HeatPulse: Thermal Attacks on Air Pollution Sensors

Natsuki Morand
Inria, INSA-Lyon, Univ. Lyon, France
natsuki.morand@etu.univ-lyon1.fr

Ahmed Boubrima
Inria, INSA-Lyon, Univ. Lyon, France
ahmed.boubrima@inria.fr

Walid Bechkit
Inria, INSA-Lyon, Univ. Lyon, France
walid.bechkit@insa-lyon.fr

Zhambyl Shaikhanov
University of Maryland, USA
zhambyl@umd.edu

ABSTRACT

Air pollution monitoring, especially using miniaturized and low-cost sensors, has been increasingly adopted across many application areas - including chemical and automotive industries, smart cities, and agriculture - to protect public health and comply with environmental regulations. In this paper, we demonstrate a new security vulnerability that enables adversaries to remotely spoof low-cost air pollution sensors via long-range and highly localized thermal attacks. Modeling the key attack characteristics, we show how adversaries can exploit the temperature-dependent internal calibration process of air pollution sensors to strategically manipulate sensor measurements through induced heat signals, thereby deceiving the sensors. Using inexpensive laser pointers and commercial Nitrogen Dioxide pollution sensors, we design an evaluation testbed and experimentally show the attack's effectiveness in both indoor and outdoor environments.

CCS CONCEPTS

• **Computer systems organization → Embedded and cyber-physical systems**; **Sensor networks**; **Embedded systems**; • **Security and privacy → Embedded systems security**;

KEYWORDS

Sensor security; sensor attacks; air pollution sensors; air quality monitoring.

1 INTRODUCTION

Recent technological advances, particularly in electronics miniaturization, have contributed to the democratization of air pollution sensors as many manufacturers are now marketing and commercializing low-cost sensors to the general public. These low-cost sensing devices help measure air quality in the immediate environment without expensive equipment and offer as a result a promising complementary solution to the traditional air quality monitoring reference stations [1].

Today, many local and governmental decisions are taken based on air quality and pollution emissions [2]. This includes, for instance, traffic restrictions in case of high pollution levels, or factory production limitations in chemical and automotive industries depending on gas and particle emission levels. As such decisions rely on data collected from air pollution measurement sensors and reference stations, studying air pollution sensor attacks ensures accurate sensing and helps prevent catastrophic public decisions.

Sensor attacks: Sensor attacks are a type of security attack where the adversary injects a fake malicious physical signal to alter the measurements of the targeted sensor [3]. The attack relies on the fact that the sensing probe blindly trusts the physical signals that are measured as there is no way to differentiate between legitimate and fake input signals. The intended outcome of the adversary attack could be either i) introducing a measurement drift (i.e. sensor spoofing attacks), which can result in triggering a false alarm for instance; or ii) causing a sensor failure by making the sensor completely blind (i.e. sensor saturation attacks) [4]. To achieve a sensor attack, the adversary can use any signal type that interferes with the measurement process of sensing probes (acoustics, ultrasound, infrared, thermal, etc).

Sensor security challenges in air pollution sensing: Low-cost air pollution sensors can be easily impacted by weather changes and mainly temperature variations [5]. To counter these negative weather effects, air pollution sensor manufacturers implement calibration methods that convert the output electrical signal of the pollution sensing probe (in *mv*) to a meaningful pollution concentration level (in *ppb* or $\mu g/m^3$) while accounting for the impact of ambient temperature levels [6]. This is achieved by embedding a temperature sensing probe in the air pollution sensor box, which allows measuring the ambient temperature level that is affecting the pollution sensing process.

Although pollution calibration methods are efficient in correcting pollution measurements in field deployments [7], the security aspect of the measurements carried out by air pollution sensors has not yet been investigated in the literature. Indeed, calibration methods assume that the sensing context (weather conditions such as temperature) varies in a natural way [8] (i.e. a natural increase or decrease in temperature levels through a typical day for instance). As a result, an adversary can remotely manipulate the output of the calibration function using a remote disruptive physical signal that affects in an unnatural way the temperature levels on the physical surface of the pollution sensing probe.

