Review of Mobile Traffic Monitoring Paper

Haystack: A Multi-Purpose Mobile Vantage Point in User Space

arXiv:1510.01419, Oct. 2016

Abbas Razaghpanah (Stony Brook University) Narseo Vallina-Rodriguez (ICSI)

Srikanth Sundaresan (ICSI) Christian Kreibich (ICSI / Lastline) Phillipa Gill (Stony Brook University)

Mark Allman (ICSI) Vern Paxson (ICSI / UC Berkeley)

Traffic Monitoring in Mobile Device

- Packet capture in Android or iPhone smartphone!
- How?
  - Root your smartphone???
  - Just install a mobile app
    - tPacketCapture,
- There is a possibility of implementing a user-layer app to capture packets on the smartphone!!!
  - Due to Android API
    - https://developer.android.com/reference/android/net/VpnService.html

Problem

- Given a proxy API on the smartphone
- Implement
  - Proxy TCP/UDP/IP protocol stack

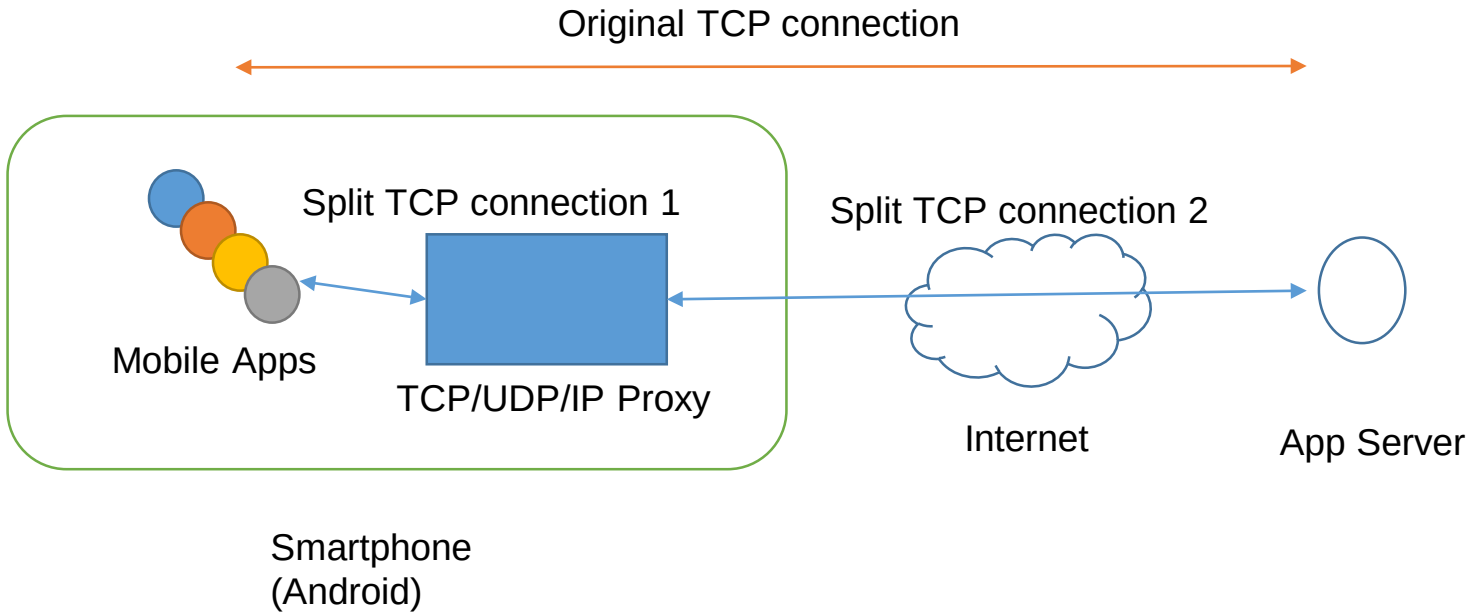

State-of-the-Art

- Haystack
  - ICSI--UC Berkeley and IMDEA Networks in collaboration with UMass and Stony Brook University
- AntMonitor: A System for Monitoring from Mobile Devices
  - UC Irvine
- ReCon: Revealing and Controlling PII Leaks in Mobile Network Traf
  - Northe
- Mobile Apps
  - Google Play. Packet Capture
  - Google Play. tPacketCapture

What is this “Haystack” Mobile App?

