sauron learn Alice's friends?

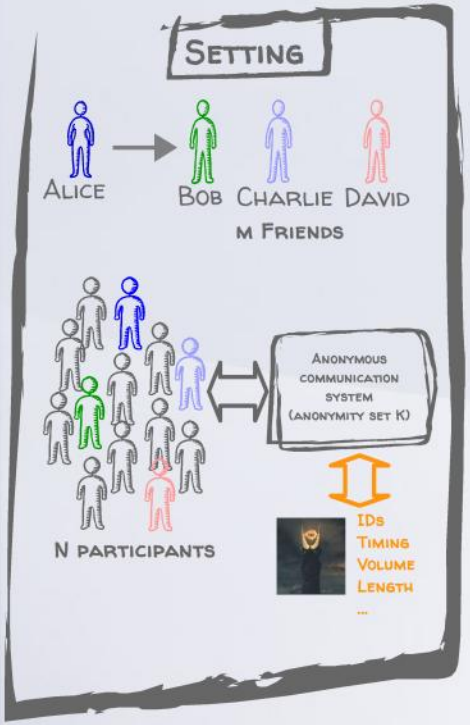

Finding persistent communications
Disclosure Attacks

In reality...

- Alice has few friends with whom she communicates often
- Alice is not always online (at least not active)

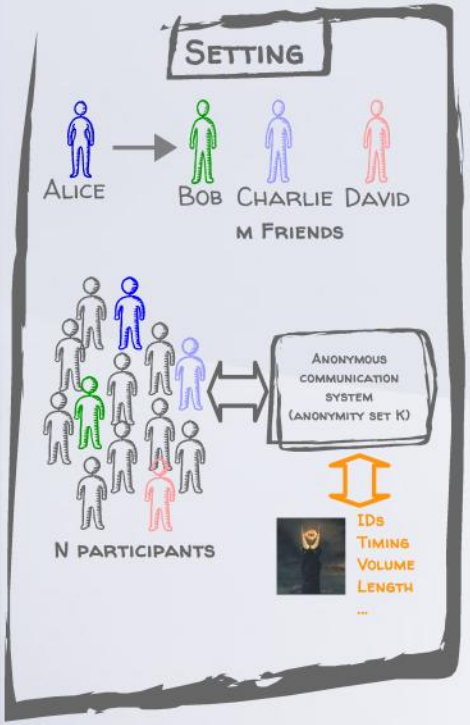

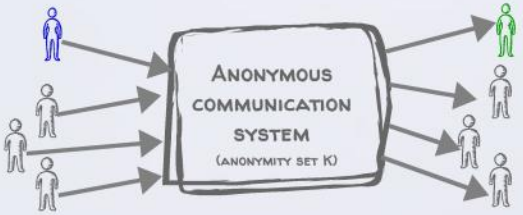

1-  sees Alice sending a single message to the system

2- Anonymity set size = K

3- Perfect! 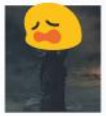

As time goes by and Alice sends more messages...

