Analyzing and Protecting Communication Metadata

Ludovic Barman

Laboratory for Data Security & Decentralized and Distributed Systems Laboratory

PhD Private Defense, 22.06.2021

Committee President:	Prof.	Mathias Payer
Advisors:	Prof.	Jean-Pierre Hubaux
	Prof.	Bryan Ford
Experts:	Prof.	Claudia Diaz
	Dr.	Nina Taft
	Prof.	Carmela Troncoso

Communication systems leak metadata

- Goal: protect sensitive information from network eavesdroppers 😈
- Encryption is used to provide confidentiality
- Often, some metadata remain unprotected:
    - the (time, size) of network packets
    - the identity of the sender or recipient
    - at what times a party is sending messages

      ...

Metadata can reveal sensitive information

- In research, metadata from network traces help to infer the contents:
    - Tor: visited web pages
    - VPN: contents and the destination server
    - Skype: English words
    - Smartphones: installed applications
    - IoT: user activity

- Practical example using source / destination:
    - NSA phone-calls metadata collection

Analyzing and Protecting Communication Metadata

Attacks ⟷ Defenses

Analyzing and Protecting Communication Metadata

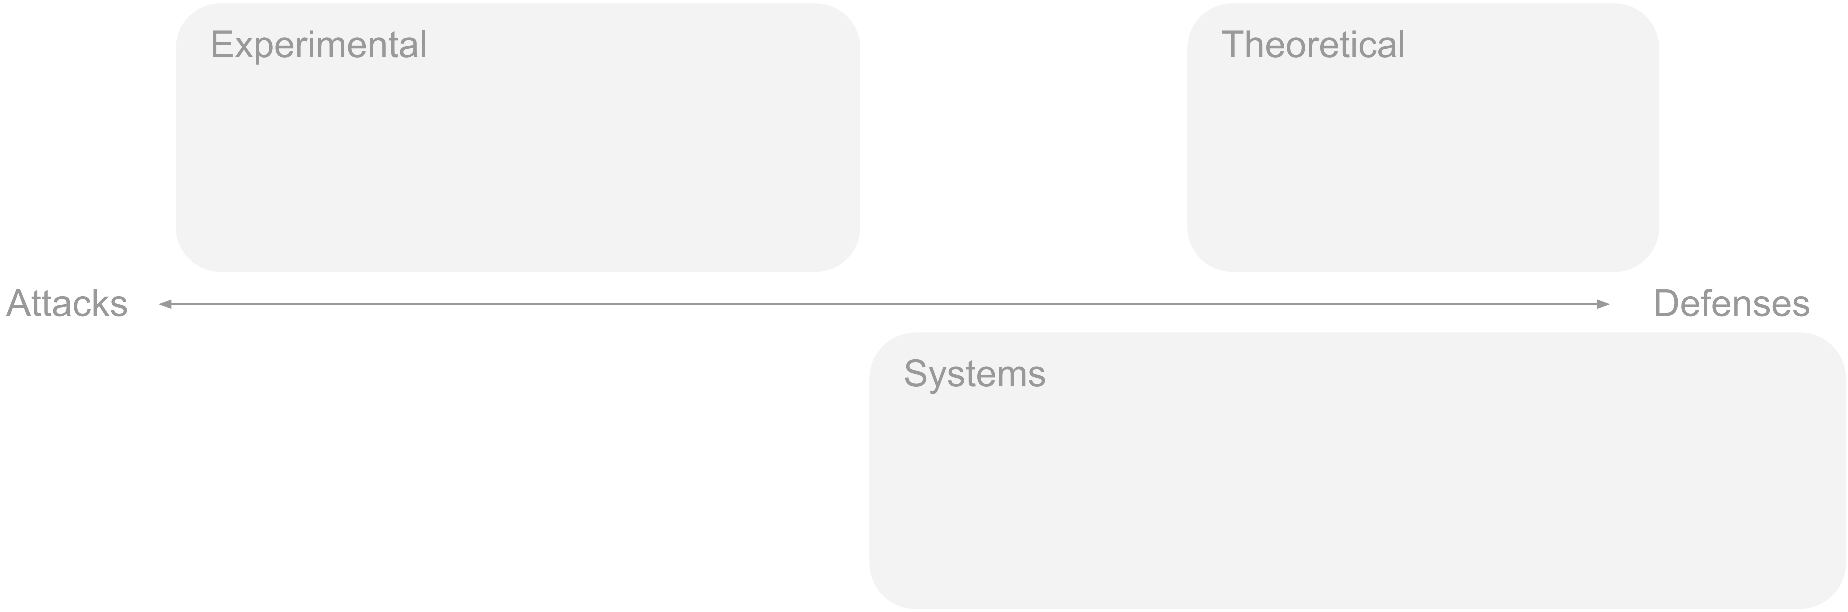

Analyzing and Protecting Communication Metadata

Experimental

Theoretical

Attacks ⟷ Defenses

Systems

Analyzing and Protecting Communication Metadata

Experimental

Every Byte Matters (Ch2) [4]

Traffic-analysis attack of wearable devices.

Theoretical

Attacks ⟷ Defenses

Systems

[4] L. Barman, A. Dumur, A. Pyrgelis, J.-P. Hubaux. **Every Byte Matters: Traffic Analysis of Bluetooth Wearable Devices.** UbiComp 2021.

Analyzing and Protecting Communication Metadata

Experimental

Every Byte Matters (Ch2) [4]

Traffic-analysis attack of wearable devices.

Theoretical

Attacks ⟷ Defenses

Systems

[4] L. Barman, A. Dumur, A. Pyrgelis, J.-P. Hubaux. **Every Byte Matters: Traffic Analysis of Bluetooth Wearable Devices.** UbiComp 2021.

Analyzing and Protecting Communication Metadata

Experimental

Every Byte Matters (Ch2) [4]
Traffic-analysis attack of wearable devices.

Theoretical

Padmé (Ch3) [2]
A padding function that efficiently hides sizes.

Attacks ⟷ Defenses

Systems

[2] K. Nikitin*, L. Barman*, W. Lueks, M. Underwood, J.-P. Hubaux, B. Ford. **Reducing Metadata Leakage from Encrypted Files and Communication with PURBs.** PETS 2019

[4] L. Barman, A. Dumur, A. Pyrgelis, J.-P. Hubaux. **Every Byte Matters: Traffic Analysis of Bluetooth Wearable Devices.** UbiComp 2021.

Analyzing and Protecting Communication Metadata

Experimental

Every Byte Matters (Ch2) [4]

Traffic-analysis attack of wearable devices.

Theoretical

Padmé (Ch3) [2]

A padding function that efficiently hides sizes.

Attacks ⟷ Defenses

Systems: Anonymous Communication Networks (ACN)

[2] K. Nikitin*, L. Barman*, W. Lueks, M. Underwood, J.-P. Hubaux, B. Ford. **Reducing Metadata Leakage from Encrypted Files and Communication with PURBs.** PETS 2019

[4] L. Barman, A. Dumur, A. Pyrgelis, J.-P. Hubaux. **Every Byte Matters: Traffic Analysis of Bluetooth Wearable Devices.** UbiComp 2021.

Analyzing and Protecting Communication Metadata

Experimental

Every Byte Matters (Ch2) [4]

Traffic-analysis attack of wearable devices.

Theoretical

Padmé (Ch3) [2]

A padding function that efficiently hides sizes.

Attacks ⟷ Defenses

Systems: Anonymous Communication Networks (ACN)

PriFi (§4.4) [1,3]

Traffic-agnostic, low-latency ACN for local-area networks.

[1] L. Barman, M. Zamani, I. Dacosta, J. Feigenbaum, B. Ford, J.-P. Hubaux, D. Wolinsky. **PriFi: A Low-latency [...] Protocol for Local-Area Anonymous [...]**. WPES 2016.

[2] K. Nikitin*, L. Barman*, W. Lueks, M. Underwood, J.-P. Hubaux, B. Ford. **Reducing Metadata Leakage from Encrypted Files and Communication with PURBs.** PETS 2019

[3] L. Barman, I. Dacosta, M. Zamani, E. Zhai, A. Pyrgelis, B. Ford, J. Feigenbaum, J.-P. Hubaux. **PriFi: Low-latency Anonymity for Organizational Networks.** PETS 2020

[4] L. Barman, A. Dumur, A. Pyrgelis, J.-P. Hubaux. **Every Byte Matters: Traffic Analysis of Bluetooth Wearable Devices.** UbiComp 2021.

Analyzing and Protecting Communication Metadata

Experimental

Every Byte Matters (Ch2) [4]
Traffic-analysis attack of wearable devices.

Theoretical

Padmé (Ch3) [2]
A padding function that efficiently hides sizes.

Attacks ⟷ Defenses

Systems: Anonymous Communication Networks (ACN)

PriFi (§4.4) [1,3]
Traffic-agnostic, low-latency ACN for local-area networks.

Rubato (§4.5) [5]
Large-scale ACN for text messaging on mobile devices.

[1] L. Barman, M. Zamani, I. Dacosta, J. Feigenbaum, B. Ford, J.-P. Hubaux, D. Wolinsky. **PriFi: A Low-latency [...] Protocol for Local-Area Anonymous [...]**. WPES 2016.

[2] K. Nikitin*, L. Barman*, W. Lueks, M. Underwood, J.-P. Hubaux, B. Ford. **Reducing Metadata Leakage from Encrypted Files and Communication with PURBs**. PETS 2019

[3] L. Barman, I. Dacosta, M. Zamani, E. Zhai, A. Pyrgelis, B. Ford, J. Feigenbaum, J.-P. Hubaux. **PriFi: Low-latency Anonymity for Organizational Networks**. PETS 2020

[4] L. Barman, A. Dumur, A. Pyrgelis, J.-P. Hubaux. **Every Byte Matters: Traffic Analysis of Bluetooth Wearable Devices**. UbiComp 2021.

[5] L. Barman, M. Kol, D. Lazar, Y. Gilad, N. Zeldovich. **Rubato: Metadata-Private Messaging for Mobile Devices**. Under submission.

Focus of this talk

Experimental

Every Byte Matters (Ch2) [4]
Traffic-analysis attack of wearable devices.

Theoretical

Padmé (Ch3) [2]
A padding function that efficiently hides sizes.

Attacks ⟷ Defenses

Systems: Anonymous Communication Networks (ACN)

PriFi (§4.4) [1,3]
Traffic-agnostic, low-latency ACN for local-area networks.

Rubato (§4.5) [5]
Large-scale ACN for text messaging on mobile devices.

[1] L. Barman, M. Zamani, I. Dacosta, J. Feigenbaum, B. Ford, J.-P. Hubaux, D. Wolinsky. **PriFi: A Low-latency [...] Protocol for Local-Area Anonymous [...]**. WPES 2016.

[2] K. Nikitin*, L. Barman*, W. Lueks, M. Underwood, J.-P. Hubaux, B. Ford. **Reducing Metadata Leakage from Encrypted Files and Communication with PURBs**. PETS 2019

[3] L. Barman, I. Dacosta, M. Zamani, E. Zhai, A. Pyrgelis, B. Ford, J. Feigenbaum, J.-P. Hubaux. **PriFi: Low-latency Anonymity for Organizational Networks**. PETS 2020

[4] L. Barman, A. Dumur, A. Pyrgelis, J.-P. Hubaux. **Every Byte Matters: Traffic Analysis of Bluetooth Wearable Devices**. UbiComp 2021.

[5] L. Barman, M. Kol, D. Lazar, Y. Gilad, N. Zeldovich. **Rubato: Metadata-Private Messaging for Mobile Devices**. Under submission.

Every Byte Matters:
Traffic Analysis of Wearable Devices

(Chapter 2)

Setting

Wearable devices communicate with a smartphone over Bluetooth

Fitness tracker

Smartwatch

Sleep tracker

ECG/BPM

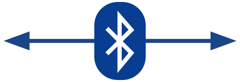

The data exchanged is personal and sensitive

Consumer devices | Medical devices

Smartwatch

Fitness tracker

Sleep tracker

ECG/BPM

The data exchanged is personal and sensitive

Consumer devices → Medical devices

Smartwatch

Fitness tracker

Sleep tracker

ECG/BPM

...

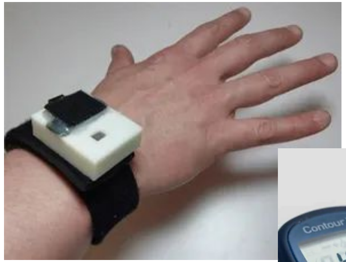

Asthma monitor

Blood sugar monitor

The data exchanged is personal and sensitive

Consumer devices → Medical devices

Fitness tracker

Smartwatch

CG/BPM

...

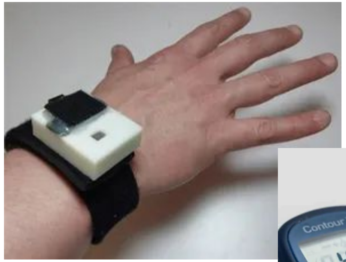

Asthma monitor

Blood sugar monitor

- List of apps (app to stop smoking, medication reminder)
- User activities (e.g., recording an insulin injection)

Privacy of Bluetooth communications

- Eavesdropping is expensive today

- Cost of eavesdropping is decreasing ↘

Privacy of Bluetooth communications

Motivation: What will eavesdroppers learn from Bluetooth wearable devices ?

Privacy of Bluetooth communications

Motivation: What will eavesdroppers learn from Bluetooth wearable devices ?

Do Bluetooth wearable devices leak metadata ?

- Simple firmware with few capabilities => easy to model & fingerprint ?
- Power-constrained devices that transmit little data => naturally protected ?
- Bluetooth network stack specifics ?

Privacy of Bluetooth communications

Motivation: What will eavesdroppers learn from Bluetooth wearable devices ?

Do Bluetooth wearable devices leak metadata ?

- Simple firmware with few capabilities => easy to model & fingerprint ?
- Power-constrained devices that transmit little data => naturally protected ?
- Bluetooth network stack specifics ?

Our contribution: Analysis of the encrypted communications of Bluetooth wearable devices

Examples of attacks

- Can an advertiser in a store recognize users/devices from encrypted Bluetooth traffic ?

Smart billboards: the real-life cookie

Sep 2, 2020 | Article

- Can a smart billboard infer nearby activities from Bluetooth traffic?

- Can a nosy neighbor infer my daily routine from wearable devices?

Test Bed

We cover popular vendors and devices:

Bluetooth Classic

Vendor	Model	OS
Samsung	Galaxy Watch	Tizen
Fossil	Explorist HR	Wear OS 2
Apple	Watch 4	watchOS 5
Huawei	Watch 2	Wear OS 2
Fitbit	Versa 2	Fitbit OS 4
Sony	MDR-XB9	-
Apple	AirPods	-

Bluetooth LE

Vendor	Model
Apple	Watch 4
Fitbit	Charge 2
Fitbit	Charge 3
Huawei	Band 3e
Mi	Band 2
Mi	Band 3
Mi	Band 4

smartwatches
headphones
step counters & fitness trackers

Phones used: Nexus 5, iPhone 8

Data collection

Challenges:

- Heterogeneous devices
- Only Wear OS can be automated
- Generating real samples is difficult (e.g., UI Fuzzing won’t create realistic traces)

Methodology: We *manually* use the devices in the intended way, recording Bluetooth traffic.

We collect a dataset of 10’700 samples (≈ 100h of recording, 30-sec samples):

- 32 actions (e.g., Add Insulin, Measure Heart Rate, Start Workout, ...)
- 80 applications (categories: Religion, Health, Lifestyle, Local newspapers, ...)

Features

We use simple, standard features (e.g., proposed in [6]):

- **General statistics**:

$$\left.\begin{array}{l}\texttt{min}\\ \texttt{mean}\\ \texttt{max}\\ \texttt{count}\\ \texttt{std dev}\end{array}\right\} \times \left\{\begin{array}{l}\text{of the list of packet sizes from}\\ \text{of the inter-packet timings from}\end{array}\right\} \times \left\{\begin{array}{l}\text{Slave -> Master}\\ \text{Master -> Slave}\\ \text{All with non-null payload}\end{array}\right.$$

- **Size histograms**: 10-byte wide “buckets” [7] that count the packets of corresponding sizes

- **Bursts**: $\text{AvgIPT (seq)} = \dfrac{\sum_i \text{time}_{i+1} - \text{time}_i}{|\text{seq}| - 1}$ [8]

[6] J. Hayes, G. Danezis. **k-fingerprinting: A Robust Scalable Website Fingerprinting Technique**. Usenix Security 2016.

[7] M. Liberatore, B. N. Levine. **Inferring The Source of Encrypted HTTP Connections.** CSS 2006.

[8] B. Saltaformaggio et al. **Eavesdropping on fine-grained user activities within smartphone apps over encrypted network traffic**. WOOT 2016.

Feature Extraction & Training

- We use standard features [9]
- We use a simple, standard model (Random Forests)

Dataset:

80% training	20% testing

[9] J. Hayes, G. Danezis. **k-fingerprinting: A Robust Scalable Website Fingerprinting Technique**. Usenix Security 2016.

Identifying devices from traffic patterns

- Paired devices can stop advertising
- No friendly names (e.g. “Ludovic’s Apple Watch 4”), no MAC address

Identifying devices from traffic patterns

- Paired devices can stop advertising
- No friendly names (e.g. “Ludovic’s Apple Watch 4”), no MAC address

Methodology:

- 2 classifiers (Bluetooth Classic + LE)
- 7 devices each

Identifying devices from traffic patterns

- Paired devices can stop advertising
- No friendly names (e.g. “Ludovic’s Apple Watch 4”), no MAC address

Methodology:

- 2 classifiers (Bluetooth Classic + LE)
- 7 devices each

Result:

- In both cases: ~96% precision / recall
- Top features: timings for Bluetooth Classic, sizes for LE

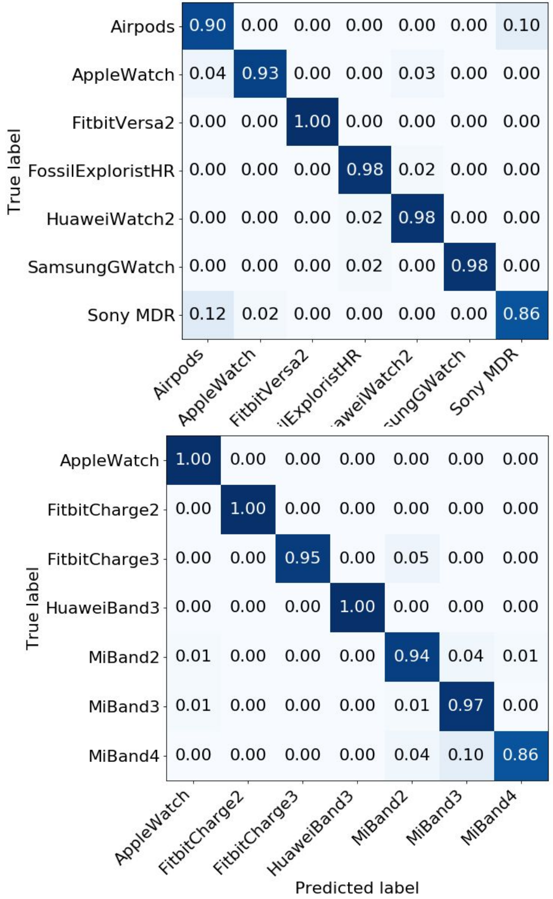

Identifying devices from traffic patterns

- Paired devices can stop advertising
- No friendly names (e.g. “Ludovic’s Apple Watch 4”), no MAC address

Methodology:

- 2 classifiers (Bluetooth Classic + LE)
- 7 devices each

Result:

- In both cases: ~96% precision / recall
- Top features: timings for Bluetooth Classic, sizes for LE

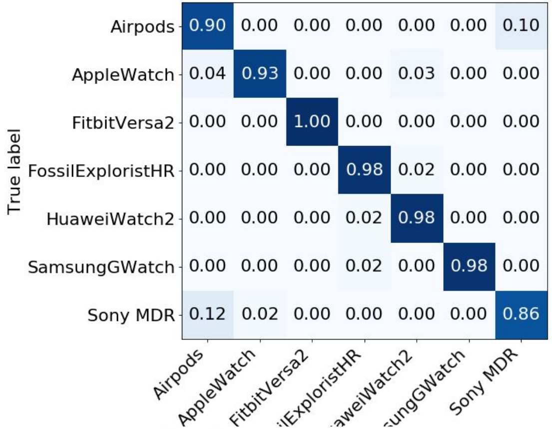

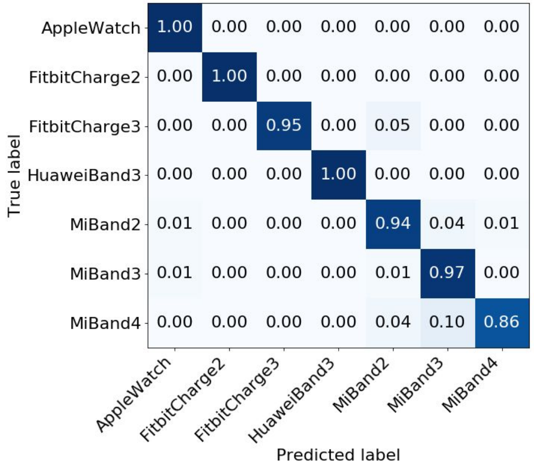

Take-away:

