You cannot hide for long: De-anonymization of real-world dynamic behaviour

George Danezis (University College London)
Carmela Troncoso (Gradiant)

Gradiant
Centro Tecnolóxico de Telecomunicacións de Galicia

Privacy beyond confidentiality

- Common belief: “if I encrypt my data, then the data is private”
  - Encryption works and gets more and more efficient!
  - But does not hide all data
    - Origin and destination
    - Timing
    - Frequency
    - Location
    - …

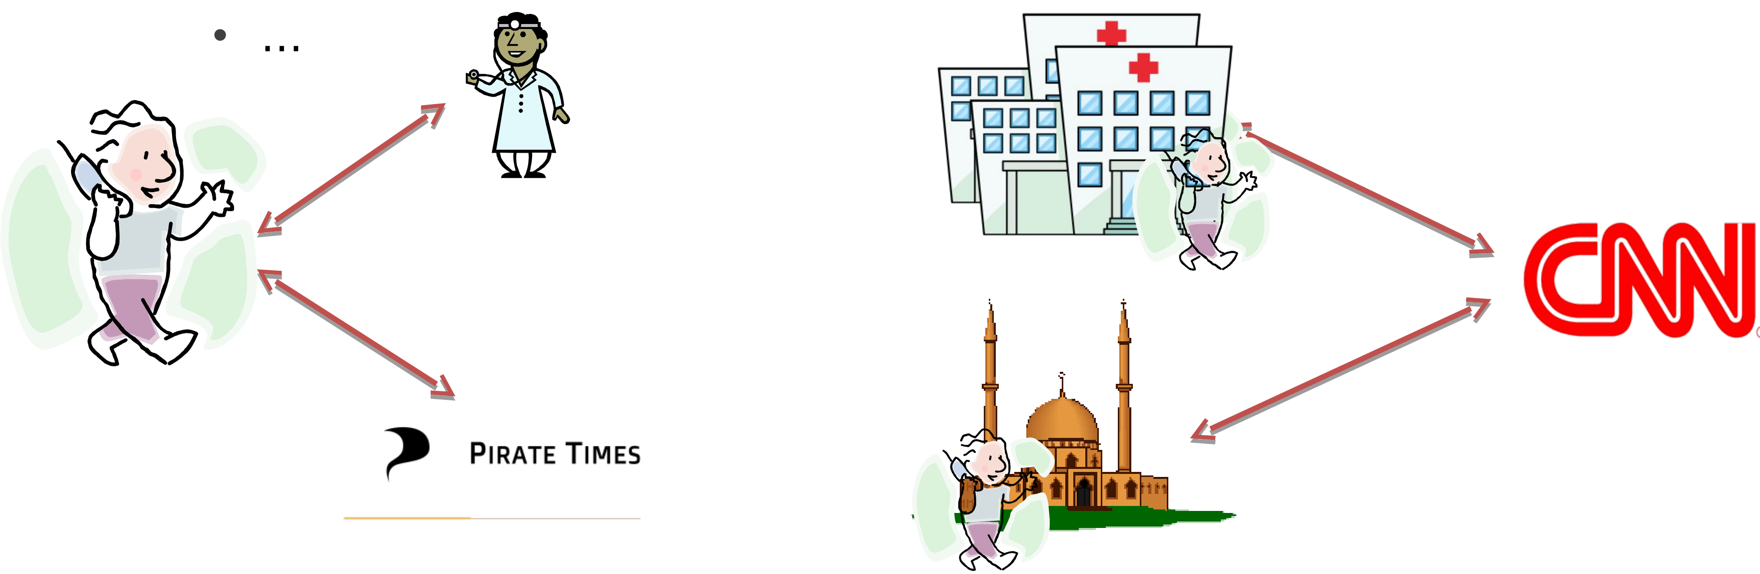

Anonymization

- Decouple user identity from actions

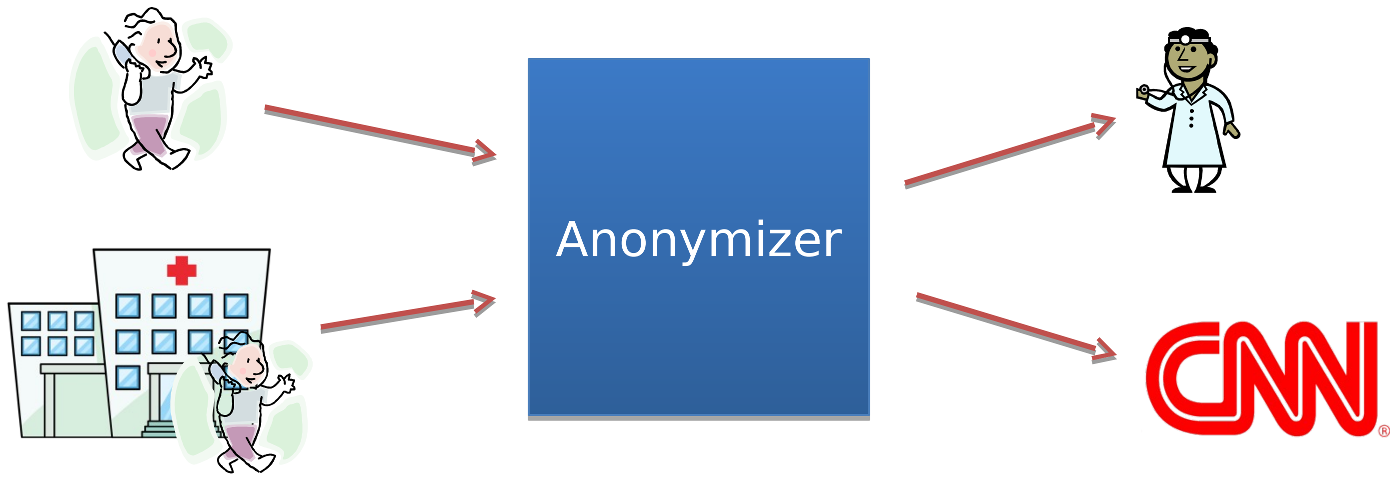

- Enabler for privacy-preserving technologies
  - Anonymous credentials
  - eVoting
  - Privacy-preserving statistics computation

Anonymity in reality

- Difficult to guarantee perfect anonymity due to constraints
  - Observations allow for inferences (e.g., behavioral profiles)

