Experimental Security Analyses of Non-Networked Compact Fluorescent Lamps: A Case Study of Home Automation Security

Temitope Oluwafemi[1], Sidhant Gupta[2], Shwetak Patel[1, 2], Tadayoshi Kohno[2]

[1]*Elec. Eng.*, [2]*Comp. Sci. & Eng.*

University of Washington

Seattle, WA

{oluwat, sidhant, shwetak, kohno}@uw.edu

ABSTRACT

Background. With a projected rise in the procurement of home automation systems, we experimentally investigate security risks that homeowners might be exposed to by compact fluorescent lamps (CFL), where the lamps themselves do not have network capabilities but are controlled by compromised Internet-enabled home automation systems.

Aim. This work seeks to investigate the feasibility of causing physical harm—such as through the explosion of CFLs—to home occupants through an exploited home automation system.

Method. We set up a model of a compromised automated home; placing emphasis on a connected Z-Wave enabled light dimmer. Four distinct electrical signals were then applied to two different brands of CFLs connected to a Z-Wave enabled light dimmer until they popped[1] or gave way[2].

Results. Three of ten CFLs on which we conducted our experiments popped, although not to the degree of explosions we expected. The seven remaining CFLs gave way with varying times to failure indicating process and design variations. We did find that it was possible to produce fluctuations at an appropriate frequency to induce seizures. We were also able to remotely compromise a home automation controller over the Internet. Due to timing constraints, however, we were only able to compromise the light bulbs via an adversary-controlled device using open-zwave libraries, and not via the compromised controller.

Conclusions. Our results demonstrated that it will be hard for an attacker to use the described methods to harm homeowners, although we do demonstrate the possibility of attacks, particularly if the homeowner suffers from epilepsy. However, and more importantly, our work demonstrates that non-networked devices—such as light bulbs—might be connected to networked devices and hence can be attacked by remote adversaries.

[1] We define popped as the visual or auditory observance of a spark in the CFL.

[2] The term "give way" refers to the normal failure of a CFL without a spark.

General Terms

Experimentation, Measurement.

Keywords

Home automation systems, cyber-physical systems, computer security, cyber-physical security, compact fluorescent lamps, CFLs

1. INTRODUCTION

To date, few experimental computer security research efforts have focused on computer systems that interact directly with the physical world—the so-called cyber-physical systems. There are, of course, some exceptions, e.g., various software attacks have been successfully demonstrated on cars, printers, robots and pacemakers with physical consequences as shown in [1, 3, 4, 10 and 11]. However, we argue that the field of experimental computer security research for cyber-physical systems is still in its infancy. This is partly due to the fact that cyber-physical systems are just emerging on the commercial market, but the greater challenge has been how to conduct research in this space. Significant, important issues can and do arise when attempting to experimentally evaluate the security of a cyber-physical system—issues that are not normally encountered, at least not in the same form—when experimenting with conventional non-cyber-physical systems. For example, is it possible to reconstruct the environment for the cyber-physical system in a laboratory setting in sufficient detail in order to ensure experimental validity? And is it possible to conduct the experiments in a way that does not jeopardize anyone's safety?

In this work we describe experiments that we conducted with an emerging class of cyber-physical systems: home automation systems. Many home automation systems already exist in the market, and recent worldwide market forecasts by Berg Insight claim that revenues generated through the sales and purchases of home automation units will grow at a compound annual rate of 33% from $2.3 billion USD in 2010 to nearly $9.5 billion USD in 2015 [19]. Home automation systems are often Internet-connected, and indeed—as an example of such connectivity—the number of cellular connections used by home automation units are expected to grow worldwide at a compound rate of 85.6% from 0.25 million in 2010 to 5.5 million connections in 2015 [19]. Home automation systems allow homeowners to control appliances—e.g., lights or ovens—from another device (such as a laptop)

within the home, or even over the Internet from a mobile device.

We began by obtaining two mainstream home automation systems and subjected them to a number of experiments. We describe briefly the totality of our work since we believe that it is important to understand the full context for our research, but foreshadow here that the bulk of this paper is focused on our experimentation with a seemingly unlikely target: light bulbs. Returning to the full context, we experimentally found that the home automation systems we acquired are vulnerable to remote attacks. We experimentally verified that an attacker—even from someplace outside a home, i.e., over the Internet—could violate the sanctity of the home by, for example, turning on or off home automation-connected devices (like light dimmers and HVAC systems) and even unlocking a home's front door or disabling a networked alarm system. We also found that an attacker could learn which devices are in a home and connected to the home automation network, thereby violating the homeowner's privacy. We also found that an attacker could control switches and dimmers in the home. While we identified and experimentally demonstrated these vulnerabilities with the home automation systems that we purchased, we note that others have made similar observations before, e.g., [5 and 8].

One of the capabilities mentioned above—that an Internet-connected attacker can remotely control switches and dimmers—may not sound significant at first. But herein lies what we believe is a fundamentally interesting property: there *is* the potential for an attacker to affect a device plugged into an outlet by maliciously controlling the outlet in certain ways. Certainly an attacker could use this capability to turn something connected to the outlet on and off or alter the brightness of a light bulb using a dimmer. While such actions might initially seem to only create nuisances for home occupants, after further contemplation, we began to speculate on whether an attacker could also use this simple capability to enact significantly more *physical* damage to the home environment. Concretely, one question we asked was: would it be possible for an attacker to make a light bulb connected to the network-controlled outlet explode?

Modern lighting solutions, such as CFLs and LED lamps, are designed to be efficient and thus increasingly make use of sophisticated electronic circuitry when compared to traditional incandescent light bulbs. We hypothesized that by altering the supply voltage characteristics to such devices, they could be made to operate beyond safe specifications of the electronic circuitry. We argue that knowing whether it would be possible to explode a light bulb remotely would be valuable for the computer security community. If possible, then defenses would need to be created before home automation systems become more ubiquitous and the risks increase.

Of course, we could not enter into an investigation of "can we experimentally explode light bulbs" lightly. In fact, we nearly did not proceed with this line of investigation because we did not know how to proceed both safely and in a cost effective manner. For example, how would we contain an explosion, should one occur? And how would we handle the leak of chemicals, should the physical damage to a light bulb cause some of its internal chemicals to leak. Fortunately, after significant research into possible options, we did derive a method. We use a glove box, which provides a controlled and well-ventilated environment to help contain and clean up hazardous materials. Another thing we learned to be cautious about during our experiments: the potential to induce seizures by fluctuating power to a light bulb.

