Security of Autonomous Systems Employing Embedded Computing and Sensors

Alexander M. Wyglinski
Xinming Huang
Taskin Padir
Lifeng Lai
Thomas R. Eisenbarth
Krishna Venkatasubramanian
Worcester Polytechnic Institute

Society is becoming increasingly dependent on embedded computing and sensor technology to enable complex networks of autonomous systems, such as robots, unmanned aerial vehicles (UAVs), self-driving cars, and unmanned underwater vehicles (UUVs). In fact, several recent developments highlight the realization of advanced autonomous systems that were deemed science fiction just a few decades ago. One example is UAVs being extensively deployed in missions around the world to perform numerous operations, including reconnaissance and intelligence gathering, war fighting, and scientific research (see Figure 1a). Another example is the DARPA Grand Challenge and DARPA Urban Challenge, which yielded several sophisticated implementations of self-driving ground vehicles that can drive long distances and challenging driving environments without a human operator's assistance. In addition, there have been public demonstrations of Google's Self-Driving Toyota Prius, which has by recent accounts logged more than 160,000 miles[1] (see Figure 1b). Nevada recently legalized the use of self-driving cars on public roads, with another half-dozen states following suit. Automobile manufacturers such as General Motors are actively investigating approaches for deploying partially autonomous automobiles,[2] including vehicular formations driving under highway conditions referred to as ''platooning.''

Overall, these networks of robotic and autonomous systems should significantly improve the quality of life in several areas of society, including search-and-rescue operations, national defense, surveillance and border protection, and ground transportation of materials and goods. Three key technologies are needed to enable these advanced autonomous systems. First, at least one embedded processor is necessary to guide and control a specific platform's activities. In many applications, the autonomous platform can possess more than one embedded processor, where each processor is assigned a dedicated task. For example, most vehicles on the road today possess at least 80 embedded processors, all of which are interconnected with each other.

Second, an array of sensors is needed to provide the autonomous system and its collection of embedded processors with situational awareness about the physical world. The embedded processors use this information to make the appropriate decisions about what actions the autonomous system should perform.

The last necessary component is a communication system, which relays information to and receives information from a command center as well as information shared by other autonomous systems that form part of the network. Embedded processors in these complex autonomous systems often do not act alone, but rather in concert with each other. Consequently, information sharing between embedded processors within the same autonomous system or between two different systems is crucial for many operations.

As with many of these complex networks of systems, it's possible for external intruders to intentionally compromise the proper operation and functionality of these systems. However, unlike complex networks such as the Internet, where the issue of security has been

Published by the IEEE Computer Society

extensively researched and funded, security issues surrounding complex networks of autonomous automotive systems haven't been as readily studied. Moreover, these systems' security vulnerabilities are increasingly being discovered and exploited.

Technical challenges and potential threats

Despite substantial investments in creating and perfecting unmanned autonomous vehicles, one key aspect is noticeably absent from these systems' design: security. To our knowledge, little research has been conducted in the area of securing unmanned autonomous platforms. Almost all of it has focused on simply encrypting all data, on both the embedded system and the wireless channels, without assessing other potential vulnerabilities. In fact, researchers have already demonstrated several other potential vulnerabilities on actual hardware platforms and published them in the open literature. For example, several researchers have explored embedded-computing and sensor-system vulnerabilities on commercial vehicles, which can be accessed by nonconventional methods such as the vehicle's entertainment system or tire-pressure sensors.[3,4] Researchers have also explored exploiting these commercial vehicles' firmware-updating mechanisms as a potential vulnerability,[5] as well as attacks carried out over the wireless channels connecting the vehicular platform to some information network.[6,7] Cryptographic attacks have also been demonstrated on these platforms,[8] and techniques are being developed to assess whether the embedded-computing and sensor-system resources are being compromised by an attack.[9] GPS spoofing is another research topic being explored, and the results of this research could impact the navigation of unmanned systems, as demonstrated recently with UAV platforms.[10]

Consequently, the investigation of all possible forms of vulnerabilities contained within these embedded processor-based autonomous systems, as well as the development of the appropriate defenses and countermeasures, is essential for guaranteeing the reliable operation of these platforms and the growth of this sector. Given these autonomous systems' sophistication and increased susceptibility to malicious activities due to their operations and interactions in an open environment, security vulnerabilities of autonomous systems go well beyond simple exploits such as information leakage of digital data. Thus, the engineering community must identify potential threats and devise a collection of corresponding and effective countermeasures.

