Monoxalyze: Verifying Smoking Cessation with a Keychain-sized Carbon Monoxide Breathalyzer

Joshua Adkins and Prabal Dutta
Electrical Engineering and Computer Science Department
University of Michigan
Ann Arbor, MI 48109
{adkinsjd,prabal}@umich.edu

ABSTRACT

We present Monoxalyze, a keychain-sized, Bluetooth-based, carbon monoxide breathalyzer that aims to enable mobile, scalable smoking cessation intervention programs. These intervention programs have been shown to greatly increase the rate of a quit attempt, which in turn decreases the rate of smoking, a major public health problem that still affects over one billion people around the world. Currently, intervention programs verify cessation compliance by requiring program participants to periodically visit clinics and exhale through large, expensive carbon monoxide breathalyzers—a practice that cannot scale to one billion smokers. Monoxalyze enables mobile cessation verification by working with a user's smartphone to establish a ring of spatio-temporal transitive trust between the Monoxalyze device, the user, and the smartphone, a concept that can be applied to many third-party monitoring applications. In Monoxalyze, the links of this trust are represented by simultaneous exhalation verification, facial recognition, and device-to-phone visible light authentication. In our evaluation, we show that Monoxalyze lasts over 80 days between charges, and has the ability to verify a Monoxalyze user. With a small user study we show that Monoxalyze determines smoking cessation with 92% accuracy, a level comparable with commercial CO breathalyzers. Further contributions describe the design decisions behind creating a low-power BLE device.

CCS Concepts

•**Security and privacy** → **Multi-factor authentication; Embedded systems security;** *Hardware-based security protocols; Malicious design modifications;* •**Human-centered computing** → *Ubiquitous and mobile computing systems and tools; Mobile devices;* Smartphones; •**Computer systems organization** → *Sensors and actuators;* Embedded hardware;

The smoking cessation determination trials presented in this paper were approved by the Intitutional Review Board of the Jersey City Medical Center–Barnabas Health in Jersey City, NJ under protocol number 15–07.

SenSys '16 November 14-16, 2016, Stanford, CA, USA

ACM ISBN 978-1-4503-4263-6/16/11.

DOI: http://dx.doi.org/10.1145/2994551.2994571

1. INTRODUCTION

Smoking has caused 20 million premature deaths over the past fifty years and annually costs the health care system over $130 billion in the United States alone [30]. While the number of smokers has and continues to decrease, worldwide over one billion people still regularly smoke cigarettes.

Fortunately, most people in developed countries are aware of the adverse health affects of smoking and therefore want to quit. Unfortunately, most people fail to successfully quit smoking due to the addictive nature of cigarettes. Of the 50% of smokers who attempt to quit smoking each year, only 4-6% succeed [30]. Enrollment in a smoking cessation intervention program can increase the long term success of a quit attempt to above 15% [41] and can increase short term success by much more [38]. These intervention programs use techniques such as personalized motivational messages, competitions, and financial incentives.

The successful implementation of intervention methods for smoking cessation requires a compliance verification mechanism to report if the participant has smoked. This is helpful for programs that provide personal motivation, but is particularly important for programs that offer a financial incentive. Self reporting of smoking cessation is a common technique used in smoking cession programs, but is known to be unreliable when monetary incentives or lottery prizes are introduced [8]. Even in clinical settings—where there is little incentive to lie—self cessation reports have shown to be unreliable [42].

One method for improving the accuracy of smoking cessation reporting mechanisms is to detect biochemical markers left behind from smoking cigarettes rather than relying on self reporting. In particular, detecting the carbon monoxide (CO) concentration of exhaled breath provides a minimally invasive and relatively inexpensive mechanism for detecting smoking activity. CO breathalyzers exist for this purpose, but current versions are large, provide limited data reporting capabilities, and include no verification mechanism, thus still requiring daily interaction with an intervention program coordinator to eliminate deceptive self reporting.

To address these problems, we present Monoxalyze, a keychain–sized, mobile, wireless CO breathalyzer. Monoxalyze increases the reliability of using exhaled CO as a smoking cessation intervention technique by attempting to establish that a specific program participant has exhaled through an authenticated Monoxalyze device. Monoxalyze then reports this reliable CO reading to the intervention program through an Internet–connected smartphone.

To use the device, participants periodically exhale a complete breath through Monoxalyze. Monoxalyze employs a pressure sensor to detect when a user is exhaling through the device and uses this event as a trigger to perform a full CO breath sample test. The device then cooperates with the participant's smartphone to perform simultaneous CO measurements for smoking detection and facial recognition for user verification. To verify that the user is exhaling through the same Monoxalyze device to which the smartphone is connected, rather than a nearby device that a non-smoker friend is using, Monoxalyze uses a visible light challenge-response protocol. This protocol requires the smartphone to send a challenge to the Monoxalyze device over BLE and receive the response from the LEDs on the front of the Monoxalyze device using the same camera that is performing facial recognition, effectively localizing the connected Monoxalyze device to the program participant's face during exhalation.

The authentication system used in Monoxalyze is an early example of ensuring data provenance without requiring professional observers to verify the authenticity of the data, and as such could be used for many applications that requires third party monitoring and data collection. These applications could include a parent verifying information about their child, a medical insurance company providing discounts for verified exercise, doctors collecting information about their patients, or law enforcement verifying court-mandated alcohol abstinence testing. All of these monitoring applications could be strengthened by requiring semi-random data collection within a specific time window, which would make complicated attacks on the authentication system difficult to perform consistently.

Critical to creating a usable and practical device is ensuring that the device fits in a reasonable form factor, and that the device can perform many breath tests without requiring the battery to be recharged. To minimize the size of the device, Monoxalyze uses a smaller CO sensor than the traditional canister sensors used in most commercial CO breathalyzers. To maximize the battery life, we employ the pressure sensor as a low-power wakeup sensor, search for the BLE radio parameter selection that minimizes energy consumption, and design the visible light challenge-response protocol with low-power transmission in mind. This response protocol could be used to perform out-of-band authentication on any low-power embedded system, even if it does not require the transitive trust properties of Monoxalyze.

2. RELATED WORK

Stemming from the considerable research into the negative health effects of smoking, there have been many efforts to develop successful smoking cessation intervention methods, including research into the success rate of such programs, self reporting methods, biochemical detection, and breathalyzer systems.

While it has been shown that intervention methods increase smoking cessation success rates [45], the relative success of intervention methods, both individually and in combination, is still a matter of debate. There is work that shows increasing intrinsic motivation while also providing support to a smoker increases the success rate of a quit attempt [9]. Furthermore, the research suggests that extrinsic motivators only increase program enrollment, rather than long term abstinence. Other work agrees with this view of financial incentive as a poor motivator for long term abstinence (exceeding 12 months) [7]. There are, however, many studies that vouch for the success of financial incentives at achieving at least short term abstinence (less than 12 months) with and without additional intrinsic motivators [11, 17, 20, 25, 38]. This is especially useful in cases of attempting to induce immediate cessation in newly pregnant mothers. On the other side of the argument, there is a single study that shows long term cessation success due solely to competition and financial incentive [41].

All efforts to study the costs and benefits of smoking cessation programs rely on a reliable smoking cessation reporting mechanism. As one would expect, the reliability of self reporting smoking cessation is low and consistently underestimates smoking prevalence [15]. While in general this underestimation appears to be relatively small at 2-6% [40,42], when there is an external incentive to be untruthful during self reporting, the reliability of self reporting drops significantly. This primarily has been researched with the moral incentive of women to quit smoking during pregnancy, where the underestimation of self reporting is shown to be 22-25% [14,36]. There has been some work detailing deception in self-reporting due to financial incentive as well, however, this does not examine the fraction of smokers who lied about smoking, but rather the amount of nonsmokers who lied about being smokers to gain financial incentive [8].

There are several methods of determining smoking cessation more reliably than self reporting by using the biochemical markers that smoking leaves behind in the body. One such chemical marker is carboxyhemoglobin in blood. It is well documented that the concentration of carboxyhemoglobin increases in a smoker [18], and furthermore that this increase can be used to determine smoking cessation [32]. The increase in blood carboxyhemoglobin concentration in turn causes an increase in the exhaled carbon monoxide concentration in breath. This is caused by the CO level in alveolar air in the lungs equilibrating with the carboxyhemoglobin in the blood during gas exchange [19]. This increase in exhaled CO concentration can therefore be used as an indicator of smoking cessation [10,21,28]. The concentration of CO in expired breath is dependent on the completeness of the equilibrium achieved with the carboxyhemoglobin in blood, which in practice is dependent on the speed of exhalation and the amount of time the exhaled breath is held [43].

