This paper appeared in *Proceedings of the 8th USENIX Workshop on Offensive Technologies* (WOOT '14), August 2014.

Green Lights Forever: Analyzing the Security of Traffic Infrastructure

Branden Ghena, William Beyer, Allen Hillaker, Jonathan Pevarnek, and J. Alex Halderman
Electrical Engineering and Computer Science Department
University of Michigan
{brghena, wbeyer, hillaker, jpevarne, jhalderm}@umich.edu

Abstract

The safety critical nature of traffic infrastructure requires that it be secure against computer-based attacks, but this is not always the case. We investigate a networked traffic signal system currently deployed in the United States and discover a number of security flaws that exist due to systemic failures by the designers. We leverage these flaws to create attacks which gain control of the system, and we successfully demonstrate them on the deployment in coordination with authorities. Our attacks show that an adversary can control traffic infrastructure to cause disruption, degrade safety, or gain an unfair advantage. We make recommendations on how to improve existing systems and discuss the lessons learned for embedded systems security in general.

1 Introduction

Traffic signals were originally designed as standalone hardware, each running on fixed timing schedules, but have evolved into more complex, networked systems. Traffic controllers now store multiple timing plans, integrate varied sensor data, and even communicate with other intersections in order to better coordinate traffic.

Studies have shown the benefits of a well coordinated traffic signal system in terms of wasted time, environmental impact, and public safety [2], but coordination has been difficult to achieve due to the geographic distribution of roadways and the cost of physical connections between intersections. Wireless networking has helped to mitigate these costs, and many areas now use intelligent wireless traffic management systems [10, 32, 33]. This allows for new capabilities including real-time monitoring and coordination between adjacent intersections. However, these improvements have come with an unintended side effect. Hardware systems that had previously been only physically accessible are now remotely accessible and software controlled, opening a new door for attackers.

To test the feasibility of remote attacks against these systems, we perform a security evaluation of a wireless traffic signal system deployed in the United States. We discover several vulnerabilities in both the wireless network and the traffic light controller. With coordination from the road agency, we successfully demonstrate several attacks against the deployment and are able to change the state of traffic lights on command.

The vulnerabilities we discover in the infrastructure are not a fault of any one device or design choice, but rather show a systemic lack of security consciousness. We use the lessons learned from this system to provide recommendations for both transportation departments and designers of future embedded systems.

2 Anatomy of a Traffic Intersection

The modern traffic intersection is an amalgamation of various sensors, controllers, and networking devices. Figure 1 shows some common devices found at intersections.

2.1 Sensors

Sensors are used to detect cars and inspect infrastructure. Induction loops (also known as in-ground loops) are frequently used to detect vehicles. These devices are buried in the roadway and detect cars by measuring a change in inductance due to the metal body of the vehicle. Video detection is also frequently used to sense vehicles at intersections. In the United States, 79% of all vehicle detection systems use video detection or induction loops [18]. Microwave, radar, and ultrasonic sensors are less common, but also used [17]. Video cameras are also commonly installed to allow remote inspection of the intersection.

2.2 Controllers

Traffic controllers read sensor inputs and control light states. The controller is typically placed in a metal cabinet by the roadside along with relays to activate the traffic lights. Sensors are typically directly connected to the controller, allowing it to combine vehicle detection information with pre-programmed timing controls in order to determine the current state of the traffic lights.

Intersections can be configured to operate in several different modes. In the simplest case, pre-timed mode, lights are controlled solely on preset timings [8]. More complicated controllers function in a semi-actuated mode where the side street is activated based on sensors and the main street otherwise runs continuously. In fully-actuated mode, both streets are serviced based on sensor input [36].

Controllers can function as isolated nodes or as part of an interconnected system. Isolated intersections maintain

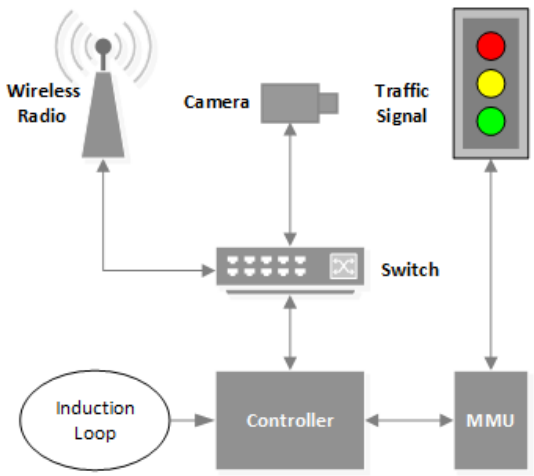

Figure 1: A typical traffic intersection. The radio connects to the switch and transmits controller diagnostics, live video feed, and other information back to the road agency. The malfunction management unit sits between the controller and lights and ensures that the lights are not put in an unsafe configuration. The controller adjusts light timings based on data from the induction loop.

