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

Yasser Shoukry, Paul Martin, Yair Yona, Suhas Diggavi, and Mani Srivastava

I. Summary

Briefly, PyCRA implicitly requires the victim sensor to select a sample rate that is higher than a threshold that depends on the non-zero physical delays present in the attacker. This threshold does not depend on the sampling rate at the attacker, and therefore there is no *"race"* between the two sample rates. Also, some of the key conclusions in your paper are based on using the simple detector which the PyCRA paper itself had recognized as having weaknesses which were addressed by the more sophisticated Chi-squared detector. Indeed, the PyCRA paper had presented the simple detector not as an alternative solution but as an intermediate step to motivate the development of the Chi-squared detector. Furthermore, your paper portrays PyCRA as requiring a sample interval at the victim sensor that exceeds the physical delays at the attacker, which is contrary to PyCRA's actual functioning.

II. Detailed Explanation

1) The basic idea of PyCRA is to give the victim sensor an asymmetric advantage over the attacker by exploiting non-zero physical delays in the attacker's system. Indeed, if the victim sensor does not sample fast enough relative to this non-zero delay at the attacker, then the victim loses this asymmetric advantage and gives the attacker a better opportunity to win the "race". Put differently, if both the attacker and the victim increase their sampling frequencies, then there exists a threshold (dictated by the physics) on the sampling frequency after which the victim will always win regardless of how fast the attacker is. Therefore, because of this fundamental physical delay, the race is actually between the victim's sampling rate and the attacker's *physical delay* and not the attacker's *sampling rate*. Moreover, because of this non-zero delay, a victim sensor with high sampling frequency will always win the race, i.e., if PyCRA operates on a sampling rate that is comparable to the attacker's physical delay, the attacker, regardless of how fast it samples, will be detected. Moreover, sampling rate beyond that required by the bandwidth of the signal dynamics is sufficient.

2) Contrary to statements in your paper, PyCRA never assumes low sampling rates. Indeed, as noted in the discussion section of our paper, in order to increase its security, PyCRA requires the victim sensor to increase its sampling rate, which we believe is a reasonable price to pay for security. Quoting from the PyCRA paper:
"Higher sampling rate comes at the cost of increased power consumption, but this is perhaps a reasonable price to pay for security."

3) Section 4.2 of your paper is devoted to analyzing the vulnerability of the simple detector. First, we would like to stress the fact that the simple detector itself is never presented in PyCRA paper as a standalone algorithm but just as a motivation to the more sophisticated Chi-squared detector. Putting this fact aside, the analysis in Section 4.2 of your paper (along with Figure 2 and 3) is based on the assumption that PyCRA samples after the transients of the sensor diminish. We would like to stress that such assumption on the sampling rate of the simple detector was never made in PyCRA.

4) Similarly to the previous issue, your paper introduces new assumptions related to how the confusion phase is functioning. Quoting from Section 2.4 in your paper:
"We note here that the confusion phase concept assumes the attacker will not stop spoofing during the confusion phase."

   We would like to stress that such assumption was never proposed by PyCRA. On the contrary, as explicitly mentioned in our paper, PyCRA assumes that the attacker is trying to detect the challenge time in order to conceal its signal and remain stealthy. The main purpose of the confusion phase is to provide an immutable non-zero probability of the

Yasser Shoukry, Paul Martin, Yair Yona, Suhas Diggavi, and Mani Srivastava are with the UCLA Electrical Engineering, Los Angeles, CA, {yshoukry,pdmartin,yair99,suhasdiggavi,mbs}@ucla.edu

attacker to miss such detection. In fact, if the confusion phase was designed carefully, the simple attack mechanism that is proposed in the experimental section of your paper would have been easily captured. Unfortunately, and because of the page limit size of the original PyCRA paper, we had to omit a lot of details about the confusion phase. However, these details are given in the extended version of PyCRA published on the ArXiv (http://arxiv.org/pdf/1605.02062v2.pdf).

5) Quoting from Section 4.1 in your paper:
"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."

We would like to comment that the scenario discussed in this statement is handled directly by the Chi-Squared detector which does not wait until the signal level reaches low. On the contrary, to the quotation above, it continuously monitors the transient of the signal as it goes from a HIGH to LOW state and checks against the physical model. Any deviation from the model within this duration is attributed to the existence of an attacker. As an experimental example, we show in the extended version of PyCRA (published on the ArXiv at http://arxiv.org/pdf/1605.02062v2.pdf) a case where the victim sensor has a $15\times$ slower delay compared to the attacker and therefore the attacker reacts before the signal reaches the LOW state. Thanks to the Chi-Square detector, such a fast attack can still be detected (given a sampling rate that is comparable to the physical delay of the attacker, again regardless of how fast the attacker's sampling rate is).

6) Quoting from Section 4.4 in your paper:
"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."

We would like to comment that some universal lower bound can be always calculated using basic laws of physics. While there are many sources of delay on the attacker side, we based our argument in PyCRA on the delay that occurs at the actuator electrical components. Recall that basic laws of physics relates how voltage and current change overtime inside electrical circuits with the characteristics of the components used to build these circuits (e.g. material type, coil length, ...). The characteristics of the state-of-the-art components can be used as a lower bound for the delay on the attacker's circuit. For example, current reported state-of-the-art in low-dimension, high Q-factor magnetic actuators results into a physical delay of $200\mu s$.

7) Quoting from Section 5.3 in your paper:
"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."

We do not agree with the claim that the *"race will never end"*, a claim seems to be based on not sufficiently distinguishing between physical delay and sampling interval. As discussed above in comment #1, the non-zero delay on the attacker reaction/actuation is fundamental and will never be zero. This physical delay at the attacker does not diminish with higher sampling rate at the attacker, i.e., regardless of how fast the attacker increases his sampling rate, it can not overcome this non-zero delay. Therefore, to end this race, the victim needs to sample relatively fast to this non-zero delay.

8) We humbly disagree with the conclusions drawn from the experimental setup presented in your paper. First, in your experiments, a sampling rate of 200 KHz is chosen based on the experiments shown in PyCRA. As with any physical phenomena, one needs to pick a sampling rate that suits its characteristics (e.g., how fast such physical phenomena changes with time). A choice of sampling rate that is designed for changes in magnetic waves will be, indeed, insufficient to monitor changes in electromagnetic waves (LED in your case). Portraying that 200 KHz is a general recommendation done by PyCRA for any type of sensor is misleading and was never stated in PyCRA's paper. As explicitly mentioned in our paper the choice of 200 KHz is based on the survey we did over the state-of-art, commercially available magnetic actuators. This survey (explicitly mentioned in PyCRA's paper) shows that a physical delay of at least $200\mu s$ is presented in such magnetic actuators and therefore the choice of 200 KHz was made. A different sensor based on different physical phenomena will indeed require different sampling rate. As supported

by your experiments, a physical delay of $2.8\mu s$ do exist in the attacker's actuators. Therefore, if the experiments shown in your paper were correctly designed, a sampling frequency of 500KHz—which is far below the state-of-art, off-the-shelf, analog-to-digital converters—should have been used.