Against the Law: Countering Lawful Abuses of Digital Surveillance

Andrew 'bunnie' Huang Edward Snowden

abstract: Front-line journalists are high-value targets, and their enemies will spare no expense to silence them. Unfortunately, journalists can be betrayed by their own tools. Their smartphones are also the perfect tracking device. Because of the precedent set by the US's "third-party doctrine," which holds that metadata on such signals enjoys no meaningful legal protection, governments and powerful political institutions are gaining access to comprehensive records of phone emissions unwittingly broadcast by device owners. This leaves journalists, activists, and rights workers in a position of vulnerability. This work aims to give journalists the tools to know when their smart phones are tracking or disclosing their location when the devices are supposed to be in airplane mode. We propose to accomplish this via direct introspection of signals controlling the phone's radio hardware. The introspection engine will be an open source, user-inspectable and field-verifiable module attached to an existing smart phone that makes no assumptions about the trustability of the phone's operating system.

Introduction and Problem Statement

Front-line journalists risk their lives to report from conflict regions. Casting a spotlight on atrocities, their updates can alter the tides of war and outcomes of elections. As a result, front-line journalists are high-value targets, and their enemies will spare no expense to silence them. In the past decade, hundreds of journalists have been captured, tortured and killed. These journalists have been reporting in conflict zones, such as Iraq and Syria, or in regions of political instability, such as the Philippines, Mexico, and Somalia.

Unfortunately, journalists can be betrayed by their own tools. Their smartphones, an essential tool for communicating with sources and the outside world–as well as for taking photos and authoring articles–are also the perfect tracking device. Legal barriers barring the access to unwitting phone transmissions are failing because of the precedent set by the US's "third-party doctrine," which holds that metadata on such signals enjoys no legal protection. As a result, governments and powerful political institutions are gaining access to comprehensive records of phone emissions unwittingly broadcast by device owners. This leaves journalists, activists, and rights workers in a position of vulnerability. Reporter Marie Colvin's 2012 death is a tragic reminder of how real this vulnerability can be. A lawsuit against the Syrian government filed in 2016 alleges she was

deliberately targeted and killed by Syrian government artillery fire. The lawsuit describes how her

location was discovered in part through the use of intercept devices that monitored satellite-dish and cellphone communications.[1]

Turning off radios by entering airplane mode is no defense; for example, on iPhones since iOS 8.2, GPS is active in airplane mode. Furthermore, airplane mode is a “soft switch”–the graphics on the screen have no essential correlation with the hardware state. Malware packages, peddled by hackers at a price accessible by private individuals, can activate radios without any indication from the user interface; trusting a phone that has been hacked to go into airplane mode is like trusting a drunk person to judge if they are sober enough to drive.

This work aims to give journalists the tools to know when their smart phones are tracking or disclosing their location when the devices are supposed to be in airplane mode.

Approach and Goals

Numerous researchers and extensive corporate resources have been dedicated to the task of building a more secure smart phone. However, smartphones are extremely complex and present a large, porous attack surface. Furthermore, even a perfectly secure phone will not save a reporter from “victim-operated” exploits such as spearphishing. Eliminating this vector is complicated by the fact that effective reporters must communicate with a diverse array of sources who may intentionally or unintentionally convey a malware payload to the reporter.

As a result, this work starts with the assumption that a phone can and will be compromised. In such a situation, a reporter cannot take the UI status at face value. Instead, we aim to provide field-ready tools that enable a reporter to observe and investigate the status of the phone’s radios directly and independently of the phone’s native hardware. We call this direct introspection.

Our work proposes to monitor radio activity using a measurement tool contained in a phone-mounted battery case. We call this tool an introspection engine. The introspection engine has the capability to alert a reporter of a dangerous situation in real-time. The core principle is simple: if the reporter expects radios to be off, alert the user when they are turned on.

Our introspection engine is designed with the following goals in mind:

1. Completely open source and user-inspectable (“You don’t have to trust us”)
2. Introspection operations are performed by an execution domain completely separated from the phone’s CPU (“don’t rely on those with impaired judgment to fairly judge their state”)
3. Proper operation of introspection system can be field-verified (guard against “evil maid” attacks and hardware failures)
4. Difficult to trigger a false positive (users ignore or disable security alerts when there are too many positives)
5. Difficult to induce a false negative, even with signed firmware updates (“don’t trust the system vendor” – state-level adversaries with full cooperation of system vendors should not be able to craft signed firmware updates that spoof or bypass the introspection engine)

7. Simple, intuitive user interface requiring no specialized knowledge to interpret or operate (avoid user error leading to false negatives; “journalists shouldn’t have to be cryptographers to be safe”)
8. Final solution should be usable on a daily basis, with minimal impact on workflow (avoid forcing field reporters into the choice between their personal security and being an effective journalist)

This work is not just an academic exercise; ultimately we must provide a field-ready introspection solution to protect reporters at work. Although the general principles underlying this work can be applied to any phone, reducing these principles to practice requires a significant amount of reverse engineering, as there are no broadly supported open source phone solutions on the market. Thus we focus on a single phone model, the 4.7” iPhone 6 by Apple Inc., as the subject for field deployment. The choice of model is driven primarily by what we understand to be the current preferences and tastes of reporters. It has little to do with the relative security of any platform, as we assume any platform, be it iOS or Android, can and will be compromised by state-level adversaries.