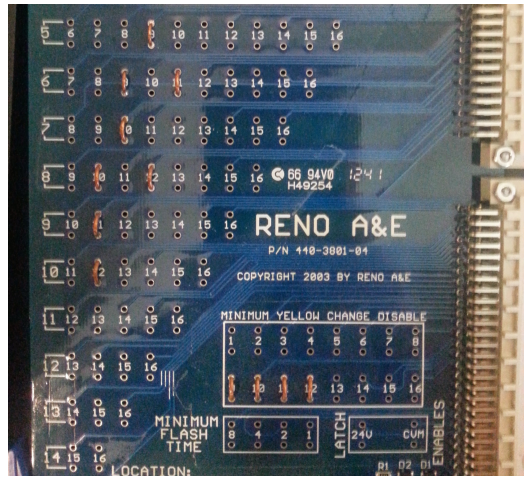

Figure 2: MMU configuration board. If two phases which are not physically connected by a jumper attempt to activate simultaneously, the output from the controller will be ignored and the intersection will be put into a known-safe configuration.

state independently from nearby intersections. Interconnected systems share timing information and may also react to sensor input from nearby intersections [8]. According to a 2003 survey by the Institute of Transportation Engineers, 62% of traffic intersections operate in an interconnected fashion [18].

2.3 Communications

Controllers may communicate with both each other and a central server in order to share information on current traffic conditions. Hard-wired communication through optical or electrical means is common in dense urban areas but becomes difficult to install and maintain when intersections are geographically distant. In this scenario, radios are frequently used in a point-to-point or point-to-multipoint configuration in order to provide connectivity [35]. Radios commonly operate in the ISM band at 900 MHz or 5.8 GHz, or in the 4.9 GHz band allocated for public safety [35].

2.4 Malfunction Management Unit

Malfunction management units (MMUs), also referred to as conflict management units, are hardware-level safety mechanisms. They act as a whitelist of light states and monitor the outputs of traffic controllers to ensure that no invalid states occur. Valid configurations are stored on a circuit board rather than in software, with safe configurations literally wired together as shown in Figure 2. If an unsafe configuration (e.g. conflicting green lights) is detected, the MMU overrides the controller and forces the lights into a known-safe configuration (e.g. blinking reds) [24]. The intersection enters a fault state and requires manual intervention to reset. The MMU also ensures that durations of lights are long enough. Too short of a yellow or red light duration will trigger a fault.

3 Case Study

Our study was performed with cooperation from a road agency located in Michigan. This agency operates just under a hundred traffic signals in primarily urban environments. However, their coverage area is very large, necessitating wireless communications to prevent the infrastructure from becoming prohibitively expensive.

The wireless capabilities of the intersections allow them to report current traffic conditions to a central server. This information can be used to make modifications to light timings at specific intersections in the event of congestion. Even though the intersections are all part of the same wireless network, they operate in an isolated mode and do not coordinate directly with one another.

While other deployments may use different wireless radios or even wired connections between intersections, we have no reason to believe that any fundamental differences exist between the network we studied and other traffic signal systems. We believe that many traffic infrastructure devices created by various vendors and installed by various transportation departments will have similar security properties due to a lack of security consciousness in the entire field. The vendor studied in Cerrudo's research on traffic systems [3] is not the same as the vendors of the controller or radios investigated in our study, but had similar vulnerabilities, reinforcing this belief.

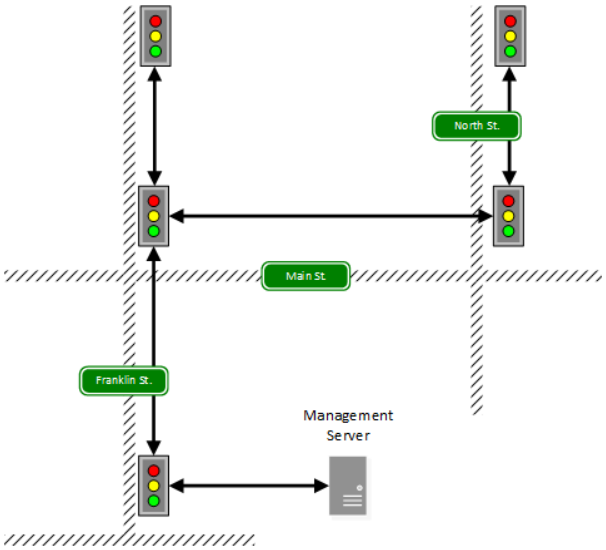

Figure 3: Example traffic signal network. Within the road agency's network, intersections are connected in a tree topology with information flowing to a root node. This root node communicates with the central management server, which can send commands to any intersection on the network.

3.1 Network

The system we investigated uses commercially available radios that operate on the ISM band at either 5.8 GHz or 900 Mhz. Figure 3 shows an example of the network topology. One intersection acts as a root node and connects back to a management server under the control of the road agency. Intersections often have two radios, one slave radio to transmit to the next intersection towards the root and one master radio to receive from one or more child nodes beyond it. All devices form a single private network and belong to the same IP subnet.