Research objectives: In this paper, we aim to demonstrate air pollution sensor spoofing attacks that remotely introduce a drift in sensor measurements and as a result either trigger false alarms or prevent true alarms. The main contributions of our work are as follows:

(i) We leverage the sensitive nature of air pollution sensors to ambient temperature levels as we design a thermal attack system while using inexpensive off-the-shelf lasers as a heating source. The presented attack is noninvasive and stealthy as it can be initiated remotely using the laser's thermal power.

(ii) We perform multiple indoor (lab environment) and outdoor (campus field environment) experiments on a lab-designed air pollution sensor while using commercial Nitrogen Dioxide (NO_2) Alphasense sensing probes [9].

(iii) We show the efficiency of the attack system while highlighting the attack outcomes and the potential hardware and software defense solutions.

Paper structure: The remainder of this paper is organized as follows. We first discuss the related work on sensor security in Section 2 and address the background of air pollution sensing in Section 3. Then, we present the air pollution sensor attack model in Section 4. Next, we present the main sensor and adversary components of our lab-designed evaluation testbed in Section 5. Finally, we analyze the evaluation results and discuss the possible hardware and software countermeasures in Sections 6 and 7.

2 RELATED WORK

Sensor security is a challenging area due to sensor heterogeneity and the fact that the sensing signal modality varies from one sensor to another [3]. As a result, sensor attacks and their corresponding defense methods depend on the sensor operation mechanism and the application area.

Shin et al. [10] study the vulnerabilities and defense methods of optical smoke detectors while using highly directional light sources to alter the operation of the sensors. In [11], Tu et al. study the vulnerability of temperature-based control systems against electromagnetic injected signals. In [12], Barua et al. demonstrate how an adversary can use an external magnetic field to spoof Hall sensors. Park et al. [13] focus on studying the security of infrared sensors in medical infusion pumps as they demonstrate how infrared drop sensors can be saturated using an additional infrared source. In [14], Sun et al. study LiDaR sensor spoofing attacks that use external lasers to alter the operation of LiDaRs.

In contrast to prior work, and to the best of our knowledge, this paper is the first to focus on the security of low-cost air pollution sensors while using remote thermal attacks to alter the output of the temperature-dependent pollution calibration functions.

3 BACKGROUND: POLLUTION SENSING

Air pollution sensors measure the concentration of a specific pollutant (gaseous pollutants such as Nitrogen Dioxide, or particulate pollutants such as PM2.5) [15]. The contact between the ambient air and the surface of the pollution sensing probe of a given pollutant creates a reaction that generates an output electrical signal (current or voltage). The output electrical voltage is proportional to the concentration of the measured pollutant in the air. Furthermore, the reaction generating the output voltage depends on the sensor's technology. For instance, Nitrogen Dioxide sensors use electrochemical cells that rely on *Redox* reactions (oxidation-reduction) [9].

3.1 Conversion of raw measurements

The conversion of the output voltage of air pollution sensors to a meaningful concentration (in ppb or $\mu g/m^3$ depending on the pollutant nature) is performed using a calibration formula that is provided by the manufacturer [16]. In lab conditions (lab temperature and humidity levels), the sensor's electrical response to pollutant concentrations is linear for most air pollution sensors. As a result, the calibration formula can be easily obtained by exposing the pollution sensing probe to different pollutant concentrations and then deriving the linear relationship between the sensor's output voltage and reference concentrations.

3.2 Impact of temperature and humidity

Compared to most other environmental sensors (such as wind, temperature, and humidity sensors), air pollution sensors are highly sensitive to multiple physical quantities and do not react only to the physical pollution signal they were designed to measure. Mainly, the electrical output of air pollution sensors is temperature dependent due to the impact of weather conditions on the reactions that are used in the sensing process of air pollution sensors [5]. To counter these effects, air pollution sensor manufacturers include in their calibration formulas a correction factor, which accounts for the changing temperature levels (and which is not necessarily linear) [6]. As a result, the concentration levels measured by air pollution sensors are a function of (i) the output electrical voltage of the pollutant sensing probe in addition to (ii) temperature levels that are measured using a separate sensing probe (which is also integrated within the air pollution sensor device).