Methods & Intermediate Results

The first step toward executing this work was to visit the Hua Qiang electronics markets of Shenzhen to collect samples and documentation for evaluation. These markets are ground zero for the trade and practice of iPhone repair; as such, it is a rich source of spare parts and repair manuals. The repair manuals frequently contain detailed blueprints of the iPhone 6, which were used to assist the reverse engineering effort.

Based on the phone model selection and available documentation, we can enumerate the radio interfaces available:

- Cellular modem – 2G/3G/4G
- Wifi / BT
- GPS
- NFC (Apple Pay)

Although our work can be extended to input systems such as the IMU (inertial measurement unit), barometer, microphone and camera, to focus the effort we restrict our exploration to only RF interfaces that can directly betray a user’s location. Note that a camera can be defeated by obscuring the lens; as such the final physical design of our battery case will likely include a feature to selectively obscure the rear camera lens.

Methods that Do Not Meet our Criteria

Numerous semi-intrusive countermeasures were considered along the way to our current solution, including but not limited to RF spectrum monitoring, active jamming, and the selective physical isolation or termination of antennae. Semi-intrusive countermeasures would require minimal

RF spectrum monitoring consists of building an external radio receiver that can detect transmissions emanating from the phone's radios. In some cases, it was hypothesized that the receiver could be as trivial as an RF power monitor within the anticipated radio bands. A simple example of such monitoring already exists in the form of novelty lights that flash based on parasitic power extracted from the GSM antennae. The problems with this approach is that 1) it can only reliably detect active transmissions from the radio, and 2) malware that passively records the user's position and delivers it as a deferred payload when the radios are intentionally activated cannot be detected. Furthermore, this approach is subject to spoofing; false positives can be triggered by the presence of nearby base stations. Such false alarms can confuse the user and eventually lead the user to be conditioned to ignore real alerts in hazardous situations.

Active jamming consists of building an external radio transmitter that attempts to inject false signals into the radios. Thus, even if malware were to activate the radios and listen for position-revealing signals, it would, in theory, report largely bogus position information. This is particularly effective against GPS, where GPS signals are very weak and thus even a weak local transmitter should be able to overpower the GPS satellites. However, active jamming was ruled out for several reasons. The jammer's emissions could create a signal that can be traced to locate the reporter; the jammer will require substantial battery power, and the user is left vulnerable once the jammer's power is exhausted. Furthermore, nearby base stations may still be detected by the receivers, as modern radio protocols have sophisticated designs to protect against unintentional jamming.

Selective physical isolation or termination of the antennae consists of inserting an electronic switch between the connectors of the logic board and the antenna. The switch, when activated, would shunt the antenna to a matched resistive load, which would greatly reduce the transmission power and receive sensitivity of the radios. However, experimental verification on the WiFi subystem indicated that removing the antenna connection and permanently terminating with a shunt resistor still leaked sufficient RF into the receivers for local base stations (e.g., within the same room) to be detected, which could be sufficient information to betray a reporter's location.

Methods that Do Meet our Criteria

Upon determining that semi-intrusive countermeasures were inadequate, we investigated options that involve measuring signals on the phone's logic board, typically via test points designed in by the manufacturer. It is no surprise that complex systems such as the Apple iPhone 6 would have test points baked into the circuit board design to assist with debugging. These are an essential part of yield and customer experience improvement; defective units from the factory and the field are sent back to the headquarters, and engineers rely on these testpoints to determine the root cause of the device's failure.

Using repair manual documentation acquired from the Hua Qiang electronics market, we cataloged a set of internal test points that were:

3. Would be difficult or impossible to disable or spoof (e.g., future-proof against adversaries aware of our research).

For the accessibility criteria (1), test points were considered viable even if they required desoldering an RF shield or the SIM card connector, and manual removal of soldermask. In our experience, a trained operator can perform these tasks with low probability of irreparable damage to the motherboard. These operations are not recommended for entry-level novices. However, our experiences in Shenzhen indicate that any technician with modest soldering skills can be trained to perform these operations reliably in about 1-2 days of practice on scrap motherboards. Thus, technicians could be trained to perform the modifications in any locale with sufficient demand for modified iPhones.

The following table is a list of test points we have accessed and have found to provide introspection data that potentially meet criteria (2) and (3).

Signal name	Signal type	Signal function	Related radios
FE2DATA	Shared serial data bus	Configure antenna switches	Cellular radio antenna multiplexer, RX diversity, antenna tuning
FE2CLK	Reference clock		
FE1DATA	Shared serial data bus	Configure Power Amplifiers and Duplexers	Cellular radios
FE1CLK	Reference clock		
BBTX	UART	Baseband to AP comms	GPS, others
BBRX	UART		
WLAN RX	UART	WLAN to AP comms	WLAN
WLAN TX	UART		
WLAN_PERST	Reset	Reset PCI bus on WLAN	WLAN
BT RX	UART	Bluetooth to AP comms	Bluetooth
BT TX	UART		
GPS_SYNC	Sync status	GPS signal quality and sync	GPS

Above: table of internal signal candidates for introspection.

Above: image of the FE1, FE2 bus probe experiment. Test points from the back side of the PCB are wired to the top side for easy probing.

