Non-invasive Spoofing Attacks
for Anti-lock Braking Systems

Yasser Shoukry[1,2], Paul Martin[2], Paulo Tabuada[1], and Mani Srivastava[2]

[1] Cyber-Physical Systems Laboratory, Dept. of Electrical Engineering
University of California at Los Angeles
`http://www.cyphylab.ee.ucla.edu`
[2] Networked and Embedded Systems Lab., Dept. of Electrical Engineering
University of California at Los Angeles
{yshoukry,pdmartin,tabuada,mbs}@ucla.edu
`http://www.nesl.ee.ucla.edu`

Abstract. This work exposes a largely unexplored vector of physical-layer attacks with demonstrated consequences in automobiles. By modifying the physical environment around analog sensors such as Antilock Braking Systems (ABS), we exploit weaknesses in wheel speed sensors so that a malicious attacker can inject arbitrary measurements to the ABS computer which in turn can cause life-threatening situations. In this paper, we describe the development of a prototype ABS spoofer to enable such attacks and the potential consequences of remaining vulnerable to these attacks. The class of sensors sensitive to these attacks depends on the physics of the sensors themselves. ABS relies on magnetic–based wheel speed sensors which are exposed to an external attacker from underneath the body of a vehicle. By placing a thin electromagnetic actuator near the ABS wheel speed sensors, we demonstrate one way in which an attacker can inject magnetic fields to both cancel the true measured signal and inject a malicious signal, thus spoofing the measured wheel speeds. The mounted attack is of a non-invasive nature, requiring no tampering with ABS hardware and making it harder for failure and/or intrusion detection mechanisms to detect the existence of such an attack. This development explores two types of attacks: a disruptive, naive attack aimed to corrupt the measured wheel speed by overwhelming the original signal and a more advanced spoofing attack, designed to inject a counter-signal such that the braking system mistakenly reports a specific velocity. We evaluate the proposed ABS spoofer module using industrial ABS sensors and wheel speed decoders, concluding by outlining the implementation and lifetime considerations of an ABS spoofer with real hardware.

Keywords: Automotive embedded systems, Cyber-physical security, Non-invasive sensor attacks, Magnetic sensors.

1 Introduction

Increased coupling between embedded computing technologies and modern control systems has opened the door for developing many engineering systems with increasing complexity. In such systems, commonly termed *cyber-physical systems* or CPS,

information from the physical world is quantized and processed using digital electronic components, and decisions taken by these "cyber components" are then applied to the physical world. Unfortunately, this tight coupling between cyber components and the physical world oftentimes leads to systems where increased sophistication comes at the expense of increased vulnerability and security weaknesses. At the heart of secure cyber-physical systems is the notion that information collected from the physical world through sensors poses a significant vulnerability risk. Although, such information is exchanged between individual components of the CPS in an encrypted fashion, the coupling with the physical world leads to new security breaches that do not exist in the traditional cyber-security domain. Thus understanding how an attacker might modify and corrupt such information from the physical part of the system becomes of critical importance in assessing the dependability and security of these systems.

Moreover, successful attacks on the information collected from sensors in a feedback control system can be even more damaging compared to open-loop systems due to the active property of control systems, where the data collected from sensors are used to decide the next actions to be taken. It's unsurprising, then, that analyzing and detecting sensor spoofing attacks in the context of cyber-physical systems is a growing concern and the subject of many recent research endeavors [1–4].

Automotive vehicles continue to be one of the most complex cyber-physical systems to date, and, with many millions of people entrusting their lives to automobiles everyday, addressing security threats in automotive systems is undoubtedly a real concern. Security threats in automotive vehicles have been examined thoroughly in [5, 6], where the authors explore how an attacker can make use of external vehicle interfaces as well as internal networks to pose a threat on the vehicle control sub-systems. This work describes an additional mode of attack in the form of modifying sensor signals directly. Modern automotive vehicles are equipped with, on average, 70 sensors classified into 21 different types. Comparing this number to the mere 24 sensors seen on a typical vehicle ten years ago [7] shows just how dramatic the growth in number of sensors deployed in automotive vehicles has been and further illustrates the growing concern for sensor-level attacks.

Sensor-level attacks can be classified into invasive and non-invasive attacks [8]. Invasive attacks are those in which the attacker has to tamper with internal components of the system (e.g. internal circuitry and wiring of the sensor or changing software dealing with processing sensor measurements). The defining characteristic of these attacks is that some part of the system is physically altered. On the contrary, non-invasive attacks do not physically alter the components of a sensor but rather make use of the information gathered from the physical environment around the sensor to infer some information about the operation of the sensor and (remotely) inject a malicious signal. In many cases, invasive attacks can be easily detected with intelligent circuit designs and robust programming. Non-invasive attacks, however, can be much more difficult to detect—here, the system designer can no longer blindly trust the output of a sensor. In effect, the system designer can no longer trust the *physical environment* that is being monitored. Shielding a system from these more sophisticated attacks requires protection in kind.