Identifying user actions from traffic patterns

Methodology:

- 49 actions across 13 devices
- classify device + user action in one step

Identifying user actions from traffic patterns

Methodology:

- 49 actions across 13 devices
- classify device + user action in one step

Result:

- 82% precision / recall / F1 score

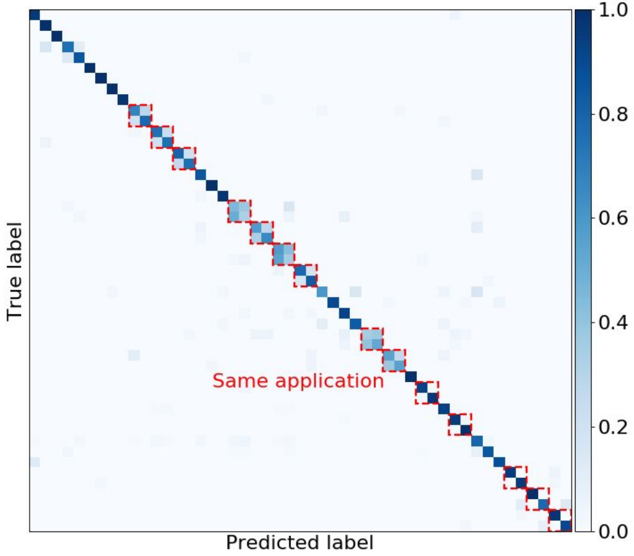

Identifying user actions from traffic patterns

Methodology:

- 49 actions across 13 devices
- classify device + user action in one step

Result:

- 82% precision / recall / F1 score

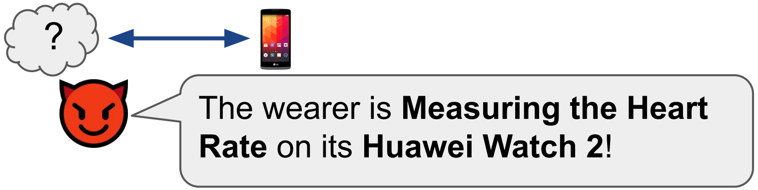

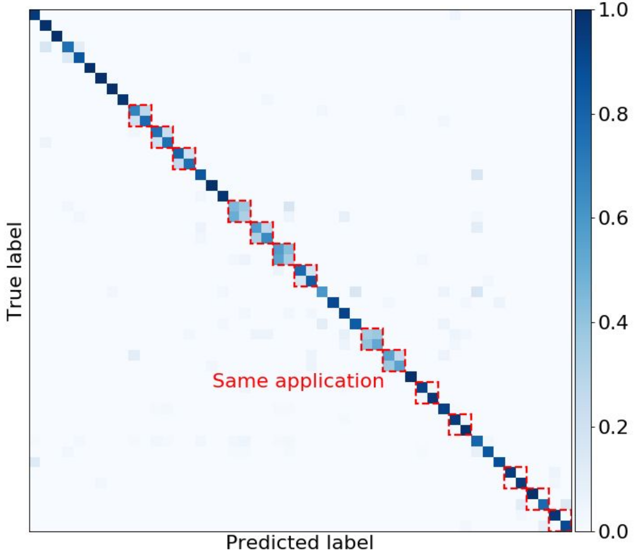

Identifying user actions from traffic patterns

Methodology:

- 49 actions across 13 devices
- classify device + user action in one step

Result:

- 82% precision / recall / F1 score

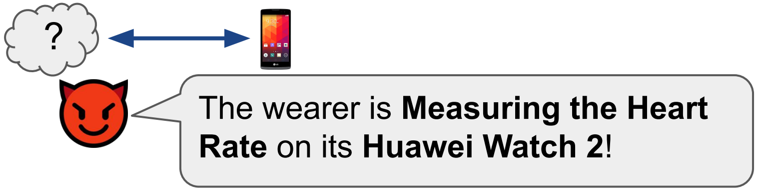

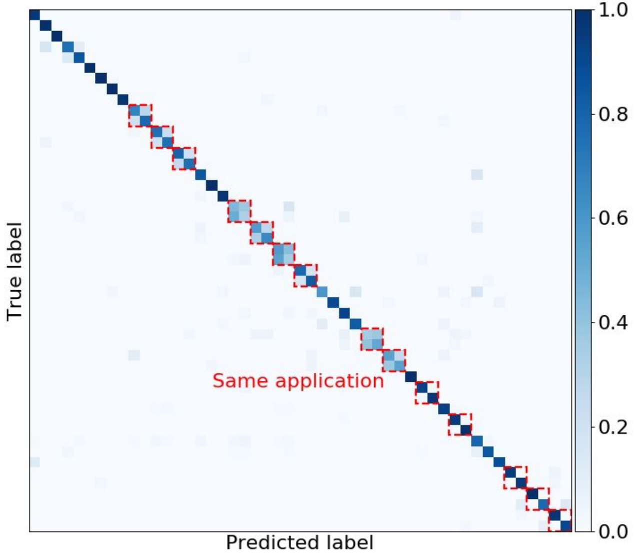

Take-away: across all devices, most user actions generate unique patterns

Identifying applications on Wear OS

Methodology:

- recognizing the opening of a particular app (e.g., DiabetesM, StopSmoking) on Wear OS
- 56 applications

Identifying applications on Wear OS

Methodology:

- recognizing the opening of a particular app (e.g., DiabetesM, StopSmoking) on Wear OS
- 56 applications

Result:

- 64% precision / recall / F1 score

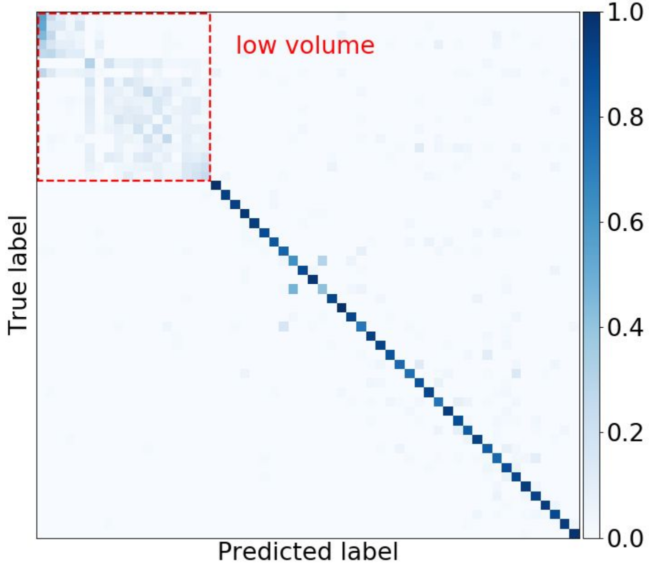

Identifying applications on Wear OS

Methodology:

- recognizing the opening of a particular app (e.g., DiabetesM, StopSmoking) on Wear OS
- 56 applications

Result:

- 64% precision / recall / F1 score

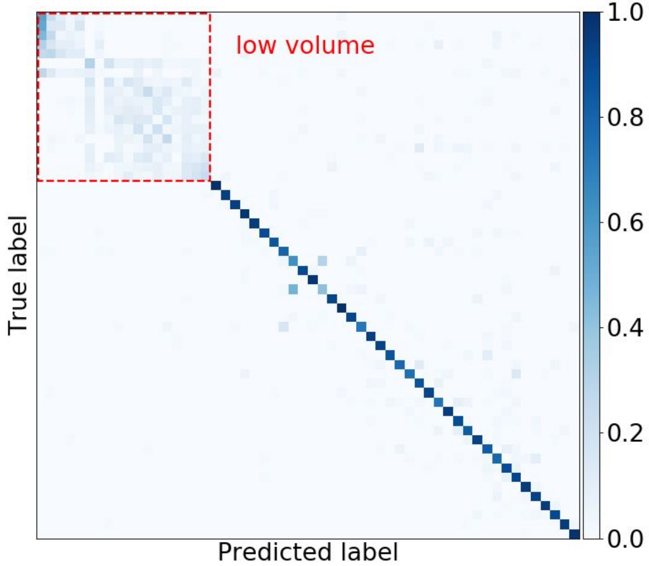

Identifying applications on Wear OS

Methodology:

- recognizing the opening of a particular app (e.g., DiabetesM, StopSmoking) on Wear OS
- 56 applications

Result:

- 64% precision / recall / F1 score

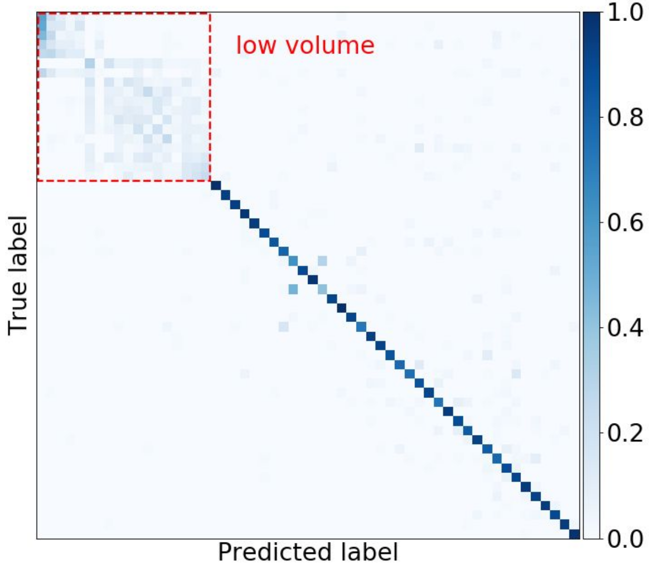

Take-away: the majority of apps can be recognized upon being opened

Identifying actions within an application

Methodology:

- 6 actions within the app `DiabetesM`

Identifying actions within an application

Methodology:

- 6 actions within the app `DiabetesM`

Result:

- 70% precision / recall / F1 score
- Top features: timings

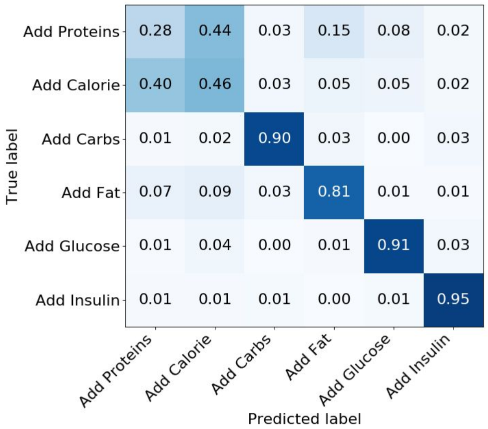

Identifying actions within an application

Methodology:

- 6 actions within the app `DiabetesM`

Result:

- 70% precision / recall / F1 score
- Top features: timings

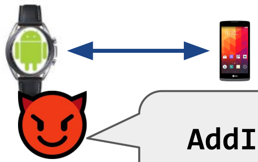

`AddInsulin` pressed on application **`DiabetesM`** !

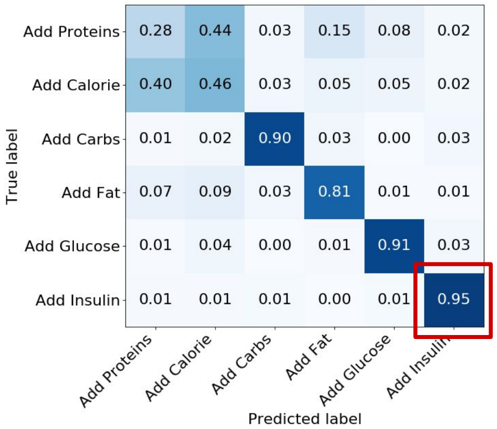

Identifying actions within an application

Methodology:

- 6 actions within the app `DiabetesM`

Result:

- 70% precision / recall / F1 score
- Top features: timings

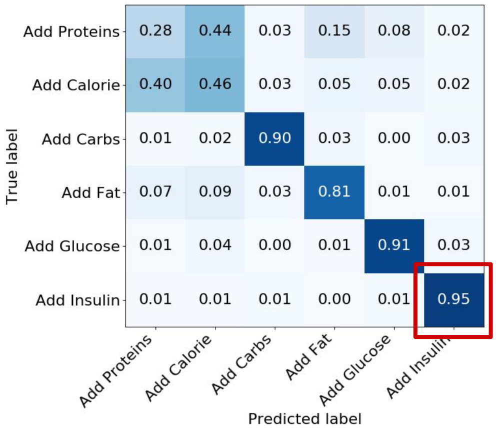

Take-away: some sensitive, medical information are fingerprintable

Highlights of other experiments

- Transferability
  - Train on 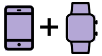 , test on 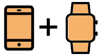 : good performance for Wear OS devices
- Model staleness
  - Small variations in accuracy over 1 month (95% ⟶ 90% mean accuracy for 38 apps)

Negative results (= good news for privacy)

- Audio
    - Phone calls / voice data use constant bit-rate (no “Skype”-like traffic-analysis attack)

- Transferability
    - Android / Apple transferability was unsuccessful

Summary of the attacks

Successful attacks :

- recognize device
- infer user action (from many wearable devices)
- infer opened app (Wear OS)
- infer action within an app (Wear OS)
- attack with model transfer
- attack with “old” dataset
- ...

Unsuccessful attacks :

- voice (phone calls + VoIP)
- transfer Android/Apple

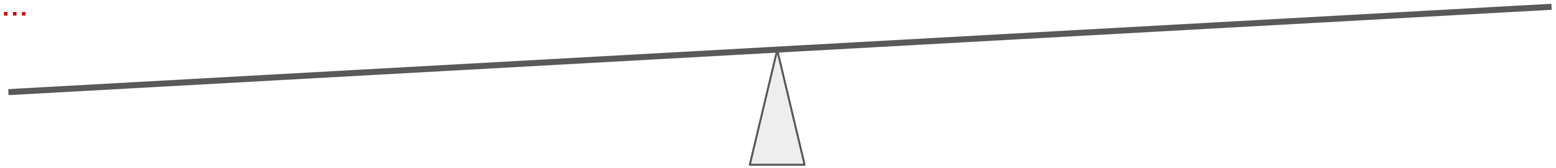

Summary of the attacks

Successful attacks :

- recognize device
- infer user action (from many wearable devices)
- infer opened app (Wear OS)
- infer action within an app (Wear OS)
- attack with model transfer
- attack with “old” dataset
- ...

Unsuccessful attacks :

- voice (phone calls + VoIP)
- transfer Android/Apple

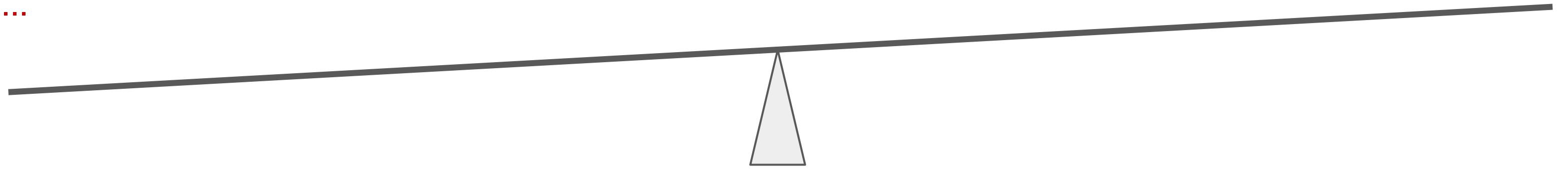

Our conclusion:

- In most cases, sensitive information can be inferred despite the encryption

Summary of the attacks

Successful attacks :

- recognize device
- infer user action (from many wearable devices)
- infer opened app (Wear OS)
- infer action within an app (Wear OS)
- attack with model transfer
- attack with “old” dataset
- ...

Unsuccessful attacks :

- voice (phone calls + VoIP)
- transfer Android/Apple

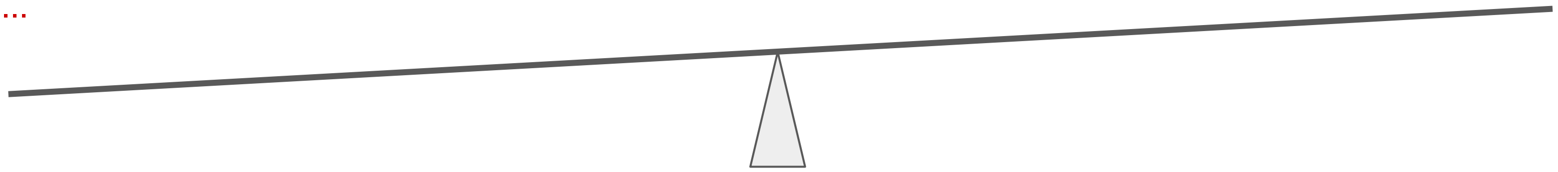

Our conclusion:

- In most cases, sensitive information can be inferred despite the encryption
- Traffic analysis defenses might be required in this setting

Defenses

- we evaluate 3 standard, orthogonal defenses
  - padding, delaying, adding dummy packets

Defenses

- we evaluate 3 standard, orthogonal defenses
  - padding, delaying, adding dummy packets

- Take-aways
  - Defenses are expensive (e.g., 200x data sent) and mildly effective (e.g., -20% accuracy)

Defenses

- we evaluate 3 standard, orthogonal defenses
  - padding, delaying, adding dummy packets

- Take-aways
  - Defenses are expensive (e.g., 200x data sent) and mildly effective (e.g., -20% accuracy)
  - No one-size-fits-all defense:
    - The “wrong” defense does not decrease the accuracy of the adversary

Defenses

- we evaluate 3 standard, orthogonal defenses
  - padding, delaying, adding dummy packets

- Take-aways
  - Defenses are expensive (e.g., 200x data sent) and mildly effective (e.g., -20% accuracy)
  - No one-size-fits-all defense:
    - The “wrong” defense does not decrease the accuracy of the adversary
  - Other valid strategies:
    - data minimization (low-volume apps might be protected)
    - bulk-transfers

Discussion

- No easy fix to the problem
- More awareness is needed
  - We contacted all relevant vendors & app developers with our findings

Discussion

- No easy fix to the problem
- More awareness is needed
  - We contacted all relevant vendors & app developers with our findings
- Limitation: this work is a first quantification/discussion point
- Our hope: better protect the next generation of wearable devices

Brief Highlight

Reducing Metadata Leakage from Static Files

Padmé

(Ch 3)

Padmé

- The **size** is a stable & important feature in traffic analysis

- What is a good *generic* defense ?