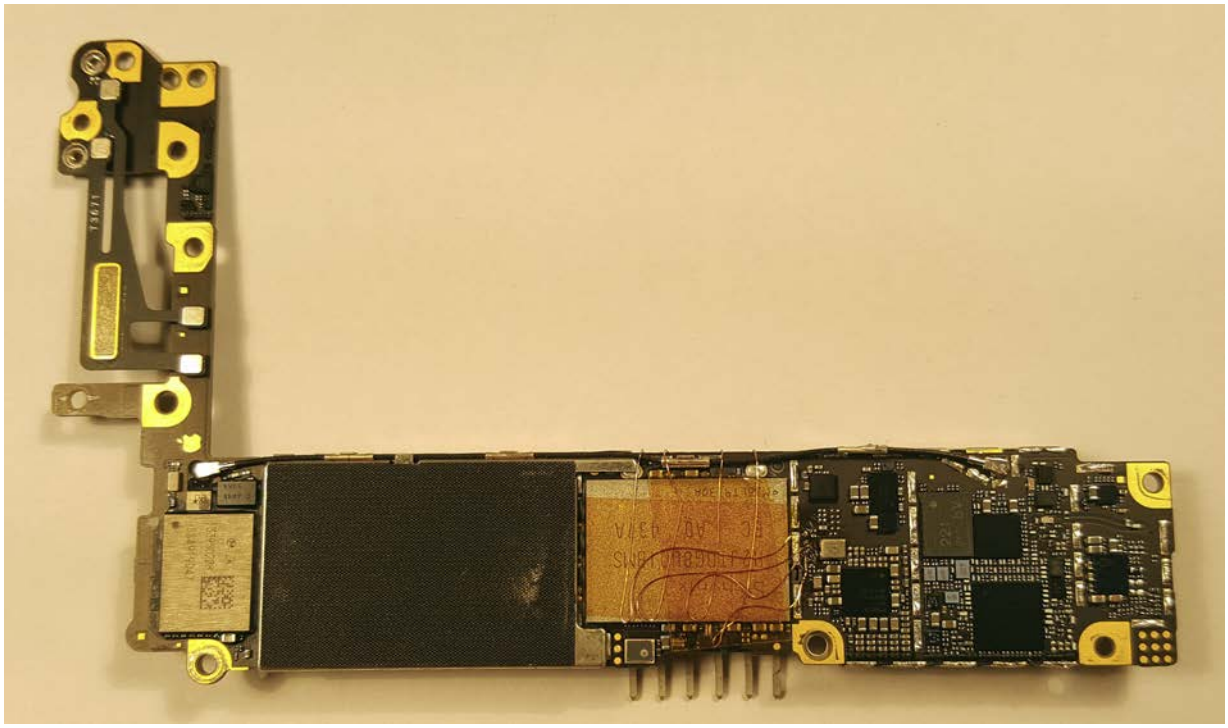

Above: image of the backside of the FE1, FE2 probe experiment. The test points are located adjacent to the NAND Flash, underneath an RF shield which was removed for this experiment. The test points were covered with soldermask, which was removed through mechanical abrasion.

Above: image of the UART and GPS sync probing experiment. The majority of the test points are located underneath the SIM card connector, which was removed for this experiment.

