

Is your smoke detector working properly?

Robust fault tolerance approaches for smoke detectors

Arjun Tambe*
Microsoft Research, India
arjuntambe1@gmail.com

Akshay Nambi
Microsoft Research, India
akshayn@microsoft.com

Sumukh Marathe
Microsoft Research, India
t-summar@microsoft.com

ABSTRACT

Billions of smoke detectors are in use worldwide to provide early warning of fires. Despite this, they frequently fail to operate in an ongoing fire, risking death and property damage. A significant fraction of faults result from drift, or reduced sensitivity, and other faults in smoke detectors' phototransistors (PTs). Existing approaches attempt to detect drift from the PT output in normal conditions (without smoke). However, we find that drifted PTs mimic the output of working PTs in normal conditions, but diverge in the presence of smoke, making this approach ineffective.

This paper presents two novel approaches to systematically detect faults and measure and compensate for drift in smoke detectors' PTs. Our first approach, called FallTime, measures a PT "fingerprint," a unique electrical characteristic with distinct behavior for working, drifted, and faulty components. FallTime can be added to many existing smoke detector models in software alone, with no/minimal hardware modifications. Our second approach, DriftTestButton, is a mechanical test button that simulates the behavior of smoke when pressed. It offers a robust, straightforward approach to detect faults, and can measure and compensate for drift across the entire smoke detector system. We empirically evaluate both approaches and present extensive experimental results from actual smoke detectors deployed in a commercial building, along with custom-built smoke detectors. By conducting tests with live smoke, we show that both FallTime and DriftTestButton perform more effectively than existing fault tolerance techniques and stand to substantially reduce the risk that a smoke detector fails to alarm in the presence of smoke.

CCS CONCEPTS

• Computer systems organization → Embedded and cyber-physical systems; • Hardware → Fault tolerance.

KEYWORDS

Smoke detectors; Phototransistors; Fault tolerance; Drift detection

*Work done while the author was an intern at Microsoft Research India.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiSys '21, June 24– July 02, 2021

© 2021 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/21/xx... \$15.00
<https://doi.org/10.1145/1122445.xxxxxx>

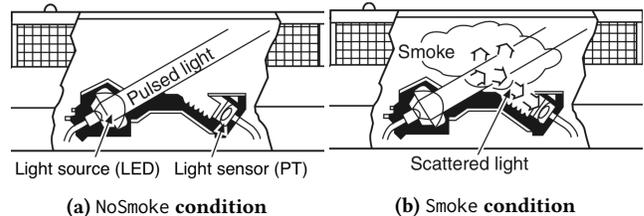


Figure 1: Working principle of a Photoelectric Smoke Detector [49].

1 INTRODUCTION & RELATED WORK

Smoke detectors (or smoke alarms/sensors) are currently being integrated into smart homes, hospitals and shopping malls, and constitute one of the most important, safety-critical applications of networked sensor systems. These smoke detectors are deployed at scale to provide early warning of fires, and are responsible for precipitous declines in fire deaths [12, 13]. Despite this, smoke detectors frequently fail to operate in an ongoing fire, risking death, injury and property damage. In the US alone, smoke alarms failed to operate in 15% of all fires, contributing to nearly 2,000 fire casualties and property damage of \$10 billion per year [13]. Improving the reliability of smoke detectors can significantly reduce the rates of death, injury, and property damage, which represents a key application of mobile systems.

Photoelectric detectors are the most common type of smoke detector, and are more responsive to smoldering fires [7], which are frequent in household and commercial settings [13]. Photoelectric smoke detectors employ a light source (LED) and a phototransistor (PT), which converts light signals into electrical signals. Figure 1 shows the working principle of photoelectric detectors. As smoke enters the chamber and crosses the path of the light beam, light is scattered by smoke particles toward the PT. Figure 2 shows a simplified circuit diagram of a smoke detector. A microcontroller (MCU) reads the PT output; if the PT output exceeds a pre-calibrated threshold, the MCU triggers an alarm.

A smoke detector may fail to alarm promptly for numerous reasons. Significant progress has already been made to address a subset of failure scenarios such as, (i) failure of a power source; (ii) improper placement of the unit, preventing smoke from flowing towards the detector; and (iii) basic electrical failures, such as broken connections. For (i), most detectors issue light indications or audio commands when the battery is low [49]. For (ii), government regulations (such as NFPA 72 [15]) mandate the number and position of smoke alarms to be installed in a given building. For (iii), modern smoke detectors annunciate loose or broken connections, which is required by NFPA 72 [15, 24].

Detecting damage or faults in smoke detectors' phototransistors remains an unsolved problem, despite great attention to improving

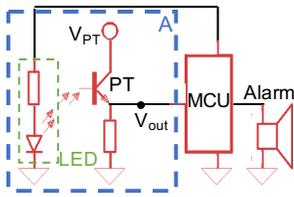


Figure 2: Circuit diagram of a smoke detector.

smoke detector reliability. This is because detector with faulty components generally does not output data that differs obviously from that with a functioning component—until an actual fire occurs (see Section 2). Phototransistors (PTs) are a key problem area because they are a major cause of smoke detectors’ failure to alarm, as we verify in Section 2.1. While fault detection techniques exist for other smoke detector components like the LED [39], techniques for phototransistors remain highly limited [14]. PTs present a particular challenge for fault detection because they do not consume power, making other common fault detection techniques inapplicable.

PTs undergo two classes of faults: *drift*, in which PT responsiveness to light is reduced, and *catastrophic faults*, in which the PT no longer responds to light at all. Wear and tear of PTs due to continuous operation and deployment for several years, as well as exposure to harsh conditions (high temperature in case of a fire event, humidity, dust, etc.) produces drift that gradually increases over time, and may eventually cause a catastrophic fault [13, 31]. As the PT experiences drift, the smoke detector’s sensitivity is reduced, leading to delayed detection of smoke [31].

Faults or drift typically go undetected for several months at a time, as the detector is only triggered during a fire event or maintenance, both of which occur infrequently [14]. Thus, drift or faults may occur “silently” - i.e., without warning - which can be disastrous in a real fire [30]. Existing methods to test the PT’s health manually or automatically are incapable of testing for drift, as we show in Section 3. New, automated detection techniques are needed to improve smoke detector fault tolerance and reduce fire risk. This paper presents a novel approach to detect catastrophic faults and drifts in smoke detector PTs. We present experimental evidence for a key finding in phototransistor physics that the root causes of PT drift also increase the time required for the PT to switch from a “high” to a “low” state (its fall time), resulting in a strong correlation between the fall time and PT drift. We use this finding to devise a phototransistor “fingerprint” we name *FallTime* that differs between working and faulty sensors, and can be used to quantify drift.

The *FallTime* fingerprint has the following key characteristics: (i) *Distinct* for a working, drifted and faulty component; (ii) *Environment-agnostic* and robust to variations in temperature, humidity and dust accumulation; (iii) *Quantifies the amount of drift*; (iv) *Requires no or minimal hardware modification* to existing smoke detector models and can be automatically measured using lightweight APIs. A baseline fingerprint is first determined “in the factory” before deployment to characterize a working component. The fingerprint is measured periodically in a deployed device and compared against the baseline. Depending on the level of deviation, the component can either be classified as working, catastrophically

faulty, or drifted. If it is drifted, the quantity of drift can be estimated to apply appropriate compensation so that the detector is as sensitive as, and alarms in a fire as quickly as, a functioning detector.

To show the efficacy of the proposed fingerprints, we first show that the fingerprint effectively detects faults and quantifies drift in two commercial smoke detector models, namely, *Firelite detector* [29] (12 instances) and *System Sensor* [1] (9 instances). We also built three new smoke detectors (Custom1- 6 instances, Custom2- 7 instances, Custom3- 6 instances) using components commonly used in commercial smoke detectors, and showed that our fingerprinting-based approach can be applied to these components. Finally, we test our approach with real smoke, using the fingerprint to compensate for drift in sensors with drifted PTs. We show that the detectors with our approach operate within desired bounds despite PT drift, while smoke detectors employing the conventional approach frequently fail to alarm before reaching lethal concentrations of smoke.

As an alternative to the fingerprint approach, we also present a mechanical *DriftTestButton* that can be incorporated into a smoke detector’s physical body to detect faults and drift in its components. Intuitively, if existing approaches fail to predict the device’s behavior in Smoke conditions, then a method to simulate the behavior of smoke offers a promising and more robust test. Our *DriftTestButton* protrudes into the sensing chamber when activated, reflecting light from the LED into the PT in the same way that smoke would. We test the proposed button with live smoke and show that it, too, is more effective than the conventional approach in compensating for drift. The key contributions are:

- We show that phototransistor drift is an understated contributor to smoke detector failure, and identify a phenomenon in phototransistor physics in which phototransistor fall time is related to its drift.
- We develop a novel “fingerprint” exploiting this phenomenon to measure and compensate for phototransistor drift.
- We design a mechanical *DriftTestButton* that simulates smoke conditions to measure and compensate for drift.
- We show that these two approaches outperform existing methods with five smoke detector models in real Smoke conditions.

2 UNDERSTANDING SMOKE DETECTOR FAULTS

In this paper, we focus on faults that directly threaten fire safety, or cases in which a fire is not detected quickly, or not detected at all, and the device fails to alarm quickly enough for people to evacuate from the fire.

