

Haystack

A multi-purpose mobile
vantage point in user space

Christian Kreibich

Abbas Razaghpanah, Narseo Vallina, Srikanth Sundaresan,
Phillipa Gill, Mark Allman, Vern Paxson

International Computer Science Institute
Stony Brook University

Part I

Background

Smartphones are everywhere



...but can we trust them?

- Privacy violations
- Malicious apps (ransomware, spyware, ...)
- App permission overuse
- Insecure operation

Investigation implies trade-offs



Tradeoffs: ISP traces



- Large scale
- Real-world traffic



- No context
- Encryption a problem


Tradeoffs: instrumented devices



- Device context
- Real-world activity
- Comprehensive analysis
- Small scale
- Tricky setup

[OSDI'10, IMC'13, CoNEXT'13]

Tradeoffs: static analysis



- Large scale
- Sufficient for some analyses
- No organic user activity

[NDSS'11, CCS'12-13, MobiSys'15]


Tradeoffs: proxy MITM



- Real-world traffic
- Comprehensive analysis
- No device context
- Detoured routes
- Higher trust hurdle

[CoNEXT'12, C2BID'15]

Tradeoffs: crowdsourced active measurement



- Large scale
- Comprehensive analysis

- No organic user activity

[CoNEXT'14, MobiSys'15, HotMiddlebox'15]

Can we do
better ?

Part II

The Haystack app

A few observations

- We want to run on the device
 - Best access to device context and the user
- We do not want to root the device
 - Drastically limits and skews deployment
- We like crowd-sourced measurement
 - Demonstrated large scale in Netalyzr

МНММ...

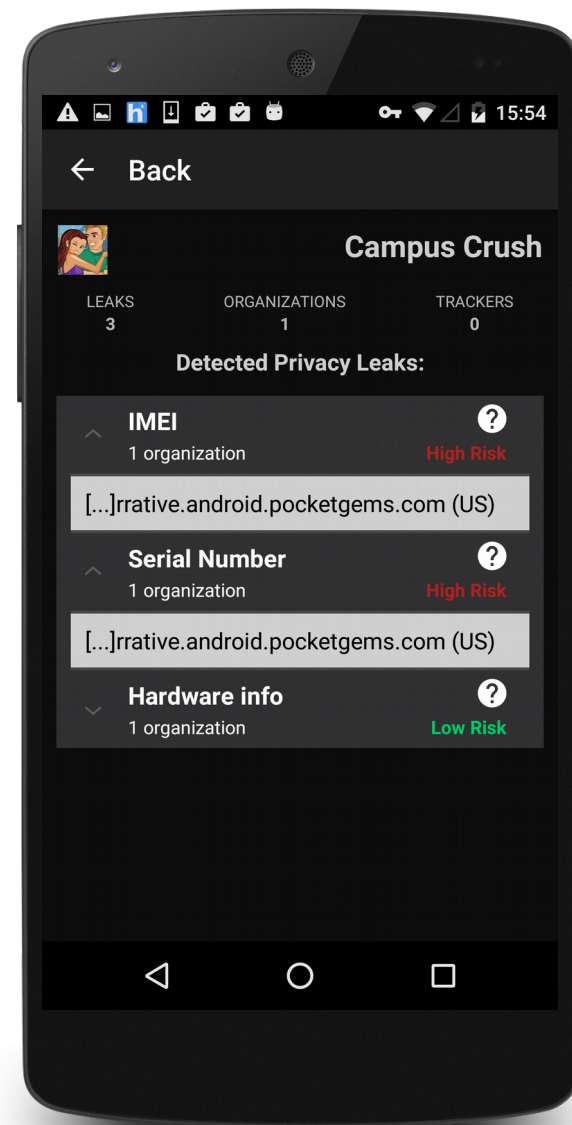
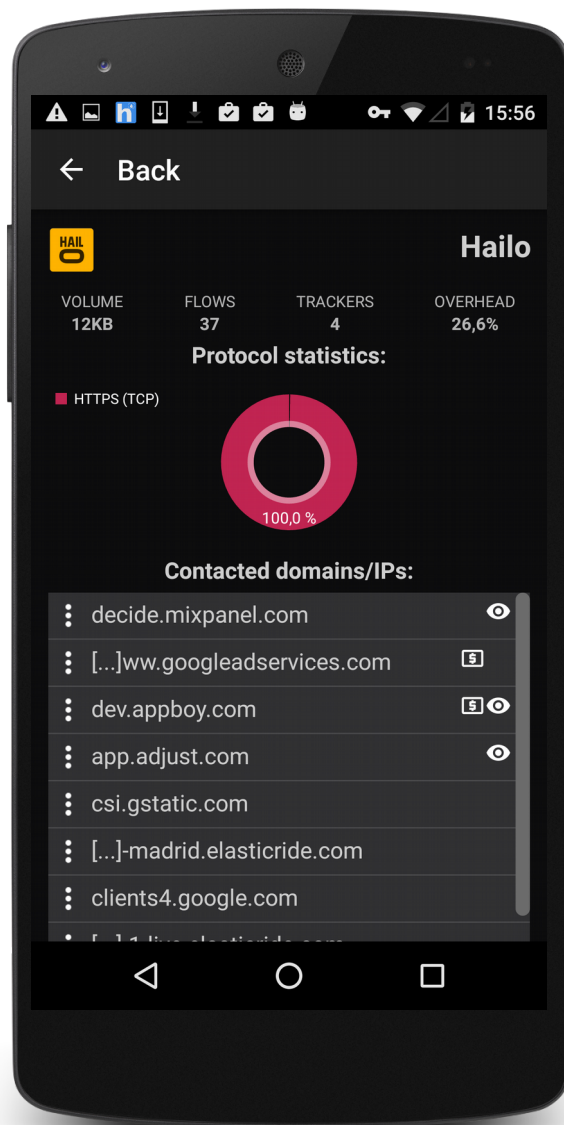