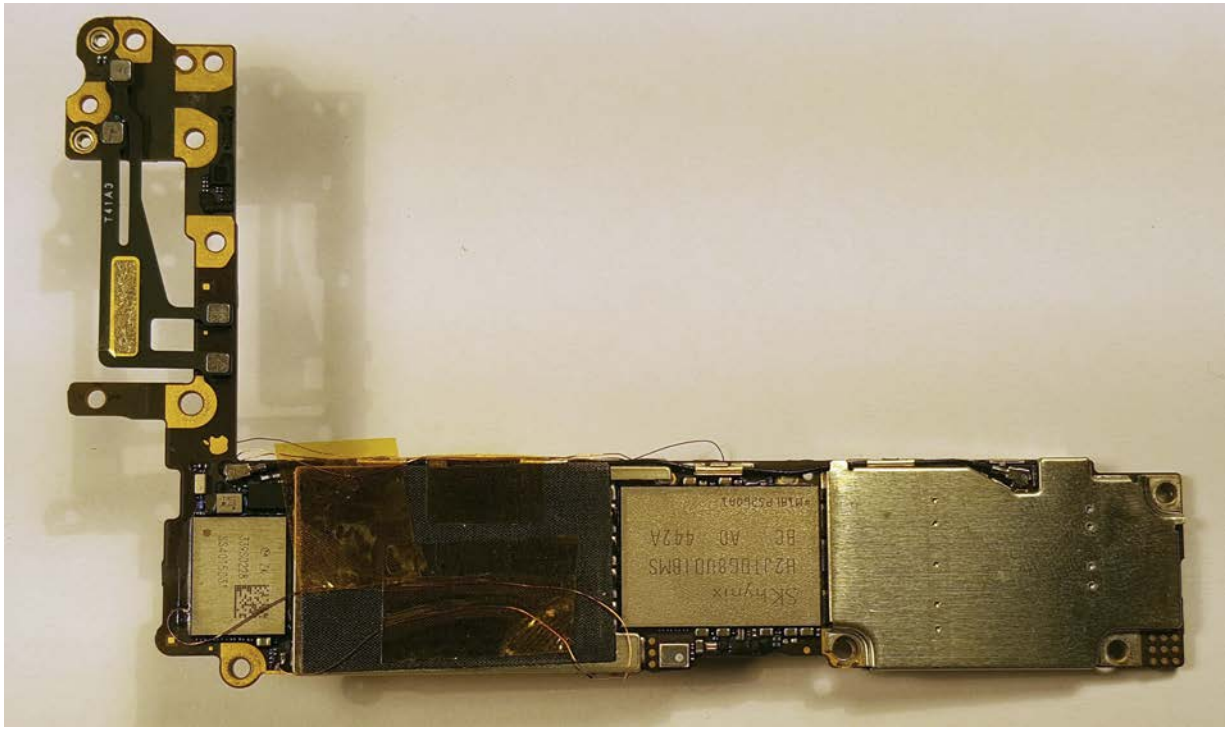

Above: image of the back side of the UART and GPS sync probing experiment. A pair of wires are run to break out WLAN_PERST and power-related signals for monitoring.

Cellular Modem Introspection

The FE1 and FE2 serial buses run at 20MHz, with a 1.8V swing. This bus is used primarily to configure the cellular modem radios. When the radios are on, there is constant traffic on these buses. When in airplane mode, the traffic completely ceases.

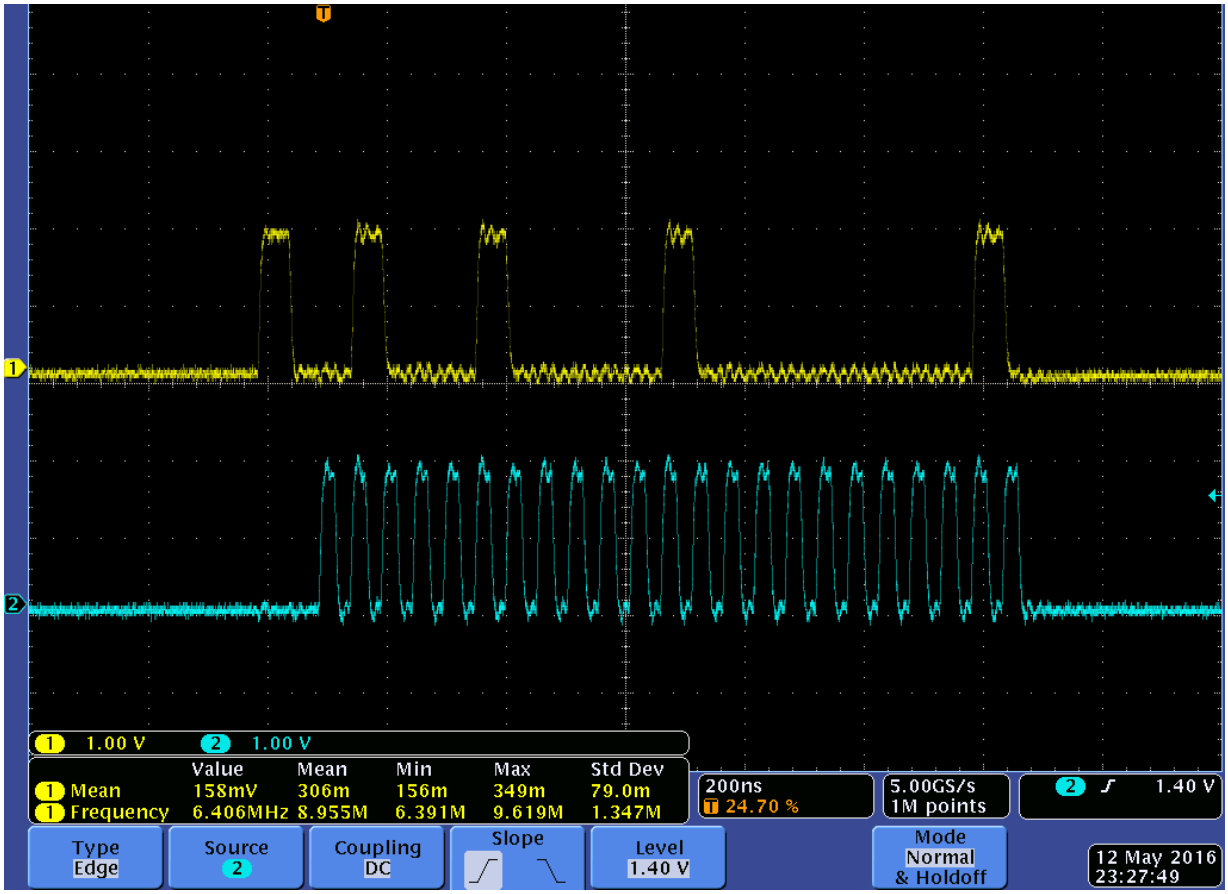

Above: example of bus traffic on the FE1 bus.

Cellular radios operate in a complex environment, and require constant adaptation of the antennae, power amplifiers, and band selection for proper operation. It is hypothesized that an attempt to even passively scan for base stations without transmitting will require traffic on this bus; at the very least, the antenna switches must be powered on and configured to receive. Therefore, cellular modem introspection may be as easy as noting if there is any activity on the FE buses during airplane mode.

