A Smart Home is No Castle: Privacy Vulnerabilities of Encrypted IoT Traffic

Noah Apthorpe, Dillon Reisman, Nick Feamster

Computer Science Department & Center for Information Technology Policy, Princeton University

Abstract

- Internet-of-things (IoT) devices for "smart" homes are rapidly growing in popularity
- Many IoT devices have always-on sensors that constantly monitor users' home environments and transmit sensor readings to the cloud
- Encrypting communications between devices and cloud servers does not preserve privacy
- **A network observer can identify smart home devices and infer sensitive user behaviors from Internet traffic rates alone**
- A combination of firewalling, tunneling, shaping and injecting network traffic could protect privacy if properly implemented

Threat Model

- Passive network observer
  - E.g., an Internet service provider (ISP) or Wi-Fi eavesdropper
  - Can record network traffic
  - Can obtain and analyze IoT devices for supervised inference
- **Packet contents encrypted**
- Metadata available
  - IP & transport layer packet headers
  - DNS queries
  - Send/receive rates

3 Step Privacy Attack

1) Separate traffic into individual device flows
2) Identify device generating each flow
3) Infer user behaviors from traffic rate changes leveraging known device function

Data Collection

- Laboratory smart home with several commercially-available IoT devices
- Recorded network traffic using customized Raspberry Pi Wi-Fi access point

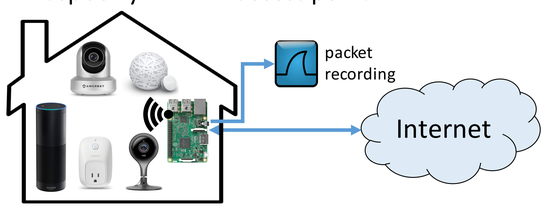

ISPs and Wi-Fi eavesdroppers can identify IoT devices inside a smart home

DNS queries

- DNS queries often contain device manufacturer name or other identifier, allowing easy device identification

Device	Informative DNS queries
Sense Sleep Monitor	messeji.hello.is ntp.hello.is sense-in.hello.is time.hello.is
Nest Cam Indoor Security Camera	nexus.dropcam.com oculus519-vir.dropcam.com
WeMo Switch	prod1-api-xbcs-net-889336557.us-east-1.elb.amazonaws.com
Amazon Echo	device-metrics-us.amazon.com ntp.amazon.com pindorama.amazon.com softwareupdates.amazon.com

Machine Learning Classification

- DNS queries may not be available
  - A Wi-Fi eavesdropper without WPA2 password cannot read DNS queries in encrypted 802.11 frames
- **Simple supervised machine learning can identify device types using only network traffic send/receive rates**
  - Divide traffic into ≈3 min windows and extract feature vectors
  - Devices cluster well by mean and standard deviation of traffic volumes within 3 second samples of windows
  - *k*-nearest-neighbors classifier has >95% accuracy with tested IoT devices for a range of window sizes, sampling rates, and values of *k*
- Additional data, higher-dimensionality feature vectors, and more complex classifiers likely to be effective for more devices