This paper focuses on measuring the CO concentration in exhaled breath to determine and report smoking cessation. Devices that use electrochemical CO sensors and nondispersive infrared CO detectors to measure the CO concentration of exhaled breath are common. Nearly all such devices are commercial devices such as the Smokerlyzer products [5] and COSlueth [6]. These devices are all too large to comfortably carry in a pocket or on a keychain. The smallest of these devices, the Smokerlyzer iCo, was released in March 2015, and is only slightly smaller than a toilet paper roll [2]. The iCo is also the only commercial device that is able to connect to a phone or tablet, making it the only commercial carbon monoxide breathalyzer that can also report its data to a mobile device. It has been previously shown that a mobile phone based breathalyzer is possible [27], and the iCo is the first commercial offering of such a device. Monoxalyze is specifically designed with both size and mobile phone connectivity in mind, and is smaller than the iCo, while also including wireless data transmission.

While the devices discussed above succeed at measuring CO concentration of exhaled breath and increase smoking cessation reporting reliability in a clinical setting, they do not verify that a user is exhaling through them and they do not verify the identity of the user. Therefore a trusted third party must watch and verify the CO measurement for accurate reporting, and this requirement prevents the use of current CO breathalyzers in large, scalable smoking cessation programs outside of a clinical setting. Some research has tried to address this problem by requiring program participants to record a daily video of themselves exhaling into a CO breathalyzer, then remotely validating the CO reading by watching the video [37], however this puts a high burden on intervention providers, and still does not verify that the program participant is actually exhaling through the breathalyzer.

Monoxalyze solves this authentication problem by combining smartphone facial recognition with device authentication. Authentication is performed by leveraging prior work in out-of-band channel communication to verify location [34]. The out-of-band channel used in Monoxalyze is visible light communication (VLC). Prior

Device	Dimensions	Weight	Connectivity	Exhalation Verification	Identity Verification
Smokerlyzer piCO+ [4]	120x75x45 mm	200 g	None	None	None
Smokerlyzer Micro++ [3]	138x77x44 mm	250 g	None	None	None
Smokerlyzer COmpact USB [1]	115x50x33 mm	100 g	USB	None	None
Smokerlyzer iCo [2]	100x40x29 mm	75 g	Mobile Phone Audio Jack	None	None
CO Sleuth [6]	Unknown	Unknown	None	None	None
Mobile Research Device [27]	51x51x32 mm	Unknown	Mobile Phone Audio Jack	None	None
Monoxalyze (this work)	**81x23x13 mm**	**15 g**	**BLE**	**Pressure Sensor**	**Facial Recognition and VLC Authentication**

Table 1: Comparison of Monoxalyze to other CO breathalyzers. Monoxalyze is the smallest CO breathalyzer by volume, the first to employ a wireless connection, and the first to provide exhalation and identity verification mechanisms. These features allow Monoxalyze to perform as a mobile smoking cessation verification device that does not require daily visits to an intervention clinic.

work shows that VLC can be received by an unmodified smartphone camera [22, 33], and our work extends this with the unique property of transmitting the visible light communication from a low-power embedded device, requiring a protocol that minimizes the power drawn by the transmitter.

3. DESIGN REQUIREMENTS

Monoxalyze is designed to be a mobile smoking cessation verification device, and as such, it is designed with specific requirements in mind. These requirements focus on the usability, data validity, and success of smoking cessation detection. Monoxalyze's specific design goals are:

Usability. Monoxalyze was designed to be a fully mobile carbon monoxide breathalyzer. As such Monoxalyze should fit easily in a handbag or pocket, or on a keychain. Monoxalyze should also have as long a battery life as possible, but no less than a week. For ease of interface, Monoxalyze should respond quickly to intentional user interaction. This includes responsive wakeup from sleep and quick formation of wireless connections.

Data Validity. Monoxalyze must validate with reasonably high probability that a specific user has not smoked. To provide this validation, Monoxalyze along with its smartphone application, must verify the identity of the Monoxalyze user and also verify that the smartphone is receiving data from the Monoxalyze device through which the identified user is exhaling. Monoxalyze must also verify that the user is exhaling sufficient positive pressure through the Monoxalyze device that is providing the carbon monoxide reading.

Cessation Determination. To verify that a person has not smoked, Monoxalyze should provide a carbon monoxide reading under normal use scenarios. While the overall accuracy of these carbon monoxide readings is useful, the ability of Monoxalyze to distinguish between smokers and non-smokers independent of accuracy is more important. Monoxalyze should be able to detect failure of a carbon monoxide sensor.

To achieve the usability requirements, we design Monoxalyze to be a low power embedded system with pressure sensitive wakeup. To achieve cessation determination requirements, we consider and evaluate the accuracy of the carbon monoxide sensor as well as recognize the effects of human breath on this accuracy. To achieve the data validity requirements, we specify the threat model below, and then design Monoxalyze to defend against this model.

3.1 Threat Model

To verify the data validity of a CO reading, we must consider the specific threat model attempting to undermine the validity of the data. In this threat model we will assume that the only attacker of the system is a user without technical knowledge. We believe that many threats which require technical knowledge such as attacking the wireless communication link may be solved through proven methods of encryption and authentication, so exploring these threats is not necessary in this paper.

This average user is therefore assumed to be incapable of modifying or appending to the Monoxalyze circuitry. They are also assumed to be incapable of breaking open the physical casing and modifying the characteristics of the pressure chamber. In addition to these requirements, this user is assumed to be incapable of intercepting and re-transmitting the visible light communication in any way.

A user may still, however, attempt to act maliciously by changing their interaction with the exterior workings of the system. This includes any method by which they do not exhale a high concentration of CO through the system. They may attempt this by modifying how and when they exhale through the device, by changing how exhaled air exits the device, and by changing the identity of the person exhaling through the device.

Monoxalyze will therefore attempt to mitigate this threat model by verifying that a specific user is exhaling their breath through the Monoxalyze device that is sending CO breath measurements to the smartphone. It should be noted that the goal is threat mitigation rather than full data validation, so that it is reasonably difficult for a user to trick Monoxalyze into sending a CO reading that is lower than the actual CO concentration of that user's exhaled breath. Because much of the attack surface depends on a user's physical interactions with Monoxalyze, complete data validation is considered more ambitious than necessary. The goal is to make the difficulty of deceiving Monoxalyze high enough that the effort of deception significantly outweighs the reward.

4. SYSTEM OVERVIEW

This section presents an overview of the Monoxalyze system and how it meets the requirements presented in Section 3. We do so by walking through the step-by-step use of Monoxalyze. The steps below refer to the numbers shown in Figure 1.

0) Register. Before using Monoxalyze, a user is required to register with an intervention program. This includes registering training pictures with the dedicated application that runs on the user's smartphone. The pictures used for registration me be confirmed as the program participant either through another mechanism such as ID scanning or by a human. This registration will only occur once.

1) Open. When a user opens the Monoxalyze app, their smartphone starts scanning for Monoxalyze devices over BLE.

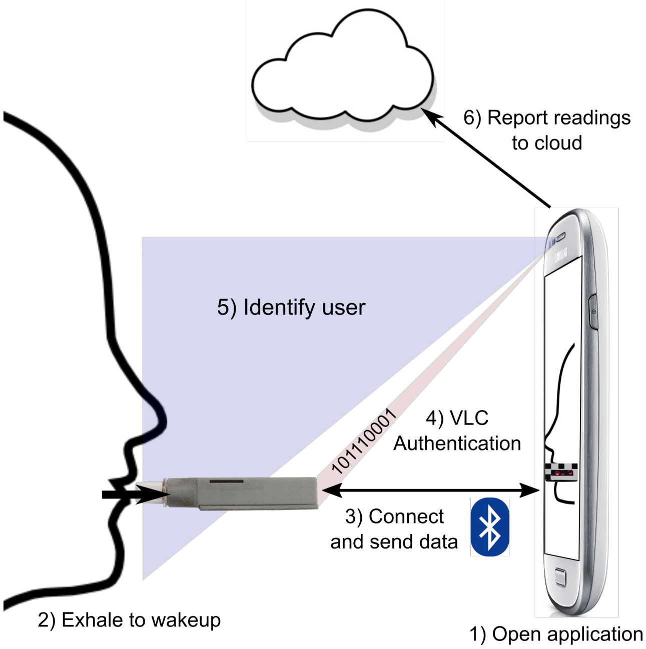

Figure 1: An overview of the normal use ofMonoxalyze. 1) *Open* the smartphone application so that it starts attempting to connect to the Monoxalyze device. 2) *Wakeup* Monoxalyze by taking a deep breath and exhaling slowly into the device. 3) *Connect* to the smartphone using Bluetooth Low Energy. Start reporting carbon monoxide measurements and verification that the user is exhaling to the smartphone. 4) *Authenticate* and localize the Monoxalyze device to the user's face with a visible light challenge response protocol. 5) *Identify* the Monoxalyze user with facial recognition. 6) *Report* the carbon monoxide reading and verification that the enrolled user exhaled through Monoxalyze to the intervention provider.