In this paper we assume the role of an attacker, attempting to exploit non-invasive vulnerabilities in one important class of sensors found on modern vehicles–inductive magnetic field sensors used to control Anti-Lock Brake Systems (ABS). We demonstrate that attacks on even a small subset of car sensors can have very serious consequences in terms of safety. The rest of this paper is organized as follows. Section 2 introduces the operation of ABS sensors and discusses the different types of attacks that can be mounted on these sensors. Internal details of the developed ABS Hacker module are presented in Section 4. Evaluation of the proposed system through practical tests is presented in Section 5. Finally, we offer some concluding thoughts in Section 7.

2   Attacking ABS Sensors

Anti-Lock Braking Systems (ABS) have become a standard active safety technology in current vehicles. Because the friction force on car wheels during lock-up events considerably decreases, ABS is designed to prevent the wheels from locking when the brakes are applied. In order to avoid lock-up and achieve maximum adhesion between tires and road surface, ABS measures the speed of each individual wheel, sends this information to the electronic control unit (ECU) which compares each individual wheel speed versus the lateral car speed, and if a mismatch is found the ECU starts to decrease the brake torque to prevent wheel lock-up.

As a motivating example, we examine the sequence of a driver taking a turn and applying the brakes. Here, the ABS computer reads the individual speeds and applies a braking torque to make sure that no wheel is slipping, thus stabilizing the vehicle. If even a single wheel comes under attack, the malicious attacker can spoof the sensor such that a wheel that is actually slipping is perceived to be operating normally. Since these measurements are used for instantaneous decisions, the ABS computer will then apply an incorrect torque which can destabilize the vehicle. We will revisit this example again later in the text in order to demonstrate that such cases can occur.

The purpose of the work presented in this paper is to demonstrate a small electronic attacking module that is capable of changing the readings of an ABS sensor without tampering with the sensor hardware itself. By attaching this module to a position in close proximity to any typical ABS wheel sensor, the module will start to alter the physical environment around the wheel speed sensor in order to inject precisely the desired (and incorrect) speed. Although the work presented in this paper focuses on ABS sensors, the general developed concepts and proposed methodologies still apply to other types of wheel speed sensors used in other applications (for example motor encoders) and in fact any similar inductive sensor—ABS systems serve merely to illustrate the potential severity of these attacks.

2.1   Types of ABS Sensors

ABS systems use magnetic speed sensors to measure individual wheel speeds. Magnetic speed sensors are typically used because of their ability to accommodate harsh operating environments like dust and temperature while optical speed sensors could fail.

(a)   (b)
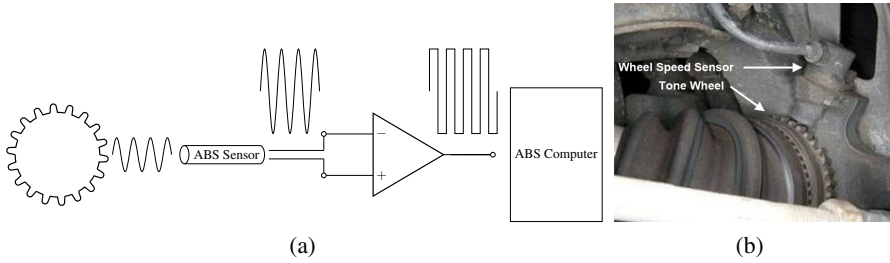
Fig. 1. (a) Basic speed sensor operation for ABS systems, (b) An exposed ABS speed sensor and tone ring

In this section, we will review different types of ABS sensors, describe sensor interfaces that provide encoded wheel speeds to the ECU, and discuss how an attacker can mount an attack in order to spoof the measurements of a magnetic wheel speed sensor.

ABS sensors found on today's vehicles come in two varieties—passive and active. Both types rely on the existence of a ferromagnetic toothed gear (also called the tone ring) rotating in front of a permanent magnet. Figure 1a shows the basic operation of magnetic wheel speed sensors. As a gear tooth of the rotating wheel passes in front of the magnet, the magnetic flux density (generated by the permanent magnet) is at a maximum. When the tooth moves away and an air gap is presented in front of the magnet, the magnetic flux density drops to its minimum value. The result is a time-varying magnetic flux with a variation rate that is proportional to the gear speed, the diameter of the tone ring, and the number of teeth on the ring. In a typical ABS setup, the triggering gear is located on the axle of the wheel.

Conventional ABS passive sensors (also called Variable Reluctance or VR sensors) are composed of a copper wire wrapped around a permanent magnet, forming a pick-up coil. The output of these sensors is a sinusoidal wave (shown in Figure 1a) whose frequency is proportional to the wheel speed (multiplied by number of teeth). The output of the passive ABS sensor then passes through a comparator circuit which produces the typical rotary encoder signal (a square wave where the frequency of the transitions is proportional to the speed of the rotary object).

The more advanced active ABS sensors rely on a cluster of three hall effect sensors separated in space. Each sensor in the cluster measures the time-varying magnetic flux of the tone ring, and then the internal DSP inside the sensor uses information from all three hall effect sensors to improve the accuracy of the measured time-varying magnetic flux, especially for slow moving objects. The DSP then generates the rotary encoder signal representing the wheel speed.

