Sampling Race: Bypassing Timing-based Analog Active Sensor Spoofing Detection on Analog-digital Systems

Hocheol Shin[1], Yunmok Son[1], Youngseok Park[1,2], Yujin Kwon[1], and Yongdae Kim[1]

[1]Korea Advanced Institute of Science and Technology
{*h.c.shin, yunmok00, dbwls8724, yongdaek*}*@kaist.ac.kr*
[2]NAVER Labs
ys.park@navercorp.com

Abstract

Sensors and actuators are essential components of cyber-physical systems. They establish the bridge between cyber systems and the real world, enabling these systems to appropriately react to external stimuli. Among the various types of sensors, active sensors are particularly well suited to remote sensing applications, and are widely adopted for many safety critical systems such as automobiles, unmanned aerial vehicles, and medical devices. However, active sensors are vulnerable to spoofing attacks, despite their critical role in such systems. They cannot adopt conventional challenge-response authentication procedures with the object of measurement, because they cannot determine the response signal in advance, and their emitted signal is transparently delivered to the attacker as well.

Recently, PyCRA, a physical challenge-response authentication scheme for active sensor spoofing detection has been proposed. Although it is claimed to be both robust and generalizable, we discovered a fundamental vulnerability that allows an attacker to circumvent detection. In this paper, we show that PyCRA can be completely bypassed, both by theoretical analysis and by real-world experiment. For the experiment, we implemented authentication mechanism of PyCRA on a real-world medical drop counter, and successfully bypassed it, with only a low-cost microcontroller and a couple of crude electrical components. This shows that there is currently no effective robust and generalizable defense scheme against active sensor spoofing attacks.

1 Introduction

Sensors observe environments by measuring physical quantities and converting them to electrical signals, typically used for automation in *sensing and actuation systems*, such as self-driving cars and drones. In a sensing and actuation system, a processor uses the output data from the sensor to control system actuators. For example, gyroscopes and accelerometers measure angular velocities and accelerations to determine the position and direction of a drone system. In addition, radars and lidars gauge distances by calculating time differences between an emitted signal and its echo.

To ensure the correct operation of sensing and actuation systems, *sensing systems*, composed of sensors and environment characterization algorithms using the output of those sensors must be robust not only against the various naturally arising sources of noise and errors but also against the intentional fabrication of environments, such as sensor spoofing attacks. However, most developers of sensing systems have considered only natural disturbances or errors originating from the environment, and intentional attacks by adversaries have been largely ignored. Therefore, it is important to devise defense schemes that can protect sensors against such intentional fabrication. Such defense schemes should be robust with the guarantee that the attacker cannot bypass them. Furthermore, the universality of such defensive measures is also important because defense schemes that depend on the characteristics of a specific type of sensor cannot be applied to sensors without them. Lastly, the cost for such defenses should be in a reasonable range. Even if a scheme is both robust and universal, impractical cost that requires far more resources than the sensing function itself cannot be taken seriously by manufactures.

Shoukry et al. introduced PyCRA [27], an authentication scheme to detect spoofing attacks against active sensors, based on an analog challenge and response mechanism at the 2015 ACM CCS, which is the only work claimed to be robust and generalizable. Active sensors are a special type of sensors that emit physical signals and listen to echoes or variations of the transmitted signal to measure the target. The basic idea of PyCRA is that spoofing attempts can be detected if an active sensor receives an echo when no signal has been transmitted for a random amount of time. Attackers cannot impersonate a legitimate challenge and response process, because they cannot predict when the challenge will be transmitted. As a result, a nonzero delay lower bounded by the physical delay of the attacker will always accompany their responses to challenges.

Despite this, we discovered a fundamental vulnerability of PyCRA, and designed new attacks exploiting this vulnerability and verified the result both analytically and experimentally. Analytically, we show that PyCRA reduces to a race of sampling rate between the attacker and the victim, unless the victim's sampling rate is high enough to detect the fundamental physical delay of the attacker. We further derive a sufficient condition for the attacker to avoid detection. We also show experimentally that a victim's sampling rate of over 350 kHz can be easily bypassed by an attacker, even with a low-cost microcontroller and crude additional circuitries [1]. Existence of our attack shows that design of robust and generalizable defense mechanism for active sensors remains as an *open problem*. Finally, we discuss several anti-spoofing measures for active sensors, and also examine their own limitations to becoming a robust and universal active sensor spoofing defense. In summary, the main contributions of this research are as follows:

- We analyze the security of PyCRA in theory, and derive a sufficient bound for attack success.
- We experimentally prove that these limitations can be exploited in practice with low-cost hardware.

The rest of this paper is organized as follows: Section 2 provides background concerning active sensors, sensor spoofing, various defense forms against active sensor spoofing, and the PyCRA authentication mechanism. Section 3 describes the attack model for sensor spoofing detection. Our analysis and experiments addressing the limitations of PyCRA are described in sections 4 and 5, respectively. We discuss other spoofing defense approaches for active sensors and their limitations in Section 6, and related works in Section 7. Section 8 concludes the paper.

2 Background

In this section, we explain the definition and classification of active sensors, and review previous sensing attacks on active sensors. A few defense mechanisms for sensor spoofing attacks follow, with PyCRA as an example detection mechanism.

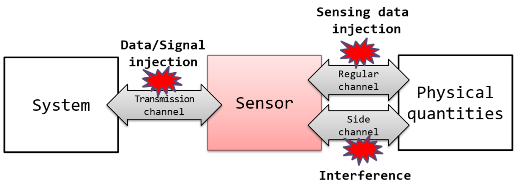

Figure 1: Three attack vectors for sensor systems

2.1 Active Sensor

Sensors can be categorized into passive and active ones, depending on the mechanism they use to measure the environment. Passive sensors simply receive natural emissions from the surroundings. Owing to the energy attenuation to be measured, they are generally used for relatively short-distance applications. For example, typical passive sensors such as gas, ambient light, and pressure sensors can measure physical quantities only in their vicinity. In contrast, active sensors have emitters to transmit physical signals, and passive sensors (receivers) to receive the response of the measured entity (via the channels between emitters and receivers). They measure physical quantities based on the difference between the emitted and received signals.

