OS-level Side Channels without Procfs: Exploring Cross-App Information Leakage on iOS

Xiaokuan Zhang[1], Xueqiang Wang[2], Xiaolong Bai[3],
Yinqian Zhang[1] and XiaoFeng Wang[2]

[1]The Ohio State University, [2]Indiana University Bloomington, [3]Tsinghua University

THE OHIO STATE UNIVERSITY

INDIANA UNIVERSITY
BLOOMINGTON

清華大學
Tsinghua University

Mobile Side-Channel Attacks

- Side-channel Attack: make use of seemingly harmless information to infer sensitive information

Cache Side Channels

Sensor-based Side Channels

OS-level Side Channels

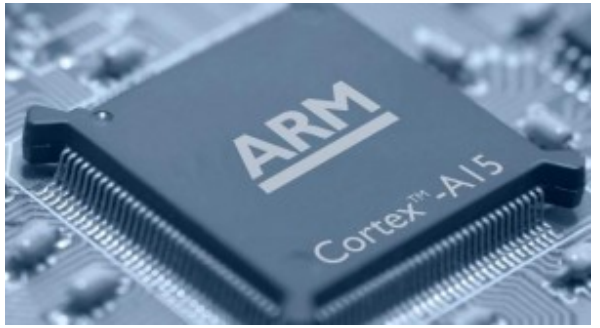

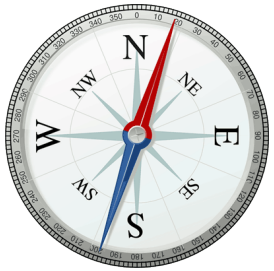

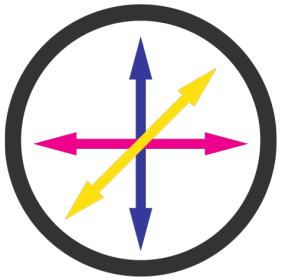

OS-level Side-Channel Attacks on Android

- Malicious app running in the background, calling APIs

- Procfs: system statistics
  - virtual/physical memory, network traffic, CPU usage info, …

```
zxk@zxk-VirtualBox:~$ ls /proc
1     1498  1776  1957  2055  2226  2421  4    526  65  769  96          fb           locks       stat
10    15    18    1961  2061  2230  2476  401  53   66  77   97          filesystems  mdstat      swaps
1056  150   1870  1962  2064  2245  2499  47   54   67  78   acpi        fs           meminfo     sys
11    154   1881  1966  2090  2246  25    471  55   68  8    asound      interrupts   misc        sysrq-trigger
1102  1542  1886  1967  2099  2251  2524  475  555  69  82   buddyinfo   iomem        modules     sysvipc
1134  155   19    1980  21    2255  2535  48   56   693 866  bus         ioports      mounts      thread-self
1197  156   1911  1984  2129  2271  2544  49   561  7   870  cgroups     irq          mtrr        timer_list
12    157   1912  2     2143  2277  2545  493  57   70  877  cmdline     kallsyms     net         timer_stats
1221  158   1913  20    2164  23    2558  5    58   71  878  consoles    kcore        pagetypeinfo tty
1234  16    1916  2041  2176  2364  26    50   59   714 881  cpuinfo     keys         partitions  uptime
1286  1655  1921  2045  2189  2373  28    503  6    72  9    crypto      key-users    sched_debug version
13    169   1925  2046  2198  2387  29    507  60   726 938  devices     kmsg         schedstat   version_signature
1308  17    1929  2047  22    2399  3     51   61   73  945  diskstats   kpagecgroup  scsi        vmallocinfo
1333  170   1931  2048  2202  24    30    517  62   74  95   dma         kpagecount   self        vmstat
14    1704  1941  2051  2205  2404  31    52   63   75  951  driver      kpageflags   slabinfo    zoneinfo
148   1774  1954  2054  2207  2411  397   525  64   76  956  execdomains loadavg      softirqs
```

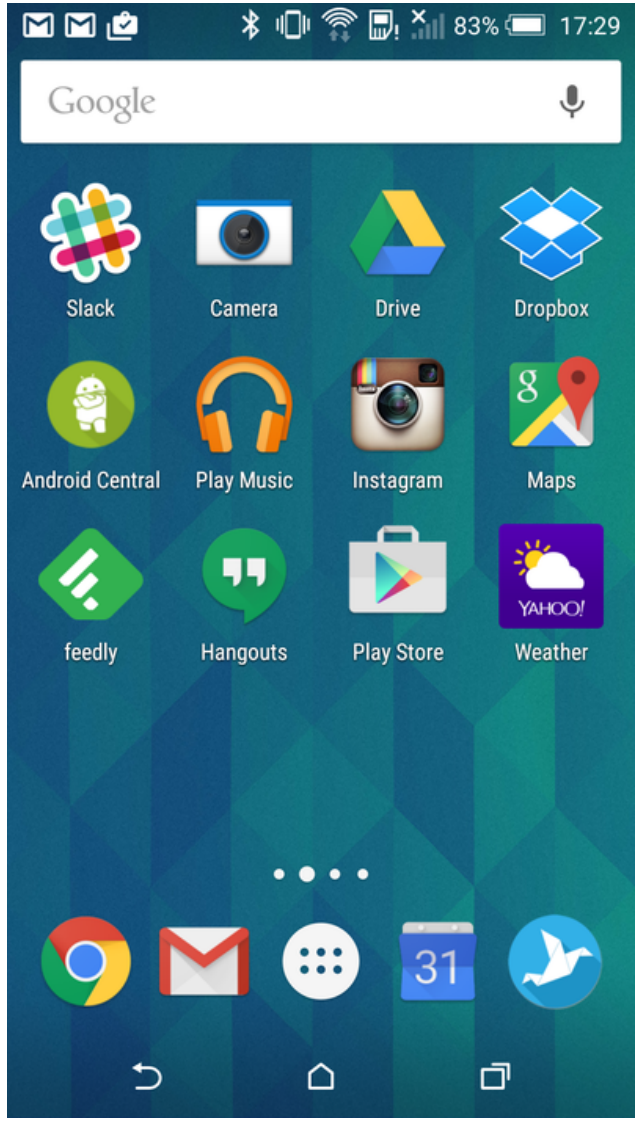

OS-level Side-Channel Attacks on iOS

- No Procfs providing system stat

```
zxk@zxk-VirtualBox:~$ ls /proc
1     1498  1776  1957  2055  2226  2421  4    526  65   769  96          fb           locks        stat
10    15    18    1961  2061  2230  2476  401  53   66   77   97          filesystems  mdstat       swaps
1056  150   1870  1962  2064  2245  2499  47   54   67   78   acpi        fs           meminfo      sys
11    154   1881  1966  2090  2246  25    471  55   7    8    asound      interrupts   misc         sysrq-trigger
1102  1542  1886  1967  2099  2251  2524  475  555  [illegible]  bus         iomem        modules      sysvipc
1134  155   19    1980  21    2255  2535  48   56   [illegible]  [illegible]  ioports      mounts       thread-self
1197  156   1911  1984  2129  2271  2544  49   561  7    [illegible]groups  irq          mtrr         timer_list
12    157   1912  2     2143  2277  2545  493  57   70   [illegible]ndline  kallsyms     net          timer_stats
1221  158   1913  20    2164  23    2558  5    58   71   [illegible]soles   kcore        pagetypeinfo tty
1234  16    1916  2041  2176  2364  26    50   59   7[illegible]51  [illegible]fo  keys         partitions   uptime
1286  1655  1921  2045  2189  2373  28    503  6    [illegible]9  cr[illegible]  key-users    sched_debug  version
13    169   1925  2046  2198  2387  29    507  60   [illegible]6  938  devic[illegible]  kmsg         schedstat    version_signature
1308  17    1929  2047  22    2399  3     51   [illegible]73  945  diskst[illegible]  kpagecgroup  scsi         vmallocinfo
1333  170   1931  2048  2202  24    30    517  [illegible]  74   95   dma         kpagecount   self         vmstat
14    1704  1941  2051  2205  2404  31    52   63   75   951  driver      kpageflags   slabinfo     zoneinfo
148   1774  1954  2054  2207  2411  397   525  64   76   956  execdomains loadavg      softirqs
```

- No unauthorized cross-app query

Is it possible to conduct OS-level side-channel attacks on iOS?

Outline

1. Side-channel Attack Vectors on iOS
2. Attack 1: Classifying User Activities
3. Attack 2: Detecting Sensitive In-App Activities
4. Attack 3: Bypassing Sandbox Restrictions
5. Practical Issues
6. Countermeasures
7. Conclusion

Threat Model

- Monitoring app:
  - User downloads it from App Store
  - Audio player

New Attack Vectors

- Host_statistics64(): Gl

```
kern_return_t host_statistics64(host_t
host_priv, host_flavor_t flavor,
host_info64_t host_info64_out,
mach_msg_type_number_t
*host_info64_outCnt);
```

- Getifaddrs():

```
int getifaddrs(struct ifaddrs **ifap);
```

- [NSFileManager fileExistsAtPath:]: The existence of a file/directory

```
- (BOOL)fileExistsAtPath:(NSString *)path;
```

Outline

1. Side-channel Attack Vectors on iOS
2. Attack 1: Classifying User Activities
3. Attack 2: Detecting Sensitive In-App Activities
4. Attack 3: Bypassing Sandbox Restrictions
5. Practical Issues
6. Countermeasures
7. Conclusion

Classifying User Activities --- Example Trace

- Calling APIs to get time series A
  - Host_statistics64()
  - Getifaddrs()
- Plotting diff series: A[i] – A[i-1]

Time series leak information!!!

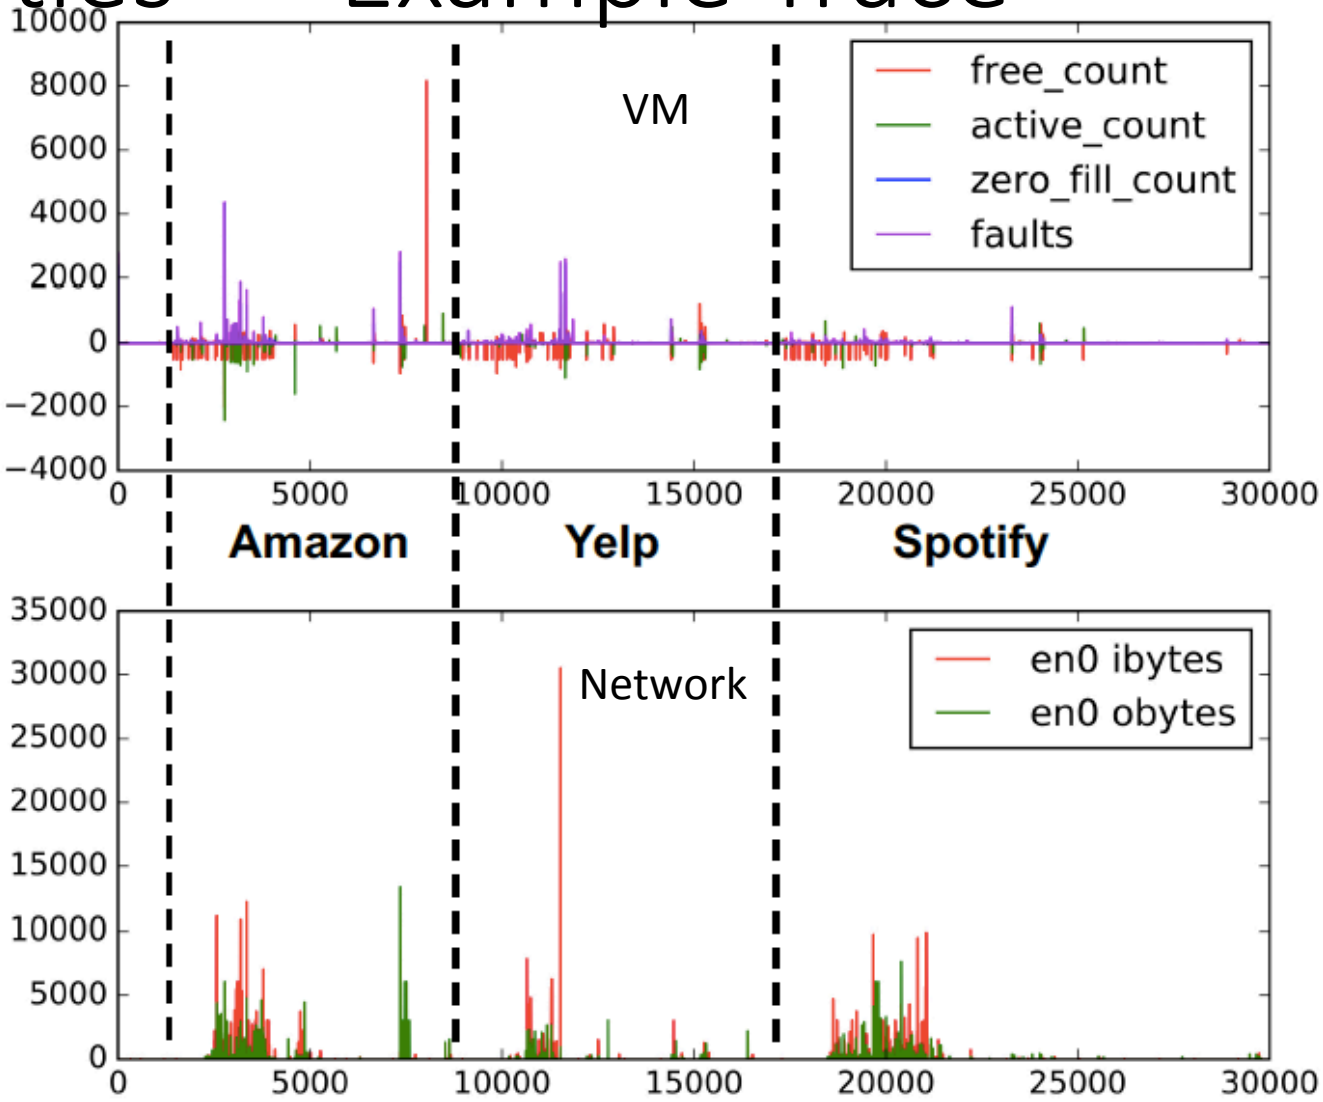

Classifying User Activities --- Example Trace

How to combine multiple time series to perform inference attacks?