4 SECURITY ANALYSIS

4.1 Sensor design vulnerabilities

Due to the low-cost nature of air pollution sensors, both the pollutant and temperature sensing probes are usually directly exposed to the ambient air [15]. This opens the door for a security vulnerability in the design of air pollution sensors as an adversary can change the temperature of the pollutant sensing probe through a localized thermal attack without altering the measurements of the temperature sensing probe. In such a case, the manufacturer's calibration function fails in correcting the temperature-dependent output electrical signal as there will be no way to identify the temperature change that is affecting the pollutant sensing probe.

4.2 Assumptions

We consider the following assumptions:

- The adversary cannot manipulate the ground-truth pollutant concentrations. This means that we exclude the case where a pollutant is released in the air to trigger a pollution peak alarm for instance.
- The attack must be non-invasive and stealthy by being initiated from a distance. This means that the adversary is not allowed to tamper with the hardware or the software of the sensor. Therefore, we exclude attacks where the adversary can get in touch with the sensor (i.e. only remote attacks are possible). We also exclude, for instance, the attacks that aim to change the software code of the pollution calibration functions.

4.3 Attack model

We consider the following attack model:

- **Adversary:** The adversary is a third-party entity that can be either transported by a person (as illustrated in Fig. 1) or mounted on a remotely-controlled drone.
- **Intent of the adversary:** The adversary aims to introduce a drift (positive or negative) in the pollution measurements that are determined using the sensor's internal calibration formulas.
- **Attack signal:** The adversary uses remote localized heating (via a strong enough laser for instance) to cause a long-range thermal impact on the surface of the pollutant sensing probe. This, as a result, alters the output of the manufacturer's temperature-dependent calibration function as there will be no way to identify the temperature change that is affecting the pollutant sensing probe.

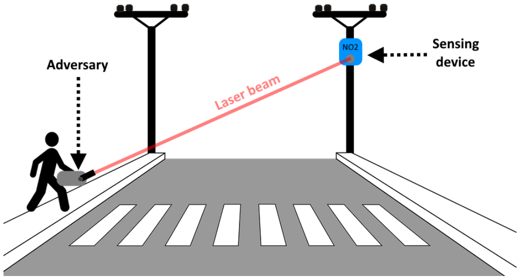

Figure 1: HeatPulse attack scenario

5 EVALUATION TESTBED

In this section, we present the components of our designed evaluation testbed, which consists of two main parts: (i) a low-cost air pollution sensor using the widely-adopted commercial Alphasense Nitrogen Dioxide sensing probes [17]; and (ii) a remote thermal attack system using inexpensive commercial laser pointers as a heat source.

5.1 Low-cost air pollution sensor

To perform our security analysis, we designed a high resolution Nitrogen Dioxide sensor (depicted in Fig. 2) based on an Arduino Yún board and an Alphasense NO2-B43F sensing probe [9]. The Alphasense sensing probe is an electrochemical cell, which relies on *Redox* reactions (oxidation-reduction) to output an electrical current that is proportional to Nitrogen Dioxide concentrations in the air.

The NO_2 sensing probe has two main output electrodes (working and auxiliary) and is connected to an ISB board, which converts the output currents into voltages. The converted voltages vary mostly between $200mV$ and $500mV$ depending on the NO_2 concentrations, with the minimum value corresponding to an offset voltage that is introduced by the ISB board to minimize the impact of the electronic noise. In order to measure the NO_2 voltages, we use a 16-bit analog-to-digital converter with a $0.015mV$ resolution, which allows us to get 1ppb-resolution pollution measurements.

In order to account for the temperature-dependent nature of air pollution sensors, we also integrate in our sensor device a DHT22 sensing probe that measures temperature (with a 0.1^oC resolution) in addition to relative humidity (with a 1% resolution).