We note for the sake of completeness that it may be possible for an attacker to statically configure the antenna, channel, and power amplifier settings and convert the device into a radio beacon that blasts out a signal that is inconsistent with the cellular modem standard but detectable through other means. In this mode, one would observe no traffic on the FE buses, but one could, in theory, triangulate the location of the transmitter with modified base stations or specially deployed receivers. This scenario can be mitigated by doing deep packet inspection and noting the addresses that should be hit to power down the cellular modem systems. If any devices are skipped during the power-off sequence, that would be flagged as a potentially hazardous condition.

However, this scenario would require modifications to the cellular modem transport specifications, and as such one would need to deploy modified base stations across the territory to gain adequate surveillance coverage. This would likely require extensive cooperation of both the baseband radio vendors and cellular providers to craft and effectively deploy such an exploit. Because of the difficulty, we imagine such an exploit would be available only to well-organized government-level

sends random “NOP” packets over the FE buses during airplane mode to force false positives and make this technique less effective. Again, in such a case deep packet inspection could help to discard chaff from signal. Although future hardware versions could encrypt this bus to foil observation, we believe it is not possible to introduce bus encryption with a software-only change: the peripheral devices on this bus lack loadable firmware. Thus, at least for current phone models, deep packet inspection should be robust.

WiFi & Bluetooth Introspection

The WiFi subsystem interfaces to the CPU through multiple buses, namely, PCI-express and a UART; the Bluetooth subsystem interfaces to the CPU through a UART, with a separate UART channel for coexistence. Because of the Bluetooth subsystem’s relatively simple interface, it should be possible to robustly detect Bluetooth activity by simply monitoring the BT UART signals.

The WLAN UART signals seem to carry configuration and status information regarding WiFi configuration, as evidenced by the UART trace below.

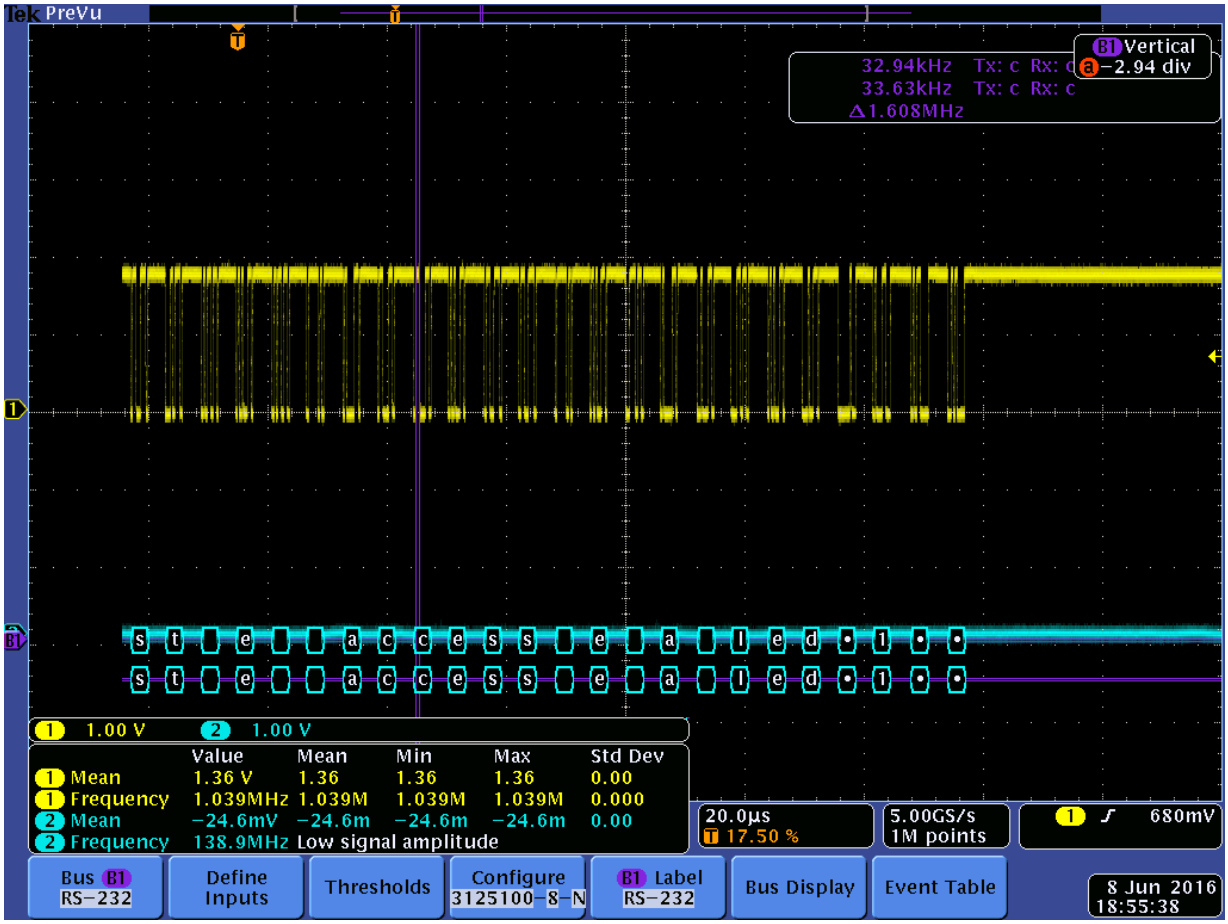

Above: example data on the Wifi UART as decoded by a Tek MDO4014B.