VPNs!

Android's VPN API

- VPN apps don't require rooting
- Access to all packets sent by the device
- Nobody says you have to tunnel them!
- Instead inspect & rewrite, and interact directly with intended destination

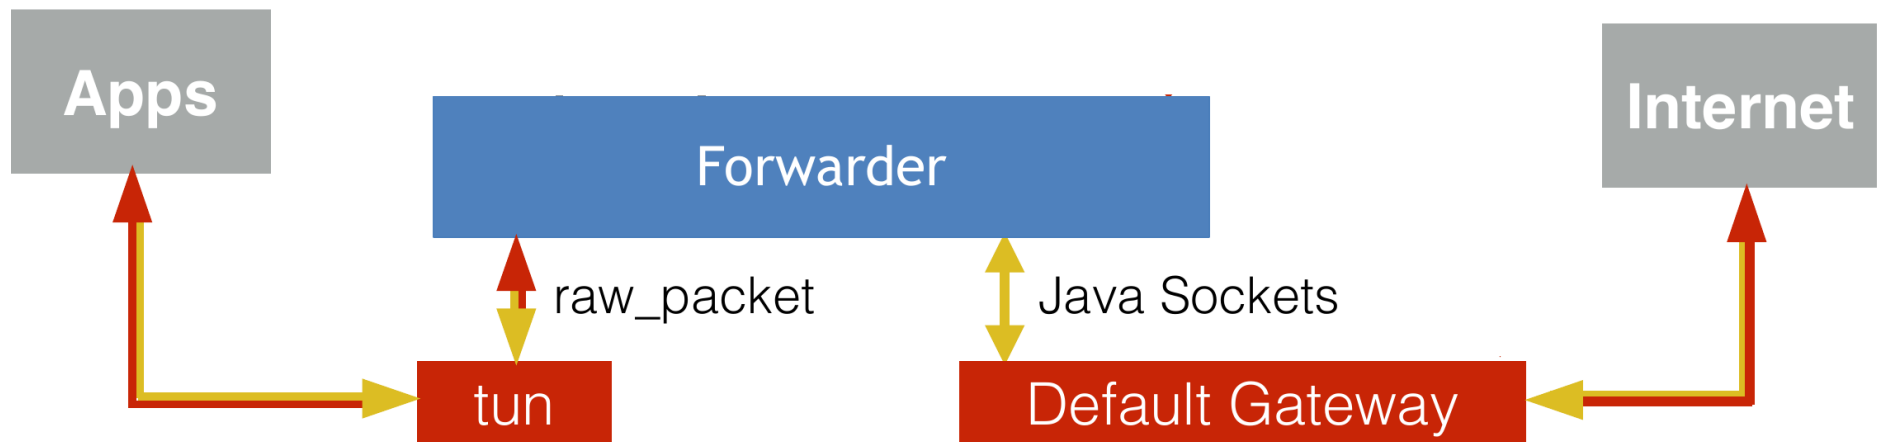
The Haystack app

A user-centric, on-device measurement platform platform that intercepts and inspects network traffic and app activity in user-space.