A major cause of failure that has received inadequate attention is *sensitivity drift*, which is caused by “wear and tear” or aging of sensing components, and is highly prevalent in smoke detectors [13, 31]. Specifically, phototransistors (PTs) accumulate damage over time through hot carrier injection, which reduces the device’s forward-bias current gain [18, 36, 61]. This presents a substantial challenge because smoke detectors are expected to last for 10 years of continuous operation, or about 88,000 hours, so even relatively minor levels of drift per year will pose safety risks, compromising the “margin of safety” for a smoke detector to alarm in time for

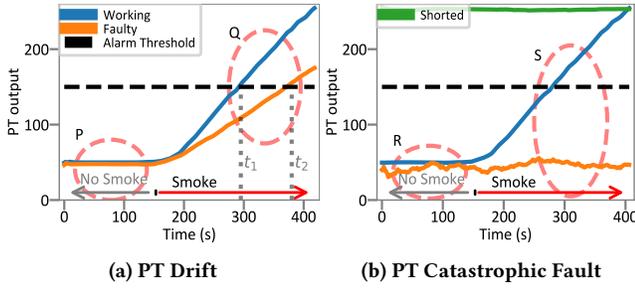


Figure 3: Faults observed in smoke detector PTs

people to evacuate [30]. Yet, methods to measure drift are lacking [14].

To illustrate the effects of various faults, Figure 3 compares the PT output of working sensors with that of faulty sensors. The x-axis indicates the time in seconds; smoke is introduced and begins accumulating after 150 seconds of data accumulation. The black horizontal line represents an alarm threshold.

Under NoSmoke conditions the PT of a working sensor does not directly receive light from the LED, resulting in a *low* PT output as shown in Figure 3(a) region P. When smoke enters the sensing chamber, light emitted by the LED is scattered by the smoke particles and the intensity of light incident on the PT increases, which is converted into an electrical signal, resulting in a *high* output, as shown in Figure 3(a) region Q. Crucially, higher densities of smoke will scatter more light and increase the PT output. When the output of the PT is above a set threshold, an alarm is triggered.

Case 1: Drift: Figure 3(a), region Q, illustrates that PT drift in a faulty sensor causes a delay in alarm. The PT must wait for higher smoke densities before crossing the alarm threshold and activating an alarm. Thus, a working sensor alarms at time t_1 while a drifted sensor alarms later, at t_2 . This threatens fire safety.

Case 2: Catastrophic faults: Some PT faults can cause the output to “float” (or take on arbitrary values due to being electrically un-grounded) in the normal output range (region R), without rising in the presence of smoke (region S, Figure 3(b)).

Key challenge: *The PT output of a drifted or faulty PT component is not distinguishable from the output of a working sensor in NoSmoke conditions (regions P, R) and deviates only in smoke conditions (regions Q, S).*

While LED faults or drift also cause failure, existing approaches can already detect it, as we describe in Section 3. *Conversely, the effect of PT drift/fault remains poorly studied.* In the next section, we address this gap by conducting natural and controlled experiments that demonstrate PT drift is highly prevalent and impactful on smoke detector operation, highlighting the need for drift detection techniques for PTs. *The tests are conducted with sensors already using a conventional, industry-standard fault compensation approach, indicating that existing approaches fail to correct for drift.*

2.1 Case study: Impact of PT drift in smoke detectors

Natural experiments. We collected 12 *Firelite detector* (see Section 6.1 for details), one of which was brand-new and 11 of which had been in use in a commercial building for 1 - 5 years, compared to the detectors’ rated lifetime of 10 years. To determine whether

their PTs were drifted, we measured the response of the PT to varying light intensities. An LED [46] is shone straight into the PT from 2 cm away, in the absence of ambient light, and the PT has a 15K Ω load resistor.

Figure 4(a) shows the response for four of the detectors’ PTs (collector current in mA) to increasing light intensities (measured in Watts per steradian, W/sr), which corresponds to increased smoke densities. Increasing the light intensity increases the PT output, as expected. However, the PT output at any given light intensity varies considerably between PTs, indicating drift. For instance, the sensor corresponding to the bottom orange line would require higher smoke densities than the sensor corresponding to the top blue line, to cross the alarm threshold.

To quantify the level of drift, we define a “drift score” ranging from 0 to 1 for each PT. The score quantifies how far the behavior of a drifted PT deviates from a new PT, and is proportional to the area between the curves for a new PT and for a drifted PT in Figure 4(a). It is defined as follows:

$$D_{PT} = 1 - \frac{\sum_{i \in I} PT_{test}(i)}{\sum_{i \in I} PT_{new}(i)}, \quad (1)$$

where $i \in I$ represents the set of light intensities tested, $PT_{test}(i)$ and $PT_{new}(i)$ represents the PT’s output at light intensity i for the PT under test and the new PT, respectively. The new PT has a score of 0 and a higher score, up to 1, indicates more drift. Figure 4(b) uses a bar-and-whisker plot to show the distribution of drift scores, between 0 to 0.3 with a mean of about 0.14, for the 12 *Firelite detector* deployed in the field.

To study the impact of PT drift on detector sensitivity, we conducted a series of 20 smoke tests in which the *Firelite detector* is exposed to an increasing density of smoke. We recorded the time and the smoke density at which the detector alarms. We conducted the test with five of the 12 detectors in our smoke chamber setup (see Figure 17(a)) because each must be connected to a Fire Alarm Control Panel, limiting the number of detector used in each test.

Figure 4(c) shows the mean results from all the smoke tests. The x-axis shows the drift score. The left y-axis shows the mean smoke density (in obscuration %/ft [17]), at which each detector alarmed. The right y-axis shows the time at which each detector alarmed, normalized so that the fastest sensor to alarm has a value of 1. The drift score correlates positively with smoke density and time required to trigger the alarm. Detectors whose PTs had more drift needed higher smoke densities and more time before alarming. For example, detectors with drift scores between 0.2 and 0.3 (rightmost bar) alarm at smoke densities 10% points higher and 1.5 times slower than fully functional detectors (leftmost bar). *This represents a dangerous delay in alarm in the drifted sensors.*

We compute the coefficient of determination for the correlation between drift and density level, and between drift and the time at which a detector alarms, across all 20 smoke tests. The drift score explains 67% of the variance in both measures. The impact is also very substantial: a single standard deviation increase in the drift score prolongs a detector’s time taken to alarm by 16.4%, and increases the smoke density level at which it alarms by 3.94%/ft, which is dangerously high [21].

Controlled experiments. To strengthen the above finding, we conducted experiments on 9 newly purchased commercial smoke detectors, *System Sensor* [1]. We manually drifted the PTs in all

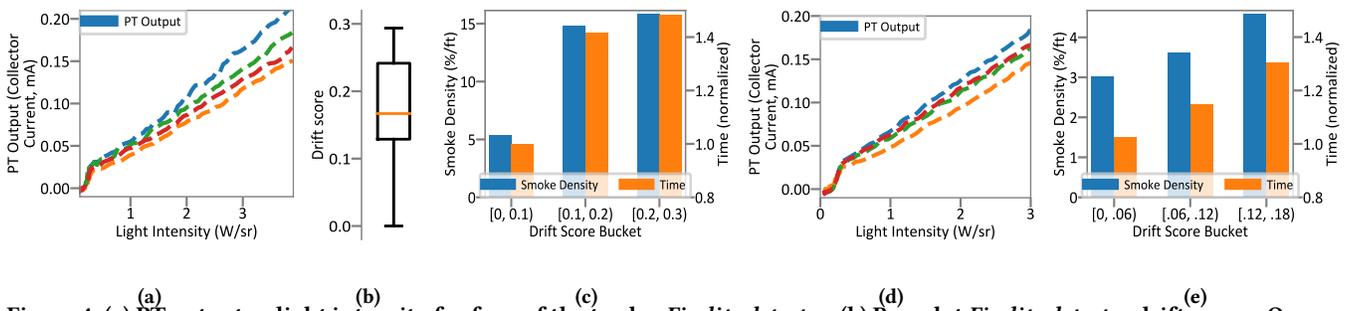


Figure 4: (a) PT output vs light intensity for four of the twelve *Firelite detector*, (b) Box plot *Firelite detector* drift scores. Orange line, box edges, and whisker edges mark the mean, 25th/75th percentiles, and min/max of the drift scores, respectively, (c) Smoke density and time to alarm for drifted *Firelite detectors*, (d) PT output vs light intensity for four of the nine *System Sensor*, (e) Smoke density and time to alarm for *System Sensors*.

but two of the *System Sensor* using the procedure explained in Section 6.2. Figure 4(d) shows the PTs’ responses for 4 detectors to different light intensities. The drifted PTs’ (orange line) have a lower response than the working PT (blue line). Further, the manually drifted PTs’ responses from *System Sensor* closely matches the responses of the naturally drifted PTs from *Firelite detector* (see Figure 4(a)). In this controlled experiment, the PTs are the only components that experience degradation while the remainder of the circuit remains intact. *This allows us to isolate PT drift as the single cause of delayed alarm.*

We conducted another series of 20 tests in which the 6 *System Sensor* are exposed to an increasing smoke density, and recorded the time and smoke density at which each sensor alarms. Figure 4(e) shows the results of smoke tests on *System Sensor*. As before, detectors with drifted PTs need much higher smoke density levels and time than functional detectors to alarm. In this case, PT drift explains 54% of the variance in the time and smoke density level at which a detector alarms. A modest increase in the PT’s drift score by 0.1 points increases the time to alarm by 18% and the smoke density at which a sensor alarms by 1%/ft.

Finding: *Phototransistor drift has a substantial effect on smoke detector’s sensitivity, and greatly reduces the detector’s speed in responding to smoke.*

These results are notable because it is typically assumed that LED drift accounts for the majority of sensitivity drift, as in optocouplers [2]. As a result, research on smoke detectors has devoted limited attention to phototransistor drift. We show that even modest PT drift can have a significant impact on device performance, even when no other part of a smoke detector has been changed. Drift detection for phototransistors is thus a critical part of the overall reliability of smoke detector circuits.

3 EXISTING FAULT DETECTION METHODS

In this section, we highlight the shortcomings of existing fault detection methods.