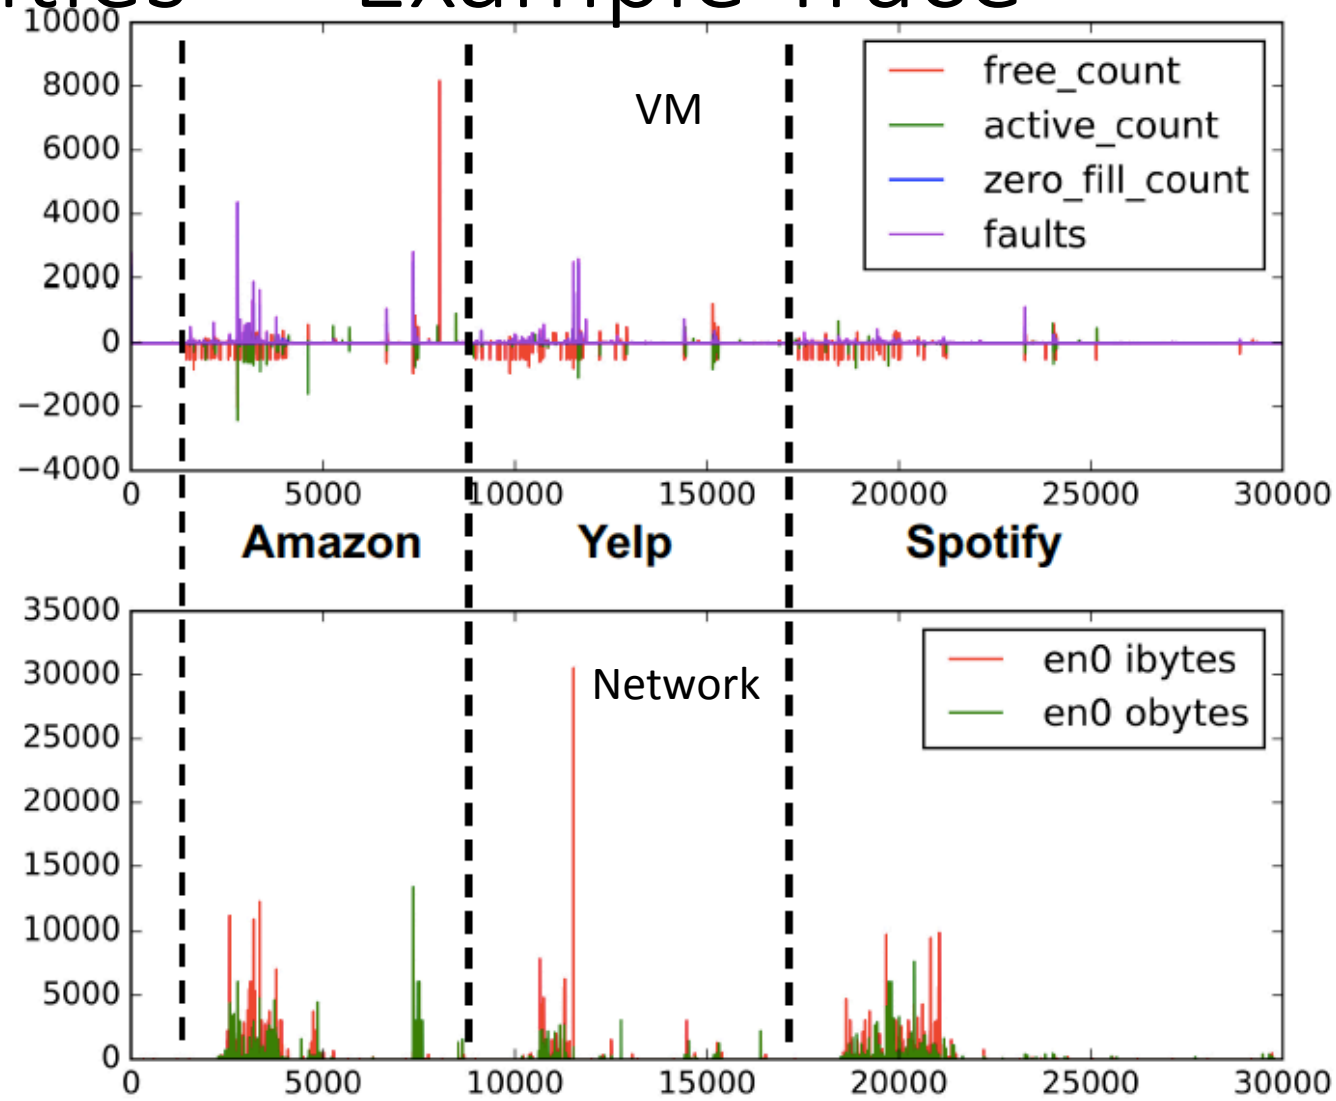

Classifying User Activities --- Example Trace

How to combine multiple time series to perform inference attacks?

- Requirements:
  - Combining multiple time series
  - Reducing the dimension
- Major components:
  - SAX (Keogh et al., 2002)
  - BOP (Lin et al., 2009)
  - LibSVM (Chang et al., 2011)

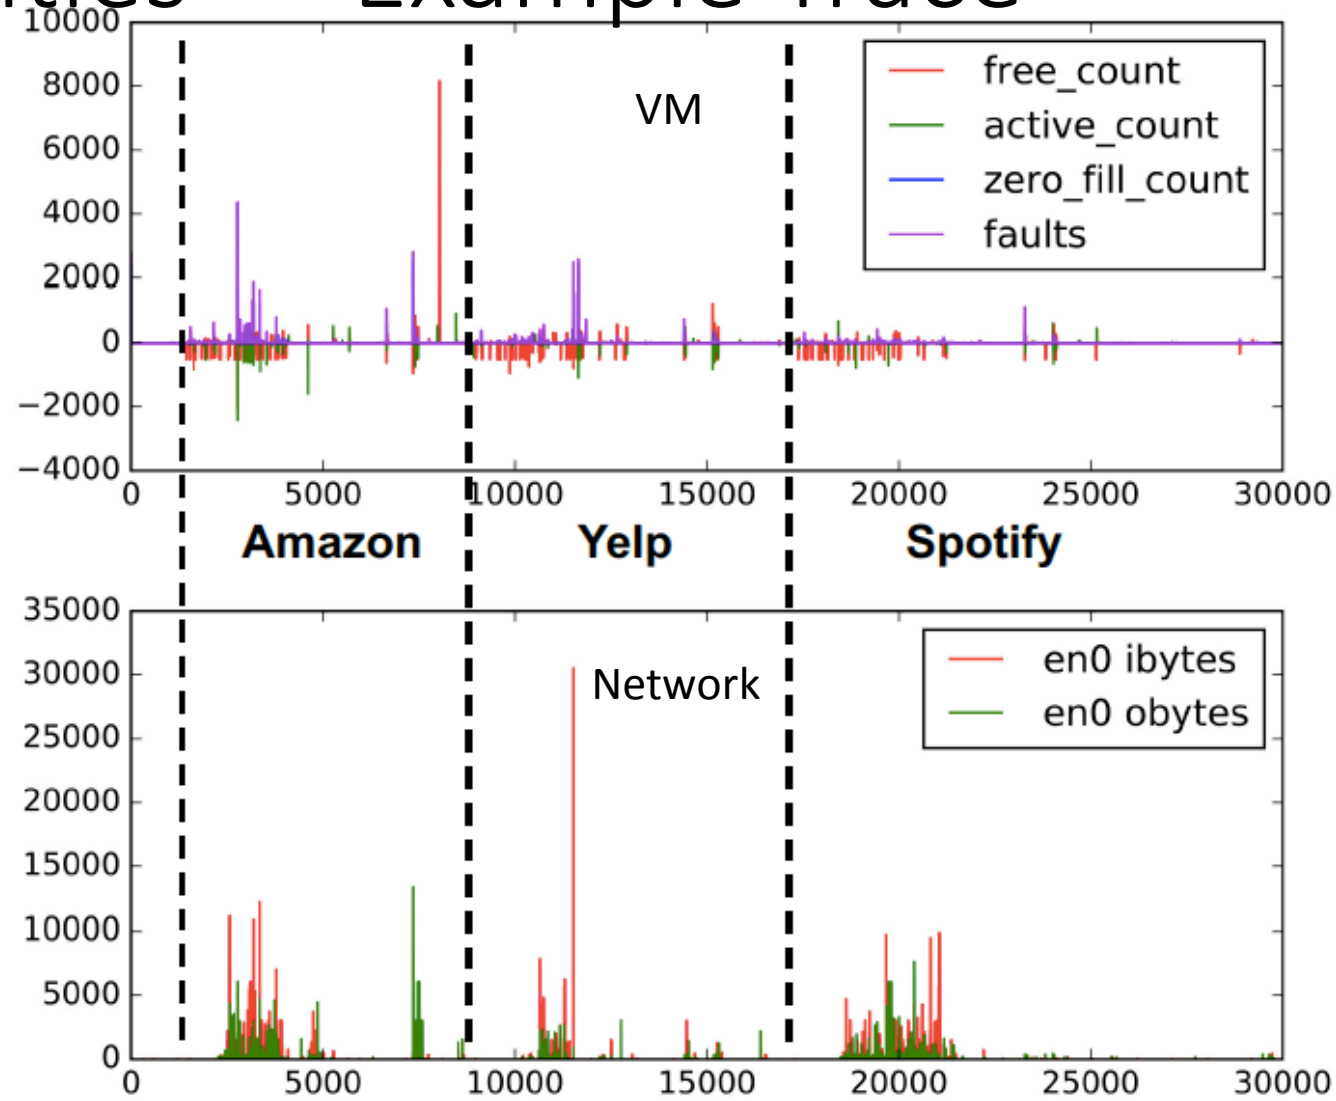

Classifying User Activities --- Case Studies

- Device: jailbroken iPhone 7 with iOS 10.1.1

- Automated using Cycript

- Monitoring app:
  - running in the background
  - calling APIs at a rate of 1000/s

Classifying User Activities --- Case Studies

- Foreground Apps:
  - 100 apps from Top Charts + 20 pre-installed apps
  - Top N accuracy: the percentage of the test samples being correctly labeled by one of the top N predicted classes by the classifier