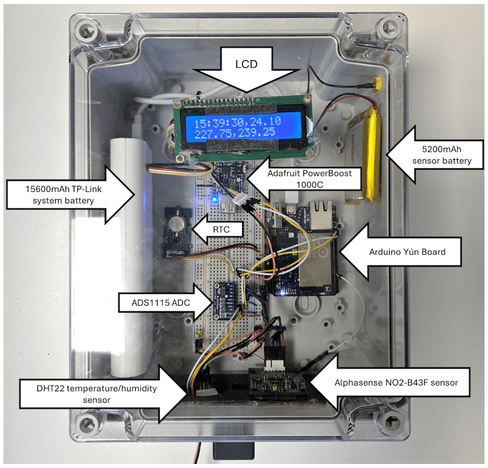
LCD
15:39:30,24.10
227.75,239.25
5200mAh sensor battery
Adafruit PowerBoost 1000C
15600mAh TP-Link system battery
RTC
Arduino Yún Board
ADS1115 ADC
DHT22 temperature/humidity sensor
Alphasense NO2-B43F sensor

Figure 2: Testbed pollution sensor and its components

NO_2 **calibration function:** To convert the electrical voltages delivered by the NO_2 sensing probe, and following the guidelines of the manufacturer [6][9][16], we implement the pollution calibration function as:

$$P = \frac{V - V_0 - (T - T_{ref}) \cdot T_{factor}}{V_{sens}},$$

where P is the converted pollution concentration in ppb, V is the voltage (in mV) delivered by the working electrode of the NO_2 probe, V_0 is the working electrode's voltage offset corresponding to a $0ppb$ gas concentration level, T is the temperature (in ^oC) measured by the DHT22 sensing probe, T_0 is the reference temperature level of the NO_2 sensing probe and is equal to 25^oC, T_{factor} (expressed in $mV/^oC$) is the temperature correction factor and corresponds to the voltage change caused by a 1^oC difference in temperature levels, and finally V_{sens} (expressed in mV/ppb) is the NO_2 sensitivity parameter and corresponds to the sensor voltage change due to a $1ppb$ difference in NO_2 concentrations.

Note that although the NO_2 sensing probe has two output signals, we only use the working electrode (the electrode that is exposed to the polluted air). Indeed, the auxiliary electrode (which is not exposed to the polluted air) can also be used in the calibration formula but its integration is challenging because it reacts differently to temperature changes and also ages differently over time compared to the working electrode.

5.2 Attack system

We build our attack system (illustrated in Fig. 3) using laser diodes as a heat source. Compared to other heating solutions such as heating fans and infrared heaters, lasers deliver highly localized and long-range heating, which allows the stealthy adversary to manipulate the calibration function by heating the pollution sensing probe without affecting temperature measurements.

To demonstrate the easy accessibility of the attack, we use an inexpensive laser pointer coupled with a biconcave lens and connected to an Arduino board that is powered using a $5200mAh$

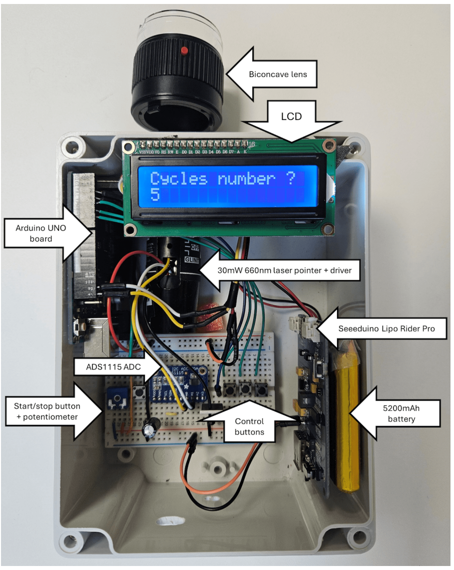

Figure 3: Testbed attack system and its components

lithium-ion battery. The selected laser pointer is made of a $30mW$-rated $650nm$ laser diode. The biconcave lens allows us to better focus the laser output, increasing its thermal impact on the sensor's surface.

We modified the laser pointer to replace its original AA battery with a $5V$ power supply and also added an NPN mosfet to control the duty cycle of the laser output using the Arduino board. This allows us to study the impact of both short-term thermal pulses and long-term heat attacks. We also implemented a user interface with an LCD screen and a set of buttons that allow us to demonstrate the attack's effectiveness in real time. The designed attack system has an average current consumption of $51mA$ when idling and $190mA$ when the laser is turned on, which allows us to achieve over 20 hours of continuous thermal attacks using the integrated battery.