Further exploration of the data contained within the signals is necessary to determine if it is possible for an adversary to perform access point scans, which is an effective means of geolocation, without invoking the UART. Unfortunately, the WiFi power remains on even in airplane mode, so

holding WLAN_PERST low prior to power-on and throughout boot, WiFi will fail to enumerate on the PCI bus. iOS will continue to boot and is fully usable, but in the Settings panel, WiFi will appear to be off and cannot be switched on. Attempts to switch on Bluetooth fail, and GPS, although active, cannot access its antenna as the antenna for GPS is shared with WiFi. Note that forcing WLAN_PERST low during normal operation forces a phone reboot, so disabling WiFi using this technique effectively necessitates a reboot.

This is a simple but effective method to force several critical subsystems to be off, with no chance for an updated firmware to bypass a WiFi hardware reset. However, the failure of Bluetooth and GPS subsystems to activate may be due to firmware-only dependencies. It is hypothesized that these systems rely on WiFi to initialize before activating the respective antenna switches for these subsystems, since they all share a common antenna port. Thus it may be possible for an exploit to be developed to force Bluetooth and GPS to be on even if WiFi is in reset. Furthermore, it may be possible for malware to fingerprint systems where the WiFi has failed to initialize, and flag these users for further monitoring.

Thus, depending on the user's threat model, the WLAN_PERST defeat may be a simple but effective method to defeat several radios with a single signal, but it may also give away information to advanced adversaries on the presence of an introspection engine. Because of the effectiveness of the WLAN_PERST trick, we would present users with the option to activate this, but not require it.

Significantly, repair manuals indicate that the WiFi/Bluetooth module includes a hardware "RFKILL" pin. Apple leaves this pin unconnected and very difficult to access through mods, but if phone vendors wanted to support efforts like this, future revisions of phones could break such pins out to offer a more graceful defeat that doesn't require rebooting the phone or leave a measurable signature while disabling these radios.

GPS Introspection

To date, we have identified three possible methods for detecting GPS activation. One is to look for activity on the BB UART bus. When GPS is active, coordinate data seems to be transmitted over the BB UART bus. A second is to look at the GPS_SYNC signal. When GPS is active, the GPS_SYNC signal pings the base band at a rate of about once per second, with a pulse width inversely proportional to the quality of the GPS lock. A very wide pulse indicates a high degree of uncertainty in the GPS signal. Finally, the GPS has an independent power regulator which is turned off when the GPS is not active, to save power.

NFC Introspection/Defeat

For NFC, we decided that the risk/reward of selectively enabling and monitoring Apple Pay is not worth it. In other words, we do not expect journalists operating in conflict zones to be relying on Apple Pay to get their work done. Therefore, to simplify the effort, we opt to fully disable Apple Pay

this screw and separating the antenna from the main logic board, we hope to substantially and selectively reduce the sensitivity of the NFC radio. Further testing is required to determine if this is sufficient to guard against attacks by adversaries using high-power amplifiers to query the Apple Pay NFC feature. If found inadequate, further countermeasures, including but not limited to permanently removing the Apple Pay NFC RF front end chip from the mainboard, are options to prevent exploitation of the radio without leaving a clear signature that can be detected by an adversary.

