How to Phone Home with Someone Else's Phone: Information Exfiltration Using Intentional Sound Noise on Gyroscopic Sensors

Benyamin Farshteindiker, Nir Hasidim, Asaf Grosz and Yossi Oren
Faculty of Engineering Sciences, Ben Gurion University of the Negev, Israel

Abstract

We show how a low-power device, such as a surveillance bug, can take advantage of a nearby mobile phone to exfiltrate arbitrary secrets across the Internet at a data rate of hundreds to thousands of bits per second, all without the phone owner's awareness or permission. All the attack requires is for the phone to browse to an attacker-controlled website. This feat is carried out by exploiting a particular characteristic of the phone's gyroscope which was discovered by Son et al. in [11]. We discuss the theoretical principles behind our attack, evaluate it on several different mobile devices, and discuss potential countermeasures and mitigations. Finally, we suggest how this attack vector can be used benevolently for the purpose of safer and easier two-factor authentication.

1 Introduction

An increasing number of people are finding themselves branded as **intelligence targets**. Intelligence targets are entities which are of interest to state-sponsored signals intelligence agencies, or similarly powerful malicious adversaries. These adversaries presume their victims are in possession of some secret information, and attempt to acquire this information using various stealthy techniques. Most intelligence targets are tracked by huge scale, bulk collection efforts which target the Internet's routing backbone and data centers; for higher-value targets, the malicious adversaries tend to make use of custom hardware implants, or "bugs".

Our work suggests a way of designing a particularly stealthy and effective implant. Before we describe our design in detail, we first describe the architecture of implants in general. As described in Table 1, an implant has three main functional components: First, it must **collect** secret information from its victim; next, it must **exfiltrate** this secret payload by connecting to a central command and control (C&C) server, an activity colloquially referred to as "phoning home"; finally, the implant requires some sort of **power supply** to power its computation and communication functions.

Component	Example Instantiations
Secret Collection	Microphone, camera, side-channel probe
Secret Exfiltration	RF backscatter, acoustical coupling, this work
Power Supply	Passive power, battery

Table 1: Components of a general implant device

Secret collection can be achieved by various methods. Most trivially, an implant can use onboard sensors such as microphones or cameras to spy on the victim. If the implant is placed near the victim's computer or mobile phone, the implant can mount a side-channel attack on the phone, using techniques similar to Genkin et al. [4], to recover secret information such as encryption keys and bitcoin wallets. In some cases (such as supply-chain interdiction, as described below), implants have direct access to interesting data lines, such as the bus between a computer and its video display.

Secret exfiltration is a bigger engineering challenge. Implants have to operate in an **adversarial setting**, meaning that they should be as difficult as possible to detect by the victim. Combined with their very limited power budget this implies they cannot contain a cellular modem, satellite radio or other forms of long-range radio transmitter, since these functions are all power-hungry and easily detectable. Instead, implants tend to use various low-power, short-range transmission schemes based on RFID backscatter, short-range radio networks, acoustical coupling, etc. To collect this exfiltrated data, the intelligence agency is thus required to deploy a field agent, equipped with a sophisticated **collection device**, to the vicinity of intelligence target [10]. This endeavor is both costly and risky, an aspect which limits the amount of implants in practical use. In this work we challenge this limitation and propose a low-cost, low-risk exfiltration method.

The **power supply** is the third main component of any implant. In many cases it is impossible to provide the implant with an external power supply, requiring it to survive on battery power for as long as possible. Some implants do without a power supply altogether, harvesting electromagnetic energy from an external radiation source provided by the field agent's collection device.

The intelligence agency needs to place the implants in proximity to its victims. In some cases,

a field agent will break and enter into the victim's property to place the implant in the walls of the victim's residence or hide it among the victim's belongings. Less romantically but more practically, agencies use **gifts and souvenirs**: malicious agencies may distribute free accessories, such as USB sticks, screen protectors or phone cases, which contain embedded implants, with the hope that at least one of them will end up in proximity to an intelligence target. Implants deployed using this method are required to be very effective in hiding their malicious intent. A famous example of this method was used between 1945 and 1952 to listen in on the US Embassy in Moscow [12, p. 162]. Another well known method is called **supply-chain interdiction**: according to reports published in the German magazine Der Spiegel in January 2015 [1], intelligence agencies routinely intercept deliveries of hardware which is on its way to its targets, physically attach an implant into the hardware, then forward the implanted hardware onwards toward its intended destination. In many cases the agency has only approximate knowledge about which particular item of hardware from a certain batch of shipments will actually be delivered to the intelligence target. Therefore, implants may be installed on hundreds or thousands of devices, but only a small subset of them may be ultimately activated and used. Implants delivered via interdiction have more generous operating conditions, since they may tap directly into interesting data lines or external power sources; However, they still must operate under the risk of discovery by the victim.

Our discussion will focus on an implant which specifically targets mobile phones and personal computers. Specifically, our implant is designed to be as close as possible to the victim's mobile phone, but not to be connected to any of its interfaces. Examples for such attack settings include maliciously modified phone cases, screen protectors and, quite ironically, the privacy stickers security-conscious users are increasingly using to cover their phone's cameras. We assume that the secret to be exfiltrated has already been collected and focus on the question of exfiltration. Specifically, we describe how an implant in close proximity to a mobile phone can exfiltrate a reasonable amount of data over unbounded distances through abuse of the phone, all without the awareness or permission of the user and without exploiting any software vulnerability on the phone. To do so, we exploit a particular characteristic of the micro-electromechanical (MEMS) gyroscope sensor found on virtually all phones, and on most portable personal computers.