- “Application-layer” tcpdump in Android
  - Tcpdump
    - Capture packet and inspect the payload
    - Usually need “root” privilege
  - Mobile devices hacking is difficult
    - “rooting” is not popular to “average joe” users
  - Tcpdump as a mobile app!!!
- Why do we need the traffic monitoring app?
  - Many security and privacy incidents on the mobile devices
  - Monitor privacy leakage
- Android implementation
  - Mobile app in Java

Haystack Mobile App

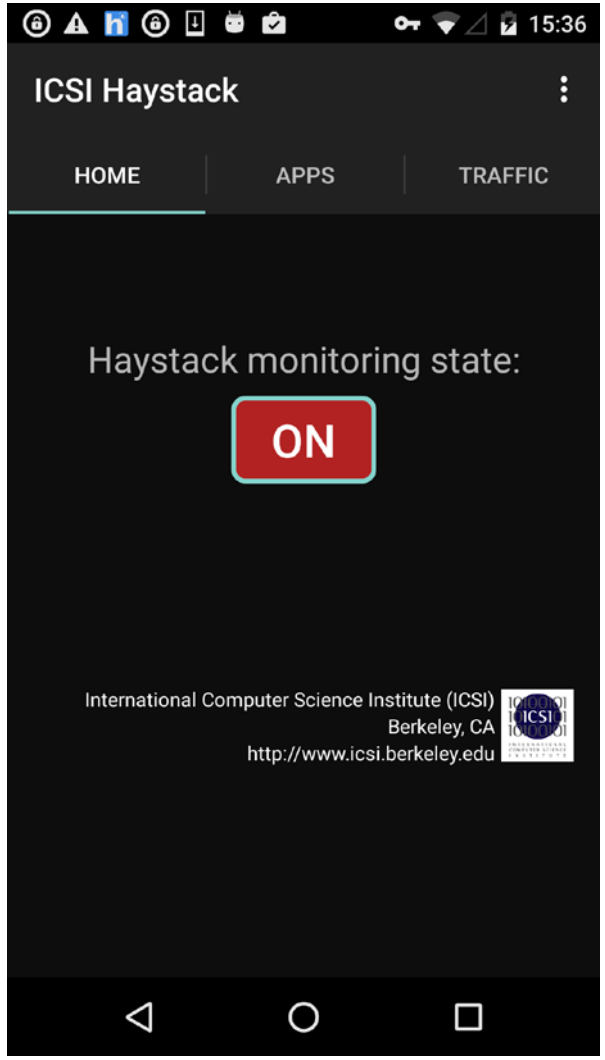

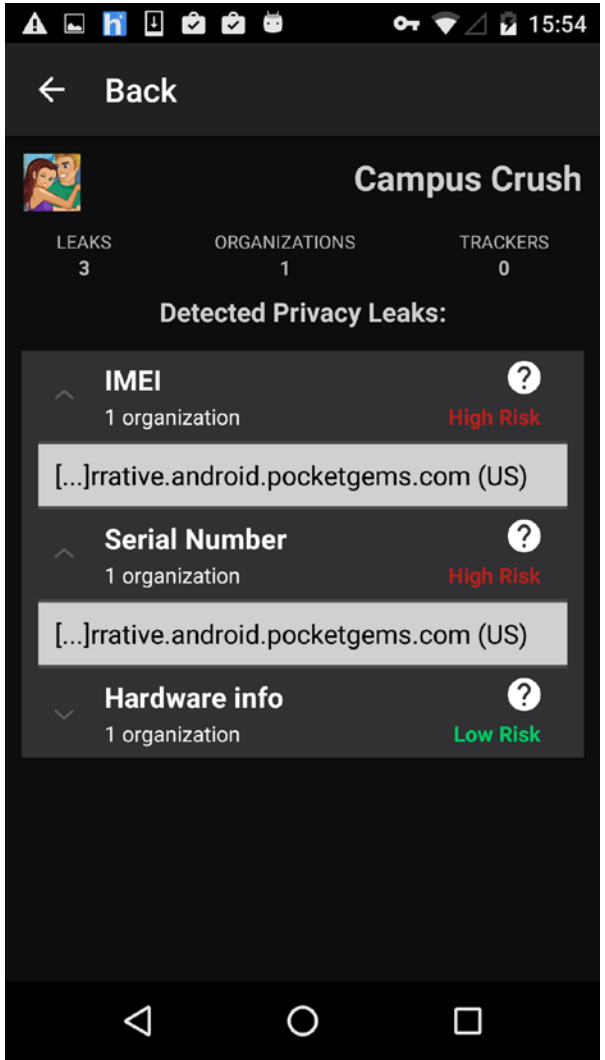

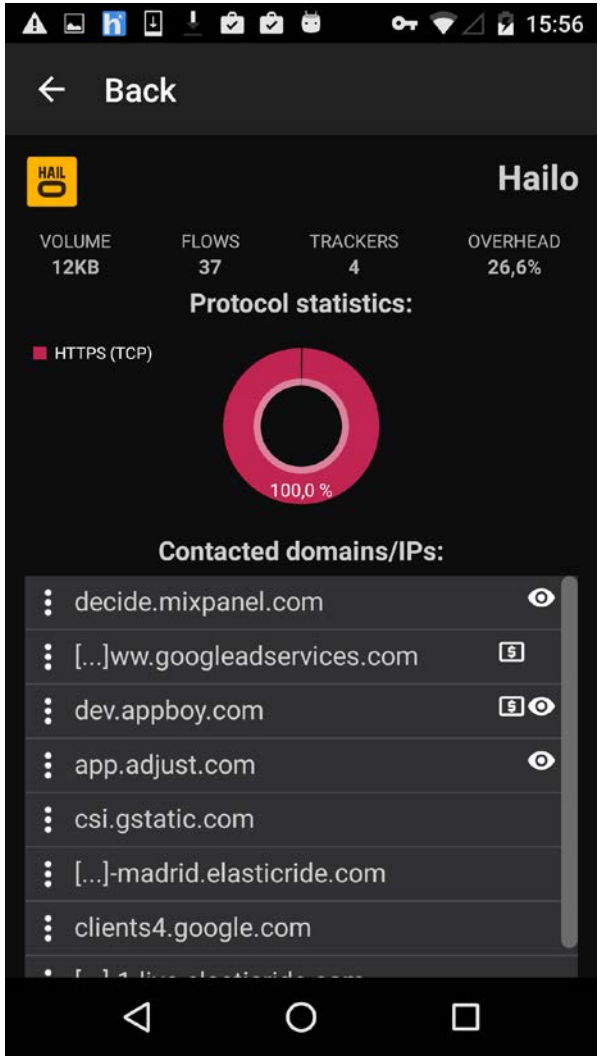

System Design

- Traffic Analyzer(TA)
  - Intelligence Service, Aho-Corasick Parsers
- Forwarder
  - TLS Proxy