While we did end up making light bulbs "pop", the "pops" were not nearly as significant as the worst-case explosions that we had feared. We also found that an attacker can cause the light bulbs to oscillate at a frequency known to cause seizures. From a security perspective—and the perspective of the homeowner—these results provide valuable insight into how secure these systems and connected devices are and the scale of attacks that can be mounted against them. The results indicate that it will be harder for an attacker to use these exact techniques to harm homeowners, such as exposing the occupants of the home to the mercury content of CFLs. However, on a more serious note, the results clearly demonstrate that it *is* possible for a remote attacker to compromise something as simple as a light bulb—a technology that, by design, has no network connectivity itself. We view this observation as an important contribution of the paper, with the other main contributions being the experimental methodologies we discuss. This observation is an important contribution because it provides proof of plausibility that—in the future—other devices *without network connections* might be found vulnerable to *network-based* compromise in the future.

Stepping back, we observe that cyber-physical systems are becoming increasingly prevalent. As such, we expect to see increasing interest in experimentally evaluating the security properties of such cyber-physical systems. But, if these systems are vulnerable to security compromises, then the experiments—if successful—have the potential to cause harm to the experimental environment, and possibly even to the experimenter. Hence, we believe that the foundations we lay in this paper may be of value for future cyber-physical systems researchers.

In the following section, we define the problem we aim to solve. Section 3 gives a brief background on home automation systems, CFLs and related work. Section 4 presents a detailed analysis of the security vulnerabilities discovered in the home automation systems we examined. Section 5 presents an overview of the method describing our work with CFLs. We will conclude this paper by

examining the results from our experiments and discuss potential future directions for all stakeholders.

2. PROBLEM BEING SOLVED

The purpose of our research is to gauge the level of impact that a remote attacker might have on the inhabitants of an automated home, particularly in regards to manipulating appliances, like CFLs, that do not have network capabilities. We specifically sought to explore the possibility of causing physical harm through the application of known electric signals to CFLs controlled through wireless light dimmers.

Our hypotheses is based on the incorrect use of non-dimmable CFLs with wirelessly controlled light dimmers. Since most CFLs cannot be dimmed using a traditional triac-based dimmer[3], manufacturers may not test against such specifications (using a CFL in a dimmer) and/or guard against these situations. There are newer CFLs that can be dimmed and these bulbs sense the dim level and internally regulate the power to the bulb. Our focus in this paper is on standard CFLs. Clearly one could simply use a non-dimming appliance module, but our assumption is a person might use a dimmable module without knowing the consequences. Consequently, there is a possibility of a current spike in non-dimmable CFLs used with dimmers that can result in fires. Hence, we hypothesize that an attacker can mount an attack to cause an explosion with the possibility of starting a fire and/or releasing harmful mercury contents of CFLs when connected to remotely compromised and controlled light dimmers.

To investigate the plausibility of our hypothesis, we describe experiments using open-zwave libraries to control Z-Wave enabled light dimmers with connected CFLs. We conduct these experiments in a glove box to provide a shield from shattered glass and to contain the mercury content of CFLs, in the event of an explosion.

As with most computer security vulnerability efforts, the results of this research can inform the design of future home automation systems and/or light bulbs (and other devices that might connect to home automation systems). We believe that now is the time to perform such research, before these systems become ubiquitous and the risks of any (possibly unknown) vulnerability increases.

3. BACKGROUND AND RELATED WORK

3.1 Home Automation

Most home automation systems consist of a primary controller that controls a variety of connected secondary nodes which include but are not limited to, door locks, alarm systems, HVAC and sprinkler systems, light modules (dimmers), and energy monitoring nodes as shown in Figure 1. Although some home automation systems use WiFi for communications between secondary nodes, data is usually sent through a low power and low data rate wireless communication standard such as Z-Wave or ZigBee. It is also usually the case that most primary controllers are equipped with both WiFi and Z-Wave or ZigBee for added connectivity to the Internet.

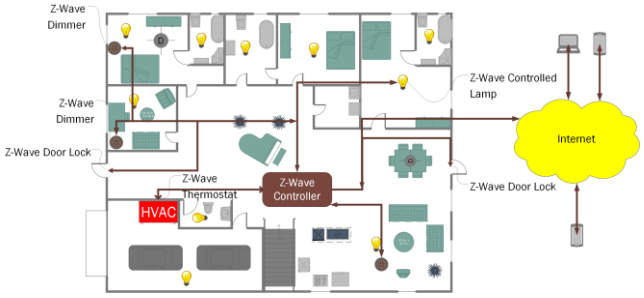

Figure 1: Home automation model.[4]

3.2 Compact Fluorescent Lamps

CFLs are fast becoming the standard for electric bulbs as countries around the world are beginning to phase out incandescent bulbs due to their power inefficiency. CFLs provide about seventy-five percent savings in energy when compared to incandescent bulbs. Newer CFLs are mostly integrated with electronic ballasts, while older models use large and heavy magnetic ballasts [15].

Most common CFLs integrated with electronic ballasts have more complex circuitry and active electronic components than incandescent light bulbs. Standard CFLs are also not supposed to be used along with dimmers as the current drawn by the lamps increases by a magnitude of about five times their normal operation [6]. There have been instances of fires caused by using CFLs with dimmers [18 and 20]. Furthermore, most CFLs contain about 3-5mg of mercury, which is harmful and constitutes environmental waste.

Given the composition of CFLs and their mode of operation, they appear to be appealing targets for a potential attacker with an intent to physically harm occupants of automated homes.

Moreover with automated homes, attackers have the ability to remotely control CFLs connected to dimmers and light switches by sending arbitrary signals, as we demonstrate in Sections 4 and 5. We again stress that such adversarial capabilities are possible *even though* the CFLs themselves do *not* have any built-in network capabilities.

3.3 Related Work

As mentioned, others have investigated the security and privacy of other classes of cyber-physical systems. For example, Checkoway and others in [1] were able to demonstrate the remote compromise of automobiles, providing attackers with the ability to remotely disable brake systems and eavesdrop on in-vehicle conversions. They were able to highlight consumer safety concerns and

[3] See Section 3.2.

[4] The house is assumed to be retrofitted with a variety of automated appliances, sensors, actuators and controllers.

motivate car manufacturers to develop more secure and robust defenses against such attacks.