5.8 GHz radios are preferred as they provide higher data rates. They communicate using a proprietary protocol and utilize point-to-point or point-to-multipoint connections. The proprietary protocol is similar to 802.11 and broadcasts an SSID which is visible from standard laptops and smartphones but cannot be connected to. In order to properly connect, a slave radio must use the proper protocol and know the network SSID. The wireless connections are unencrypted and the radios use factory default usernames and passwords. The configuration software for these radios accepts customized credentials but assumes that the same username and password are used across all radios on the network.

900 MHz radios are used when there is not a direct line-of-sight connection to the next intersection due to obstructions such as buildings, bridges, or terrain. These radios use a proprietary protocol with frequency hopping spread-spectrum (FHSS) and point-to-point connections. In order to make a connection, both the master and slave radios must know the network name and a 16-bit slave ID value. As with the 5.8 GHz radios, the connections between 900 MHz radios are unencrypted and the radios use default usernames and passwords. The configuration software for these radios assumes the default username and password will be used. If they are modified, the software is no longer able to connect to the device.

At an intersection, the radios, controller, and video camera connect to a commercial switch via an Ethernet connection. The switch does not implement any security features and utilizes its default username and password.

3.2 Controller

A single controller at each intersection reads sensor data and controls the traffic lights and pedestrian signs. Many settings on the controller are programmable, including light timing parameters. The controller does support profile-based access control, but this is disabled by default and is not used in the setup we studied.

In addition, the controller model used at the majority of the intersections in our study contains a logic processor that permits traffic engineers to have fine-grained control over the intersection. Basic if-then-else logic is used to execute arbitrary commands, such as changing the state of a particular light or freezing lights in their current state. An example logic statement can be seen in Figure 4. Without user profiles, this programming ability is open to anyone with access to the controller.

All of the settings on the controller may be configured via the physical interface on the controller, but they may also be modified though the network. An FTP connection to the device allows access to a writable configuration database. This requires a username and password, but they are fixed to default values which are published online by the manufacturer. It is not possible for a user to modify the FTP username or password.

The controller conforms to NTCIP Protocol 1202, which specifies communication standards for traffic signal systems [27]. Via this protocol, commands can be sent to the controller which equate to button presses on the front panel, allowing the user to walk the menus, change timing plans, and request the contents of the screen.

The controller runs the VxWorks 5.5 real-time operating system. The default build settings for this version of VxWorks leave a debug port open for testing purposes. This port is meant to be disabled for build environments but is so commonly left enabled that it has been marked as a vulnerability by ICS-CERT [15]. Connecting to the port requires no password and allows for arbitrary reading and writing of device memory among other capabilities. This vulnerability can be easily fixed by modifying build arguments to close the debug port, and the vendor has addressed the issue in a software patch. However, the con-

```
IF   COORD CYCLE LENGTH      IS     0

THEN CTR SET STIME ALL           ON

ELSE CTR SET STIME ALL           ON
```

Figure 4: Example logic statement. This example prevents controlled lights from changing state. The controller is given a boolean statement and corresponding actions. Selecting the same action in both cases forces the controller to enter a chosen state regardless of the condition. In the example shown, the controller is set to "Stop Time", a state where the lights are frozen in their current state.

trollers examined in this study used an older version of the controller software and were vulnerable to this attack.

3.3 Other Devices

Our study focuses on attacking the network and controller, but other devices on the network may be targets of adversaries as well. One example is live video data from cameras placed at the intersections. These cameras record color video, and include pan, tilt, and zoom control. They also connect to the switch, and their data is sent back to the management server so that traffic engineers can remotely monitor intersections in real time. Whenever attached to the network rather than directly to the controller, sensors may be targets of attacks as well. Spoofed sensor data could be used to manipulate light states, as has been shown in prior work [3]. We did not investigate attacks against these devices due to a lack of available hardware.

3.4 Findings

To summarize, we discovered three major weaknesses in the road agency's traffic infrastructure deployment:

1. The network is accessible to attackers due to the lack of encryption.
2. Devices on the network lack secure authentication due to the use of default usernames and passwords.
3. The traffic controller is vulnerable to known exploits.

4 Attacks

This section describes attacks we performed against the traffic signal control system and the impact of those attacks. We begin by discussing how to access the traffic signal network and then go into detail on methods to access and attack the controller. With cooperation from the road agency, we successfully demonstrate several attacks at a real intersection in the deployment. An image of the intersection can be seen in Figure 5.

4.1 Threat Model

The threat model we consider is an attacker infiltrating the traffic network through its wireless infrastructure. We

Figure 5: Intersection used for field tests. The targeted infrastructure controls a T intersection. The failsafe state of the intersection is a blinking yellow light on the main road and a blinking red light on the side road. At left is the traffic controller cabinet, which contains the controller, switch, and MMU. Radios and a video camera are located at the top of the pole.

assume the attacker has sufficient resources and motivation to monitor the network for extended periods of time and the funds to purchase radio equipment.