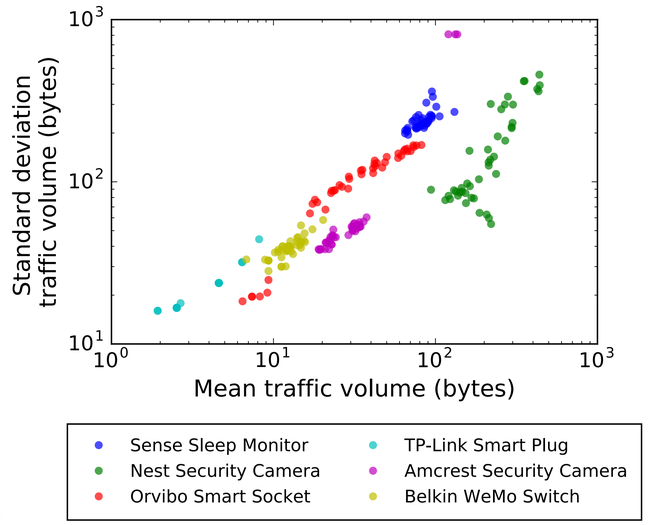

IoT devices reveal private user behavior in network traffic rates

Sleeping habits – Sense Sleep Monitor

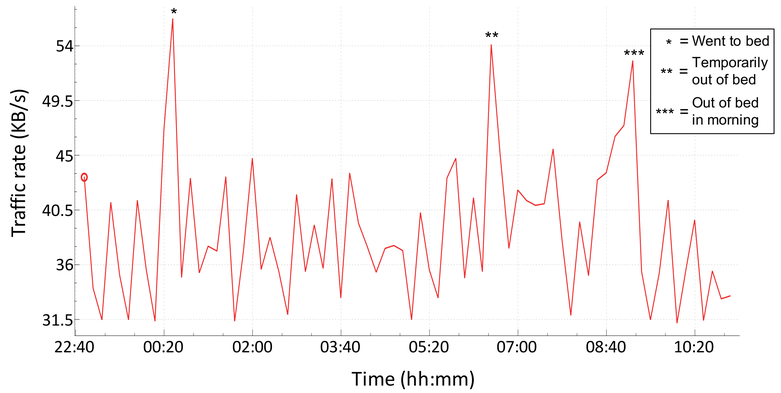

Security monitoring – Nest Camera

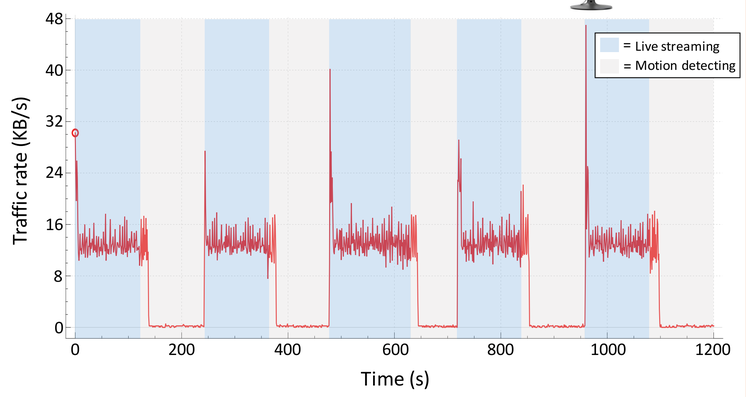

Appliance power cycle – Belkin WeMo Switch

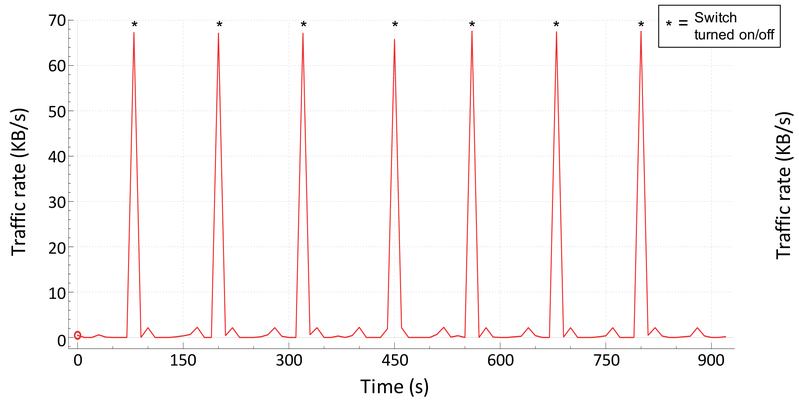

Motion in home – Nest Camera

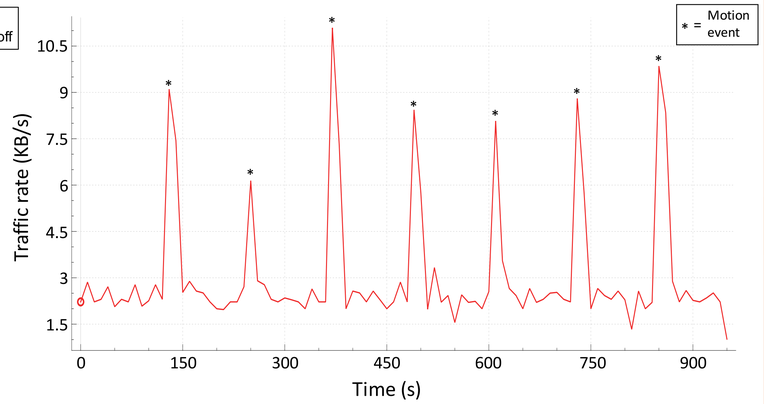

Proposed Solutions

- **Firewall devices** to prevent adversary from collecting traffic rate data
  - Difficult to determine which encrypted flows are essential for device function and which can be safely blocked
  - Generating effective firewall rules would require manufacturer support
- **Tunnel traffic** through VPN to prevent adversary from separating flows from individual devices
  - Pushes problem to VPN exit point
  - Ineffective vs Wi-Fi eavesdropper
- **Shape traffic** to prevent accurate behavior inference
  - Less time-sensitive devices can delay cloud updates
- **Inject traffic** mimicking user behaviors to reduce adversary's confidence in behavior inferences

Acknowledgements

This research was funded by the Department of Defense through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program, a Google Faculty Research Award, and the National Science Foundation

PRINCETON UNIVERSITY