Although active sensors tend to be more complex, they are generally more effective. First, they can selectively amplify signals from the objects of interest by illuminating them with their emitters. With this property, they can even measure objects that cannot be measured by only receiving ambient signals. Similarly, active sensors can suppress the noise power to maximize the signal-to-noise ratio, and measure objects of interest even in a noisy environment. Second, because the transmitted signal is well known, they can extract more information of the received signal by comparing the transmitted and received signals. For example, radars, lidars, and sonars gauge distances to objects by comparing transmitted signals (radio waves, light, and acoustic waves) and received their echoes, which would not be possible unless the sensor fully recognizes the transmitted signal.

2.2 Sensor Spoofing

Sensor spoofing refers to the injection of a malicious signal into the victim sensor, so that the victim believes the injected signal to be legitimate. This can be accomplished using various attack vectors, according to the type and characteristics of the victim sensor.

Attack Vectors for Sensor Spoofing Sensor spoofing can be categorized by the particular channel exploited. As shown in Figure 1, there are three major channels exposed to spoofing attacks. The first one is called the *regular channel*, and is the very physical interface the victim

[1] In PyCRA, the authors used 30 kHz of sampling rate for their experiment, and stated that the range of 200 kHz is considered to be a sampling rate for high-end microcontrollers. Our experiment uses worse condition for attackers from the perspective of PyCRA.

sensor depends on. For example, the regular channel for ambient light sensors is light.

The second one is the *transmission channel*. Sensors generally do not operate isolated; their output is typically used to support higher-level functions. Therefore, the sensor output must be transmitted to the remainder of the system via transmission channels. Such transmission channels can assume various forms: wired/wireless, and analog/digital. Sometimes, however, these transmission channels can be influenced by external physical stimuli, an effect by which this attack vector is anchored. For example, Foo Kune et al. [14] succeeded in injecting a voice waveform into a Bluetooth headset by intentional EMI in the wire connecting the microphone with the remainder of the system. As a result, the injected voice waveform was treated as if it had been a real voice, even though there was no sound nearby.

The last one is called the *side channel*. Every sensor comprises a transducer, to translate the physical signal of interest into another type of signal (generally electric). Sometimes, transducers are affected by physical stimuli other than those they are supposed to sense. We call such paths side channels. For instance, the internal structures of MEMS gyroscopes are known to be affected not only by movement but also by acoustic noise [4, 6, 7]. Son et al. [28] exploited this property to incapacitate a drone.

All three aforementioned attack vectors have identical outcomes: modification of the sensor output. Furthermore, owing to structural limitations, the victim sensors cannot distinguish the legitimate input and the deceiving input given via these attack vectors. This leaves the detection of spoofing attacks entirely up to the system behind the sensor. Therefore, attackers can exploit any of these vectors to spoof the victim sensor output.

Active Sensor Spoofing via the Regular Channel Of the three attack vectors, active sensors are more prone to be exposed to regular-channel attacks, especially when they are used in remote sensing applications. In such applications, the channel connecting emitters and receivers is publicly exposed, and attackers can freely access it. Furthermore, to maximize sensitivity, these sensors generally use directional receivers, oriented toward the direction of the incoming echoes of the transmitted signal. This allows attackers to easily infer the most effective attack direction. Consequently, without any defense mechanisms, attackers can easily influence the victim's input signal without any authentication or authorization by simply generating the same type of physical quantities used by the victim sensor.

2.3 Defense Categorization

We consider that defenses for sensor spoofing attacks generically fall into one of the following categories:

Spoofing Detection In this type of defense, the defender can detect the malicious deceiving signal. However, the sensor still remains vulnerable to the spoofing attack. Once the defender detects a spoofing attempt, it can activate its defensive measures.

Signal Integrity In contrast to *Spoofing Detection*, this type of defense cannot detect spoofing attempts. However, it is resilient against spoofing attacks.

Signal Recovery In this type of defense, the defender identifies the spoofing signal in the received waveform, and removes it. Although it is the most difficult to achieve, this type of defense is the most complete form of defense against sensor spoofing attacks, encompassing both *Spoofing Detection* and *Signal Integrity*.

2.4 Concept of PyCRA

PyCRA was devised as a generalizable regular-channel active sensor spoofing detection and recovery scheme, which falls into the categories of *Spoofing Detection* and *Signal Recovery*. We focus on bypassing detection mechanism only, because its signal recovery process has no effect when the attacker can avoid detection.

To detect spoofing attempts, PyCRA sets several traps to identify ongoing spoofing attempts: a simple detector, a confusion phase, and a χ^2 detector. The simple detector [27, § 3.1] is the most basic detector, based on the fundamental idea of PyCRA. The confusion phase [27, § 3.2] is an additional trap, devised to make spoofing more difficult. Lastly, the χ^2 detector [27, § 3.4] is a more advanced detector, designed to identify spoofing attempts during the transition time. We briefly introduce them here, to show how PyCRA works.

Simple Detector On active sensors equipped with simple detectors, the emitter signal is turned off at random instants to verify the existence of any spoofers. Right after the emitter is turned off, the receiver carefully monitors the intensity of incoming signals. If there are no spoofing attempts, it will receive nothing. Any spoofing attempts will appear at the receiver with nontrivial signal intensity, because the attacker cannot predict the instant of the random challenge. In other words, no matter how fast the attacker responds, a nonzero physical delay in turning off the spoofing signal is unavoidable. As a result, spoofing attempts can be detected under following conditions. First, the physical delay of the attacker is long enough so that the attacker cannot turn the spoofing signal off even before the transition time finishes. Second, the victim sensor has sufficient time precision to detect the minute interval of nonzero signal during the attacker's physical delay. Finally, if the victim sensor is digital, its sampling interval is shorter than the physical delay of the attacker.