1.1 Our contribution

In this work we present the first experimental evidence of the disruptive effect of ultrasonic vibrations on the gyroscope sensors of mobile phones and laptops. Specifically, we show how a specially-crafted audio stimulus, which was first discussed in a different context by Son et al. [11] , can cause these sensors to vibrate at their resonant frequency and falsely report that the device is being rotated rapidly. We build on this ability to construct a **stealthy communications channel** between a malicious device adjacent to the victim's phone and a command-and-control server on the Internet. As illustrated in Figure 1, the implant **modulates** the secret to be exfiltrated using ultrasonic vibrations, these vibrations affect the phone's sensor, and software running on the phone uses the phone's connectivity to finally exfiltrate the secret to the adversary's command and control server. We show that the communication channel we describe does not require any unprivileged code to run on the phone; specifically, it can be deployed in the form of an untrusted webpage. This setup does away with the requirement of deploying a field agent to collect the secret data, thus lowering the operational risk, and therefore raising the potential deployment rate, of these stealthy devices. We extensively characterize this communication channel, both empirically and theoretically, and measure its performance under various data rates and victim activity profiles. We discuss both malicious and benevolent uses of this communication channel and finally propose countermeasures at various levels that can be applied to reduce its potential for harm.

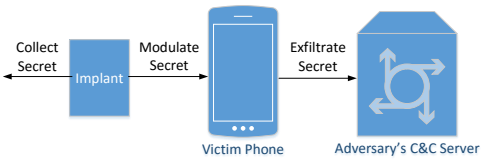

Figure 1: Attack model

Document Structure: In Section 2 we provide background about the nature and design of micro-electromechanical gyroscopes and their susceptibility to harmonic vibrations. We also discuss the access levels provided to untrusted websites and native applications who wish to access the gyroscope under different operating systems. In Section 3 we present a thorough evaluation of our proposed attack vector for multiple devices and under various environmental conditions. We conclude with a discussion of the implications of this attack, possible countermeasures and open questions in Section 4.

2 Gyroscopes on Personal Devices

Modern personal devices such as mobile phones, tablets and personal computers are equipped with various sensors such as ambient light detectors, rotation sensors, motion sensors, location monitors and so on. This Section discusses one particular sensor, the gyroscope rotation sensor, and describes

its design and the permissions model it exposes to applications.

2.1 Micro-electromechanical (MEMS) Gyroscopes

An excellent introduction to MEMS gyroscopes is given by Michalevsky et al. [9]. As stated in that work, MEMS gyroscopes sense the angular rate of rotation by measuring the magnitude of the Coriolis effect force which is acting on a moving mass within them. The mass is moving at a constant frequency named the driving frequency (f_{drive}). To improve the sensor's sensitivity and reduce its power consumption, f_{drive} usually equals to the mass mechanical resonance frequency in the driving direction, $f_{sens-res}$.

As the gyroscope is rotated, the Coriolis effect generates a force, orthogonal to the direction of the driving and the rotation. This force causes the mass to vibrate in this direction with a frequency equal to the driving frequency and an amplitude which is directly related to the angular rotation rate. The modulated vibration amplitude is then converted to voltage, typically by a capacitive or a piezo-electric sensor, and demodulated back to baseband by an analog or digital lock-in amplifier which is synchronized with f_{drive}.

2.2 Gyroscope Vulnerability Mechanisms

We consider a scenario where a vibration source is physically connected to a structure which the sensor is anchored to, and is located in close proximity to it. Several previous works demonstrated that an acoustic signal may generate false readings at the gyroscope output [11]. In those works it was assumed that the sensor is subjected to an acoustic noise from a far source. Naturally, these assumptions are not valid in our scenario and therefore we found it necessary to describe the possible mechanisms that have the potential to generate false readings at the sensor output:

A rotational mechanism: The MEMS gyroscope is soldered to a relatively long PCB which can slightly bend like a beam or a wing, depending on how and where it is anchored to the phone. Vibrations can generate a bending moment in the PCB which may rotate the MEMS gyroscope.

A linear acceleration mechanism: Similarly to the first case, vibrations can generate a linear acceleration in the MEMS gyroscope sensing direction. Modern gyroscopes possess a differential sensing mechanism that mitigates the effect of linear accelerations in the sensing direction. However, the non-ideality of the differential measurement (due to slight differences in the MEMS structure or the analog electronic circuitry) will allow part of the signal to be sensed as a valid rotational movement.