Halperin et al. [11] and Gollakota et al. [10] demonstrated how an implantable cardiac defibrillator could be remotely compromised using software radios; Denning and others in [4] demonstrated some cyber-physical exploits in robots used in homes.

Printers were also shown to be vulnerable by Cui et al. in [3] through the remote modification of their firmware and resulting compromise to cause possible fire outbreaks, in addition to providing exfiltration capabilities of privately printed documents.

In the context of the home, and the ever increasing array of connected appliances with sentimental value, Denning et al. [5] investigated and highlighted various entry points for the tech savvy criminal to infiltrate the home. Fouladi et al. [8] also demonstrated exploits taking advantage of vulnerabilities in the Z-Wave protocol stack. Similar to these, [12 and 21] also identified some flaws and mounted some attacks in both Z-Wave and ZigBee implementations respectively,

All of these findings stress the need for more emphasis to be placed on the security and privacy of cyber-physical systems. This is due to the fact that these systems, unlike most traditional computing systems have the capability to effect changes in the physical world. We argue that home automation systems are as critical as the aforementioned cyber-physical systems, as these classes of systems are in direct and prolonged contact with humans in the comfort of their homes.

4. REMOTE COMPROMISE AND CAPABILITIES

We now describe several successful attempts at remotely compromising both of the home automation systems that we purchased. The vulnerabilities we uncover are a result of not following standard security best practices, so the vulnerabilities themselves are not novel contributions. However, we include these vulnerabilities because they underscore an important point: that future home automation systems may be vulnerable to compromise, and that it is important to follow-through with understanding the implications of those compromises and explore opportunities for defense-in-depth so that, if compromised, the damages can be mitigated. As noted, others have also evaluated the security of home automation systems, e.g., [5, 8, 12 and 21].

We have notified the relevant manufacturers about the vulnerabilities so that the vulnerabilities can be patched. Since the vulnerabilities are not novel, and since we have no reason to believe that other home automation systems are more secure, we have chosen not to mention product makes and models in this paper.

4.1 Experimental Setup

We chose two brands of Z-Wave enabled home automation systems for consistency. We will refer to the first of the two products as product A while the second will be referred to as product B. Product A requires an external Z-Wave module to be connected to it and exposes a web interface which allows the homeowner to connect to the system over the Internet. Once connected, the homeowner can control lights, door locks, thermostats, and other connected devices.

Product B on the other hand, used a built-in Z-Wave module. Product B also exposes a web interface for monitoring web cam feeds and the alarm system. Remote control of Z-Wave enabled appliances is also possible through the provided web interface.

Remote connectivity to these systems is achieved through port forwarding on the homeowner's router. Both systems have premium services to provide this feature.

In addition to these controllers, we had Z-Wave enabled door locks, thermostats, light dimmers and binary switches all connected to these controllers to closely simulate the use of these appliances in the home. Moreover, it was important to analyze how these nodes are affected in the event of a security breach.

4.2 XSS Vulnerability

Through extensive investigations, we found that we were able to embed persistent JavaScript tags in the logs page of product A. This was possible because product A kept a log of all login attempts, including the username, without properly parsing and sanitizing the username input. Hence, in place of a valid username, an attacker can enter JavaScript code that will be included in the logs of the system. The consequence of this is, whenever the homeowner views the log page, persistent JavaScript code executes and the attacker can do whatever he or she wishes. Moreover, the attacker can mount a covert attack by erasing the logs afterwards.

We wrote some JavaScript code to exploit this vulnerability. The embedded JavaScript code, when executed, will create a new user with arbitrary credentials and escalated privileges. We ensured the covertness of our attack by embedding the core functionality of our exploit in an *iframe* not visible to the homeowner. We also cleared the logs to erase any trace of a newly added user. For security reasons, we chose not to publish this exploit and informed the manufacturers of product A.

Extensive work has been done to exploit XSS vulnerabilities as illustrated in [13 and 14].

4.3 Insecure HTTP

Using plain HTTP on all pages was also a prevalent problem we noticed in product A. Every communication that we observed with the unit is sent in the clear whether the homeowner accesses the controller on his or her home network or over the Internet. An attacker can eavesdrop on

credentials including usernames, passwords and other valuable information. Because it is easy for an attacker to intercept wireless communications, if a user logs into product A over the Internet such as via the wireless Internet at a coffee shop, then an attacker at the same location may be able to intercept those wireless communications, learn the user's credentials, and then use those credentials for him or herself in the future.

4.4 Miscellaneous Attack Vectors

Product A also had a VNC server enabled by default with a password of "admin" running on a fixed high-numbered port. This service was however only enabled for LAN access and could not be remotely accessed. Nevertheless, a local attacker could gain access to this service.

Furthermore, product A allows developers to design various plugins, scripts or applications to enhance functionality of their units. While this provides room for innovation on many fronts, there are apparent security and privacy risks associated with this model. The question of how well vetted these scripts, plugins or apps are before being distributed to their respective application stores, begs to be asked. Can a developer with malicious intents distribute packages on a large scale through app stores? With product A, we suspect this to be possible since there does not appear to be a vetting process to ensure that apps do not infringe on the security (digital and physical) and privacy of homes and individuals using their product; we did not, however, experimentally attempt to distribute a malicious app.

As for product B, it stored a very simple and predictable authentication cookie on the user's browser which was not associated with any session id or expiration time frame. As a result, by adding this cookie to the browser, we were able to by-pass the authentication page and had direct access to the control panel. Hence, the only hurdle left for an attacker to gain access to the control panel of the system is obtaining the IP address of product B or the IP address of the homeowner's router and the specific port that product B is bound to. The latter option depends on port forwarding being enabled on the router of the homeowner.

4.5 Implications of Vulnerabilities

We experimentally verified that—after compromising products A and B—an adversary would be able to control other Z-Wave connected devices in the home. For example, we experimentally verified that an attacker could lock and unlock a Z-Wave door lock that we purchased. We also found that an attacker could turn on and off power-hungry devices, such as HVAC systems and home appliances, if connected to a Z-Wave switch. There are clear negative consequences to such capabilities, ranging from allowing an attacker easy access to a house (a broken door or scratches inside a door lock would provide evidence of forced entry, whereas a door unlocked via a remote exploit may not provide such evidence) to allowing an attacker to control power-hungry devices in the home (and potentially impacting the homeowner financially).