2) Wakeup. The user takes a deep breath and begin exhaling into the device. A pressure sensor inside the device wakes the device from sleep mode into advertising mode.

3) Connect. The Monoxalyze device advertises over BLE until a connection is made with the smartphone. It then sends carbon monoxide readings and verifies to the smartphone that the user is exhaling.

4) Authenticate. The smartphone sends a challenge to the Monoxalyze device over BLE, and Monoxalyze sends back the challenge to the phone's camera using the front facing LEDs. This verifies that in the next steps the phone is receiving data from the same Monoxalyze device that the person it identifies is exhaling through, as opposed to a second Monoxalyze that a non-smoker is exhaling into out of view.

5) Identify. The smartphone verifies the user's identity against the training data setup at registration using facial recognition.

6) Report. The smartphone sends the carbon monoxide reading and verification information to the intervention provider when network connectivity becomes available.

During actual Monoxalyze use, these steps can be parallelized. Wakeup and connection usually occur in under a second after which simultaneous authentication and identification can be performed by analyzing the same camera frames being used to authenticate the device to also identify the user. Data collection over BLE can also occur in parallel to authentication and identification. This means that a Monoxalyze user should be able to open the smartphone application, take a deep breath, and slowly exhale while the smartphone camera points at the user's face. It will not be necessary for the user to explicitly stop at each of these steps.

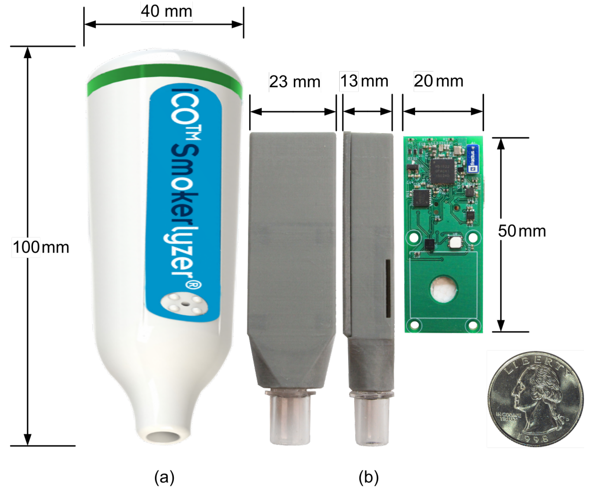

Figure 2: Relative size of Monoxalyze and the CO breathalyzer by Bedfont Scientific [2]. (a) shows the newly released Bedfont Smokerlyzer iCo. This is a picture taken from their website that has been morphed to fit the size specification in their datasheet. The iCo is oval shaped with a called out depth of 29 mm. (b) Multiple views of Monoxalyze. Monoxalyze is just 13 mm thick, making it easy to fit on a keychain or put in a pocket. The Monoxalyze circuit board is also shown and vertically positioned where it sits in the case. Monoxalyze is smaller than any carbon monoxide breathalyzer previously created. This shows that verifying smoking cessation is viable in a mobile form factor.

5. SYSTEM DESIGN

Monoxalyze, shown in Figure 2, provides verifiable mobile readings of the carbon monoxide present in exhaled breath. Monoxalyze achieves this by combining contributions in exhalation verification and BLE connection models, while applying concepts in out-of-band channels to location verification on an embedded system, and leveraging the field of mobile facial recognition. We present the system design by going step-by-step through the Monoxalyze use-case described in Section 4.

5.1 Pressure Sensitive Wakeup

While not required of a mobile carbon monoxide breathalyzer, we use the pressure sensor present in Monoxalyze for the verification purposes described in Section 5.4 to provide a user-transparent wakeup from sleep. This means the only interaction a user has with the Monoxalyze device is when they exhale through the device, a much cleaner user experience than a wakeup button or switch.

To achieve this, the sensed pressure is compared with a pressure threshold value. When the sensed value crosses the threshold, an interrupt wakes up the microcontroller. For this to be successful, sensed pressure must change by a reasonable amount when a user exhales through Monoxalyze. The work on exhalation verification discussed in Section 5.4, a key link in the ring of transitive trust, can be used to determine the appropriate pressure threshold.

After doing this analysis, we decide that a 300 pa pressure wakeup

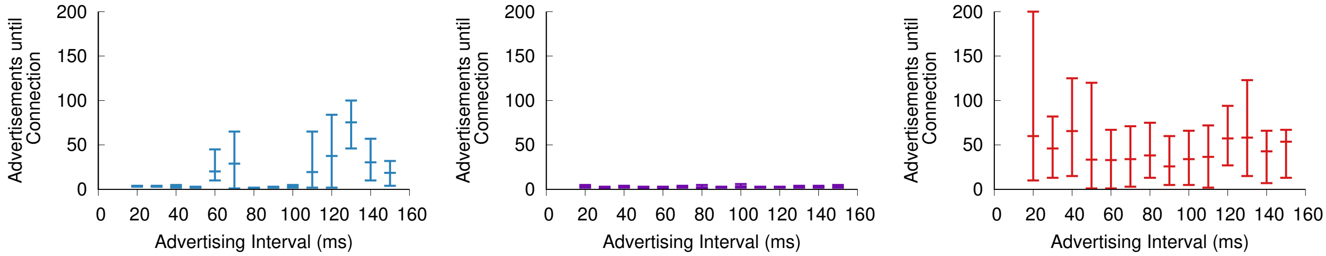

(a) Samsung Galaxy S5 (b) iPhone 6 (c) Motorola Moto X

Figure 3: A comparison of the number of advertisements until connection on three smartphone models. The error bars denote the maximum and minimum advertisements until connection. (a) shows that on the Samsung Galaxy S5 some advertising intervals have much better performance than others. We speculate that this is caused by intermittent scanning. (b) shows that the iPhone 6 has good performance for all advertising intervals. We speculate this is caused by constant scanning. (c) shows relatively poor performance for all advertising intervals on the Motorola Moto X. We think this is caused by either very low duty cycle scanning or poor radio performance. This comparison motivates the choice of an advertising interval and maximum number of advertisements for Monoxalyze. It also shows that total power used for a BLE interaction is very phone dependent; this leads to phone specific power optimizations for BLE peripherals.

threshold is reasonable for Monoxalyze. To account for external factors that may trigger Monoxalyze to wakeup, and to ensure that Monoxalyze never enters a state in which a user cannot induce wakeup, we institute a recalibration policy that causes Monoxalyze to wakeup on both negative and positive pressure threshold events. Monoxalyze always resets its reference pressure before entering sleep. These precautions enable users to still wakeup Monoxalyze even if they take actions that significantly lower their ambient air pressure, such as increasing their elevation. We analyze the impact of pressure sensitive wakeup and the potential for false wakeups has on power in Section 7.1.

5.2 BLE Connection Model

Monoxalyze uses BLE to wirelessly transfer carbon monoxide readings and exhalation verification to the smartphone. To the best of our knowledge, Monoxalyze is the first wireless carbon monoxide breathalyzer. The existing commercial and research devices that connect to mobile phones do so through the headset jack [23]. BLE was chosen because it is an increasingly common low-power communication protocol that allows Monoxalyze to connect to a smartphone.

The BLE connection model we design aims to transfer the necessary data with the lowest power possible. We can break down BLE energy usage into two sub categories: energy used for advertising and energy used during a connection.

To minimize advertising energy we would like to form a connection with the smartphone in as few advertisements as possible. Advertisements can be sent at different intervals, and to determine the interval that sends the fewest advertisements before connection, we start three smartphones, a Samsung Galaxy S5, a Motorola Moto X, and an iPhone 6 in active scanning mode, then record the number of advertisements until connection at a range of advertising intervals. From this data, shown in Figure 3, it is evident that the advertising interval chosen does matter, and is phone dependent. Using this data we choose an interval of 40 ms, and advertise a maximum of 8 seconds before either successfully connecting to the device or going back to sleep. This advertising interval achieves fast connection on both the Samsung Galaxy S5 and the iPhone 6, and the 8 s of maximum advertisements is the minimum time that we believed would help to ensure connection on phones that are not expedient at forming connections, such as the Motorola Moto X.