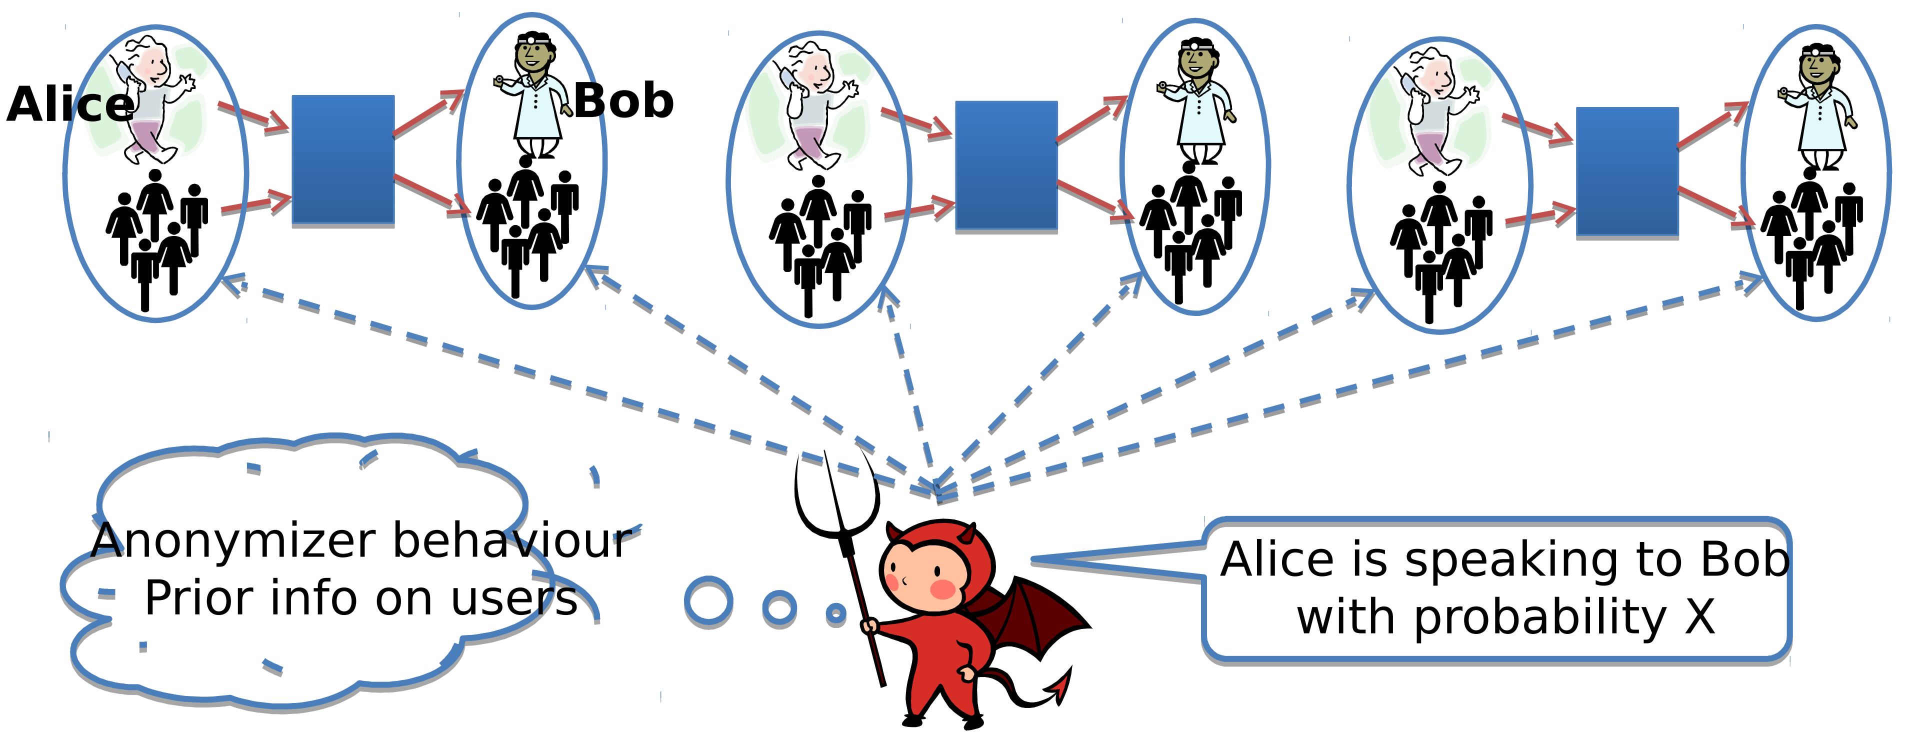

State of the art limitation: static behavior

A model for dynamic behaviour

- Users

 Sends messages to  at rate λ_{AB}

Sends messages to  at rate λ_{AO}

 Send messages to  at rate λ_{OB}

Send messages to  at rate λ_{OO}

Dynamism:
Epochs **t** of stationary behaviour
Profile evolution probability

$$\lambda_{AB}^{t} \xleftarrow{\$} E(\lambda_{AB}^{t} | \lambda_{AB}^{t-1})$$

- Anonymizer
  - Divided in batches (n batches per epoch)
  - Perfect anonymity

Given observation...
What is λ_{AB}?

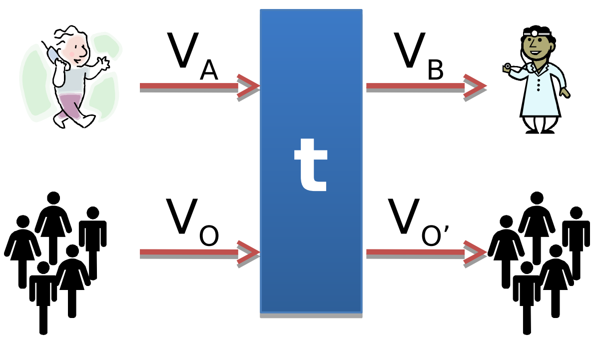

Visible

$$V_A \leftarrow \mathrm{Pois}(\lambda_{AB} + \lambda_{AO})$$
$$V_O \leftarrow \mathrm{Pois}(\lambda_{OB} + \lambda_{OO})$$
$$V_B \leftarrow \mathrm{Pois}(\lambda_{AB} + \lambda_{OB})$$
$$V_{O'} \leftarrow \mathrm{Pois}(\lambda_{AO} + \lambda_{OO})$$