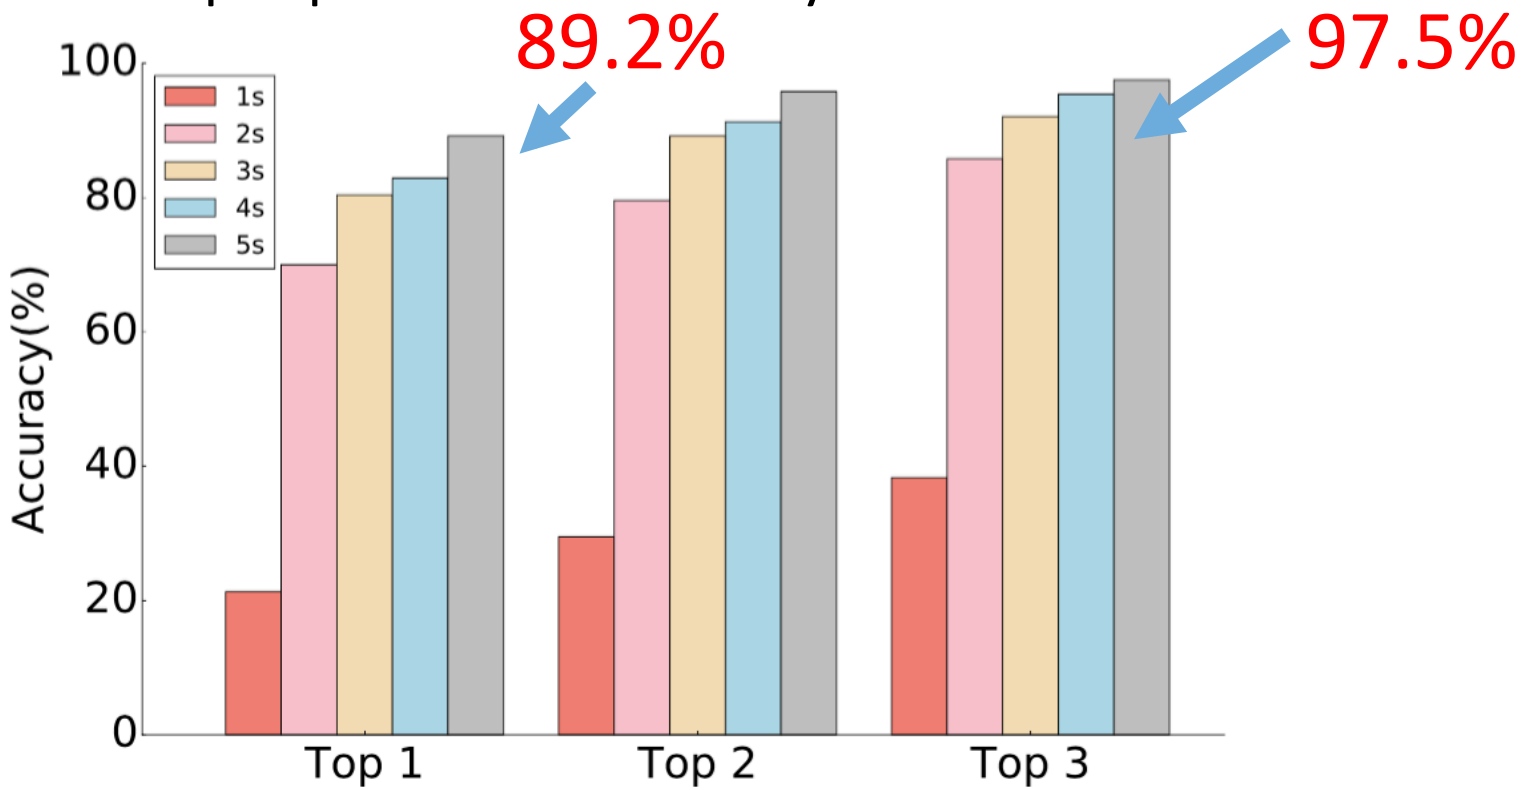

Classifying User Activities --- Case Studies

- Safari Websites

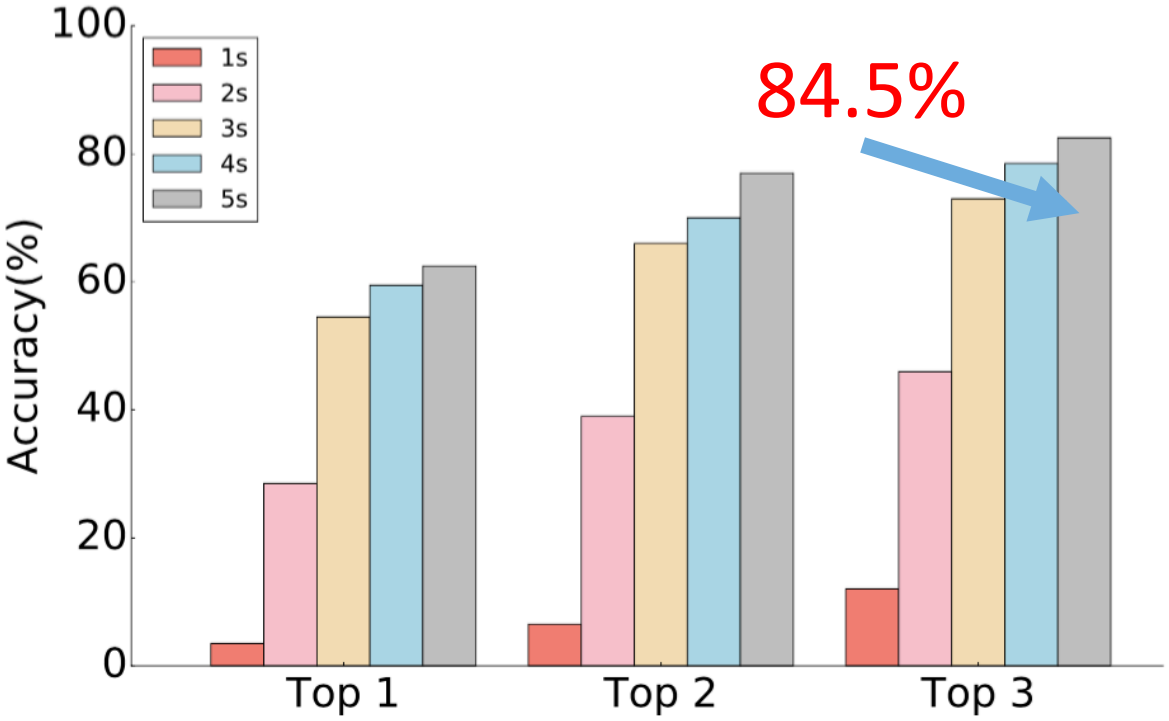

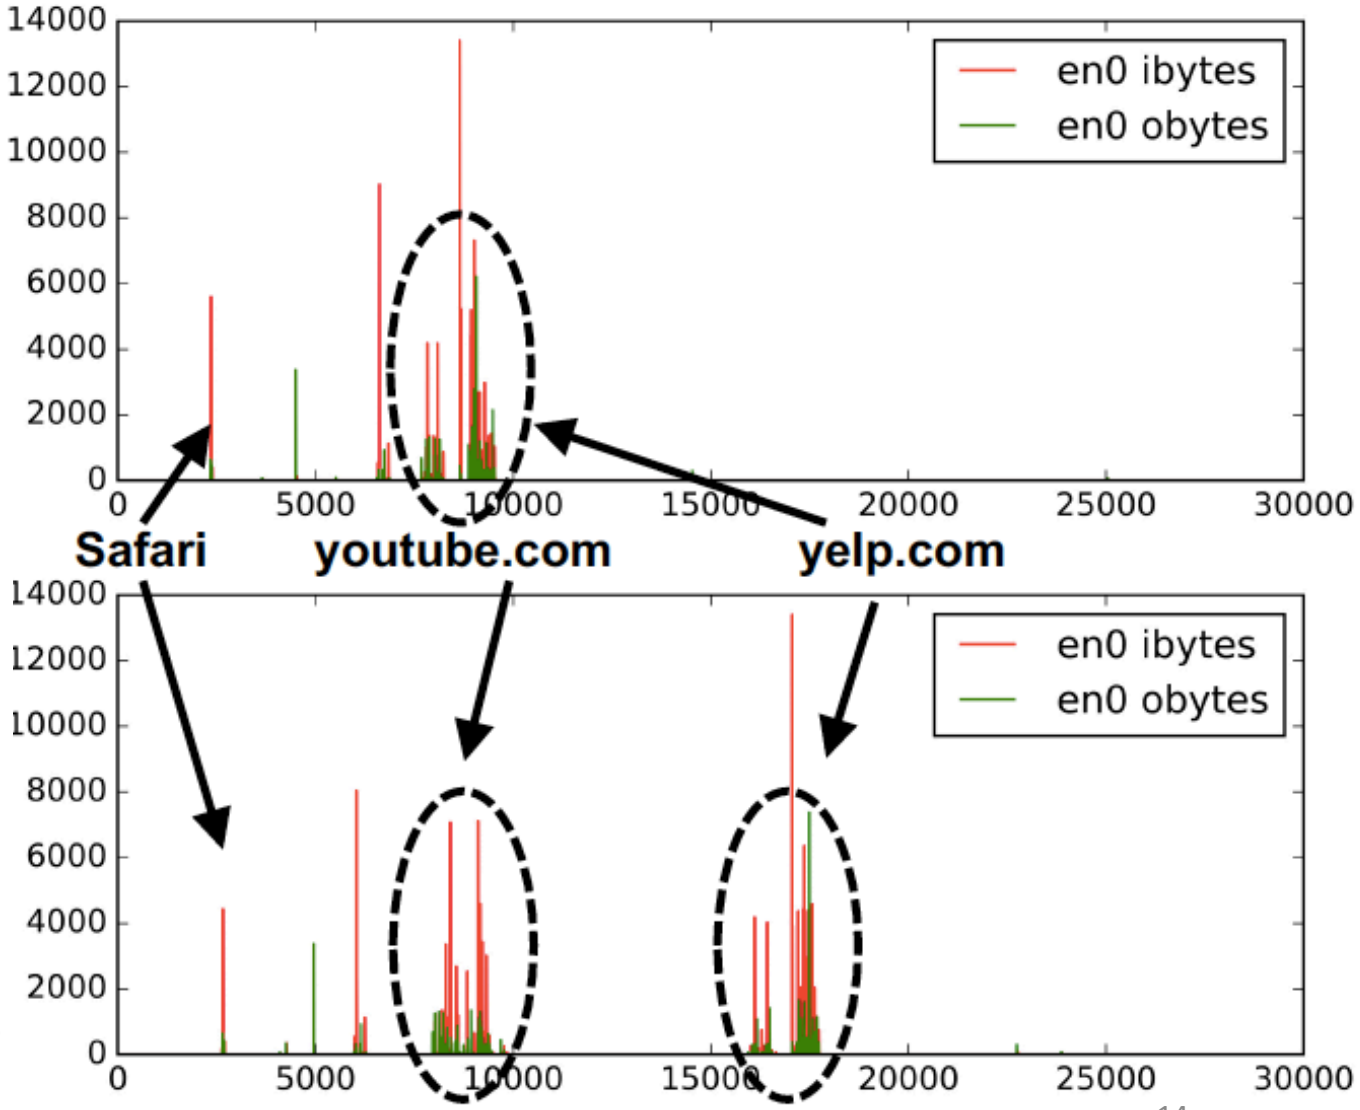

Outline

1. Side-channel Attack Vectors on iOS
2. Attack 1: Classifying User Activities
3. Attack 2: Detecting Sensitive In-App Activities
4. Attack 3: Bypassing Sandbox Restrictions
5. Practical Issues
6. Countermeasures
7. Conclusion

Detecting Sensitive In-App Activities

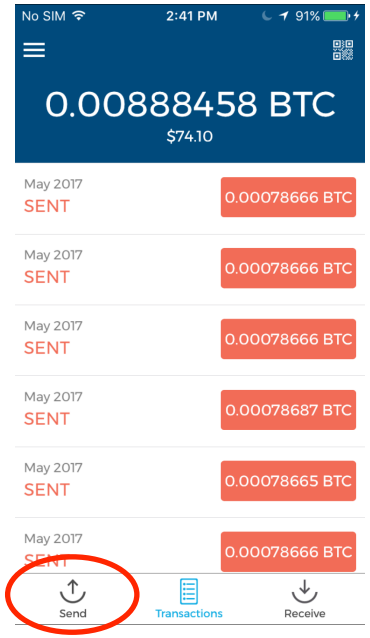

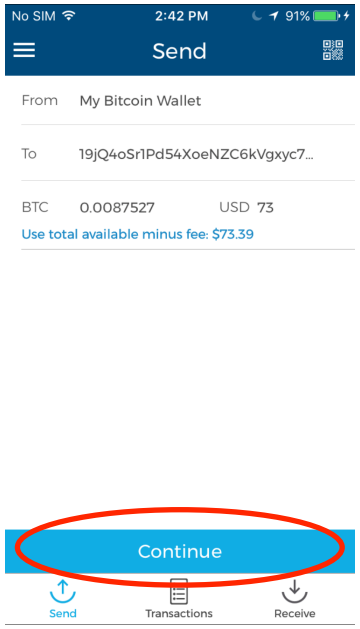

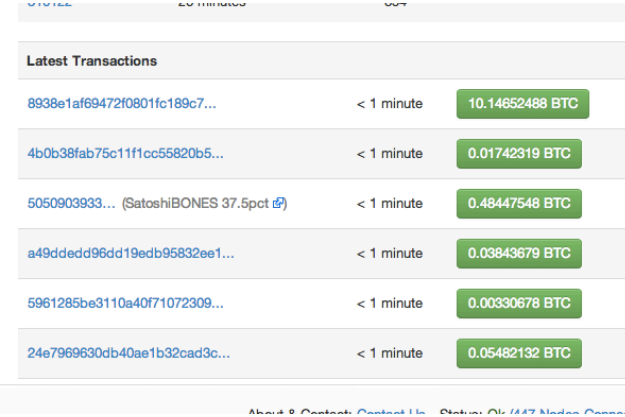

Blockchain.info

Detecting Sensitive In-App Activities --- Attack Methods

- Identify critical events

- Correlates with public records

Detecting Sensitive In-App Activities --- Case Studies

- Target: *Blockchain Wallet* App

- Goal: identify *payment* event (idx: 0)

$$d(\vec{X}_t, \vec{S}_t) = \sum_{k=1}^{l} \frac{1}{w_k} \cdot \mathrm{DTW}\left(\vec{X}_t^k, \vec{S}_t^k\right)$$