We do not consider the above capabilities any more. For the rest of this paper, we focus on what an attacker might be able to do to a perhaps surprising target—a non-dimmable CFL. The CFL has no network connectivity itself. For this study, we assume that the homeowner has physically plugged a standard CFL into a Z-Wave-connected dimmer. We note that such CFLs should not be plugged into dimmers and hence our analyses explicitly take the bulbs outside their intended operating environment. We do not know how many homeowners will plug CFLs into dimmers, though there have been instances of such incorrect usage as evidenced by [20]. A key question that we ask ourselves is whether it would be possible to use vulnerabilities in a home automation system to attack a device that, by itself, does *not* have any network connectivity—the light bulb.

5. APPROACH AND METHOD

We now describe our approach to experimentally analyze the range of attacks that homeowners may be exposed to via CFLs and compromised home automation systems.

5.1 Safety Measures

We needed to make sure we were working in a safe environment since we anticipated a chance of glass shattering or a more severe scenario in which we would have been exposed to the mercury content of CFLs. We initially decided to design an enclosure from Plexiglas as shown in Figure 2. Our initial assessment was that this would be effective in protecting against shattered CFLs, but that it would not adequately protect against mercury vapor; hence, we did not use this enclosure for our experiments. We also considered using gas masks to ensure our safety, but cleanup of residual mercury vapor is a non-trivial task since the vapor could be persistent.

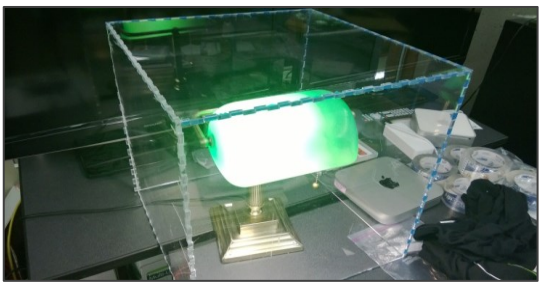

Figure 2: Plexiglas structure.

Being computer scientists and electrical engineers, we did not immediately know how to proceed and contemplated not being able to conduct our experiments. However, upon further research, and following EPA's recommended guidelines for cleaning up a broken CFL [2], we settled on using a properly ventilated glove box, shown in Figure 3. The glove box ensured that mercury vapor and shattered glass, if any, would be well contained and properly cleaned up. Another challenge was to figure out how to supply the CFL with electricity in the glove box while ensuring that it remained airtight. We improvised by drilling conduits within rubber stoppers shown in Figure 4 and carefully

sealed it up with adhesives to ensure that there would be no vapor leakage.

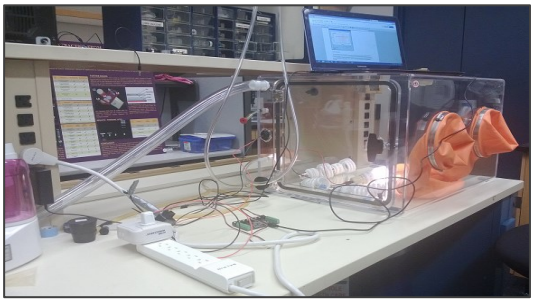

Figure 3: Glove box used to contain shattered glass and mercury.

Due to fire safety concerns, we had to be physically present when conducting our experiments. We did not have the luxury of many computer science experiments where tasks could be left to run with results viewed at a convenient time.

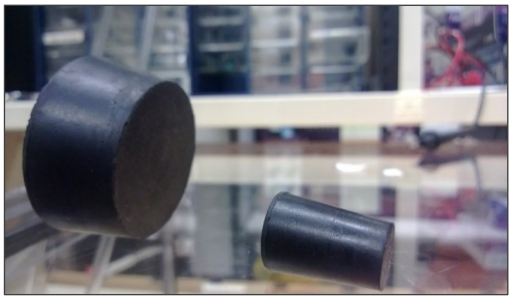

Figure 4: Rubber Stoppers.

5.2 CFL Current Monitor

In conducting our experiments, it was necessary to know how much current (RMS) was flowing through the CFL because this information would help us keep track of operation anomalies and help us recognize patterns of failure in the CFL. We acquired a Phidgets current sensor and an interface kit shown in Figure 5, with the capability of providing 125 samples per second. We also designed a graphical user interface shown in Figure 6 to aid visualization.

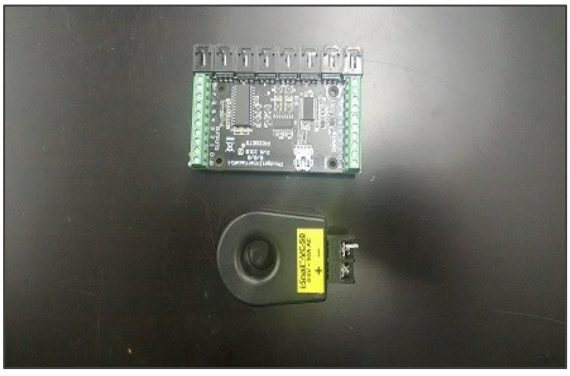

Figure 5: Phidgets Interface Kit and Current Detector.

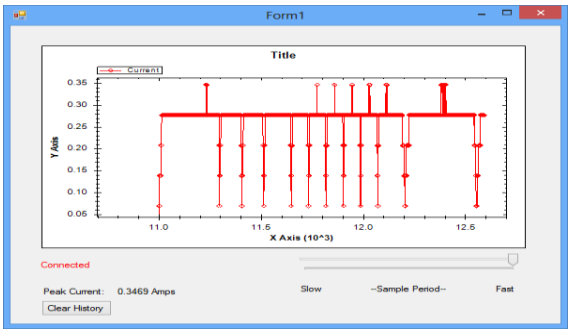

Figure 6: Real-time Plotting Utility Reporting Current Consumption.

5.3 AC Box

Similar to the need to measure current flowing through the CFL, we also logged the voltage waveform driving the CFL. Unlike current, where we log the RMS, the entire voltage waveform was recorded since shape of the waveform can change drastically depending on the load and dim level rather than the amplitude.

To safely measure the voltage, we galvanically isolated the measurement equipment from the AC-line using a step-down transformer (Triad Magnetics part F12-090-C2-B) with an approximate coil ratio of 1:16 under full load (a 75 ohm resistor was placed across the secondary terminals to load the output). This allowed us to safely connect a bench oscilloscope to the secondary of the AC transformer. Figure 11 shows one instance of the recorded waveform using this approach.