Confusion Phase The confusion phase is an additional trap to confuse the attacker. When challenging, the victim sensor's transmitted power is first lowered to a level that slightly exceeds the noise level; this is called the confusion phase. It lasts a random period, and after that period, the emitter signal is completely turned off and enters the silent phase. If the attacker continues to spoof the victim even after the confusion phase starts, there will be a nonzero probability that the attacker misses the change in the emitter signal, which results in the detection of the spoofing attempt. We note here that the confusion phase concept assumes the attacker will not stop spoofing during the confusion phase. We will later discuss if it is possible to bypass this mechanism (Section 4.2).

χ^2 **Detector** The main limitation of the simple detector is that it does not consider the case when the attacker has comparably shorter transition time than the victim. The χ^2 detector tries to solve this by adopting the mathematically modeled dynamic characteristics of the victim. The main idea of this detector is explained below, using a modified version of the mathematical expressions, without the loss of significant ideas from the original.

First, an accurate sensor model should be acquired, so that the real input and output signals of the victim sensor can be modeled as in Eq. (1).

$$x(t+\delta) = f(x(t)) + w(t) \tag{1a}$$

$$y(t) = h(x(t)) + v(t) \tag{1b}$$

Here, $x(t)$ and $y(t)$ denote the transmitted and received signals of the victim sensor, respectively. Note that $x(t)$ includes the random challenges (i.e., output modulation), and $x(t+\delta)$ represents the next emitter output immediately after $x(t)$. f and h are precisely modeled transition functions, where f determines how $x(t)$ evolves over time, and h denotes the input-output transfer function of the sensor. $w(t)$ and $v(t)$ are mismatch terms, responsible for filling up the gap between modeling and reality.

Apart from real values, we can derive the estimated emitter signal, $\hat{x}(t)$ from the modeling by solving $\hat{x}(t+\delta) = f(\hat{x}(t))$, and $\hat{y}(t)$ as $h(\hat{x}(t))$. The residual denoting difference between estimation and measurement is derived as $z(t) = y(t) - \hat{y}(t)$. Based on these notations, the χ^2 detector operate as follows.

1. Select a random instant: $t_{challenge}$, and a random time period: $t_{confusion}$.
2. Start entering the confusion phase at $t_{challenge}$, and completely turn off the emitter after $t_{confusion}$.
3. Measure how much the real output deviates from the estimation by deriving the sum of squared residuals:

   $$g(t) = \frac{1}{T}\int_{\tau=t-T}^{t} z^2(\tau)d\tau \tag{2}$$

   where T is a preset time interval whose duration is equal to the transition time of the victim, and t is the instant the victim emitter completely turns off.
4. Alert when $g(t)$ exceeds a preset alarm threshold.

3 Attack Model

In our attack model, attackers attempt to spoof the regular channel through the same physical media sensed by the victim sensor, while the defender tries to detect or incapacitate those spoofing attempts.

3.1 Victim System

We assume the following victim system:

- The victim system is an active sensor system composed of an emitter and a receiver.
- Neither the emitter nor the receiver can be shielded from the external environment to ensure correct operation.
- Although the sensor output is analog, it will be sampled and quantized into digital form.
- The victim system may adopt a regular-channel spoofing detection system (e.g., PyCRA) to detect spoofing attacks against it.

The second assumption encompasses the case when the location of the measured entity relative to the sensor is not fixed, as is the case in radars or sonars. The third assumption is required to take the sampling rate into consideration. In purely analog systems, the equivalent sampling rate is infinite, and comparing the sampling rates of the attacker and the victim system becomes pointless. However, in modern cyber-physical systems, analog-digital systems are dominant because implementing complex functions is much more difficult without the aid of digital processors.

3.2 Capabilities of Attackers

Based on the characteristics of the victim system, we assume that attackers have the following capabilities:

Attackers are capable of *transparently* receiving and transmitting the physical signal from/to the victim's emitter/receiver. Note that *transparency* excludes applications where shielding can be applied, as assumed in the victim system model. When both the emitter and the receiver are properly shielded from the external environment, attackers cannot receive or inject signals.

We assume that the attacker has more resources than the victim sensor, as in many other attack models. The capability of a sensor is often limited owing to production and maintenance cost. However, the attacker does not have such limitations.

This assumption stems from a fundamental asymmetry between attackers and defenders. Because sensors are generally produced in large quantities, sensor manufacturers cannot adopt expensive microcontrollers or sensing structures irrelevant to the sensors' measuring capability. However, attackers can invest all their resources into the fabrication of one advanced attacking device. For example, once the sampling rate is fast enough for measurement, increasing the sensor's maximum sampling rate will be a difficult decision for sensor manufacturers; meanwhile, attackers can implement an attacking device with a much higher sampling rate.

In addition to these two assumptions, we also directly adopt assumptions A1–A4 from PyCRA [27, § 2.4], which include the following:

- Non-invasiveness: Attackers do not have direct access to the sensor hardware.
- Trusted measured entity: The physical entity to be measured by the sensor is trusted and incapable of being compromised.
- Physical delay: Adversaries require physical hardware with inherent physical delays.
- Computation delay: Adversaries may have superior computing power.

The second assumption means that the attacker cannot fabricate the measured entity itself. Taking an example of fire alarm system, we do not assume that the attacker starts an actual fire to spoof the system. In essence, we make the exact same assumptions as PyCRA.

4 Security Analysis of PyCRA

Even though PyCRA was proposed as a generally applicable solution to detect analog sensor spoofing, we point out that it has a fundamental vulnerability. Using this vulnerability, attackers with the profile introduced in Section 3 can spoof active sensors avoiding PyCRA's detection mechanisms. In this section, we first analyze its fundamental vulnerability, and show how it can be used to break all defensive measures of PyCRA. We also derive the sufficient condition for detection avoidance.