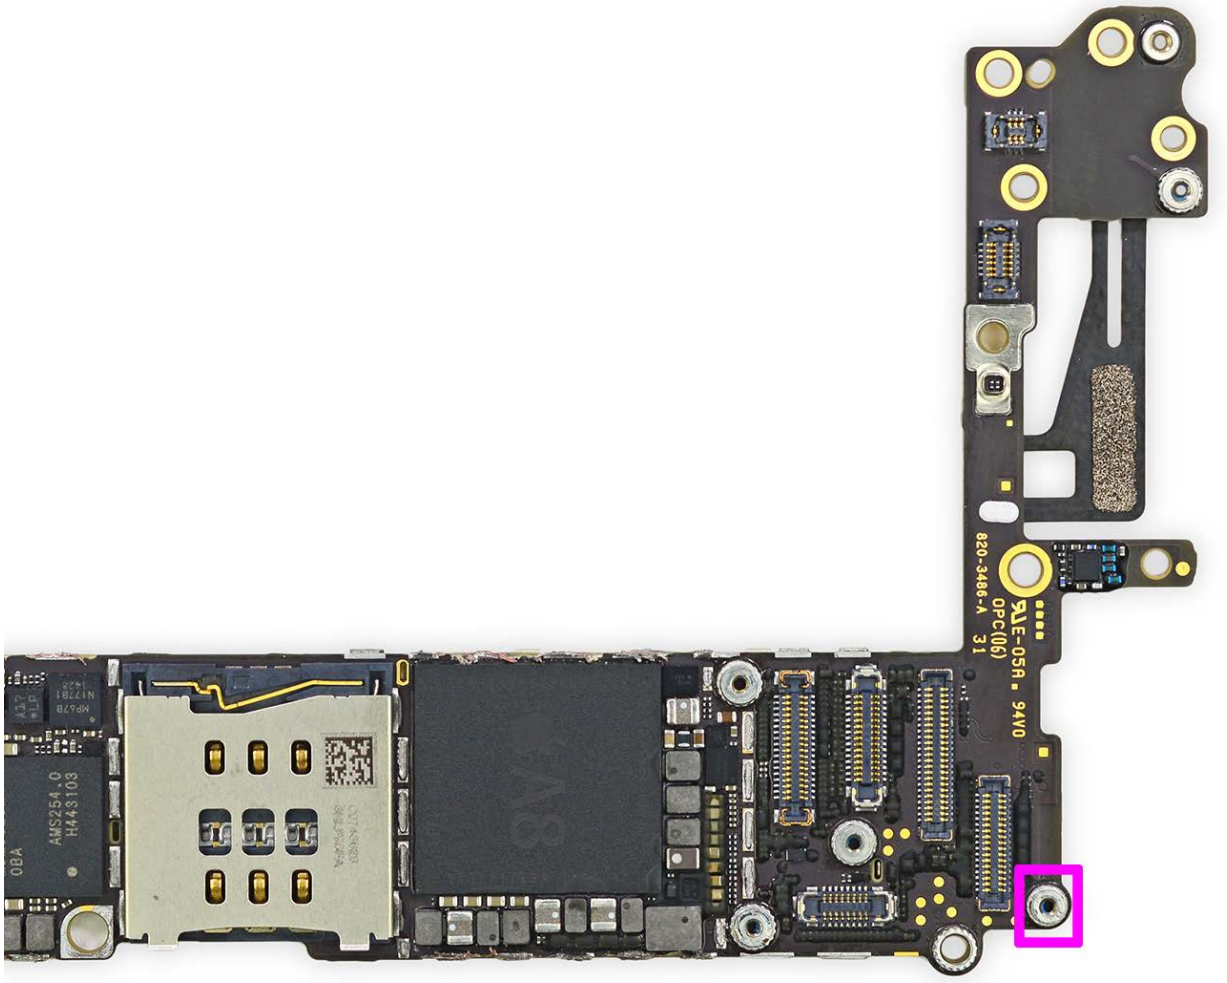

Above: location of the Apple Pay antenna connection, highlighted in pink. Original image courtesy iFixit, CC-BY-NC-SA licensed.

Next Steps and Field Deployment

Now that a set of viable signals has been identified for introspection, the next step is refining the system for field deployment.

From the outside, the introspection engine will look and behave like a typical battery case for the iPhone 6. However, in addition to providing extra power to the iPhone 6, the case will contain the introspection engine's electronics core. The electronics core will likely consist of a small FPGA and

introspection engine.

Above: Conceptual rendering of a "battery case" style introspection engine, piggybacked on an iPhone6.

The battery case/introspection engine will also feature an independent screen to update the user on radio status; for example, it can inform the user on time elapsed since the last traffic was detected on any radio bus. Thus, users can field-verify that the bus taps are in place by briefly bringing the system out of airplane mode in a safe location. Any radio that does not report traffic out of airplane mode would indicate a hardware failure of the introspection engine. Of course, the system will also feature an audible alarm that can be set to trip in case any activity is seen on any set of radios. It might also be desirable to incorporate a "kill switch" feature which forcibly disconnects power to the phone in the case that a radio is found to be errantly transmitting.

be designed with contacts pre-loaded at signal test point locations. This will streamline phone modifications while making the final product more robust. As the SIM card has to be removed for access to key test points, the FPC will also connect to the SIM card signals. An additional FPC will then exit via the existing SIM card port, making available to the introspection engine both the bus taps and the SIM card signals.