Hidden

Sequential Monte Carlo aka. Particle Filters

- Inferring hidden parameters of sequential models
  - Our case: modeling λ_{AB} at t depends on λ_{AB} at t-1
- Core idea:
  - Particles representing sample hidden states (λ_{AB} , λ_{OB})
  - $\underbrace{\text{Distributed following posterior distribution given evidence } (V_x)}_{\text{Allow for statistic computation (mean, std, …) of hidden variables}}$
  - From Bayes theorem

Likelihood of obs. given hidden state

Prob evolving to current λ_{AB}

$$\Pr[(\lambda^t_{AB}, \lambda^t_{OB}) \mid V_*] \propto L(V_* \mid \lambda_*) E(\lambda^t_{AB} \mid \lambda^{t-1}_{AB}) \Pr[(\lambda^{t-1}_{AB}, \lambda^{t-1}_{OB})]$$

Prob at epoch t

Prior (epoch t-1)

Toy example

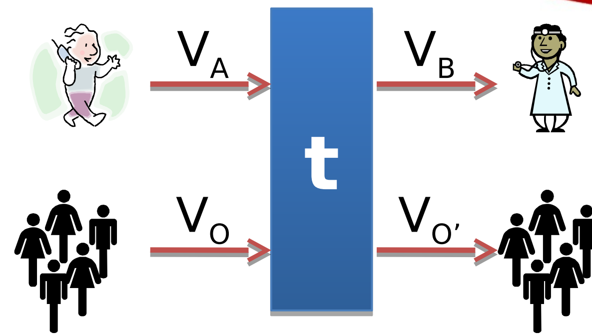

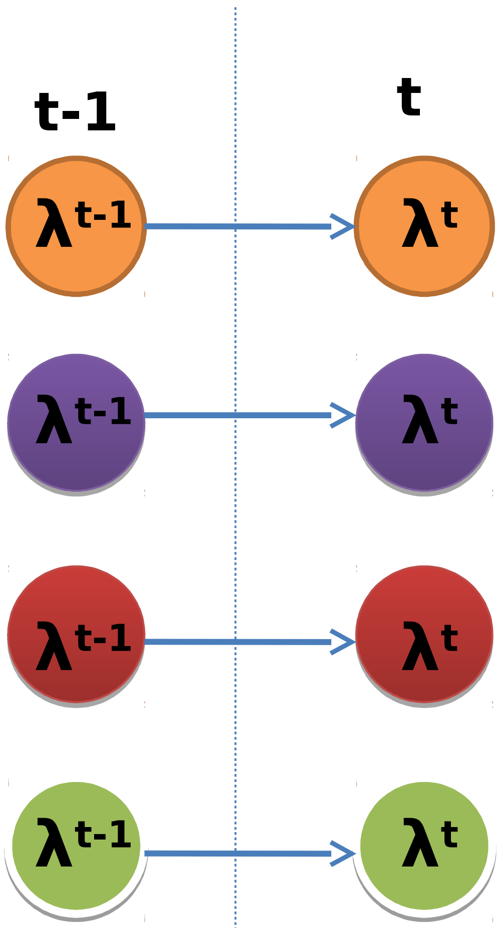

Weight particles:

i. **Likelihood**
ii. **Evolution**
iii. **Proposal**

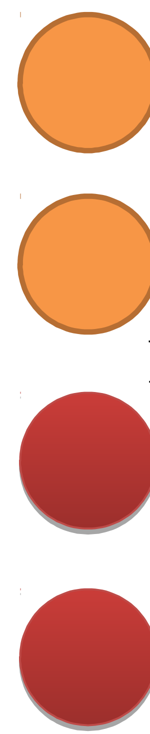

$$\Pr[(\lambda^t_{AB}, \lambda^t_{OB}) \mid V_*]$$

1. Propose new particles

2. Likelihood given Obs and previous state

3. Re-sample

In pseudocode

Take obs in all epochs

function SSDFILTER($V^{(t,n)}_{A,O,B,O'}$)

for all particles i do

$(\lambda^0_{ABi}, \lambda^0_{OBi}) \sim priors;$

end for

Initialize particles

for all epochs t do

for all particles i do

$\lambda_{Ai}, w_{Ai} \xleftarrow{\$} \text{Gamma}(V^t_A + 1, 1);$

$\lambda_{Oi}, w_{Oi} \xleftarrow{\$} \text{Gamma}(V^t_O + 1, 1);$

$\lambda_{Bi}, w_{Bi} \xleftarrow{\$} \text{Gamma}(V^t_B + 1, 1);$

$\lambda'_{ABi}, w_\theta \xleftarrow{\$} \text{Mixture } \mathcal{M}$

if $\lambda'_{ABi} > \lambda_{Ai}$ or $\lambda'_{ABi} > \lambda_{Bi}$ then

reject & continue;

end if

$\lambda'_{OBi} \leftarrow \lambda_{Bi} - \lambda'_{ABi};$

Propose current state given observation

All types of samples

$w_{ABi} \leftarrow L(V^{(t,n)}_{\{A,B,O,O'\}} | \lambda^t_{\{\star\}}) \cdot E(\lambda'_{ABi} | \lambda^{t-1}_{ABi});$

$w_i \leftarrow w_{ABi} / w_{Ai} \cdot w_{Oi} \cdot w_{Bi} \cdot w_\theta;$

end for

Likelihood of observation given current and previous state

Reweighting of proposal likelihood given proposal distributions

for all particles i do

$(\lambda^t_{ABi}, \lambda^t_{OBi}) \leftarrow \text{Re-sample } (\lambda'_{ABi}, \lambda'_{OBi}) \sim w_i;$

end for

end for

return $(\lambda^{\max t}_{ABi}, \lambda^{\max t}_{OBi});$

end function

Resampling to obtain new particles according to posterior

The likelihood function $L(V_* \mid \lambda_*)$

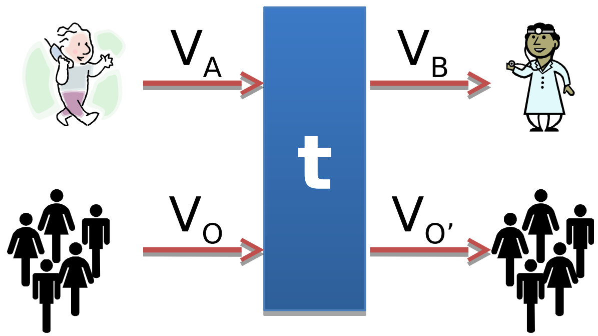

Visible

$$V_A \leftarrow \text{Pois}(\lambda_{AB} + \lambda_{AO})$$
$$V_O \leftarrow \text{Pois}(\lambda_{OB} + \lambda_{OO})$$
$$V_B \leftarrow \text{Pois}(\lambda_{AB} + \lambda_{OB})$$
$$V_{O'} \leftarrow \text{Pois}(\lambda_{AO} + \lambda_{OO})$$

Hidden

- How likely is an observation V_* given sending rates λ_*

$$L = \Pr[V_A^t; \lambda_{AB}^t + \lambda_{AO}^t] \cdot \Pr[V_O^t; \lambda_{OB}^t + \lambda_{OO}^t] \prod_{n=1}^{N} L^n$$

Prob of each of the rounds

Prob of total volume in epoch given λ_* (just Poisson)