We assume that the attacker does not have physical access to any part of the traffic infrastructure. With direct access to the traffic cabinet, an attacker would be able to remove fail-safe equipment and perform dangerous attacks (e.g. four-way green lights) in addition to the attacks described in this paper. This would come with a higher risk of detection due to the necessity of accessing the traffic cabinet in view of the road and any video cameras.

4.2 Accessing the Network

In order to attack devices in the traffic system, the adversary must first gain access to the network. The process for gaining network access varies between radio types and configurations. Once the network is accessed at a single point, the attacker can send commands to any intersection on the network. This means an adversary need only attack the weakest link in the system.

5.8 GHz Radios None of the radios used in the network we studied made any attempt to conceal or encrypt their traffic. In the case of the 5.8 GHz radios, any attacker with a wireless card capable of 5.8 GHz communication is able to identify the SSIDs of infrastructure networks.

While the proprietary protocol used by the radios in our study could potentially be reverse engineered to allow any 5.8 GHz radio to communicate with the deployed radios, we chose to circumvent this issue and use the same model radio that was deployed in the studied network for our attack. While these radios are not usually sold to the

public, previous work has shown social engineering to be effective in obtaining radio hardware [38].

Due to the lack of encryption, any radio that implements the proprietary protocol and has knowledge of the network's SSID can access the network. When testing on the deployment, we successfully accessed the network using a radio produced by the same vendor as the deployed radios. The radio we connected to was over half a mile down the road from our location and allowed us to access a networked controller located at a different intersection.

900 MHz Radios Attacking the 900 MHz network requires knowledge of both the network name and the 16-bit slave ID value. Sniffing for the network name is made more difficult by both the proprietary communication protocol used and FHSS. The frequency hopping done by the radios makes sniffing packets more difficult, but with the appropriate hardware this becomes a solvable problem. Previous work has shown that FHSS is not inherently secure [9, 14]. The proprietary protocol could be reverse engineered, but can also be overcome by purchasing compatible hardware.

With the correct 16-bit ID, the attacker's radio can imitate the existing slave radio and communicate with the master. A brute force approach to determining the ID is possible. Given the time required to make automated configuration changes to the radio (on the order of several seconds), it could theoretically take up to several days to determine the ID value. In practice, however, operators commonly select low-numbered ID values, and a successful attack may take only minutes. For the deployment we studied, all radio ID values were less than one hundred.

We did not test obtaining network names on the 900 MHz network due to a lack of proper hardware. Assuming a known network name, however, we were successful in spoofing radio ID values in a laboratory setting.

4.3 Accessing the Controller

Once on the network, there are two primary methods of accessing the controller: the operating system's debug port, or the remote control capabilities of the controller.

The first method of access takes advantage of an open debug port in the VxWorks OS. The debug port gives the attacker the ability to read and write arbitrary memory locations, kill tasks, and even reboot the device. A sophisticated attacker could use this weakness to reprogram the controller as desired. We created a program which used this vulnerability to dump the entire contents of memory from the controller.

The second access method uses the remote control functionality built into the controller itself, as defined by the NTCIP 1202 standard [27]. By reverse engineering the communications protocol, the attacker can craft a UDP packet which has the same effect as key presses on the front panel of the controller. The vendor has released a program to allow operators to remotely configure a controller through this protocol. By sniffing packets sent between the controller and this program, we discovered that communication to the controller is not encrypted, requires no authentication, and is replayable. Using this information, we were then able to reverse engineer parts of the communication structure. Various command packets only differ in the last byte, allowing an attacker to easily determine remaining commands once one has been discovered. We created a program that allows a user to activate any button on the controller and then displays the results to the user. We also created a library of commands which enable scriptable attacks. We tested this code in the field and were able to access the controller remotely.

4.4 Controlling the Lights

After gaining access to the controller, an adversary has a number of methods to attack the device. We focus only on attacks made possible by accessing the controller through its remote control functionality, as they are applicable even against controllers which have patched the VxWorks debug port issue. The two primary attack vectors are malicious logic statements and modified light timings.

The logic processor on the controller allows an operator to plan actions that will be executed when conditions are met. These commands are set up using simple if-then-else logic. Possible commands include switching the state of a specific light or even freezing the state of the entire intersection as shown in Figure 4. An attacker could choose to advance lights at a much faster or slower pace than normal if desired. Programmed statements remain present in the controller until removed by operators.

Controller operation can also be modified by changing the timing values of light states. The traffic controller we examined allows the minimum and maximum times for each light state to be altered. An attacker could use this method to shorten or lengthen the timings of specific light states, either to cause congestion or diminish safety.

While these attacks are capable changing controller outputs, the MMU maintains safety by disallowing many attacks. Any attack that causes conflicting green or yellow lights or which causes too short of a duration for yellow or red lights results in the MMU activating and setting the intersection to a known-safe configuration instead. Not all possible attacks are prevented however, as those which utilize only safe configurations are still possible (all-way red lights, short duration green lights, etc.).