Regardless of the induced motion mechanism, the gyroscope is sensitive to those vibrations especially around two frequencies: the driving frequency, f_{drive}, and the mechanical resonance frequency in the sensing direction, $f_{sens-res}$. While vibrations in f_{drive} will induce signals that will be directly demodulated to baseband, vibrations in $f_{sens-res}$ will appear in baseband after a more complex route. First, the vibrations will be dramatically amplified by the ultra-high quality factor of the mechanical system. Second, the induced signal will be demodulated to a frequency equal to $|f_{drive} - f_{sens-res}|$. Third, the signal will be filtered by an analog low pass filter (we note that since f_{drive} and $f_{sens-res}$ are relatively close, and since the mechanical amplification in the first stage is extremely high, it is reasonable to assume that in some cases the analog low pass filter will not be able to sufficiently filter out this signal). Finally, the remaining signal will be sampled and aliased into the sensor bandwidth, appearing as a valid signal at the gyroscope output.

The exact mechanism that leads to the appearance of a signal at the gyroscope output may very between different phone and gyroscope models. Since this mechanism has no effect on the principles underlying our method, we did not investigate it thoroughly; therefore, we define the vibration frequency which generates a maximum signal as the *responsive frequency*.

2.3 The Gyroscope Programming and Permission Model

In contrast to other sensors, such as the microphone or the GPS-based location sensor, the rotation sensor is not considered as a sensitive component and thus no special privileges are required to access it. Specifically, any web page running on a modern browser can register for the `ondevicemotion()` and `ondeviceorientation()` events and subsequently be notified whenever the device is rotated. The web browser on iOS, Android and Windows devices enables this behavior without asking for the user's permission – in fact, it even does so without displaying any notifications that the gyroscope is being interrogated. Similarly, both on iOS and on Android any native app which the user downloads and installs from the first-party app store has immediate and full access to the gyroscope without any form of notification or confirmation.

As shown by [9], the *sampling rate* at which the gyroscope can be interrogated differs between web pages and native applications. This sampling rate is 60 Hz for the Chrome and Safari browsers, 100 Hz for Firefox and 20 Hz for the stock android browser. In contrast, native apps achieve a sampling rate of 100 Hz for iOS and 200 Hz for Android.

Many websites and mobile app derive income through advertising, either by embedding iframes containing ads into their web content, or by link-

ing their native apps together with third-party advertising libraries which load ads over the web on demand. As a consequence of the gyroscope's permission model, third-party ads of both types are always allowed to query the gyroscope.

3 Our Attack

As shown in [11], intentional acoustic vibrations can induce an undesired signal at the output of most MEMS gyroscopes through their several mechanisms, as we discussed in Section 2. The central point of our attack is the use of this induced signal to **carry data**: the implant can **modulate** a secret over the gyroscope channel by intentionally varying the amplitude, frequency or phase of this undesired signal over time. A program running on the mobile device can subsequently pick up this modulated signal and pass it on to the C&C server. There are two unique advantages to this exfiltration channel, in comparison to other sensor-based exfiltration schemes such as [3, 7]. The first advantage relates to the lax security model imposed on the gyroscope sensor, especially in contrast to other sensors such as the microphone or camera. This weak security makes deploying an attack based on gyroscopes very simple. In fact, as stated in Subsection 2.3, the victim merely has to browse to an unprivileged web page for the attack to succeed. The second advantage relates to the gyroscope's enhanced sensitivity at its responsive frequency. Due to this sensitivity, a relatively weak audio signal (as low as several microwatts in power, as we show in Subsection 4.1) is sufficient to trigger the phone's sensors, allowing even a small battery-powered implant to make use of this exfiltration method.

Figure 2 provides a brief demonstration of our attack, based on real lab measurements. The top of the figure shows the baseband bit sequence that the implant wishes to exfiltrate. The bit sequence is transmitted to the phone by an on-off keying modulation of the audio signal, with the frequency of the carrier wave set near the gyroscope's responsive frequency. The bottom of the Figure shows the absolute values of real-time readings from the victim phone's gyroscope (an iPhone 5S in this case) as it receives the audio signal, captured using JavaScript code running within an unprivileged web page. It can be seen, even with the naked eye, that the readings from the gyroscope experience strong fluctuations when the audio signal is being sent, but are relatively quiet during other periods. Thus, the gyroscope readings contain an encoding of the transmitted bit sequence. These readings can then trivially be sent to a C&C server, providing a very effective exfiltration channel. We describe our results in more detail in the following Section.

3.1 Attack Model

Our attack assumes that the adversary has managed to place an implant in close proximity to the gyroscope sensor located in the victim's phone or mobile device, and that this implant has some secret it wishes to exfiltrate to the adversary's command and control (C&C) server. We furthermore assume that the attacker has the ability to make the victim's mobile device display a website, or otherwise run some unprivileged code. This assumption can be achieved by using one of the following methods:

- By purchasing a advertisement, containing the attacker's JavaScript code, which will then be displayed on one of the victim's favorite websites or native applications. Malicious advertisements are a well-known risk to the advertising ecosystem [14]. As we stated in 2.3, displaying an ad that interacts with the gyroscope does not require special permissions from the hosting webpage or native app.

- By inducing the user to download and install a "repackaged" application – an innocent-looking native application modified to include additional malicious components [15]. Note again that the additional functionality requires no extraordinary permissions, making it a good candidate for repackaging attacks .

- By replacing the contents of an innocent webpage the victim is attempting to view with an infected version containing the malicious functionality, through the use of state-actor capabilities such as man-in-the-middle or man-on-the-side attacks [5, 8].