4.1 Vulnerability: The Sampling Race

In typical challenge-response authentication, prover and verifier share a secret, and the prover demonstrates to the verifier that it knows the secret, normally using nonce and cryptographic primitives such as encryption and digital signatures. If the nonce is not fresh or the cryptographic primitives are not secure, the challenge-response mechanism is vulnerable to replay or spoofing attacks.

However, in PyCRA, the only private information shared between the emitter and the receiver is the *timing*[2] of the emitter signal level changes (i.e., LOW, confusion phase, and HIGH).If this information is leaked, the attacker can defeat the challenge-response authentication. The fundamental assumption behind PyCRA is that the attacker's hardware has a non-zero physical delay due to its dynamic characteristics [15, pp. 25-31]. This physical delay is the time period between the time to sense the victim's challenge and the time to react to it. However, the victim system also has a delay, the duration between the time to emit its challenge and the time to recognize it. Therefore, the victim system cannot notice the spoofing signal, if an attacker can react faster than the victim can notice. In other words, the delays in both systems are in competition. For the victim to permanently win the competition, key factors are if there exists nontrivial[3] lower bound in the physical delay in a certain active sensor application and if the victim achieves more fine-grained time precision than it. This lower bound of physical delay is discussed in PyCRA [27, § 5.6], but it was limited to a specific type of active sensor, i.e. magnetic encoder. The core issue in determining the effectiveness of PyCRA for generalizable active sensor spoofing detection was neither discussed nor considered.

In analog-digital systems, analog sensing outputs must be converted to digital, which requires sampling of the analog information. If attackers have a sufficiently faster sampling rate than the victim, and the victim's time precision is insufficient to cover the minimal physical delay of the attacker, they can win the above competition. Moreover, because PyCRA's authentication is based on the sudden drop of signal levels in the silent phase, attackers can easily sense the start of the falling edge in challenges issued by the victim and react before the signal level even reaches the LOW state.

4.2 Attacks on the Simple Detector and the Confusion Phase

Figure 2a and 2b show two different cases of the victim and spoofing signals under PyCRA authentication. Figure 2a shows how PyCRA detects spoofing attempts. In this case, the physical delay of the attacker is larger than the victim's sampling interval. The attacker notices the sudden drop of the signal, and attempts to react to it. However, she fails to bypass the authentication, because the victim has already noticed the existence of the spoofing signal before the attacker can react to the challenge. In contrast, in Figure 2b, the attacker's physical

[2] Because these timings depend on the channel between the emitter and the receiver, PyCRA is applicable only to active sensors with fixed channels, where such timings become predictable (e.g. magnetic encoders). However, this does not hold for active sensors for ranging: radars, lidars, and sonars. In such applications the channel is basically assumed to be flexible.

[3] The existence of nontrivial lower bound in physical delay is essential for establishing *practical* defense. As any change in signal cannot be propagated faster than speed of light, trivial lower bounds always exist. However, such trivial lower bounds will necessitate the defender *impractical* time precision. As mentioned in Section 3.2, adoption of sampling rate much higher than the one required for measurement itself is a hard decision for sensor manufacturers.

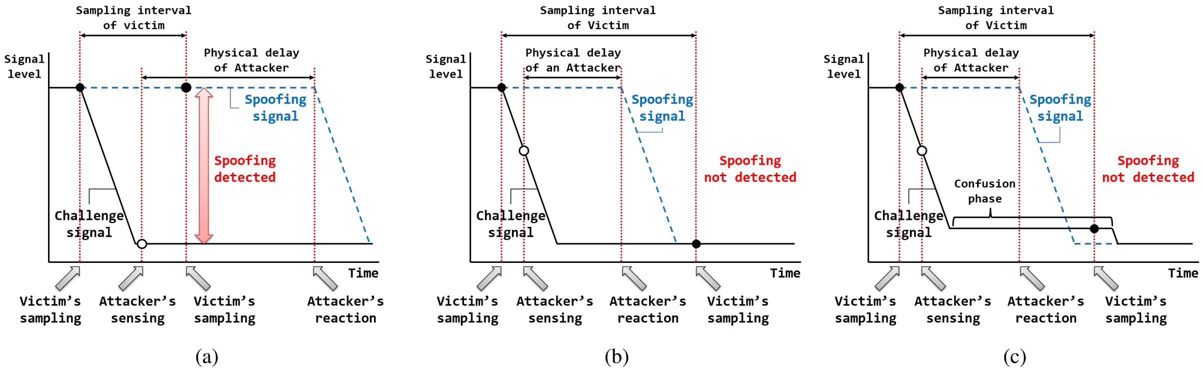

Figure 2: Illustration of PyCRA challenge under various physical delays of the attacker. (a) With long physical delays of the attacker, spoofing is detected by PyCRA; (b) Detection fails when the attacker's physical delay is sufficiently small; (c) Spoofing against the victim with the confusion phase. The attacker's sampling moments are marked with white dots, and the sampling moments of the victim's receiver are marked with black dots.

delay is shorter than the victim's sampling interval, and the attacker can detect the falling edge of the challenge as in the previous case. In the second case, however, the attacker successfully bypasses the authentication, because she can react to the challenge before the victim's next sampling moment. This result also shows that the confusion phase of PyCRA is completely useless under spoofing attacks by attackers who can react faster than the victim, as illustrated in Figure 2c.

Based on the diagrams depicted on Figure 2, we can derive the condition for deceiving the simple detector. Figure 3 shows the worst case (from the attacker's point of view) when attacking the simple detector. In this worst case, the attacker samples just before the victim challenges. As a result, it takes a full sampling period, T_A for the attacker to recognize that the challenge has started. Therefore, to successfully bypass the authentication the attacker should meet the condition below.

$$T_A + t_{f,A} + t_{p,A} \leq T_V \tag{3}$$

Note that, this also applies to the confusion phase. The only difference compared to the case of the simple detector is that the attacker turns off its emitter after the victim enters the confusion phase, not the silent phase.