4.5 Types of Attacks

The methods described above allow an attacker to gain access to the network and the controller. In this section we describe several possible attack scenarios and examine what kind of damage could be done by a dedicated adversary. This is by no means an exhaustive list.

Denial of Service A denial of service attack in this context refers to stopping normal light functionality. The most obvious way to cause a loss of service is to set all lights to red. This would cause traffic congestion and considerable confusion for drivers. Alternatively, the attacker could trigger the MMU to take over by attempting an unsafe configuration. This would cause the lights to enter a safe but suboptimal state. Since this state can be triggered remotely, but cannot be reset without physical access to the controller, an adversary can disable traffic lights faster than technicians can be sent to repair them. These attacks are overt and would quickly be detected by road agency personnel, who would be left with the recourse of disabling network connections between intersections.

Traffic Congestion More subtly, attacks could be made against the entire traffic infrastructure of a city which would manipulate the timings of an intersection relative to its neighbors. The effect would be that of a poorly managed road network, causing significant traffic congestion but remaining far less detectable than overt actions. This type of attack could have real financial impacts on a community. One study by the city of Boston calculated that simply reconfiguring the timings of 60 intersections in one district of the city could save $1.2 million per year in person-hours, safety, emissions, and energy costs [2].

Light Control An attacker can also control lights for personal gain. Lights could be changed to be green along the route the attacker is driving. Since these attacks are remote, this could even be done automatically as she drove, with the lights being reset to normal functionality after she passes through the intersection. More maliciously, lights could be changed to red in coordination with another attack in order to cause traffic congestion and slow emergency vehicle response.

5 Recommendations

There are several practical ways that transportation departments, traffic light operators, and equipment manufacturers can increase the security of their infrastructure. Good approaches to improving security include enabling encryption on wireless networks, blocking non-essential traffic from being sent on the network, and regularly updating device firmware. The simplest solution with the greatest amount of impact would be to change the default credentials on all network devices.

5.1 Wireless Security

Many of the issues we discovered during our investigation were problems with the wireless network configuration.

The 5.8 GHz radios used in the deployment are more vulnerable to attack than the 900 MHz radios. SSID broadcasting should be disabled on this network. While this does little to deter a determined adversary, it prevents casual observers and curious teenagers from noticing that the networks exist. The 5.8 GHz radios support WPA2 encryption and it should be enabled in the field.

Although we did not focus our attacks on the 900 MHz radios, an attacker with the proper resources could determine the frequency hopping pattern of the radios and reverse engineer the communication protocol. The radios we examined only allow for encryption types of WEP or WPA to be enabled. Ideally, radios in the field should support WPA2 encryption, but enabling WPA is better than no encryption at all.

Enabling encryption on either type of radio requires establishing a set of encryption keys. These keys should be unique across each master and slave radio pair, but even a single key used across the entire network would provide more security than the current setup.

5.2 Firewalls

If an attacker does manage to gain access to the network (wirelessly or physically), there are many vulnerabilities left open. To mitigate these attacks, we recommend reconfiguring the radios and switch to put restrictions on network traffic whenever possible. Only necessary communications should be allowed through. Unused ports should be blocked entirely, preventing attacks such as accessing the debugger functionality on the controller.

5.3 Firmware Updates

Firmware for embedded devices should be kept up to date. For field-deployed devices, this may not always be possible. In Cerrudo's study, vulnerable sensors are in some cases buried in the roadway, making firmware updates that require physical access impossible [3]. In these cases vendors should be clear with their customers as to what weaknesses exist, so that adequate measures can be taken. Other devices, such as traffic controllers and radios, are accessible and should be kept up to date.

The controllers we examined had outdated firmware, and were missing a critical security update that protected them against the VxWorks debug port issue. However, vendor update logs did not mention the debug port, and we only learned that the issue had been resolved after directly contacting the vendor. Transparency on the part of the vendor could have encouraged customers to update sooner. We encourage device vendors to maintain accurate information regarding changes made to their software and device users to update their firmware accordingly.

5.4 Changing Default Credentials

All of the devices in the deployment we studied used the default credentials that came built into the device. Past work has shown that this is a problem across a number of embedded devices [5], and this network is no exception. These default credentials are often available on the Internet and provide no security whatsoever. An adversary with network access could easily alter settings on

these devices and cause additional problems for traffic engineers, possibly locking them out of their own systems. Ideally all credentials on a device should be changed before it is deployed, and vendors should work to make this a requirement of the basic setup process.

At a minimum, we encourage device manufacturers to allow credentials to be changed. The traffic controllers used in the network we studied provided access to an FTP server to upload and download configuration files but did not allow the default username and password to be changed. Similarly, the configuration software for the 900 MHz radios only worked if the default device credentials were used, forcing administrators to choose between security and usability.