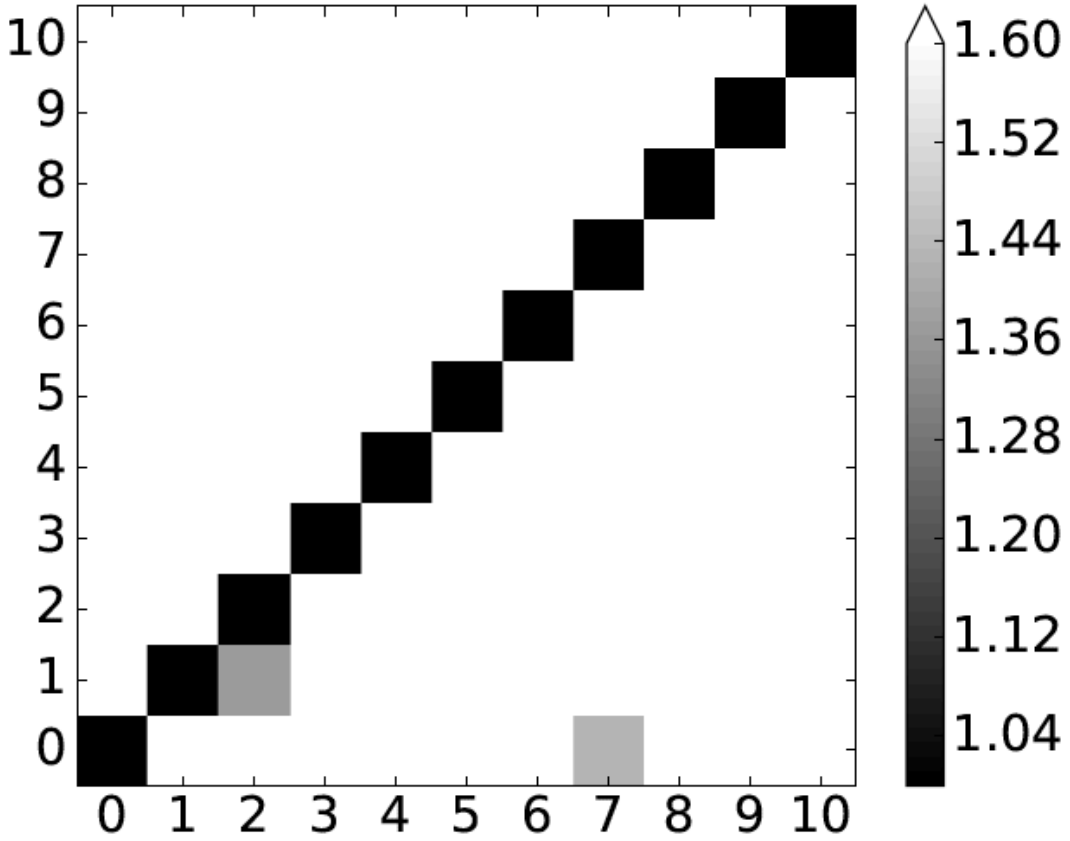

Detecting Sensitive In-App Activities --- Case Studies

- Target: *Blockchain Wallet* App

- Goal: identify *payment* event (idx: 0)

- Normalize the distance per row
  using cell(i,i) as the base (diagonal)

$$d(\vec{X}_t, \vec{S}_t) = \sum_{k=1}^{l} \frac{1}{w_k} \cdot \mathrm{DTW}\left(\vec{X}_t^k, \vec{S}_t^k\right)$$