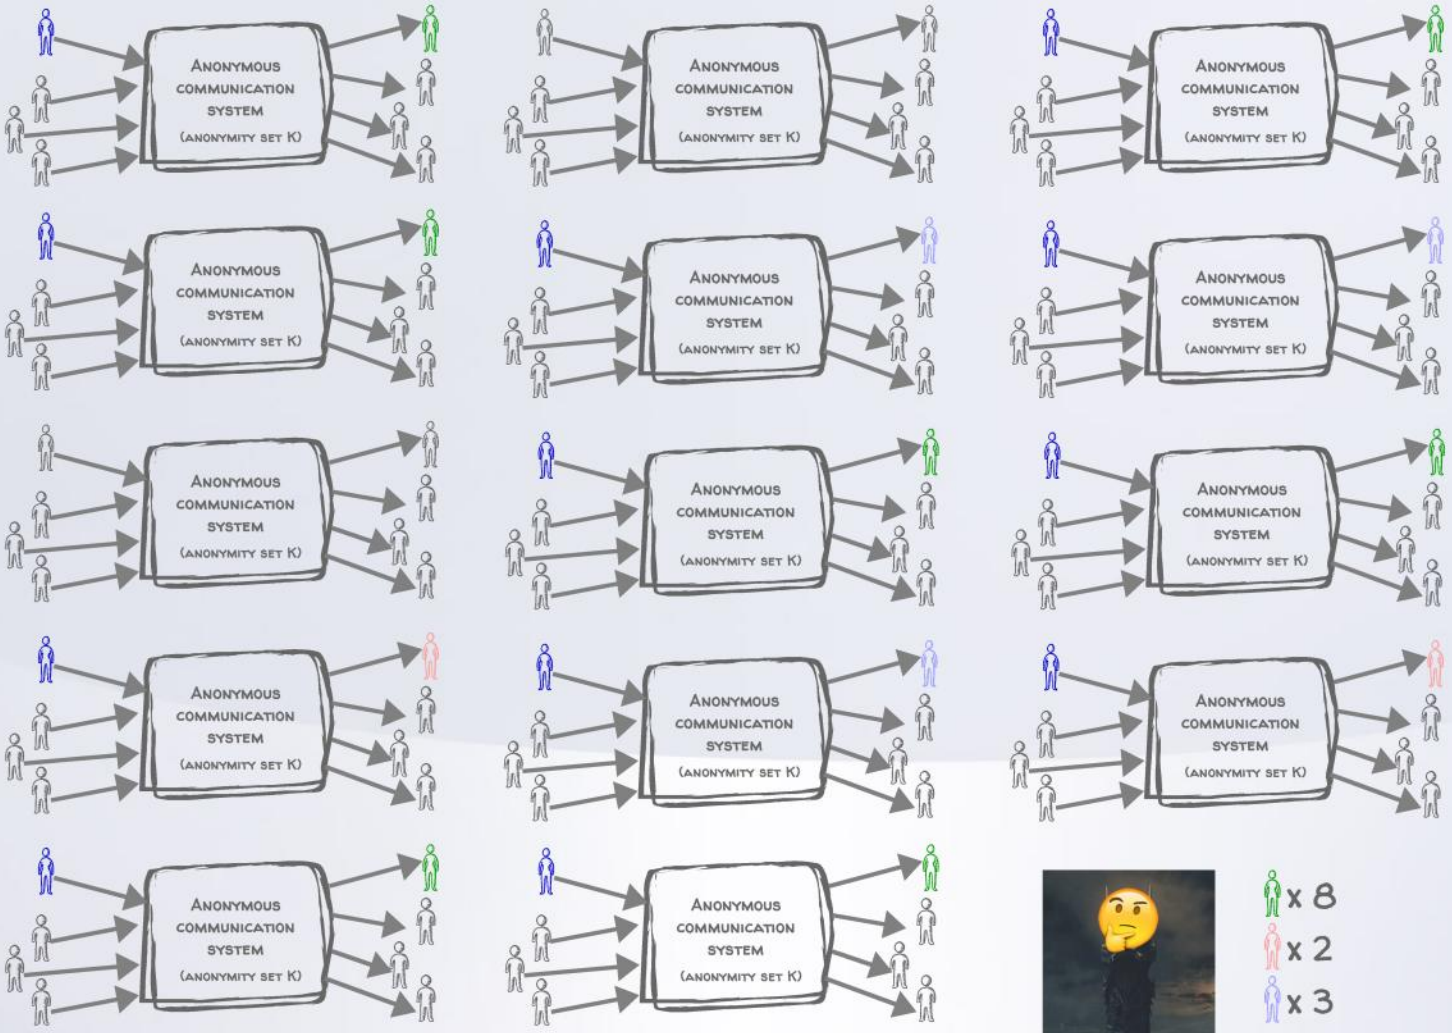

Let's "do" the math

Approach 1: Statistical Disclosure Attack

- Alice's friends will be in the sets more often than random receivers. How often?
  Expected number of messages per receiver after t rounds:
  - $\mu_{other} = (1 / N) \cdot (K-1) \cdot t$
  - $\mu_{Alice} = (1 / M) \cdot t + \mu_{other}$