After forming a connection, Monoxalyze communicates CO readings and verification information back to the smartphone. The interval at which this communication occurs can be negotiated by the two parties in a connection according to the BLE specification. While the optimal connection interval would allow communication at a rate just exceeding the maximum data rate of Monoxalyze, in practice we find that many phones do not allow for connection interval negotiation, and always choose an interval that well exceeds the data rate of Monoxalyze.

5.3 Visible Light Authentication

Visible light authentication protocol provides a low-power out-of-band challenge response protocol for verifying the location of a Monoxalyze device. Specifically, it enables the smartphone to ensure that the Monoxalyze device it is connected to is also in its camera view while it performs facial recognition on the user. This verifies that the user identified by facial recognition is the same person who is exhaling through Monoxalyze. To implement the visible light authentication protocol, Monoxalyze has three front facing LEDs that are visible through holes in the case. These LEDs are shown in Figure 4.

Using out-of-band channels with a limited communication domain as a means of verifying location is not a new idea [34], and multiple devices have been created that show visible light communication to a phone's camera is possible [22, 33], however communicating from a low power embedded system to the phone's camera using small LEDs imposes a new set of constraints. These constrains include limiting the time that the LEDs are active to decrease total energy used during an authentication event. To achieve these goals we design a generalized BLE service for out-of-band challenge response and optimize the service to fit the needs of Monoxalyze.

In the generalized BLE challenge response service that we design, the peripheral exposes data rate, challenge data and start fields. To issue a challenge, the challenger writes the data rate and challenge data to their respective fields, and when it is ready to receive the challenge, it writes to the start field and then listens for the challenge on the appropriate out-of-band channel.

To optimize this service for Monoxalyze, we solve a few additional problems, which focus on lowering the power of the VLC communication protocol including making it easier for the phone to find and decode the LEDs and lowering the power of the data transmission itself.

To ease in decoding the transmitted signal, it is advantageous for the smartphone to find Monoxalyze in its view before issuing the challenge start command. This same mechanism allows the

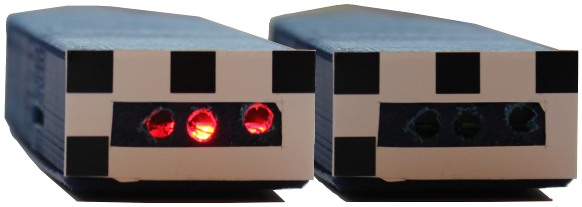

Figure 4: An end view of Monoxalyze with and without the center LED active. The black and white pattern allows the smartphone camera to localize and orient Monoxalyze in the camera view without requiring the LEDs to be powered. The smartphone then sends the challenge and expects a response from the LEDs. This shows how we reduce the energy use of visible light communication and make it feasible for small, battery-powered embedded systems to communicate with a smartphone camera.

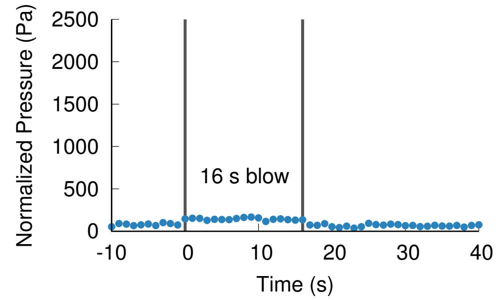

(a) Dry Air

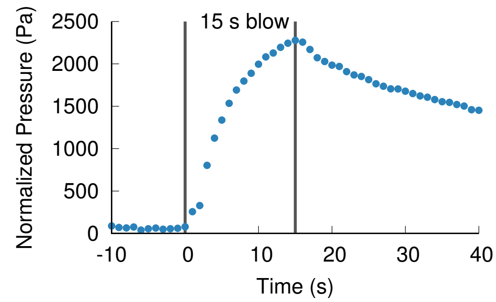

(b) Natural Breath

Figure 5: Comparing the pressure change induced by the movement of dry vs humid air through Monoxalyze. The vertical lines denote the start and stop of air movement respectively. This shows that natural humid breath causes significantly greater pressure increase (about 2 kPa) than dry air (about 50 Pa). This is caused by the increase in partial pressure of water vapor in the humid air. We use this significant change in pressure caused by humidity to verify exhalation during a breath test and wake Monoxalyze up from sleep in a user-transparent way.

smartphone to know if the challenge failed because there is no Monoxalyze in view or because there is a device that failed the challenge. Originally we implemented this pre-challenge alignment by repeating a known pattern on the LEDs, however, to save power and make template matching easier on the smartphone, we change this to be a static checkered pattern that also allows for orientation determination. This static pattern can be seen in Figure 4. The addition of this static pattern around the LEDs also makes it easy for the smartphone to find the state of the LEDs when only a single LED is illuminated. When using the LEDs for both data transmission and Monoxalyze localization in the image frame, it is difficult to determine which LED is illuminated if only one is turned on.

To synchronize the video frames with the data being transmitted, we could transmit synchronization information such as data rate and start conditions over BLE, and use inter-frame timing to maintain synchronization. Unfortunately this is difficult on many smartphones due to the varying frame rates and lag in frame rate reporting. This means that to synchronize the video processing with the data, the data signal must have an embedded clock. Manchester encoding is a well known scheme for self-clocking signals, however we observe that using three LEDS enables a lower power self-clocking signal encoding.

This signal encoding method uses On-Off-Keying and sends bits in three bit frames, in which each LED represents a single bit. The encoding mandates that every three bit data frame be different from the previous three bit data frame, with the start condition being represented by any non-zero data frame. This encoding scheme allows for seven possible frames in every time period (8 possible combinations of three LEDs, with the previous frame being invalid), and effectively communicates $log_2(7) = 2.8$ bits of information per frame. The downside of this system is that it does not enable the transmission of arbitrary data, however this limitation is not detrimental to our authentication system, and can still match the generic authentication protocol outlined above given that the smartphone only sends the Monoxalyze device challenges that meet this format. If in the future the smartphone application can record video at a more consistent frame rate, it could send arbitrary challenges to the device, and use this frame rate for signal recovery. It should not be difficult for the smartphone to generate a random challenge which meets these criteria at relatively short challenge lengths, such as the 12 bit challenges that we use, considering that 59% of 12 bit integers meet the limitations of the encoding. We evaluate the implications of this encoding scheme, its reliability, and its improvements over Manchester encoding in Section 7.4.

Given these optimizations, the final process is as follows: 1) the phone will process front camera frames until it detects the static pattern on Monoxalyze, 2) the phone will generate random numbers until it generates a valid 12 bit challenge, 3) the phone will send this challenge to Monoxalyze and send the start command, 4) Monoxalyze will blink back the challenge on the front LEDs using 3-bit frames of OOK data, 5) the phone will record the challenge response and process it to extract the challenge data, and 6) If the challenge fails then restart from step 1.

In our current application, this challenge response is unencrypted, and we believe that this is sufficient for Monoxalyze because challenges are random and not reused, and therefore are not secret.

5.4 Exhalation Verification

More important than being a wakeup mechanism, the pressure sensor inside of Monoxalyze is used to verify that a user is actually exhaling for a sufficient amount of time through the device. If this verification fails, a user could simply go through all the steps of using Monoxalyze without exhaling and receive a CO reading of the surrounding air. For the pressure sensor to accomplish this task, a person exhaling must sufficiently change the pressure inside the Monoxalyze gas chamber.

To ensure that the pressure inside the gas chamber changes during exhalation, we design the case so that the gas exit orifice is smaller than the input opening of the mouthpiece, however upon doing a test exhalation with an "artificial lung" (inflated bag with an exit tube), we find that the pressure inside the gas chamber does not increase. Testing exhalation with an actual person results in a great increase in pressure. The data from this experiment is shown in Figure 5b.

We hypothesize that most of the pressure increase is due to an increase in the humidity in the gas chamber. The pressure sensor reports a higher pressure with increasing humidity because the partial pressure of water vapor is higher than that of dry air. While this does not allow for a direct mapping of exhalation speed to sensed pressure, it does verify that a human is exhaling through Monoxalyze instead of an air hose or other fake gas flow.