Binomial

$$L^n = \sum_{k=0}^{\min(V_A^{(t,n)}, V_B^{(t,n)})} \Pr_b[k; V_A^{(t,n)}, p_{ab}^t] \cdot \Pr_b[V_B^{(t,n)} - k; V_O, p_{ob}^t]$$

p_{ab} is just the probability A sent to B $p_{ab} = (\lambda_{AB} / \lambda_{AB} + \lambda_{OB})$

The profile evolution probability $E(\lambda_{AB}^{t} | \lambda_{AB}^{t-1})$

- Probability of λ_{AB} at t given λ_{AB} at t-1

- Two stages
  1) Probability transitions silent-communication
  2) Probability of given difference: mixture with heavy tails

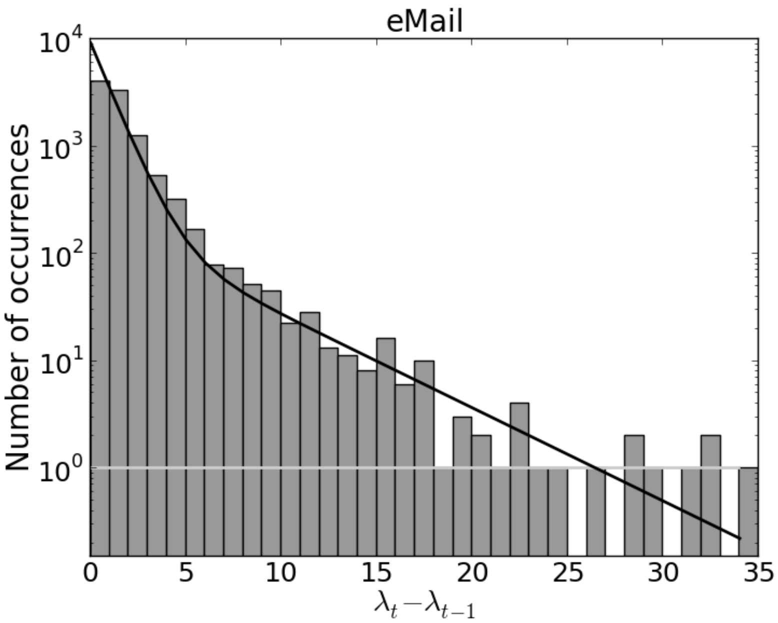

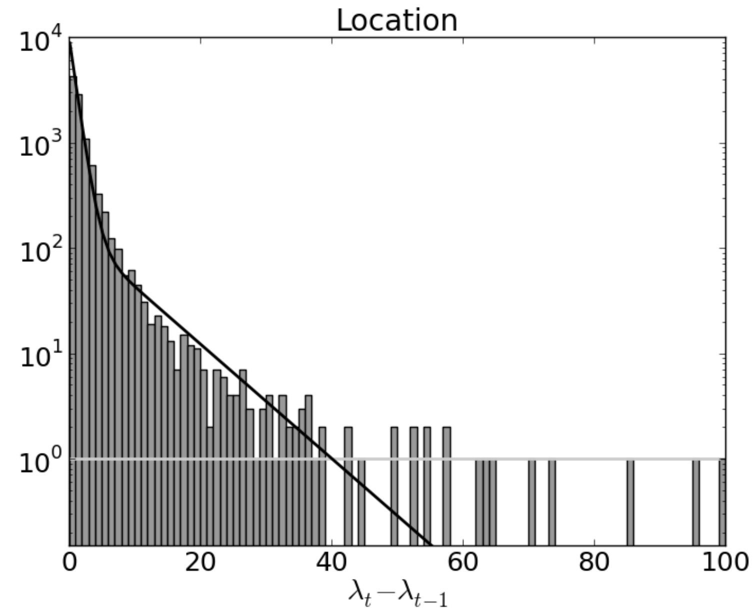

Evaluation

- Three datasets:
  - eMail: Enron dataset ~0.5M emails, 150 users.
  - Mailing list: Indymedia ~300K posts from 28237 senders to 693 lists
  - Location: Gowala dataset ~6.5M checkins from ~200K users
- Parameters empirically inferre

	p_Z	p_{PZ}	p_{ZZ}	p	λ_s	λ_e
eMail	0.958	0.04	0.995	0.88	1.0	4.0
Mailing list	0.982	0.02	0.998	0.97	1.0	22.0
Location	0.993	0.06	0.999	0.87	1.0	7.0

Prior silent

Transitions
Stop talking
Stay silent

Mixture evolution

- Two sets
  - Communication
  - Silent
- Anonymity system
  - 1 day delay (anonymity vs delay trade-off given 1 week epochs)
  - Thresholds: eMail/Mailing ~100 Location ~15K

Evaluation - an example trace (Avg(Batch)= 244)

- State of the art: Statistical Disclosure Attack
  - Background traffic:  messages 
  - Use background to estimate to  volume in her rounds
- Assumes static behaviour: short and long term