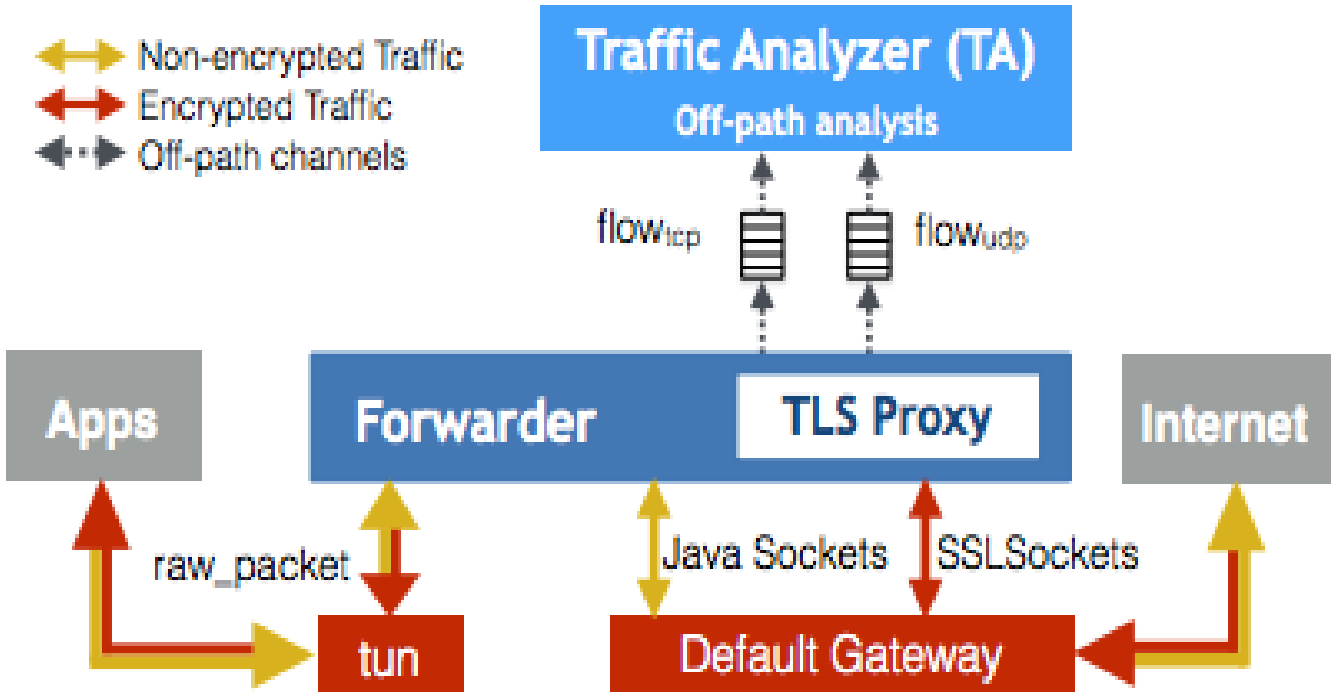

Ethical Considerations

- Best case
  - Do not deal with the private information
  - Collect the necessary information
- IRB at UC Berkeley
  - Need ethical consideration in the research

- SSL Decryption
  - User agreement
  - CA

Forwarder state machine

- Forwarder
  - TCP/UDP proxy
    - Split TCP/UDP connection
  - App ←→ Haystack ←→ Internet
- Tunneling interface
  - App ← → Haystack
- Nio interface
  - Haystack ←→ Internet servers