5.4 Signal Generation

For the purpose of our experiments, we assume a naive homeowner has upgraded his home with a home automation system and has connected a CFL to a dimmer with remote control capabilities. We also assume the attacker has compromised the system through one of the aforementioned vulnerabilities and is intent on physically harming the occupants of the home by causing CFLs connected to dimmers to explode. Our experiments are designed to gauge what, if anything, an attacker might be able to accomplish. As additional background knowledge: lights fluctuating at certain frequencies can be dangerous to people with photosensitive epilepsy; CFLs contain mercury; and an exploded light bulb could result in shattered glass or possibly a fire outbreak [18 and 20].

To study this threat experimentally, we utilized an Aeon Z-stick® static update controller which uses the Z-Wave protocol for low data rate communications as shown in Figure 7. Additionally, we utilized open-zwave libraries to remotely control Z-Wave-enabled light dimmers, and connected to the dimmers were two different groups of CFL brands. We then generated four distinct signals and extensively tested them out on the CFLs until they either gave way or produced an anticipated result like a dramatic pop. Since the only parameter we could alter from a remote perspective was the Z-Wave dimmer level and considering

how time-intensive the experiments were, we were unable to experiment with a large number of signal types. We therefore chose the four signal types that we thought would cause the CFL to operate outside normal operating conditions.

Figure 7: Aeon Z-stick®.

Figure 8 shows periodic triangular pulses that were applied to the Z-Wave dimmer. With this mode of operation, a peak voltage level was chosen (as described below) and the voltage applied to the CFL was varied from zero to the chosen peak level and back to zero at a refresh rate from at least every second to about every 60 milliseconds. While the timing in addition to the signal, were chosen to closely simulate an individual physically varying the brightness of the CFL, we had upper bounds on the refresh rate due to the low data rate constraint of the Z-Wave protocol. The peak dimmer level shown in Figure 8 is arbitrary and can be set to any value between 0 and 100 (the range 0 and 100 correspond to the levels allowed by the dimmer). The peak voltage was chosen by observing the voltage level at which the CFL became unstable, i.e., at the onset of visual fluctuations. The level at which the CFL became unstable was largely affected by process and design variations.

From our experience, instability usually kicked in when the dimmer level was set to about 20% of the maximum brightness of the lamp. The reason why the peak voltage selection was important is that we observed through repeated experimentation that the CFL was more likely to fail at a faster, however inconsistent rate, when the selected voltage level induced visual fluctuations in the lamp. We again stress, however, the limited sample size of our experiments.

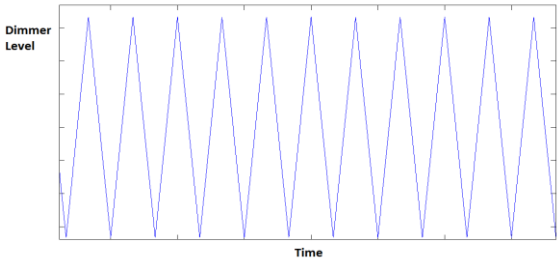

Figure 8: Periodic triangular pulses applied to Z-Wave Dimmer. Plot of Z-Wave dimmer level versus time.

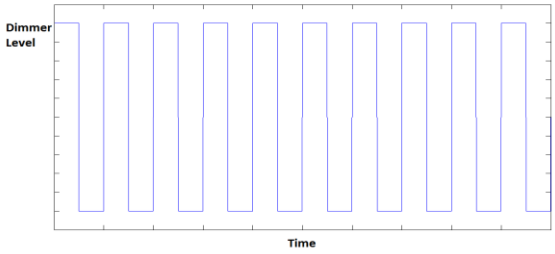

Figure 9: Periodic rectangular pulses applied to CFL. Plot of Z-Wave dimmer level versus time.

For the second signal, we toggled the applied voltage level between a peak voltage of our choice and zero at a refresh rate from at least every second to about every 60 milliseconds as shown in Figure 9. Again, the peak dimmer level shown in Figure 9 is arbitrary and can be set at any level between 0 and 100. In this case, we selected the peak voltage to be maximum. We selected this waveform as a simple variant of the periodic triangular pulses, though we acknowledge that other wave forms are possible too. Our original intentions with this signal was to cause the CFL to pop, but we soon realized that this signal might cause the light to flash at a seizure-inducing frequency (see results section).

For the third signal, we wanted to gauge whether a random signal might be effective at damaging the bulbs. Hence, we decided to add randomness by generating Gaussian distributed random numbers and decided against a certain threshold to either increase or decrease the applied voltage. An example plot of the applied random signal is shown below in Figure 10.

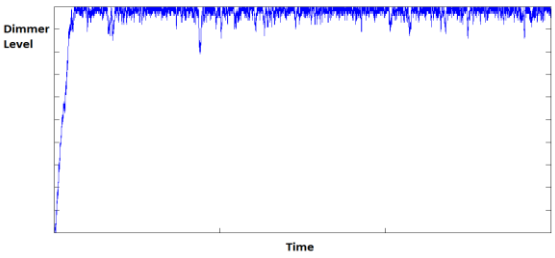

Figure 10: Random Gaussian distributed signal applied to CFL. Plot of Z-Wave dimmer level versus time.

Finally, we combined the triangular pulses shown in Figure 8 with some randomness from a Gaussian distributed random number generator similar to the signal shown in Figure 10. We also set the peak voltage as defined for the triangular pulses described earlier.

Table 1 has labels “Signal A”, “Signal B”, “Signal C” and “Signal D” attached respectively to the applied periodic triangular and rectangular pulses, the random Gaussian distributed signal and the periodic triangular-random Gaussian signal combo.

Table 1: Summary of Applied Signals.

Label	Characteristics
Signal A	Periodic Triangular Pulses
Signal B	Periodic Rectangular Pulses
Signal C	Random Gaussian Distributed Signal
Signal D	Periodic Triangular Pulses + Random Gaussian Distributed Signal

The effect that these applied signals have on the CFL are shown in Figures 11-13. The dimmer generates a pulse width modulated signal whose width is controlled by the applied dimmer level. Figure 11 shows the voltage plot across the terminals of the CFL when the dimmer level is set to 8, while Figures 12 and 13 show voltage plots across the terminals of the CFL with the dimmer set to levels 50 and 100 respectively.