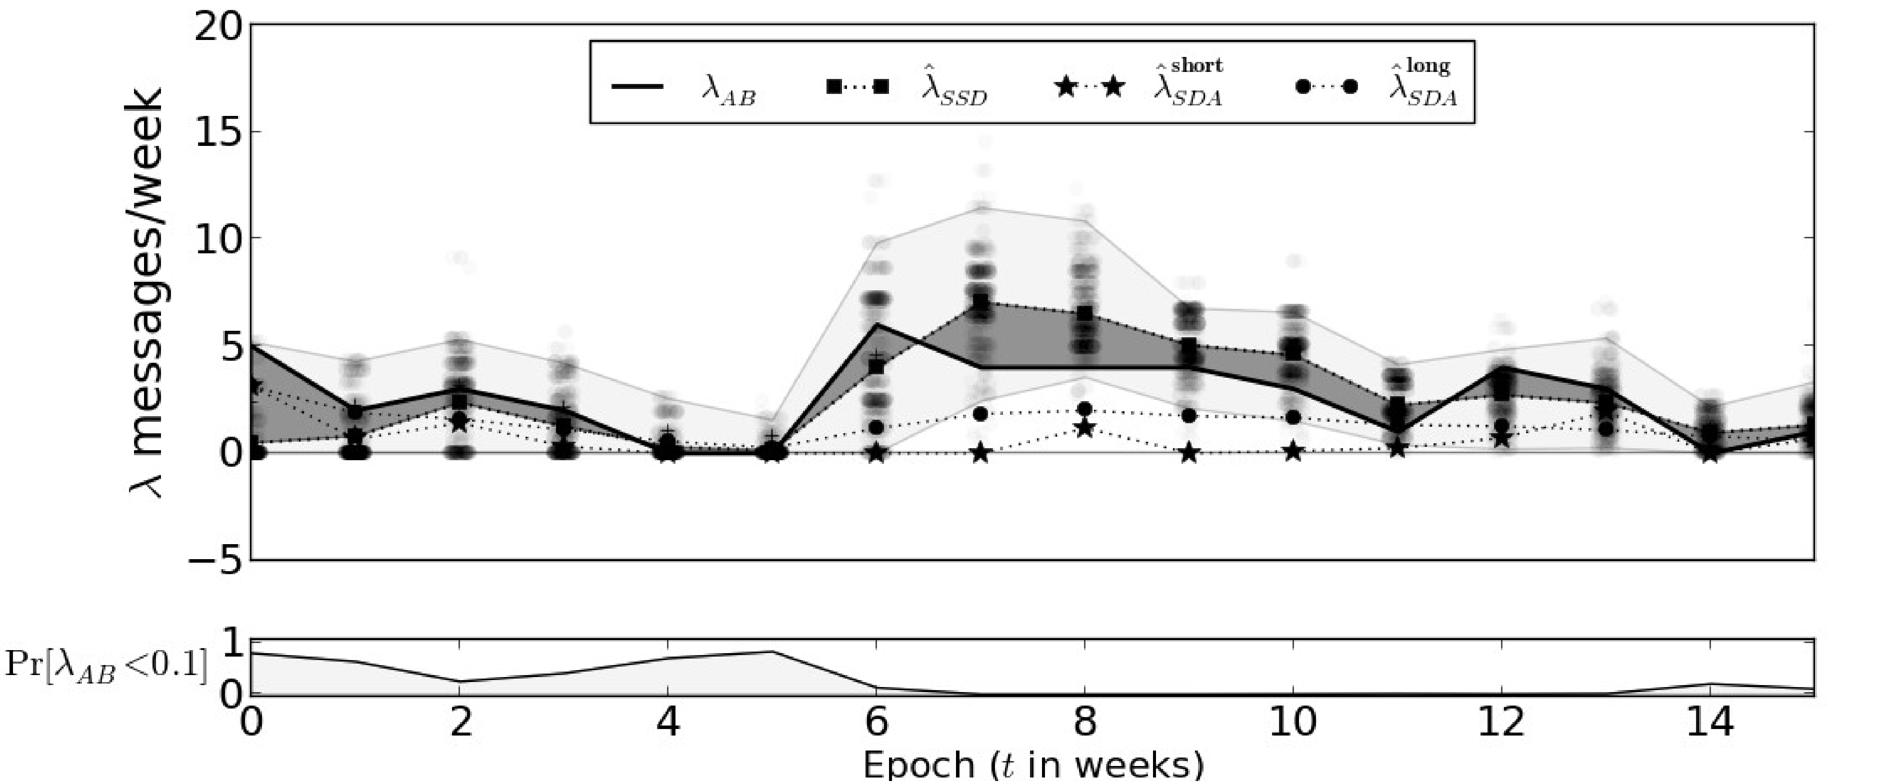

Evaluation – estimation accuracy as Squared error

MSE_{Comm}	112	13	20	84	3.7K	83K
MSE_{Silent}	0.7	2.8	0.8	1.2K	2.3K	360

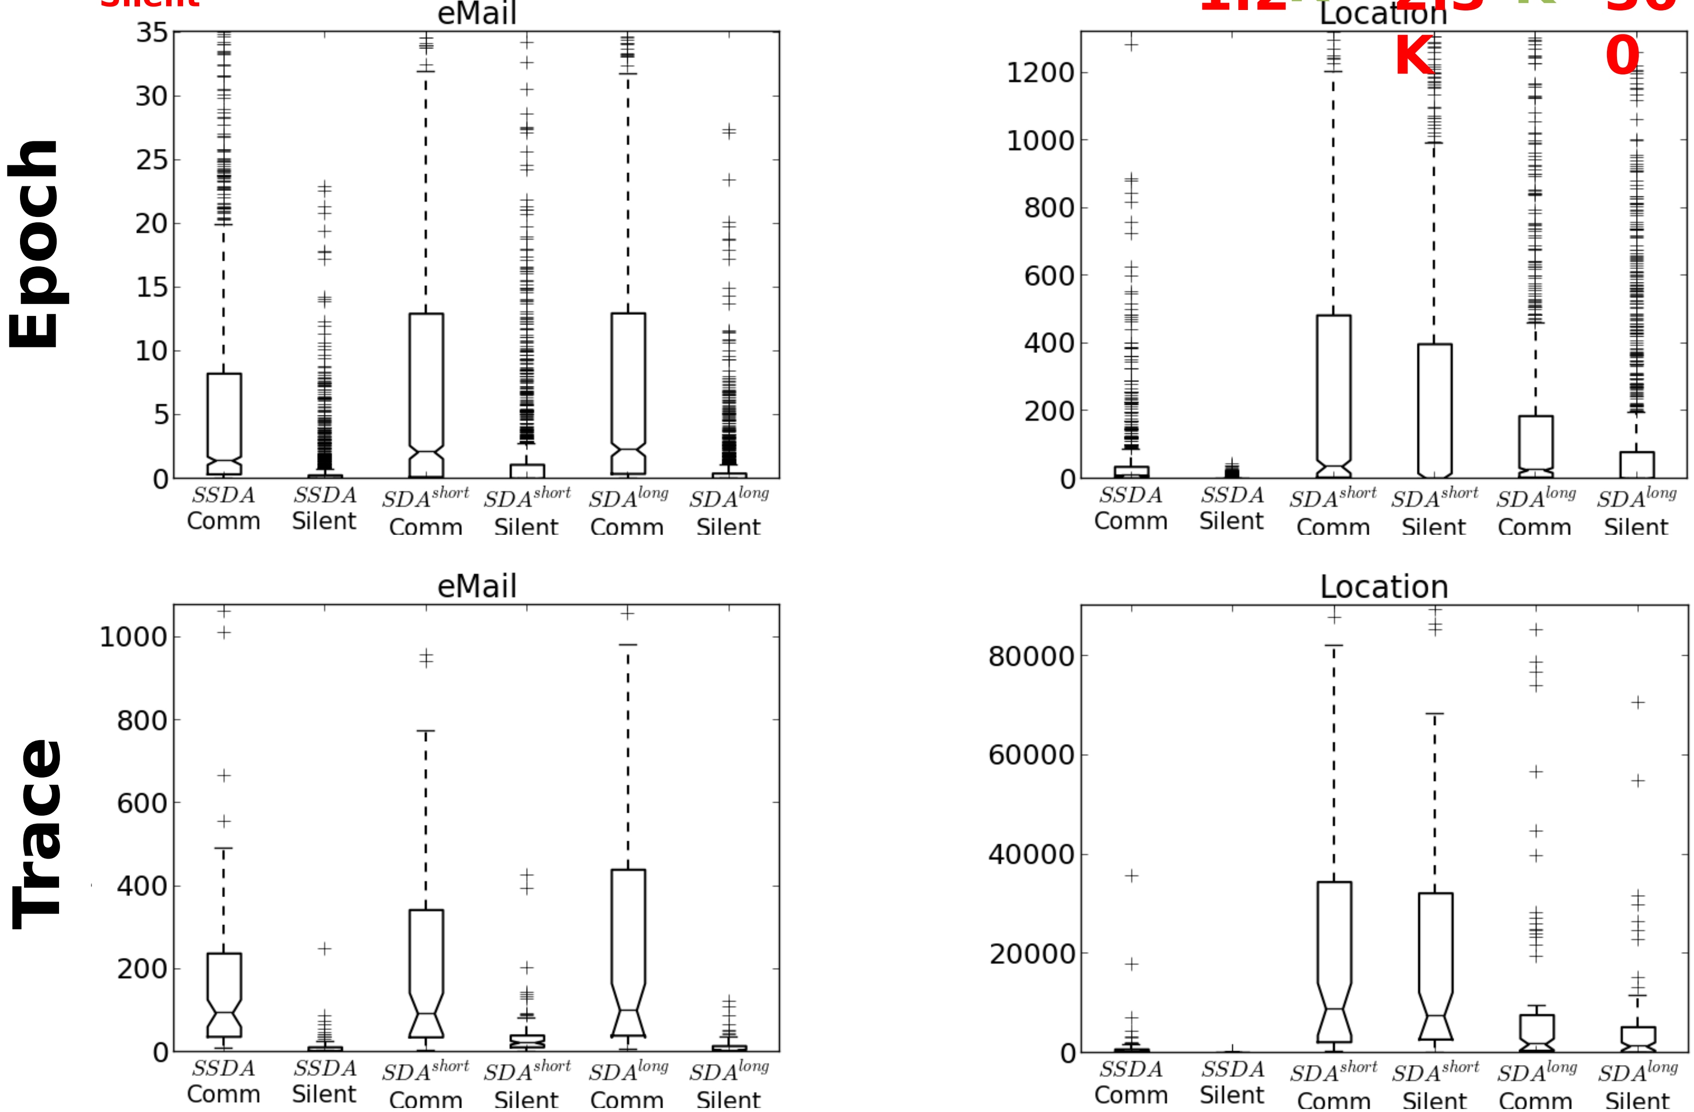

Evaluation – communication detection

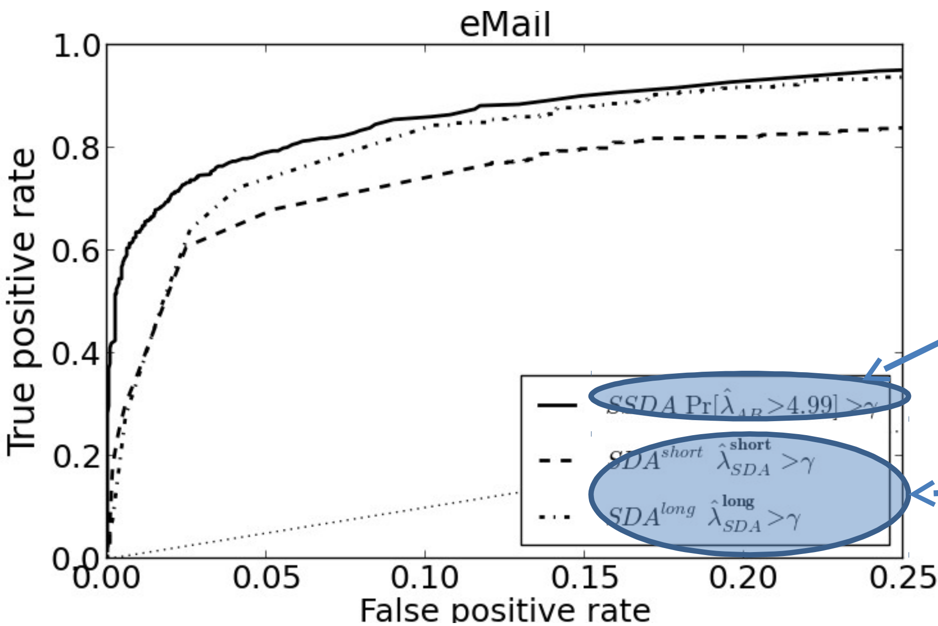

Are Alice and Bob communicating?