(a)

(b)

Figure 1. Two examples of autonomous vehicles deployed in real-world operations. A NASA Global Hawk unmanned aerial vehicle (UAV) platform conducting atmospheric measurements (a). A pair of Global Hawk UAV platforms recently performed close-flight-formation maneuvers that could support autonomous high-altitude operations such as in-flight refueling.[11] A Google driverless car operating on a testing path (b). (NASA figure courtesy of Carla Thomas. Google car photo by Steve Jurvetson.)

A case study: Automotive computing and sensing

Autonomous automotive systems require accurate situational awareness of the physical environment around them to function properly. As a result, we

Table 1. Electronic control units (ECUs) commonly found in commercial automobiles.

ECU	Functionality
Powertrain	Essentially the ''brain'' of the engine control system.
	Commonly controls more than 100 factors in a car or truck.
	Handles charging system, transmission, various emission controls, and communications with other onboard control modules.
Safety systems	Responsible for handling various driver safety features of the vehicle.
	Operations include collision avoidance, airbag deployment, and active braking.
Body control	Responsible for monitoring and controlling various electronic accessories in a vehicle's body.
	Operations include power windows and mirrors, air conditioning, immobilizer system, and central locking.
Data communications	Handles data communications between different components within the vehicle, between two or more vehicles, and between the vehicle and another device or roadside infrastructure.
	Operations include Bluetooth connectivity, cellular access, and DSRC safety communications.

must leverage sensory information and control data together with available embedded-computing resources between the different platforms to enable reliable autonomy for these systems. With researchers continuously discovering new security vulnerabilities for these platforms, research activities focusing on automobiles' physical security have the potential to be transformative and high-impact given our growing reliance on this technology and its evolution into semi- and fully autonomous platforms. Furthermore, research activities focusing on secure autonomous automotive systems span numerous research concentrations in electrical engineering, computer engineering, and computer science, with many of the issues and their solutions lying at the boundaries of embedded processing, hardware security, robotics, sensor systems, cyber-physical systems, and communications. Every action performed by today's automotive vehicles is increasingly being handled by a computing device called an *electronic control unit*. ECUs are responsible for various operations, from power locks, seat adjustments, and automotive stereo systems to power steering, fuel injection, and emissions control. Table 1 gives a list of typical ECUs. Given the growing threat of attacks against these systems, researchers have been working on identifying security threats toward these ECUs, as well as automotive sensor systems, with most experiments in this area involving conducting physical security attacks against an actual road vehicle.[3,12] Similarly, several researchers have devised ''hardware-in-the-loop'' computer simulations that incorporate the input and output data paths of several automotive components.[13]

Anatomy of an autonomous automotive platform

According to a recent *IEEE Spectrum* article, a typical car today includes 70 to 100 microprocessor-based ECUs, which control the engine, power train, transmission, brakes, body, doors, dashboard, tires, and heating, ventilation, and air conditioning (HVAC).[1] An autonomous vehicle requires an array of sensors, such as laser, radar, light detection and ranging (LiDAR), GPS, and computer vision systems, to gather information. It also requires more advanced computer systems to process information to real-time navigation and controls.[14] Figure 2 illustrates the various functions that could potentially be employed on an autonomous automotive platform. Table 2 lists the key ECUs identified by functionality. Typically, these ECUs are connected using the CAN (controller area network) bus[15] or FlexRay[16] network. Hence, ECUs can communicate with each other and coordinate their actions. Some vehicles are also equipped with telematics for remote diagnostics, monitoring, and additional communications. Previous studies have shown that an autonomous vehicle possesses software consisting of over 100 million lines of code.

Autonomous vehicles rely heavily on their onboard sensors to navigate in an environment consisting of static and dynamic obstacles. Onboard sensors also enable an autonomous vehicle to localize itself relative to a geographic map, monitor vehicle health, and diagnose system faults. An autonomous vehicle typically incorporates the following sensor systems:

- A differential GPS receiver that provides absolute position data with submeter accuracy. This data can also be used to derive the vehicle velocity.
- An inertial measurement unit that provides velocity, acceleration, and orientation data using a combination of accelerometers and gyroscopes. The data can be integrated to calculate vehicle position.
- LiDAR sensors that can generate a map of the vehicle's environment that is used for localization, obstacle avoidance, and navigation.
- Cameras (stereo- or monovision) that provide static and dynamic obstacle detection, object recognition, and 360° surrounding information when fused with other sensors.
- Ultrasonic or infrared-range sensors that provide redundancy and are effectively used for parking, detecting vehicles, and obstacles.

Figure 2. An autonomous automotive system with an array of sensors, embedded computing devices, and communication systems. One or more of these components could be compromised in a physical security attack.