The malicious functionality embedded into the website or the app is very simple – it simply queries the gyroscope as quickly as possible and uploads its reading to a central server. The implant will use intentional acoustic vibration to selectively corrupt the readings of the gyroscope as they are being read by the attacker's code, therefore modulating the secret to be exfiltrated. [1]

3.2 Evaluation Setup

We designed and carried out an experiment to evaluate the data-bearing potential of the intentional acoustic vibration channel. The hardware setup of our experiment is indicated in Figure 3. As shown in the Figure, a Keysight 33622A Waveform Generator was connected via an RG-58 coaxial cable to a PUI Audio APS2509S-T-R piezoelectric transducer, which was placed on the victim device as

[1] We must assume that the implant knows to start transmitting precisely when the malicious code is running on the phone. Synchronizing the implant and the code can either be done by sending a signal from the phone that is picked up by the implant's sensors, or by fixing a predetermined time of day at which the implant always transmits its payload.

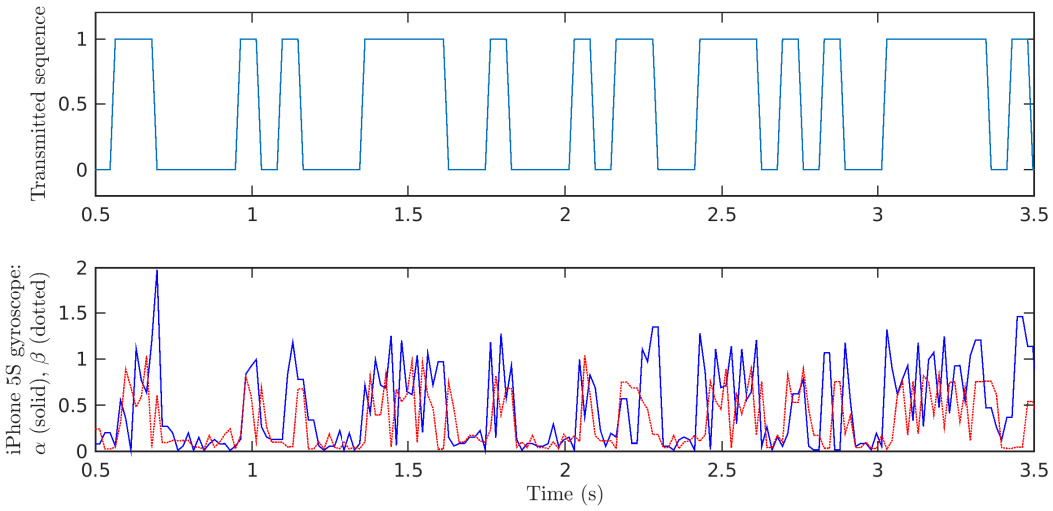

Figure 2: A transmitted PRBS sequence is received by an iPhone 5S gyroscope

Figure 3: Experiment Setup

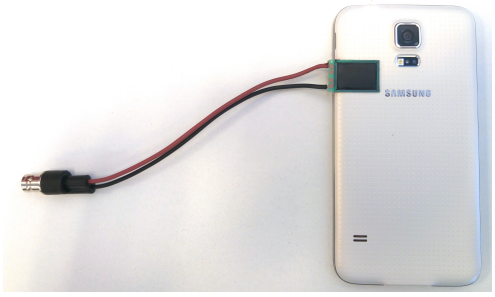

Figure 4: A phone with an attached transducer

close as possible to the location of the device's internal gyroscope. Figure 4 shows a photograph of the victim phone and the attached piezoelectric transducer.

While the PUI Audio piezoelectric transducer's data sheet states that its highest working frequency is 20 kHz, we were consistently able to use it to generate tones at frequencies of up to 30 kHz. We determined the exact responsive frequency for each device by generating a sine-sweep signal in the 25-29 kHz range, looking for anomalies in the gyroscope response, then gradually reducing the span of the sweep until we arrived at the exact responsive frequency. To determine the optimal location for the piezoelectric transducer, we referred to publicly-available tear-downs of the victim devices and attempted to locate the speaker as close to the gyroscope as possible. If tear-downs were not available, we manually moved the piezoelectric transducer across the device, while vibrating at the responsive frequency, until we detected a strong response at the gyroscope's output.

The waveform generator was configured to create an on-off keying-modulated signal at its output. The carrier frequency of this output was a sine wave at a frequency close to the responsive frequency of the gyroscope of the victim device (typically between 26 kHz and 28 kHz) and an amplitude of 10 V_{pk-pk}. The modulating signal was the standard pseudorandom bit-sequence (PRBS) PN7, which is created with 7 bits of state and the generating polynomial $G(X) = x^7 + x^6 + x^0$.

The devices and software environments used in our experiment are listed in Table 2. As the table shows, we successfully evaluated devices from multiple hardware vendors, running multiple operating systems (Windows, iOS and Android) and using both native applications and web browsers.