- Naïve approaches:
  - constant-block-size padding
  - padding to the next power of two

Intuition behind Padmé

Leakage

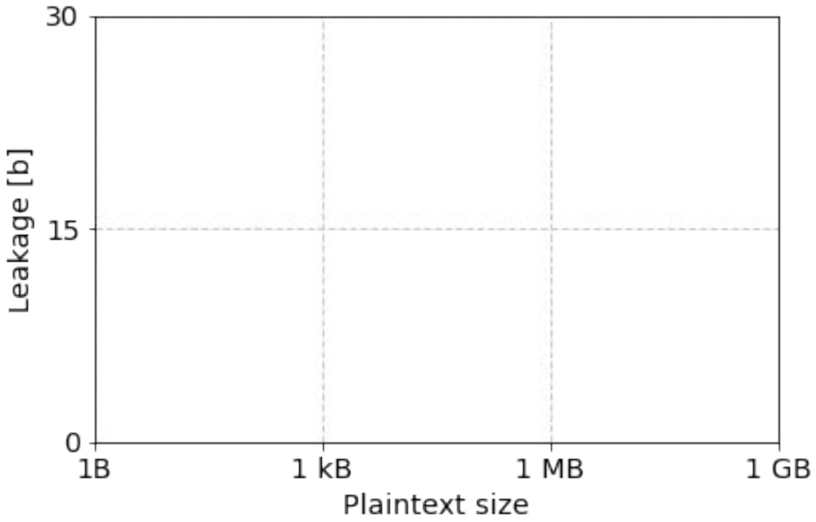

Overhead

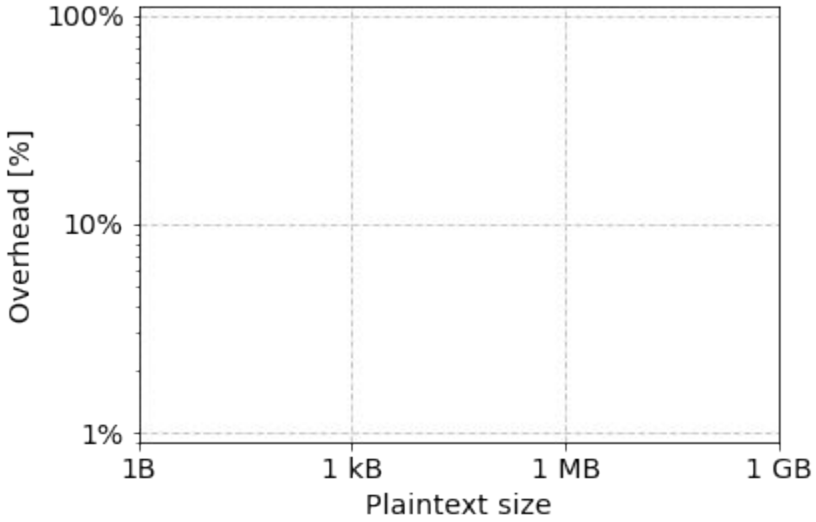

Intuition behind Padmé

Leakage

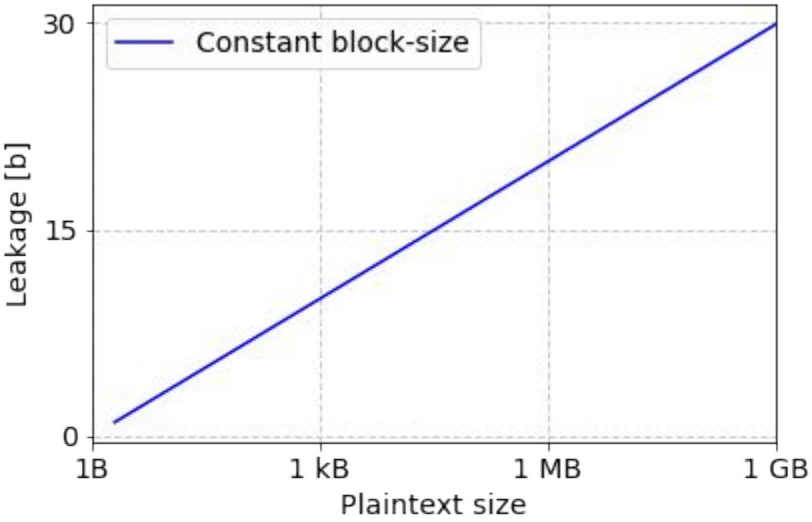

Overhead

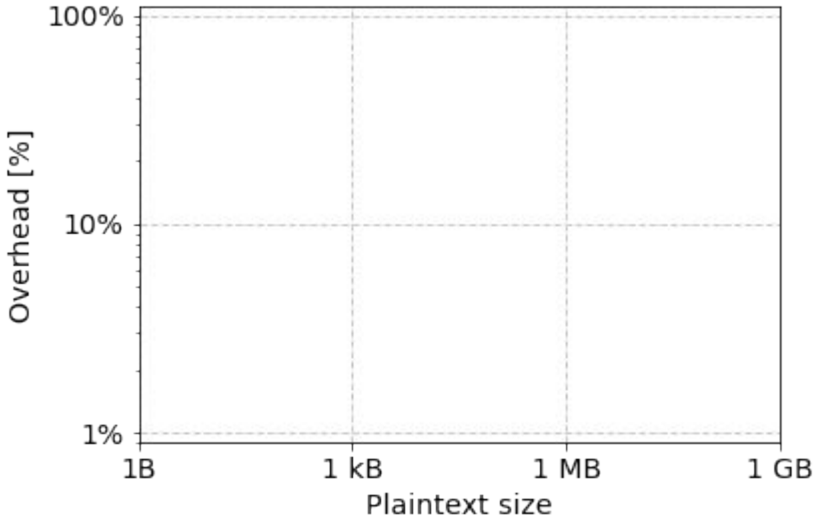

Intuition behind Padmé

Leakage

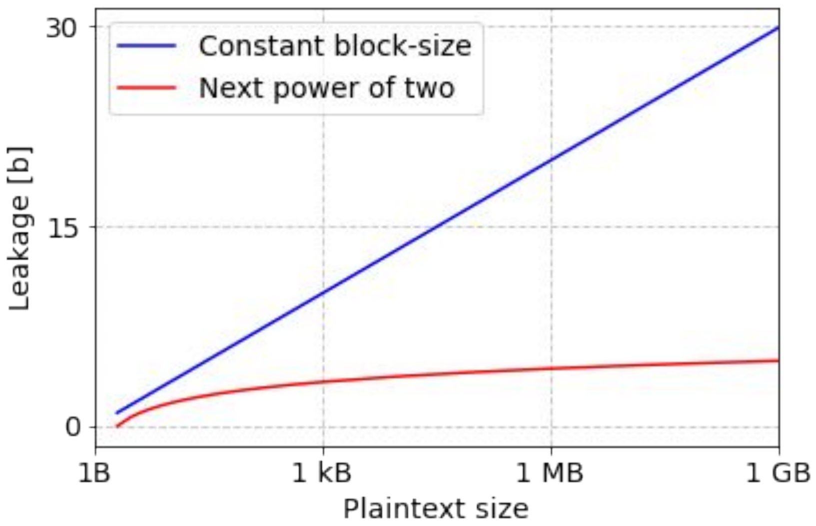

Overhead

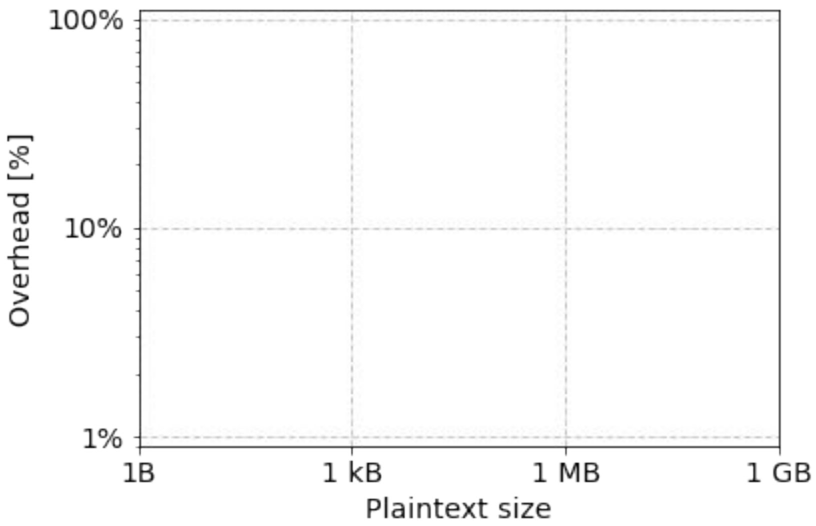

Intuition behind Padmé

Leakage

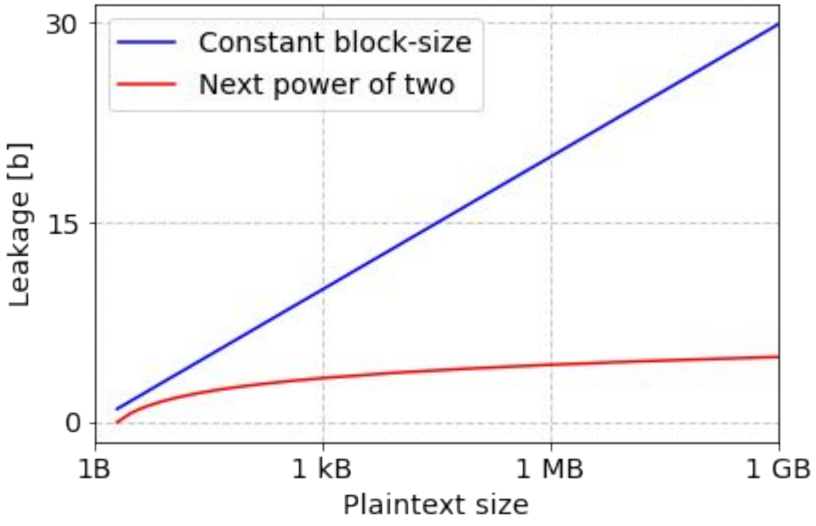

Overhead

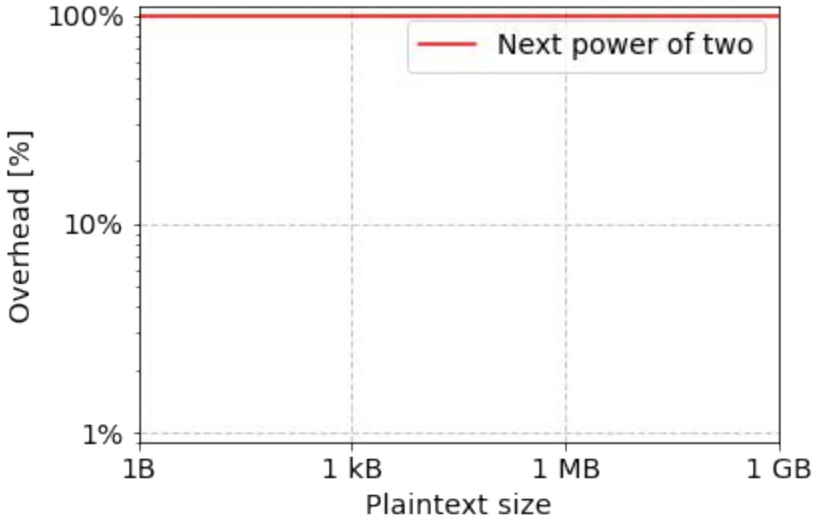

Intuition behind Padmé

Leakage

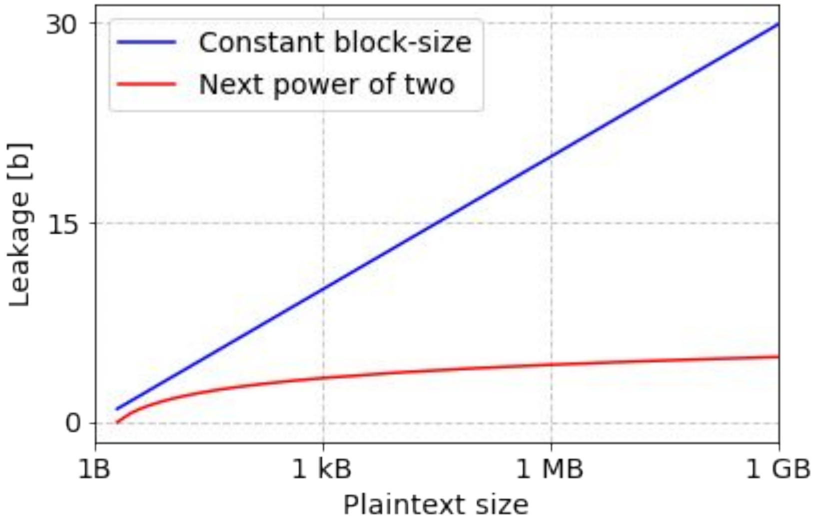

Overhead

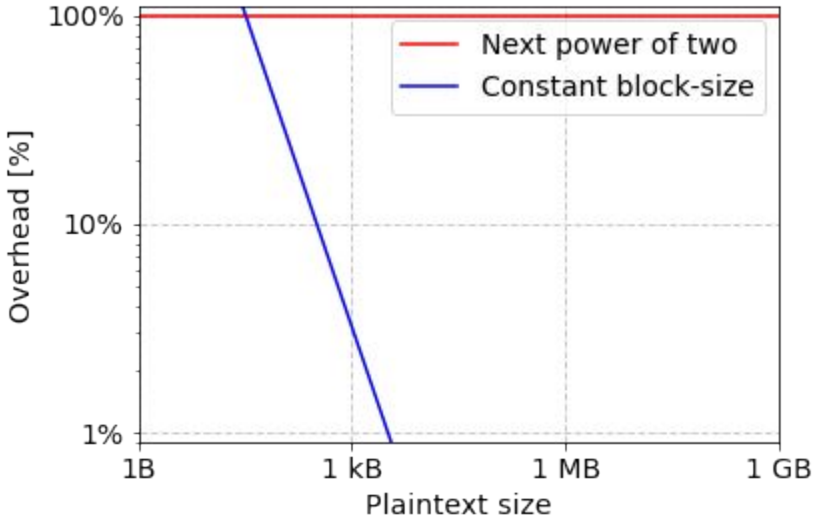

Intuition behind Padmé

Leakage

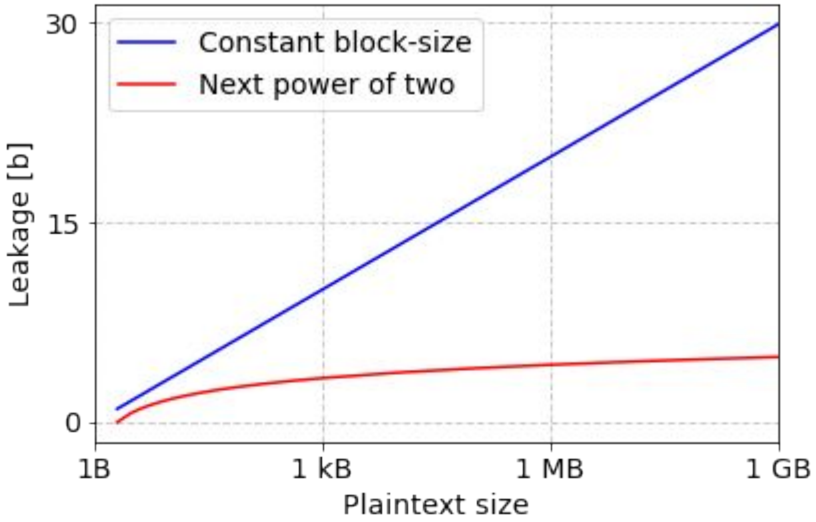

Overhead

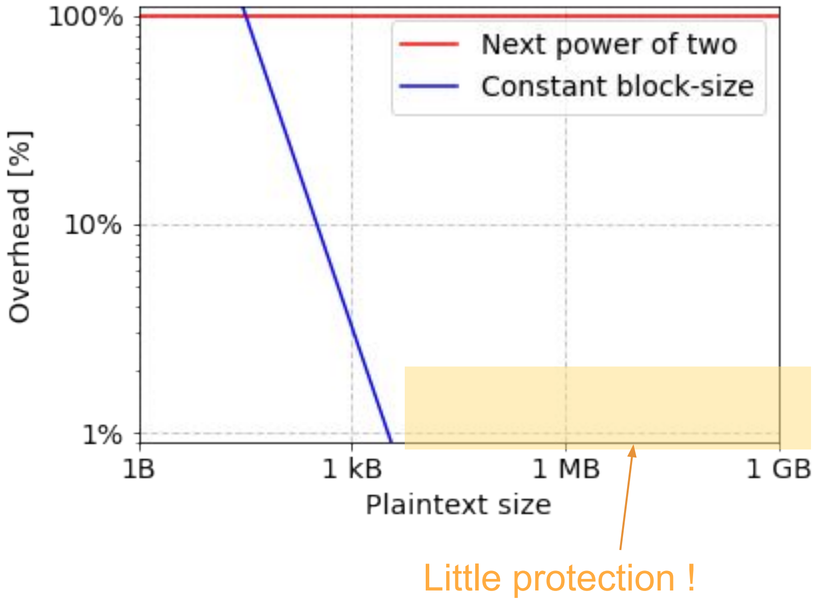

Intuition behind Padmé

Leakage

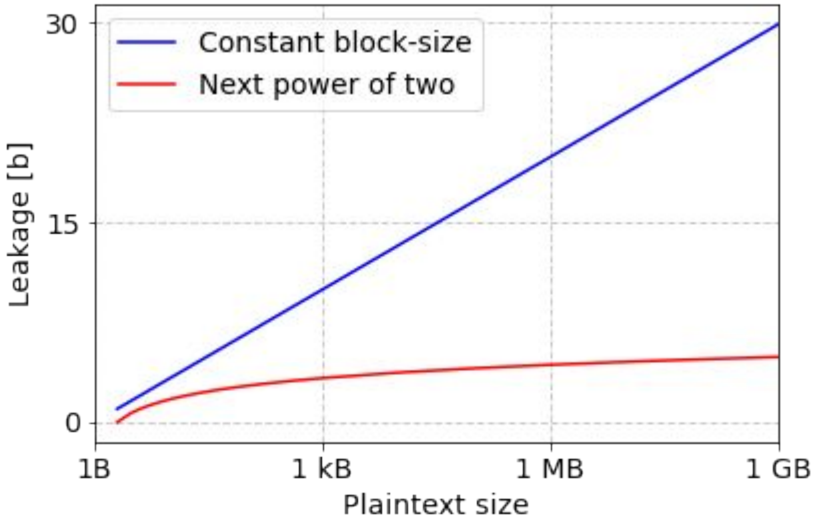

Overhead

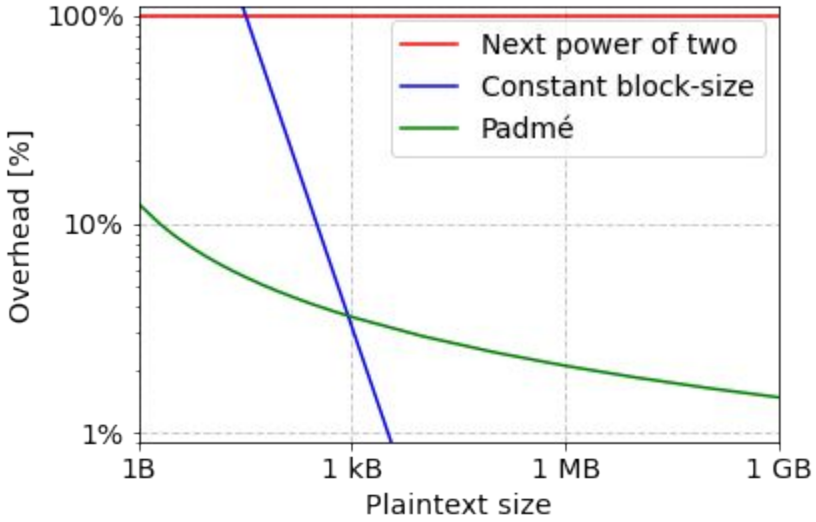

Insight: a slowly-decreasing overhead is more practical

Intuition behind Padmé

Leakage

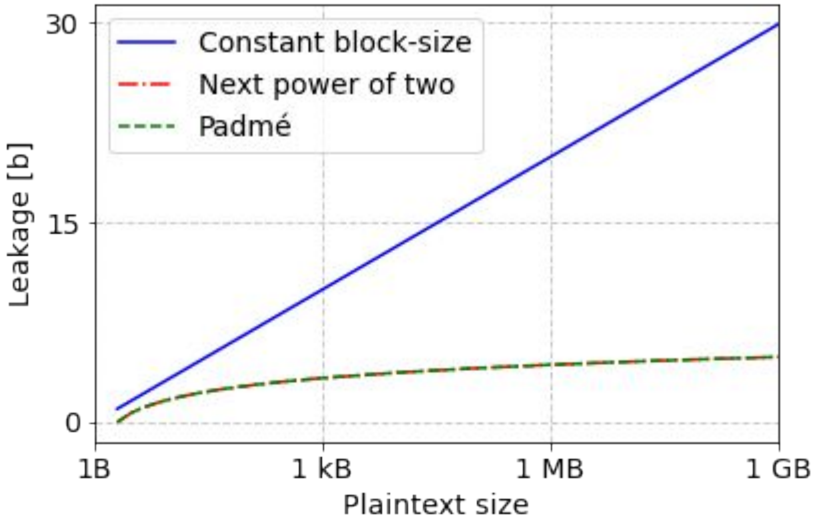

Overhead

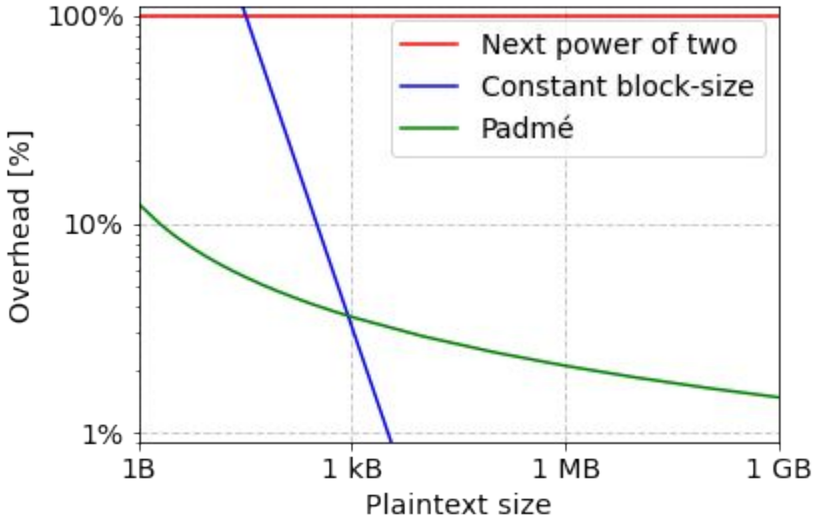

Take-away: same leakage as next power of 2

Padmé

Overhead:

- max +12% ∀L
- max +6% ∀L > 1 MB
- max +3% ∀L > 1 GB

Padmé

Overhead:

- max +12% ∀L
- max +6% ∀L > 1 MB
- max +3% ∀L > 1 GB

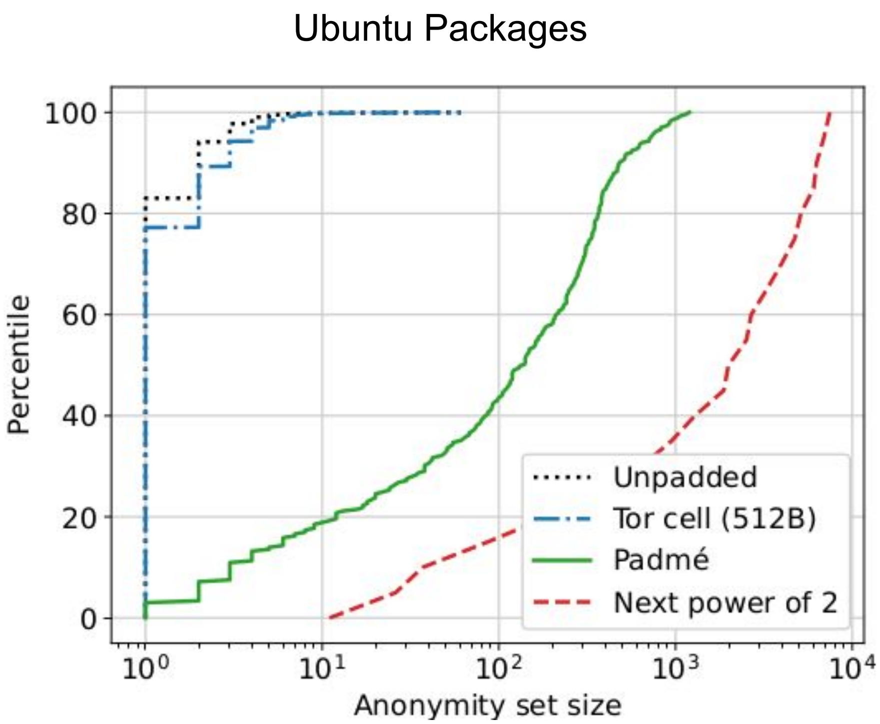

Take-away: low overhead + good hiding properties

Brief Highlight

PriFi

A low-latency ACN for LANs and WLANs

(§4.4)

PriFi

- Problem:
  - Risk of targeted attacks in loosely trusted, sensitive WLANs (e.g., NGOs)

- Goal:
  - Hide the traffic of key individuals

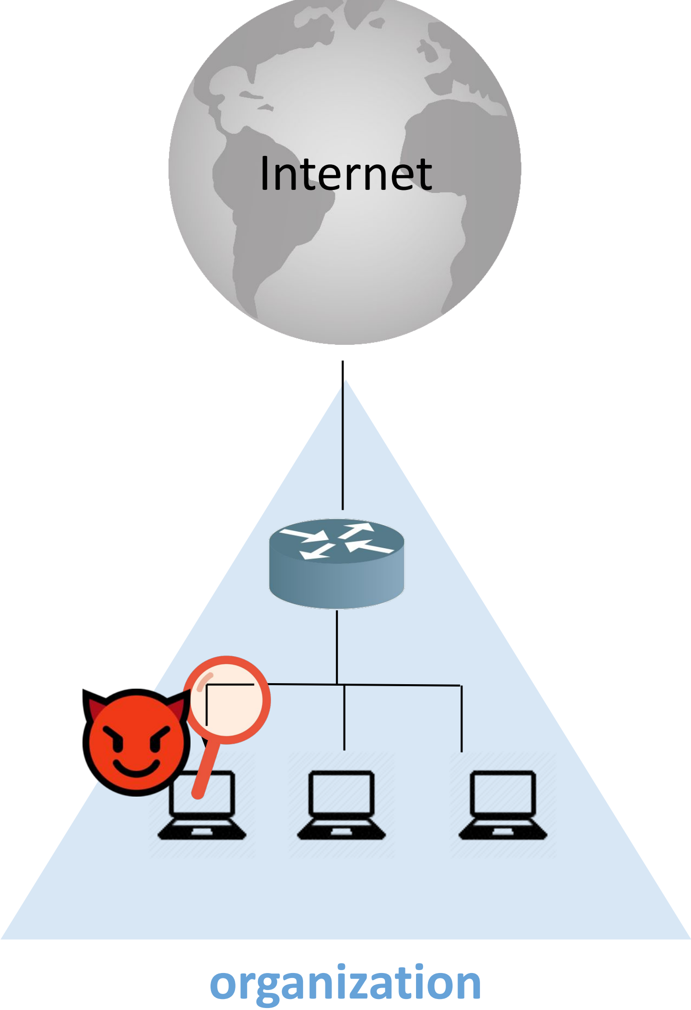

ACNs are poorly suited to LANs

- Tor / Mixnets add extra hops = extra latency
- Traffic leaves the organization

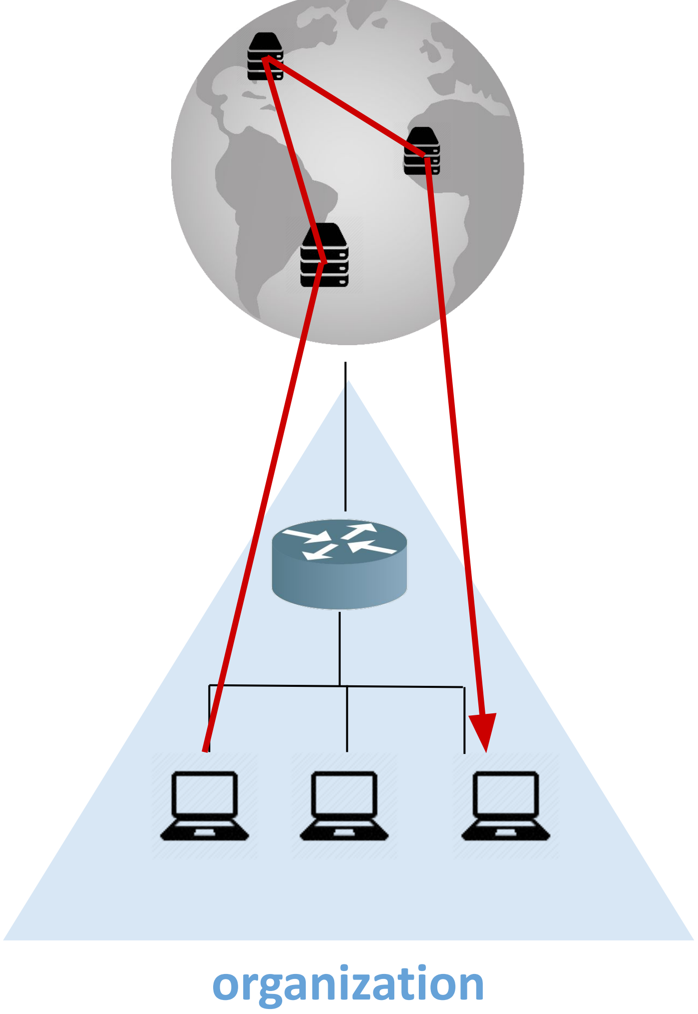

ACNs are poorly suited to LANs

- DC-nets can avoid this
- In practice, they don't [10]
- At each round, "chatty" protocol with the servers [10]

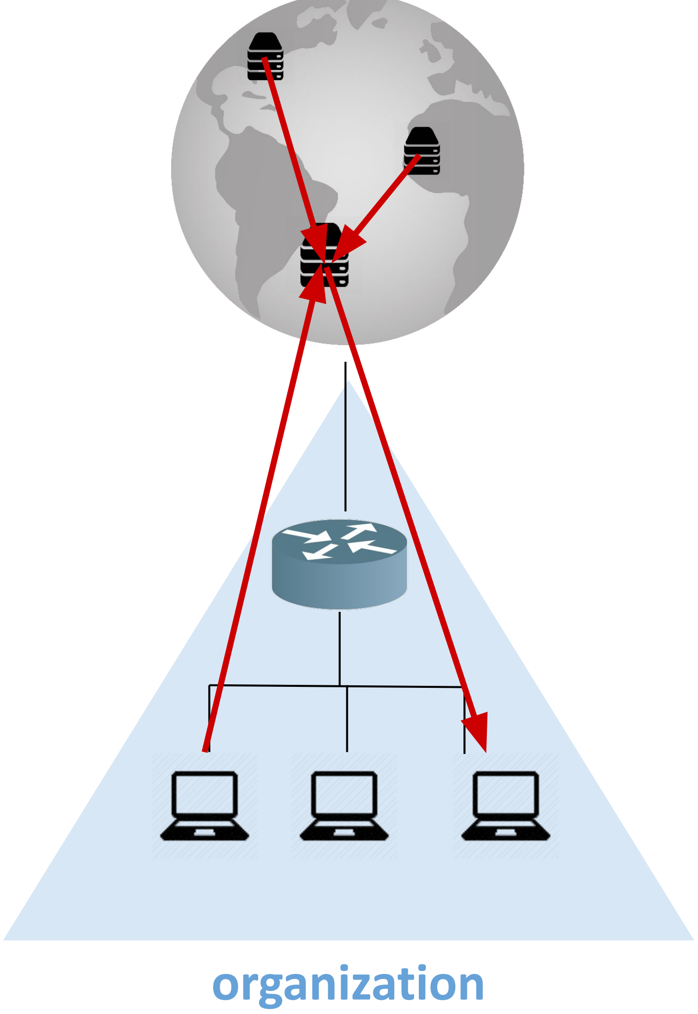

[10] D. Wolinsky, H. Corrigan-Gibbs, B. Ford. **Dissent in Numbers: Making strong anonymity scale.** OSDI 2012

PriFi

- New topology for DC-nets
- Redesign of the protocols
  - servers contributions are sent in advance
  - avoid server-to-server messages

=> Latency to the servers is not important

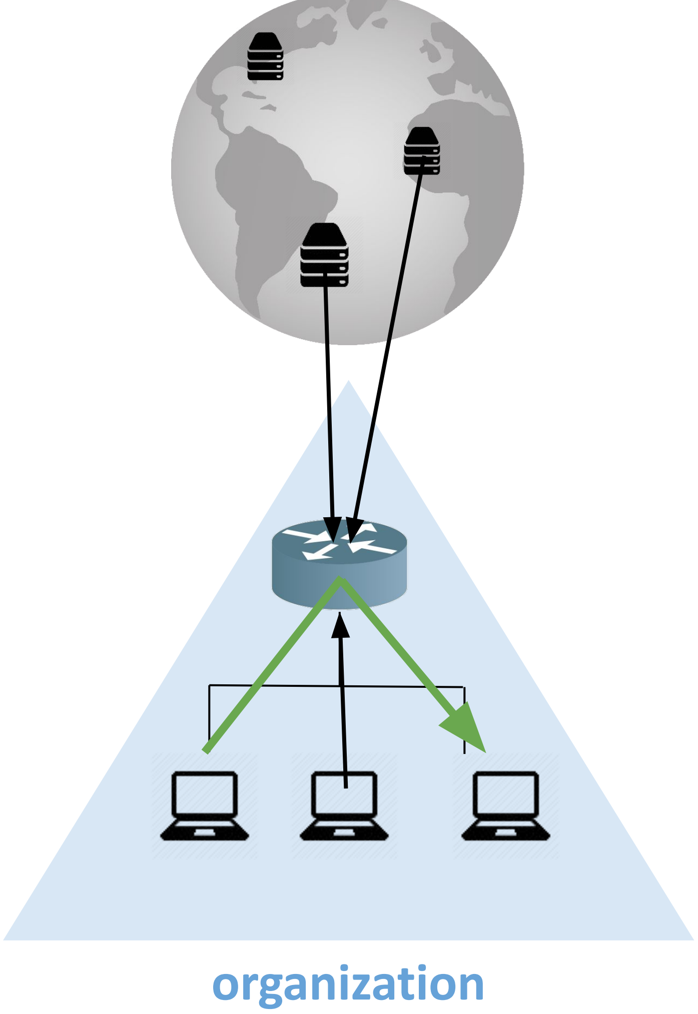

PriFi

- New topology for DC-nets
- Redesign of the protocols
  - servers contributions are sent in advance
  - avoid server-to-server messages

=> Latency to the servers is not important

=> “on-path” anonymity

=> cheap broadcast in WLANs

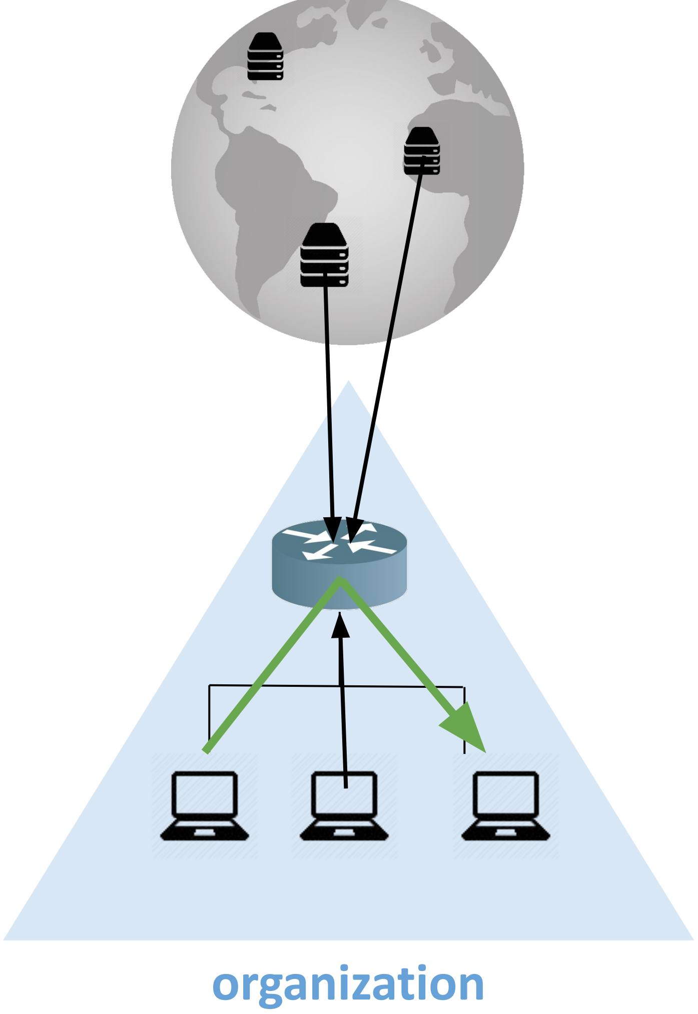

PriFi

- New topology for DC-nets
- Redesign of the protocols
  - servers contributions are sent in advance
  - avoid server-to-server messages

=> Latency to the servers is not important

=> “on-path” anonymity

=> cheap broadcast in WLANs

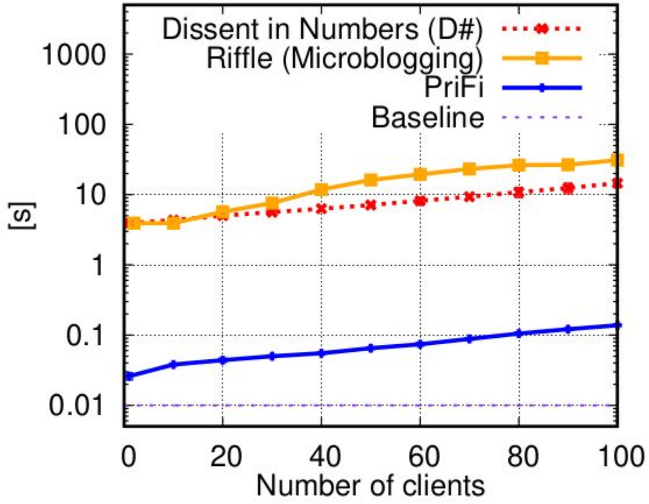

Rubato: Metadata-Private Communications for Mobile Devices

(§4.5)

System for text communication on phones

System for text communication on phones

Model

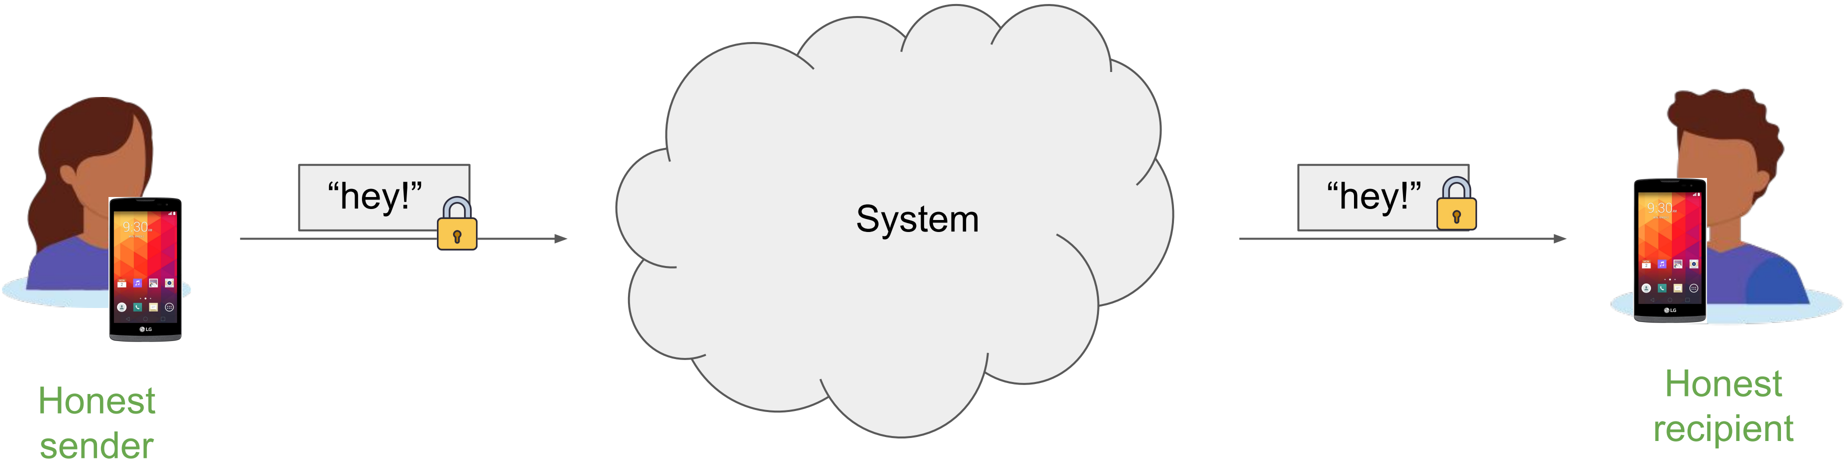

Global Active Adversary

- observes all network communication
- can edit/drop/inject any message
- controls a fraction of the entities