2.2   Types of Attacks

The attacks described in this work are all non-invasive, attempting to deceive an ABS sensor about the actual wheel speed without tampering with the internal circuitry of the sensor and/or the connection with the ABS controller. ABS sensors are exposed from

underneath the vehicle body (as shown in Figure 1b), making them an easy target for an attack. By placing an electromagnetic actuator in the air-gap between the ABS sensor and the ABS tone wheel (only a few millimeters in width), an attacker is able to modify the magnetic field measured by the ABS sensor. Two types of attacks can be mounted:


Fig. 2. Schematic of two proposed attacks: (a) Disruptive attack and (b) Spoofing attack. The black objects are the original components and signals in the ABS system while the red objects represents the external malicious components and signals injected by the attacker.

Disruptive Attacks. In a simplistic attack, the actuator placed near the ABS sensor is used to superimpose a malicious magnetic field on the original one. The resulting magnetic field will be different from the original one, but the attacker will not be able to precisely control the measured wheel speed, because the original magnetic field from the tone ring still has a considerable effect on the output of the speed sensor. Figure 2a offers a visual representation of this attack.

Spoofing Attacks. In order to deceive an ABS system into thinking a wheel is spinning at a precise speed, the attacker first needs to shield the sensor from the original magnetic field such that the gear rotation does not affect the sensor anymore, allowing the attacker to apply a new synthetic magnetic field corresponding to the new erroneous speed. The idea of this attack is summarized in Figure 2b.

As with the speed sensors themselves, magnetic shielding comes in both passive and active varieties. In passive shielding, a high permeability ferromagnetic material is used to provide a return path and thus significantly decrease the magnetic flux reaching the sensor. The main disadvantage of using this type of shielding in spoofing ABS speed sensors is that the small air-gap between the sensor and the rotating tone ring prohibits the use of thick shielding materials. In active shields a control feedback loop is used to sense the magnetic field and generate an opposing & canceling magnetic field.

Accordingly, in order to implement an active shield the attacker needs to implement a complete feedback loop—that is, the speed sensor spoofer needs to be equipped with a magnetic sensor, actuator, and controller. Only the sensor and actuator need be installed in the air-gap between the ABS sensor and the rotating gear, while all other components can be installed away from the ABS sensor. In the physical implementation of these
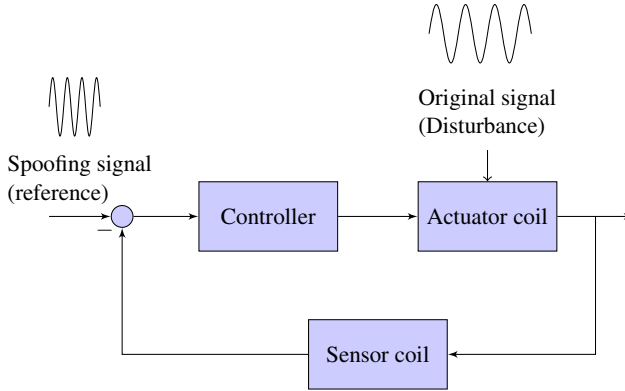
Fig. 3. Feedback control loop used in advanced attacks. The original magnetic field is modeled as a disturbance that needs to be rejected while the spoofing signal is modeled as a reference signal which the output should track.

spoofers, the sensors and actuators are realized as thin coils on a PCB. These coils are stacked one on top of the other and can easily be placed inside the air-gap.

Traditional work in active shielding often deals with suppressing static or slowly varying stray magnetic fields [9]. In this context the suppression must have better dynamic characteristics since the ABS sensor is just one block of a larger control loop. In other words, when the ABS hacker starts to suppress the magnetic field and spoof the ABS sensor, this information will be propagated to the ABS controller which will take action leading to a change in the very magnetic field which the ABS hacker is trying to suppress.

One way to negate the original magnetic field is to model it as a disturbance for which the feedback controller should compensate and force to zero. The spoofing signal is then modeled as a reference signal which the final output of the system should track precisely. The resulting control loop is shown in Figure 3, where both disturbance and reference signals are sinusoidal signals with varying frequency.

A final remark is that in order to attack the active ABS sensor, three feedback control loops are needed—one control loop for each sensor inside the sensor cluster. Accordingly, in this paper we will be presenting an ABS Hacker module which can be used to precisely spoof passive ABS sensors. However, these results can be extended directly to the case of the active ABS sensor.

3   ABS Spoofing Algorithm

As discussed in Section 2.2, in order to spoof the ABS signal, a feedback loop is required to suppress the original magnetic field and then apply a new synthetic one. In the feedback control literature, this problem is called the "error feedback output regulation problem" which we now discuss in this section.

3.1 Error Feedback Output Regulation

Both the disturbance and the reference signal to be tracked are assumed to be sinusoidal signals. We model these as an output of a harmonic oscillator which we call the *exo-system*. The dynamics of each harmonic oscillator can be written as:

$$\dot{w} = \begin{pmatrix} 0 & \omega \\ -\omega & 0 \end{pmatrix} w \tag{1}$$

where $w \in \mathbb{R}^2$ is the vector of the states for the harmonic oscillator, $\omega = \omega(t) \in \mathbb{R}$ is the frequency of the harmonic signal which changes with time, and the output of this exo-system is the first state. The amplitude and phase of the harmonic signal depends on the initial state of this exo-system. In the context of spoofing ABS sensors, we have two harmonic frequencies—one to reject denoted ω_o (representing the original magnetic field) and one for the attack signal denoted ω_a. Accordingly the dynamics of the exo-system can be written as:

$$\dot{w} = Sw, \quad w \in \mathbb{R}^4 \tag{2}$$

where

$$S = \begin{pmatrix} 0 & \omega_o & 0 & 0 \\ -\omega_o & 0 & 0 & 0 \\ 0 & 0 & 0 & \omega_a \\ 0 & 0 & -\omega_a & 0 \end{pmatrix} \tag{3}$$