Manual approaches: Many smoke detectors have a “test button,” but this merely closes a switch to activate an alarm, testing only the siren and not the sensing components [49]. Other test buttons activate a test mode that employs the *IndustryApproach* [44], which we address below. Another common approach is the use of “canned smoke” or aerosols that can be sprayed into the sensing

chamber to mimic smoke [47]. This cannot detect drift because the amount of aerosol entering the detector cannot be precisely controlled. It is thus discouraged as an effective testing mechanism by various fire protection agencies [35, 47, 54]. In general, manual intervention is disfavored because it severely limits the frequency of testing.

New hardware signature-based approaches: New techniques monitor the current consumption of the sensor component to determine if it is working or faulty [38], [39]. These approaches are designed for high-power-consuming components such as LEDs. However, the PT in the smoke detector does not consume significant power, so its current consumption characteristics are not affected by faults or drifts. Thus, the current-monitoring approach is unable to tackle PT faults and drift. [19] uses the shape of a sensor’s “fall curve,” or its voltage output after the device is switched off, to detect faults in simple analog sensors. This technique only detects catastrophic faults and classifies devices as either working or faulty. As such, it cannot detect or quantify drift in phototransistors. Quantifying PT drift is essential in the context of smoke detectors because drift substantially delays the time at which a detector alarms (Section 2.1). Thus, new fault tolerance techniques are needed in our context.

The existing Industry Standard approach: A common approach implemented in commercial smoke detectors, which we call the *Industry Standard Approach* or *IndustryApproach*, uses the response from the PT in the NoSmoke condition to estimate drift. Most “intelligent” smoke detector models available in the market use this approach [1, 28, 29, 34], and it is the leading method cited in NFPA 72 [15], recent publications [10], and patents [20]. This method saves a baseline PT output, and compares the measured PT output (typically over several days) in “normal” smoke conditions against the baseline. A response below the baseline indicates the PT is drifted, and corresponding gain is applied to the PT to offset the drift or to issue a warning to the user.

We argue that this fault detection scheme has inherent limitations because the output of the PT in low light intensities (i.e., NoSmoke) cannot accurately estimate its behavior at higher light intensities (i.e., Smoke). Intuitively, Figure 4(a) and (d) show that drifted sensors do not have a significantly different output at low light levels (i.e. lines are very close together on the left side of the graph), even if they have substantially different outputs at high light levels (lines diverge on the right side).

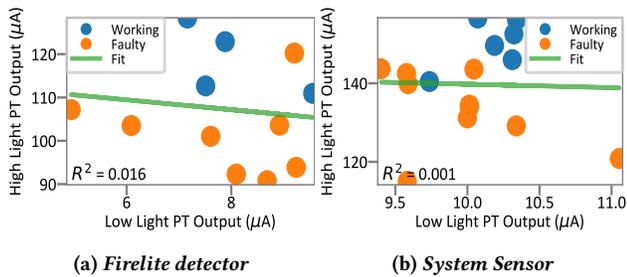


Figure 5: Fit between High and Low Light Intensities. Value at low light intensity is a poor predictor of value at high light.

To quantify this claim, we train a simple linear model to estimate the PT’s output in high light intensity (Smoke) using the PT’s output in low light intensity (NoSmoke). We used 12 instances of *Firelite detector* and 6 instances of *System Sensor* with components of varying drift scores. The low light intensity is chosen to match the typical PT output in NoSmoke and the high light intensities is in the PT’s active range at 1.5 W/sr. We take roughly 100 measurements at each intensity.

Figures 5(a), 5(b) shows the result. Each point represents data from one instance of one PT. Each point is positioned according to the PT’s response at low light intensity (x-axis) and at high light intensity (y-axis). Blue dots correspond to working sensors and orange dots correspond to drifted sensors. Visually, drifted sensors (orange) have lower high-light responses (lower on the y-axis) than working sensors, but do not appear to have distinguishable low-light responses (x-axis). Concretely, the fit (green line) is very poor, with r^2 values [16] of 0.016 for *Firelite detector* and 0.001 for *System Sensor*.

We repeat this test at six different high light intensities ranging from an intensity corresponding to a smoke density of 1%/ft, to the intensity that saturates the PT output. The same pattern holds for these trials for both *Firelite detector* and *System Sensor*, with r^2 [16] never exceeding 0.07. This implies that the PT output in low light is a very poor predictor of PT output in high light intensities. It is not possible to accurately distinguish working from faulty sensors based on their low-light output alone. IndustryApproach is thus incapable of diagnosing drift in these sensors.

Table 1 gives an overview of the detection capabilities of a few popular smoke detector models. While the datasheet of *Firelite detector* and *System Sensor* claims to perform drift compensation, it did not appear to work in our real-world tests. Both the *Firelite detector* collected from the live deployment and the *System Sensor* detectors whose PTs we manually drifted failed to issue a warning about the condition of the detector, despite the fact that their PTs were substantially drifted. Further, as noted in Section 2.1, we conducted smoke tests on these sensors with drifted PTs, and the drift compensation algorithm provided by the manufacturer did not adequately offset the drift. The drifted detectors still alarmed substantially later, when smoke reached dangerously high density levels of 10%/ft and higher.

4 SENSOR FINGERPRINTING APPROACH

In this section, we propose a “fingerprint” that captures the electrical characteristics of the PT. As described in Section 3, traditional

Detector Model	LED Fault	Led Drift	PT Fault	PT Drift
GST [4]	No	No	No	No
First Alert [3]	No	No	No	No
Kidde P9050 [5]	Yes	No	Yes	No
Firelite SD355 [29]	Yes	No	Yes	No*
System Sensor 3150E [1]	Yes	No	Yes	No*

Table 1: Detection capabilities for popular detectors.

*The datasheet claims to use drift compensation; our experimental data show that it has limited effectiveness.

techniques monitor the PT output to estimate drift, which is ineffective. The FallTime fingerprint estimates drift more accurately by monitoring the PT’s intrinsic electrical characteristics. Our experiments also reveal fundamental insights into phototransistor physics, and contribute to the field’s understanding of how phototransistors age.

Central idea: The FallTime fingerprint measures the time required for the PT output to fall from high to low after the base input current (generated by light incident on the PT’s) falls to zero as the incident light is switched off. The length of the fall time is used to determine the PT’s “health,” or classify if it is working, faulty, or drifted, and estimate the quantity of drift. Due to parasitic capacitance in the PT, the output signal takes a few microseconds to decrease to zero after incident light is switched off [19, 45]. Our approach measures the time taken to fall from 90% of its high-state value to 10% of that value [59] (“90-10% fall time”). The fingerprint measured in the field is compared to a baseline, which is measured “in the factory” for a fully functional sensor of the same model, to determine the sensor’s health. For instance, we collect a fingerprint only once for a fully functional PT model, which can then be used as a baseline fingerprint for all individual PTs of the same model type. Further, the fingerprint collection can be easily implemented with lightweight firmware updates to existing commercial smoke detector models that can support a sampling rate adequate to measure FallTime (updates typically must be made by the manufacturer, not by the user), or can be implemented with minimal hardware changes to smoke detectors that do not support such sampling rates as described next.

4.1 FallTime characteristics

We now describe FallTime’s key characteristics.

① Distinct for working, drifted and faulty components.

Faults/drift in a PT result from physical degradation of the device, which also changes its FallTime. FallTime can therefore be measured as a proxy for drift/faults. Figure 6(a) shows the voltage response of two working and three drifted PTs. The graphs show the raw output values. The two working PTs have similar voltage response and FallTime, but the FallTime of a drifted PT or a faulty PT is distinct from that of the working PT. Deviation in the length of a sensor’s FallTime can be used to distinguish working from drifted PTs. Faulty PTs’ outputs may “float” in the active region as shown in Figure 6(a), and can easily be distinguished by their irregular fall times.

② Environment-agnostic.

The FallTime of a PT is agnostic to environmental changes, such as temperature or humidity. In all smoke detector designs, the PT is housed in a closed chamber, which blocks ambient light and reduces the impact of temperature and

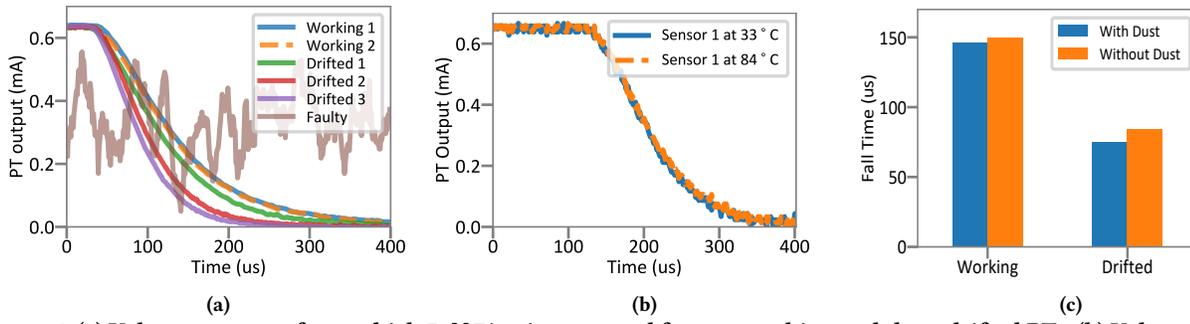


Figure 6: (a) Voltage response from which FallTime is measured for two working and three drifted PTs, (b) Voltage response at two different temperatures, (c) Fall time evaluation on working/drifted detectors with and without dust.

humidity changes. FallTime is also robust to those changes. We measured FallTime fingerprints of a working PT in two different temperature conditions. Figure 6(b) shows that across different temperatures, the FallTime remain same.

FallTime is also operable in the presence of dust. Dust accumulation typically increases the reflectivity inside the smoke chamber. Because FallTime measures fundamental device properties rather than the reflectivity (as explained in the next subsection), it is more robust to dust than IndustryApproach, which measures reflectivity in order to infer device properties. As shown in Figure 6(c), dust affects the fall time by a much smaller degree than even modest levels of PT drift (drift score = 0.34). The presence of dust changes the fall time negligibly for either a working or drifted sensor, so FallTime can quantify drift even in the presence of dust.