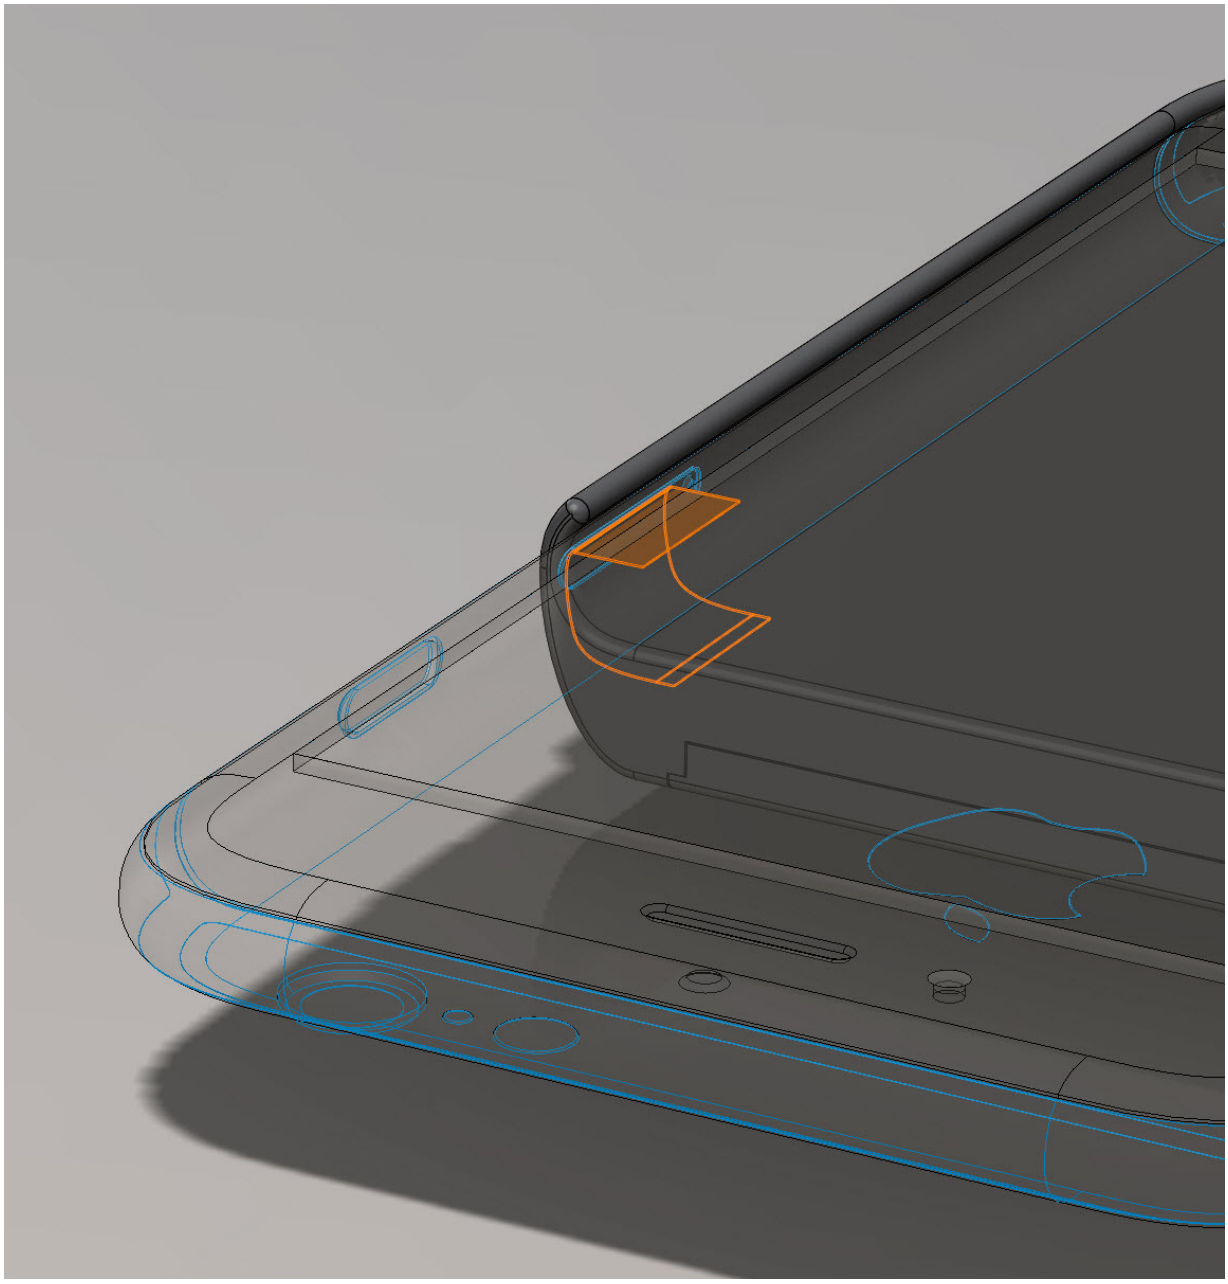

Above: The orange highlighted part is a proposed FPC which exits via the SIM card port and routes signals from the modified iPhone6 mainboard to the introspection engine's electronics.

This architecture opens the possibility of the introspection engine featuring multiple SIM card slots. Although the system will still need to be rebooted when switching SIMs, it can be convenient for certain users to be able to switch SIMs rapidly without the use of any extra tools or worry of dropping and losing the tiny SIM cards. This is especially problematic, for example, when switching SIM cards during transit on unpaved, bumpy roads. It should be noted that changing SIM cards is

Over the coming year, we hope to prototype and verify the introspection engine's abilities. As the project is run largely through volunteer efforts on a shoestring budget, it will proceed at a pace reflecting the practical limitations of donated time. If the prototype proves successful, the FPF may move to seek the necessary funding to develop and maintain a supply chain. This would enable the FPF to deploy modified iPhone 6 devices for field service among journalists in high-risk situations.

The techniques developed in this work should also be applicable to other makes and models of phones. Pervasive deployment of radio introspection techniques could be assisted with minimal cooperation of system vendors. By grouping radio control test points together, leaving them exposed, and publishing a terse description of each test point, direct introspection engines can be more rapidly deployed and retrofitted into future smartphones.

Furthermore, direct introspection may be extendable beyond the radio interfaces and into the filesystem layer. We theorize an introspection engine attached to the mass storage device within a phone; for example, an FPGA observing the SD bus between the CPU and the eMMC in a typical Android phone implementation. This introspection engine could observe, in real time, file manipulations and flag, or even block, potentially suspicious operations. With further system integration, the introspection engine could even perform an off-line integrity check of the filesystem or disk image. The efficacy of filesystem introspection is enhanced if the system integrator chooses to only sign OS-related files, but not encrypt them. As core OS files contain no user data or secrets, baring them for direct introspection would not impact the secrecy of user data while enabling third-party attestation of the OS's integrity.

References

[1] Dana Priest. *Washington Post.* [http://wpo.st/5W2l1] (http://wpo.st/5W2l1)