- Just count the number of messages per receiver when Alice is sending!
  - $\mu_{Alice} > \mu_{other}$

Danezis, George. "Statistical disclosure attacks." Security and Privacy in the Age of Uncertainty, 2003.
Danezis, George, Claudia Diaz, and Carmela Troncoso. "Two-sided statistical disclosure attack." PETS, 2007.
Mathewson, Nick, and Roger Dingledine. "Practical traffic analysis: Extending and resisting statistical disclosure." PETS, 2004
Troncoso, Carmela, Benedikt Gierlichs, Bart Preneel, and Ingrid Verbauwhede. "Perfect matching disclosure attacks." PETS, 2008

Let's "do" the math

Approach 1: Statistical Disclosure Attack

- Alice's friends will be in the sets more often than random receivers. How often? Expected number of messages per receiver after t rounds:
  - $\mu_{other} = (1 / N) \cdot (K-1) \cdot t$
  - $\mu_{Alice} = (1 / M) \cdot t + \mu_{other}$
- Just count the number of messages per receiver when Alice is sending!
  - $\mu_{Alice} > \mu_{other}$

N=20 M=3 K=5 T=45
Alice's Friends={[0, 13, 19]}

Round	Receivers	SDA
1	**[15, 13, 14, 5, 9]**	**[13, 14, 15]**
2	[19, 10, 17, 13, 8]	[13, 17, 19]
3	[0, 7, 0, 13, 5]	[0, 5, 13]
4	[16, 18, 6, 13, 10]	[5, 10, 13]
5	[1, 17, 1, 13, 6]	[10, 13, 17]
6	[18, 15, 17, 13, 17]	[13, 17, 18]
7	[0, 13, 11, 8, 4]	[0, 13, 17]
8	[15, 18, 0, 8, 12]	[0, 13, 17]
9	[15, 18, 15, 19, 14]	[13, 15, 18]
10	[0, 12, 4, 2, 8]	[0, 13, 15]
11	[9, 13, 14, 19, 15]	[0, 13, 15]
12	[13, 6, 2, 16, 0]	[0, 13, 15]
13	[1, 0, 3, 5, 1]	[0, 13, 15]
14	[17, 10, 14, 11, 19]	[0, 13, 15]
15	[12, 14, 17, 13, 0]	[0, 13, 17]
16	[18, 19, 19, 8, 11]	[0, 13, 19]
17	[4, 1, 19, 0, 19]	[0, 13, 19]
18	[0, 6, 1, 18, 3]	[0, 13, 19]
19	[5, 1, 14, 0, 5]	[0, 13, 19]
20	[17, 18, 2, 4, 13]	[0, 13, 19]
21	[8, 10, 1, 18, 13]	[0, 13, 19]
22	[14, 4, 13, 12, 4]	[0, 13, 19]
23	[19, 13, 3, 17, 12]	[0, 13, 19]
24	[8, 18, 0, 10, 18]	[0, 13, 18]

Danezis, George. "Statistical disclosure attacks." Security and Privacy in the Age of Uncertainty, 2003.
Danezis, George, Claudia Diaz, and Carmela Troncoso. "Two-sided statistical disclosure attack." PETS, 2007.
Mathewson, Nick, and Roger Dingledine. "Practical traffic analysis: Extending and resisting statistical disclosure." PETS, 2004
Troncoso, Carmela, Benedikt Gierlichs, Bart Preneel, and Ingrid Verbauwhede. "Perfect matching disclosure attacks." PETS, 2008