On the software side, we wrote a simple webpage that constantly queries the gyroscope using JavaScript and uploads the measurements on demand to a web server. We also wrote a native Android app which queried the gyroscope at the highest possible rate and uploaded its measurements to the same web server. The web server, which we implemented in node.js, simply time-stamped each batch of measurements and saved them to disk. Finally, we analyzed the measurements using custom scripts written in Matlab R2015a. In the analysis step, we determined the optimal phase for detection by cross-correlating the gyroscope signal with a locally-generated PN7 sequence, then applied a simple threshold-based detector to determine the values of each bit. Finally, we calculated the bit error rate by counting how many bits were incorrectly decoded by our method. As we state in Subsection 4.1, it is certainly possible to improve this modulation scheme and increase the channel's capacity while reducing its error rate.

Device Name	Gyroscope Hardware	Software Environment	Max. Sampling Rate
Apple iPhone 5s	Unknown (STMicroelectronics?)	iOS 9.3 (Safari)	60 Hz
Samsung Galaxy S5	Invensense MP65M	Android 5.1 (Chrome)	60 Hz
-	-	Android 5.1 (Native App)	200 Hz
Microsoft Surface Pro 3	Unknown (Bosch Sensortec?)	Windows 10 (IE Edge)	60 Hz

Table 2: Devices under test

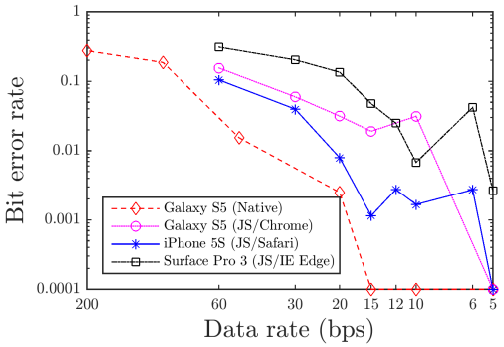

Figure 5: Bit error rates for different devices

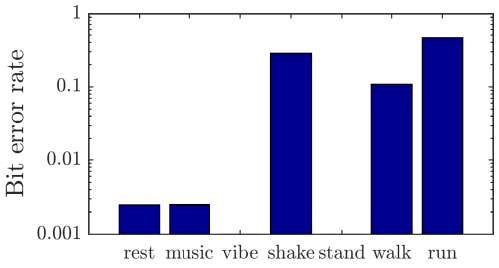

Figure 6: Bit error rates for various activities

3.3 Evaluation Results

Figure 5 shows the bit error rates we achieved using our basic decoder. The horizontal axis shows the bit rate chosen for the modulating bit sequence, while the vertical axis shows the average achievable bit error rate for this bit rate using our decoding setup. All bit error rates were measured when the devices were at rest on a flat surface. Our decoding setup achieved practical error rates on all three evaluated hardware platforms, even at the highest sampling rates supported by the software setup. For example, a 60 bps sequence sent to the iPhone 5S was received with an error rate of 11%. As expected, increasing the amount of samples per data bit, either by reducing the data rate or by increasing the sampling rate (by moving from a webpage to native code) resulted in a lower overall bit error rate. As we discuss in Subsection 4.1, the physical characteristics of the channel indicate that much higher bit rates can be achieved using advanced modulation techniques and a better decoder. It is important to note that the victim cannot detect that exfiltration is in progress – the frequency of the audio signal generated as part of our attack is far beyond the human hearing range, and its amplitude is too low to be detected as motion.

Figure 6 shows the bit error rate exhibited by our stealthy channel under various user activity profiles. All of the measurements were carried out on a Samsung Galaxy S5 device with a piezoelectric transducer glued to its plastic back cover, running the native app at a data rate of 20 bits per second. Bit error rates were first measured when the phone was at rest on a flat table. The experiment was repeated while the phone was playing music through its speakerphone. Then, it was repeated while the phone was vibrating as a result of an incoming call. Next, the bit error rate was measured while the phone was being shaken vigorously. Finally, the phone was placed in the front pants pocket of an experimenter and the data rates were recorded while the experimenter was standing idly, walking at 2 km/h and running at 6 km/h. As the results show, we achieve virtually error-free communications under low to moderate amounts of phone motion, but vigorous motions such as running or shaking make our scheme less practical to use, at least using the basic decoder evaluated in this work.

4 Discussion

Our results show how a low-powered implant can take advantage of a mobile phone's sensors and connectivity to exfiltrate secret data. We believe that the methods we discuss in this paper are more troubling than conventional exfiltration methods such as radio backscatter, since they allow the intelligence agency to monitor many implants at the same time at a low cost, with no risk of exposure to their field agents. State actors typically spread thousands of implants through supply-chain intervention or other methods, but only interrogate a few dozens due to the operational costs and risks involved with signal collection. This new attack vector changes the economics of state-sponsored attacks, and may induce malicious intelligence agencies to activate all of the implants they distribute,

not only a selected few, thus drastically raising the amount of people targeted by hardware-based spying methods.

4.1 Capacity and Power Bounds

The channel capacity according to Shannon-Hartley theorem is:

$$C = BW \log_2 (1 + SNR)$$

Where BW is the bandwidth of the channel in Hz and SNR is the signal power to noise power ratio. To estimate the error-free capacity limit of the Samsung Galaxy S5 phone piezo-gyro channel, we start by measuring its bandwidth. A Keysight 33509B function generator was operated as a voltage source to excite the piezoelectric crystal with a 10 V amplitude sine sweep. The noise was measured in a quiet room during night. To further suppress outer vibration interference, the phone was placed on top of a passive vibration isolation platform (model 25BM-4 made by Minus k Technology). The noise measurements results are presented in Table 3 for the average of 20 measurements of 100 seconds each[2].