Model

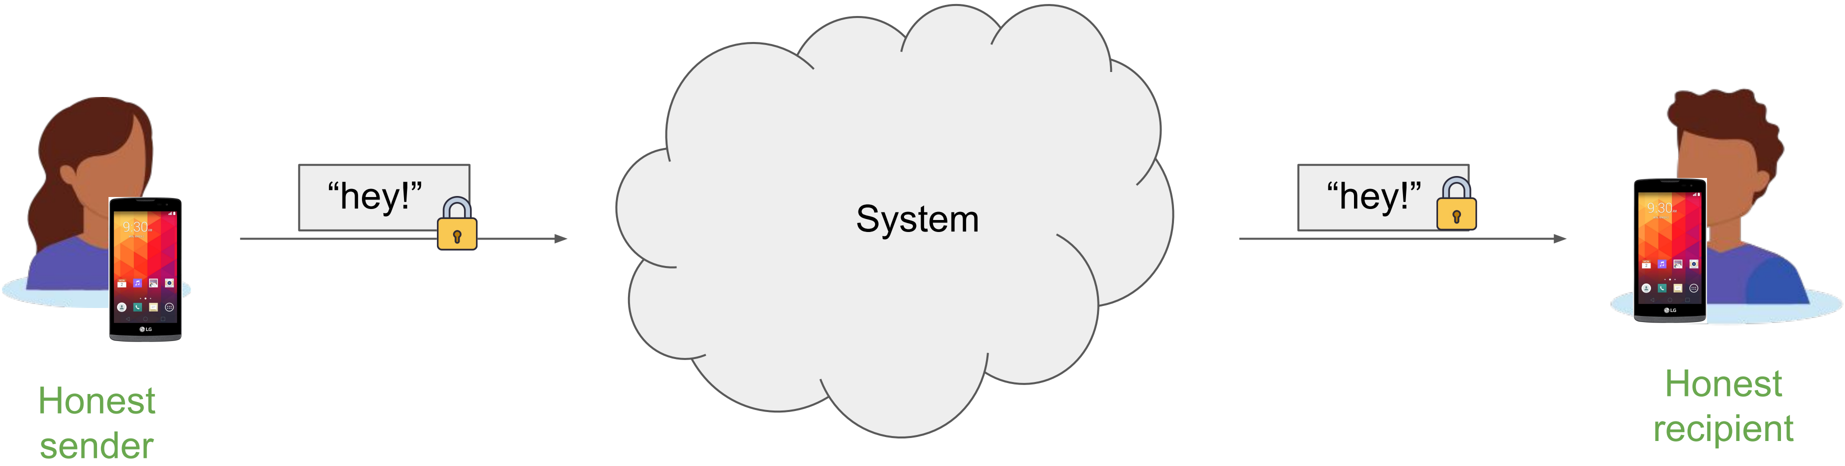

Global Active Adversary

- observes all network communication
- can edit/drop/inject any message
- controls a fraction of the entities

Security notion:

$$\mathtt{obs}_A(\mathtt{Alice} \leftrightarrow \mathtt{Bob}) \cong \mathtt{obs}_A(\mathtt{Alice} \not\leftrightarrow \mathtt{Bob})$$

Current deployed systems are unsafe

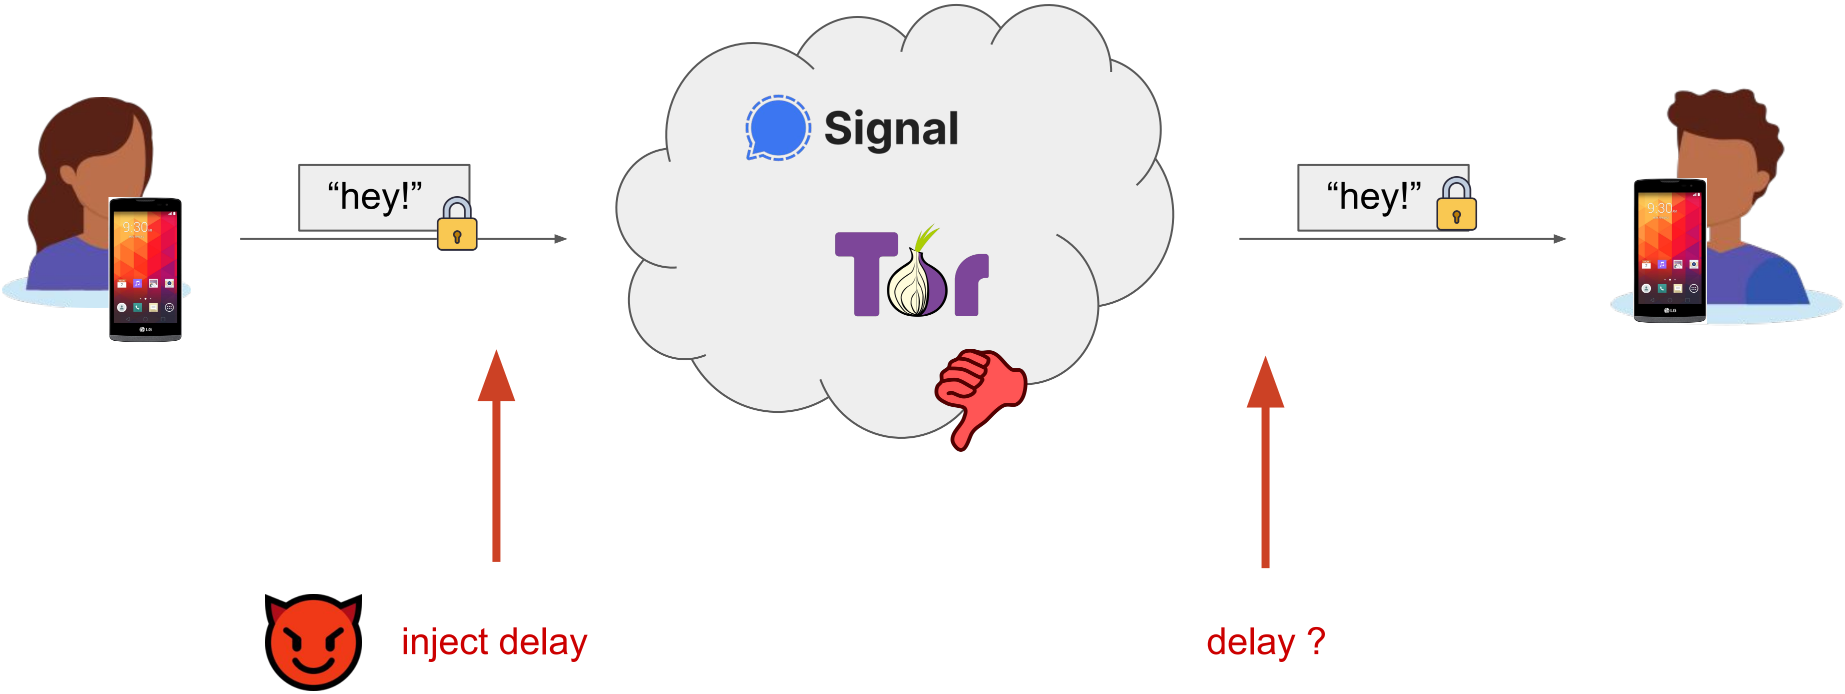

Loopix / Miranda [11,12]

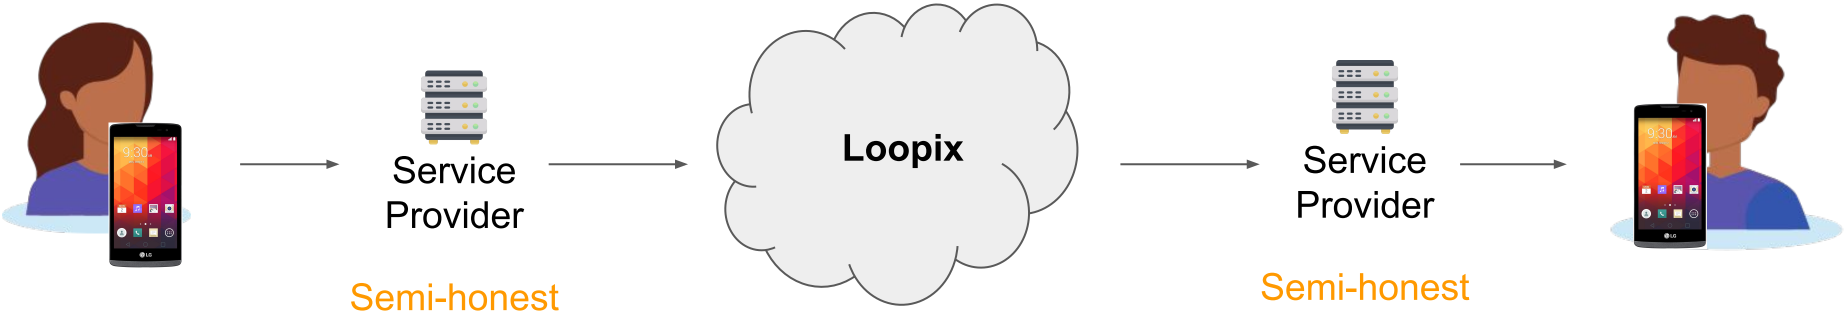

[11] A. Piotrowska, J. Hayes, T. Elahi, S. Meiser, G. Danezis, **The Loopix anonymity system.** Usenix Security 2017.

[12] H. Leibowitz, A. Piotrowska, G. Danezis, A. Herzberg. **No right to remain silent: isolating malicious mixes.** Usenix Security 2019.

Loopix / Miranda [11,12]

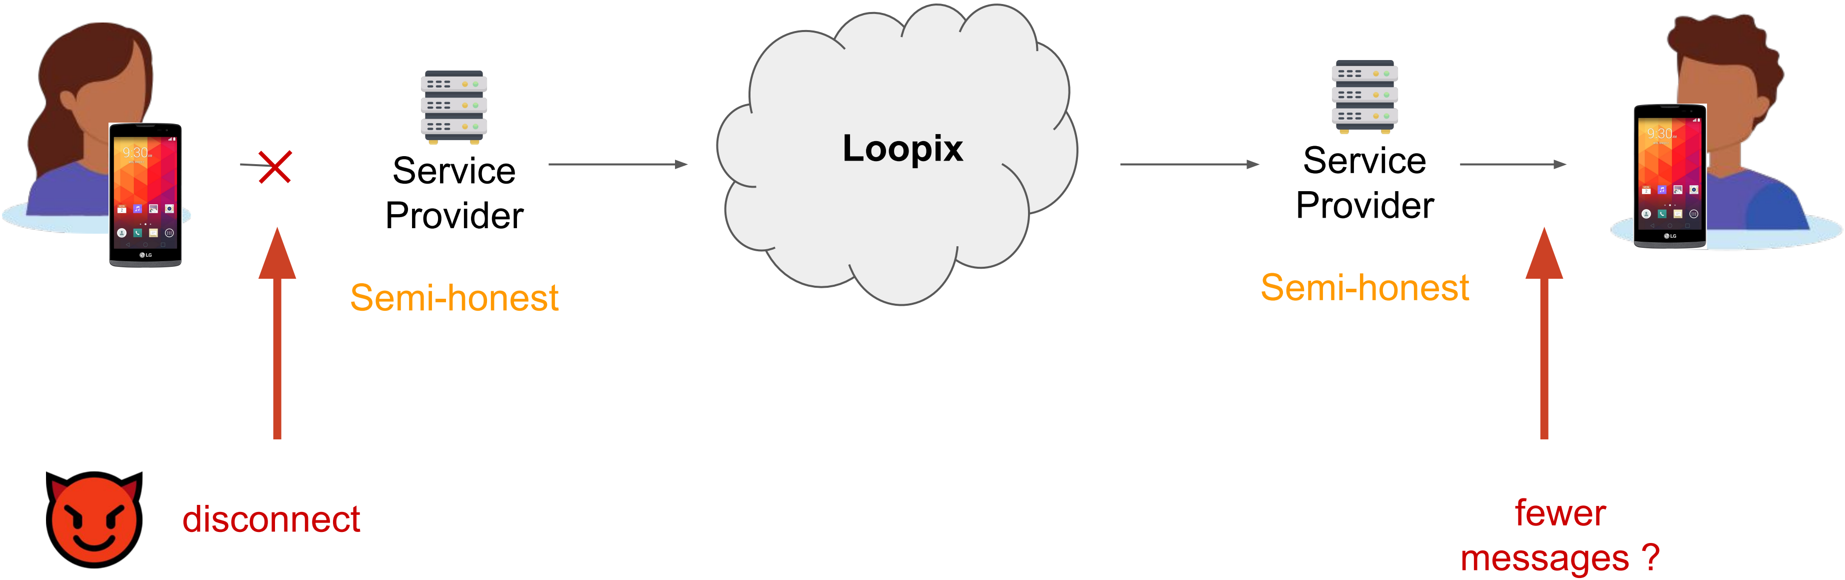

[11] A. Piotrowska, J. Hayes, T. Elahi, S. Meiser, G. Danezis, **The Loopix anonymity system.** Usenix Security 2017.

[12] H. Leibowitz, A. Piotrowska, G. Danezis, A. Herzberg. **No right to remain silent: isolating malicious mixes.** Usenix Security 2019.

Mixnets with constant-rate communications

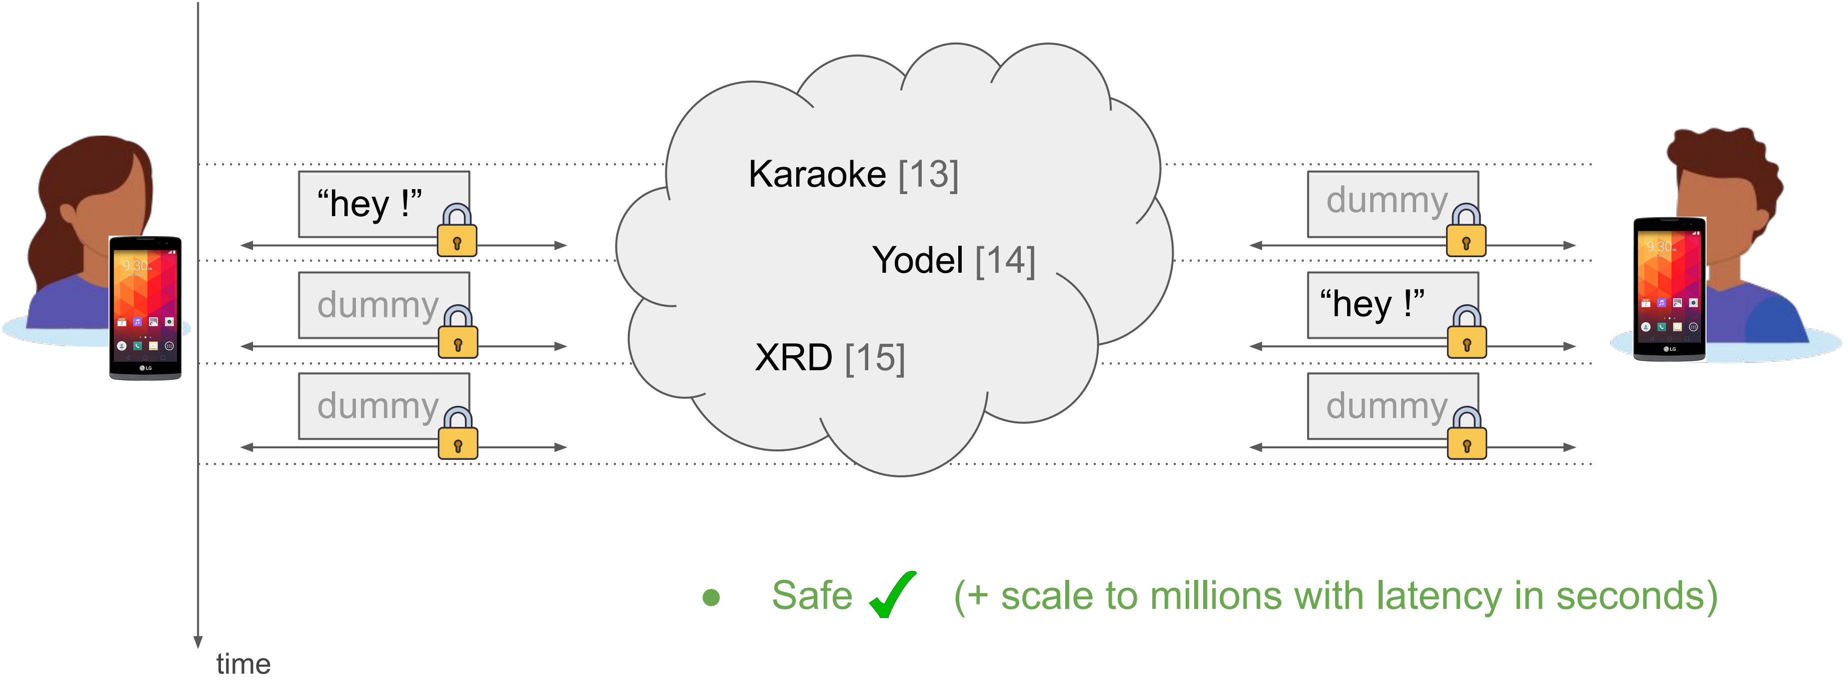

- Safe ✓ (+ scale to millions with latency in seconds)

[13] D. Lazar, Y. Gilad, N. Zeldovich. **Karaoke: Distributed Private Messaging Immune to Passive Traffic Analysis.** OSDI 2018

[14] D. Lazar, Y. Gilad, N. Zeldovich. **Yodel: strong metadata security for voice calls.** SOSP 2019

[15] A. Kwon, D. Lu, S. Devadas. **XRD: Scalable Messaging System with Cryptographic Privacy.** NSDI 20

State of the art: mixnets with constant-rate communications

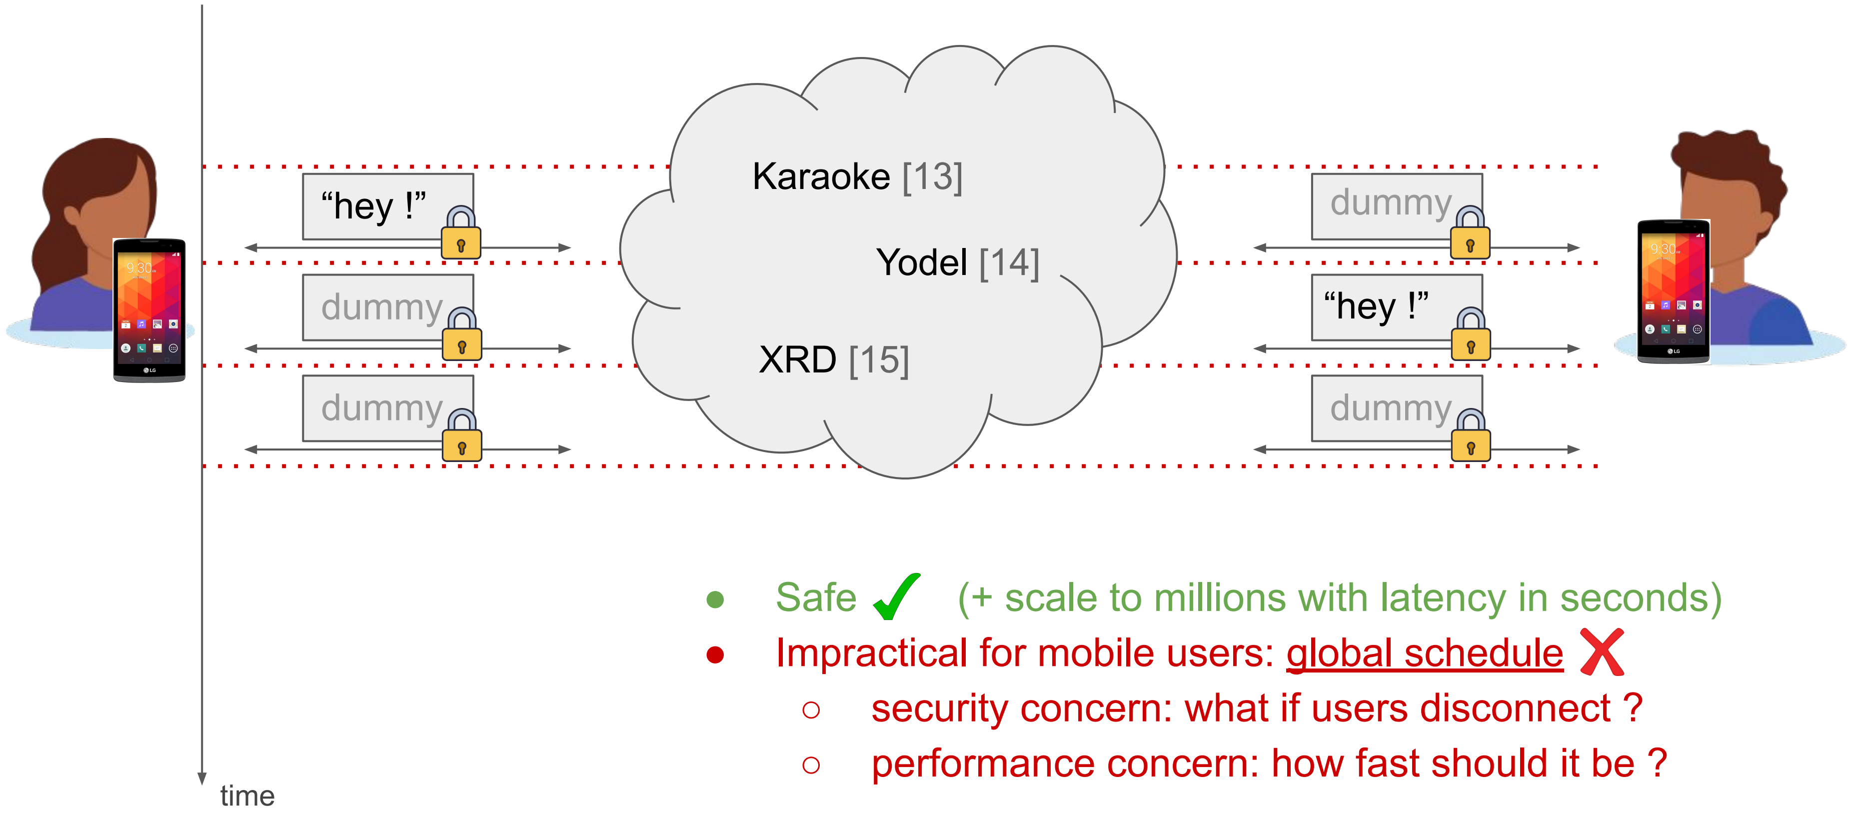

- Safe ✓ (+ scale to millions with latency in seconds)
- Impractical for mobile users: <u>global schedule</u> ✗
  - security concern: what if users disconnect ?
  - performance concern: how fast should it be ?