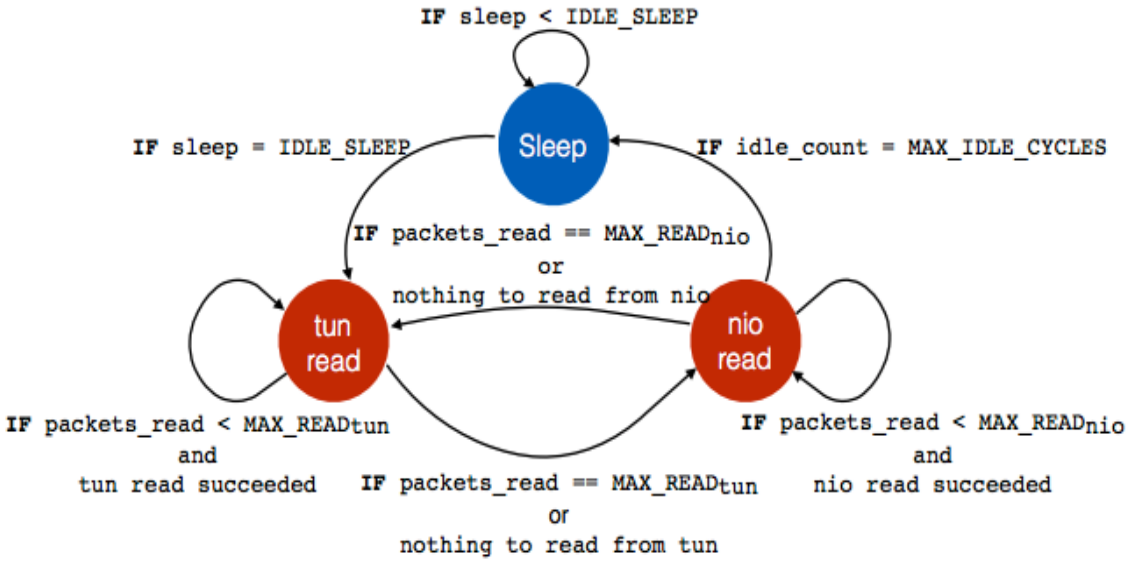

Figure 2: Haystack's Forwarder state machine. It controls read/write operations and transitions between **tun** interface, Java NIO socket, and sleep states. The idle count variable increments when both **tun** and NIO do not succeed, *i.e.*, there is nothing to read. Each read operation from the **tun** interface potentially becomes a write operation for a NIO socket and vice versa.

TLS Interception

- Man-in-the-middle(MITM) proxy on the TLS transaction
- Need self-signed Haystack CA
  - User agreement
- Decryption

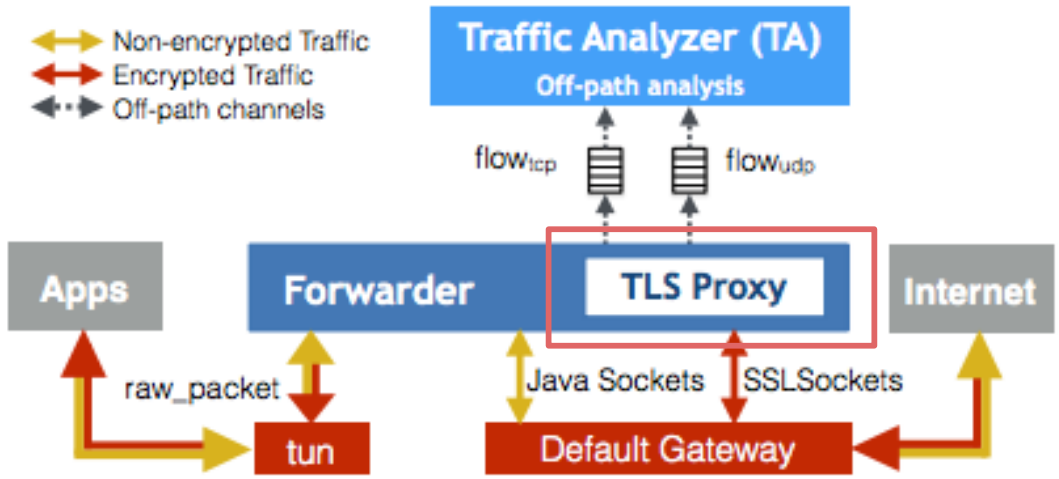

Dealing with failed TLS interception

- Failure of TLS proxy
  - Strong security policy used in app
    - E.g., facebook, google
    - Certificate pinning
- Bypass proxy