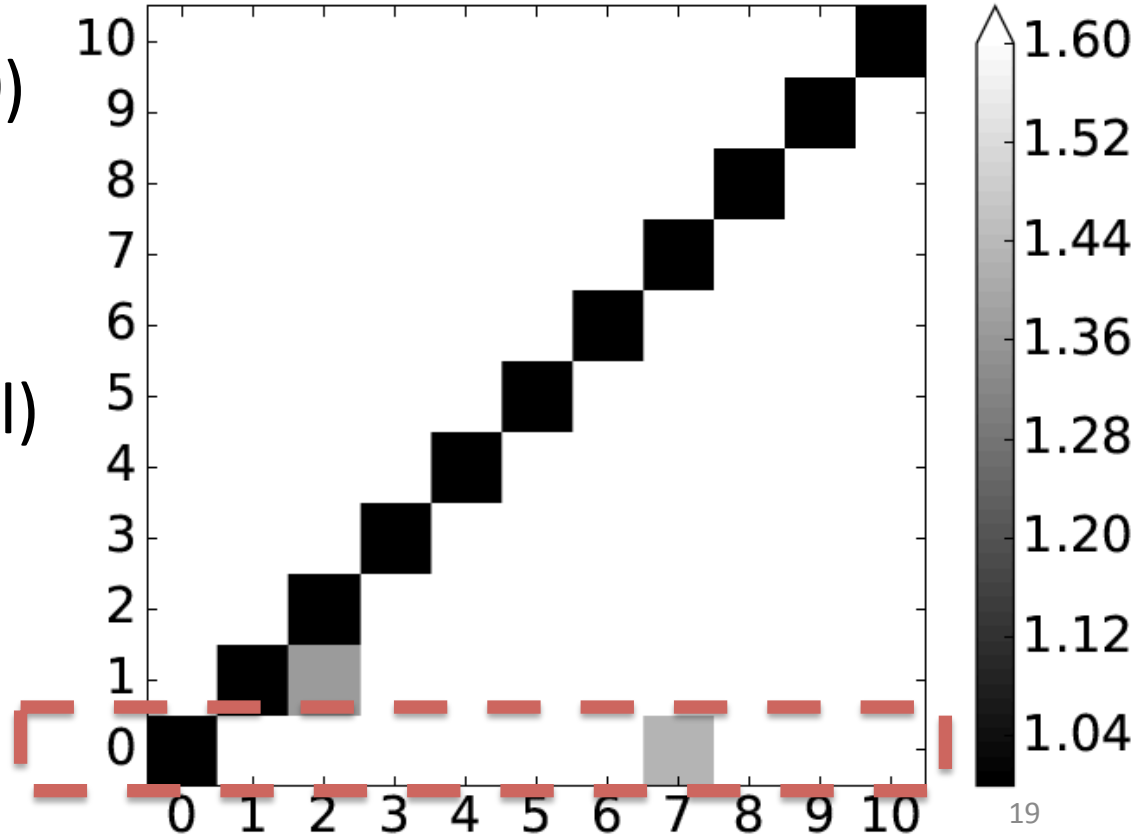

Detecting Sensitive In-App Activities --- Case Studies

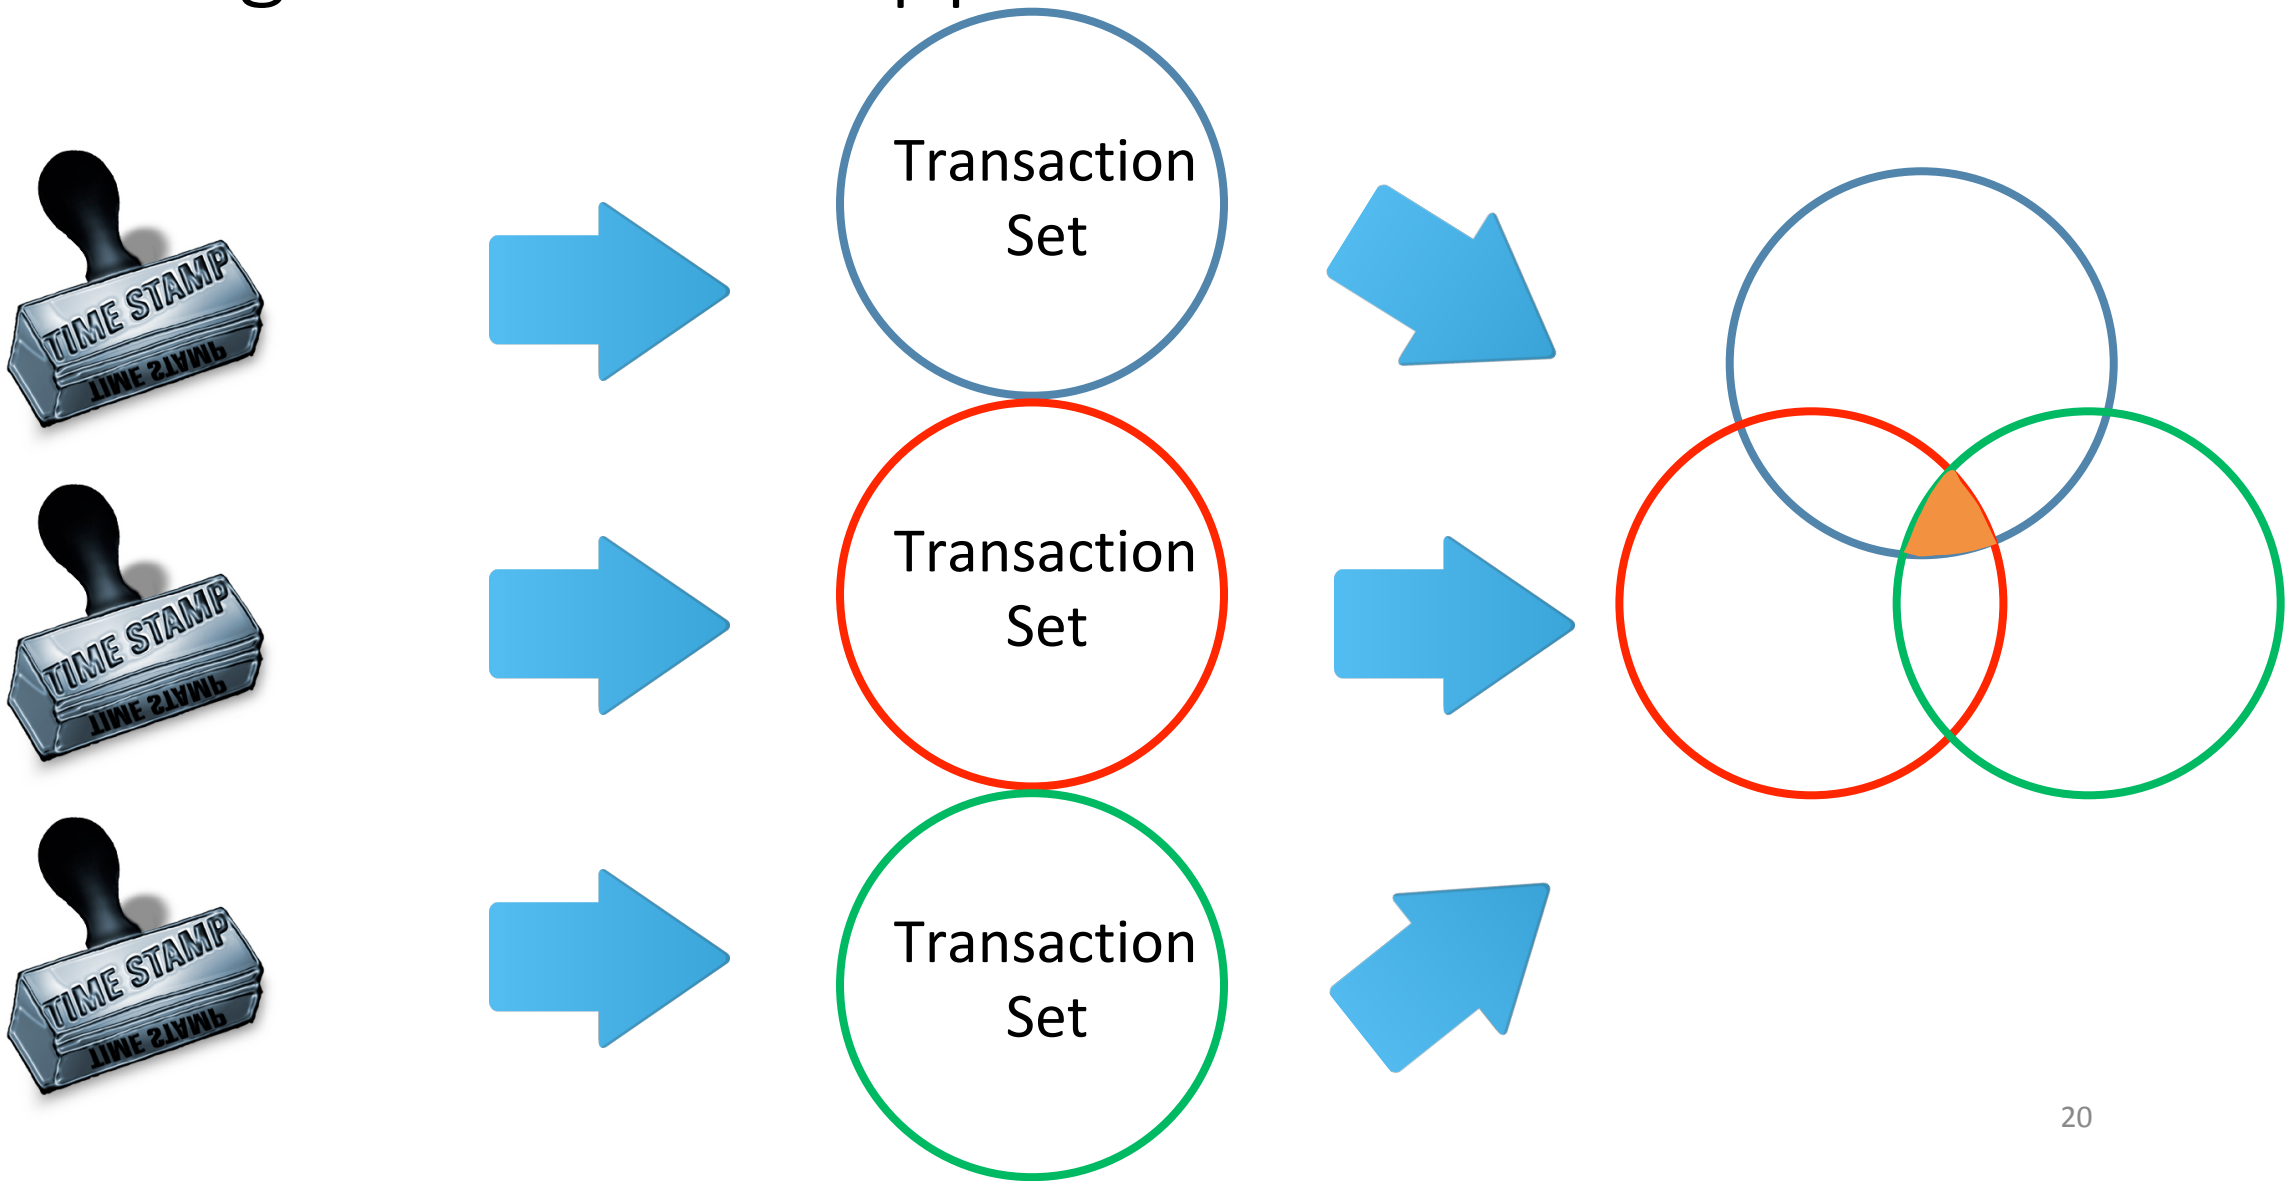

Detecting Sensitive In-App Activities --- Case Studies

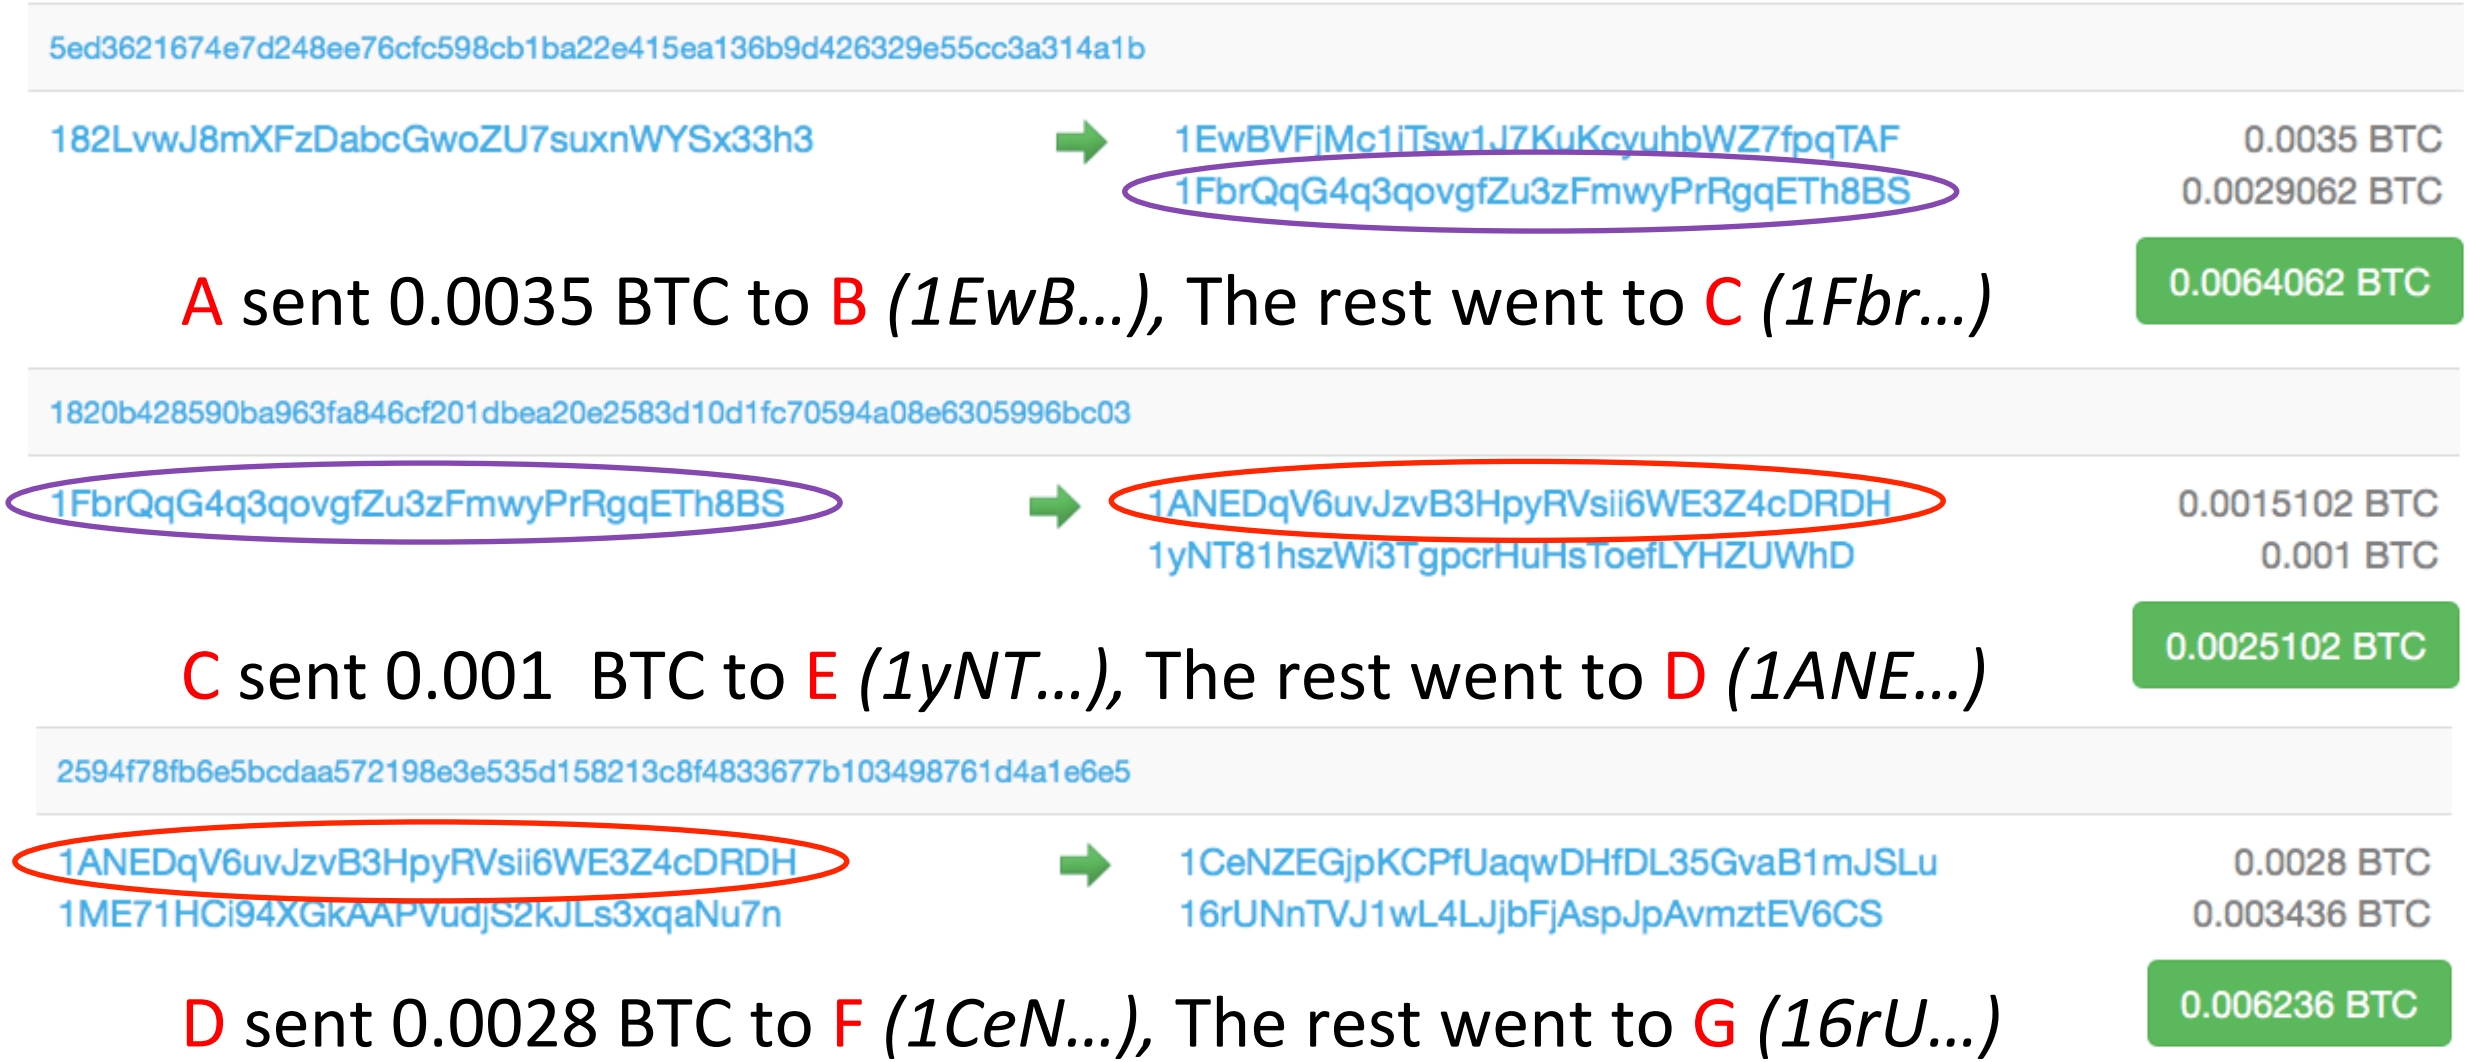

Detecting Sensitive In-App Activities --- Case Studies

- Other Targets: *Venmo / Twitter*

Outline

1. Side-channel Attack Vectors on iOS
2. Attack 1: Classifying User Activities
3. Attack 2: Detecting Sensitive In-App Activities
4. Attack 3: Bypassing Sandbox Restrictions
5. Practical Issues
6. Countermeasures
7. Conclusion

Bypassing Sandbox Restrictions --- Attack Methods

- Device: non-jailbroken iPhone 7 with iOS 10.2.1
- Execution time of FileExistAtPath

Huge Difference!!!

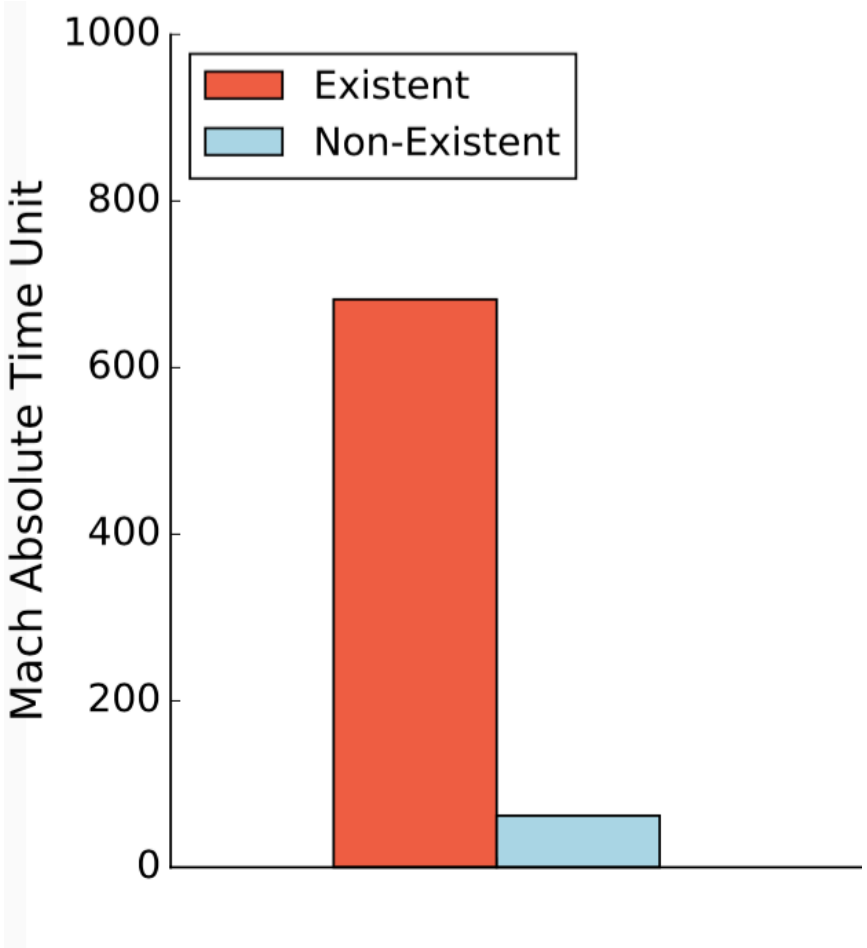

Bypassing Sandbox Restrictions --- Case Studies

- Detect whether an app has been installed

DivorceForce

AsthmaMD

Pregnancy+

Sugar Sense

Bypassing Sandbox Restrictions --- Case Studies

- Push notifications:
  - .pushstore file with the bundle identifier as its name will be created in a specific directory
  - (/var/mobile/Library/SpringBoard/PushStore/com.g
    Gmail app)
- Dynamically registered home screen quick actio
  - .plist file with the bundle identifier as its name will b
    var/mobile/Library/SpringBoard/Application Shortcu
    Gmail app)
- Top 150 apps in App Store's "Top Charts" (Aug. 2
  - Push notification: 67 (44.7%)
  - dynamically registered home screen quick actions: 44 (31.3%)

Bypassing Sandbox Restrictions --- Case Studies

- Other cases: number of photos/memos

- Generic approach to detect files

Outline

1. Side-channel Attack Vectors on iOS
2. Attack 1: Classifying User Activities
3. Attack 2: Detecting Sensitive In-App Activities
4. Attack 3: Bypassing Sandbox Restrictions
5. Practical Issues
6. Countermeasures
7. Conclusion

Practical Issues

- App Store Vetting
  - Disguised as an *Audio Player*
  - Passed the vetting

- Power Consumption
  - Device: jailbroken iPhone 7 with iOS 10.1.1
  - 60 min: 5% battery was consumed