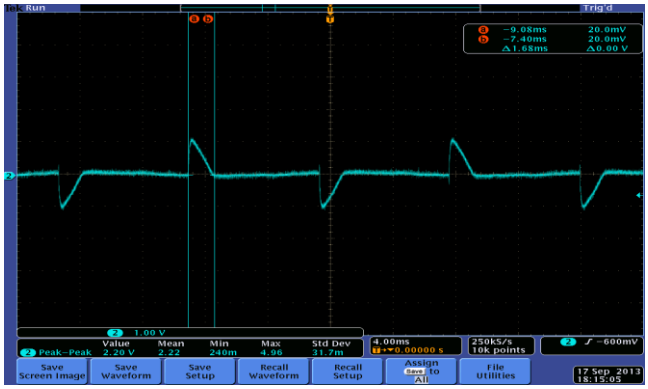

Figure 11: Plot of voltage across the terminals of the CFL with dimmer level set to 8.

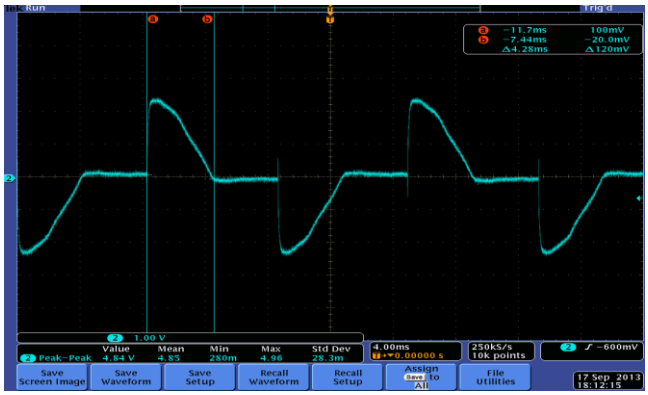

Figure 12: Plot of voltage across the terminals of the CFL with dimmer level set to 50.

Figure 13: Plot of voltage across the terminals of the CFL with dimmer level set to 100.

In real time, there is a progressive increase of the pulse width from the minimum to the maximum (Figure 13) when Signal A is applied and the peak dimmer level is set to 100. Once the voltage across the CFL reaches its maximum width, it shrinks and flattens out to zero. This is repeated until the CFL pops or gives way.

Similarly, the pulse width of the voltage plot across the terminals of the CFL changes between two values—maximum and minimum pulse widths—when Signal B is applied and the peak dimmer level is set to 100. For Signal C, the pulse width randomly increases or decreases depending on the set threshold, dimmer level and previous CFL voltage. Finally, Signal D is simply a combination of the effects Signals A and C have on the CFL.

6. RESULTS

6.1 Data and Analysis

Our experiments yielded a wide variety of results, including inconsistent times to popping the CFLs. We conducted several preliminary experiments to determine the most effective and safest way (from our perspective as researchers) to get the CFLs to fail. Table 2 shows some of the results we obtained through the application of the signals defined in Table 1. Through repeated experimentation, we found out that Signal A was the most effective in causing CFLs to fail, Signal B had a side effect of possibly triggering seizures, and Signal C had to be combined with Signal A for it to be as effective. We acknowledge that our sample sizes are small, however—an artifact of the resource intensiveness of conducting experiments with this class of cyber-physical systems.

We conjecture that the inconsistent times to failure is largely attributed to process and design variations among similar and different CFL brands. Even though the lifespan of the devices were ultimately reduced, the time to failure varied to a large extent. It is also important to note that we did not conduct this particular set of experiments over the Internet, but limited the scope to a local control of the Z-Wave controller using open-zwave libraries. We hope to experimentally evaluate an end-to-end attack as an

extension to the work in the future. We got some results that we believe the community would be interested in, as a component within the electronic ballasts of some of our test CFLs, specifically a bipolar junction transistor (BJT) dramatically burnt out with a "pop". This left some charring on the device as shown in the Figure 14 identified by the circular ring.

Table 2: Time to failure for CFLs. *Popped after direct connection to electricity without the dimmer.

Lamp Tag	**Brand**	**Time (Hours)**	**Signal Type(s) Applied**	**CFL Status**
#1	Walmart Great Value	0.3	Signal A	Gave way
#2	Walmart Great Value	0.8	Signal A	Gave way
#3	Walmart Great Value	7	Signal A	Gave way
#4	**GE**	**7**	**Signals A, B and C applied in no particular order**	**Popped**
#5	Walmart Great Value	3	Signal A	Gave way
#6	GE	0.6	Signal A	Gave way
#7	Walmart Great Value	4	Signal A	Gave way
#8	**GE**	**0.7**	**Signal A**	***Popped**
#9	Walmart Great Value	Over 6	Signal C	Settled in a state consisting of visual fluctuations.
#10	**GE**	**1.5**	**Signal D**	**Popped**

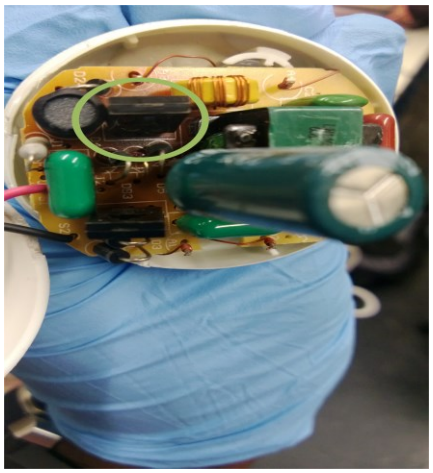

Figure 14: Charred CFL.

The recorded results shown in the Table 2 do not include several preliminary test runs we had, to determine the feasibility of inducing failures in CFLs.

For the recorded set of experiments, we initially started out by applying Signal A to the Walmart Great Value brand, which only resulted in the CFLs giving way at, however, inconsistent times. We also experimented with signals A, B and C by randomly applying them to the same CFL (lamp 4) in no particular order. This resulted in the first pop we observed after seven hours of experimentation.

After applying Signal A to lamp 8 (highlighted in orange in Table 2) for about 42 minutes, we noticed it was beginning to fail. To confirm its failure, we connected lamp 8 directly to the power source without the dimmer and heard a pop, indicating that a component (BJT) had given way in its ballast. The current spike that resulted from connecting the CFL directly to the power source is shown below in Figure 15. Depending on the kind of lighting fixture or shade around the light bulb, the heat generated from the failing bulb may pose a fire hazard. CFLs failing in this manner have been reported to cause major fire damage based on past recall reports [16 and 17]. No fires were ignited in our experiments, however.