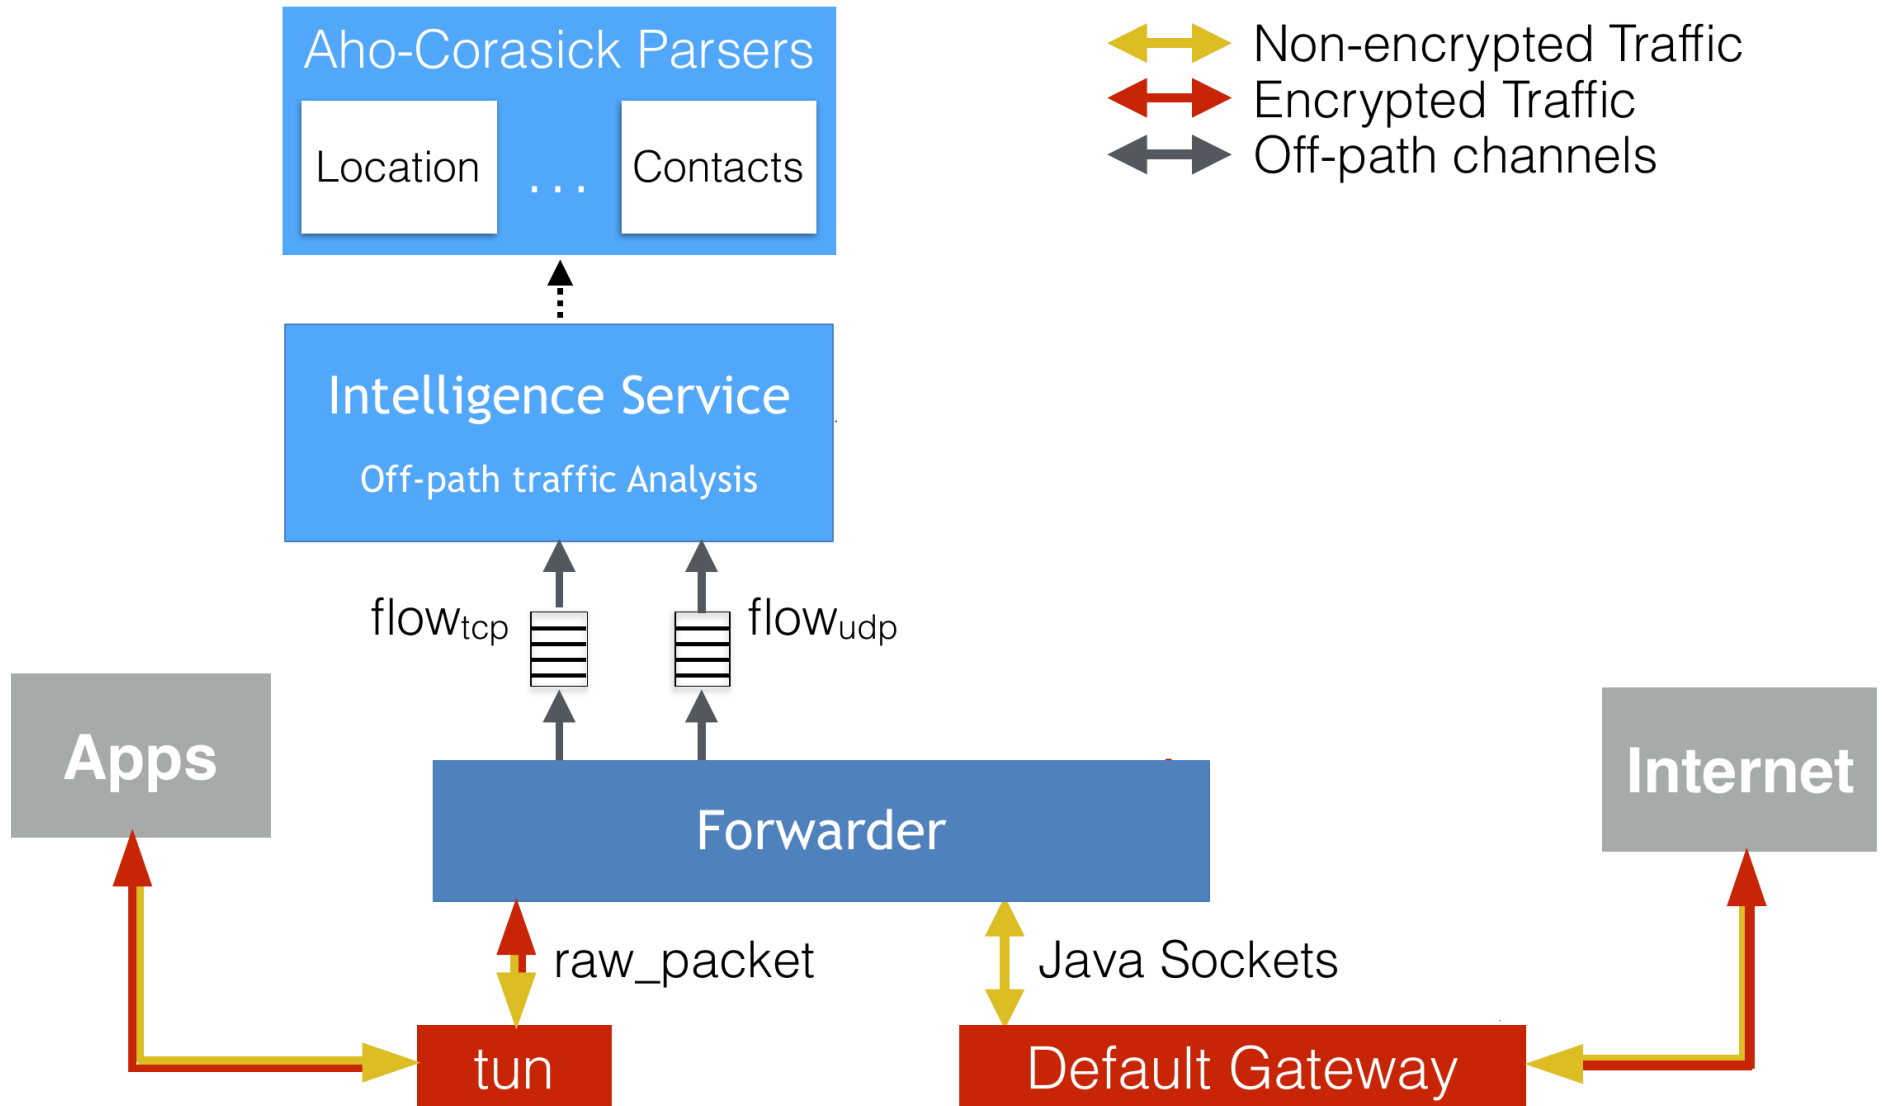


Architecture overview

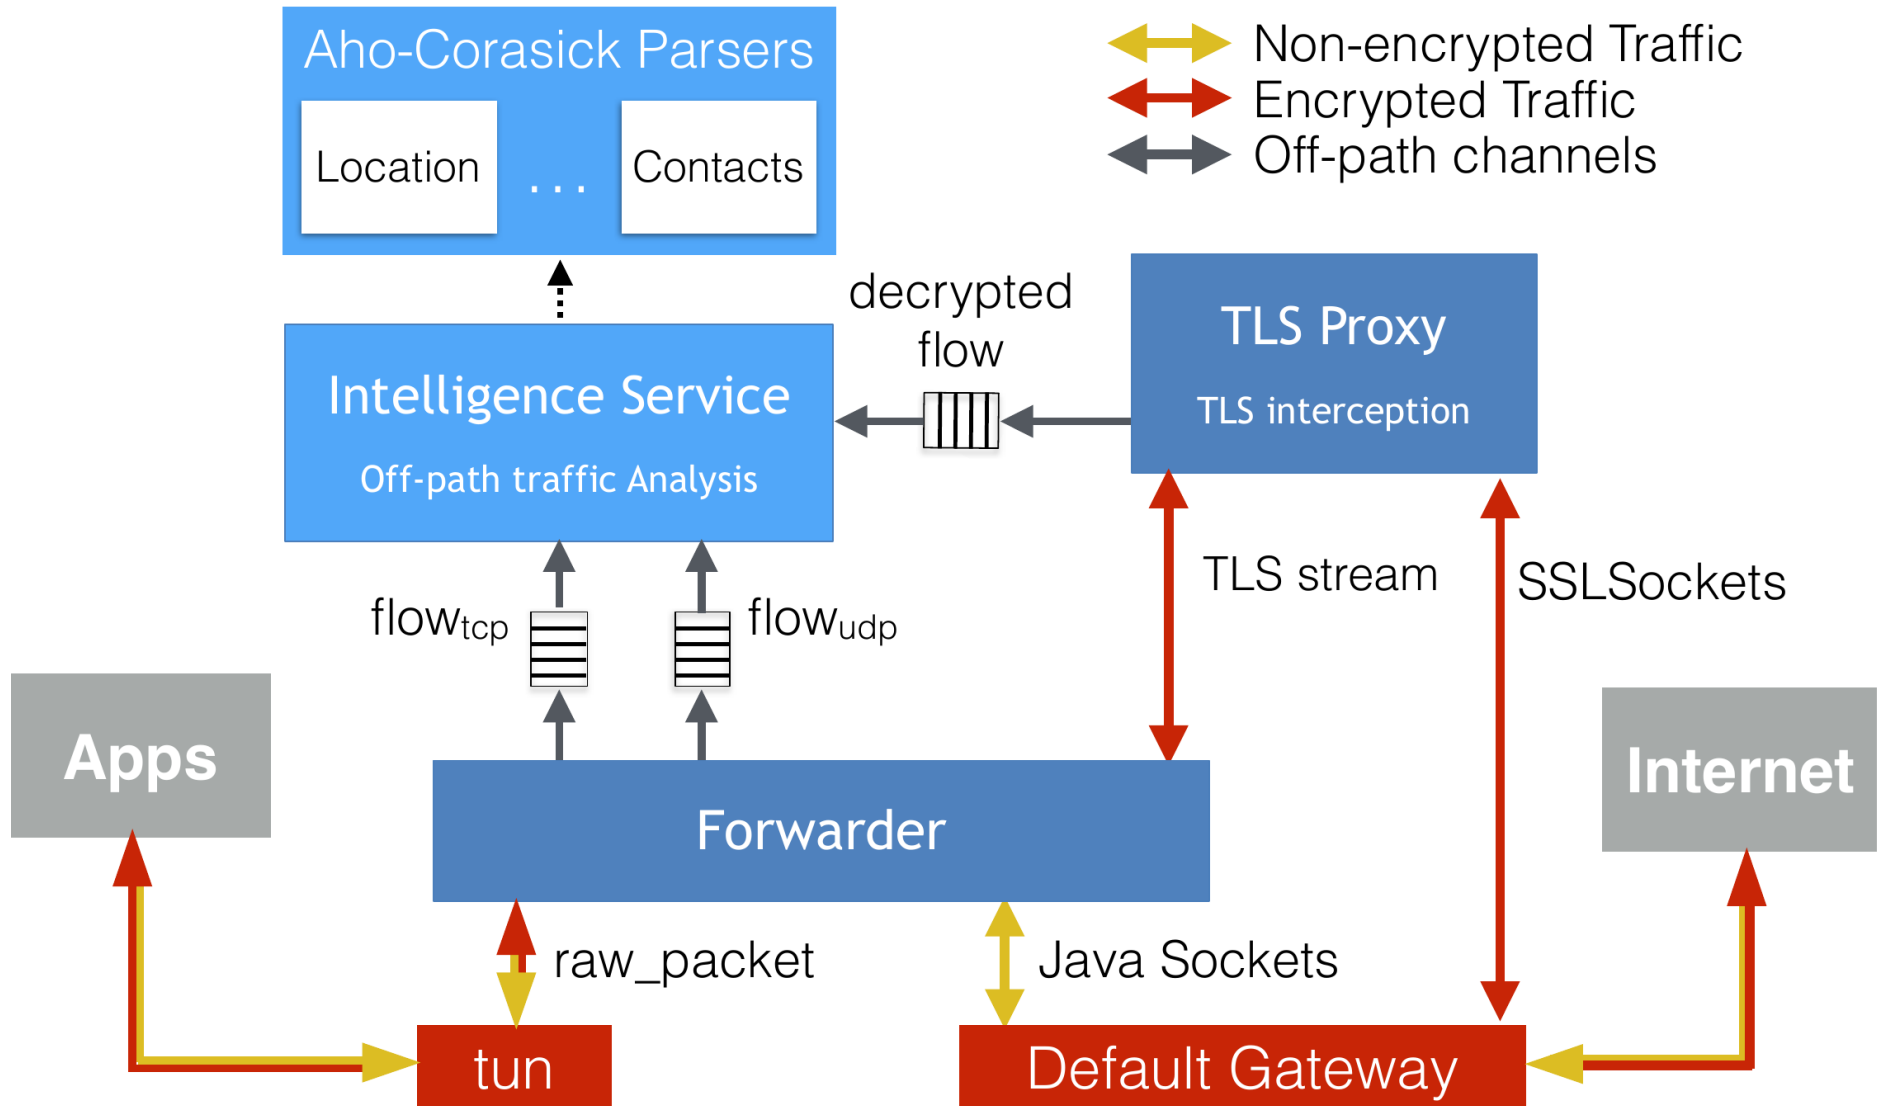
Architecture overview




Architecture overview




Architecture overview



Polling state machine




Polling state machine




Part III

Evaluation


CPU and power overhead



CPU and power overhead

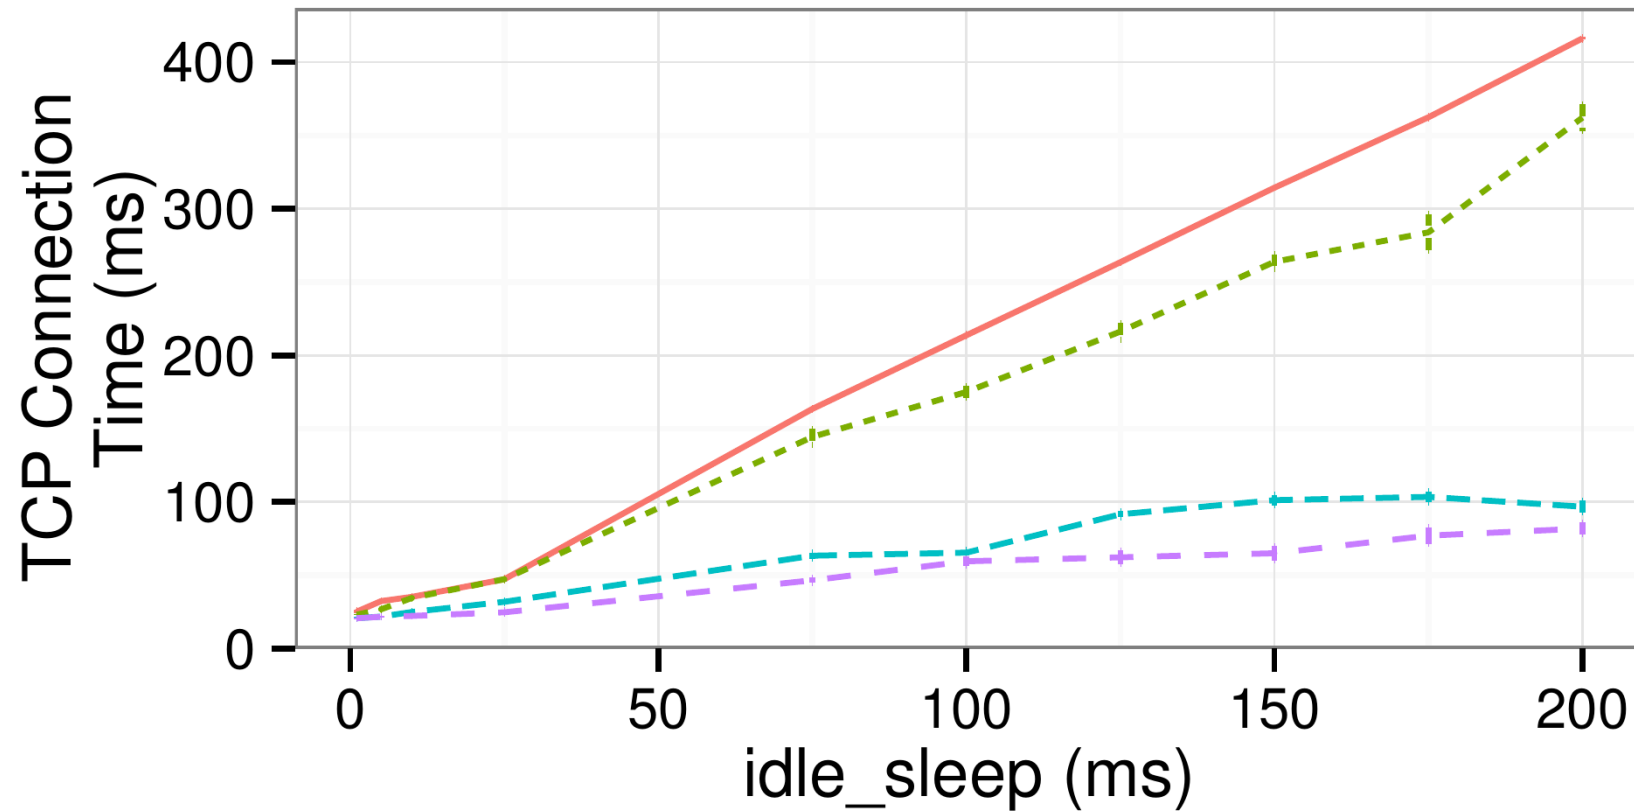