During our testing we also observe a drop-off in reported pressure immediately after exhalation ceases because the humidity inside of the of gas chamber begins to equilibrate with the humidity in the surrounding air. We notice that this drop-off is proportional to the size of the exit orifice, which makes sense as a larger exit orifice allows a higher rate of equilibration. The drop off rate shown in Figure 5b is proportional to the exit slit that can be seen on Monoxalyze in Figure 2.

With these two considerations, exhalation verification is performed by mandating that pressure increases for the beginning of the exhalation period and remains stable for the rest of the exhalation period. The end of exhalation is denoted by detecting the drop off in pressure readings.

A user could attack this part of the system in several ways, including blocking the exit hole to artificially increase the pressure inside the gas chamber, or attempting to exhale at such a low velocity that they do not sufficiently empty their lungs. In the first case, the pressure behaves erratically and the processing algorithm flags the exhalation has ended early, which could later be flagged as an invalid test by the intervention program provider. In the later case, we observe that when exhaling at a velocity so low that the user has not completely exhaled in the 20-30s of required exhalation, the humidity does not build up sufficiently to cause the pressure sensor to report a pressure increase in the gas chamber.

5.5 CO Sensor Design Challenges

While a CO breathalyzer is not a new device [2, 5, 6, 27], most currently available CO breathalyzers are commercial products that often use expensive medical grade canister sensors, are large, and we have not found any open publications that discuss the complications of sensing a person's breath.

This may be because at first sight sensing CO concentration is not technically difficult. CO sensors are electrochemical sensors. They induce a current that is directly proportional to the CO concentration, which can be converted to a voltage, amplified, and read by an ADC. Indeed, this is exactly what our sensor does. However, this simplicity quickly breaks down when the CO sensor is exposed hot and humid exhaled breath.

High levels of humidity cause problems with both the CO sensors and the electrical components in Monoxalyze. To prevent the problems with the electrical components, we design a better case that seals the electrical components (except for the pressure sensor) from the air chamber using O-rings.

To address the problems presented by the cross-reactivity of humidity with the CO sensor, we take preventative measures to keep condensation and humidity from reaching the gas sensor element. Most commercial sensors do this with vapor-filtering mouth pieces, which are themselves larger than Monoxalyze. Instead, we create a stack of filters consisting of carbon film, active charcoal mesh, and a hydrophobic PTFE membrane. This design decreases exposure of the gas sensor to humidity without changing its response to dry CO test gas. It is future work to explore different filter stacks, and find the best combination of filters.

6. IMPLEMENTATION

We implement several iterations of the electronics and enclosures to evaluate Monoxalyze.

System. The circuit implementations are centered around the Nordic NRF51822, a Bluetooth SoC. This IC implements the BLE radio stack, serves as the ADC for the CO measurements, and performs all other communications and processing. All power calculations for the BLE connection models were based on the BLE stack implementation developed for the NRF51822 by Nordic Semiconductor.

Gas Sensor. The Monoxalyze design presented in this paper uses the Microcel CF CO Sensor [26]. This sensor was chosen for its small form factor. It presents slightly higher sensitivity to humidity than traditional canister CO sensors, and costs approximately the same. Ideally we would use the smaller and significantly less costly gas sensor discussed in Section 8.1, and indeed, several implementations of Monoxalyze were produced with this sensor. We used the TI LMP91000 in an unbiased, 3-lead amperic mode as the analog front end for the CO sensor. To keep the Microcel CO sensor from biasing during sleep and requiring a high stabilization time after wakeup, we short the internal P-Fet on the LMP91000 before sleep.

Pressure Sensor. We use the STMicroelectronics LPS25HB pressure sensor. Most pressure sensors function as described in Section 5.1 and Section 5.4, however a nontrivial fraction of the pressure sensors exhibit an entirely different response, by which their output appears to saturate. This could be due to an assembly process error such as excessive heat exposure. It may be advantageous to look into using a different pressure sensor in the future.

Image Processing. For image processing, we currently offload challenge-response videos to a computer. The processing algorithm uses multi-scale normalized cross-correlation to find the static template on the back of Monoxalyze and thresholds the values of the expected LED positions after finding the template. While the processing of videos will be slower on a mobile phone, the current challenge response protocol which transmits frames at 10Hz only requires the algorithm to process about 12 frames of data. On a compute,r the unoptimized processing takes approximately 1s, so a phone should easily be able perform this processing within a single exhalation test. There are also well known computer vision libraries for mobile phones such as OpenCV, which implement optimized version of all of the functions used in our processing algorithm.

Firmware. The software is a bare-metal implementation that runs on the Cortex-M0 in the Nordic Semiconductor nRF51822. The software is largely based off of Nordic's provided libraries.

Enclosure. The case is designed in Solidworks 2015 and has been 3D printed on both a Makerbot Replicator and Zortrax m200. The enclosure is designed to be compatible with BACTrack mobile breathalyzer mouthpieces.

Cost. The total cost of the prototype sensor is approximately $125 each. This cost includes all electrical components, the gas sensor, and the PCB. This does not include manufacturing of the case. This also does not take into account any price breaks that would come from scaling up production or negotiating with manufacturers.

Open Source. All firmware, image processing software, circuit designs, PCB layouts, and case designs are open source and available at www.github.com/lab11/monoxalyze.

7. EVALUATION

We evaluate Monoxalyze on three main dimensions: usability, data validity, and determination of smoking cessation. To evaluate usability, we consider user burden, battery life, size, and reliability of the pressure sensitive wakeup. To evaluate the validity of data collection, we explore the reliability of our user authentication. Finally, we enlist a third part to run a small pilot study with thirteen people who are attempting to quit smoking and compare our sensor's determination with ground-truth smoking status.

7.1 Energy Cost and Model

To evaluate the battery life of a Monoxalyze device we measure the energy associated with each event that occurs during its various operating modes. We predict the rate at which these events occur during normal use and use these predictions to estimate battery lifetime. We begin by evaluating the active use of Monoxalyze and then evaluate sporadic events such as false wakeups and pressure recalibration events. A summary of the energy cost of each event is presented in Table 2.

Event	Energy/ Event	Events/ Test	Energy/ Test
Advertisement (E_{adv})	190.5 µJ	2-200[a]	0.381-38.1 mJ
Connection Interval(E_{conn})	176.1 µJ	1500[a]	264.2 mJ
Pressure Sample(E_{psamp})	15.9 µJ	750	11.9 mJ
Pressure Read(E_{pread})	50.7 µJ	750	38.9 mJ
Visible Light Authentication(E_{vla})	55259 µJ	1	55.3 mJ
Gas Sensor(E_{gas})	1380 µJ	1	1.380 mJ

[a] This is phone dependent. We present an average or range of what we observed on the phones that we tested.

Table 2: Energy cost of each primitive of the system. These are the basic events that occur on Monoxalyze. Each event occurss at a different rate or number of times while Monoxalyze is active, and we assume Monoxalyze is active for 30 s/test to calculate the total energy/test. This breakdown allows us to calculate an expected system battery life.

Advertising. We measure the energy cost of a single advertisement to be $E_{\text{adv}} = 190.5$ µJ. From our exploration of phone connection rates, we choose a 40 ms advertising interval and advertise for a maximum of 8 seconds or 200 advertisements. Across five trials, we find that the iPhone 6 and Samsung Galaxy S5 reliably connect in only 2–3 advertisements, however the Motorola Moto X requires a median of 53 advertisements and a worst-case of 125 advertisements to connect. In the best case, the energy cost from 2 advertisements is $E_{\text{advmin}} = 381$ µJ, while the worst case, the 200 advertisement cutoff, incurs an energy cost of $E_{\text{advmax}} = 38.1$ mJ.

Connection. We measure the energy cost of a single BLE connection interval to be $E_{\text{conn}} = 176.1$ µJ. It should be noted that this is a standard Monoxalyze connection interval and that the energy for a connection interval is data dependent. On the phones that we test, connection intervals occur at 25-100 Hz. In Table 2 we choose an average connection interval of 50 Hz to calculate the total energy for a 30 s Monoxalyze test. The possible range of connection intervals that we observe provides a best case energy cost of $E_{\text{conntestmin}} = 132.0$ mJ and a worst case energy cost of $E_{\text{conntestmax}} = 528.3$ mJ over the course of a standard Monoxalyze reading.