4.3 Attacks on the χ^2 Detector

PyCRA introduces χ^2 detector to defend against spoofing attacks with shorter physical delays than that of the victim sensor. As discussed in Section 2.4, χ^2 detector adopts the mathematically modeled dynamic characteristics of the victim, to detect spoofing attacks even when the attacker's signal completely vanishes before the end of the victim transition time. More precisely, the adopted mathematical approach is to calculate the residual during the transient state of the sensor. However, it can also be useless in digital systems with finite sampling rates, particularly for applications where the transition time of the attacker can be reduced below the victim's sampling interval. It is important to note that the key transition time here is not that of the victim but that of the attacker, who willingly invests all her available resources to reduce this amount of time.

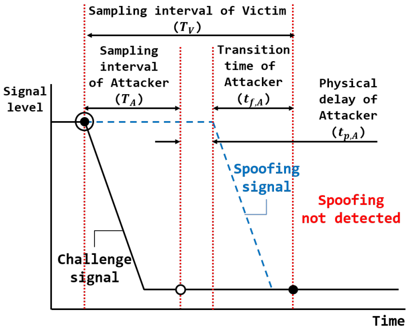
Sampling interval of Victim
(T_V)
Sampling interval of Attacker
(T_A)
Transition time of Attacker
($t_{f,A}$)
Physical delay of Attacker
($t_{p,A}$)
Signal level
Spoofing signal
Spoofing not detected
Challenge signal
Time

Figure 3: Worst case on attacking a simple detector. Attacker samples just before the victim emitter signal drops.

Here, we show this in greater detail. For digital systems, Eq. (1) and (2) become Eq. (4) and (5), respectively:

$$x[n+1] = f(x[n]) + w[n] \tag{4a}$$

$$y[n] = h(x[n]) + v[n] \tag{4b}$$

$$g[n] = \frac{1}{N} \sum_{m=n-N+1}^{n} z^2[m] \tag{5}$$

where the residual is $z[n] = y[n] - \hat{y}[n]$, and N denotes the number of victim's sampling intervals, approximately equal to the victim's transition time.

When we assume that the victim's signal completely turns off at $n = n_0$, $g[n_0]$ becomes close to zero (i.e.,

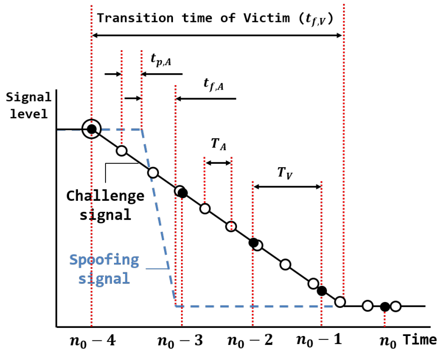

Figure 4: Illustration of the χ^2 detector bypass. Symbols are directly borrowed from Figure 3.

$g[n_0] \approx 0$) under faster attackers, as depicted in Figure 4. This is because $y - \hat{y} = z \approx 0$ even on the first sample of the transition time, $n = n_0 - N + 1$. Needless to say, the same applies to all following samples, i.e., $z^2[i] \approx 0$ $(i = n_0 - N + 1, n_0 - N + 2, \cdots, n_0)$. We note that, the same inequality (3) also applies to the χ^2 detector bypass as in bypassing simple detector.[4]

4.4 Summary

The fundamental idea of PyCRA is reasonable and simple, but it ignores critical problems: whether the lower bound of the physical delay can be universally determined for every active sensor, whether such lower bound is in practical range, and whether the victim can achieve the corresponding time precision. When the two sampling intervals satisfy inequality (3), attackers can always bypass all PyCRA detection mechanisms. Furthermore, the χ^2 detector designed for the case when the attacker's physical delay is less than that of the victim becomes useless when the above inequality holds.

5 Experiments

In the previous section, it was shown that PyCRA cannot detect attackers who react faster than the victim's sampling interval. What remains to be proved is if, in practice, it is possible to reduce the reaction delay of the attacker to a level much lower than the sampling interval of most high-performance microcontrollers. In this section, we experimentally show that, with only low cost devices, the delay can be reduced to much less than 5 μs, the sampling period corresponding to the 200 kHz sampling rate of most high-end microcontrollers. We also note that the 200 kHz of sampling rate was not arbitrarily chosen, but taken from the original PyCRA paper [27, § 5.6]. This experiment, along with the aforementioned theoretical limitations, shows that PyCRA is vulnerable even to low cost attacks.

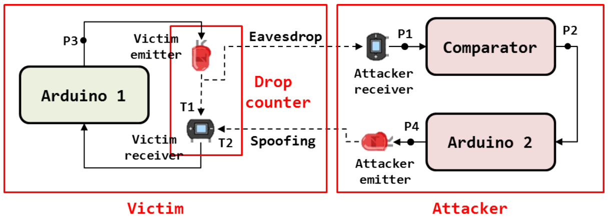

Figure 5: Experimental setup. T1 and T2 indicate the moments when the victim and attacker emitters are turned off. The time difference, T2 − T1 is repeatedly derived to measure the attacker's delay. Dotted and solid lines indicate optical and wired channels, respectively.

5.1 Design

We implemented an elementary infrared (IR) PyCRA evading system, which attempts to avoid detection by the victim. As a victim sensor, an IR drop counter installed on a commercial infusion pump [5] was used. Drop counters are used to count the number of droplets passing through the emitter and the receiver. We note that a real-world active sensor was used here, instead of a custom-built one, and that the victim sensor was not modified at all during the experiment.