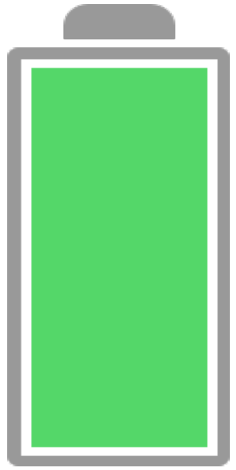

Practical Issues --- Cross-device Attack Feasibility

training device: Device A
iOS 10.1.1

testing device: Device B
Non-jailbroken iOS 10.2.1

Practical Issues --- Cross-device Attack Feasibility

- Test set: Randomly select 20 third-party apps
- Redo Foreground Apps Experiment

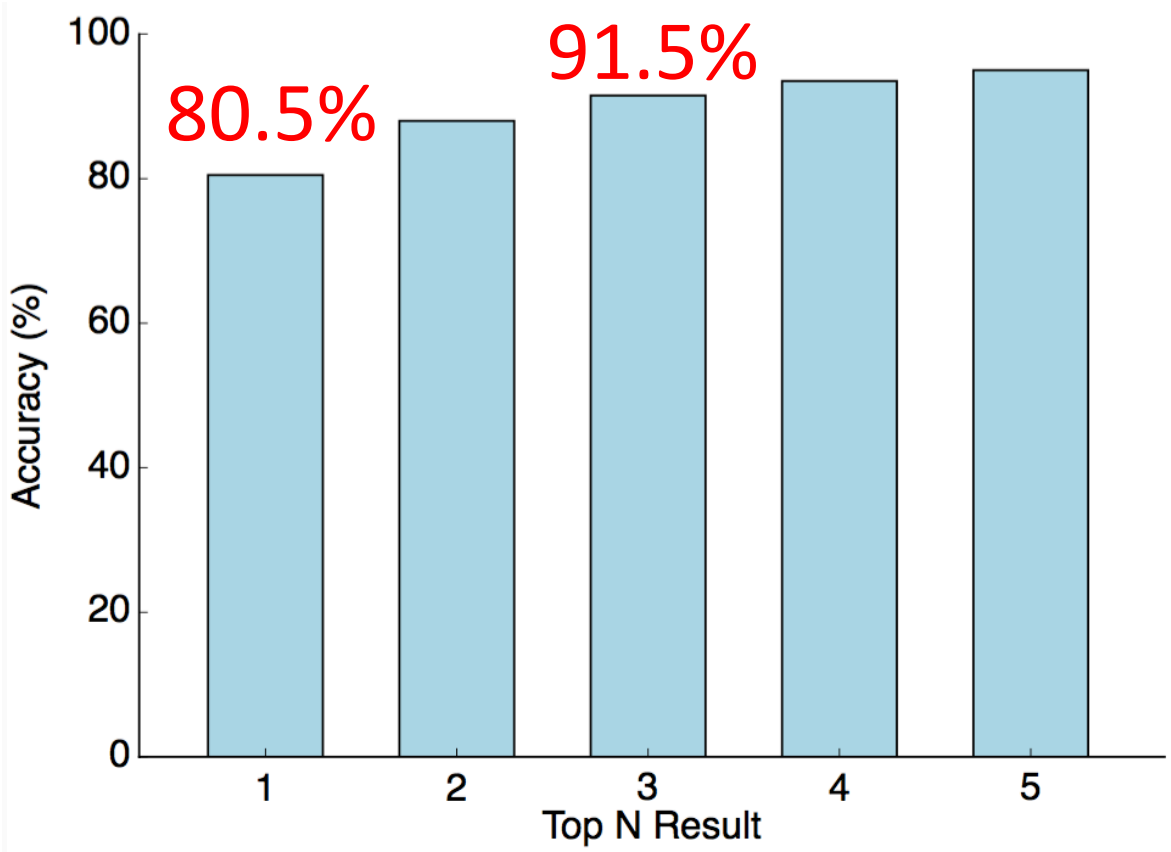

Practical Issues --- Cross-device Attack Feasibility

- Target: *Blockchain Wallet*

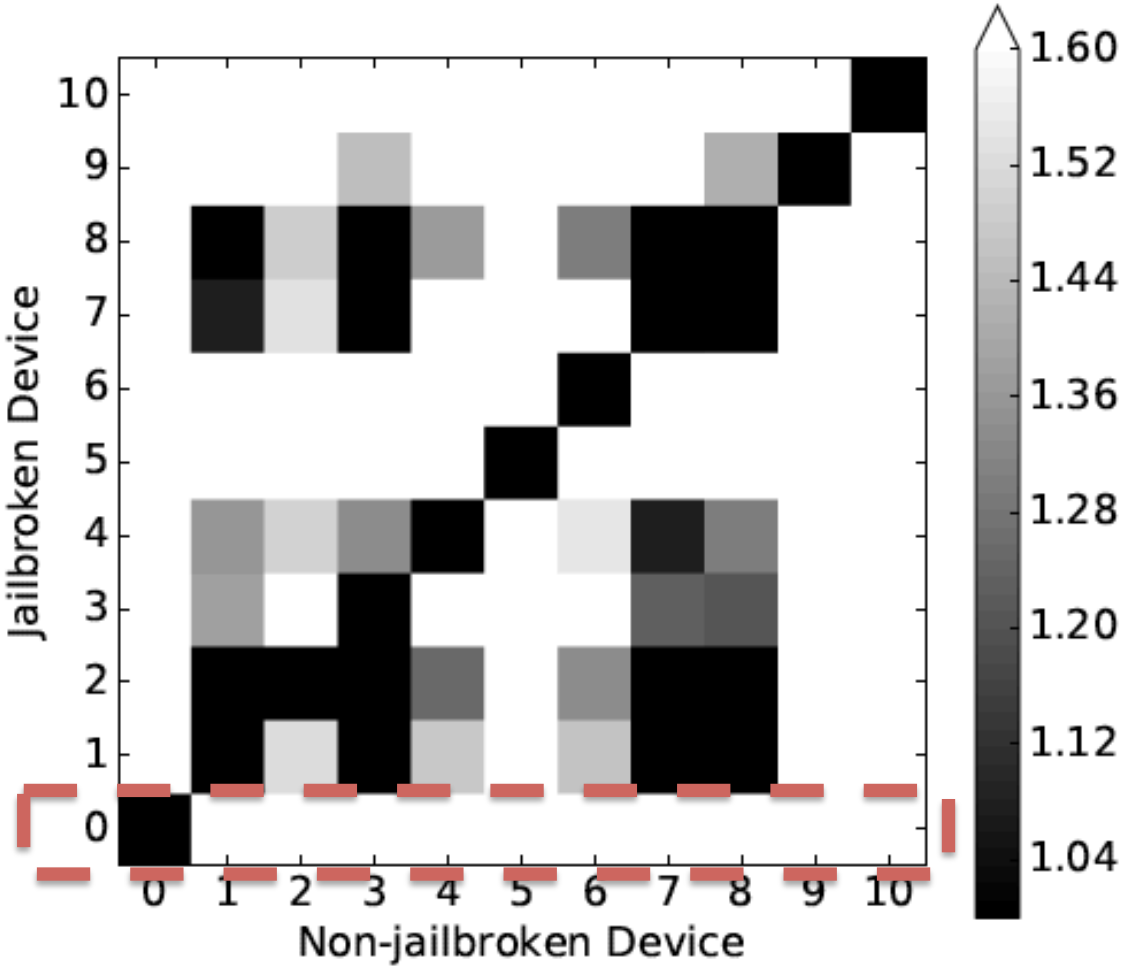

Outline

1. Side-channel Attack Vectors on iOS
2. Attack 1: Classifying User Activities
3. Attack 2: Detecting Sensitive In-App Activities
4. Attack 3: Bypassing Sandbox Restrictions
5. Practical Issues
6. Countermeasures
7. Conclusion

Countermeasures

- Rate Limiting: limit the sampling rate
  - Filter the data and only keep every *(1000/N)th* data point
  - Re-evaluate the foreground app classification

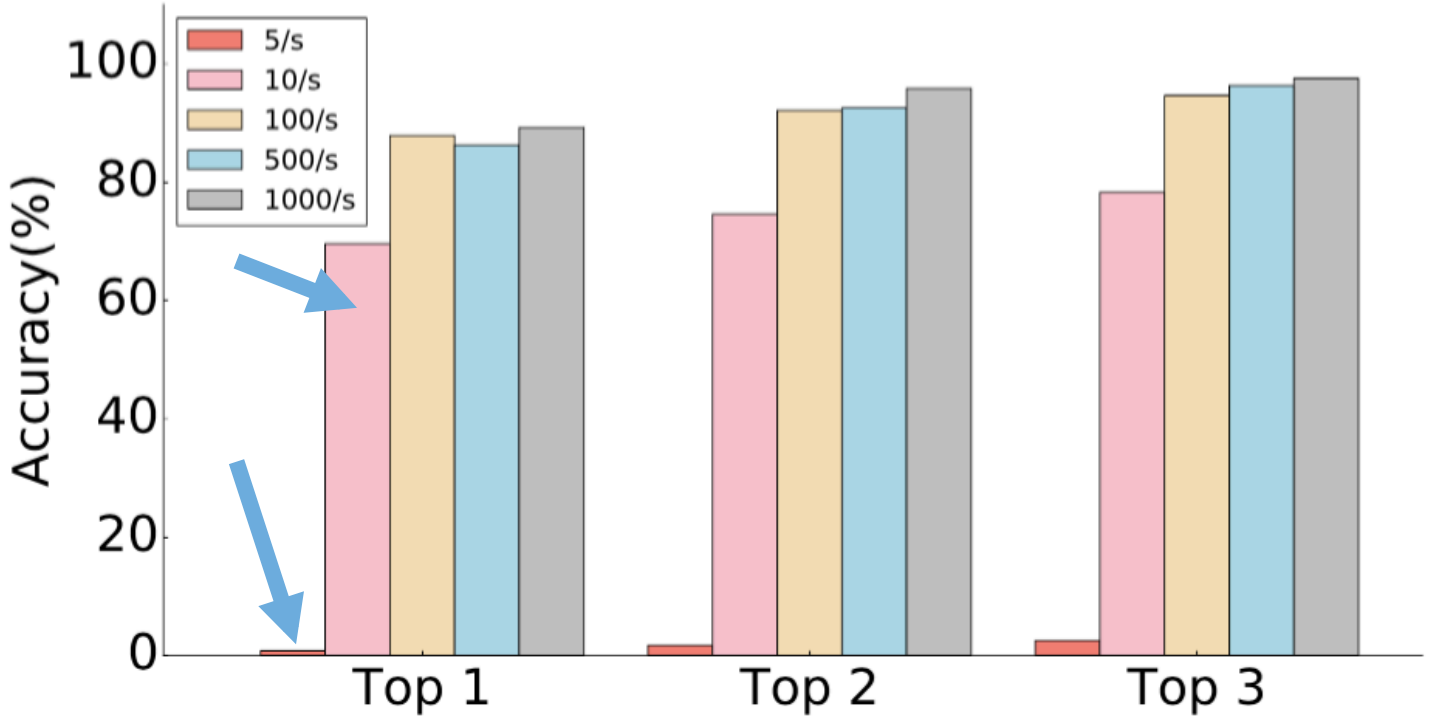

Implemented in iOS 11.1
for host_statistics64(): 2/s

Countermeasures

- Coarse-grained return values: masking the digits of return values
  - Mask 1/2/3 digits of all 6 features
  - Re-evaluate the foreground app classification

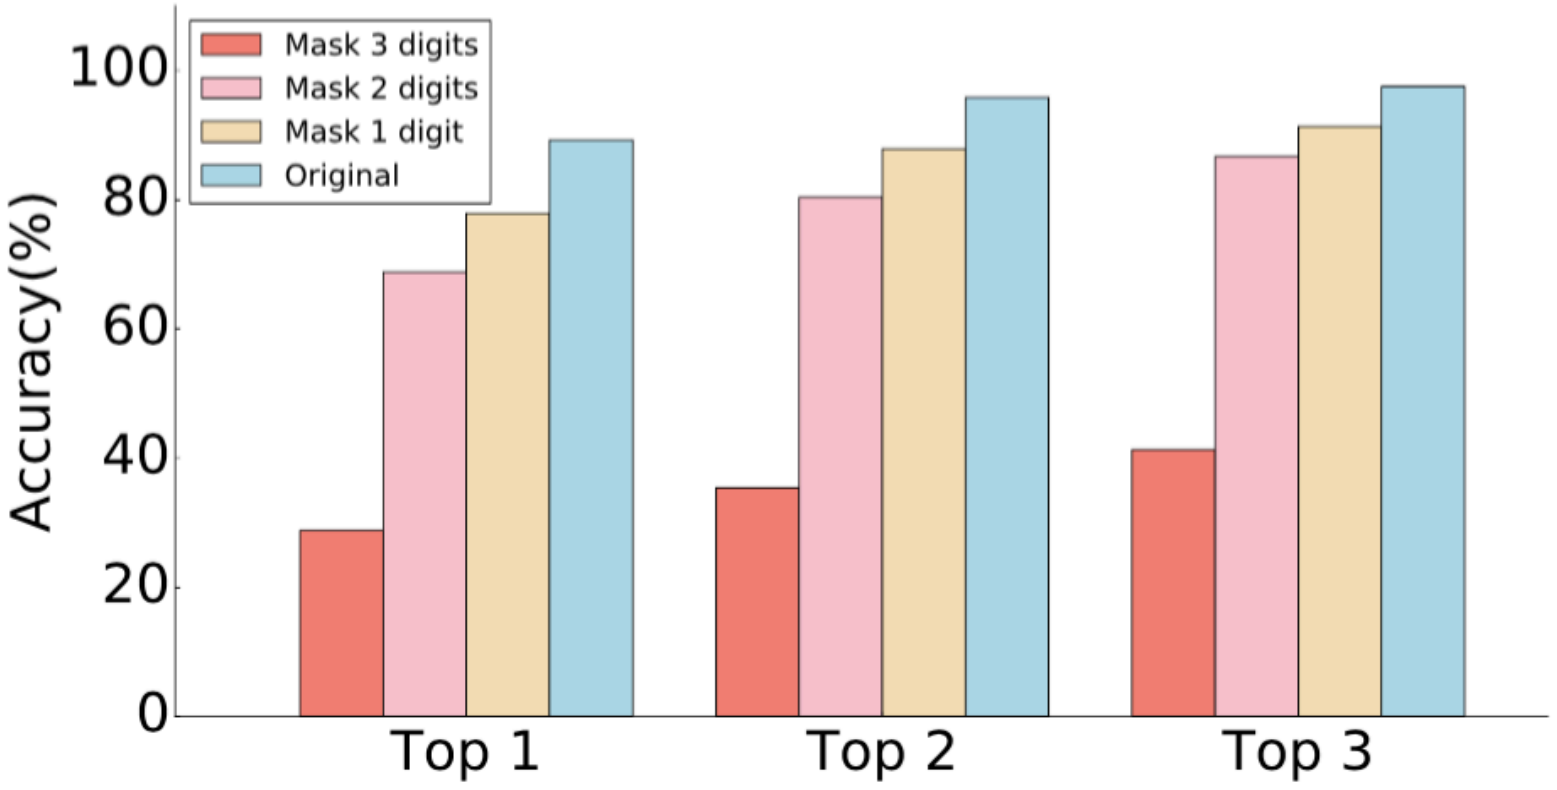

Original: 1234

Mask 1 digit: 1230

Mask 2 digits: 1200

Mask 3 digits: 1000

Countermeasures

- Coarse-grained return values: masking the digits of return values
  - Mask 1/2/3 digits of all 6 features
  - Re-evaluate the foreground app classification