Let's "do" the math

Approach 1: Statistical Disclosure Attack

- Alice's friends will be in the sets more often than random receivers. How often? Expected number of messages per receiver after t rounds:
  - $\mu_{other} = (1 / N) \cdot (K-1) \cdot t$
  - $\mu_{Alice} = (1 / M) \cdot t + \mu_{other}$
- Just count the number of messages per receiver when Alice is sending!
  - $\mu_{Alice} > \mu_{other}$

N=20 M=3 K=5 t=45
Alice's Friends={[0, 13, 19]}

Round	Receivers	SDA
1	**[15, 13, 14, 5, 9]**	**[13, 14, 15]**
2	[19, 10, 17, 13, 8]	[13, 17, 19]
3	[0, 7, 0, 13, 5]	[0, 5, 13]
4	[16, 18, 6, 13, 10]	[5, 10, 13]
5	[1, 17, 1, 13, 6]	[10, 13, 17]
6	[18, 15, 17, 13, 17]	[13, 17, 18]
7	[0, 13, 11, 8, 4]	[0, 13, 17]
8	[15, 18, 0, 8, 12]	[0, 13, 17]
9	[15, 18, 15, 19, 14]	[13, 15, 18]
10	[0, 12, 4, 2, 8]	[0, 13, 15]
11	[9, 13, 14, 19, 15]	[0, 13, 15]
12	[13, 6, 2, 16, 0]	[0, 13, 15]
13	[1, 0, 3, 5, 1]	[0, 13, 15]
14	[17, 10, 14, 11, 19]	[0, 13, 15]
15	[12, 14, 17, 13, 0]	[0, 13, 17]
16	**[18, 19, 19, 8, 11]**	**[0, 13, 19]**
17	[4, 1, 19, 0, 19]	[0, 13, 19]
18	[0, 6, 1, 18, 3]	[0, 13, 19]
19	[5, 1, 14, 0, 5]	[0, 13, 19]
20	[17, 18, 2, 4, 13]	[0, 13, 19]
21	[8, 10, 1, 18, 13]	[0, 13, 19]
22	[14, 4, 13, 12, 4]	[0, 13, 19]
23	[19, 13, 3, 17, 12]	[0, 13, 19]
24	[8, 18, 0, 10, 18]	[0, 13, 18]

Danezis, George. "Statistical disclosure attacks." Security and Privacy in the Age of Uncertainty, 2003.
Danezis, George, Claudia Diaz, and Carmela Troncoso. "Two-sided statistical disclosure attack." PETS, 2007.
Mathewson, Nick, and Roger Dingledine. "Practical traffic analysis: Extending and resisting statistical disclosure." PETS, 2004
Troncoso, Carmela, Benedikt Gierlichs, Bart Preneel, and Ingrid Verbauwhede. "Perfect matching disclosure attacks." PETS, 2008

Let's "do" the math

Approach 1: Statistical Disclosure Attack

- Alice's friends will be in the sets more often than random receivers. How often? Expected number of messages per receiver after t rounds:
  - $\mu_{other} = (1 / N) \cdot (K-1) \cdot t$
  - $\mu_{Alice} = (1 / M) \cdot t + \mu_{other}$
- Just count the number of messages per receiver when Alice is sending!
  - $\mu_{Alice} > \mu_{other}$