Consequently, leveraging the information obtained from a combination of multiple onboard sensors, various decision-making processes of an autonomous vehicle can choose the appropriate actions in real time over a wide range of operating environments.

Table 2. Key ECUs in an autonomous car.

Functionalities	Level of importance
Navigation control module (NCM)	✓✓✓✓
Engine control module (ECM)	✓✓✓
Electronic brake control module (EBCM)	✓✓✓
Transmission control module (TCM)	✓✓✓
Telematics module with remote commanding	✓✓✓
Body control module (BCM)	✓✓
Inflatable restraint module (IRM)	✓✓
Vehicle vision system (VVS)	✓✓
Remote door lock receiver	✓
Heating, ventilation, and air conditioning (HVAC)	✓
Instrument panel module	✓
Radio and entertainment center	Nonessential

Classification of vulnerabilities on autonomous vehicles

Identifying potential security vulnerabilities associated with the embedded computing and sensor systems of unmanned ground vehicles (UGVs), UAVs, and complex networks of cooperating UAVs and UGVs deployed to carry out a specific mission is technically challenging. For example, it's possible to employ active attack techniques targeting the sensor technologies commonly used in autonomous platforms, such as ultrasound sensors, infrared sensors, and hall-effect sensors.[17] An autonomous platform decides its action on the basis of inputs received from these sensors. By attacking the sensors, a malicious user can cause autonomous platform suicides and vandalism, as well as perform denial-of-service (DoS) attacks, and can even gain full control of the autonomous platform itself. Furthermore, a malicious user can compromise an autonomous system to eavesdrop on the wireless transmissions between vehicles and other information infrastructure in order to obtain information about potential actions, and can remotely access systems on the platforms themselves to create confusion among a network of autonomous platforms operating in concert with each other on a collaborative project.

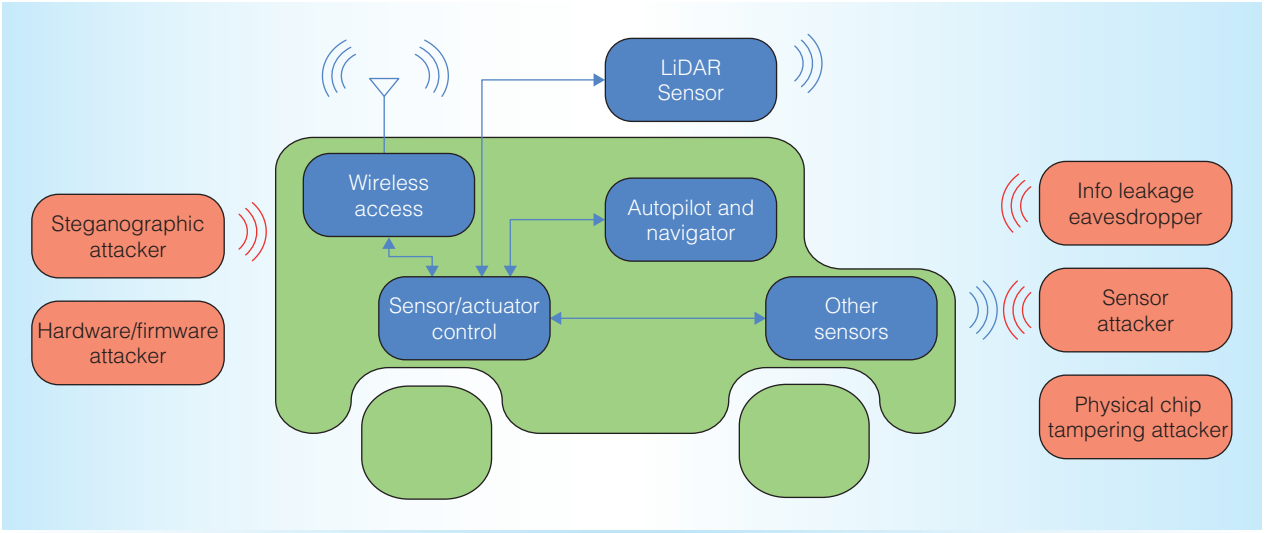

Figure 3. Various embedded computing and sensor systems employed in a typical automotive platform and the associated forms of physical security attacks.

Finally, in several instances, a malicious user can physically tamper with or attack the actual embedded computing and sensor hardware.

We can categorize these attacks as either physical attacks (for example, physically accessing or tampering with the systems), close-proximity attacks (obfuscating sensor information and injecting faults in embedded processors within ranges of approximately 10 meters), or remote attacks (accessing embedded-computing and sensor-system resources via unsecure data network paths over remote wireless access). Figure 3 illustrates these types of attacks.