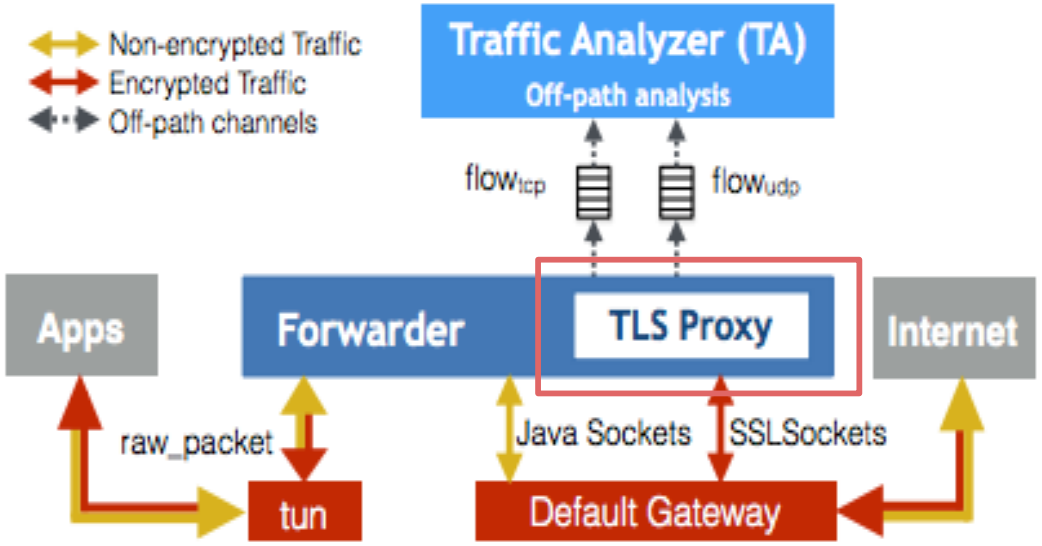

Traffic Analyzer

- Packet analysis
  - Parsing TLS, HTTP, DNS
- Off-path analysis
- Application and entity mapping
- Tracking DNS transaction
  - Non-HTTP flow: QUIC, HTTPS

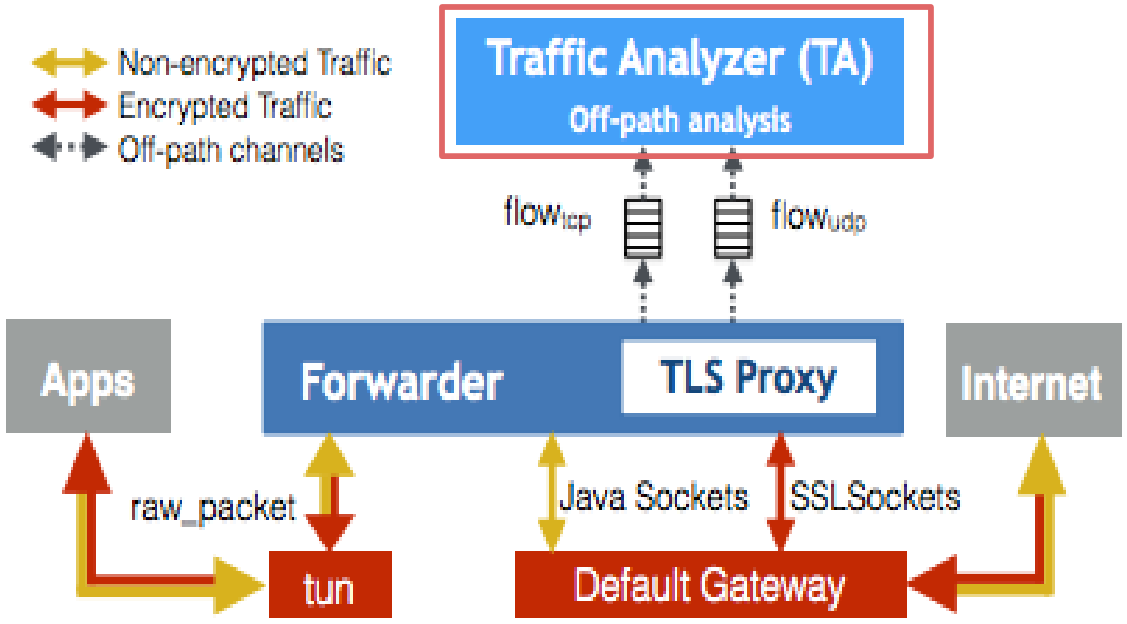

Testbed and Measurement Apparatus

- Nexus 5
- 5 GHz 802.11n link (wireless access point)
- Simple UDP and TCP echo packets
- max idle cycles, idle sleep, max readtun, max readnio