CPU and power overhead







Power: +3.1% when idle, +9.1% when busy

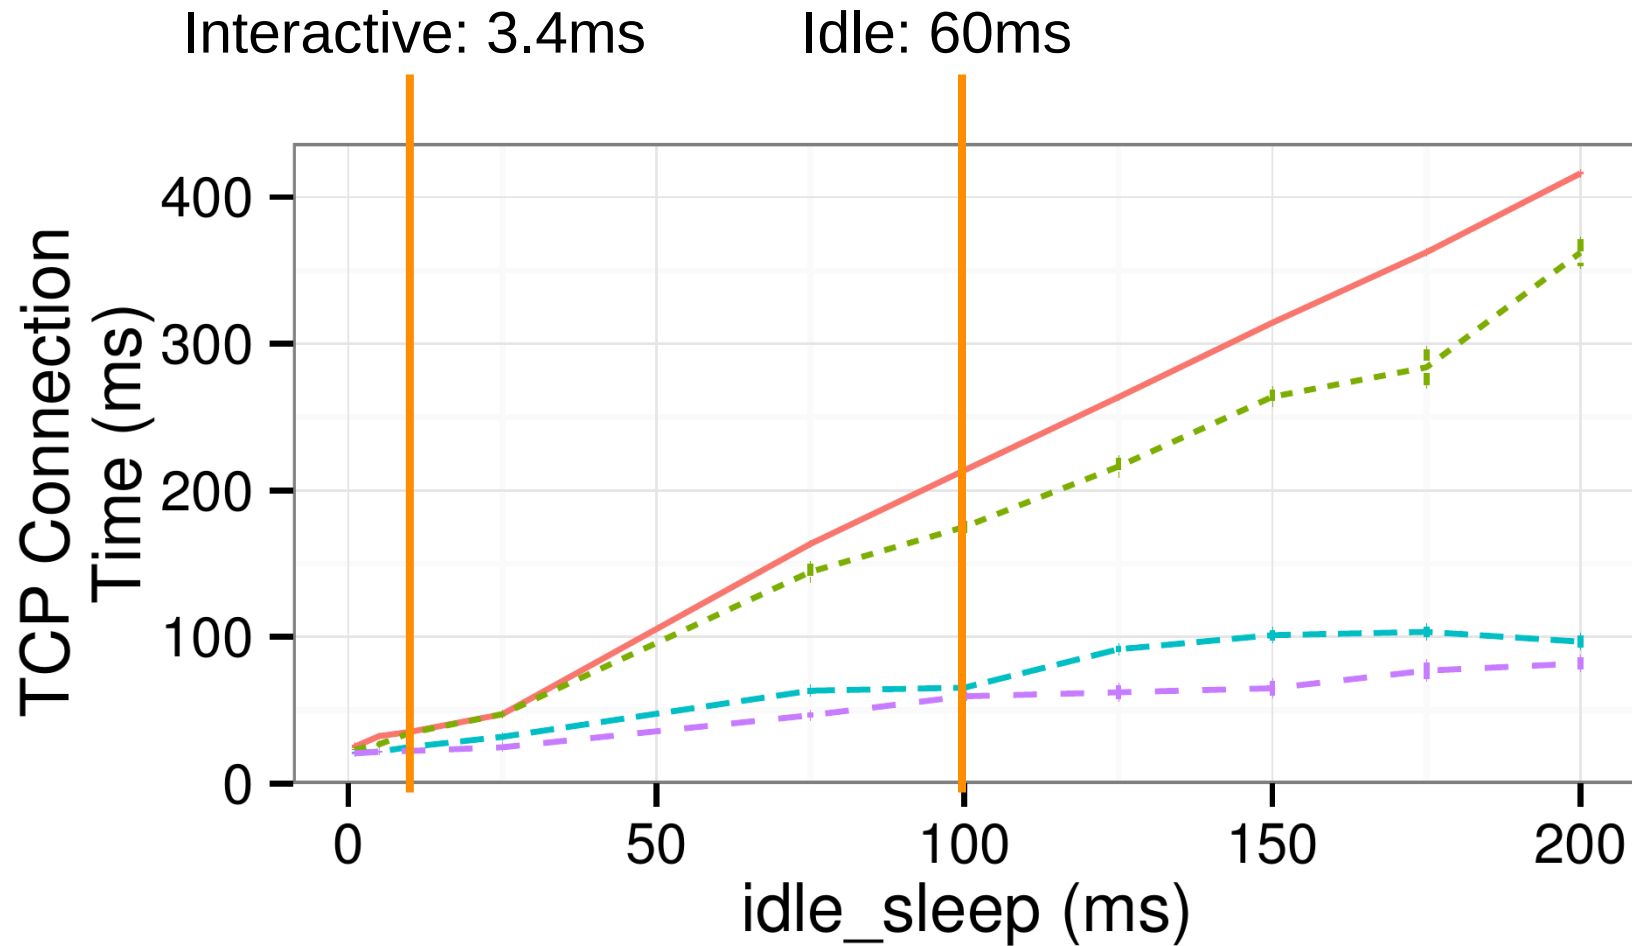
Latency overhead



max_idle_cycles


	1		10		100		200
--	---	---	----	---	-----	---	-----

Latency overhead



max_idle_cycles 1 10 100 200

Throughput




Part IV

Use cases

Pilot study

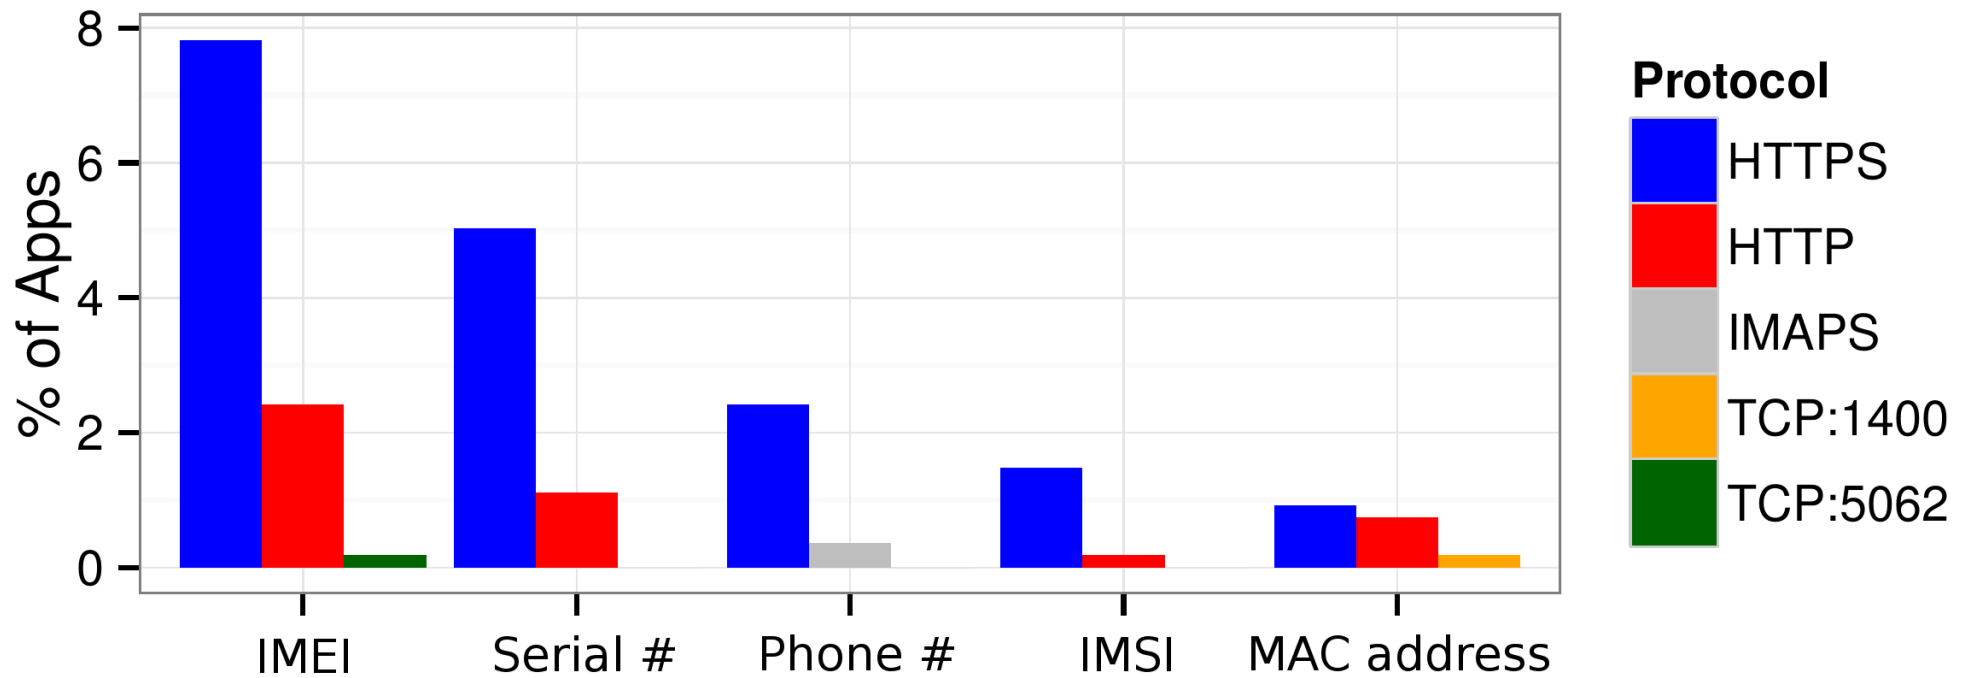
450 users, 1340 apps, 6 months,