We have found that since the data output frequency is considerably lower than the analog bandwidth around the excitation frequency, the communication channel's bandwidth is limited by the gyroscope sampling rate. As indicated in the Table, the theoretical capacity of the gyroscope channel is more than 1 kbps, even using a low sampling rate of 60 Hz, and it grows to over 4 kbps as the sampling rate increases. This compares well with other sensor-based exfiltration schemes based on the phone's microphone or magnetometer [7, 3]

We note that the results at Table 3 are given only for the the single gyroscope channel which produced the highest SNR. Combining the outputs of multiple channels may further improve the signal-to-noise ratio. To approach the theoretical capacity in Table 3, one can for example use high order modulation schemes or OFDM, combined with a high efficiency code such as turbo or LDPC.

In the second part of the study, we examined the effect of the excitation voltage and the power consumption of the piezoelectric crystal on the channel capacity. The crystal was excited by a sine sweep with different amplitudes while the gyroscope response was measured at a sampling rate of 200 Hz. The current was measured by a low-noise current probe (model i-prober 520 made by Aim-TTi Instruments). The results are presented in Figure 7.

As shown in the Figure, the gyroscope-based transmission channel achieves good data rates even at power levels as low as -21 dBm (7μW). This suggests that it should be possible to power the exfiltration device using passive power harvested from the phone or from other nearby radiation sources, allowing an implant to operate without a battery.

[2]The sensor readings returned by Firefox were multiplied by the software by a constant factor. This does not affect the ultimate SNR calculations.

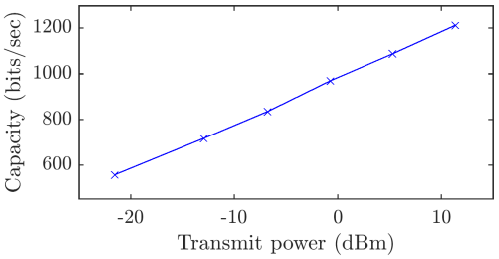

Figure 7: Channel capacity as a function of transmit power (60 samples per second)

4.2 Countermeasures

Several aspects of the phone's attack model work together to make the attack possible. Disrupting any of them can be an effective countermeasure to the attack we described.

The most critical contributor to the effectiveness of our attack is the fact that untrusted apps and web pages are able to access the gyroscope at will. A natural response to this risk would be to require permission to access the gyroscope. While effective, we believe this additional security step will not solve the problem altogether: in contrast to security warnings for issues such as expired or revoked certificates, which are reliable indicators of risky situations, it is quite reasonable for a webpage (such as a game) to ask permission to use the gyroscope. Therefore, users may not have enough information to decide whether gyroscope access should be allowed or denied in certain situations. Nevertheless, even if this mechanism will only partially mitigate the risk, we still consider it worth deploying.

Another possible mitigation would be to prevent web pages from accessing the gyroscope if they are judged risky according to some heuristic. Web pages already have a method to limit the permissions of embedded content using the iframe sandbox attribute [13, §4.7.2], and an extension to this attribute to limit sensor access would be a good addition. In another promising move in this direction, Google announced in September 2015 that they are intending to limit access to several powerful features, including device motion and orientation, to web pages delivered from insecure origins [2]. The first powerful feature which was removed was the geolocation API, which is not available to insecure origins starting from Chrome version 50 (deployed April 2016). While there is no currently planned deployment date for the restriction to the gyroscope, the current version of the Chrome browser already displays a warning message in the developer console whenever the gyroscope is accessed using JavaScript from an insecure origin. Once this countermeasure is in place, an adversary will be required to obtain a certificate for its

Sampling rate (Hz)	Software	Signal (Rad/s)	Noise (mRad/s)	SNR (dB)	Capacity (bps)
60	Chrome	2.09	0.853	67.7	1351
100	Firefox	112	53	66.4	2209
200	Native	2.17	0.92	67.4	4481

Table 3: Channel capacity for different sampling rates

C&C server from an external certification authority before it can access the gyroscope, a fact which will substantially increase the cost and risk of creating and deploying the command and control server.

A countermeasure suggested by Michalevsky et al. [9] to counter other gyroscope-based attacks was to lower the rate at which apps can sample the on-board gyroscope. Unfortunately the attacks described in our paper are possible even if the sampling rate of the gyroscope is throttled to its lowest practical value of 20 Hz. Similarly, reducing the signal-to-noise ratio of the gyroscope by filtering the signal or intentionally jittering the output will only reduce the capacity of the channel but not eliminate it altogether.

Physical modifications to the gyroscope itself can also mitigate the problem – analog anti-aliasing filters at the entrance to the gyroscope's internal sampling circuits can reject the high-frequency oscillations at the responsive frequency, while mechanical damping or sound isolation can reduce the effects of acoustic noise on the gyroscope. Sadly, these countermeasures tend to increase the cost, the power consumption and the physical dimensions of the gyroscope, making it highly unlikely that they will be integrated into gyroscopes destined for phones or other consumer devices.