Physical-access vulnerabilities. The large number of ECUs and the CAN-type networks and software running on the ECUs provide several opportunities for attacking autonomous vehicles, such as hardware vulnerabilities that can potentially remove, destroy, or replace the hardware modules in a car. Two such security vulnerabilities are bus tapping vulnerability on the CAN bus and ECU reprogramming through the CAN bus.

Close-proximity vulnerabilities. Autonomous vehicles have many types of sensors, such as stereo-vision cameras, LiDAR sensors, and additional ultrasonic and infrared range finders, GPS receivers, an inertial measurement unit (IMU), and encoders. This creates numerous opportunities for a malicious user. For instance, a malicious user could just keep sending the vehicle false signals to disorient, hijack, or even restrict it to within a virtual fence.[17]

Intelligent vehicles are manufactured with certain safety measures to ensure that the sensor and control system are reliable, fault-tolerant, and effective in navigating dynamic and clustered environments. However, these safety measures are against nonmalicious faults and usually aren't equipped with adversary detection and prevention mechanisms. Once a vehicle is manufactured and tested against natural errors, it's expected to conform to its design specifications unless there's a mechanical or electrical failure. Three possible close-proximity vulnerabilities involving the onboard sensor systems are interference with onboard ultrasonic sensor measurements, LiDAR sensor measurements, and vision systems.

Remote-access vulnerabilities. Although the CAN bus and the embedded microprocessor systems it supports aren't physically connected to anything outside the vehicle itself, the vehicle is still becoming increasingly connected to various information infrastructure, as well as to other vehicular platforms, via wireless data transmission techniques. Consequently, wireless connectivity provides an opportunity to enable new applications, such as the exchange of position data and operational information between vehicles, performing software and firmware updates for the various embedded microprocessor systems onboard the vehicle remotely via the wireless network, and providing increased driver awareness, range, and perception of the environment to further enhance road safety.[18]

The opportunities offered by wireless systems also make them highly susceptible to various forms of malicious activities. Three possible vulnerabilities aimed at vehicular communication systems and hardware are exploiting standardized wireless interfaces for DoS attacks and eavesdropping via wireless side-channel attacks, introducing malicious code to an embedded computing device onboard the vehicle, and inducing incorrect behavior in a distributed network of autonomous automotive platforms operating together on the road.

Human-vehicle interface

Although we can employ technology to detect numerous possible vulnerabilities involving autonomous vehicles, we can't ignore the significant human element—specifically, how to deal with a nontechnical human operator of an autonomous vehicle who is being compromised by a malicious user. We must ask three questions.

First, when a malicious user compromises a vehicle, can the vehicle enter a ''safe mode'' operation in order to allow the vehicle to ''fail soft'' such that the human operator's safety is enhanced? Is it possible for the vehicle to decide to enter a safe mode in real time?

Second, when malicious activity is detected, how should the human operator be informed of the situation, and how much control over the vehicle should the operator be given? Is there a universal approach to inform the operator that the vehicle is being compromised, and are there protocols for handling such situations when they are encountered?

Third, given that many successful security breaches of autonomous systems occur whenever the system's situational awareness is lacking (for example, stealth attacks), are there mechanisms that can collect sensory and other diagnostic information into a single decision-making process in order to flag a suspicious or potentially harmful situation? Should the system employ a ''deny all'' approach to rejecting possible attacks, followed by an investigation? How does the human operator get involved in the decision-making process? Given the interface between the autonomous vehicle and the human operator, it's important to ensure that the latter is brought into the loop whenever the former experiences a scenario that doesn't conform to normal specifications.

These technical challenges and issues surrounding autonomous system security have opened many research opportunities that can be pursued to make these systems safer, more reliable, and resilient to malicious attacks. Several potential directions for research exist.

One topic is investigating real-time human-in-the-loop control algorithms for supporting multiple unmanned platforms via a single human operator. Leveraging traditional approaches for controlling multiple platforms, such as wireless networking and localization information obtained from GPS devices, these unmanned platforms will also use their sensor systems to extract information about the actions of the other unmanned platforms within the network and extract from this information their updated role in the mission. In challenging environments where an adversary is attacking the traditional mechanisms for control of a large network of autonomous systems, these other sources of information will be used to continue operations uninterrupted.

Another direction is in creating efficient, lightweight cryptographic algorithms for realizations requiring a balance of computing, memory, communications, and energy to reliably protect these unmanned systems from attacks designed to compromise the system or network of systems by exploiting the leakage of critical information, such as power, electromagnetic emissions, and execution time.