6 Broader Lessons

Our findings also carry lessons beyond their immediate implications for traffic infrastructure. In this section, we take a broader view and consider what this case study can teach us about embedded systems security in general.

Network Trust Our study shows futher evidence of a lack of layered security in embedded systems. It is not enough to trust that the network will be secure or physically isolated. The CAN bus within cars is considered trusted, but attackers can penetrate it with physical access to the vehicle [20] or by remotely exploiting vehicle subsystems [4]. SCADA systems often speak unencrypted protocols, yet many critical deployments have been found attached to the public Internet [16]. In our case study, an attacker can gain full control of traffic management devices over a wide deployment by compromising any point in a poorly secured wireless network.

The trusted network assumption fails when even a single vulnerability or misconfiguration exists anywhere in the network, leading to brittle defense that can fail dramatically in practice. A far safer security posture assumes that the network is open and adversarial, just like the public Internet. With this in mind, devices should apply end-to-end encryption and authentication using cryptographic transports such as SSH or TLS. Key management should be implemented in such a way that the compromise of a single device does not allow the compromise of communication with other devices. Deployments should still take steps to secure the network, but these should be regarded as only one part of a layered defense.

Hardware Failsafes While the deployment we studied was vulnerable to damaging attacks, the danger would have been far greater without the protections provided by the malfunction management units. Traffic signal systems use MMUs as a safety feature to ensure that lights enter a safe mode if the controller has a fault. Implementing such protection has two requirements: the ability to reliably detect that the controller is commanding an unsafe state, and the ability to override the controller to instead enter a failsafe mode. Designers of other embedded controllers should consider adding similar dedicated failsafe hardware when these prerequisites can be met.

Safety features do not necessarily provide increased security however. The MMU was conceived as a safety feature, and its value to security results from particularly conservative engineering. In another embedded system we have studied [25], software on the controller has the ability to override electronic safety controls. While this override facility would be unlikely to be engaged during a naturally occurring fault, an attacker could deliberately trigger it. In order to protect security as well as safety, failsafe circuitry must be completely isolated from any malicious software running in the controller.

Security Phase Changes The historical development of traffic control systems follows an unfortunately familiar pattern. Devices that were once purely electrical became first computer controlled and then eventually networked. Each advance seemed like the natural next step in the evolution of the technology, but when the devices grew into a ubiquitously computerized and networked system, they were also exposed to an array of new security risks. As our results suggest, the developers, customers, and regulatory framework were caught unprepared.

Other classes of embedded systems have undergone similar security "phase changes," including modern automobiles [4, 20], electronic voting machines [11, 19], and medical devices [13, 21]. In each case, the incremental addition of computer control and network connectivity led to large and weakly protected attack surfaces with the potential for catastrophic security failures.

Embedded system designers should be wary of this trend before it is repeated in other application domains. Rather than applying weak incremental defenses, leaving security for future work, or shirking responsibility entirely, designers should take a proactive approach to security. Fundamental changes to the design of the systems, protocols, and operating procedures may be necessary to address the emergent security risks.

7 Related Work

Our study is the first publicly available security evaluation of a deployed traffic infrastructure system.

There are various studies on the creation of adaptive traffic light control systems [7, 30, 35]. Few of these studies include any mention of security, much less an in-depth evaluation. There have been studies into protecting transportation infrastructure from both physical and computer-based attackers [6, 37]. These studies tend to be general in nature and do not evaluate specific computer-based security weaknesses. In contrast, our project evaluates an existing deployment of adaptive traffic lights.

Goodspeed previously reverse engineered the database

stored on an Econolite ASC/3 traffic controller [12]. Given access to the system, the database could be used to modify timing parameters and change light policies. This work is limited to attacks against the controller rather than against the entire infrastructure.

A recent study by Cerrudo of IOActive showed that wireless sensors can be spoofed to manipulate traffic light timings [3]. By reverse engineering the unsecured wireless protocol used by an in-ground vehicle detector, Cerrudo was able to gain full control of the sensor. Our study finds similar insecurities, but focuses on the wireless network and traffic light controller instead of the sensors.

The city of Austin, Texas performed a security evaluation of its traffic light system in 2011 [28]. Unfortunately, the results were considered confidential and were not disclosed to the public [29].

Security in Critical Infrastructure With regard to the transportation sector, the security of automobile systems has been studied by multiple groups, and various flaws in the design and implementation of these systems have been discovered [4, 20, 23]. Other critical infrastructure components such as the power grid have also received attention in recent years, and numerous problems have been exposed [1, 22]. Our work examines security issues in the same manner and spirit as some of these studies.

8 Future Work

We believe that the types of security issues discovered in this study are systemic across many manufacturers. An important area of research is the security of other critical infrastructure, such as the power grid and public water system. Much of this infrastructure has also undergone a phase change from independent nodes to a networked system and may have similar weaknesses.