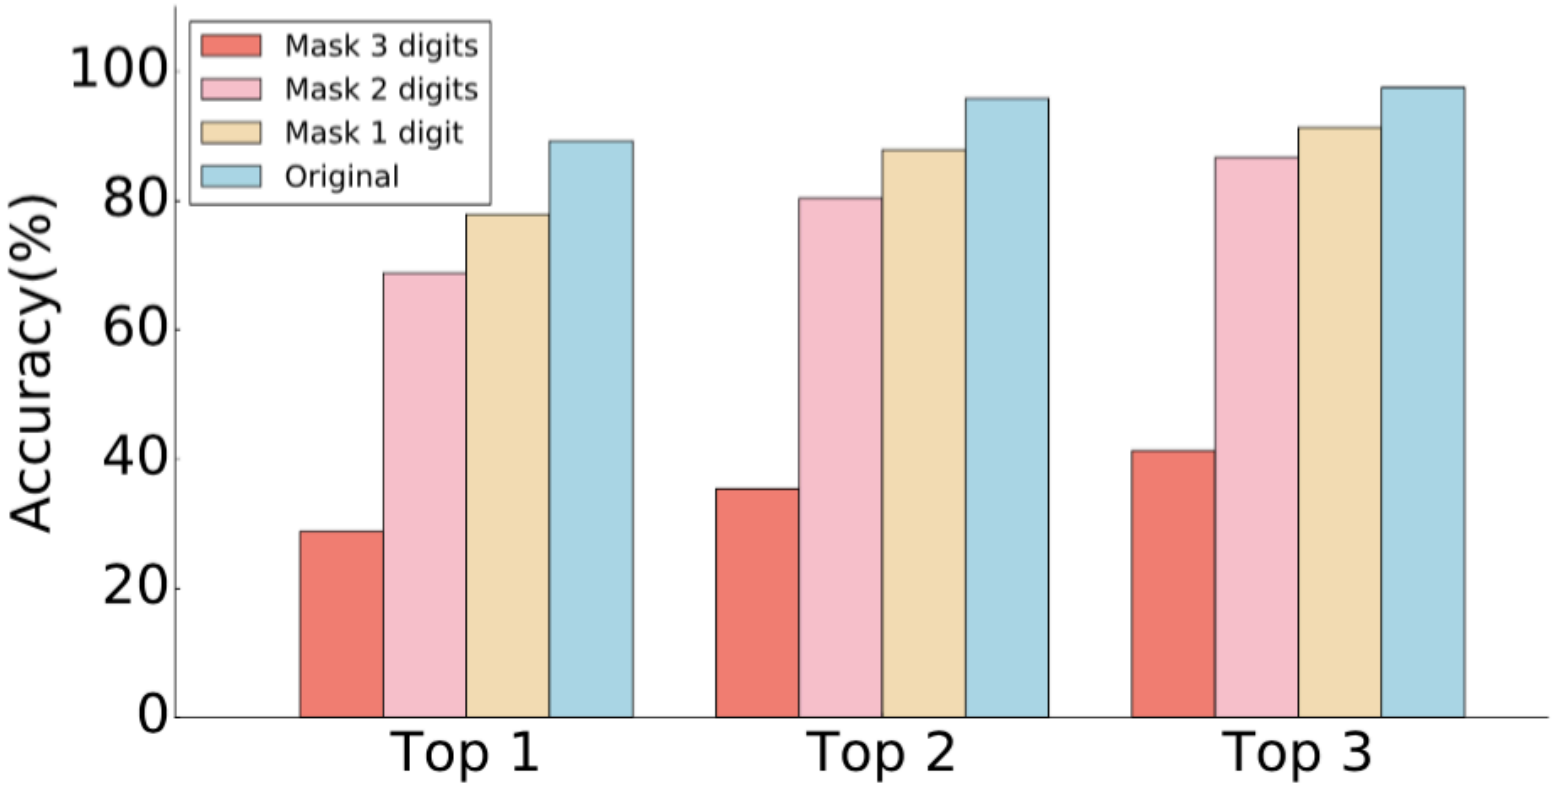

Implemented in iOS 11 for getifaddrs(): Round to 1KB

Countermeasures

- Eliminating the attack vectors
- Runtime detection
- Privacy-preserving statistics reporting
- Removing the fileExistsAtPath timing channel

fileExistsAtPath timing channel has been eliminated in iOS 11

Outline

1. Side-channel Attack Vectors on iOS
2. Attack 1: Classifying User Activities
3. Attack 2: Detecting Sensitive In-App Activities
4. Attack 3: Bypassing Sandbox Restrictions
5. Practical Issues
6. Countermeasures
7. Conclusion

Conclusion

- First exploration of OS-level side channels on iOS
- Three categories of side-channel attacks
- Proposed countermeasures integrated in iOS and MacOS

Xiaokuan Zhang
zhang.5840@osu.edu

Detecting Sensitive In-App Activities --- Attack Methods

- Time is short (<0.5s)
- Difference is subtle

Detecting Sensitive In-App Activities --- Attack Methods

- Pattern Matching: compare two multi-dimensional data traces
  - Sample: $\vec{X}_t = \{\vec{X}_t^1, \vec{X}_t^2, \cdots, \vec{X}_t^l\}$, where $\vec{X}_t^i = (X_{t_1}^i, X_{t_2}^i, \cdots, X_{t_{n_i}}^i)$
  - Signature: $\vec{S}_t = \{\vec{S}_t^1, \vec{S}_t^2, \cdots, \vec{S}_t^l\}$
  - Goal: measure the distance $d(\vec{X}_t, \vec{S}_t)$
  - Extended DTW (DTW_I): (w_k: normalization factor)

$$d(\vec{X}_t, \vec{S}_t) = \sum_{k=1}^{l} \frac{1}{w_k} \cdot \mathrm{DTW}\left(\vec{X}_t^k, \vec{S}_t^k\right)$$

iOS Attacks

JUN 17, 2015 @ 10:51 AM 24,925

Apple App Securi Vulnerable To 'D

Thomas Fox-Brewst

I cover crime, privacy and s forms. **FULL BIO**

It's become almost axiomatic th and the apps on them are more competition. But researchers co notion and today a group of aca the security protections in Mac not only possible to create malw Store, but it's also feasible to lau using rogue software to steal the data around, from iCloud passw to dodgy selfies and more.

The attacks, known as unauthor access or XARA, expose design to access critical pieces of data i Apple has struggled to fix the is released today from Indiana University Bloomington, Peking University and the Georgia Institute of Technology.

JUL 28, 2016 @ 09:40 AM

How Hackers Photos

Thomas Fox-Bre

I cover crime, privacy forms. **FULL BIO**

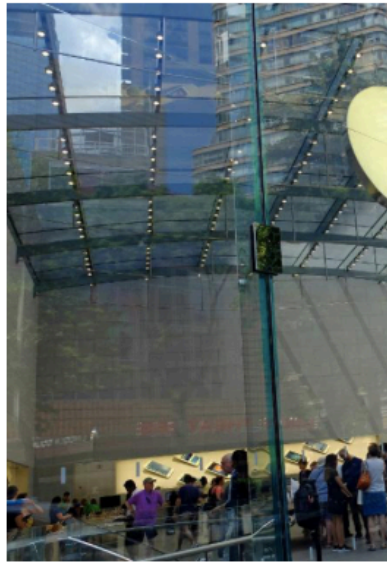

NEW YORK, NY - JULY 27: The Apple

JAN 17, 2018 @ 07:36 PM 19,800

The Little Black Book of Billionaire Secrets

Dangerous iPhone Bug Hiding in iMessage Is Causing Chaos

Ewan Spence, CONTRIBUTOR

FULL BIO

Opinions expressed by Forbes Contributors are their own.

Apple is facing another blow to its reputation for security on the iPhone. A flaw in iMessage has been discovered that allows a single message to lock up and potentially crash your handset. And you don't even have to read the message for it to activate.

The bug itself is relatively easy to explain. When iMessage receives a message with a URL embedded, it will go online and generate a small thumbnail preview of the link. If the metadata is much larger than normally accepted (on the order of hundreds of thousands of characters), then iMessage will lock up the device. The hacker who announced this bug demonstrated it to BuzzFeed News through a poisoned page hosted on Github:

Paper	Vector	Impact
Chen et al., Security'14	/proc/pid/ statm	UI inference attacks (stealing login credentials, photos)
Diao et al., Oakland'16	/proc/ interrupts	Interrupt timing analysis (cracking unlock patterns)

Classifying User Activities --- Attack Methods

- Requirements:
  - Combining
  - Reducing t
- Major comp
  - Symbolic
  - Bag-of-Pa
  - Support Vector Machine (LibSVM) (Chang et al., 2011)

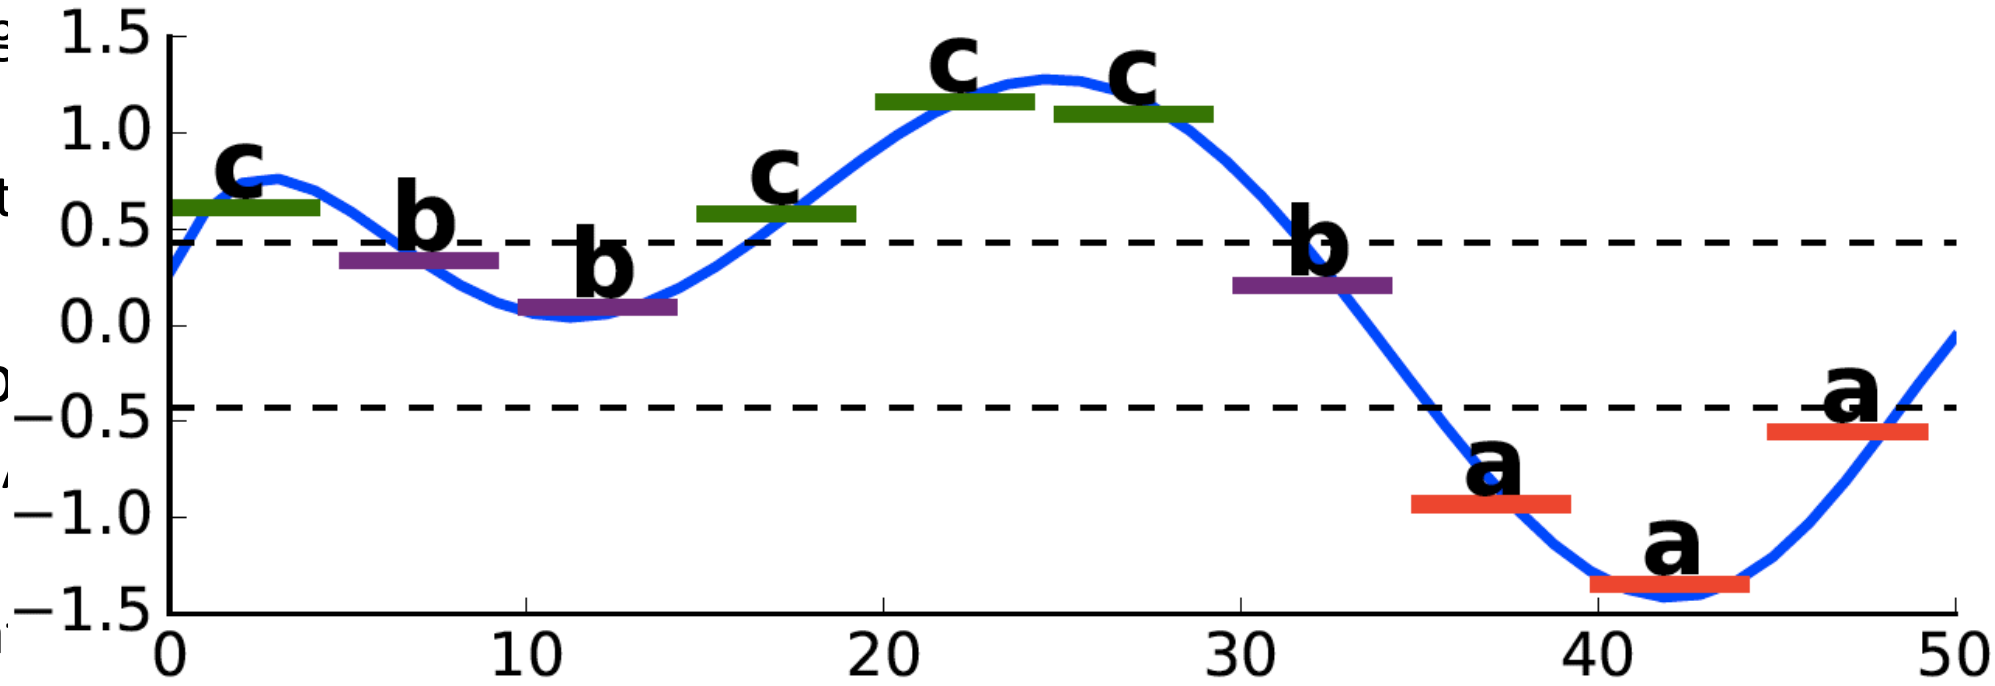

{cbb:1, bbc:1, bcc:1, ccc:1,
ccb:1, cba:1, baa:1, aaa:1}

Classifying User Activities --- Case Studies

- Top N Accuracy Example

Sample	True Class	SVM Prediction (Probability Model)		
A	1	4	2	1
B	2	2	5	4
C	3	3	1	2
D	4	1	4	2
E	5	5	2	4