6 EXPERIMENTAL EVALUATIONS

6.1 Evaluation setup

We conducted multiple experiments in an indoor lab environment and an outdoor campus field in Lyon, France during May and June 2024. The indoor setup is air-conditioned with temperature levels oscillating between 24 and 26 degrees ^oC. Regarding the outdoor experiments, temperature levels were in the range of 23-28 degrees ^oC during the period of the experiments. In both indoor and outdoor experiments, the attack system was placed $30cm$ away from the sensor device as illustrated in Fig. 4. Placing the attack system at a much larger distance is also possible but requires more powerful lasers depending on the desired distance.

We used an Envea Cairsens sensor [18] to determine the reference NO_2 concentrations in both the indoor lab environment and the outdoor campus environment. The reference concentrations were on average $5ppb$ indoors and $15ppb$ outdoors, which correlates very well with expected NO_2 concentrations [19]. Compared to the reference baseline, our designed low-cost sensor device is well correlated and achieves a normalized RMSE lower than 10% (using the implemented temperature-dependent calibration function).

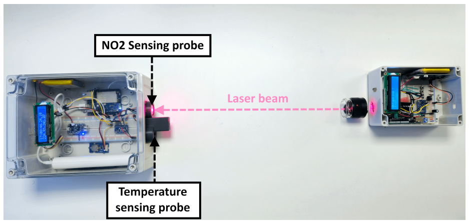

Figure 4: A top view of the evaluation setup

6.2 Evaluation results

6.2.1 Characterization of the attack's thermal power. We first determine the temperature change that is expected when the surface of the NO_2 sensing probe is exposed to the laser beam. To that end, we focus the beam of the attack system on a DHT22 temperature and humidity sensing probe. The targeted DHT22 in this scenario was uncovered to properly expose the temperature thermistor that is responsible for temperature measurements. We maintain the attack for $40min$ and depict the results in Fig. 5.

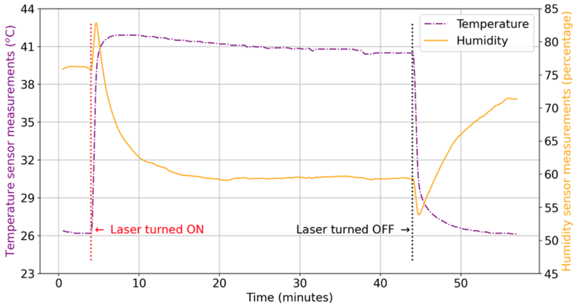

Figure 5: Heating power of our attack system

The results show that the designed attack system successfully raises the temperature of the exposed DHT22 thermistor from 26.2^oC to 41.9^oC in $3min$ before the temperature levels stabilize at around 40^oC for the remaining time of the attack. The drop to 40^oC is mainly due to the varying temperature levels in our air-conditioned lab environment. Note also that the temperature increase caused by the designed attack system is fast enough to trigger short-term heat pulses. Indeed, around 90% of the thermal impact is achieved in $1min$ as the DHT22 temperature values rise from 26.2^oC to 40.5^oC.

In addition to temperature variations, we also depict in Fig. 5 the impact of the laser thermal attack on ambient humidity levels. As expected, the increase in temperature levels results naturally in lower relative humidity measurements. We also notice that relative humidity slightly increases for a few seconds before starting to drop when the attack is initiated (with a similar pattern happening when the laser diode is turned off). This can be attributed to the low-precision humidity measurements as the DHT22 sensing probe is mainly designed for high-precision temperature sensing.

6.2.2 Indoor evaluation of the thermal attack. To demonstrate the effectiveness of the attack on air pollution measurements, we aim the laser beam toward the NO_2 sensing probe as already illustrated in Fig. 4. We maintain the attack for $20min$ and depict in Fig. 6 the resulting NO_2 measurement drift with respect to reference concentrations.