N=20 M=3 K=5 T=45
Alice's Friends={[0, 13, 19]}

Round	Receivers	SDA
1	**[15, 13, 14, 5, 9]**	**[13, 14, 15]**
2	[19, 10, 17, 13, 8]	[13, 17, 19]
3	[0, 7, 0, 13, 5]	[0, 5, 13]
4	[16, 18, 6, 13, 10]	[5, 10, 13]
5	[1, 17, 1, 13, 6]	[10, 13, 17]
6	[18, 15, 17, 13, 17]	[13, 17, 18]
7	[0, 13, 11, 8, 4]	[0, 13, 17]
8	[15, 18, 0, 8, 12]	[0, 13, 17]
9	[15, 18, 15, 19, 14]	[13, 15, 18]
10	[0, 12, 4, 2, 8]	[0, 13, 15]
11	[9, 13, 14, 19, 15]	[0, 13, 15]
12	[13, 6, 2, 16, 0]	[0, 13, 15]
13	[1, 0, 3, 5, 1]	[0, 13, 15]
14	[17, 10, 14, 11, 19]	[0, 13, 15]
15	[12, 14, 17, 13, 0]	[0, 13, 17]
16	**[18, 19, 19, 8, 11]**	**[0, 13, 19]**
17	[4, 1, 19, 0, 19]	[0, 13, 19]
18	[0, 6, 1, 18, 3]	[0, 13, 19]
19	[5, 1, 14, 0, 5]	[0, 13, 19]
20	[17, 18, 2, 4, 13]	[0, 13, 19]
21	[8, 10, 1, 18, 13]	[0, 13, 19]
22	[14, 4, 13, 12, 4]	[0, 13, 19]
23	[19, 13, 3, 17, 12]	[0, 13, 19]
24	[8, 18, 0, 10, 18]	[0, 13, 18]

Danezis, George. "Statistical disclosure attacks." Security and Privacy in the Age of Uncertainty, 2003.
Danezis, George, Claudia Diaz, and Carmela Troncoso. "Two-sided statistical disclosure attack." PETS, 2007.
Mathewson, Nick, and Roger Dingledine. "Practical traffic analysis: Extending and resisting statistical disclosure." PETS, 2004
Troncoso, Carmela, Benedikt Gierlichs, Bart Preneel, and Ingrid Verbauwhede. "Perfect matching disclosure attacks." PETS, 2008

Let's "do" the math

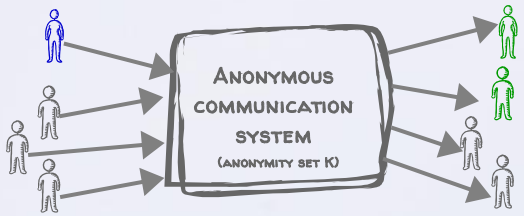

$P_{i,j}$ = probability that i sends a message to j
x^r = vector of n# of messages sent round r (x^r_i =1)
y^r = vector of n# of messages received round r (y^r_j = 2)
$H = [x^1, x^2, x^3, \dots,]$

Approach 2: Least Squares Disclosure Attack

- Maximum likelihood approach: solve a Least Squares minimizing mean squared error between real and estimated profiles

Pérez-González, Fernando, and Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Oya, Simon, Carmela Troncoso, and Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014
Perez-Gonzalez, Fernando, Carmela Troncoso, and Simon Oya. "A least squares approach to the static traffic analysis of high-latency anonymous communication systems." TIFS 2014

Let's "do" the math

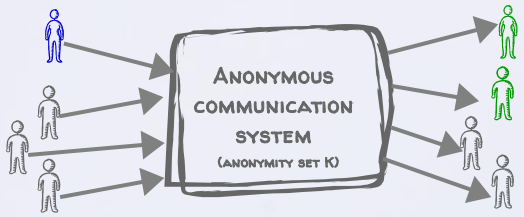

$P_{i,j}$ = probability that i sends a message to j
x^r = vector of n# of messages sent round r (x^r_i =1)
y^r = vector of n# of messages received round r (y^r_j = 2)
H = [x^1,x^2,x^3, … ,]

Approach 2: Least Squares Disclosure Attack

- Maximum likelihood approach: solve a Least Squares minimizing mean squared error between real and estimated profiles

$$\hat{p} = \arg\min_{p} \|y - Hp\| \quad \text{s.t.} \quad p_{i,j} \leqslant 1, \ \sum_i p_{i,j} = 1 \quad \Rightarrow \quad \hat{p} = (H^T H)^{-1} H^T y$$