VLC Authentication. We measure that a single LED draws a constant $P_{\text{led}} = 80.6$ mW of power. We know the average number of LEDs active in a three bit frame of our communication protocol is 1.71 LEDs. With knowledge that the bit rate of our implementation is 30 bps on 12 bit challenges, we know that 1.71 LEDs are on for 0.4 seconds, which means the average cost of a single VLC authentication is:

$$1.71 \text{ LEDs} \times 0.4 \text{ seconds} \times P_{\text{LED}} = 55.3 \text{ mJ}$$

Pressure Sensor Sampling. We measure that a single pressure sensor sample internal to the pressure sensor costs $E_{\text{psamp}} = 15.9$ µJ and that reading this pressure sample over I^{2}C costs another $E_{\text{pread}} = 50.7$ µJ. When Monoxalyze is awake and taking a reading the pressure sensor samples, and the MCU reads that sample, at 25 Hz which costs a total of

$$(E_{\text{psamp}} + E_{\text{pread}}) \times 25 \text{ Hz} \times 30 \text{ s} = 50.0 \text{ mJ}$$

over the course of a test. When Monoxalyze is sleeping the pressure sensor samples at 1 Hz waiting for a pressure wakeup event, and this 1 Hz sampling averages to a constant $P_{\text{pidle}} = 15.9$ µW of power.

Gas Sensor. If the gas sensor is active it draws a constant $P_{\text{gas}} = 46$ µW. Using the Microcel CF sensor, the gas sensor can be put in deep sleep during sleep mode.

False Wakeup Event. A false wakeup occurs when a pressure threshold event wakes up Monoxalyze, but no phone is listening to connect to Monoxalyze. A false wakeup consists of 200 advertisements. This totals to a cost of $E_{\text{wakeup}} = 38.1$ mJ.

Idle Power. The idle power draw of Monoxalyze consists of the 1 Hz pressure sensor sampling along with the sleep and leakage power of the other chips on the board. This totals to $P_{\text{idle}} = 55.9$ µW of idle power.

7.1.1 Battery Life Calculation

Using the energy costs outlined in Table 2 and their frequencies as described in Section 7.1, we construct two models of Monoxalyze energy usage—one best case model and one reasonable worst case model—and use these models for battery life approximation.

For our best case model, we assume that a user must exhale into Monoxalyze twice per day and that each test costs 372.1 mJ of energy. We assume that Monoxalyze experiences no false wakeups throughout the day, which our data shows is a reasonable assumption. This means that the total energy usage of Monoxalyze for one day in the best case is

$$\begin{array}{rlcrcl} & E_{\text{test}} & \times & 2 \text{ tests} & = & 0.744 \text{ J} \\ + & P_{\text{idle}} & \times & 86340 \text{ s} & = & 4.826 \text{ J} \\ \hline & E_{day} & & & = & 5.57 \text{ J} \end{array}$$

To provide some intuition for the lifetime range of Monoxalyze, we construct a "reasonable worst-case" scenario. The goal of this exercise is to provide intuition, as such some liberties are taken with estimates. We increase user testing to three tests per day and estimate that each test now costs 400.5 mJ because the user's phone averages 150 advertisements to connect. We also assume that Monoxalyze experiences 10 false wakeups due to elevation changes throughout the day.

The total energy usage of Monoxalyze throughout this day is

	E_{test}	$\times$	3 tests	=	1.201 J
+	E_{wakeup}	$\times$	10 wakeups	=	0.382 J
+	P_{idle}	$\times$	86340 s	=	4.826 J
	E_{day}			=	6.41 J

Monoxalyze has a 0.148 Wh battery. This is equivalent to 532 J of energy. With this battery capacity, Monoxalyze would last 95 days in the best case scenario and 83 days in the more realistic scenario. These are both acceptable lifetimes since Monoxalyze is rechargeable.

7.2 Wakeup Reliability

To test wakeup reliability we run two trials. In the first trial three users exhale into Monoxalyze 15 times. These exhalation events happen within two minutes of one another. Over this trial, we observe a wakeup success rate of 87% and an average wakeup time of 2.34 s. While this success rate might seem low, it should be noted that the trial described above is the antagonistic case. Blowing into Monoxalyze repeatedly causes humidity to build up, meaning there is less humidity change, and therefore a lower pressure change on the next exhalation event.

To show that the previous trial is the antagonistic case, we exhale into Monoxalyze 10 more times with approximately 10 minutes between each trial. Monoxalyze has a 100% wakeup success rate and an average wakeup time of 1.4 s across these tests.

The problems caused by many subsequent exhalation events should not be a factor during normal use of Monoxalyze. We expect most intervention participants to exhale through Monoxalyze 2–3 times each day.

7.3 Exhalation Verification

Evaluating exhalation verification in a quantitative way is difficult because it requires human exhalation, which is often inconsistent from person to person. However, we put Monoxalyze in several scenarios and output whether it detects an exhalation event or not based on the algorithm described in Section 5.4. To perform all of the testing below, we wait at least 10 minutes between tests for the humidity within the gas chamber to return to equilibrium with the ambient humidity.

First we exhale at full volume for 30 s into Monoxalyze 10 times. Monoxalyze correctly identifies all exhalation events and correctly indicates when the user stops exhaling. We then repeat this test by attempting to blow the minimum volume of air through Monoxalyze while ensuring that Monoxalyze reports that we are exhaling. From a qualitative viewpoint, this amount of air always nearly emptied the test subject's lungs, which is sufficient for a CO measurement.

We also test the antagonistic cases by not blowing into Monoxalyze and by blocking the Monoxalyze exit port and attempting to hold pressure in the gas chamber. Both of these tests report inconsistent behavior, as do tests in which we inhale through the Monoxalyze gas chamber instead of exhaling.

7.4 User Authentication

First we evaluate our VLC encoding protocol for its power efficiency and compare that efficiency to Manchester encoding. The reason for using Manchester encoding as a baseline is discussed in Section 5.3. To make our evaluation independent of power itself, we will compare the two encoding schemes on the metric of LEDs active/bit of information transmitted.

With our encoding scheme, on the first clock period, we expect no more than $\frac{12}{7} = 1.7$ LEDs to be active. This is the worst case because the first frame cannot be zero, so as the length of the data increases, the average number of expected active LEDs decreases. We also know that there are seven possible LED configurations per frame, which is equivalent to $log_2(7) = 2.8$ bits per clock period, or $\frac{1}{2.8}$ clock periods per bit. By multiplying these two numbers, we see that our scheme achieves $\frac{1.7}{2.8} = 0.61$ active LEDs per bit in the worst case.

In contrast, we know that Manchester encoding takes two clock periods for every one bit of information, and that every clock period has 0.5 active LEDs. This means that Manchester Encoding would have 1 active LED per bit, 64% more power than our encoding. While there may be even more power efficient encodings than the one designed for Monoxalyze, this protocol is a simple, efficient way to communicate from a low power device to a smartphone for a challenge response protocol that does not require arbitrary data to be communicated.

To evaluate the reliability of our image processing we generate 30 random 12 bit challenges, transmit those challenges at 10 Hz per frame, or 30 bps, and decode them at two different distances. With the face filling most of the image frame (phone 5-8 inches from the person's face), this evaluation yields only one failed challenge with only one bit of error, an equivalent bit error rate of 0.1%. At the longer distance, none of the challenges succeed. On one hand this shows that the image processing could be improved, however this mode of failure is an inconvenience to the user rather than a security issue, and the average use case of the phone reasonably close to the person's face is still quite usable.

We test challenges at the closer distance under different lighting conditions and as long as the conditions are not extreme with bright lights or very similar patterns in the frame, we observe these changes do not decrease rate of success.

Many people have implemented facial recognition and verification algorithms [29, 39], and there are many APIs that provide facial recognition [12, 13, 24]. A review study showed that the best facial recognition algorithms can achieve over 97% accuracy in small sample groups [16]. More recently, companies like Amazon and MasterCard have started to use mobile facial recognition to verify payments, and Amazon has even filed a patent to use biometric signals such as blinking to protect against spoofing [31, 44]. These technologies could be integrated into the Monoxalyze system to improve user authentication. Monoxalyze currently uses the Face++ API to perform facial verification, however the exact evaluation of this API is part of future work due to the already established metrics on the success of facial recognition algorithms.

7.5 Smoking Cessation Determination

We perform a preliminary test on smoking cessation determination by collecting data from 13 smokers or ex-smokers that are patients at a smoking cessation clinic. All participants' data are anonymized; neither names, nor times of visits are collected. One user reports that Monoxalyze connects to the phone, but does not report sensor readings. From the user report, we attribute this failure to an error in the phone application and not the Monoxalyze hardware. The following evaluation considers data from the remaining 12 participants.