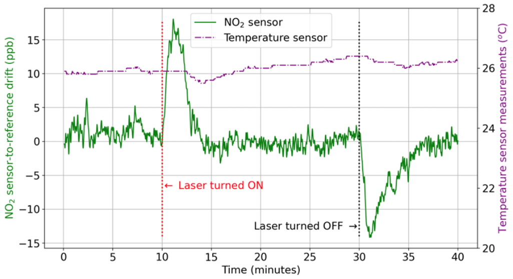

Figure 6: Attack's impact on indoor NO_2 measurements

We first notice that the attack results in a measurement drift that cannot be corrected in pollution calibration functions as the temperature sensing probe is not exposed to the laser beam and reports measurements of around 26^oC throughout the whole experiment. The measurement drift caused by the thermal attack is variable and can be classified into 3 different types:

(i) **Transient heating short-term impact:** When the thermal attack is first triggered, the NO_2 calibration function output increases exponentially and reaches a difference (with respect to the reference sensor) of around $15ppb$ within the first minute of the attack. This positive drift is due to the impact of temperature transients on the electrochemical cell (that is responsible for measuring NO_2 concentrations) according to the sensor manufacturer [6]. Indeed, and as already demonstrated in the previous test, most of the thermal impact caused by the designed attack system is achieved in around $1min$. We also varied in our lab tests the distance between the attack system and the sensor device and concluded that the positive drift caused by the transient heating impact depends on the thermal attack power and can be further increased using a more powerful laser diode.

(ii) **Transient cooling short-term impact:** We also notice that upon turning off the laser beam, the NO_2 measurements drift in a similar pattern but decrease rather than increasing and reach a negative drift of around $-15ppb$ within one minute of disabling the attack system. Similarly to transient heating, this short-term sensor drift is due to negative temperature transients according to the manufacturer [6].

(iii) **Steady-state heating long-term impact:** The final impact of the studied thermal attack, although less visible, can be observed right after the first transient heating short-term impact. Indeed, the NO_2 sensor measurements slightly decrease after the first peak and reach an average negative drift of around $-1ppb$ for the remaining time of the laser attack duration. Similarly to the short-term impacts, our tests also showed that the attack distance and laser power determine the amount of this long-term negative drift.

Finding: Depending on the intent of the adversary, the thermal attack power and the attack duration need to be adapted to favor either (i) a short-term positive drift (to cause a pollution alarm), (ii) a short-term negative drift (to prevent a true alarm), or (iii) a long-term negative drift (to maintain pollution measurements within the regulatory threshold for instance).

6.2.3 *Outdoor evaluation of the thermal attack.*

We perform the same previous experiment in our campus field in Lyon, France, and depict the results in Fig. 7. Similarly to the indoor test results, the thermal attack causes first a positive short-term drift, followed by a long-term small negative drift, and finally a short-term negative drift upon disabling the attack system. However, due to the less efficient laser propagation in outdoor environments, the attack results in an $11ppb$ short-term positive drift and a $-9ppb$ short-term negative drift, which is 30% lower than the indoor test results.

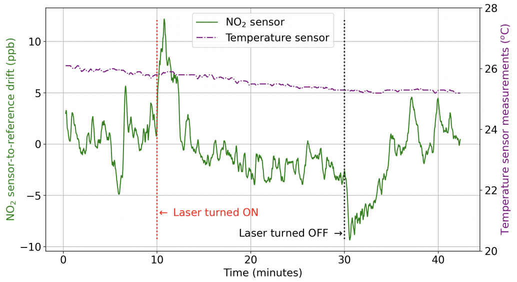

Figure 7: Attack's impact on outdoor measurements

7 COUNTERMEASURES

Hardware solutions: Shielding the pollution and temperature sensing probes while using an internal fan for airflow circulation can make thermal attacks detectable. The temperature sensing probe can be used in this case to identify unnatural variations in temperature levels. However, this increases the cost and power consumption of the low-cost sensor and heavily reduces its response time and ability to detect quick pollution emission events.

Software solutions: Besides hardware solutions, we are currently exploring in our ongoing work the use of machine learning techniques to train the sensors on the expected electrical signal variations that are due to thermal attacks. Our objective is to improve the calibration function by filtering out the electrical signal variations that do not result from an actual pollution event. Another software countermeasure would be to use sensor fusion by leveraging the measurement correlations between neighboring sensors.

8 CONCLUSION

In this paper, we demonstrate a new security vulnerability that enables adversaries to spoof low-cost air pollution sensors via localized and long-range heating. To that end, we leverage the sensitive nature of pollution sensors to temperature variations. We perform multiple indoor and outdoor experiments to show the efficiency of the thermal attack while highlighting the attack's short and long-term impacts.