③ **Quantifies drift.** The central characteristic is that as a PT experiences drift, its FallTime correlates well with its drift score. Figure 7(a), (b) shows this result for PTs from Custom1 and Custom3 detectors (see Section 6.1). As the PT’s drift score (x-axis; see Eq. 1) increases, the FallTime (y-axis) decreases with strong correlation, shown by the best fit line (r^2 scores are given in the graphs’ corners). Measuring the FallTime allows us to quantify the level of drift, and then compensate for that drift so that the detector alarms in the presence of smoke as quickly as a working one.

④ **No or minimal hardware modification required.** To measure the FallTime, we continue to sense the PT output for a short period of time after incident light is switched off. Smoke detector circuits pulse the LED on and off, so the PT fall time can be measured without changing the sensor’s basic operation.

Measuring FallTime also requires sampling the PT’s analog output value at a rate of at least 1 megasamples per second (1Msps). Many smoke detector models integrate microcontrollers with this capability [53, 56], or simpler microcontrollers with separate 1Msps sample-and-hold analog-to-digital converters (ADC) [55]. In such cases, fingerprint collection can be easily implemented with lightweight firmware updates. Smoke detectors lacking these components must make one of three hardware modifications to collect FallTime fingerprint, all of which can be made with minimal cost overhead (typically in the range of a few cents to US\$1) and intrusion: (i) replacing the microcontroller with a newer class of microcontrollers with a >1Msps ADC sampling, (ii) adding a separate >1Msps ADC with sample-and-hold capability, or (iii) use of a larger load resistor to the photo-sensing circuit to greatly increase

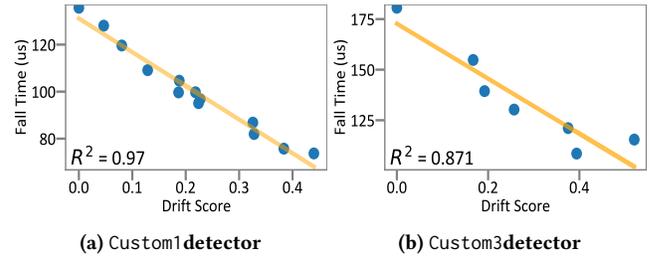


Figure 7: Fall time vs drift score for different PTs.

the fall time, thus allowing existing slower-sampling ICs to still measure FallTime.

4.2 Underlying physics and insights into PT aging and lifecycle

We now describe the physics underlying the change in FallTime. FallTime results from parasitic capacitance in the PT component. The base current will charge the parasitic capacitance between the base and emitter when it is switched on, and when switched off, the capacitor will discharge into the emitter region, prolonging the FallTime. It is well-known that phototransistors accumulate damage over time from hot-carrier injection [11, 25, 57, 58], degrading their current gain. We argue that this also reduces base-emitter diffusion capacitance in NPN phototransistors [27, 33, 61]. Lower capacitance (due to aging) will reduce the time required for it to discharge, thereby reducing the fall time.

To validate and generalize this result, we use a circuit simulator to characterize PT behavior. We model the PT as an ideal transistor with parasitic capacitances between its junctions, and use a current source to simulate incident light, which injects current into the base of a real PT [32]. Figure 8(a) shows the circuit diagram used in our simulation. We matched parameters to the datasheet of a real phototransistor (Vishay BPV 11 [51]) where available, and experimentally derived other values. We simulate our hypothesized effect of hot-carrier injection by decreasing C_{BE2} from 2 nF to 20 pF, and feed the circuit to a SPICE simulator [6, 48]. We compared the simulated voltage response and FallTime of a working sensor against that of a drifted sensor (with C_{BE2} reduced).

Figure 8(b) shows the voltage response of working and drifted sensors, from both simulated and real PTs. Each graph shows when the base current is turned off for reference (orange line). The vertical dotted lines in red indicate the beginning and end of the FallTime

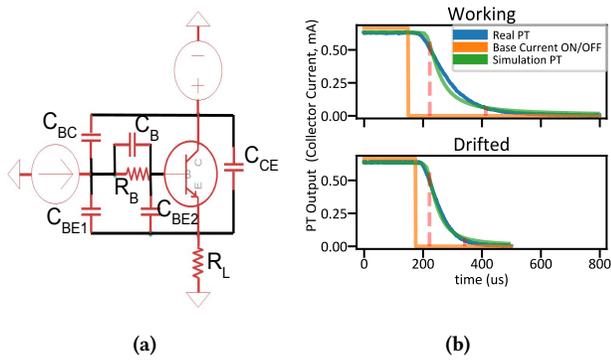


Figure 8: Simulation results. (a) PT circuit used in simulation. C_{BC} : 19pF, C_{BE1} : 1.5nF, C_B : 500nF, R_B : 30 Ω , C_{BE2} : 2nF, C_{CE} : 15pF, R_L : 15K Ω . (b) Voltage response for simulated & real PT.

(90–10% of the high-state value). The voltage response and FallTime for the simulated and real PT match closely in both working and drifted conditions, confirming that reduced C_{BE2} models the effect that degradation has on a PT’s fall time.

This finding is important because it ❶ strengthens our understanding of the physical causes of phototransistor degradation, which can contribute to designing chips with longer lifetimes [42], ❷ identifies a hitherto unknown effect as phototransistors degrade, which can help improve the long-term performance of systems sensitive to changes in phototransistor switching times, and ❸ can be used for fault detection in the same manner as this paper in the hundreds of other applications of phototransistors [43].

4.3 Fault tolerance algorithm

Once FallTime is measured, we use the following procedure to compensate for drift and detect faults. First, “in the factory,” or before the sensor is deployed, the manufacturer models the relation between the fingerprint and drift score for a given smoke detector model type. This relationship is modeled generically for a given PT part number, and applies to all PTs of the same part number, rather than being measured for each individual PT. The relation is modeled simply by selecting m and b to fit a line $k_{FallTime} = m \cdot d_{PT} + b$ for a set of PTs of the same model with varying drift scores, as shown in Figure 7. d represents the PTs’ drift scores and k represents the PTs’ FallTime. b corresponds to the “baseline,” or the fall time of a fully functioning sensor, and m is the rate at which fall time decreases relative to drift score.

When a given detector is deployed in the field, we periodically measure the FallTime fingerprint, and use the linear model above to obtain \hat{d} , an estimate of the sensor’s drift score, without needing to directly measure it. To avoid applying gain to “false positive” cases of drift that reflect noisy measurements instead of a genuine decrease in FallTime, we only apply this procedure if the measured fall time is below 2 standard deviations of the baseline fall time. We take the mean of several FallTime measurements to reduce the impact of measurement error.

The gain required to offset a given drift score follows from Equation 1 and results in the gain:

Algorithm 1: Sample smoke detection and drift compensation routine

```

static m, b;                                ▶ set in the factory
static alarm_threshold;                    ▶ set in the factory
static tolerable_margin;                  ▶ set in the factory
gain = 1;
Do once a month:                          ▶ drift compensation routine
    turn_LED_on();
    turn_LED_off();
    k = read_PT_fall_time();
    if abs(k - b) >= tolerable_margin then
        | warn_user();
    else
        | gain = m / (m - k + b);
Do every 15 seconds:                      ▶ smoke detection routine
    turn_LED_on();
    output = read_PT_output() * gain;
    turn_LED_off();
    if output >= alarm_threshold then activate_alarm();
    
```

$$G = \frac{m}{m - \hat{k}_{FallTime} + b} \quad (2)$$

where $\hat{k}_{FallTime}$ represents the FallTime measured in the field. G is applied to the PT output (in software) so that it matches the output of a sensor in good condition, in order to compensate for drift. The FallTime fingerprint should be measured periodically (e.g., weekly or monthly) to adjust the gain as drift increases over time.

Algorithm 1 shows our procedure to update the gain from FallTime measurements alongside the procedure for smoke detection. The gain is set to 1 for a fully functional detector.

Catastrophic fault detection: Detectors with measured fall times outside of a “tolerable margin” can be classified as having catastrophic faults, which cannot be compensated for. The margin is defined by experimental evidence on the bounds for FallTime. For instance, in Figure 7, Custom1 sensors have a minimum fall time close to 80 μ s. Sensors never reach a lower fall time or greater drift score without undergoing a catastrophic fault. We set the tolerable margin accordingly. An audio warning (like low battery warnings found in most detectors) can be issued to the user in this case.

5 MECHANICAL DRIFTTESTBUTTON TO TEST SENSOR COMPONENT HEALTH

This section exhibits a novel design for a physical test button, independent of the previous fingerprinting approach, that enables simple and robust fault and drift diagnosis and compensation.

The basic shortcoming of the Industry Standard Approach (described in Section 3) is that the output of the PT in NoSmoke does not predict the output in Smoke. Intuitively, a method to simulate the behavior of smoke offers an elegant and robust test of sensor fault and drift. Thus, as an alternative to FallTime, we present a mechanical DriftTestButton that precisely simulates the behavior of smoke.

The proposed DriftTestButton protrudes into the sensing chamber when activated/pressed, reflecting light from the LED into the

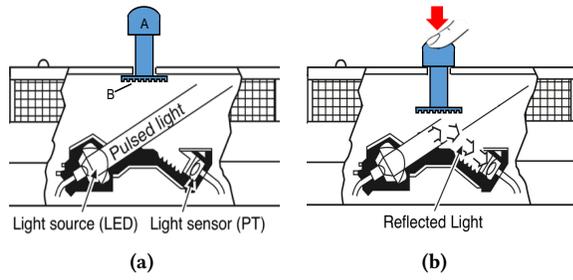


Figure 9: Schematic of the DriftTestButton: (a) Button Inactive/Default, (b) Button Activated/Pressed.