As an attacking emitter and receiver, we used an IR light emitting diode (LED) and an IR phototransistor, respectively, whose targeted wavelength is from 850 nm to 950 nm. Two Arduino UNO [1] boards were used to implement the victim and attacker processors, and for the attacker's side only, a simple self-implemented crude comparator was built, with an operational amplifier.

Figure 5 shows the overall experimental setup. Arduino 1 (together with the drop counter) plays the role of a victim, while Arduino 2 (on the right side) is used as an attacker. The victim drop counter is directly connected to Arduino 1, and both the LED and the phototransistor are connected to Arduino 2. Arduino 1 was programmed to turn off its emitter at random instants, and Arduino 2 was programmed to react to the challenge as fast as possible, to bypass authentication. In this experiment, the attacker receives the IR signal from the emitter via optical channel. With an oscilloscope, we measured the time difference between the turn-off instants of each emitter, while applying various attacker-side circuitries, to minimize the attacker delay.

[4] In practice, the condition for bypassing χ^2 detectors will be looser than inequality (3). As no modeling can completely fit the reality, a χ^2 detector have a *margin*—an alarm threshold—to detect spoofing attempts as remarked in Section 2.4. This margin does not only prevent the sensor from frequent false alarms, but helps the attacker to be undetected by ignoring small deviations from the modeling.

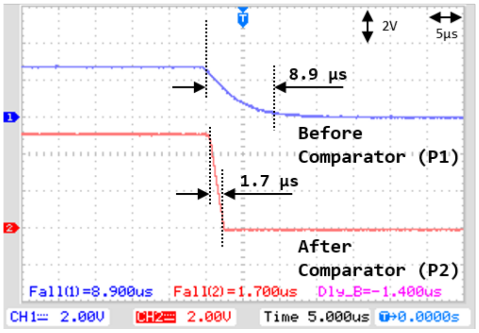

Figure 6: Comparison between the comparator input and output (P1 and P2 in Figure 5). The upper graph is the transition signal of the victim emitter (i.e., the comparator input) and the other is the comparator output.

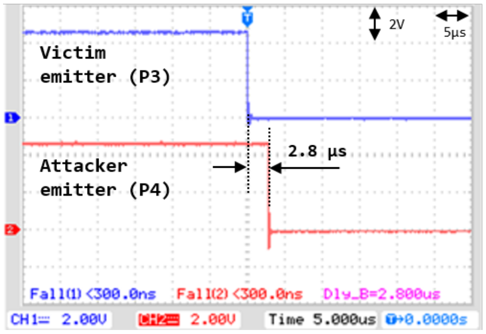

Figure 7: Total attacker's delay versus the victim challenge. The upper graph is the victim challenge signal, and the other is the attacker emitter input (measured at P3 and P4 in Figure 5) respectively.

5.2 Minimizing the Attacker Delay

In this experiment, the primary goal is to reduce the attacker delay, which is measured by deriving the time difference, T2 − T1 in Figure 5. We now describe the process of minimizing this time difference, to show that the attacker delay can be reduced to less than 5 μs, which is equivalent to the sampling rate over 200 kHz.

Reducing the Transition Time We first attempted to minimize the phototransistor delay. Phototransistors are generally used with a resistor to convert the photocurrent into voltage, thus constituting a photodetector. According to various application notes on phototransistors [29, 18], the resistance used with the phototransistor and the agility of the photodetector are negatively correlated. Therefore, we reduced the resistance as much as possible to 100 Ω in order to maximize the agility of the attacker receiver. Note that this resistance cannot be reduced infinitely, because it is positively correlated to the resulting photodetector sensitivity.

As discussed in Section 4.3, the attacker does not have to wait until the victim's emitter completely turns off. Instead, for faster response, the attacker can turn off her emitter, whenever the victim's emitter signal drops below a certain threshold. Therefore, the next step was to attach a comparator with a carefully chosen threshold in front of the attacker's processor, to further reduce the transition time of the attacker. As a result of the combined effect of these two design decisions, the transition time was drastically reduced, as shown in Figure 6.

Reducing the Processor-side Delay We first used one of analog pins of the Arduino with a sampling rate of 10 kHz, to react to the victim challenge. However, it was too slow to bring the reaction time under 5 μs. Contrary to the victim sensor, which must precisely measure the emitter signal to perform its sensing function, the attacker only needs to determine whether the victim is or is not challenging at a given time. Therefore, we changed from an analog pin to a digital input pin, and to further reduce the attacker delay, we used the digitalReadFast and digitalWriteFast libraries [8], which directly use port commands, instead of the Arduino APIs.

5.3 Results

As an intermediate result, we were able to reduce the transition time to under 10 μs, by setting the photodetector resistance to the optimal value of 100 Ω. When the comparator was added, the delay was again reduced to about 1.7 μs, as shown in Figure 6. We note that this value can be reduced even further, if a dedicated comparator IC is adopted instead of our crude comparator.

After all manipulations mentioned above, the total attacker delay was reduced to 2.8 μs. Figure 7 shows the time difference between the start of the victim's challenge and the attacker's reaction. As shown, the delay between two graphs is about 2.8 μs, which is approximately equivalent to the sampling rate of 358 kHz.

So far, we have experimentally shown that it is possible to react faster than 200 kHz, even with several low-cost devices. 200 kHz of sampling rate is possibly not an absolute criterion for determining the effectiveness of a PyCRA detector. However, our experiment is a conceptual one showing that even a 200 kHz of sampling rate can readily be bypassed with some crude equipment. Serious attackers will be equipped with much more advanced devices. In summary, as long as the victim system has a finite sampling rate, this race will never end until the sampling interval of the victim becomes much shorter than the best transition time achievable with contemporary technologies.

6 Discussion

So far, we have shown the limitations of PyCRA as a generalizable active sensor spoofing detection scheme both theoretically and empirically. Because no active sensor defense scheme is both robust and generalizable to our knowledge, we list several alternative defense approaches. However, none of them are both robust and generalizable.