ACKNOWLEDGMENTS

This work has been funded and supported by the French National Research Agency (ANR) under the project ANR-21-CE25-0003 (ANR DRON-MAP).

REFERENCES

[1] Ye Kang, Lu Aye, Tuan Duc Ngo, and Jin Zhou. Performance evaluation of low-cost air quality sensors: A review. *Science of The Total Environment*, 2022.

[2] Wei Zhang, C-Y Cynthia Lin Lawell, and Victoria I Umanskaya. The effects of license plate-based driving restrictions on air quality: Theory and empirical evidence. *Journal of Environmental Economics and Management*, 2017.

[3] Anomadarshi Barua and Mohammad Abdullah Al Faruque. Sensor security: Current progress, research challenges, and future roadmap. In *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design*, 2022.

[4] Dohyun Kim, Mangi Cho, Hocheol Shin, Jaehoon Kim, Juhwan Noh, and Yongdae Kim. Lightbox: Sensor attack detection for photoelectric sensors via spectrum fingerprinting. *ACM Transactions on Privacy and Security*, 2023.

[5] Balz Maag, Olga Saukh, David Hasenfratz, and Lothar Thiele. Pre-deployment testing, augmentation and calibration of cross-sensitive sensors. In *EWSN*, 2016.

[6] Alphasense LTD. Environmental changes: temperature, pressure, humidity, 2022. Application Note AAN 110/SEP22.

[7] Ahmed Boubrima and Edward W Knightly. Robust environmental sensing using UAVs. *ACM Transactions on Internet of Things*, 2021.

[8] Yun Cheng, Olga Saukh, and Lothar Thiele. Sensorformer: Efficient many-to-many sensor calibration with learnable input subsampling. *IEEE Internet of Things Journal*, 2022.

[9] Alphasense LTD. NO2-B43F Nitrogen Dioxide sensor, 2023. Technical specifications datasheet.

[10] Hocheol Shin, Juhwan Noh, Dohyun Kim, and Yongdae Kim. The system that cried wolf: sensor security analysis of wide-area smoke detectors for critical infrastructure. *ACM Transactions on Privacy and Security*, 2020.

[11] Yazhou Tu, Sara Rampazzi, Bin Hao, Angel Rodriguez, Kevin Fu, and Xiali Hei. Trick or heat? manipulating critical temperature-based control systems using rectification attacks. In *ACM SIGSAC Conference on Computer and Communications Security*, 2019.

[12] Anomadarshi Barua and Mohammad Abdullah Al Faruque. Hall spoofing: A non-invasive dos attack on grid-tied solar inverter. In *29th USENIX Security Symposium*, 2020.

[13] Youngseok Park, Yunmok Son, Hocheol Shin, Dohyun Kim, and Yongdae Kim. This ain't your dose: Sensor spoofing attack on medical infusion pump. In *10th workshop on offensive technologies*, 2016.

[14] Jiachen Sun, Yulong Cao, Qi Alfred Chen, and Z Morley Mao. Towards robust lidar-based perception in autonomous driving: General black-box adversarial sensor attack and countermeasures. In *USENIX Security Symposium*, 2020.

[15] Anamika Sharma, Brijesh Mishra, Ronak Sutaria, and Rajesh Zele. Design and development of low-cost wireless sensor device for air quality networks. In *IEEE TENCON Region 10 Conference*, 2019.

[16] Alphasense LTD. Correcting for background currents in four-electrode toxic gas sensors, 2023. Application Note AAN 803-05/JAN23.

[17] Nima Afshar-Mohajer, Christopher Zuidema, Sinan Sousan, Laura Hallett, Marcus Tatum, Ana M Rule, Geb Thomas, Thomas M Peters, and Kirsten Koehler. Evaluation of low-cost electro-chemical sensors for environmental monitoring of ozone, nitrogen dioxide, and carbon monoxide. *Journal of occupational and environmental hygiene*, 2018.

[18] Envea Global. Miniature solution for real-time continuous pollution monitoring, 2022. Technical specifications datasheet.

[19] Heidi Salonen, Tunga Salthammer, and Lidia Morawska. Human exposure to no2 in school and office indoor environments. *Environment international*, 2019.