Base rate fallacy!

Use particles distribution

Use rate directly

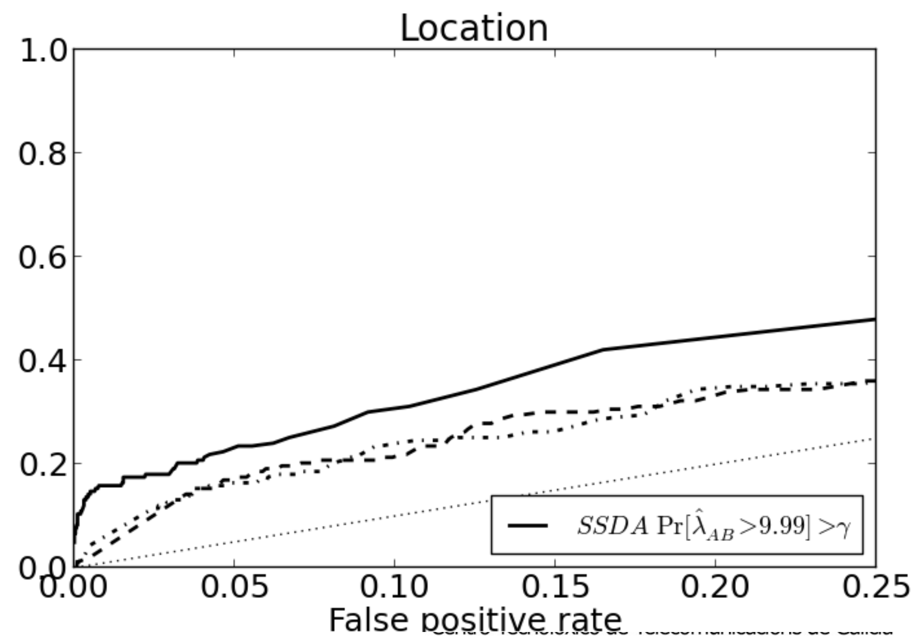

Conclusions

- Structured model for traffic analysis based on known Bayesian inference techniques
  - easy to extend
  - allow assessment of inference quality
  - avoid base rate fallacy
- Attacks on real world traces
  - can be effective for rather low action rates
  - can be effective over a much shorter period of time than previously thought
  - can be effective for secure configurations of the anonymity system
- Rethink current evaluations and figures of merit

Thanks!!

ctroncoso@gradiant.org
g.danezis@ucl.ac.uk