app-focused data collection

Traffic properties



- Less than 20% of apps only send cleartext
- 22% of flows are encrypted
- 59% of TLS-using apps allow MITM
- 40 apps generate local IoT traffic

Privacy-related leakages



App properties

- 15% of apps do not come from Google Play
 - Pre-installed or from other stores
 - They create 22% of the observed traffic
- 78% use third-party trackers
 - Advertising, analytics, social net interactions, ...


https://haystack.mobi/panopticon

The ICSI Haystack Panopticon

An interactive map of tracking activity
on mobile apps

- [Home](#)
- [Blog](#)
- [Android App](#)
- [What is a panopticon?](#)

🔍 Use touchpad/mouse wheel for
zoom control



Help us to find more
trackers!



© 2016, ICSI, CA

Future work

- More direct user involvement
 - Notify of leakages as they happen
 - Highlight third-party footprint
- Alter / block traffic
 - Suppress third-party trackers
- Reactive measurement
 - Active measurement can give context or inform traffic alterations

Summary

The Haystack app

A user-centric, on-device measurement platform based on the Android VPN API



Access to organic user activity



Optionally inspects TLS



Has full device context



Enables user interaction



No rooting required, thus scalable



(Modest) performance overheads



Subject to crowdsourcing biases

Thanks!

<https://haystack.mobi>

christian@icir.org

[@ckreibich](#)