The dynamics of the ABS Hacker system (including the actuator coil, sensor coil, sensor filters, and all supporting electronics) can be expressed as :

$$\dot{x} = Ax + Bu + Pw \tag{4}$$
$$\dot{w} = Sw \tag{5}$$
$$e = Cx - Qw \tag{6}$$

where $x \in \mathbb{R}^n$ is the state vector, $A \in \mathbb{R}^{nxn}$, $B \in \mathbb{R}^{nx1}$, $C \in \mathbb{R}^{1xn}$ represent the dynamics of the system, $u \in \mathbb{R}$ is the control input, $P = [1\ 0\ 0\ 0]^T \in \mathbb{R}^4$, $Q = [0\ 0\ 1\ 0]^T \in \mathbb{R}^4$, $w \in \mathbb{R}^4$, represents the vector of the exogenous inputs which in turn represents both the disturbance to be rejected and the signal to be tracked, $e \in \mathbb{R}$ is the tracking/regulated output, and it is required to find a controller $u = f(x, e)$ such that $\lim_{t \to +\infty} e(t) = 0$.

3.2 Comparison between Different Techniques

Solutions for the problem of asymptotically tracking/rejecting uncertain exogenous inputs of unknown or varying frequencies without measuring the disturbance have received increased attention in recent years within the control systems literature (see for example references in [10–12]).

Marino et al. in [10], applies results from indirect adaptive control theory. The dynamics of the system are transformed into the adaptive observer form, after which an observer is constructed to estimate the unknown frequency of the harmonic signal. The estimated frequency is then used to generate a sinusoidal signal with a 180 degree phase offset in order to reject the original signal.

Landau in [11] proposed a direct adaptive control scheme based on the internal model principle and the use of the Youla-Kucera parametrization. Instead of estimating the signal frequency and then changing the controller parameters, the Youla-Kucera parametrization allows to adaptively change the controller directly without the intermediate step of estimating the frequency explicitly. The resulting controller uses a technique called pole placement with independent objectives to separate the dynamics of disturbance rejection from the tracking dynamics [13]. In this technique, the poles of the regulation loop are kept fixed regardless of the value of the unknown frequency of the harmonic disturbance. The disturbance rejection uses a simple gradient-descent parameter adaptation algorithm to update the controller. An additional filter is then used to invert the dynamics of the regulation loop. The controller used to track the spoofing signal is easier to design and implement and thus we leave this simple exercise to the reader.

On the other hand, Isidori in [12] applies techniques from non-linear high-gain observer theory to design a robust non-linear observer and controller which is able to suppress the unknown harmonic signal without the conventional adaptation schemes.

In order to choose a suitable algorithm from among [10–12], we implemented all three. The metrics used to select the appropriate algorithm are the size of the constructed controller (measured by number of states) and the complexity of the algorithm in terms of process-hungry operations like online matrix inversion. These metrics lead to the selection of an algorithm which can fit within the computational power in the designed ABS hacker system.

The nonlinear algorithm presented in [11] requires an 8th order controller and observer. The main disadvantage of this algorithm is the requirement of an online inversion of an 8x8 matrix at each sampling period. Moreover, due to the usage of high-gain observers, the numerical values presented in the matrix to be inverted are quite large, leading to many challenges in resource- and processor-constrained microcontroller architectures.

The indirect adaptive observer presented in [10] is more complex due to the necessity of multiple transformations before the system is represented in the adaptive observer. This leads to an observer of size $= 2n + 6$. We will soon see that our proposed system has $n = 6$ (resulting from system identification experiments) , which results in an adaptive observer with order $= 18$. The algorithm in [10] also requires an online inversion of a 9×9 matrix.

The direct adaptive internal model algorithm presented in [11] uses three fixed linear digital filters, one adaptive parameter, and no matrix inversion operations. The complexity of the final design varies according to each specific design—the particular system described in this work requires a 12^{th} order linear controller. Based on this discussion, we adopt the algorithm presented in [11] for use in the design of the ABS hacker. Table 1 summarizes the results discussed above. Details of the chosen algorithm is reviewed in Appendix A at the end of this paper.

Table 1. Results of evaluating different error-feedback output regulation algorithms

	Indirect Adaptive Method [10]	Direct Adaptive Method [11]	Nonlinear High gain Observer [12]
Number of states	18	12	8
Matrix Inversion	9×9	0	8×8

4   ABS Hacker Hardware

In this section, we outline the development of the various hardware and software blocks necessary to implement an ABS hacker system. The final system is capable of launching both simplistic and advanced attacks, as described in Section 2.2. The schematic of the implemented ABS hacker is shown in Figure 4. The following is an overview of the main blocks of the proposed "ABS Hacker".