Classifying User Activities --- Case Studies

- Top N Accuracy Example

Sample	True Class	SVM Prediction (Probability Model)		
A	1	4	2	1
B	2	2	5	4
C	3	3	1	2
D	4	1	4	2
E	5	5	2	4

Classifying User Activities --- Case Studies

- Top N Accuracy Example

Sample	True Class	SVM Prediction (Probability Model)		
A	1	4	2	1
B	2	2	5	4
C	3	3	1	2
D	4	1	4	2
E	5	5	2	4

Top 1 Accuracy: 3/5 = 60%

Classifying User Activities --- Case Studies

- Top N Accuracy Example

Sample	True Class	SVM Prediction (Probability Model)		
A	1	4	2	1
B	2	2	5	4
C	3	3	1	2
D	4	1	4	2
E	5	5	2	4

Classifying User Activities --- Case Studies

- Top N Accuracy Example

Sample	True Class	SVM Prediction (Probability Model)		
A	1	4	2	1
B	2	2	5	4
C	3	3	1	2
D	4	1	4	2
E	5	5	2	4

Top 2 Accuracy: (3+1)/5 = 80%

Classifying User Activities --- Case Studies

- Top N Accuracy Example

Sample	True Class	SVM Prediction (Probability Model)		
A	1	4	2	1
B	2	2	5	4
C	3	3	1	2
D	4	1	4	2
E	5	5	2	4

Classifying User Activities --- Case Studies

- Top N Accuracy Example

Sample	True Class	SVM Prediction (Probability Model)		
A	1	4	2	1
B	2	2	5	4
C	3	3	1	2
D	4	1	4	2
E	5	5	2	4

Top 3 Accuracy: (2+1+2)/5 = 100%

Detecting Sensitive In-App Activities

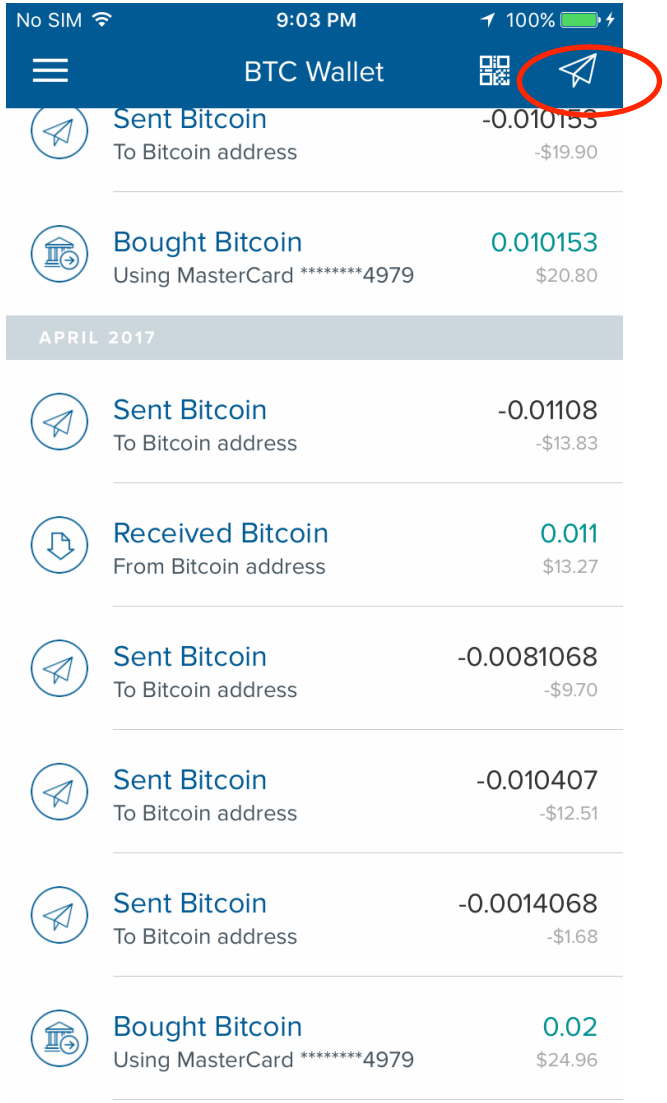

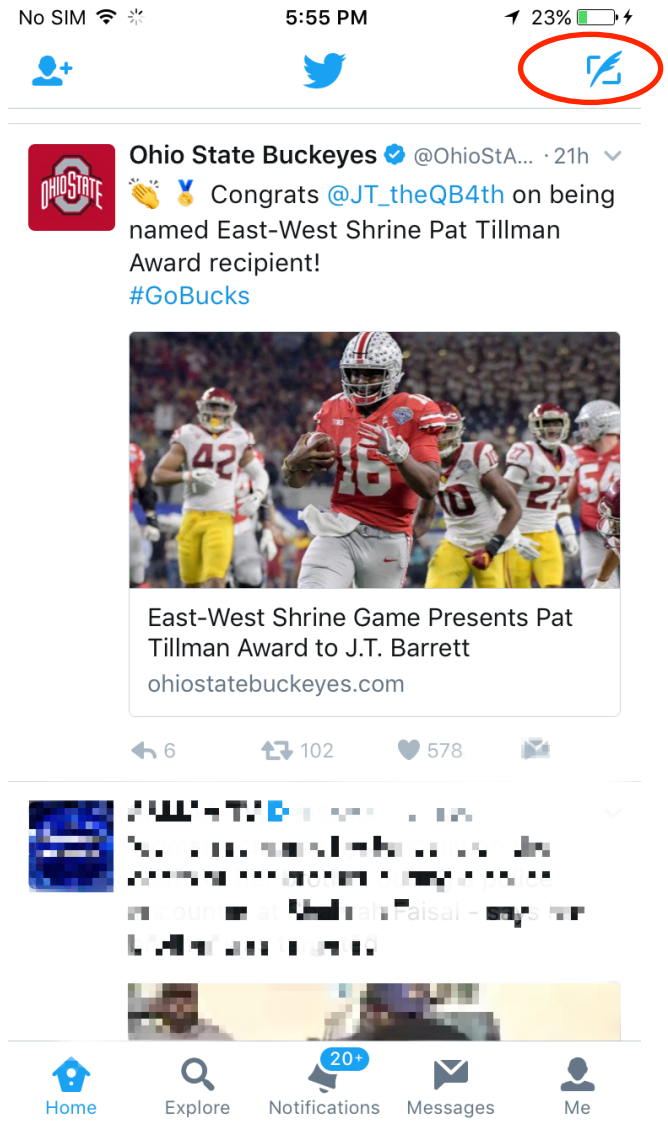

Detecting Sensitive In-App Activities --- Attack Methods

- Identify critical events

- Correlates with public records

Detecting Sensitive In-App Activities

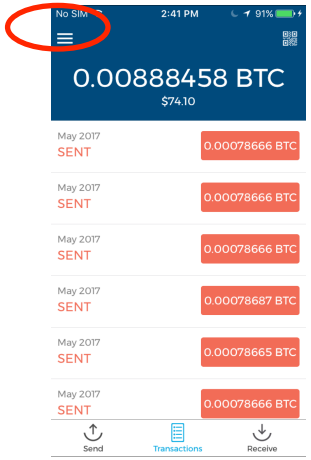

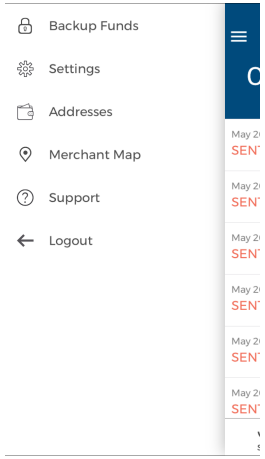

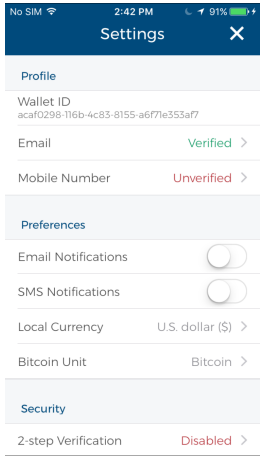

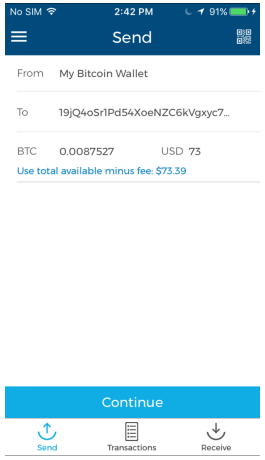

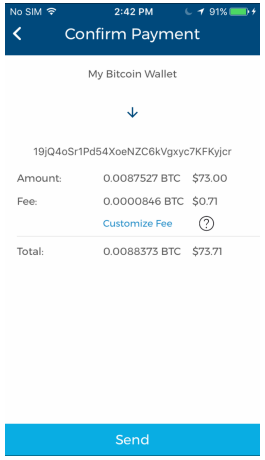

Classifying User Activities --- Case Studies

- Device: jailbroken iPhone 7 with iOS 10.1.1
- Automated using Cycript

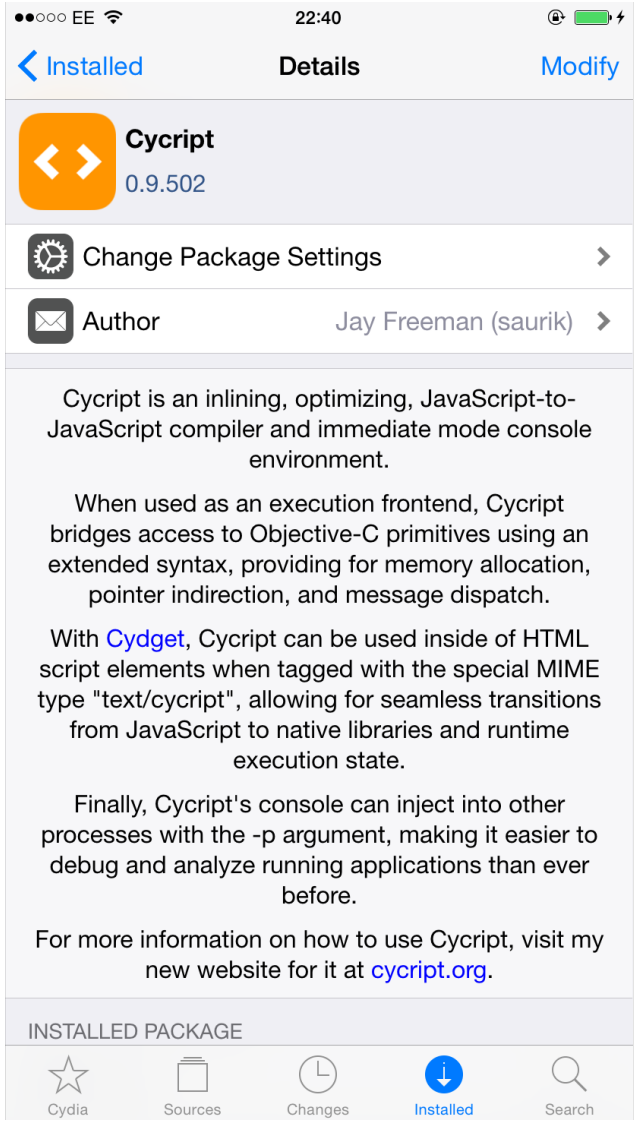

Why global stat can work?

- iOS itself suspends apps when they run in the background, unless the app specially requests background permissions

- iOS is relatively quieter than Android, which greatly facilitates side-channel attacks

Run Background Apps on iOS

- *AUDIO* background mode

- [NSTimer scheduledTimerWithTimeInterval: target: selector: userInfo: repeats:]

Detecting Sensitive In-App Activities