Mobile devices require asynchronicity

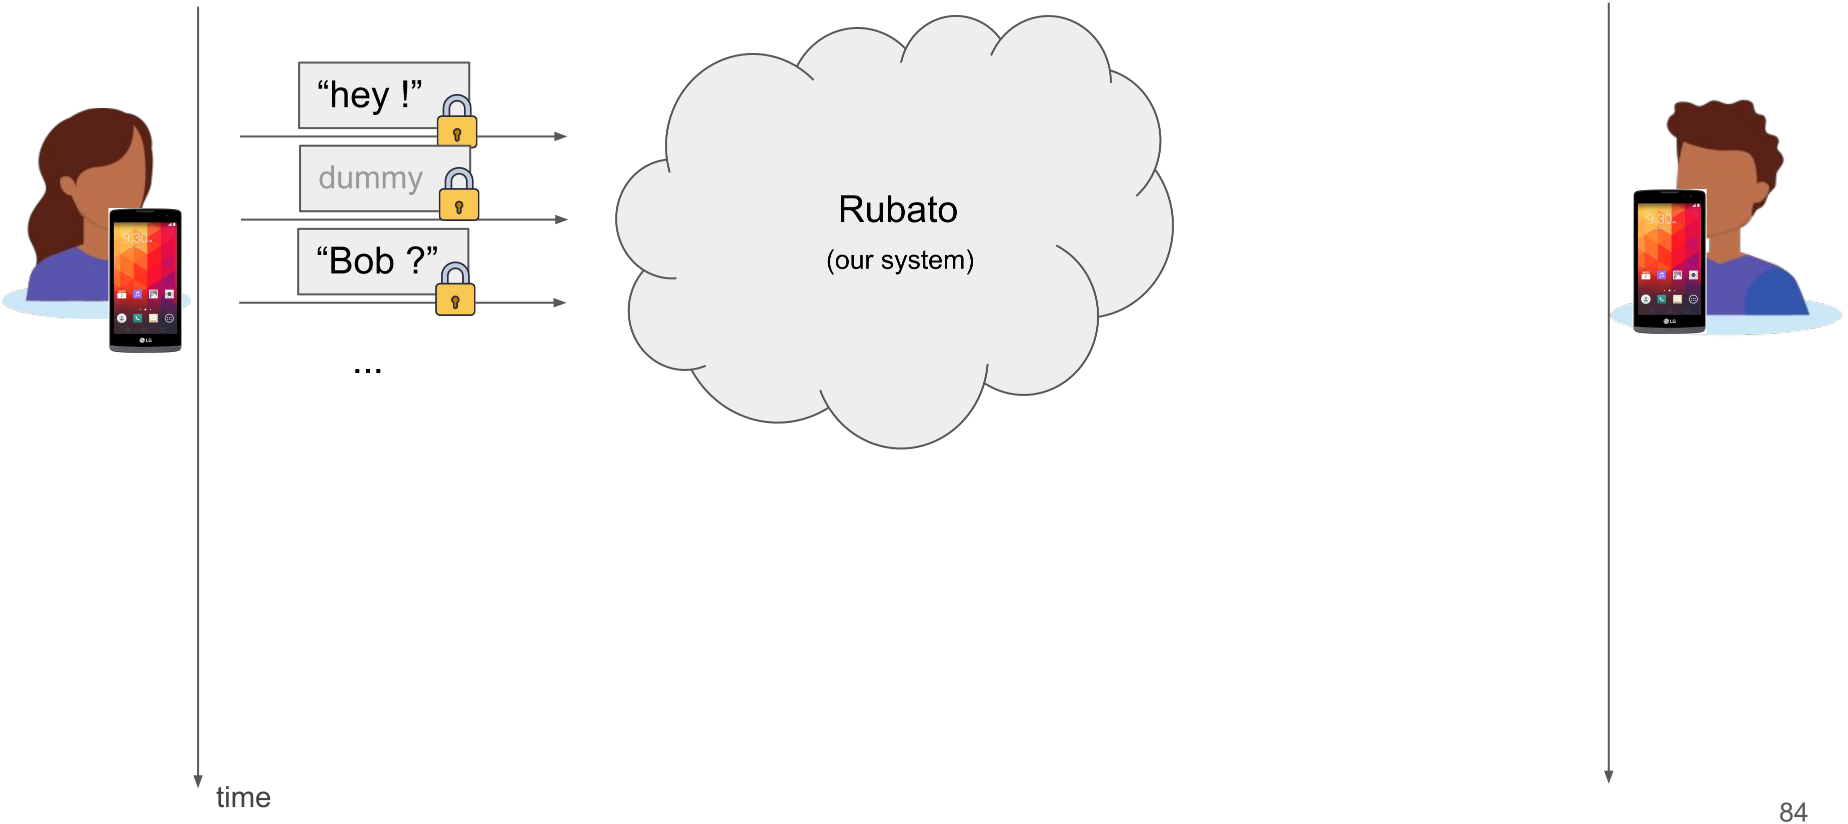

Mobile devices require asynchronicity

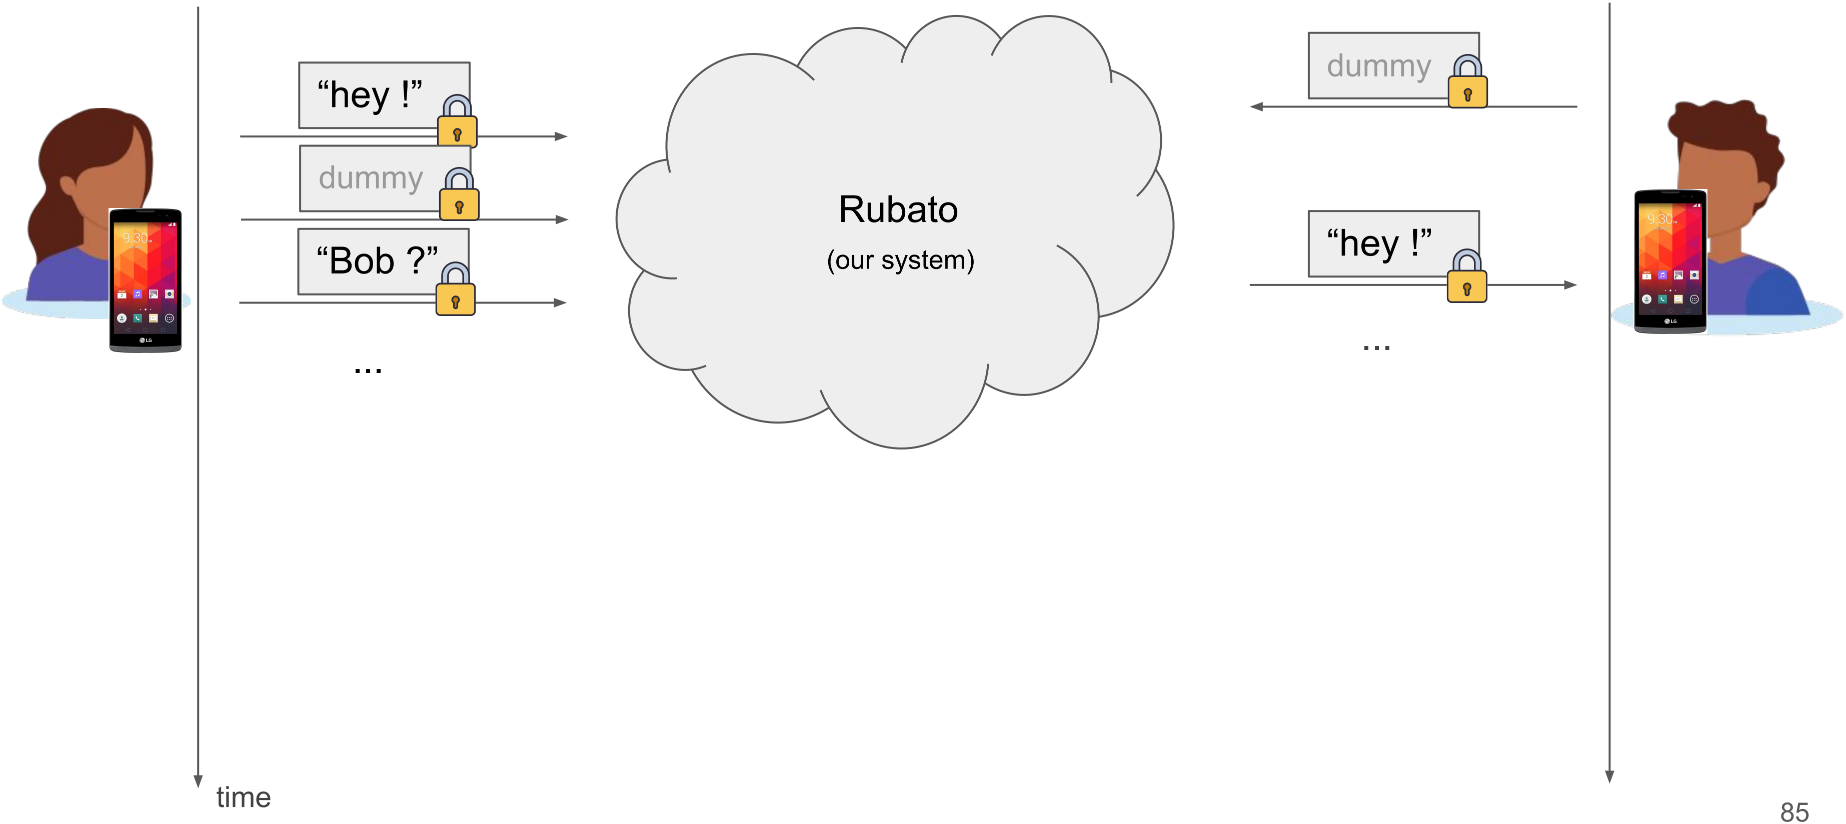

Mobile devices require asynchronicity

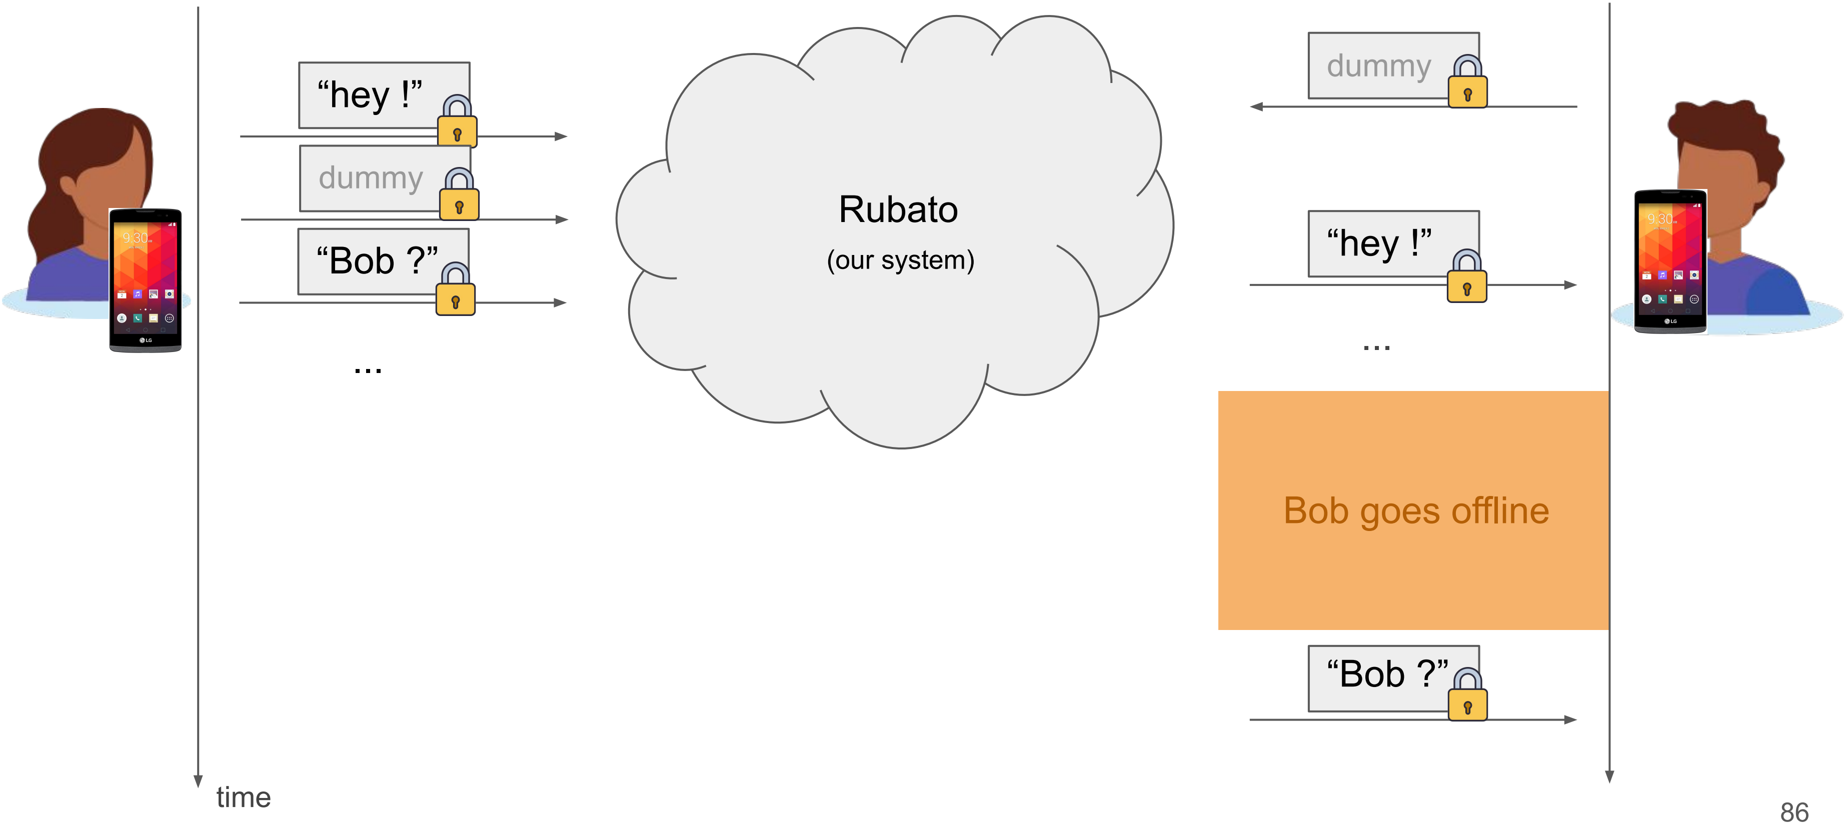

Mobile devices require asynchronicity

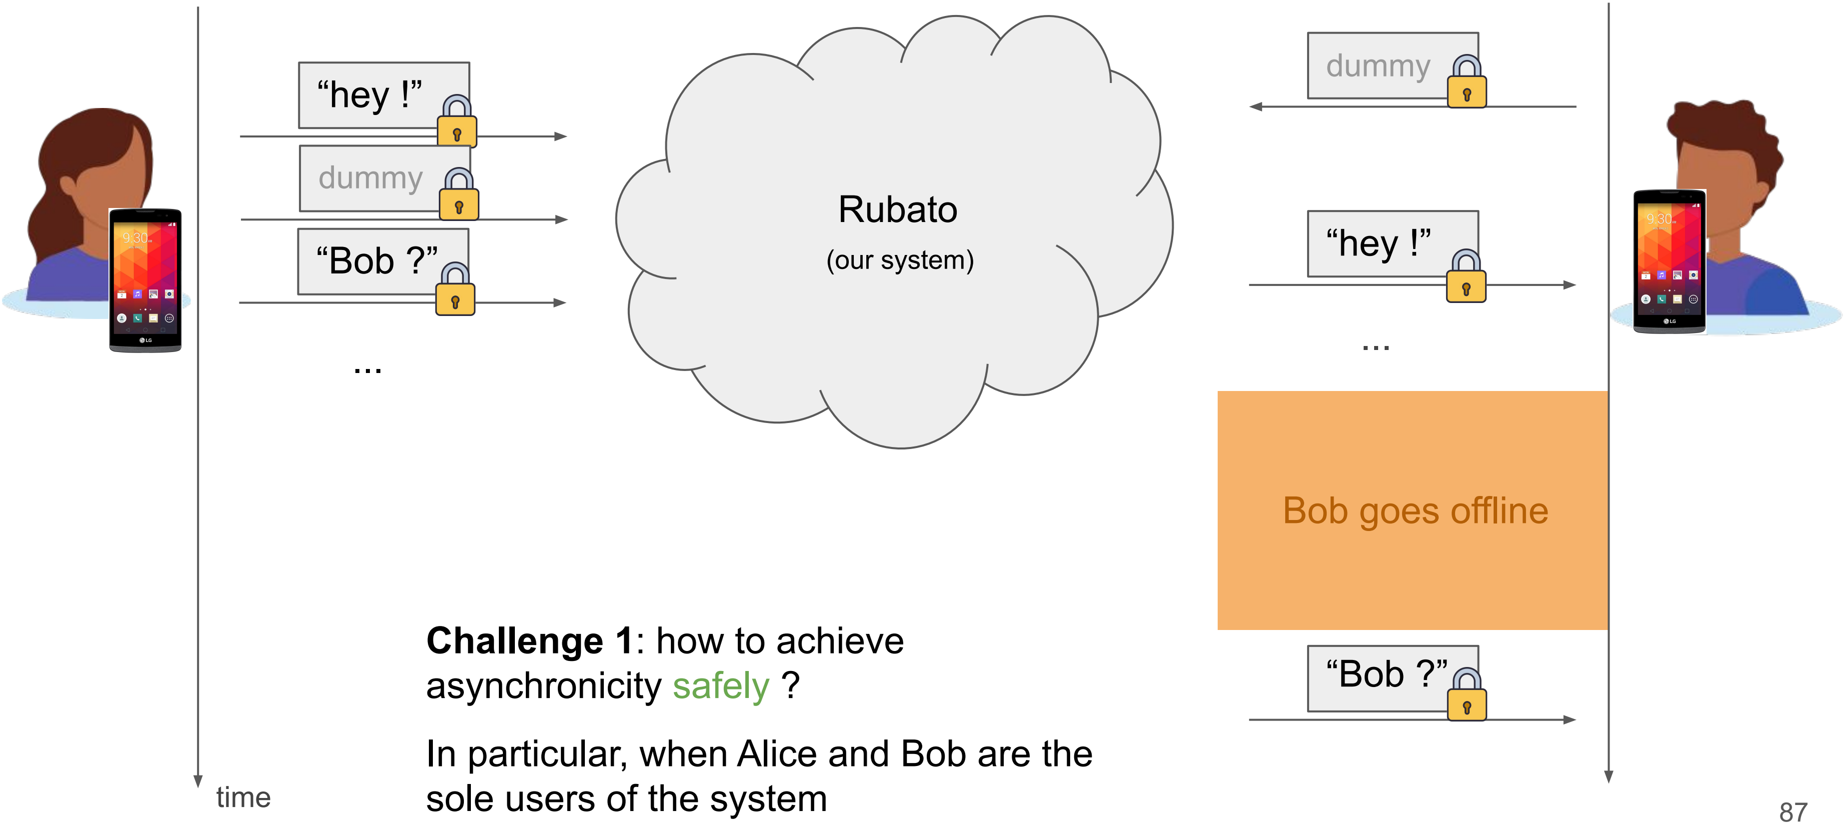

Challenge 1: how to achieve asynchronicity safely ?