(a)   (b)

Fig. 4. (a) Schematic of the proposed ABS Hacker, (b) Final hardware implementation of the ABS Hacker, showing the sensing and actuation coil (top) and the driving circuitry (bottom)

4.1   Sensor and Actuator

The main components required to mount an attack on the ABS sensor are the sensing and actuation blocks. For actuation, we use a flat PCB coil driven by a high current op-amp. The usage of flat coils leads to a design which can fit within the small air-gap while still generating a magnetic field with the same amplitude of the original field. In order to maximize the magnetic field generated by the actuator, multiple flat coils are stacked on top of each other and placed electrically in series. This increases the effective number of turns for the aggregate coil without increasing the width of the PCB, and it has the added benefit of decreasing current consumption and increasing system lifetime.

In order to sense the magnetic field, we considered two different techniques. The first is to use a hall effect sensor while the other is to use a flat coil as a magnetic field pickup. After several preliminary experiments, the latter technique proved adequate for the ABS Hacker system. Three factors lead to this conclusion: 1) Size: a flat coil fits

better in the constraint of the air-gap. 2) Hall effect sensors generate a voltage which is proportional to the magnetic field density while a flat coil output is proportional to the change in the magnetic flux which is the same mode of operation as the ABS sensors. From the active shielding point-of-view, this leads to a simpler dynamical model to be used in the feedback loop.

4.2   Filtering

In an attempt to reduce the effects of noise from the various sources of EMI within the automotive body, the output of the flat coil sensor is used in differential mode connected to an instrumentation amplifier with high common-mode-rejection. The output is then filtered using an elliptic low-pass filter with a corner frequency at 500 Hz, corresponding roughly to a car speed of 100 mph (for a car with standard wheel size and a tone ring with 33 teeth).

4.3   Processing Elements and Interface

The ABS hacker operates in two modes: "waiting" and "spoofing." In the waiting mode, a wireless radio interface is duty-cycled until a spoofing attack command is received. Upon receiving this command, the ABS Hacker changes its mode and starts to spoof the magnetic field around the ABS sensor to change its measurements.

In order to reduce the current consumption in "waiting" mode and thus prolong the battery life, the designed system adopts a heterogeneous processor architecture. The first processor is a low power MSP430F2410, used to poll the radio interface until the attack command is received. Once such a command is received, the MSP430 cold boots the main processor—a high power ARM Cortex M4 STM32F407—and all corresponding peripherals. The higher power ARM has floating-point support and higher speeds needed to accomplish the DSP computations for accomplishing the active shielding in real-time.

The nature of such malicious attacks dictates that the hardware realization be as discreet as possible. The final hardware must be small enough to remain unseen, and it must also be able to fit within the small air-gap between the ABS sensor and the tone ring. The final system consists of the two parts shown in Figure 4b. The first part includes only the sensor and actuator to be placed within the air-gap, and the second part holds all supporting circuitry. Splitting the system into two like this allows for the sensor/actuator to remain small enough to fit within the air gap while the remaining bulkier circuitry can be placed in a distal location out of view. In order to maximize the effect of the actuator and reduce the required current drive, several coils are placed in series on successive layers. The resulting board contains 4 actuating coils and 1 sensing coil on a 6-layer PCB. The second part of the ABS hacker system is equipped with the radio for wireless activation, the low power MSP430, the powerful ARM Cortex M4, a high power amplifier to drive the coil actuator, and an instrumentation amplifier to condition the signal from the sensing coil.

5   Evaluation Results

5.1   Testbed

In order to test the proposed ABS Hacker, the testbed shown in Figure 5 was built. This testbed consists of two Mazda RX7 ABS sensors attached to a Mazda Rx7 tone ring. One of the two ABS sensors is used to provide the ground truth while the other one is used to simulate the sensor under attack. The tone ring is attached to a DC motor which emulates the action of the wheel shaft. The output of the two ABS sensors are connected to a MAX9926U evaluation kit which includes an ABS sensor interface capable of converting the raw sinusoidal wave into the encoded square wave. The output of the ABS sensors as well as the outputs of the MAX9926U evaluation kit are monitored by a real-time xPC Target system connected to MATLAB.

Fig. 5. Evaluation testbed consists of two Mazda car ABS sensors and a Mazda car tone ring and all signals are connected to MATLAB real-time xPC target for monitoring

5.2   Dimensions and Lifetime

The dimensions of the final actuator and sensor PCB are $40.64 \times 0.95 \times 24.13$ mm, fitting well within the typical air-gap for ABS sensors, while the driver PCB containing processing and amplification circuitry measures $25.4 \times 76.2 \times 1.524$ mm. The latter can still be easily tucked away and concealed from view. The lifetime of the system can be calculated for both idle and attacking modes. For an idle system, the quiescent current draw is 6.18 mA, dominated by the power-down current of the high power current amplifier. Given an 800 mAh battery, this gives an idle lifetime of 5.4 days. This can be dramatically increased by power-gating the high current op amp. During attack, the ARM processor consumes 109 mA while the peak attack current is 163 mA rms, giving an attack duration of 3 hours from a fresh battery, or $\frac{800 - 6.18 t_{idle}}{272}$ hours after waiting t_{idle} hours before the attack begins.