6.1 Shielding

When applicable, shielding is the simplest and robust defense approach. Once the actuator, the receiver, and the measured entity are tightly shielded so that the measured physical media cannot be penetrated from the outside, the active sensor will no longer be affected by external stimuli. This is guaranteed as long as the shielding takes effect, and the shielded active sensor is thus immune to attack attempts outside the shielding. Drop counters and magnetic encoders are good examples of active sensors that can be shielded. Indeed, shielding is already being applied in these cases [24, 2], with the caveat, however, that it is not designed for security against spoofing attacks. This type of defense can be categorized as *Signal Integrity* type defenses.

However, despite its strong resilience against external spoofers, shielding is not generalizable. As mentioned above, we cannot shield when measured entity is not fixed: radars, lidars, and sonars. Active sensors spanning large volumes of space such as optical beam smoke detectors are also difficult to be shielded properly, because a massive area of shielding would be required to enclose the emitter, receiver, and measured entity.

6.2 Redundancy and Sensor Fusion

A sensing and actuation system might improve its resilience against sensor spoofing attacks by employing redundancy. This includes adopting not only an identical type of sensors, but also multiple types of sensors (e.g., camera and radar for detecting road objects). To accommodate the acquired multiple sensor output streams, sensor fusion techniques can be used, further enhancing the overall system resistance. Although there are not many works suggesting sensor fusion as a defensive mechanism against sensor spoofing [23], many previous works have addressed the resulting improvements in precision [25, 19, 22] or reliability [3, 17] of the overall sensor system, which might also improve security against spoofing attacks.

Despite its advantages, this line of approach cannot guarantee its robustness, because an attacker might still compromise the sensing system by simultaneously conducting sensor spoofing attacks against multiple sensors.

7 Related Work

Making secure sensors is not easy, because most sensors have limited computational resources and simple functionalities, directed only to measurement. Consequently, only a few existing works are directly related to the topic of sensor security, and most of them concentrate on establishing non-ideal defenses. In this section, we classify the existing works into three broad categories: sensor spoofing attacks, defenses against them, and sensor reliability enhancement.

Sensor spoofing attacks: As mentioned in Section 2.2, there are three attack vectors for sensor spoofing. Shoukry et al. demonstrated a sensor spoofing attack through the regular channel of an automotive magnetic encoder [26]. They put a magnetic actuator in front of an anti-lock braking system sensor, whose base is a magnetic encoder, and falsified the wheel speed of a vehicle. Son et al. showed that a commercial drone can be rendered uncontrollable by a side-channel spoofing attack against MEMS gyroscopes [28]. They first found the resonant frequencies of the MEMS structure, and showed that several MEMS gyroscopes behave abnormally under acoustic noise at their resonant frequencies. Using this phenomenon, they successfully forced the victim drone to drop to the ground. Finally, Foo Kune et al. attacked a wired transmission channel connecting an analog sensor and its backend system [14]. They successfully injected fake sensor outputs by generating intentional EMI within the wire. In addition to these, sensor spoofing attacks to bypass biometric authentication schemes such as fingerprint recognition [20, 10], facial recognition [9], and automatic speaker verification [11] have also been proposed.

Defenses for sensor spoofing attacks: The work most directly related to defenses is PyCRA [27], which has already been deeply discussed (Section 2.4). Additionally, there have been several works based on redundancy using multiple sensors or additional resources. Park et al. proposed a detection algorithm based on sensor fusion, which detects malfunctioning sensors on an unmanned ground vehicle [23]. Montgomery et al. demonstrated a global positioning system (GPS) spoofing detection method with dual antennas on a GPS receiver, using antenna multiplexing [21]. An upper bound on the detectable number of corrupted sensors in multisensor systems was characterized for cyber-physical control systems by Fawzi et al. [13]. Shielding has been also mentioned as a defense in three of the aforementioned attack papers [26, 28, 14].

Reliability enhancement for sensor systems: A number of works have proposed the use of multiple sensors to enhance the precision and reliability of sensor systems, although not for security purposes. Caron et al. [3] developed an algorithm fusing GPS and inertial measurement unit (IMU) using a Kalman filter, to enhance the reliability of location based applications. Nützi et al. [22] and Martinelli [19] fused IMU and vision sensors for accurate pose estimation. For reliable pedestrian navigation, sensing data from an inertial sensor, an image sensor, and a barometer were utilized [12, 16]. A number of works also exist to improve the reliability and precision of sensor networks [25, 17].

8 Conclusion

This paper focuses on breaking PyCRA, the only authentication mechanism to detect spoofing attacks against active sensors, claimed to be robust and generalizable. We show in theory as well as in practice, PyCRA is insecure. In theory, we derive a sufficient condition for the attacker to avoid detection. We also show experimentally that PyCRA can be easily bypassed by an attacker, even with a low-cost microcontroller and crude additional circuitries. Existence of our attack shows that design of robust and generalizable defense mechanism for active sensors remains as an open problem.

Acknowledgment

This work was supported by Institute for Information & communications Technology Promotion (IITP) grant funded by the Korea government (MSIP) (No.B07171601160003003, Development of information leakage prevention and ID management for secure drone services)

References

[1] Arduino UNO Specification. `https://www.arduino.cc/en/Main/ArduinoBoardUno`.

[2] Avtron Industrial Automation. Eliminating Magnetic Encoder Interference from Brakes & Motors. `http://www.nidec-avtron.com/encoders/documents/white-papers/wp-elim-mtr-shft-crnts.pdf`. [Online; accessed 9-May-2016].

[3] Caron, F., Duflos, E., Pomorski, D., and Vanheeghe, P. GPS/IMU data fusion using multisensor Kalman filtering: introduction of contextual aspects. *Information fusion 7*, 2 (2006).