An interesting countermeasure could be based on **sensor fusion** – the phone has multiple location and orientation sensors (accelerometer, gyroscope, magnetic compass, GPS, etc.). Intentional acoustic vibration corrupts only the readings from the gyroscope, but leaves all other readings unaffected. It may be possible to design a mechanism that correlates readings from multiple sensors and suppresses the readings from the gyroscope if they do not agree with outputs from other sensors on the phone.

From the user's standpoint, perhaps the best countermeasure would be to consider the close perimeter of the phone to be as sensitive as the phone itself. Thus, security-conscious users should be careful of using screen protectors, phone cases or privacy stickers with a questionable pedigree.

4.3 Benevolent use of Gyroscope-based Modulation

The very features of the gyroscope communications channel which make it so desirable for malicious adversaries – namely, its ubiquity, its minimal power requirement, and its stealthiness – make it ideal for beneficial uses, most immediately for **two-factor authentication**. In a common two-factor authentication scenario, users are given an secure device, called an authenticator, which displays a constantly-changing numeric code, and are expected to type in this code in addition to their password as they log in to a high-security service such as a bank or health care provider. The fact that users must manually copy the numeric code from the authenticator into the login page is a cause for user errors and frustration. In addition, the digits displayed by the authenticator must be large to be read by human users. This places a lower bound on the size of the authenticator and exposes the scheme to snooping attacks, either by a shoulder-surfing adversary or by a camera.

Using gyroscope communications instead of manual key entry is an excellent way of addressing these flaws – instead of a digital display, the authenticator can contain a small piezoelectric transducer which transmits amplitude-modulated data at the gyroscope's responsive frequency. A human is not required to read the audio output, resulting in a device which can be as small and light as desired. In addition, the use of gyroscope-based communications can reduce the opportunity for user error and reduce the risk of outside snooping. Most importantly, no hardware or software changes must be made to currently deployed phones to enable this behavior.

The authors of [7] suggested several uses for their magnetic-based communications scheme which can also be applied here – namely, our scheme can also be used as a replacement for QR codes, and it can also be used for device pairing, assuming both devices are equipped with high-frequency speakers as well as gyroscopes.

4.4 Responsible Disclosure

A preliminary draft version of this report has been shared with the vendors of the hardware and software we evaluated. Browser vendors were hesitant to limit programmatic access to the gyroscope, since this step would potentially break the functionality of many existing webpages while providing a security benefit only in very limited cases. The sensor working group at the web standards body W3C were more receptive, and notified us that the next generation sensor API will include support for limiting sensor access to secure contexts. In addition, the new sensor API is planned to offer user more explicit control over sensor permissions.

4.5 Related Work

Several existing works have explored the susceptibility of gyroscope sensors to external noise and their potential use for malicious intent. Son et al. [11] demonstrated how intentional acoustic vibrations can corrupt the gyroscope readings in a remote-control drone, causing it to crash. The authors also characterized a large variety of MEMS gyroscopes, showing that the precise effect of intentional acoustic vibration depended on the make and model of the MEMS gyroscope. The work of Son et al. did not consider the data bearing capacity of the intentional noise channel. Michalevsky et al. presented a work titled "Gyrophone" [9], in which a mobile phone's gyroscope was treated as a low-frequency microphone and used to record and recognize speech. The experimental evaluation in Michalevsky et al.'s paper was done with an externally-powered standalone speaker system with a peak power rating of 50 watts, playing back recorded speech at a high sound power level. In contrast, the very short distance between the piezoelectric transducer and the gyroscope in our attack, combined with the gyroscope's enhanced response at its responsive frequency, allows us to use an extremely low-power audio signal, on the order of several microwatts, as we showed in Subsection 4.1. This allows our exfiltration mechanism to be battery powered, or even passively powered using energy harvesting techniques similar to those used by RFID tags. We note that while several of the countermeasures suggested by Michalevsky et al. will also be effective in mitigating our attack, reducing the sampling rate of the gyroscope will only slow down the exfiltration process by some factor, but not prevent it altogether.

In 2014 Jiang et al. published a work discussing how a phone's magnetic compass can be used for low-rate communication purposes [7]. Their work is similar to ours since both use a mobile phone's sensors for short range communications. However, the work of Jiang et al. did not consider the adversarial setting we discuss in this work, where a malicious party is using the sensor to communicate without the victim's awareness or permission; in fact, due to the Android permission model, the magnetometer cannot be used for this purpose since using it requires additional app permissions. In general, the modulating magnetic field can be generated from a greater distance from the phone's sensor. Conversely, if the magnetic field modulator is placed in close proximity to the phone's magnetometer it can be made much smaller than a piezoelectric transducer. However, the signal to noise ratio of the magnetic channel is considerably lower than that of the gyroscope-based channel, resulting in a lower potential bit rate. The magnetometer is also more sensitive to noise generated by power supplies, engines, and other similar devices.