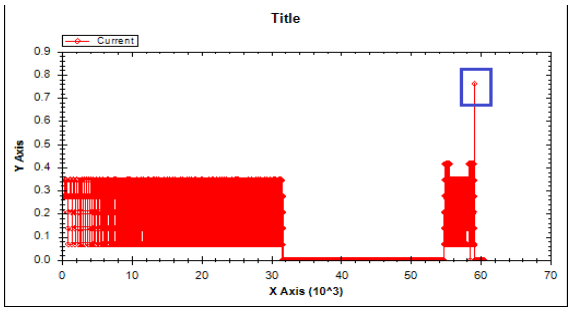

Figure 15: Resultant spike from current surge in CFL.

Taking this result into consideration, we tweaked Signal C by combining its mode of operation with Signal A to yield Signal D, The purpose of Signal D was to randomly cause a spike in the current flowing through the CFL at various

points during the experiment. This was necessary, as applying Signal C solely to the CFL was not yielding the desired result of either popping or giving way. As indicated in the previous section, we purposely set the peak voltage for Signals A and C to be reasonably low, to ensure that the CFL was in an unstable, flickering state. The voltage level that the CFL was set to varied from one lamp to the other and was due to process and design variations. The result of the Signal D was intermittent spikes in current from sporadically setting the dimmer to its maximum at times determined by the result of a Gaussian distributed random variable. We achieved the same result we got applying Signal A to the CFL with Signal D as evidenced with lamp 10.

Even though Signal B was intended to cause the CFL to pop, it may have a side effect of possibly causing a seizure; we did not experiment with this extensively after learning that the bulb was oscillating at a dangerous frequency [7]. Moreover, we did not initially anticipate that Signal B might be at a seizure-inducing frequency, but began to investigate that frequency after experiencing some discomfort from applying this signal to the CFL. For safety reasons, we did not run this experiment extensively, and when we did we took safety precautions (see Section 7 for details).

In summarizing our results, we set out to experimentally cause two different brands of CFLs to pop remotely by applying the signals shown in Figures 8-10 through a Z-Wave enabled light dimmer. Our results indicate that we were able to cause a reduced life-span, though inconsistent failure times, in the CFLs. More interestingly, we were able to cause some CFLs of the GE brand to pop with the BJT burning out. In our limited experiments, none of the pops caused serious damage to the external environment. Lastly, although we set out to pop CFLs using Signal B shown in Figure 9, we noticed a side-effect of possibly triggering seizures at the operated frequency of oscillation.

7. DISCUSSION

We stress that our demonstrated CFL attacks are not end-to-end. We demonstrated the ability for an attacker to remotely compromise and control two home automation system controllers, and from there we did confirm the ability of an attacker to do simple device manipulations, like unlock doors and turn on or off appliances. And we explored the feasibility of an attacker, connected to a wireless home automation network, to control a network-connected dimmer and thereby affect the CFLs plugged into the dimmer. However, we did not mount our attacks against the CFLs over the Internet to an uninstrumented home-automation ecosystem. A fundamental limitation was timing—using our current compromises to the home automation controllers, we were unable to send packets to the dimmer fast enough. Nevertheless, we argue that our current results are important because there are ways in which an adversary might be able to obtain internal access to a home automation system's internal wireless network. For example, more sophisticated code-injection attacks could be found against home automation controllers (e.g., full code injection rather than JavaScript injection). A nearby attacker might also attempt to attack the home automation system's wireless protocols directly, and thereby gain direct wireless access to the dimmers. An attacker might also produce Trojan home automation hardware, and unsuspecting users may connect that Trojan hardware to their home automation systems' wireless networks. The fundamental conclusion, therefore—that a network-based attacker might be able to affect a device that, by itself, is not designed to be networked (the CFLs)—remains true.

During the course of our experiments, we found out that there was no convenient and cost-effective way to detect mercury spillage. As a result, we are yet to experimentally verify the amount of mercury vapor, if any, leaked as a result of our experiments. Our glove box and ventilation system was, however, borrowed from a wet lab and was designed to deal with such vapors, whether detectable or not. Additionally, due to process and design variations, the failure times for the CFLs were very inconsistent. In certain cases, we were able to either get the CFL to fail with or without a pop in as little as eighteen minutes or as long as over seven hours. This is reflected in Table 2. We did not experiment with placing the CFLs next to lamp stands or accessories.

As mentioned, due to fire safety concerns, we had to be physically present when conducting our experiments. We did not have the luxury of most computer science experiments where tests could be left to run with results viewed at a convenient time. Also for safety, we needed to shield ourselves from staring directly at the fast switching Signal B shown in Figure 9, as it is in a frequency range that may induce a seizure in an observer [7]. While Signal B was not as effective as Signals A, C and D in terms of causing CFLs to pop, it caused discomfort to the observer; we implemented the safety precautions after experiencing this discomfort and realizing that the light was pulsing at a potential seizure-inducing frequency. Specifically, we covered the glove box with opaque black plastic bags to shield us from staring directly into the lamp. Future security research on cyber-physical systems must identify potential safety risks proactively, rather than reactively; proactive identification in all cases, however, may be fundamentally challenging if not impossible.

Sample size for cyber-physical systems research is another issue that the research community must address in the future. Some studies—such as past work on automobiles [1]—experimented with only two artifacts. We experimented with more light bulbs, but—given our limited resources—not nearly as many as we would have liked. For safety, our experiments required manual supervision, as noted above. This need for manual supervision is comparatively rare in computer science, and differentiates cyber-physical systems research from some other classes of

computer security research. Each experiment took significant time, further contributing to the small sample size. However, we acknowledge that our sample size is small and encourage future follow-on work to repeat our experiments with larger sample sizes, more signal variety, and more bulb types.

With all of these findings, it is necessary to take a step back to examine the consequences from the perspective of industry stakeholders, homeowners and also researchers.

From the perspective of industry stakeholders, it is important to stress the need for the design of more robust and secure home automation systems. This should encompass every party in the ecosystem ranging from those involved in the physical layer design to application developers who may unintentionally introduce vulnerabilities into the system. For instance, product A created a scenario like this, as the web interface was prone to XSS attacks as discussed.

For homeowners, there is an apparent trade-off between the convenience factor that home automation systems provide and security and privacy of the home. To what extent are homeowners willing to compromise security and privacy of the home for the ability to remotely control physical actuators around the home? Should homeowners be worried about inherent security flaws in the design of home automation systems and as such give the industry some time to mature and overcome these issues?