Publicly Accessible Devices The traffic light controllers examined in this study return an SNMP banner that uniquely identifies them. We discovered several devices responding with this banner on the public Internet through Shodan [31]. Future work is needed to uniquely identify other traffic controller models and to determine if any traffic infrastructure is publicly accessible.

Connected Vehicles As complex as traffic infrastructure is today, the roadway of the future will be vastly more interconnected, sensor rich, and safety critical. The U.S. Department of Transportation (USDOT) has invested heavily in vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) systems since 2002, and has recently started the Connected Vehicle Safety Pilot program as a way of testing these systems in the field [26]. The University of Michigan's Transportation Research Institute (UMTRI) became involved with this program and has since outfitted nearly 3,000 vehicles with wireless devices that transmit vehicle information such as speed and location to one another [34]. The current deployment only tests a single V2V system that alerts drivers of a potential collision, but one can easily imagine vehicles coordinating with one another to allow autonomous driving, improved resource management, and more.

This paper shows that these types of systems often have safety in mind but may forget the importance of security. Depending on how information is transferred in these systems, adversaries may be able to inject falsified data into the V2V network. Even if such systems are designed with safety in mind, a lack of security could have dangerous consequences.

9 Conclusion

While traffic control systems may be built to fail into a safe state, we have shown that they are not safe from attacks by a determined adversary. With the appropriate hardware and a little effort, an adversary can reconfigure a traffic controller to suit her needs. She can execute a denial of service attack to cripple the flow of traffic in a city, cause congestion at intersections by modifying light timings, or even take control of the lights and give herself clear passage through intersections.

We have identified practical solutions which can be deployed by transportation departments and manufacturers to guard against some of these threats. Careful changes to existing networks can mitigate many of the risks we discovered without waiting for manufacturers to provide permanent solutions.

The real problem, however, is not any individual vulnerability, but a lack of security consciousness in the field. A clear example can be seen in the response of the traffic controller vendor to our vulnerability disclosure. It stated that the company, "has followed the accepted industry standard and it is that standard which does not include security." The industry as a whole needs to understand the importance of security, and the standards it follows should be updated to reflect this. Security must be engineered into these devices from the start rather than bolted on later. Until these systems are designed with security as a priority, the security of the entire traffic infrastructure will remain at serious risk.

Acknowledgments

We are extremely grateful to the personnel at the road agency who allowed us access to their network and hardware, but who, in the interest of protecting their infrastructure, wish to remain anonymous. We would also like to thank Zakir Durumeric for advice on the project. 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.

References

[1] R. Anderson and S. Fuloria. Who controls the off switch? In *IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 96–101, 2010.

[2] Boston Transportation Department. The benefits of retiming/rephasing traffic signals in the Back Bay, 2010. www.cityofboston.gov/images_documents/The%20Benefits%20of%20Traffic%20Signal%20Retiming%20Report_tcm3-18554.pdf.

[3] C. Cerrudo. Hacking US (and UK, Australia, France, etc.) traffic control systems, Apr. 2014. blog.ioactive.com/2014/04/hacking-us-and-uk-australia-france-etc.html.

[4] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno. Comprehensive experimental analyses of automotive attack surfaces. In *USENIX Security Symposium*, 2011.

[5] A. Cui and S. J. Stolfo. A quantitative analysis of the insecurity of embedded network devices: Results of a wide-area scan. In *Proceedings of the 26th Annual Computer Security Applications Conference*, pages 97–106. ACM, 2010.

[6] Department of Homeland Security. National infrastructure protection plan, 2013. www.dhs.gov/sites/default/files/publications/NIPP%202013_Partnering%20for%20Critical%20Infrastructure%20Security%20and%20Resilience_508_0.pdf.

[7] D. DeVoe and R. W. Wall. A distributed smart signal architecture for traffic signal controls. In *IEEE International Symposium on Industrial Electronics (ISIE)*, pages 2060–2065, 2008.

[8] M. M. Dobersek. *An Operational Comparison of Pre-Timed, Semi-Actuated, and Fully Actuated Interconnected Traffic Control Signal Systems*. PhD thesis, Marquette University, Milwaukee, Wisconsion, Dec. 1998.

[9] T. Elhabian, B. Zhang, and D. Shao. Fast de-hopping and frequency hopping pattern (FHP) estimation for DS/FHSS using neural networks. In *Advances in Neural Networks-ISNN 2004*, pages 248–253. Springer, 2004.

[10] ENCOM Wireless Data Solutions. Success stories. www.encomwireless.com/about-encom-wireless/success-stories.

[11] A. J. Feldman, J. A. Halderman, and E. W. Felten. Security analysis of the Diebold AccuVote-TS voting machine. In *Proceedings of the USENIX Workshop on Accurate Electronic Voting Technology*, 2007.

[12] T. Goodspeed. Reversing the Econolite ASC/3 traffic light controller. In *ToorCon Seattle*, 2008.