The use of ultrasonic audio as a covert communications method was suggested and evaluated by Hanspach and Goetz in 2013 [6]. Also in 2013, security researcher Dragos Ruiu reported on a type of malware named BADBIOS which uses ultrasonic communications as a command and control channel. In 2014 Deshotels demonstrated a covert channel between two mobile devices based on audio waves in the 18 kHz - 19 kHz frequency range, using the phone's internal speaker and microphone [3]. This covert channel achieved a bit rate of over 300 bits per second at distances of over 10 meters. We note that the iOS security model does not allow untrusted web pages to play audio unless the user performs some action first (such as touching the screen), while untrusted applications are by default prevented from using the microphone at all.

4.6 Conclusion

In this work we demonstrated and evaluated a low-cost exfiltration method based on intentional acoustic vibration. This method allows an implant to take advantage of the gyroscope sensor of an adjacent mobile device to exfiltrate secrets to a command and control center. This method has the potential of reducing the operational risk involved in operating implants, a fact that may dramatically expand their use by malicious state agencies. Security-conscious users should not allow questionable accessories, such as phone cases, to be in close physical contact with their phones.

References

[1] APPELBAUM, J., GIBSON, A., GUARNIERI, C., MÜLLER-MAGUHN, A., POITRAS, L., ROSENBACH, M., RYGE, L., SCHMUNDT, H., AND SONTHEIMER, M. The digital arms race: NSA preps america for future battle. *Der Spiegel 1*, 17 (Jan 2015).

[2] CHROMIUM SECURITY TEAM. Deprecating powerful features on insecure origins. Online at https://www.chromium.org/Home/chromium-security/deprecating-powerful-features-on-insecure-origins.

[3] DESHOTELS, L. Inaudible sound as a covert channel in mobile devices. In *8th USENIX Workshop on Offensive Technologies, WOOT '14, San Diego, CA, USA, August 19, 2014.* (2014), S. Bratus and F. F. X. Lindner, Eds., USENIX Association.

[4] GENKIN, D., PACHMANOV, L., PIPMAN, I., TROMER, E., AND YAROM, Y. ECDSA key extraction from mobile devices via nonintrusive physical side channels. *IACR Cryptology ePrint Archive 2016* (2016), 230.

[5] HAAGSMA, L. Deep dive into QUANTUM INSERT. Online at https://blog.fox-it.com/2015/04/20/deep-dive-into-quantum-insert/.

[6] HANSPACH, M., AND GOETZ, M. On covert acoustical mesh networks in air. *JCM 8*, 11 (2013), 758–767.

[7] JIANG, W., FERREIRA, D., YLIOJA, J., GONÇALVES, J., AND KOSTAKOS, V. Pulse: low bitrate wireless magnetic communication for smartphones. In *The 2014 ACM Conference on Ubiquitous Computing, UbiComp '14, Seattle, WA, USA, September 13-17, 2014* (2014), A. J. Brush, A. Friday, J. A. Kientz, J. Scott, and J. Song, Eds., ACM, pp. 261–265.

[8] MARCZAK, B., WEAVER, N., DALEK, J., ENSAFI, R., FIFIELD, D., MCKUNE, S., REY, A., SCOTT-RAILTON, J., DEIBERT, R., AND PAXSON, V. An analysis of china's "great cannon". In *Free and Open Communications on the Internet (FOCI)* (2015), USENIX.

[9] MICHALEVSKY, Y., BONEH, D., AND NAKIBLY, G. Gyrophone: Recognizing speech from gyroscope signals. In *Proceedings of the 23rd USENIX Security Symposium, San Diego, CA, USA, August 20-22, 2014.* (2014), K. Fu and J. Jung, Eds., USENIX Association, pp. 1053–1067.

[10] SANGER, D. E., AND SHANKER, T. NSA devises radio pathway into computers. *The New York Times 1*, 15 (Jan 2014).

[11] SON, Y., SHIN, H., KIM, D., PARK, Y., NOH, J., CHOI, K., CHOI, J., AND KIM, Y. Rocking drones with intentional sound noise on gyroscopic sensors. In *24th USENIX Security Symposium, USENIX Security 15, Washington, D.C., USA, August 12-14, 2015.* (2015), J. Jung and T. Holz, Eds., USENIX Association, pp. 881–896.

[12] UNITED STATES DEPARTMENT OF STATE BUREAU OF DIPLOMATIC SECURITY. *History of the Bureau of Diplomatic Security of the United States Department of State*. Global Publishing Solutions, 2011.

[13] WORLD WIDE WEB CONSORTIUM. The iframe element. Online at https://www.w3.org/TR/html5/embedded-content-0.html#attr-iframe-sandbox.

[14] ZARRAS, A., KAPRAVELOS, A., STRINGHINI, G., HOLZ, T., KRUEGEL, C., AND VIGNA, G. The dark alleys of madison avenue: Understanding malicious advertisements. In *Proceedings of the 2014 Internet Measurement Conference, IMC 2014, Vancouver, BC, Canada, November 5-7, 2014* (2014), C. Williamson, A. Akella, and N. Taft, Eds., ACM, pp. 373–380.

[15] ZHOU, W., ZHOU, Y., JIANG, X., AND NING, P. Detecting repackaged smartphone applications in third-party android marketplaces. In *Second ACM Conference on Data and Application Security and Privacy, CODASPY 2012, San Antonio, TX, USA, February 7-9, 2012* (2012), E. Bertino and R. S. Sandhu, Eds., ACM, pp. 317–326.