In particular, when Alice and Bob are the sole users of the system

Messaging with multiple independent devices

In practice, users have multiple devices!

The adversary can:

- Prevent synchronisation
- Equivocate

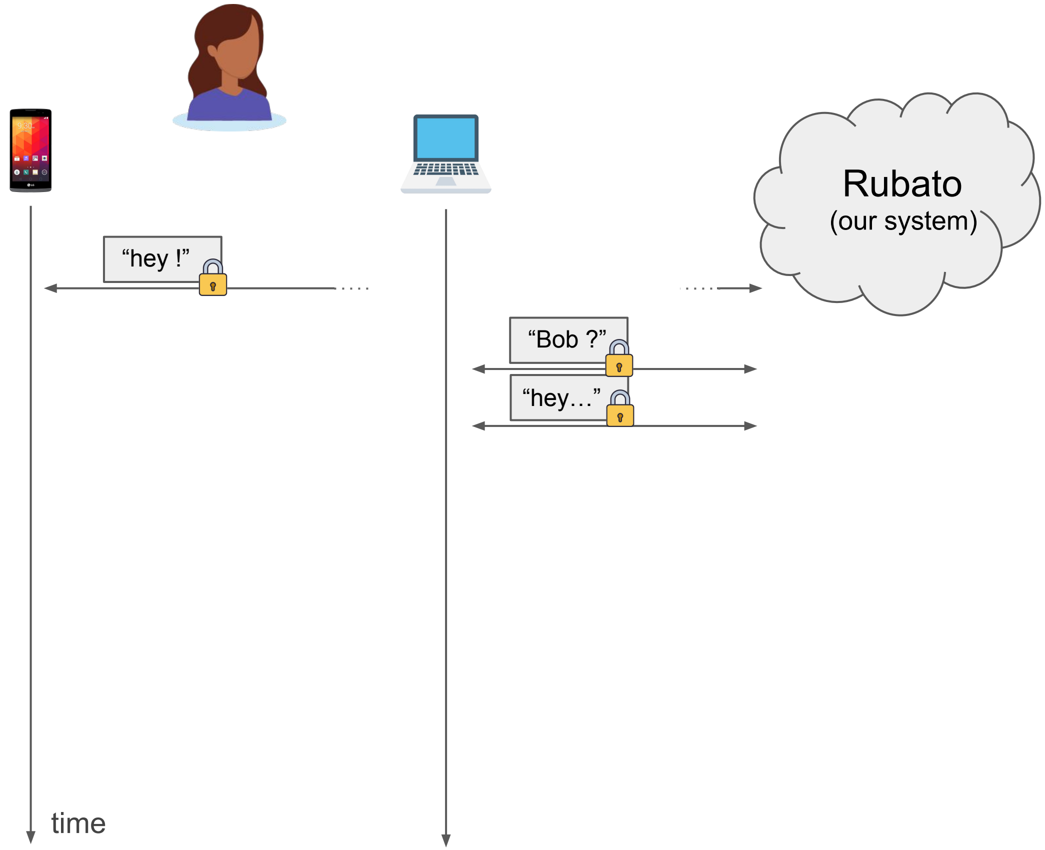

Messaging with multiple independent devices

In practice, users have multiple devices!

The adversary can:

- Prevent synchronisation
- Equivocate
- Disconnect devices

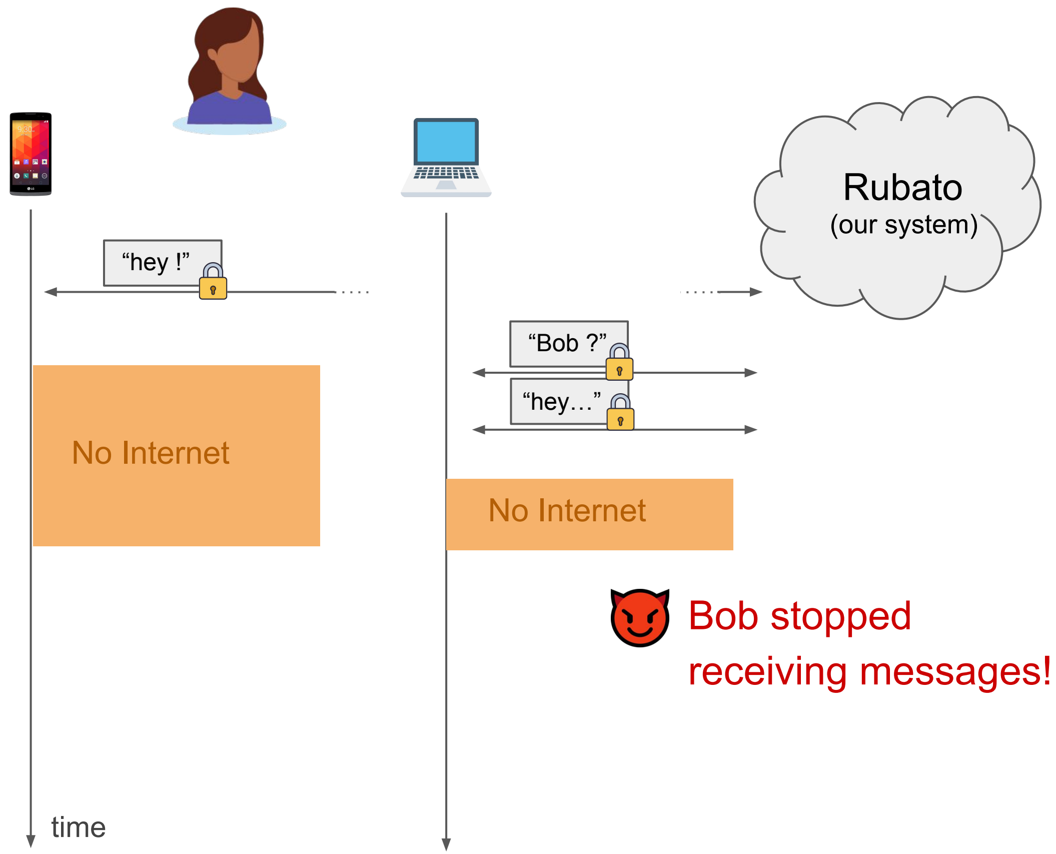

Messaging with multiple independent devices

In practice, users have multiple devices!

The adversary can:

- Prevent synchronisation
- Equivocate
- Disconnect devices
- Partition devices and observe more messages than intended

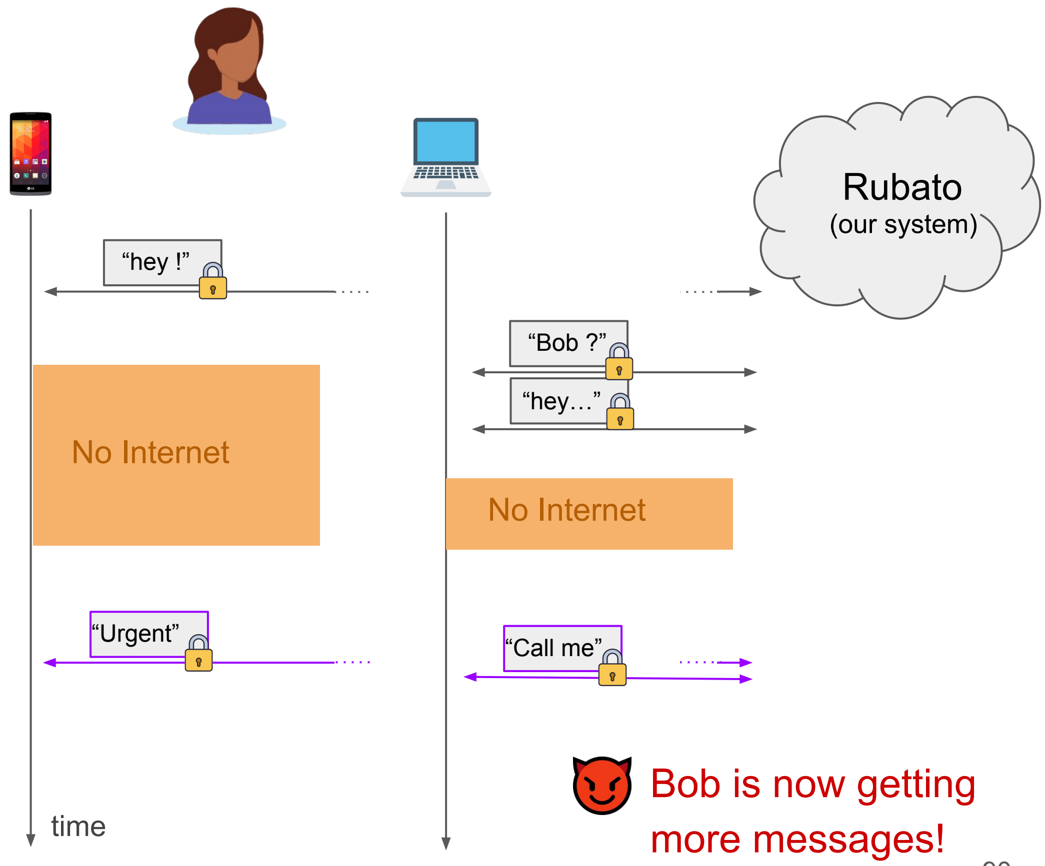

Messaging with multiple independent devices

In practice, users have multiple devices!

The adversary can:

- Prevent synchronisation
- Equivocate
- Disconnect devices
- Partition devices and observe more messages than intended

Challenge 2: how to support multiple independent, asynchronous devices safely ?

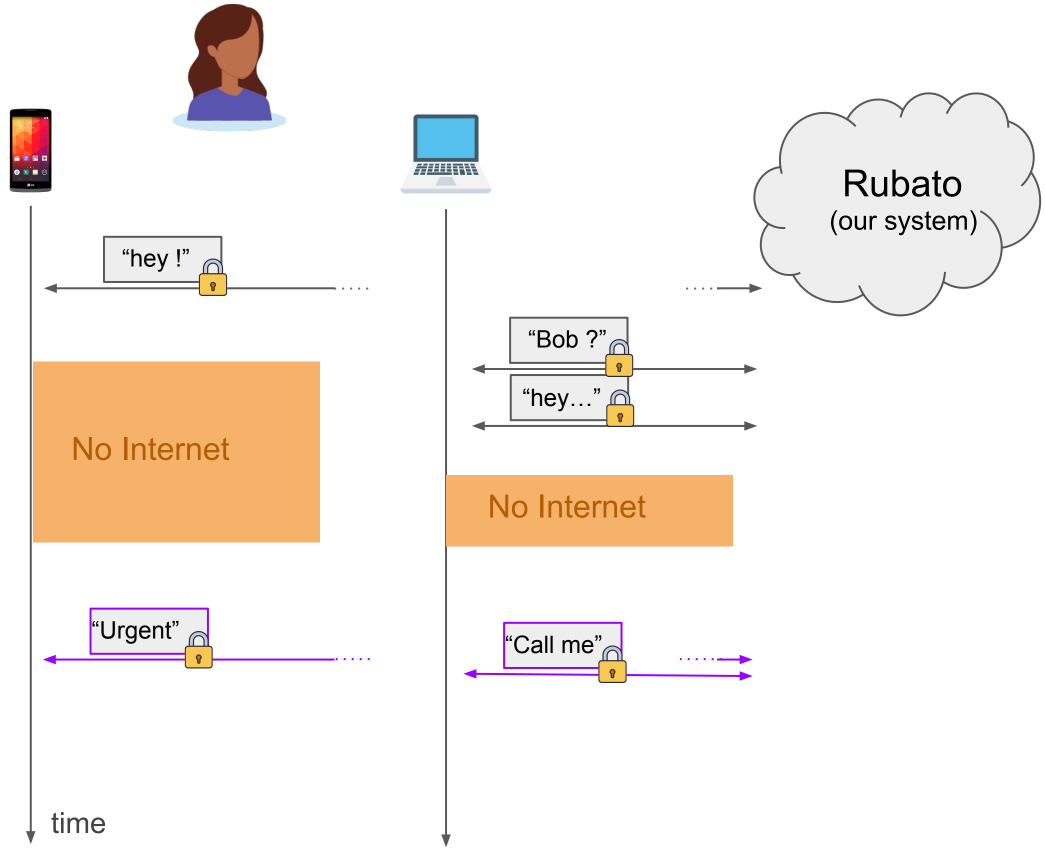

Rubato

- Rubato is a large-scale ACN for text communications
- It advances the state-of-the-art...
  - Multi-devices (that only synchronize through the untrusted network)
  - Devices can have their own communication patterns
  - … and thus it better supports mobile devices.
- … by using new techniques:
  - “Primed” circuits through a mixnet
  - Path selection across devices, Circuit tagging techniques
  - Efficient “Fetch” protocol (not presented)