[13] D. Halperin, T. S. Heydt-Benjamin, B. Ransford, S. S. Clark, B. Defend, W. Morgan, K. Fu, T. Kohno, and W. H. Maisel. Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses. In *IEEE Symposium on Security and Privacy (SP)*, pages 129–142. IEEE, 2008.

[14] R. Havelt. Yes it is too Wi-Fi, and no its not inherently secure. *Black Hat Europe*, Mar. 2009.

[15] ICS-CERT. Vulnerability note VU#362332: Wind River Systems VxWorks debug service enabled by default, 2010. www.kb.cert.org/vuls/id/362332.

[16] ICS-CERT. Project shine. *ICS-CERT Newsletter Monthly Monitor,* October-December, 2012.

[17] Institute of Transportation Engineers. Detection trends, 2003. ite.org/standards/SignalSurvey.xls.

[18] Institute of Transportation Engineers. Results from the traffic signal systems requirements survey. www.ite.org/standards/surveyresults072003.pdf, July 2003.

[19] T. Kohno, A. Stubblefield, A. D. Rubin, and D. S. Wallach. Analysis of an electronic voting system. In *Proceedings of the 2004 IEEE Symposium on Security and Privacy*, pages 27–40. IEEE, 2004.

[20] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, et al. Experimental security analysis of a modern automobile. In *2010 IEEE Symposium on Security and Privacy (SP)*, pages 447–462, 2010.

[21] C. Li, A. Raghunathan, and N. K. Jha. Hijacking an insulin pump: Security attacks and defenses for a diabetes therapy system. In *13th IEEE International Conference on e-Health Networking Applications and Services (Healthcom)*, pages 150–156, 2011.

[22] J. Liu, Y. Xiao, S. Li, W. Liang, and C. Chen. Cyber security and privacy issues in smart grids. *Communications Surveys & Tutorials*, 14(4):981–997, 2012.

[23] C. Miller and C. Valasek. Adventures in automotive networks and control units, 2013. illmatics.com/car_hacking.pdf.

[24] Minnesota Department of Transportation. Traffic signals 101 - controller operations, Jan. 2014. www.dot.state.mn.us/trafficeng/publ/signals101/2014/06_Controller_Ops.pdf.

[25] K. Mowery, E. Wustrow, T. Wypych, C. Singleton, C. Comfort, E. Rescorla, S. Checkoway, J. A. Halderman, and H. Shacham. Security analysis of a full-body scanner. In *23rd USENIX Security Symposium (USENIX Security 14)*, Aug. 2014.

[26] National Highway Traffic Safety Administration. U.S. Department of Transportation announces decision to move forward with vehicle-to-vehicle communication technology for light vehicles, Feb. 2014. www.nhtsa.gov/About+NHTSA/Press+Releases/2014/USDOT+to+Move+Forward+with+Vehicle-to-Vehicle+Communication+Technology+for+Light+Vehicles.

[27] National Transportation Communications for ITS Protocol. 1202 - NTCIP object definitions for ASC, 2007. www.ntcip.org/library/standards/default.asp?documents=yes&qreport=no&standard=1202.

[28] Office of the City Auditor - Austin, Texas. Status report: Performance audit of traffic signal synchronization and security, Feb. 2011. www.austintexas.gov/sites/default/files/files/Auditor/au11109m.pdf.

[29] Office of the City Auditor - Austin, Texas. Traffic signal security audit, Mar. 2011. www.austintexas.gov/sites/default/files/files/Auditor/au11109.pdf.

[30] R. Sen and B. Raman. Intelligent transport systems for

Indian cities. In *6th USENIX/ACM Workshop on Networked Systems for Developing Regions*, 2012.

[31] Shodan. SHODAN - computer search engine. www.shodanhq.com/.

[32] Siemens. Traffic control via the Siemens private cloud, Mar. 2014. www.siemens.com/press/pool/de/events/2014/infrastructure-cities/2014-03-intertraffic/background-private-cloud-e.pdf.

[33] Tennessee Department of Transportation. TDOT smartway. www.tdot.state.tn.us/tdotsmartway/.

[34] University of Michigan Transportation Research Institute. What is the safety pilot model deployment?, Aug. 2012. safetypilot.umtri.umich.edu/index.php.

[35] E. Vorakitolan, J. P. Havlicek, M. Atiquzzaman, and R. D. Barnes. Exploiting trunked radio to support ITS network expansion and redundancy. In *22nd IEEE International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC)*, pages 761–766, 2011.

[36] B. W. Yarger. Fully actuated vs. semi-actuated traffic signal systems, 1993. www.yargerengineering.com/fully_vs_semi-actuated.pdf.

[37] D. Zeng, S. S. Chawathe, H. Huang, and F. Y. Wang. Protecting transportation infrastructure. *Intelligent Systems, IEEE*, 22(5):8–11, 2007.

[38] K. Zetter. Hackers can mess with traffic lights to jam roads and reroute cars. *Wired*, Apr. 2014. www.wired.com/2014/04/traffic-lights-hacking/.