CPU load

- max_idle_cycles
  - ok
- idle_sleep has impact on CPU
  - Less than 10~25ms
- Optimal idle_sleep
  - 100ms

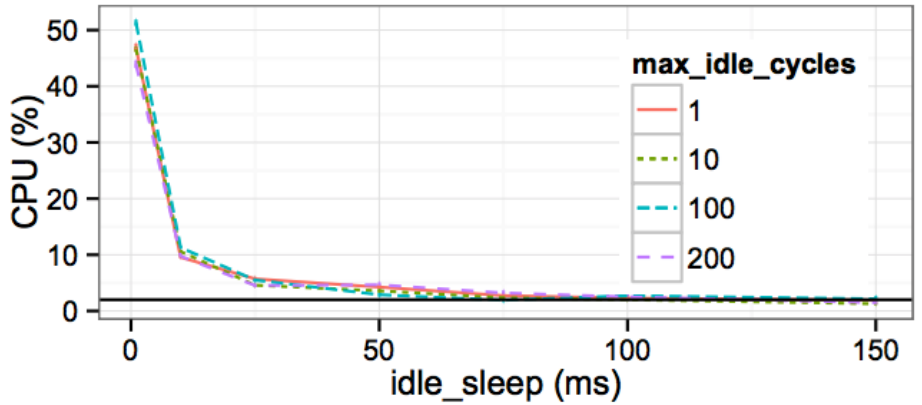

Figure 3: Haystack's CPU overhead for different *max_idle_cycles* and *idle_sleep* configurations. The horizontal line indicates the aggregated average CPU load of all apps running on the background for reference.

Power consumption

- Monsoon Power Monitor
- worst case
  - max_idle_cycles : 100
  - idle_sleep = 1ms
- 3-9% power usage increase

Test Case	Power(mW) Mean/SD	Increase
Idle	1,089.6 / 125.9	+3.1%
Idle (Haystack)	1,123.8 / 150.4	
YouTube	1,755.3 / 35.5	+9.1%
YouTube (Haystack)	1,914.4 / 16.1	

Table 2: Power consumption of Haystack when *max_idle_cycles* is 100 cycles and *idle_sleep* is 1 ms in different scenarios. The percentage indicates the increase when running Haystack.