[4] Castro, S., Dean, R., Roth, G., Flowers, G. T., and Grantham, B. Influence of Acoustic Noise on the Dynamic Performance of MEMS Gyroscopes. In *ASME International Mechanical Engineering Congress and Exposition* (2007).

[5] Changsha JianYuan Medical Technology Co., Ltd. Peristaltic Pump Infusion Pump (JSB-1200). `http://www.made-in-china.com/showroom/jympumplisha/product-detailEXAmZfvcHnWd/China-Peristaltic-Pump-Infusion-Pump-JSB-1200-.html`, 2015. [Online; accessed 9-May-2016].

[6] Dean, R., Castro, S., Flowers, G., Roth, G., Ahmed, A., Hodel, A., Grantham, B., Bittle, D., and Brunsch, J. A Characterization of the Performance of a MEMS Gyroscope in Acoustically Harsh Environments. *IEEE Transactions on Industrial Electronics 58*, 7 (2011).

[7] Dean, R., Flowers, G., Hodel, A., Roth, G., Castro, S., Zhou, R., Moreira, A., Ahmed, A., Rifki, R., Grantham, B., Bittle, D., and Brunsch, J. On the Degradation of MEMS Gyroscope Performance in the Presence of High Power Acoustic Noise. In *IEEE International Symposium on Industrial Electronics* (2007).

[8] digitalWriteFast API in Arduino Library. `https://code.google.com/p/digitalwritefast/`.

[9] Duc, N. M., and Minh, B. Q. Your face is NOT your password - Face Authentication Bypassing Lenovo–Asus–Toshiba. *Black Hat Briefings* (2009).

[10] Espinoza, M., and Champod, C. Risk evaluation for spoofing against a sensor supplied with liveness detection. *Forensic science international 204*, 1 (2011).

[11] Evans, N., Kinnunen, T., and Yamagishi, J. Spoofing and countermeasures for automatic speaker verification. In *INTERSPEECH* (2013).

[12] Fallon, M. F., Johannsson, H., Brookshire, J., Teller, S., and Leonard, J. J. Sensor Fusion for Flexible Human-Portable Building-Scale Mapping. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on* (2012).

[13] Fawzi, H., Tabuada, P., and Diggavi, S. Secure Estimation and Control for Cyber-Physical Systems Under Adversarial Attacks. *IEEE Transactions on Automatic Control 59*, 6 (2014).

[14] Foo Kune, D., Backes, J., Clark, S., Kramer, D., Reynolds, M., Fu, K., Kim, Y., and Xu, W. Ghost Talk: Mitigating EMI Signal Injection Attacks against Analog Sensors. In *IEEE Symposium on Security and Privacy* (2013).

[15] Fraden, J. *Handbook of Modern Sensors: Physics, Designs, and Applications*. Springer Science & Business Media, 2004.

[16] Gadeke, T., Schmid, J., Zahnlecker, M., Stork, W., and Muller-Glaser, K. D. Smartphone Pedestrian Navigation by Foot-IMU Sensor Fusion. In *Ubiquitous Positioning, Indoor Navigation, and Location Based Service (UPINLBS), 2012* (2012).

[17] Kreibich, O., Neuzil, J., and Smid, R. Quality-Based Multiple-Sensor Fusion in an Industrial Wireless Sensor Network for MCM. *IEEE Transactions on Industrial Electronics 61*, 9 (2014).

[18] Marktech–Optoelectronic. Application notes—Photo Sensor Application Notes. `http://www.marktechopto.com/photo-sensor-application-notes.cfm`. [Online; accessed 9-May-2016].

[19] Martinelli, A. Vision and IMU data Fusion: Closed-Form Solutions for Attitude, Speed, Absolute Scale, and Bias Determination. *IEEE Transactions on Robotics 28*, 1 (2012).

[20] Matsumoto, T., Matsumoto, H., Yamada, K., and Hoshino, S. Impact of Artificial “Gummy” Fingers on Fingerprint Systems. In *Electronic Imaging* (2002).

[21] Montgomery, P. Y., Humphreys, T. E., and Ledvina, B. M. Receiver-Autonomous Spoofing Detection: Experimental Results of a Multi-antenna Receiver Defense Against a Portable Civil GPS Spoofer. In *Proceedings of the ION International Technical Meeting* (2009).

[22] Nützi, G., Weiss, S., Scaramuzza, D., and Siegwart, R. Fusion of IMU and Vision for Absolute Scale Estimation in Monocular SLAM. *Journal of Intelligent & Robotics Systems 61*, 1-4 (2011).

[23] Park, J., Ivanov, R., Weimer, J., Pajic, M., and Lee, I. Sensor Attack Detection in the Presence of Transient Faults. In *Proceedings of the ACM/IEEE Sixth International Conference on Cyber-Physical Systems* (2015).

[24] PASCO Capstone. PASPORT High Accuracy Drop Counter. `https://www.pasco.com/prodCatalog/PS/PS-2117_pasport-high-accuracy-drop-counter/index.cfm`. [Online; accessed 9-May-2016].

[25] Sekkas, O., Hadjiefthymiades, S., and Zervas, E. A Multi-level Data Fusion Approach for Early Fire Detection. In *2nd International Conference on Intelligent Networking and Collaborative Systems* (2010), pp. 479–483.

[26] Shoukry, Y., Martin, P., Tabuada, P., and Srivastava, M. Non-invasive Spoofing Attacks for Anti-lock Braking Systems. In *Cryptographic Hardware and Embedded Systems* (2013).

[27] Shoukry, Y., Martin, P., Yona, Y., Diggavi, S., and Srivastava, M. PyCRA: Physical Challenge-Response Authentication For Active Sensors Under Spoofing Attacks. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security* (2015).

[28] 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* (2015).

[29] Van, N. T. CEL APPLICATION NOTE—Phototransistor Switching Time Analysis. Tech. rep., California Eastern Laboratories, 2009. [Online; accessed 9-May-2016].