5.3 Disruptive Attack

The results of applying the disruptive attack are shown in Figure 7. By comparing the measured and the original unmodified wheel speeds (Figure 7a), it is obvious that the hacked wheel speeds is indeed different from the original unmodified wheel speed, but they are far from what the attacker was intending.

5.4 Spoofing Attacks

The dynamics of the system (including the actuator, high gain current amplifier, sensors, and signal conditioning circuit) are identified using standard system identification methods. We applied four different pseudo random binary sequences (PRBS) to the system, collected the output, and then applied prediction error techniques in order to build models of increasing complexity. Finally we used both whiteness tests and correlation tests to assess the quality of the obtained model [13]. One should also note that the physics of the inductive sensor implies the existence of a pure differentiator in the model. This observation can be used to simplify the system identification process by considering the differential of the PRBS input signal instead of the input PRBS itself. We choose the sampling frequency to be 5 times the max frequency in the disturbance which results in a sampling frequency of 2.5 kHz. The resulting model has $n_{A_d} = 5$, $n_{B_d} = 3$, and $d = 4$.

The presence of the pure delay is a side effect of using a digital low-pass filter which adds some delay in the processing from the point at which the signal is sampled until the time instance where the output is produced. In order to reduce this delay, the elliptic low-pass filter runs at a higher sampling rate. The frequency response of the resulting identified system is shown in Figure 6a, showing one vibration mode centered at 122 Hz.

Adaptive Controller Tuning. As discussed in Appendix A, the adaptive controller starts as a central fixed controller that is designed to insure the base-line specifications of the closed loop. This central controller has been designed using pole placement where the roots of the identified model are fixed in the same positions. By having the closed-loop poles the same as the open-loop poles, we generate a flat frequency response for the disturbance sensitivity function (the transfer function between the output and the disturbance input) as shown in Figure 6b (note that the scale of the magnitude and phase responses are 10^{-12} and 10^{-11}, respectively). This flat frequency response simplifies the adaptation of the internal model since all frequencies have the same gain. The complexity of the designed controller filters are $n_{S_0} = 7$ and $n_{R_0} = 4$.

The adaptation gain $F(t)$ is initialized with a value of 1000. The update of the adaptation gain is then done using the variable forgetting factor combined with the constant trace with $\lambda_0 = 0.95$, $\lambda_{threshold} = 3x10^{-9}$ and $\lambda_2 = 1$.

The tracking filter $T(q^{-1})$ is designed as discussed in Appendix A where the dynamics of the closed loop poles are inverted except the pure delay d and the zeros on the unit circle. This lead to a first order tracking filter.

Fig. 6. (a) Frequency response of the identified system, (b) Frequency response of the closed–loop disturbance sensitivity function

Experimental Results. The results of applying the spoofing attack are shown in Figure 7. Notice that the measured wheel speed shown in Figure 7b where the measured speed is almost exactly as specified by the attacker. These results show that a malicious attacker can precisely spoof ABS sensors to a specified wheel speed. In the next subsection we are going to evaluate the effect of such spoofing on the behavior of the vehicle.

5.5 Attack Consequences

Here we revisit the motivating example shown in Section 2, showing that the described ABS attack can lead to life-threatening situations. The simulation shown in Figure 8 shows the effect of applying the ABS sensor spoofing attack. This simulation is carried out by a high-fidelity, industrial-level simulator named "CarSim". The simulated scenario is as follows. First, the driver is heading in a straight path when he faces a patch of ice. Upon seeing the ice patch, the driver starts to apply the brakes. At this moment, the ABS spoofing attack begins on the right rear wheel. Due to the attack, the ABS controller receives an incorrect wheel speed (equal to zero in this case). Accordingly, the ABS controller mistakenly does not apply any brake to the right rear wheel. The consequence of this is that all other wheels start to slow down while the right rear wheel continues to spin, and the car slips off of the road. Figure 8 shows the position of the car at different snapshots in time.

6   Discussion and Future Work

We have shown that an attacker is able to precisely spoof ABS sensors and thus arbitrarily compromise a vehicle's ability to handle situations where wheels begin to lose traction. As noted earlier, the techniques illustrated in this work can very easily be applied to similar sensors, opening vulnerable systems to a slew of new attacks. What is not readily apparent, however, is how a system designer might go about protecting against these relatively new modes of attacks. A few distilled thrusts for future work include determining:
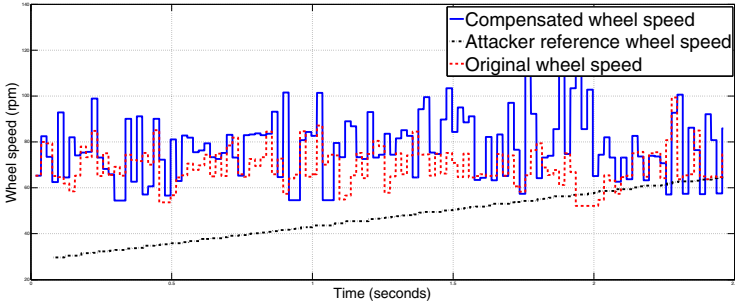
(a) Disruptive attack

(b) Spoofing attack

Fig. 7. Results of the disruptive attack (top) and the spoofing attack (bottom) showing the corresponding wheel speed detected by the output of the hacked ABS sensor (blue) versus the ground truth wheel speed measured by the un-attacked ABS sensor(green) along with the reference of the spoofing signal (black)

Fig. 8. The consequence of applying the ABS sensor spoofing attack while braking over ice. This simulation shows the position of the attacked car over multiple snapshots of time.