For this trial, participants are directed to exhale through a Smokerlyzer commercial breathalyzer [5], then exhale through Monoxalyze. Table 3 shows that Monoxalyze has a sensitivity + specificity of 1.75 at an abstinence threshold of both 6 and 7 ppm. This result shows that Monoxalyze is capable of performing smoking cessation determination. Note that this sensitivity + specificity metric is com-

Abstinence Threshold PPM	1	2	3	4	5	**6**	**7**	8	9	10	11
Sensitivity	0.75	0.75	0.75	0.75	0.75	**0.75**	**0.75**	0.25	0.25	0.25	0.25
Specificity	0.88	0.88	0.88	0.88	0.88	**1.00**	**1.00**	1.00	1.00	1.00	1.00
Sensitivity+Specificity	1.63	1.63	1.63	1.63	1.63	**1.75**	**1.75**	1.25	1.25	1.25	1.00

Table 3: This table shows Monoxalyze's ability to differentiate between smokers and non-smokers. We take preliminary CO measurements from 13 participants. Of these participant, one user failed to collect data due to a phone application error. Of the remaining participants, there are 4 smokers and 8 non-smokers. We then apply a calibration to the raw CO measurements and analyze the data against the ground truth provided by the Smokerlyzer. We determine that an abstinence threshold of both 6 ppm and 7 ppm yield the best sensitivity+specificity, classifying one smoker as a non-smoker. This success at classification is only slightly lower than the classification success of the Smokerlyzer and custom devices presented in [27], and it shows the potential of Monoxalyze to successfully classify smokers from non-smokers with more testing.

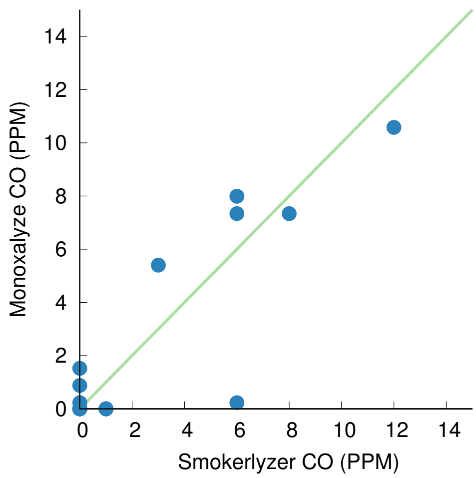

Figure 6: A comparison of Smokerlyzer CO measurements and Monoxalyze CO measurements. taken from a trial of 13 participants at a smoking cessation clinic. One of 13 trials failed to collect data, so 12 data points are plotted. The diagonal line shows the ideal correlation between the Smokerlyzer and Monoxalyze. This shows that Monoxalyze is reasonably accurate compared to the Smokerlyzer, which is considered the gold standard for smoking cessation determination.

monly used to determine the quality of CO breathalyzers, and that the Smokerlyzer only achieves a sensitivity + specificity of 1.83 in a comparable study [27].

In Figure 6 we compare the accuracy of Monoxalyze against the Smokerlyzer as ground truth. We find that the Monoxalyze readings track those of Smokerlyzer well except for one of the readings. In Section 8.1 we consider additional passive and active means of reducing the impact of humidity on our sensor readings.

8. DISCUSSION

In this section, we discuss current issues and future work on Monoxalyze. We also present alternatives to current design decisions that may improve performance of Monoxalyze.

8.1 Improving Cessation Determinism

The size of our user study is rather small, and a larger user study could improve the evaluation of smoking cessation determination in Monoxalyze. There is also possible future work on selecting humidity filters and characterizing how changes in chamber humidity affect the output signal of the gas sensor. If the cross-reactivity with humidity could be better characterized, we may be able to use a humidity sensor to correct for these effects.

8.2 Smaller and Lower Cost CO Sensors

Several of the explicit goals of Monoxalyze are to be smaller and more accessible than commercial CO breathalyzers. The only other CO sensor on the market that meets our size constraints is the SPEC CO sensor [35], which is also a small fraction of the cost of the Microcel sensor we use. Using the SPEC sensor would allow Monoxalyze to be more affordable hence pervasive. Unfortunately all of our testing with the SPEC sensor has shown it to be an order of magnitude less sensitive and just as cross-reactive to humidity as the Microcel sensor, making it poor at determining smoking cessation.

Ideally, future work could demonstrate the viability of using the SPEC sensor in Monoxalyze, and we have ongoing experiments that are attempting to use a fusion of pressure, temperature, and humidity sensors to account for the contributions of humidity and pressure to the CO sensor readings. We believe that actively accounting for these environmental variations and deploying sensors with the correct set of filters may allow the SPEC sensors to be used in future versions of Monoxalyze.

8.3 Authentication

Several improvements could be made to our proposed transitive trust authentication. Specifically, improvements could be made to the visible light authentication processing algorithms. There are also improvements to be made in the verifiable facial recognition space. We could improve our evaluation by allowing users to freely attack the system, and this is part of our proposed future work.

Visible light authentication could be improved by increasing the performance of our processing algorithms so that they can be run in real time on the mobile phone, rather than incurring some delay before the test is authenticated. While the current system is not particularly inconvenient to the use, improved processing would allow for multiple challenge-response attempts before declaring the test failed, and could mitigate the small fraction of time that we incorrectly receive the transmitted response.

When using facial recognition to identify users in real-world settings, several attacks could be employed, such as using a mask to conceal identity. While some of these threats may be addressed by using biometric markers such as blinking [44], we believe that other techniques, such as moving the phone in a semicircle around the face to calculate facial structure, could eliminate some of the potential attack vectors.

Lastly, we should allow users and other researchers to attack the system. While in theory the proposed Monoxalyze design protects against our threat model, allowing others to attack the system is likely to reveal opportunities to improve the security of Monoxalyze.

9. CONCLUSIONS

Smoking cessation intervention programs need a reliable and scalable reporting mechanism to validate smoking cessation. Existing methods of cessation detection do not extend to deployments outside of a clinical setting. We present Monoxalyze to solve these problems. Monoxalyze is a keychain-sized carbon monoxide breathalyzer that coordinates with a smartphone to verify that an individual has exhaled through a specific device. This is achieved by verifying exhalation with a pressure sensor, identifying the user with facial recognition, and localizing a Monoxalyze device to the user with a low-power visible light challenge-response protocol. Solving the verifiability problem allows for deployments outside of an in–patient, monitored clinic. We also leverage a pressure sensor to provide user-transparent wakeup and analyze BLE connection patterns to optimize power for Monoxalyze, resulting in a lifetime of at least 80 days. Within these limitations for the system and its smallest-in-class form factor, Monoxalyze is still capable of detecting smokers with a 92% success rate. With the new capabilities provided by this device, researchers can undertake new studies in smoking cessation, leading to health benefits for individuals worldwide.

10. ACKNOWLEDGMENTS

This work was supported in part by the TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP), a Semiconductor Research Corporation program sponsored by MARCO and DARPA. This material is based upon work partially supported by the National Science Foundation under grant CPS-1239031, and by the NSF/Intel Partnership on Cyber-Physical System (CPS) Security and Privacy under award proposal title "Synergy: End-to-End Security for the Internet of Things," NSF proposal No. 1505684. We would also like to the thank Matthew Bars, CEO of IntelliQuit and Director of the New York City Fire Department Tobacco Treatment Program for helping to collect samples and validate the Monoxalyze sensor.

11. REFERENCES

[1] Bedfont Scientific Ltd. Compact smokerlyzer. http://covita.net/pdfs/Datasheets/compactdatasheet_iss1.pdf.

[2] Bedfont Scientific Ltd. ico smokerlyzer. http://www.bedfont.com/shop/smokerlyzer/ico_smokerlyzer.

[3] Bedfont Scientific Ltd. Micro smokerlyzer. http://www.bedfont.com/shop/smokerlyzer/micro.

[4] Bedfont Scientific Ltd. Pico smokerlyzer. http://www.bedfont.com/shop/smokerlyzer/pico.

[5] Bedfont Scientific Ltd. Smokerlyzer. http://www.bedfont.com/smokerlyzer.

[6] Breath E-Z Systems Inc. Co slueth. http://www.testbreath.com/co.asp.

[7] K. Cahill and R. Perera. Competitions and incentives for smoking cessation. *The Cochrane Library*, 2008.