For researchers, a lot more needs to be done in this field to ensure that industry partners develop robust and secure home automation systems. Furthermore, with more heavy-duty home appliances increasingly connected to the Internet, detailed analysis of added connectivity benefits and resulting costs to security and privacy have to be carried out.

8. CONCLUSION

While home automation systems undoubtedly provide immense benefits in terms of convenience, more work needs to be done to ensure robust and secure designs of these systems. Furthermore, there is a need for all stakeholders involved—ranging from industry and research partners to homeowners—to fine-tune our understanding of whatever flaws these systems possess. We hope our work will further catalyze interest in discovering and fixing vulnerabilities in home automation systems, and their surrounding ecosystems, and enlighten end users to be cautious with their adoption and mode of use.

Of particular interest, we believe, is the fact that devices *not* designed for network connectivity (e.g., light bulbs) may be connected to other devices that *do* have network connectivity. Such connections may expose the former devices to risks that the designers of those devices never anticipated. The designers of the latter (networked) devices (like dimmers or entire home automation systems) may not know which other devices will connect to them in a home deployment, and hence providing sufficient protection mechanisms on the latter devices may be challenging. We encourage further research and design on secure home automation systems.

9. ACKNOWLEDGMENTS

We thank Professor Karl F. Böhringer for providing the lab space and necessary apparatus to conduct our experiments. We thank Dr. Carrie Gates for shepherding this paper and Karl Koscher for his help in conducting the experiments. We thank the numerous anonymous reviewers for their valuable feedback and recommendations. This work was supported by the Intel Science and Technology Center for Pervasive Computing.

10. REFERENCES

[1] S. Checkoway, D. McCoy, D. Anderson, B. Kantor, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno, “Comprehensive Experimental Analysis of Automotive Attack Surfaces,” in Proceedings of the USENIX Security Symposium, San Francisco, CA, August 2011.

[2] Cleaning Up a Broken CFL, (N.D.), from U.S. Environmental Protection Agency. Retrieved June 26, 2013 from the U.S. Environmental Protection Agency: http://www2.epa.gov/cfl/cleaning-broken-cfl#instructions

[3] Cui, A., Stolfo, S. “Print Me If You Dare: Firmware Modification Attacks and the Rise of Printer Malware,” in The 28th Chaos Communication Congress, December 27, 2011.

[4] Denning, T., Matuszek, C., Koscher, K., Smith, J. R., and Kohno, T. A spotlight on security and privacy risks with future household robots: attacks and lessons. In Ubicomp '09: Proceedings of the 11th international conference on Ubiquitous computing (2009), pp. 105-114

[5] Denning, T., Kohno, T., and Levy, H. M. Computer security and the modern home. Commun. ACM, 56(1):94–103, Jan. 2013.

[6] Elliot, R. Should There be a Ban on Incandescent Lamps?, February 22, 2007 from Elliott Sound Products. Retrieved June 26, 2013 from Elliott Sound Products: http://sound.westhost.com/articles/incandescent.htm#dim

[7] Photosensitivity and Seizures, (N.D.), from the Epilepsy Foundation. Retrieved June 26, 2013 from the Epilepsy Foundation: http://www.epilepsyfoundation.org/aboutepilepsy/seizures/photosensitivity/

[8] Fouladi, B., Ghanoun, S. Security Evaluation of the Z-Wave Wireless Protocol. In Black hat USA (2013).

[9] OpenZwave. (n.d.). Retrieved June 2013, from openzwave Google code site: https://code.google.com/p/open-zwave/

[10] Gollakota, S., Hassaneih, H., Ransford, B., Katabi, D., Fu, K. They can hear your heartbeats: non-invasive security for implantable medical devices: Proceedings of the ACM SIGCOMM 2011 conference, August 15-19, 3011, Toronto, Ontario, Canada.

[11] Halperin, D., Heydt-Benjamin, T.S., Ransford,B., Clark, S. S., Defend, B., Morgan, W., FU, K., Kohno, T., and Maisel, W. H. Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses. In *IEEE Symposium on Security and Privacy* (2008).

[12] Kennedy, D., & Simon, R. (2011). Pentesting over Power lines. Defcon 2011

[13] A. Kiėzun, P. J. Guo, K. Jayaraman, and M. D. Ernst, Automatic creation of SQL injection and cross-site scripting attacks, in ICSE'09, Proceedings of the 30th International Conference on Software Engineering, Vancouver, BC, Canada, May 20-22, 2009.

[14] Klein, A. Cross Site Scripting Explained, June 2002 from Sanctum Security Group. Retrieved June 29, 2013 from Stanford University: http://crypto.stanford.edu/cs155/papers/CSS.pdf

[15] Learn About CFLs, (N.D.), from Energy Star. Retrieved June 26, 2013 from the Energy Star: http://www.energystar.gov/index.cfm?c=cfls.pr_cfls_about

[16] OFPC Safety Alert, Compact Fluorescent Light Bulbs (CFL's) February 9, 2011 from the New York State Division of Homeland Security and Emergency Services. Retrieved July 6, 2013 from New York State New York State Division of Homeland Security and Emergency Services: http://www.dhses.ny.gov/ofpc/news/press/documents/2011_safety_alert_cfl_actual.pdf

[17] Recalls, October 5, 2010, from the United States Consumer Product Safety Commission. Retrieved July 6, 2013 from the United States Consumer Product Safety Commission: http://www.cpsc.gov/en/Recalls/2011/Trisonic-Compact-Fluorescent-Light-Bulbs-Recalled-Due-To-Fire-Hazard/

[18] Rocznik, K., Consumerwatch: CFL Bulb Safety, March 25, 2013, from CTV News. Retrieved March 25, 2013 from CTV News: http://winnipeg.ctvnews.ca/consumerwatch-cfl-bulb-safety-1.1210152

[19] Smart Homes and Home Automation, July 2011, from Berg Insight. Retrieved July 1, 2012 from Berg Insight: http://www.berginsight.com/ReportPDF/ProductSheet/bi-sh1-ps.pdf

[20] Spradlin K., Blaze Underscores need for CFL bulb Education, April 30, 2008, from Cumberland Times-News. Retrieved June 26, 2013 from Cumberland Times-News: http://times-news.com/archive/x1540421978

[21] Wright, J. (2011). Practical ZigBee Exploitation Framework. Toorcon 2011.