– In what other ways the physical environment can be spoofed
– What additional systems have potential vulnerabilities in this regard
– To what extend knowledge of the physical world can allow a system designer to reject erroneous, malicious signal injections

The answers to these questions are non-obvious and require careful consideration in future work.

7   Conclusions

Non-invasive attacks on cyber-physical systems pose considerable threats in situations that can be, at times, life critical. Such attacks are harder to detect at the sensor level and thus require higher level detection mechanisms. Using vehicle anti-lock braking systems, we have demonstrated both simplistic and advanced methods of non-invasive attacks on sensor subsystems. The advanced attack illustrates a very capable method for isolating sensors from the surrounding environment using results from adaptive feedback control theory before injecting a spoofed signal.

The proposed methodology has been evaluated for ABS sensors, where a small electronic module is designed and implemented to show the feasibility of the idea. We explored several aspects of designing such a module, and results obtained in real time from industrial ABS hardware lend credence to the efficacy of the attack and the threat that similar attacks pose.

References

1. Fawzi, H., Tabuada, P., Diggavi, S.: Secure state-estimation for dynamical systems under active adversaries. In: 2011 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton), pp. 337–344 (September 2011)
2. Cárdenas, A.A., Amin, S., Sastry, S.: Research challenges for the security of control systems. In: Proceedings of the 3rd Conference on Hot Topics in Security, HOTSEC 2008, pp. 6:1–6:6. USENIX Association, Berkeley (2008)
3. Igure, V.M., Laughter, S.A., Williams, R.D.: Security issues in SCADA networks. Computers and Security 25(7), 498–506 (2006)
4. Dorfler, F., Pasqualetti, F., Bullo, F.: Distributed detection of cyber-physical attacks in power networks: A waveform relaxation approach. In: Allerton, Allerton, IL, USA, September 2011, pp. 1486–1491 (2011)

5. Checkoway, S., McCoy, D., Kantor, B., Anderson, D., Shacham, H., Savage, S., Koscher, K., Czeskis, A., Roesner, F., Kohno, T.: Comprehensive experimental analyses of automotive attack surfaces. In: Proceedings of the 20th USENIX Conference on Security, SEC 2011, pp. 6–6. USENIX Association, Berkeley (2011)
6. Koscher, K., Czeskis, A., Roesner, F., Patel, S., Kohno, T., Checkoway, S., McCoy, D., Kantor, B., Anderson, D., Shacham, H., Savage, S.: Experimental security analysis of a modern automobile. In: 2010 IEEE Symposium on Security and Privacy (SP), pp. 447–462 (May 2010)
7. Fleming, W.: New automotive sensors - a review. IEEE Sensors Journal 8(11), 1900–1921 (2008)
8. Roosta, T., Shieh, S., Sastry, S.: Taxonomy of security attacks in sensor networks and countermeasures. In: The First IEEE International Conference on System Integration and Reliability Improvements, Hanoi, pp. 13–15 (2006)
9. Hilgenfeld, B., Strahmel, E., Nowak, H., Haueisen, J.: Active magnetic shielding for biomagnetic measurement using spatial gradient fields. Physiological Measurement 24(3), 661 (2003)
10. Marino, R., Santosuosso, G., Tomei, P.: Robust adaptive compensation of biased sinusoidal disturbances with unknown frequency. Automatica 39(10), 1755–1761 (2003)
11. Landau, I.D., Constantinescu, A., Rey, D.: Adaptive narrow band disturbance rejection applied to an active suspension-an internal model principle approach. Automatica 41(4), 563–574 (2005)
12. Isidori, A., Marconi, L., Praly, L.: Robust design of nonlinear internal models without adaptation. Automatica 48(10), 2409–2419 (2012)
13. Landaue, I.D., Lozano, R., M'Saad, M., Karimi, A.: Adaptive Control: Algorithms, Analysis and Applications. Communications and Control Engineering. Springer (June 2011)

A   Direct Adaptive Controller: A Recursive Least Square Filter

Details of the "Direct Adaptive Controller" are discussed in this appendix for the special case where the disturbance has only a single frequency (The algorithm presented in [11] can be applied to the case where the disturbance consists of multiple harmonics).

Since the designed controller will be implemented on a digital processor, it is convenient to express the controller in the discrete-time domain instead of the continuos-time domain. The designed controller consists of three digital filters which can be described by their polynomial representation, $R(q^{-1})$, $S(q^{-1})$, and $T(q^{-1})$, aptly named an RST-type controller [13], such that:

$$R(q^{-1}) = r_0 + r_1 q^{-1} + ... + r_{n_R} q^{-n_R}, \tag{7}$$

$$S(q^{-1}) = 1 + s_0 q^{-1} + ... + s_{n_S} q^{-n_S}, \tag{8}$$

$$T(q^{-1}) = t_0 + t_1 q^{-1} + ... + t_{n_T} q^{-n_T} \tag{9}$$

where the notation q^{-1} denotes time shift operator (i.e. $x(n-k) = xq^{-k}$). The proposed algorithm in [11] separates the problem into two parts, one for rejection and the other for tracking. The first two filters control the disturbance rejection dynamics while the final one separates the dynamics of the disturbance rejection from the dynamics of the tracking signal by means of dynamics inversion.
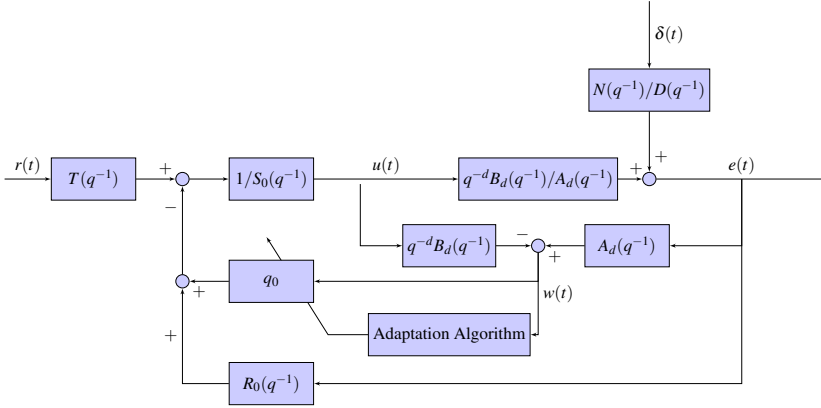
Fig. 9. R-S-T controller structure of the direct adaptive algorithm

The problem of rejecting the sinusoidal disturbance is formalized as a system identification problem where the goal is to identify the disturbance frequency by recursively minimizing the following cost function at each time step:

$$\arg\min_{\hat{\omega}} [\varepsilon]^2 \tag{10}$$

subject to the closed loop dynamics, and ε is the predicted error representing the difference between the predicted system output and the measured output (ε is a filtered version of the signal $w(t)$ shown in Figure 9). The controller filters $R(q^{-1})$ and $S(q^{-1})$ are then redesigned using the internal model principle in order to perfectly reject the sinusoidal disturbance. However, the authors of [11] found that it is more effective to adapt the controller directly without passing through the intermediate step of identifying the sinusoidal frequency by slight modifications in the equations of the system identification. In order to do that, the disturbance rejection filters are parametrized using Youla-Kucera (Q-parametrization) into the form:

$$R(q^{-1}) = R_0(q^{-1}) + A_d(q^{-1})Q(q^{-1}) \tag{11}$$

$$S(q^{-1}) = S_0(q^{-1}) - q^{-d}B_d(q^{-1})Q(q^{-1}) \tag{12}$$

where $A_d(q^{-1})$ and $B_d(q^{-1})$ are polynomials of order n_{A_d} and n_{B_d}, respectively. These two polynomials along with the pure delay q^{-d} denote the discrete-time representation of the system dynamics in the complex Z-domain with the understanding that z^{-1} is equivalent to q^{-1}. The central controller comprised of $R_0(q^{-1})$ and $S_0(q^{-1})$ is fixed and can be computed by pole placement and should be designed to give the closed-loop specifications required in the absence of the disturbance. $Q(q^{-1}) = q_0$ compensates the

effect of the varying coefficients such that the closed-loop poles (denoted as $P(q^{-1})$) remain fixed. Accordingly, the optimization objective (10) can be re-written as:

$$\arg\min_{\hat{q}_0} [\varepsilon]^2 \tag{13}$$

which can be solved by using adaptive filters utilizing a gradient descent algorithm with a variable adaptation gain of the following form:

$$q_{0_{n+1}} = q_{0_n} + F_n \phi_n \varepsilon_{n+1}, \tag{14}$$

where ϕ_n is the regressor vector containing filtered input and output measurements. The constraints of the closed loop dynamics are used to build both ϕn and ε_{n+1} as described in [11].

The variable adaptation gain F_n is designed such that the it does not reach zero and consequently moves in the optimal direction [13]:

$$F_{n+1} = \frac{1}{\lambda_{1_n}} \left[F_n - \frac{F_n \phi_n \phi_n^T F_n}{\frac{\lambda_{1_n}}{\lambda_2} + \phi_n^T F_n \phi_n} \right] \tag{15}$$

$$\lambda_{1_n} = \begin{cases} \lambda_0 \lambda_{1_{n-1}} + 1 - \lambda_0 & \text{if } \lambda_{1_n} > \lambda_{threshold} \\ \lambda_{threshold} & \text{otherwise} \end{cases} \tag{16}$$

where $\lambda_0, \lambda_1 n, \lambda_2, \lambda_{threshold} \in]0,1]$, denote the forgetting factors of the adaptation gain.

The last controller filter is the tracking filter $T(q^{-1})$. In order to achieve perfect tracking, $T(q^{-1})$ is designed to invert the dynamics of the disturbance rejection loop which can be written as:

$$T(q^{-1}) = \frac{P(q^{-1})}{B_d(q^{-1})} \tag{17}$$

However, one should take care that since $B_d(q^{-1})$ may contain unstable zeros, the design of $T(q^{-1})$ can invert only the stable zeros, and then flip the steady state frequency response of the remaining part. The same design technique can be used to get rid of the effect of the pure delay in the system. Figure 9 shows the schematic for the described controller.