Figure 10: DriftTestButton integration with physical smoke detectors: (a) Without DriftTestButton (b) Button Inactive/Default state, (b) Button Activated/Pressed state.

PT in the same way as smoke. The button reflects a precise, pre-determined amount of light into the PT to mimic a smoke density close to the sensor’s chosen alarm threshold. When a user periodically activates the button, the device compares the PT output against a baseline (measured before deployment) to test for faults and drift in both the LED and PT. Unlike existing “canned smoke” techniques, which cannot detect drift (Section 3), our technique can detect and measure drift because the amount of light reflected is precisely controlled. This solution is particularly robust because it tests the entire sensing pathway at once, and essentially *measures* drift directly rather than estimating it. It thus obviates the need for any alternative fault detection techniques.

Previous button designs differ from our design in several ways. Unlike [40], our button is designed to reflect a pre-determined amount of light into the PT, rather than an uncalibrated, unspecified amount. Unlike [23], which focuses solely on detecting faults, our DriftTestButton measures and compensates for drift in addition to faults. We also provide experimental results to support the effectiveness of the DriftTestButton in Smoke conditions.

5.1 Physical design

In our DriftTestButton design, the button hangs down from the smoke detector’s chamber and protrudes through the detector’s plastic casing (Figure 9(a)). The button’s “head” (Figure 9(a) and 11(a), label B) is attached to the top of the button, and is raised towards the PT and LED when the user presses the button. We 3D printed the test button; Figure 11(a) shows the assembly of the button in CAD and Figure 10 shows the actual DriftTestButton integration with *Firelite detector*. We now describe its key characteristics:

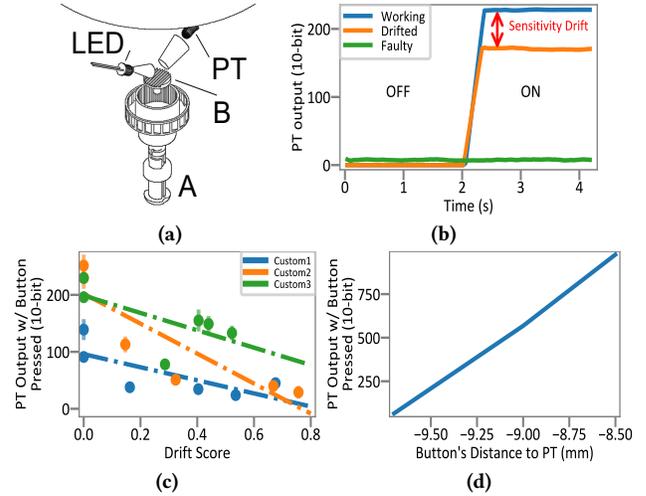


Figure 11: (a) Exploded view of button, (b) PT output when button is activated, (c) Response of button to drift, (d) Relationship between PT output and button position when activated.

❶ **DriftTestButton simulates smoke by increasing reflectivity in the sensing chamber.** Smoke reflects light from the LED into the PT. To simulate this, the button also reflects light from the LED into the PT, increasing the PT output, as shown in Figure 9(b). However, if the reflectivity is too high, the PT output saturates at its high value whether the components are drifted or working, giving a false indication that the sensor is working. Therefore, the button must avoid saturating the PT output. To accomplish this, the head of DriftTestButton has a grated, black surface; we verified that this keeps its reflectivity well below saturation.

❷ **DriftTestButton detects faults and quantifies drift.** To detect faults and drift, we can compare the PT’s response (when the button is pressed) against a “baseline,” or the response expected for a working sensor, as shown in Figure 11(b). The level of deviation is used to quantify drift. In Figure 11(c), we show the button’s behavior in relation to drift. Each point represents one phototransistor, with its drift score on the x-axis and the response it gives when the button is pressed on the y-axis. The graph shows that greater drift decreases the PT response when the button is pressed. When the button is activated, the PT response is affected by the cumulative effect of PT and LED drift, allowing us to measure and compensate for both drifts simultaneously, without needing to isolate the faulty component.

❸ **DriftTestButton has low measurement error.** In order to measure drift precisely, with low error, the button must produce a consistent response to a fixed amount of drift. To verify this, we used a 3D printer to produce 5 copies of the DriftTestButton. We tested each button’s behavior across 100 button presses on 17 different PTs (of three different models, Custom1, Custom2, Custom3: see Section 6.1) with varying levels of degradation. Figure 11(c) shows the results of this trial. Each point represents the mean response of one PT across 500 total button presses (100 presses for each of 5 buttons). The PT output is shown on the y-axis, and the standard deviation for each PT across 500 presses are shown as error bars. We can observe that the standard deviation is very small or negligible. The drift we

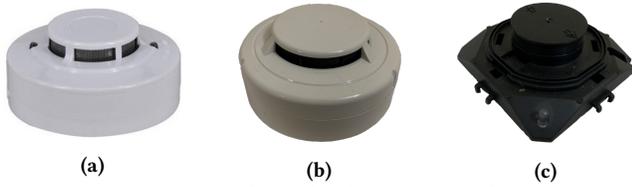


Figure 12: Commercial smoke detectors used in our experimental setup: (a) Firelite detector (b) System Sensor (c) Custom

seek to detect is typically at least an order of magnitude larger than this deviation, indicating that the button gives a response consistent enough to measure drift precisely.

④ **DriftTestButton design can be easily adapted to fit any smoke detector.** The `DriftTestButton` can easily be integrated into existing or new smoke detector designs. As noted above, the `DriftTestButton` must reflect a controlled amount of light into the PT, which will depend on the LED and PT model selected. To adapt our design, a smoke detector manufacturer should select a desired PT response within the PT’s active region. To match that value, they may increase or decrease the length the button travels upward. In Figure 11(d), the x-axis shows button’s displacement, or its position when activated, relative to the surface of the PT. Lower (negative) values indicate the button’s surface is further away from the PT, and 0 indicates the button’s surface touches the PT. The y-axis shows the PT response. Designing the button so that it is closer to the PT when pressed increases the PT response. If a manufacturer’s circuit requires the button to produce a particular PT output, the manufacturer can edit the button’s design accordingly, by displacing the button closer or further to the PT.

Further, it is straightforward to design a button that is durable and will perform consistently after years of wear and tear. Our design does not require deformation of the button’s plastic, which is the primary cause of failure in such designs [60]. Plastic buttons are typically designed for a lifetime of 1 million presses; routine monthly use of our test button would require only 120 presses to outlast the typical ten-year lifetime of a smoke detector.

5.2 Button operation

Before deployment, “in the factory,” the PT response when the button is pressed is recorded as a baseline. In the field, pushing the button closes a switch that activates a test mode in the software. In test mode, the PT output is compared against the baseline to detect faults or detect, measure, and compensate for drift. This is done in exactly the same manner as in Algorithm 1, using the button response rather than `FallTime` to measure drift. Figure 11(b) shows the PT output of a working and a drifted sensor before and after the button is pressed. The working sensor represents the baseline. If the response matches the baseline, no changes need to be made. If the response is substantially lower than the baseline but within a tolerable bound, the detector can be diagnosed as drifted and gain can be applied to the PT output so that its output matches that of the working sensor. If the PT response is outside of a tolerable bound, the sensor can warn the user of a fault. While this test method is a manual process, it is far more precise and robust to error than other methods.

6 EXPERIMENTAL SETUP

6.1 Smoke detectors

We test our proposed fault detection and drift compensation methods with 5 different smoke detector models, two commercially available and three custom-built.

Commercial detectors. We selected two popular commercially available detectors: (i) 12 *Firelite detector* Model SD335 [29], 11 of which had been deployed in a commercial building for 1-5 years (Figure 12a), and (ii) 9 *System Sensor* Model 2351E [1], each purchased brand new (Figure 12b).

Custom-built detectors. The commercial detectors do not enable re-programming of the microcontroller by the user, making it impossible to implement the proposed approaches directly on these detectors. Hence, we built three custom smoke detectors circuits. The circuit design includes an LED and PT placed inside a sensing chamber obtained from *Firelite detector* to ensure the physical setup mimics commercial designs (Figure 12c). Each of the three custom smoke detectors uses a different PT model and the same LED. All chosen components are commonly used in smoke detector applications. The LED model is SFH 4550 [46]. The three PT models are BPV 11 [51], BPW 96 [50], and OP 505 [26], forming the *three custom detectors*, Custom1 (6 detectors), Custom2 (6 detectors), and Custom3 (6 detectors), respectively. We use an STM32 [9] MCU to control the device.

6.2 Degradation procedure for PTs

Our tests are designed to offset the effects of drift, which results from “wear and tear” of sensor components over time. Since we had access to a limited number of detectors (12 *Firelite detector*) deployed for 1-5 years, we manually degrade the PTs from *System Sensor*, Custom1, Custom2 and Custom3 to simulate the effects of several years of wear and tear (i.e., “Highly Accelerated Life Testing” or HALT). This is a standard technique when designing for device longevity that reliably simulates the conditions devices are exposed to over the course of several years [18, 36, 37, 42, 61].

To degrade a PT, we applied a high reverse bias (beyond datasheet ratings) until substantial current begins flowing and sensor heats up, a common HALT method [18, 36, 61]. This simulates years of natural aging by increasing the mean electron energy to increase the rate of hot carrier injection, which occurs at a lower rate in normal conditions [25]. In order to obtain different levels of drift, we apply high reverse bias for varying periods of time [37]. We validate our degradation procedure by comparing against the 12 naturally drifted *Firelite detector* from the field. The PT response to increasing light intensities of naturally degraded sensors closely matches that of the manually degraded sensors: compare Figure 4(a) (naturally degraded) against Figure 4(d) (manually degraded). We also verify that the shape of the curve representing the PT output when the LED is switched off (see Figure 6(a)) also matches that of manually degraded sensors.