[8] S. Chapman and W. Smith. Deception among quit smoking lottery entrants. *American journal of health promotion*, 8(5):328–330, 1994.

[9] S. J. Curry, E. H. Wagner, and L. C. Grothaus. Evaluation of intrinsic and extrinsic motivation interventions with a self-help smoking cessation program. *Journal of Consulting and Clinical Psychology*, 59(2):318 – 324, 1991.

[10] S. Deveci, F. Deveci, Y. AÃğik, and A. Ozan. The measurement of exhaled carbon monoxide in healthy smokers and non-smokers. *Respiratory Medicine*, 98(6):551 – 556, 2004.

[11] R. J. Donatelle, D. Hudson, S. Dobie, A. Goodall, M. Hunsberger, and K. Oswald. Incentives in smoking cessation: Status of the field and implications for research and practice with pregnant smokers. *Nicotine & Tobacco Research*, 6(Suppl 2):S163–S179, 2004.

[12] Face++. Face++ facial verification. http://www.faceplusplus.com/api-overview/.

[13] FaceR API. FaceR API. http://api.animetrics.com/documentation.

[14] R. Ford, D. Tappin, P. Schluter, and C. Wild. Smoking during pregnancy: how reliable are maternal self reports in new zealand? *Journal of epidemiology and community health*, 51(3):246–251, 1997.

[15] S. C. Gorber, S. Schofield-Hurwitz, J. Hardt, G. Levasseur, and M. Tremblay. The accuracy of self-reported smoking: a systematic review of the relationship between self-reported and cotinine-assessed smoking status. *Nicotine & Tobacco Research*, 11(1):12–24, 2009.

[16] P. J. Grother, G. W. Quinn, and P. J. Phillips. Report on the evaluation of 2D still-image face recognition algorithms. *NIST interagency report*, 7709:106, 2010.

[17] S. T. Higgins, Y. Washio, S. H. Heil, L. J. Solomon, D. E. Gaalema, T. M. Higgins, and I. M. Bernstein. Financial incentives for smoking cessation among pregnant and newly postpartum women. *Preventive medicine*, 55:S33–S40, 2012.

[18] P. Hill, N. Haley, and E. Wynder. Cigarette smoking: carboxyhemoglobin, plasma nicotine, cotinine and thiocyanate vs self-reported smoking data and cardiovascular disease. *Journal of chronic diseases*, 36(6):439–449, 1983.

[19] M. J. Jarvis, M. A. Russell, and Y. Saloojee. Expired air carbon monoxide: a simple breath test of tobacco smoke intake. *BMJ*, 281(6238):484–485, 1980.

[20] L. A. Jason, S. Jayaraj, C. C. Blitz, M. H. Michaels, and L. E. Klett. Incentives and competition in a worksite smoking cessation intervention. *American Journal of Public Health*, 80(2):205–206, 1990.

[21] M. A. Javors, J. P. Hatch, and R. J. Lamb. Cut-off levels for breath carbon monoxide as a marker for cigarette smoking. *Addiction*, 100(2):159–167, 2005.

[22] Y.-S. Kuo, P. Pannuto, K.-J. Hsiao, and P. Dutta. Luxapose: Indoor positioning with mobile phones and visible light. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking*, MobiCom '14, pages 447–458, New York, NY, USA, 2014. ACM.

[23] Y.-S. Kuo, T. Schmid, and P. Dutta. Hijacking power and bandwidth from the mobile phone's audio interface. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, SenSys '10, pages 389–390, New York, NY, USA, 2010. ACM.

[24] Lambda Labs. Lambda labs face recognition API. https://lambdal.com/face-recognition-api.

[25] K. D. Leeks, D. P. Hopkins, R. E. Soler, A. Aten, S. K. Chattopadhyay, T. F. on Community Preventive Services, et al. Worksite-based incentives and competitions to reduce tobacco use: a systematic review. *American journal of preventive medicine*, 38(2):S263–S274, 2010.

[26] C. T. Ltd. Microcel cf - carbon monoxide. http://www.citytech.com/PDF-Datasheets/microcelcf.pdf.

[27] S. E. Meredith, A. Robinson, P. Erb, C. A. Spieler, N. Klugman, P. Dutta, and J. Dallery. A mobile-phone-based breath carbon monoxide meter to detect cigarette smoking. *Nicotine & Tobacco Research*, 16(6):766–773, 2014.

[28] E. T. Middleton and A. H. Morice. Breath carbon monoxide

as an indication of smoking habit. *CHEST Journal*, 117(3):758–763, 2000.

[29] I. Naseem, R. Togneri, and M. Bennamoun. Linear regression for face recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(11):2106–2112, 2010.

[30] U. D. of Health, H. Services, et al. The health consequences of smoking-50 years of progress: a report of the surgeon general. *Atlanta, GA: US Department of Health and Human Services, Centers for Disease Control and Prevention, National Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health*, 17, 2014.

[31] A. Petroff. Mastercard launching selfie pay app, 2016.

[32] R. Pojer, J. Whitfield, V. Poulos, I. Eckhard, R. Richmond, and W. Hensley. Carboxyhemoglobin, cotinine, and thiocyanate assay compared for distinguishing smokers from non-smokers. *Clinical Chemistry*, 30(8):1377–1380, 1984.

[33] N. Rajagopal, P. Lazik, and A. Rowe. Hybrid visible light communication for cameras and low-power embedded devices. In *Proceedings of the 1st ACM MobiCom Workshop on Visible Light Communication Systems*, VLCS '14, pages 33–38, New York, NY, USA, 2014. ACM.

[34] N. Sastry, U. Shankar, and D. Wagner. Secure verification of location claims. In *Proceedings of the 2nd ACM workshop on Wireless security*, pages 1–10. ACM, 2003.

[35] S. Sensors. Spec carbon monoxide sensor. http://www.spec-sensors.com/product-category/gases/carbon-monoxide-sensors/.

[36] D. Shipton, D. M. Tappin, T. Vadiveloo, J. A. Crossley, D. A. Aitken, J. Chalmers, et al. Reliability of self reported smoking status by pregnant women for estimating smoking prevalence: a retrospective, cross sectional study. *Bmj*, 339:b4347, 2009.

[37] W. W. Stoops, J. Dallery, N. M. Fields, P. A. Nuzzo, N. E. Schoenberg, C. A. Martin, B. Casey, and C. J. Wong. An internet-based abstinence reinforcement smoking cessation intervention in rural smokers. *Drug and alcohol dependence*, 105(1):56–62, 2009.

[38] D. Tappin, L. Bauld, D. Purves, K. Boyd, L. Sinclair, S. MacAskill, J. McKell, B. Friel, A. McConnachie, L. de Caestecker, C. Tannahill, A. Radley, and T. Coleman. Financial incentives for smoking cessation in pregnancy: randomised controlled trial. 350, 2015.

[39] M. A. Turk and A. P. Pentland. Face recognition using eigenfaces. In *Computer Vision and Pattern Recognition, 1991. Proceedings CVPR'91., IEEE Computer Society Conference on*, pages 586–591. IEEE, 1991.

[40] E. Vartiainen, T. Seppälä, P. Lillsunde, and P. Puska. Validation of self reported smoking by serum cotinine measurement in a community-based study. *Journal of Epidemiology and Community Health*, 56(3):167–170, 2002.

[41] K. G. Volpp, A. B. Troxel, M. V. Pauly, H. A. Glick, A. Puig, D. A. Asch, R. Galvin, J. Zhu, F. Wan, J. DeGuzman, E. Corbett, J. Weiner, and J. Audrain-McGovern. A randomized, controlled trial of financial incentives for smoking cessation. *New England Journal of Medicine*, 360(7):699–709, 2009. PMID: 19213683.

[42] L. E. Wagenknecht, G. L. Burke, L. L. Perkins, N. J. Haley, and G. D. Friedman. Misclassification of smoking status in the cardia study: a comparison of self-report with serum cotinine levels. *American Journal of Public Health*, 82(1):33–36, 1992.

[43] R. West. The effect of duration of breath-holding on expired air carbon monoxide concentration in cigarette smokers. *Addictive behaviors*, 9(3):307–309, 1984.

[44] W. Yicong and H. Xu. Image analysis for user authentication, 2016. US Patent 9,202,105.

[45] S.-H. Zhu, T. Melcer, J. Sun, B. Rosbrook, and J. P. Pierce. Smoking cessation with and without assistance: a population-based analysis. *American journal of preventive medicine*, 18(4):305–311, 2000.