-

Pérez-González, Fernando, and Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Oya, Simon, Carmela Troncoso, and Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014
Perez-Gonzalez, Fernando, Carmela Troncoso, and Simon Oya. "A least squares approach to the static traffic analysis of high-latency anonymous communication systems." TIFS 2014

Let's "do" the math

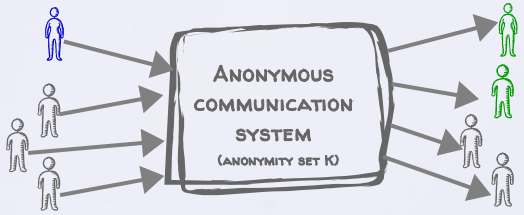

$P_{i,j}$ = probability that i sends a message to j
x^r = vector of n# of messages sent round r (x^r_i =1)
y^r = vector of n# of messages received round r (y^r_j = 2)
H = [x^1,x^2,x^3, … ,]

Approach 2: Least Squares Disclosure Attack

- Maximum likelihood approach: solve a Least Squares minimizing mean squared error between real and estimated profiles

$$\hat{p} = \arg\min_{p} \|y - Hp\| \quad \text{s.t.} \quad p_{i,j} \leqslant 1, \ \sum_i p_{i,j} = 1 \quad \Rightarrow \quad \hat{p} = (H^T H)^{-1} H^T y$$

- Analytical expressions that describe the evolution of the profiling error

$$MSE = \|p - \hat{p}\|^2 = \frac{1}{t}\left(N - 1 + \frac{1}{k}\right)\left(N - \sum_j \frac{f_j^2}{f^2 N}\right)$$

Users
rounds
Batch size
Senders that send a lot
Receivers receive from many

Pérez-González, Fernando, and Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Oya, Simon, Carmela Troncoso, and Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014
Perez-Gonzalez, Fernando, Carmela Troncoso, and Simon Oya. "A least squares approach to the static traffic analysis of high-latency anonymous communication systems." TIFS 2014

Let's "do" the math

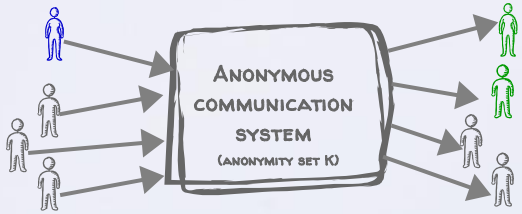

$P_{i,j}$ = probability that i sends a message to j
x^r = vector of n# of messages sent round r (x^r_i =1)
y^r = vector of n# of messages received round r (y^r_j = 2)
H = [x^1,x^2,x^3, … ,]

Approach 2: Least Squares Disclosure Attack

- Maximum likelihood approach: solve a Least Squares minimizing mean squared error between real and estimated profiles

$$\hat{p} = \arg\min_{p} \|y - Hp\| \quad \text{s.t.} \quad p_{i,j} \leqslant 1, \ \sum_i p_{i,j} = 1 \quad \Rightarrow \quad \hat{p} = (H^T H)^{-1} H^T y$$

- Analytical expressions that describe the evolution of the profiling error

Enables systematic design!

Design as ptimization problem

$$MSE = \|p - \hat{p}\|^2 = \frac{1}{t}\left(N - 1 + \frac{1}{k}\right)\left(N - \sum_j \frac{f_j^2}{f^2 N}\right)$$

Users

rounds

Batch size

Senders that send a lot

Receivers receive from many

Pérez-González, Fernando, and Carmela Troncoso. "Understanding statistical disclosure: A least squares approach." PETS, 2012.
Oya, Simon, Carmela Troncoso, and Fernando Pérez-González. "Do dummies pay off? limits of dummy traffic protection in anonymous communications." PETS, 2014
Perez-Gonzalez, Fernando, Carmela Troncoso, and Simon Oya. "A least squares approach to the static traffic analysis of high-latency anonymous communication systems." TIFS 2014