7 RESULTS

We first present the efficacy of the `FallTime` fingerprint to detect faults and measure drift. Then, we test our drift compensation technique with the three custom smoke detector models in tests with actual Smoke.

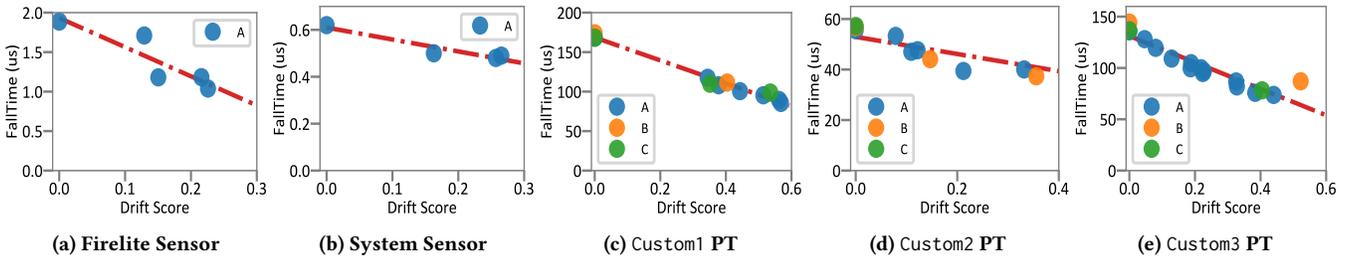


Figure 13: FallTime vs drift score for all smoke detector models.

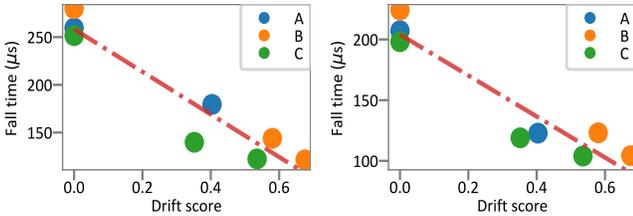


Figure 14: Drift vs fall time at 3.3V and 5V voltage levels.

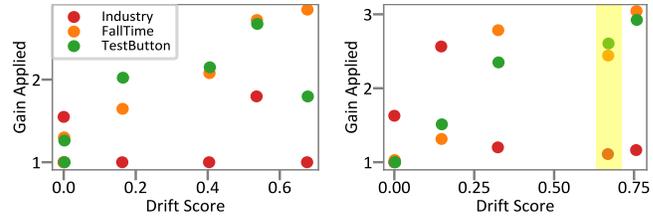


Figure 15: Gain recommended by each compensation algorithm.

7.1 Efficacy of the FallTime fingerprint

We now present extensive results on the relationship between drift and the FallTime fingerprints. We conducted experiments on all five smoke detector models, i.e., *Firelite detector*, *System Sensor*, Custom1, Custom2, and Custom3.

Figure 13(a) shows FallTime for 5 different *Firelite detector*, which are naturally drifted in the field. As the drift score increases, the FallTime decreases. We also degraded the PTs of *System Sensor*, Custom1, Custom2, Custom3 in small increments, measuring the FallTime after each step. We record the mean of 1000 FallTime measurements in each step. Figure 13(b)-(d) shows FallTime of *System Sensor* and three PT instances of each of Custom1, Custom2 and Custom3 detectors. In each graph, each color corresponds to a single sensor instance; there are several markers of each color because drift and FallTime of a single sensor instance are measured at each step as it is degraded in increments. The graph shows FallTime and drift scores are strongly correlated across all sensors (shown with the best fit line). The same pattern holds for three additional sensor instances of each model (omitted to avoid cluttering the graph). R^2 is at least 0.95 for all sensor instances.

To evaluate the consistency of FallTime measurements at different supply voltages, we measured fall time in a Custom1 sensor powered with 5V and 3.3V instead of 10V. Figure 14 shows the drift score and fall time remain highly correlated for both circuits with different supply voltages, viz., 3.3V and 5V.

Result. *The proposed FallTime fingerprint works consistently in a variety of sensors and circuits and is therefore a highly promising approach to measure PT drift accurately.*

7.2 Gain determination

We now evaluate our proposed drift compensation methods. First, we measure the FallTime fingerprints, and the PT output when the DriftTestButton is activated. Then, we compute the gain using the procedure described in Sections 4.3 and 5.2, comparing it against

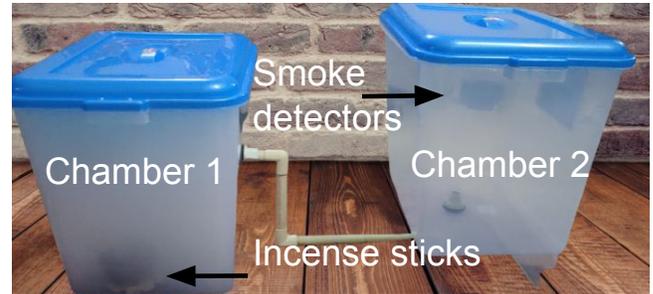


Figure 16: Smoke Chamber Experimental Setup

the IndustryApproach (Section 3). The gain indicates the factor by which the PT output is increased to keep its sensitivity close to that of a working sensor. Sensors with more drift require more gain to offset the drift.

Setup: We used 5 PTs of varying drift scores for each of the three custom detector models. The drift scores range from 0 to 0.8 for each of the three PT models. Along with drifted sensors, we also run each algorithm (IndustryApproach, FallTime, DriftTestButton) on the working sensors to ensure our algorithm does not apply excessive gain to a working PT, which would increase the probability of false alarms.

Figure 15 shows the gain recommended by each algorithm (y-axis) against the drift score (x-axis) for Custom1 and Custom2 PTs. Each point represents the gain applied to one PT and each color corresponds to a compensation algorithm. As the drift increases, the gain needed to offset the drift should also increase. The gains recommended by both FallTime and DriftTestButton (shown in orange and green) are similar, and they increase across all PTs as drift score increases. Conversely, the gain recommended by the IndustryApproach (shown in red) is inconsistent with the increase

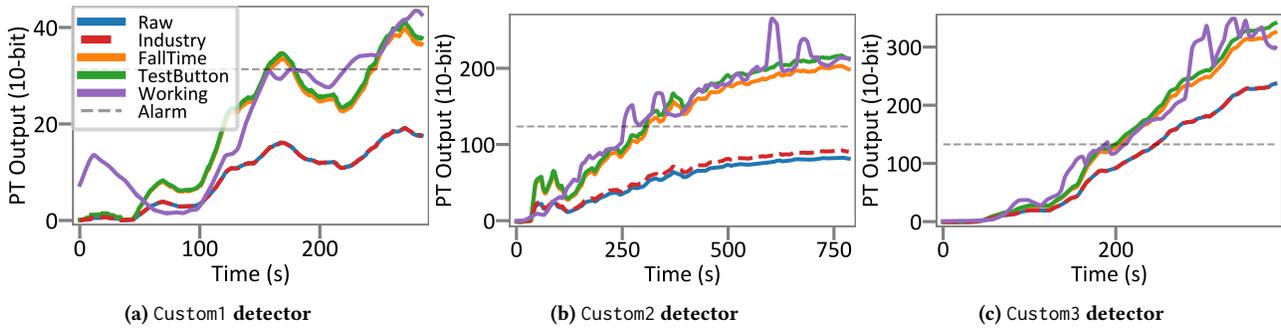


Figure 17: Impact of drift compensation using different approaches in Smoke condition.

in drift scores. For instance, the yellow highlighted region in Figure 15(b) shows one Custom2 sensor instance with a drift score of 0.67. The sensor needs considerable gain to offset its high drift score; both FallTime and DriftTestButton recommend applying a gain of about 2.5, while IndustryApproach incorrectly recommends applying almost no gain.

Result. The gain determined by both FallTime and DriftTestButton accurately compensates for the PT drift and outperforms IndustryApproach.

7.3 Smoke tests: Drift compensation

We now test our approaches in tests with real smoke.

Setup: We expose six instances of each of Custom1, Custom2, and Custom3 detectors (18 in total) to increase densities of smoke. We created a smoke chamber to ensure that a consistently increasing amount of smoke is exposed to the detectors. Figure 16 shows our smoke test setup with 2 chambers. Incense sticks are burned in Chamber 1 to generate smoke, and the smoke is pumped into Chamber 2, where smoke detectors are placed. Incense smoke matches the properties of smoke from many household materials, and is a common smoke detector test method [35, 41]. Both chambers are closed to eliminate the effect of the ambient environment.

Drift compensation occurs entirely in software. The microcontroller (MCU) records the PT's raw output, and applies gain to this value as specified by one of the drift compensation algorithms. This allows the MCU to apply all three drift compensation methods in parallel so we can compare them directly for the same smoke detector in the same experiment.

Figure 17(a)-(c) shows smoke experiments results from two smoke detectors, one working (purple line) and one with a drifted PT. For the drifted PT, we show the raw PT output (blue line) and the PT output compensated with each of IndustryApproach, FallTime and DriftTestButton approaches (red, orange, and green lines, respectively). We set the alarm threshold at 1% obscuration /ft, which represents a common sensitivity settings for smoke detectors. At the start of the experiment ($t = 2s$), we begin burning the incense sticks in Chamber 1. As smoke accumulates, the PT output should gradually increase. The raw data from the drifted sensor (shown in blue) increases much more slowly than the working sensor (shown in purple). The orange and green lines show the PT output after applying the FallTime and DriftTestButton drift compensation, respectively. Each approach applies adequate gain, and the output of the drifted sensor roughly matches that of the working sensor, and alarms as quickly as the working sensor.

The IndustryApproach (red line), on the other hand, does not apply adequate gain, and does not reach the alarm threshold quickly enough (or at all, in some cases). The drifted detectors with either FallTime or DriftTestButton compensation typically trigger an alarm almost as quickly (taking only up to 5% more time) as the working detector.