System

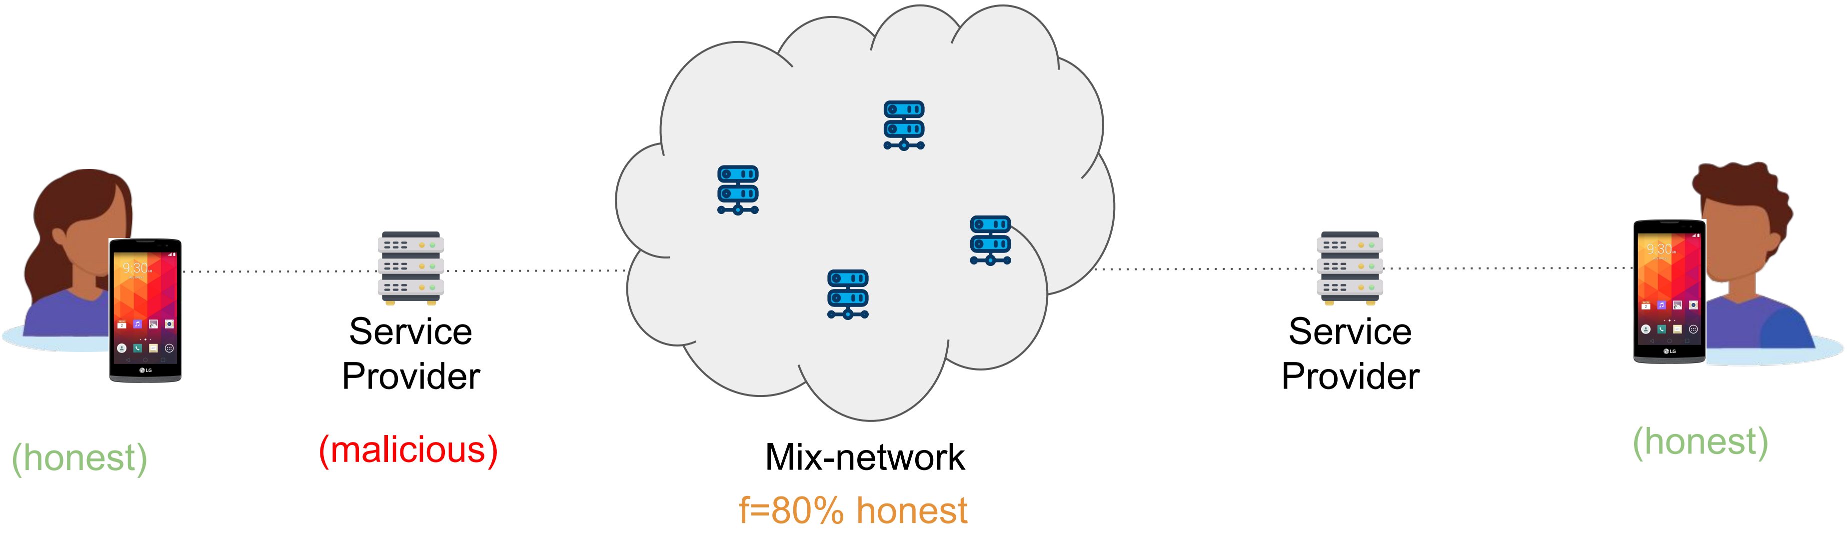

The Service Provider (SP) buffers messages from and to the synchronous mixnet

Primed Circuits

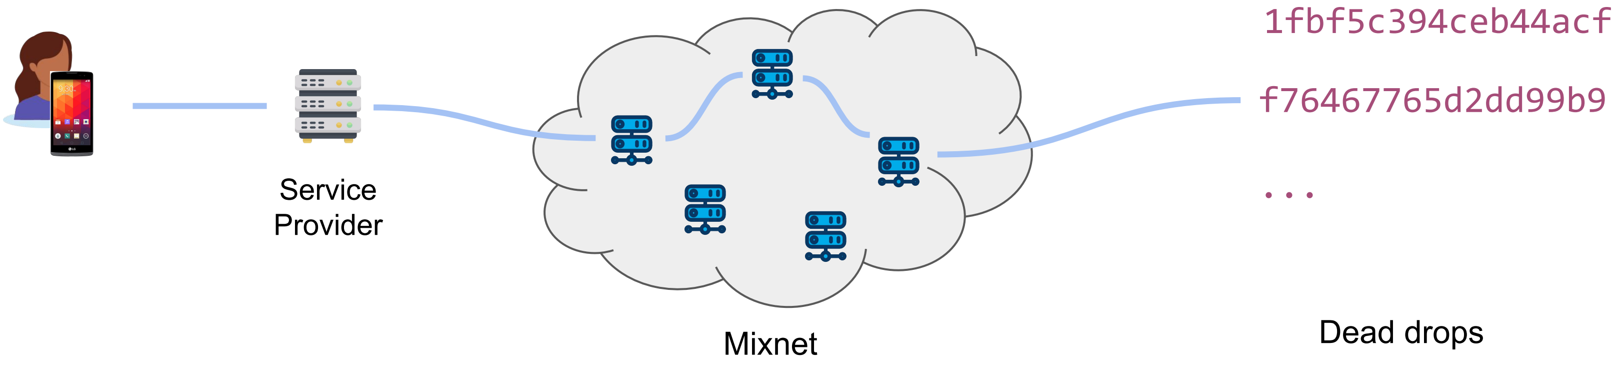

Per conversation, users build circuits: reusable, bidirectional paths

- last 1 day
- 1 msg / minute

Exchanging messages

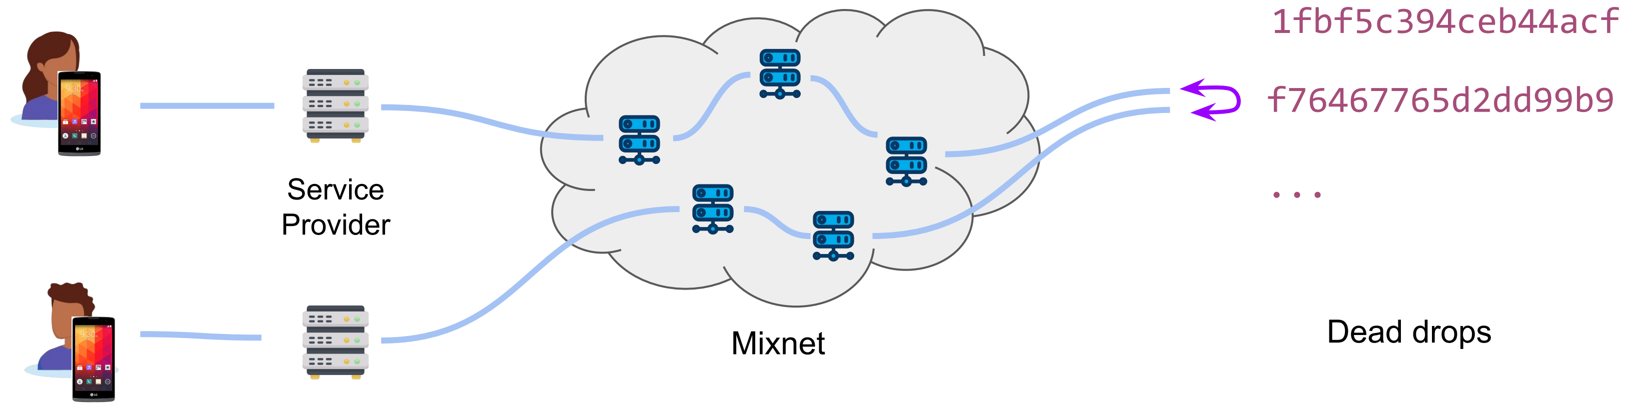

If two users pick the same dead drop, messages are swapped

Primed Circuits

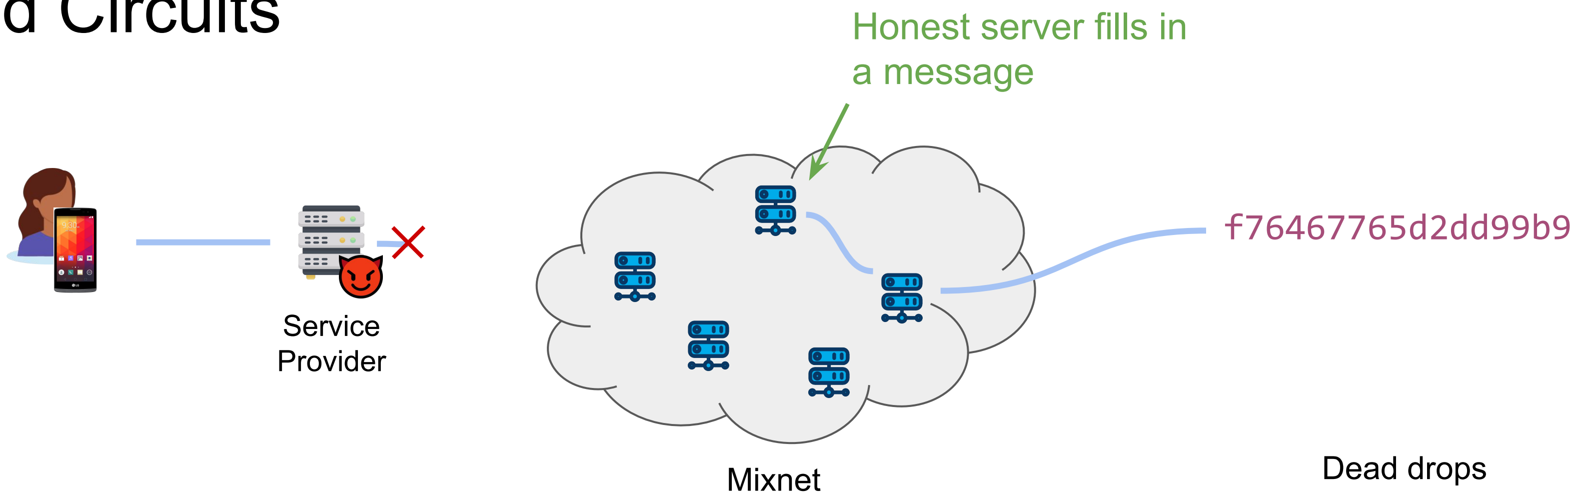

Circuits:
- Resist active attacks

Primed Circuits

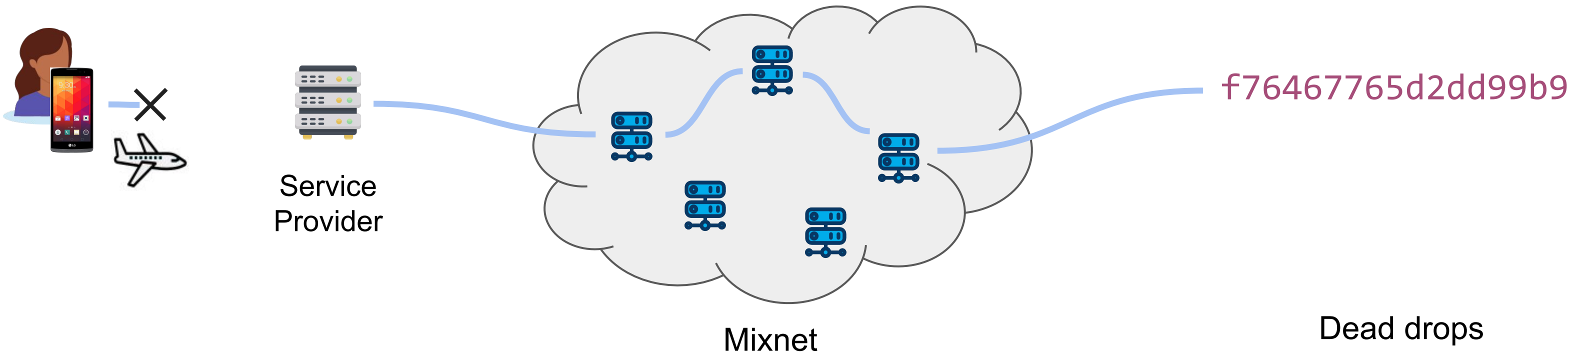

Circuits:

- Resist active attacks
- Facilitate cover traffic:
  - Every user **receives at a constant rate**, even when senders go offline

Primed Circuits

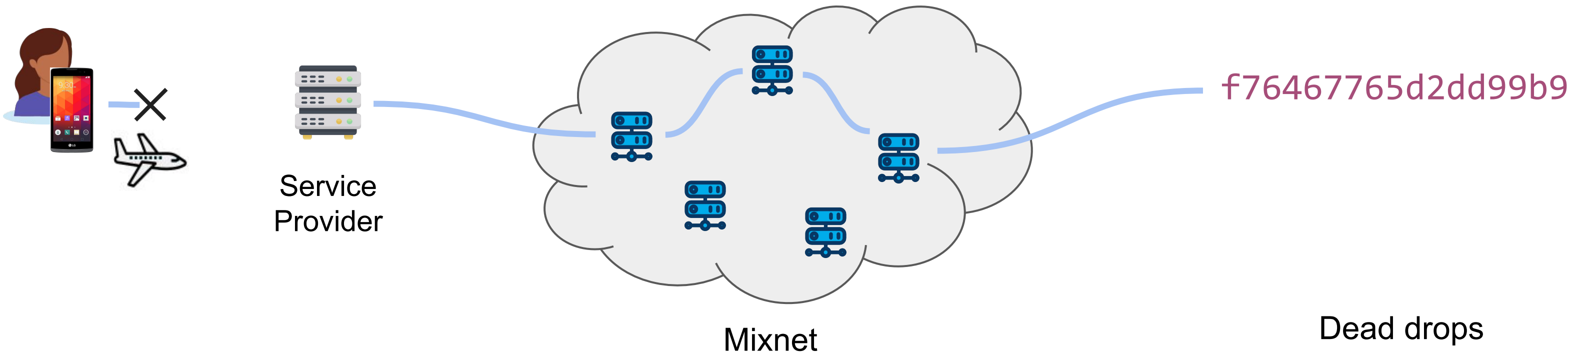

Circuits:

- Resist active attacks
- Facilitate cover traffic:
  - Every user **receives at a constant rate**, even when senders go offline

- Circuit setup is non-interactive
  - Alice uploads for ~1 month worth of circuit-setup messages

Handling many buddies

- One* circuit per friend (* actually two)
- 50 circuits = 50 friends
- Client send/fetches must not reveal which circuit is used

Upstream:

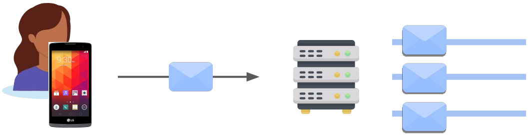

Messages are broadcasted on all (50)

Downstream:

Strawman: Download everything

Drawback: most messages are noise

Improved fetch protocol (not presented)

Multi-device safety

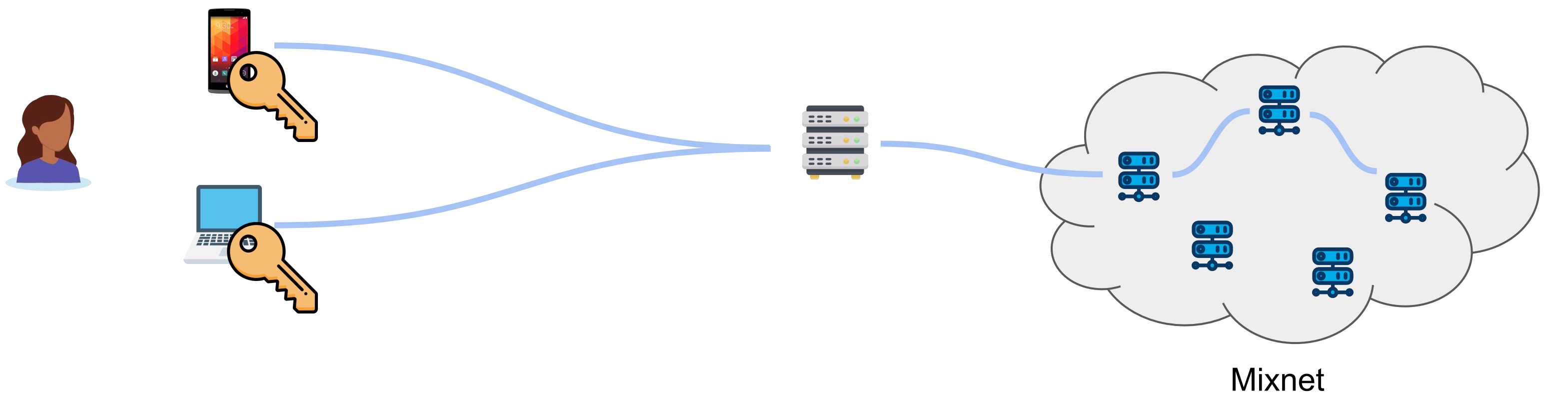

- Devices share a key `multiDeviceKey`
- Even partitioned, devices pick the same paths:

```
device_i.circuit[j].path = PRNG(multiDeviceKey, epoch, j)
                 j in [0;50]
```

- Each mix de-duplicates incoming messages with the same tag

Security properties

Two proofs:

- The mixnet provides differential privacy:

$$\Pr[\mathrm{obs}_A \mid \mathrm{Alice} \leftrightarrow \mathrm{Bob}] \leq e^{\epsilon} \Pr[\mathrm{obs}_A \mid \mathrm{Alice} \not\leftrightarrow \mathrm{Bob}] + \delta$$
$$\Pr[\mathrm{obs}_A \mid \mathrm{Alice} \not\leftrightarrow \mathrm{Bob}] \leq e^{\epsilon} \Pr[\mathrm{obs}_A \mid \mathrm{Alice} \leftrightarrow \mathrm{Bob}] + \delta$$

- Security of the service provider reduces to the mixnet

Experimental setup

- client: Pixel 4 phone
- 100 servers on AWS in 4 regions (US + EU)
- each server is a 32 core 3.1Ghz CPU, 256 GB RAM, 10 Gbps network
- 3 Mio users each with 50 conversations

Experimental results - SP + Mixnet

Bandwidth usage:

Setup: 47.5 GB / epoch / mix server
Messaging: 13 GB / round / mix server

Storage at the Service Provider for 1 month:

Setup: 2.1 MB / user
Messaging: 264 MB / user

Experimental results - SP + Mixnet

Bandwidth usage:

Setup: 47.5 GB / epoch / mix server
Messaging: 13 GB / round / mix server

Storage at the Service Provider for 1 month:

Setup: 2.1 MB / user
Messaging: 264 MB / user

Latency:

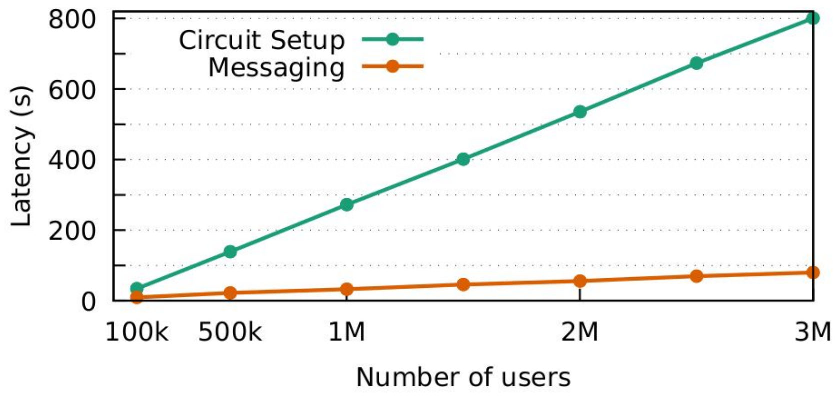

Messaging: 1M: **32s** 2M: **55s** 3M: **80s**

Experimental results - Phone

<u>Bandwidth usage:</u>

Setup: 110 KB/epoch = 100 MB/month

Messaging:

for a 1-min client schedule, SP + mixnet
latency of 32s

↑130 KB/h ↓140 KB/h = 190 MB/month

latency: between 32s and 64s

Experimental results - Phone

Bandwidth usage:

Setup: 110 KB/epoch = 100 MB/month

Messaging:

for a 1-min client schedule, SP + mixnet latency of 32s

↑130 KB/h ↓140 KB/h = 190 MB/month

latency: between 32s and 64s

Energy usage:

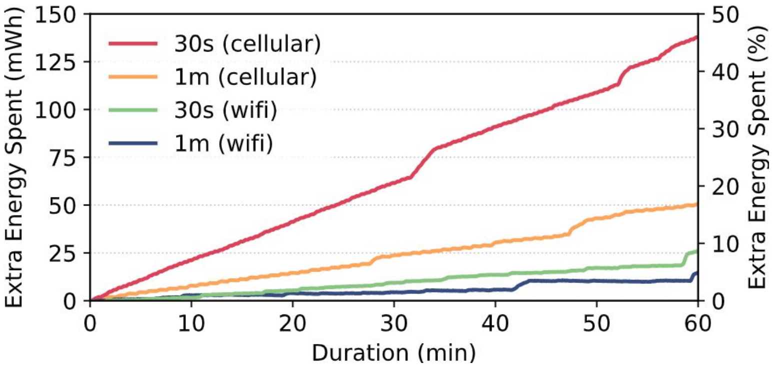

Experimental results - Phone

<u>Bandwidth usage:</u>

Setup: 110 KB/epoch = 100 MB/month

Messaging:

for a 1-min client schedule, SP + mixnet latency of 32s

↑130 KB/h ↓140 KB/h = 190 MB/month

latency: between 32s and 64s

<u>Energy usage:</u>

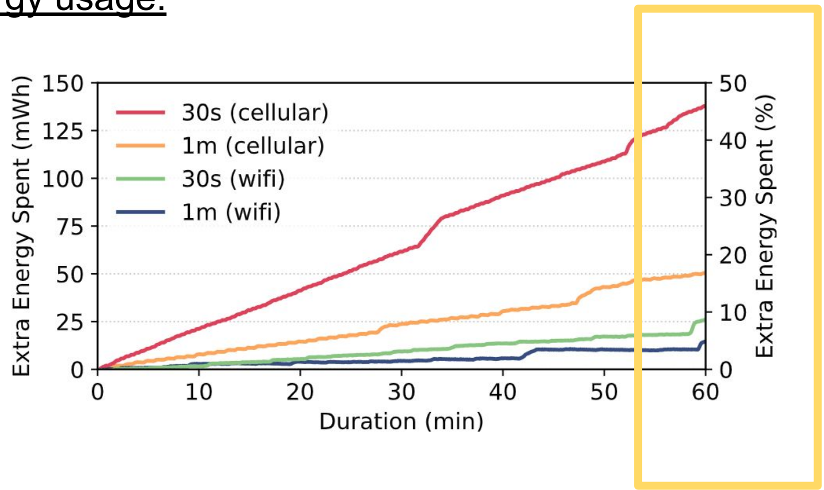

Experimental results - Phone

<u>Bandwidth usage:</u>

Setup: 110 KB/epoch = 100 MB/month

Messaging:

for a 1-min client schedule, SP + mixnet latency of 32s

↑130 KB/h ↓140 KB/h = 190 MB/month

latency: between 32s and 64s

<u>Energy usage:</u>

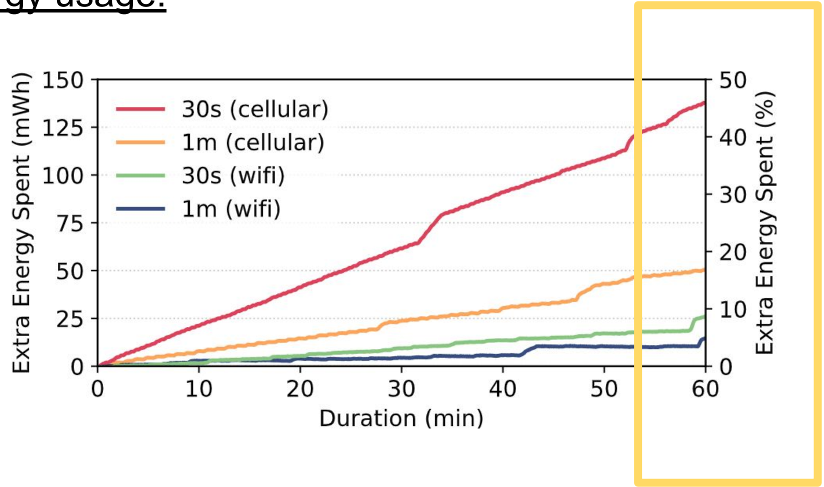

With a 5-min schedule, after 1h:
≈ +5% energy usage

Conclusion

Contributions of the thesis

- Every Byte Matters: Traffic Analysis of Bluetooth Wearable Devices (Ch 2)
  - First broad analysis of the communication metadata of wearable devices
  - We reveal a general susceptibility to traffic-analysis attacks, which can allow:
    - identifying devices, applications, user actions
    - tracking and profiling users
  - If we want to protect such information, we need defense strategies

Contributions of the thesis

- Every Byte Matters: Traffic Analysis of Bluetooth Wearable Devices (Ch 2)
    - First broad analysis of the communication metadata of wearable devices
    - We reveal a general susceptibility to traffic-analysis attacks, which can allow:
        - identifying devices, applications, user actions
        - tracking and profiling users
    - If we want to protect such information, we need defense strategies

- Padmé (Ch 3)
    - Padding function with low costs (<12%) that outperforms classic approaches asymptotically
    - In practice, we show that it has good hiding properties

Contributions of the thesis (cont')

- PriFi (§4.4)
  - Low-latency, traffic-agnostic anonymity for a small set of users (VoIP support)
  - The latency does not depend on the latency to the anytrust servers
  - "On-path" anonymization that provides low latency

Contributions of the thesis (cont’)

- PriFi (§4.4)
  - Low-latency, traffic-agnostic anonymity for a small set of users (VoIP support)
  - The latency does not depend on the latency to the anytrust servers
  - “On-path” anonymization that provides low latency

- Rubato (§4.5)
  - First large-scale ACN with multi-device, asynchronous clients (Global Active Adversary setting)
  - Each device can choose its communication frequency & costs
  - It enables mobile devices to participate at a reasonable cost

Impact outside of research

- Every Byte Matters: Traffic Analysis of Bluetooth Wearable Devices
  - Contacted ~100 vendors and manufacturers, ~10 follow-ups by email, 2 follow-up meetings with large device manufacturers
  - Received a bug bounty
- Padmé
  - Maintainers of SequoiaPGP implemented Padmé
- PriFi
  - Demos at the Red Cross (ICRC) headquarters and at EPFL (one awarded a prize)
  - Patent

Next steps for metadata privacy ?

Still an open problem:

- No one-size-fits-all defense

  => Per domain, iteratively evaluate risks

- Compared to non-metadata-private alternatives, solutions are costly

  => Increase visibility of the attacks to justify the costs

  - Open-source datasets & tooling

Building safer apps

- Could we have automated guidelines for app developers ?

- Could we have “defense strategies” provided by the OS ?

  This could be an opportunity for designing the defenses iteratively

Analyzing and Protecting Communication Metadata

Experimental

Every Byte Matters (Ch2) [4]
Traffic-analysis attack of wearable devices.

Theoretical

Padmé (Ch3) [2]
A padding function that efficiently hides sizes.

Attacks ←→ Defenses

Systems: Anonymous Communication Networks (ACN)

PriFi (§4.4) [1,3]
Traffic-agnostic, low-latency ACN for local-area networks.

Rubato (§4.5) [5]
Large-scale ACN for text messaging on mobile devices.

[1] L. Barman, M. Zamani, I. Dacosta, J. Feigenbaum, B. Ford, J.-P. Hubaux, D. Wolinsky. **PriFi: A Low-latency [...] Protocol for Local-Area Anonymous [...]**. WPES 2016.

[2] K. Nikitin*, L. Barman*, W. Lueks, M. Underwood, J.-P. Hubaux, B. Ford. **Reducing Metadata Leakage from Encrypted Files and Communication with PURBs**. PETS 2019

[3] L. Barman, I. Dacosta, M. Zamani, E. Zhai, A. Pyrgelis, B. Ford, J. Feigenbaum, J.-P. Hubaux. **PriFi: Low-latency Anonymity for Organizational Networks**. PETS 2020

[4] L. Barman, A. Dumur, A. Pyrgelis, J.-P. Hubaux. **Every Byte Matters: Traffic Analysis of Bluetooth Wearable Devices**. UbiComp 2021.

[5] L. Barman, M. Kol, D. Lazar, Y. Gilad, N. Zeldovich. **Rubato: Metadata-Private Messaging for Mobile Devices**. Under submission.