Let's "do" the math

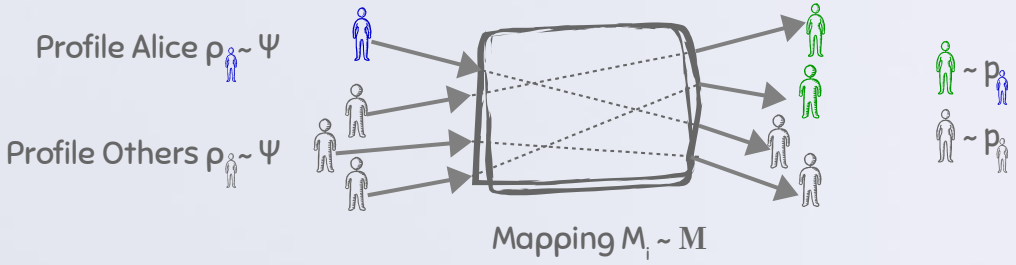

Approach 3: Disclosure attack as an inference problem

- What we are looking for: $Pr[\rho_i, \rho_j, M_i \mid O, M, \Psi]$

- More concretely, marginal probabilities & distributions
  - Pr[Alice->Bob] – Are Alice and Bob friends?
  - M_x – Who is talking to whom at round x?
  - Solve through sampling!

Profiles: $Pr[\rho_i, \rho_j \mid M_i, O, M, \Psi, K]$

(Direct sampling by sampling Dirichlet dist.)

Mappings: $Pr[M_i \mid \rho_i, \rho_j, O, M, \Psi, K]$

(Direct sampling of the matching link by link)

Danezis, George, and Carmela Troncoso. "Vida: How to use bayesian inference to de-anonymize persistent communications." PETS, 2009.

Persistent patterns Takeaways

- Near-perfect anonymity is not perfect enough!
  - High level patterns cannot be hidden for ever
  - Unobservability / maximal anonymity is needed

- Three approaches to the problem (actually I skipped the seminal work)

SDA

- Simple
- Fast!

- Best result not guaranteed
  - Only that one

LSDA

- Flexible
- Fast!

- Optimal result (MSE)
  - But only that one
- Error prediction
- Design tool!

Bayesian Inference

- Flexible
- "expensive"

- Distribution
  - Many quantities
  - Confidence intervals
- Not best solution

Agrawal, Dakshi, and Dogan Kesdogan. "Measuring anonymity: The disclosure attack." IEEE Security & Privacy, 2003
Kesdogan, Dogan, and Lexi Pimenidis. "The Hitting Set Attack on Anonymity Protocols." Information Hiding, 2004

Are we doomed? – Challenges

- **Countermeasures** – Systematic design?
  - Delay: plain batching does not seem the best
    - Pool mixes
    - Attacks can be adapted to account for more complex delay patterns
  - Dummy traffic: include "fake packets" to disorient the adversary
    - How do we make them indistinguishable?
    - Who decides about them?
  - Weaker protections suffice for other adversary models
    - e.g. Tor partial adversary
- **Privacy metric**, what is the goal?
- **Modeling adversarial knowledge**

Summary

- The Lord of The Rings is a great timeless book
- Crypto protects data, but does not always protect privacy
- Traffic analysis is the art of exploiting meta-data to extract information
- Traffic analysis can exploit a gzillion features: protecting efficiently is difficult!
  - Recovering persistent patterns, tracing messages in restricted routes
- Design privacy-preserving systems is FAR from trivial

institute imdea software

THANKS!

ANY QUESTIONS?

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

17TH PRIVACY ENHANCING TECHNOLOGIES SYMPOSIUM
JULY 18–21, 2017
MINNEAPOLIS, MN, USA

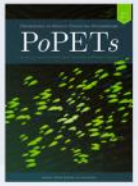

2018 BARCELONA! DEADLINES: 31 AUG, 30 NOV, 28 FEB

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

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