Result. The output of drifted sensors closely matches the working sensor when the gain is applied using FallTime and DriftTestButton, thus increasing the reliability of smoke sensors.

7.4 Smoke tests: Extensive results

We now present detailed results from many smoke tests performed on the three custom detectors.

Setup: We run the experiment ten times for each of the three smoke detector models. Each experiment contains six sensors of the same model with varying drift scores. The smoke chamber is reset after each experiment by clearing both chambers and the detectors of smoke and soot. As smoke accumulates, we measure the "ground truth" smoke density by measuring the light attenuation over a path of one foot (obscuration percent per foot). We record the true smoke density and the time at which each sensor's output crosses the alarm threshold (1%/ft), for each compensation algorithm. We report the mean value across all 10 experiments.

Figure 18(a)-(d) show our results in terms of smoke density and time for Custom2 and Custom3 detectors at a 1%/ft alarm threshold. The x-axis of each plot is the drift score and the y-axis is either the smoke density or time at which the sensor alarms. Each graph has four sets of points, colored blue, red, orange and green, corresponding to raw data, IndustryApproach, FallTime and DriftTestButton compensation methods, respectively. Each set of points are aligned on the x-axis because they correspond to a single smoke detector with the same drift score, but the y-values vary depending on the compensation algorithm. The goal of compensation is to reduce the influence of drift. An ideal compensation algorithm will eliminate the influence of drift, so its line should have 0 slope as the drifted sensors should alarm as quickly as working sensors. Figures 18(a) and (b) show that as the drift score increases, the uncompensated drifted sensors require higher smoke densities to trigger an alarm, as expected. IndustryApproach differs only slightly from the raw data, and still requires very high smoke densities to alarm. Conversely, FallTime and DriftTestButton compensation each significantly mitigate the effect of drift, with the fit lines having slope close to 0 (orange and green lines). For example, in Figure 18(a), the point on the far left corresponds to a working detector,

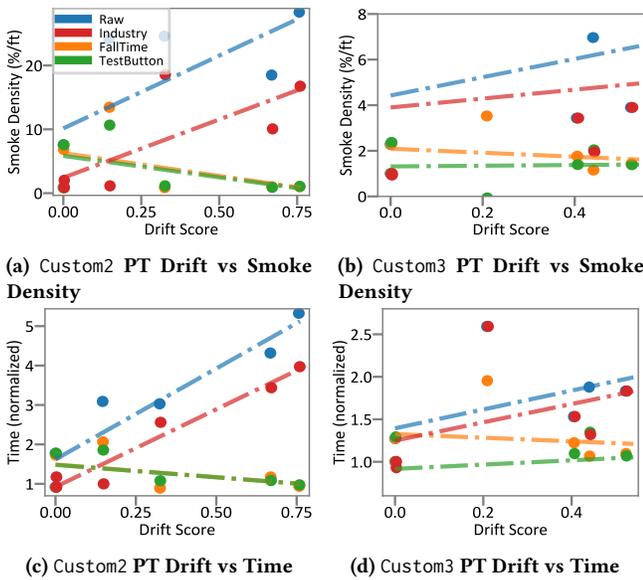


Figure 18: Influence of drift on smoke detector time and smoke density required to alarm, for each compensation algorithm. Flatter, lower lines indicate better drift compensation.

while the far right corresponds to severely drifted detectors. For the severely drifted detector, IndustryApproach would require nearly 16x the smoke density (16%/ft) to alarm, compared to a working sensor (1%/ft). On the other hand FallTime and DriftTestButton compensated detectors each trigger an alarm at roughly the same smoke density as the working detector (i.e. 1%/ft).

Similar behavior can be seen in Figure 18(c) and (d), which shows the time required to trigger an alarm as the drift score of a detector increases. The y-axis represents the time at which a sensor alarms and is normalized so that a value of 2 indicates a sensor took twice as long to alarm as the fastest sensor. For example, in Figure 18(c) detectors that are severely drifted (e.g. the point on the far right) require up to 4x the time to alarm using IndustryApproach compared to the working sensor (far left, drift score 0). However, FallTime and DriftTestButton trigger the alarm as quickly as a working sensor (near time = 1). We found pairwise comparisons between our new methods against existing methods (i.e. each of IndustryApproach and no compensation versus each of FallTime and DriftTestButton) for each alarm threshold and sensor type to be statistically significant.

Result. On average, across all the three type of smoke detectors, we found that IndustryApproach requires 12%/ft higher smoke density and takes 65% longer time to trigger an alarm compared to the FallTime fingerprint and DriftTestButton approaches. These differences are highly significant in terms of fire safety.

8 DISCUSSION & FUTURE WORK

We now present a few limitations of our fingerprint-based approach and directions for our future work:

Long-term evaluation: the compensation algorithms and the impact of drift on smoke density/time to alarm was shown only

on custom detectors (Section 7.4) rather than commercial models, because their microcontrollers are not reprogrammable. We plan to deploy our custom detectors at scale to evaluate the proposed fingerprints over long periods.

Fingerprint-based approach vs. mechanical DriftTestButton approach: the proposed fingerprint approach requires software API changes with minimal/no hardware modification, thus enabling scaled deployments. On the other hand, the DriftTestButton based approach is precise and robust, but requires a manual button press by users and hardware modification (insertion of button). The former issue could be resolved by automating the button press using a relay and a switch, and given that the button cost is low, for large-scale deployment, manufacturers can include the button in future designs.

Generalization of the proposed fingerprint beyond smoke detectors: PTs are one of the most commonly used electrical components, and are used in numerous light scattering sensors, such as particulate matter air pollution sensors [52], water turbidity sensors [22], and proximity sensors [8], across many application scenarios. The proposed FallTime fingerprint can be employed in these sensors to monitor faults and measure drift. As part of our future work, we plan to apply FallTime in a variety of sensors and make the software available for the IoT and Sensing community.

9 CONCLUSIONS

Smoke detectors have played a key role in precipitous declines in fire deaths. In spite of this, they still frequently fail to alarm in the presence of smoke due to component faults or degradation, which current techniques cannot detect. This work first presents a key finding that phototransistor (PT) component drift may be widespread in smoke detectors, contributing to greatly reduced speed in responding to fire. To address this, we present a novel FallTime fingerprint assessing key electrical properties of the PT to detect faults and measure drift in smoke detectors. We also designed a mechanical DriftTestButton to simulate smoke conditions that can test for faults and drift in the field. We conduct extensive tests with real Smoke to assess the ability of the proposed approaches to accurately compensate for drift in sensor components. We showed that the FallTime and DriftTestButton approach each outperform existing solutions, and can greatly improve fire safety. Moving forward, we will be extending the applicability of the fingerprint to other sensor types.

REFERENCES

- [1] 2005. *System Sensor Europe*. <https://www.firelite.com/CatalogDocuments/df-52384.pdf>
- [2] 2017. *Basic Characteristics and Application Circuit Design of Transistor Couplers*. <https://toshiba.semicon-storage.com/info/docget.jsp?did=13438>
- [3] 2018. *First alert 9120 Smoke detector*. <https://www.firstralert.com/product/>
- [4] 2018. *GST I-9102 Intelligent Photoelectric Smoke Detector*. <https://www.gst.com.cn/en/data/products/Manuals/I-9102%20Intelligent%20Photoelectric%20Smoke%20Detector%20Issue1.03.pdf>
- [5] 2018. *Kidde Battery Operated Photoelectric Smoke Alarm P9050*. <https://www.kidde.com/home-safety/en/us/products/fire-safety/smoke-alarms/p9050/>
- [6] 2018. *LTspice*. <http://www.analog.com/en/design-center/design-tools-and-calculators/ltspice-simulator.html>
- [7] 2019. *National Fire Protection Association, Ionization vs photoelectric*. <https://www.nfpa.org/Public-Education/Staying-safe/Safety-equipment/Smoke-alarms/Ionization-vs-photoelectric>
- [8] 2019. *Photoelectric Proximity Sensors*. <https://www.fargocontrols.com/sensors/photoelectric.html/>