Throughput of Haystack

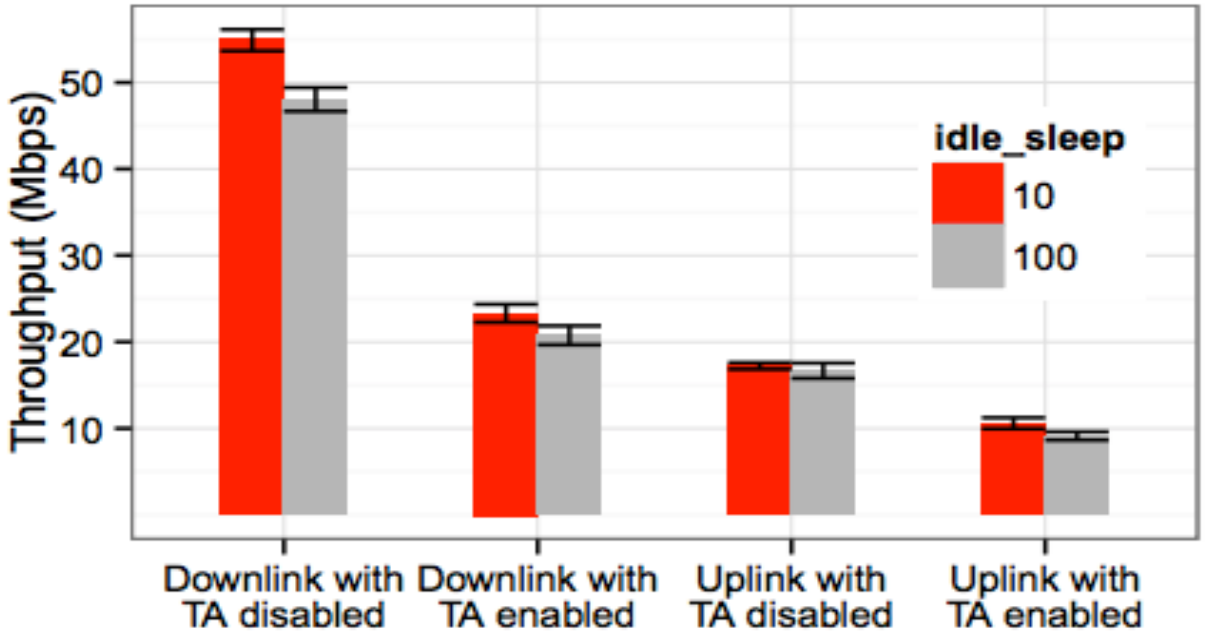

(c) TCP Throughput.

TLS Performance in Haystack

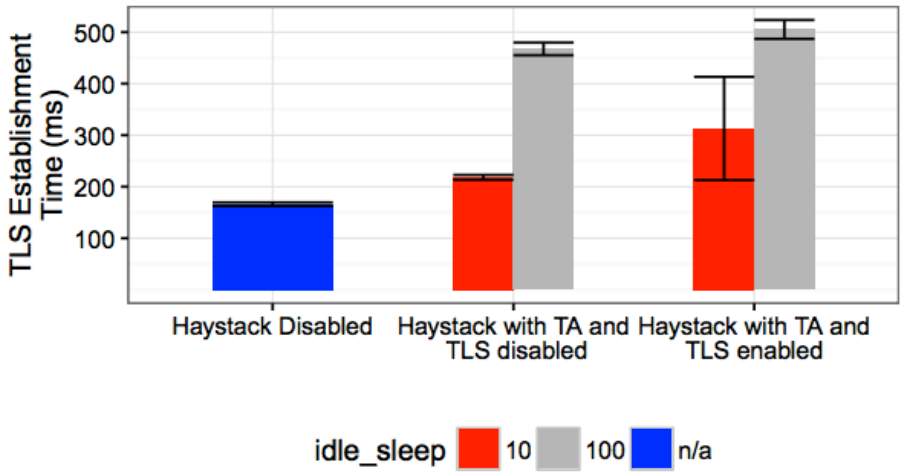

(a) TLS session establishment time.

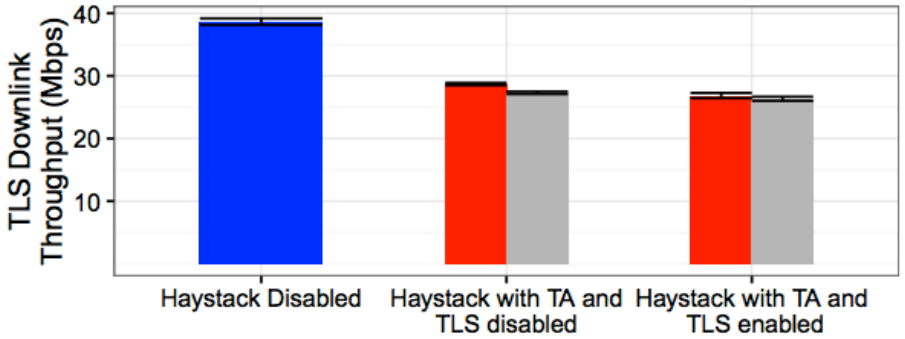

(b) TLS download speeds.

Summary

- Traffic monitoring for security in mobile device
- Need user space tool
  - Do not use “rooting”!

- Android
  - Local VPN Java class by Google
- iOS
  - Network extension library by Apple