- [9] 2019. *STM32 IoT node*. <https://www.st.com/en/evaluation-tools/b-1475e-iot01a.html>
- [10] MZ Abidin, Arnis Asmat, and MN Hamidon. 2018. Comparative Study of Drift Compensation Methods for Environmental Gas Sensors. *Earth and Environmental Science* 117 (2018).
- [11] M. G. Adlerstein and J. M. Gering. 2000. Current Induced Degradation in GaAs HBT's. *IEEE Transactions on Electronic Devices* 47 (2000).
- [12] M. Ahrens. 2009. *Smoke Alarms in U.S. Home Fires*. Technical Report. National Fire Protection Association.
- [13] M. Ahrens. 2019. *Smoke Alarms in US Home Fires*. Technical Report. National Fire Protection Association.
- [14] National Electrical Manufacturers Association. 2002. *Mission Effectiveness and Failure Rates Drive Inspection, Testing, and Maintenance of Fire Detection, Alarm, and Signaling Systems*. Technical Report. Society of Fire Protection Engineers.
- [15] National Fire Protection Association et al. 2010. *NFPA 72: National Fire Alarm and Signaling Code*. National Fire Protection Association.
- [16] James P Barrett. 1974. The coefficient of determination. *The American Statistician* 28, 1 (1974), 19–20.
- [17] J. J. Brenden. 1970. *DETERMINING THE UTILITY OF A NEW OPTICAL TEST PROCEDURE FOR MEASURING SMOKE FROM VARIOUS WOOD PRODUCTS*. Technical Report. Forest Products Laboratory.
- [18] M. S. Carrol, A. Neugroschel, and C. T. Sah. 1997. Degradation of silicon bipolar junction transistors at high forward current densities. *IEEE Trans. Electron Devices* 44 (1997), 110–117.
- [19] Tusher Chakraborty, Akshay Uttama Nambi, Ranveer Chandra, Rahul Sharma, Manohar Swaminathan, Zerina Kapetanovic, and Jonathan Appavoo. 2018. Fall-curve: A novel primitive for IoT Fault Detection and Isolation. In *Proceedings of the 16th ACM SenSys*. 95–107.
- [20] Ratzlaff L. Bush D Chandler, B. 2016. Photoelectric Smoke Detector with Drift Compensation. <https://patentimages.storage.googleapis.com/d3/f0/42/815589e7092124/US9396637.pdf> US Patent 9,396,637B2.
- [21] Thomas Cleary. 2016. *A Study on the Performance of Current Smoke Alarms to the New Fire and Nuisance Tests Prescribed in ANSI/UL 217-2015*. Technical Report. US Department of Commerce.
- [22] Shoucheng Ding and Wenhui Li. 2011. Research on photoelectric sensor turbidity detecting system. *Sensor Letters* 9, 4 (2011).
- [23] W. F. Doherty. 1977. Smoke detector with test means for simulating a predetermined concentration of smoke. US Patent No. 4144458.
- [24] K. W. Dungan. 2007. Reliability of Fire Alarm Systems. *Fire Protection Engineering* 2007 Q1 (2007), 11 pages.
- [25] J. Webster (ed.). 1999. Hot Carriers. *Wiley Encyclopedia of Electrical and Electronics Engineering*.
- [26] TT Electronics. 2016. *NPN Silicon Phototransistor OP505*. <https://www.ttelectronics.com/TTElectronics/media/ProductFiles/Optoelectronics/Datasheets/OP505-506-535-705.pdf>
- [27] A. Feinberg, P. Erslund, V. Kaper, and A. Widom. 2000. On Aging Of Key Transistor Device Parameters. Institute of Environmental Sciences and Technology, Providence, USA, 231–236.
- [28] Firebus 2014. *FB-AP Drift Compensation*. Firebus. <https://firebus.net/drift-compensation>
- [29] Firelite Alarms 2015. *SD355(A) Series Addressable Photoelectric Smoke Detectors*. Firelite Alarms. <https://www.firelite.com/CatalogDocuments/df-52384.pdf>
- [30] Joseph Fleming. 1999. Smoke Detector Technology and the Investigation of Fatal Fires. (1999). <https://www.interfire.org/features/smokedetector.asp>
- [31] Allen Hess. 2011. Smoke Detector Sensitivity Testing. http://www.cfaa.ca/Files/flash/ontario/ats2011/Smoke_Detector_Sensitivity_Testing_by_Allen_Hess.pdf "Presentation".
- [32] Chenming Hu. 2009. *Modern Semiconductor Devices for Integrated Circuits*. Pearson College Div.
- [33] S. Y. Huang, K. M. Chen, G. W. Huang, V. Liang, H. C. Tseng, T. L. Hsu, and C. Y. Chang. 2005. Hot-Carrier Induced Degradations on RF Power Characteristics of SiGe Heterojunction Bipolar Transistors. *IEEE Transactions on Device and Materials Reliability* 5 (2005), 183–189.
- [34] Kidde Fire Systems 2011. *SmartOne Devices*. Kidde Fire Systems. https://www.kidde-fenwal.com/Media/Data%20Sheets/K-76-801_screen.pdf
- [35] Shi-Kuk Kim, Hyun-Dai Yuk, Seung-Hyun Yang, Seung-Wook Jee, and Chun-Ha Lee. 2009. A Study on the Problem of Tester for the Field Inspection of the Photoelectric Smoke Detector. *Fire Science and Engineering* 23 (2009), 137–144. Issue 4.
- [36] S. L. Kosier, A. Wei, R. D. Schrimpf, D. M. Fleetwood, M. D. Delaus, R. L. Pease, and W. E. Combs. 1995. Physically based comparison of hot-carrier-induced and ionizing-radiation-induced degradation in BJT's. *IEEE Trans. Electron Devices* 42 (1995), 436–444.
- [37] Sovannarith Leng. 2017. *Identifying and evaluating aging signatures in light emitting diode lighting systems*. Ph.D. Dissertation.
- [38] Sumukh Marathe, Akshay Nambi, Nishant Shrivastava, Manohar Swaminathan, and Ronak Sutaria. 2020. Fault Diagnosis System for Low-Cost Air Pollution Sensors: Demo Abstract. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems (Virtual Event, Japan) (SenSys '20)*. Association for Computing Machinery, New York, NY, USA, 611–612. <https://doi.org/10.1145/3384419.3431191>
- [39] Sumukh Marathe, Akshay Nambi, Manohar Swaminathan, and Ronak Sutaria. 2021. CurrentSense: A novel approach for fault and drift detection in environmental IoT sensors. In *Proceedings of the 6th ACM/IEEE International Conference on Internet of Things Design and Implementation (IoTDI)*. ACM, 99–112.
- [40] A. A. Marsocci. 1975. Optical smoke detector with light scattering test device. US Patent No. 3868184.
- [41] George Mulholland. 2002. Smoke Production and Properties. In *SFPE Handbook of Fire Protection Engineering*, Phillip DiNenno, Dougal Drysdale, Craig Beyler, W. D. Walton, Richard Custer, John R. Hall, and John M. Watts (Eds.). Society of Fire Protection Engineers.
- [42] Ann Mutschler. 2017. Transistor Aging Intensifies At 10/7nm And Below. (2017). <https://semiengineering.com/transistor-aging-intensifies-10nm/#:~:text=Device%20degradation%20becomes%20limiting%20factor,%20significant%20challenge%20in%20advanced%20SoCs.&text=Transistor%20aging%20and%20reliability%20are,%20teams%20at%2010nm%20and%20below.>
- [43] Electronics Notes. 2020. Phototransistor Applications & Circuit Configurations. (2020). https://www.electronics-notes.com/articles/electronic_components/transistor/phototransistor-circuits-applications.php
- [44] NXP Semiconductor 2006. *Photoelectric Smoke Detector IC with I/O and Temporal Pattern Horn Driver*. NXP Semiconductor. <https://www.nxp.com/docs/en/datasheet/MC145012.pdf>
- [45] Bob Orwiler. 1969. Vertical Amplifier Circuits. Tektronix. (1969).
- [46] OSRAM Opto Semiconductors 2018. *SFH 4550*. OSRAM Opto Semiconductors. <https://www.osram.com/media/resource/hires/osram-dam-2496286/SFH%204550.pdf>
- [47] Corinne Peek-Asa, Jingzhen Yang, Cara Hamann, and Tracy Young. 2011. Smoke alarm tests may not adequately indicate smoke alarm function. *Journal of Burn Care & Research* 32, 4 (2011), e135–e139.
- [48] Muhammad H Rashid and MH Rashid. 1995. *SPICE for circuits and electronics using PSPICE*. Prentice Hall Upper Saddle River, NJ.
- [49] R. P. Schiffliti. 2006. *Fire Alarm Systems for Life Safety Code Users*. National Fire Protection Association, Arlington, VA. 1113 – 1138 pages.
- [50] Vishay Semiconductors. 2011. *BPW96B, BPW96C Silicon NPN Phototransistor*. <https://www.vishay.com/docs/81532/bpw96.pdf>
- [51] Vishay Semiconductors. 2013. *BPV11F Silicon NPN Phototransistor*. <http://www.vishay.com/photo-detectors/list/product-81505/>
- [52] AG Sensirion. 2017. Sensirion Particulate Matter Sensor SPS30.
- [53] ST Microelectronics 2019. *Application guide: Fire alarms and smoke detectors*. ST Microelectronics. https://www.st.com/content/ccc/resource/sales_and_marketing/presentation/product_presentation/ae/6c/ae/b5/db/88/45/a9/ulp_firealarms_appliguide.pdf/files/ulp_firealarms_appliguide.pdf/jcr:content/translations/en.ulp_firealarms_appliguide.pdf
- [54] System Sensor 2016. *System Smoke Detectors Application Guide*. System Sensor. https://www.systemsensor.com/en-us/Documents/System_Smoke_Detectors_AppGuide_SPAG91.pdf
- [55] Texas Instrument 2017. *ADS7056 Ultra-Low Power, Ultra-Small Size, 14-Bit, High-Speed SAR ADC*. Texas Instrument. https://www.ti.com/lit/ds/symlink/ads7056.pdf?ts=1617225081254&ref_url=https%253A%252F%252Fwww.google.com%252F
- [56] Texas Instruments 2018. *MSP430FR596x, MSP430FR594x Mixed-Signal Microcontrollers*. Texas Instruments. https://www.ti.com/lit/ds/symlink/msp430fr5969.pdf?ts=1617221382068&ref_url=https%253A%252F%252Fwww.ti.com%252Fproduct%252FMSP430FR5969
- [57] N. Thudsalingkarnsakul, S. Thainoi, and S. Kanjanachuchai. 2007. Characteristics of GaAs-based Phototransistors After High-Power Operation. In *ECTI International Conference*. ECTI, 141–144.
- [58] N. Toufik, F. Pelanchon, and P. Mialhe. 2001. Degradation of Junction Parameters of an Electrically Stressed NPN Bipolar Transistor. *Active and Passive Electrical Components* 24 (2001), 155–163.
- [59] Van Tran, Robert Stuart, and Hardik Bhavsar. 2009. *Phototransistor Switching Time Analysis*. California Eastern Laboratories.
- [60] Gary Wasserman. 2002. *Reliability Verification, Testing, and Analysis in Engineering Design*. Wayne State University, Detroit, Michigan.
- [61] Noah Zamdmer. 1994. *Hot-Electron Degradation of Bipolar Transistors*. Ph.D. Dissertation. Massachusetts Institute of Technology.