Researchers can also devise techniques for enabling trust in UAVs, UGVs, and cooperative networks of these platforms, as well as reliably identifying when these individual platforms or network of platforms are being attacked. Trusted cooperation and attack detection are vital when vehicles interact with each other and share information, because they're designed to minimize the risk of a malicious user inserting false information into the vehicle or network of vehicles.

It will also be important to develop robust and reliable algorithms for sensor fusion, and motion control for autonomous systems under physical attacks. Furthermore, understanding how to implement minimal yet sufficient redundancies in these embedded system designs to ensure proper operation of autonomous systems under physical attacks is a significant area of research opportunity.

Another research direction is identifying new side-channel attacks and developing countermeasures.[19] This includes evaluating trusted platform module (TPM) microcontrollers.[20] Recently, several major manufacturers have produced TPM chips. Additionally, researchers must assess and study techniques for protecting field-programmable gate array (FPGA) designs from attacks, including bitstream decoding, spoofing, Trojan horses, side-channel, and fault insertion.[21] If attackers deciphered the information-processing algorithms from the encoded bitstream, they could design attacks to spoof the algorithms. Even worse, they could implant an injected-code Trojan horse into the FPGA design. When it is activated, it can reveal the authentication code, compromise the circuits, and even break the system. It's difficult to detect an FPGA Trojan horse due to the large circuit size and complicated design flow.

Finally, researchers must enhance the firmware security of embedded computing devices on these platforms such as protecting the firmware from counterfeiting.[22] Typically, firmware can be protected using a security built-in microprocessor, which is usually more expensive, or a cryptographic authentication IC, also referred to as a security coprocessor. Furthermore, new techniques for performing secure remote firmware updates are needed.[23] For practicality and functionality, updating remote firmware is considered important. However, connecting the processor over the network greatly increases the security risk. If attackers gained access to the remote firmware update channel, the firmware's security would be broken down. MICRO

References

1. E. Guizzo, ''How Google's Self-Driving Car Works,'' *IEEE Spectrum* Automation blog, 18 Oct. 2011; http://smartmachines.blogspot.com/2007/10/reinforcement-learning-is-cool.html.
2. A. Barari, ''GM Promises Autonomous Vehicles by End of Decade,'' *Motorward,* 17 Oct. 2011; http://www.motorward.com/2011/10/gm-promises-autonomous-vehicles-by-end-of-decade.

3. K. Koscher et al., ''Experimental Security Analysis of a Modern Automobile,'' *Proc. IEEE Symp. Security and Privacy,* IEEE, 2010, pp. 447-462.
4. P. Kocher et al., ''Security as a New Dimension in Embedded System Design,'' *Proc. 41st Ann. Design Automation Conf.* (DAC 04), ACM, 2004, pp. 753-760.
5. D.K. Nilsson and U.E. Larson, ''Secure Firmware Updates over the Air in Intelligent Vehicles,'' *Proc. IEEE Int'l Conf. Communications,* IEEE CS, 2008, pp. 380-384.
6. D.K. Nilsson and U.E. Larson, ''Conducting Forensic Investigations of Cyber Attacks on Automobile In-Vehicle Networks,'' *Proc. 1st Int'l Conf. Forensic Applications and Techniques in Telecommunications, Information, and Multimedia,* ICST, 2008, no. 8.
7. U.E. Larson and D.K. Nilsson, ''Securing Vehicles Against Cyber Attacks,'' *Proc. 4th Ann. Workshop Cyber Security and Information Intelligence Research,* ACM, 2008, no. 30.
8. A. Francillon, B. Danev, and S. Capkun, ''Relay Attacks on Passive Keyless Entry and Start Systems in Modern Cars,'' *Proc. 18th Network and Distributed System Security Symp.,* The Internet Soc., 2011.
9. D.K. Nilsson, P.H. Phung, and U.E. Larson, ''Vehicle ECU Classification Based on Safety-Security Characteristics,'' *Proc. IET Road Transport Information and Control Conf.,* IEEE CS, 2008, pp. 1-7.
10. M. Mixon, ''Todd Humphreys' Research Team Demonstrates First Successful GPS Spoofing of UAV,'' 2012, http://www.ae.utexas.edu/news/archive/2012/todd-humphreys-research-team-demonstrates-first-successful-gps-spoofing-of-uav.
11. C. Kwan, ''Two Global Hawk Unmanned Aircraft Fly in Close Formation, Move AHR Program Closer to Autonomous Aerial Refueling,'' 5 Oct. 2012; http://globenewswire.com/news-release/2012/10/05/495475/10007491/en/Multimedia-Release-Two-Global-Hawk-Unmanned-Aircraft-Fly-in-Close-Formation-Move-AHR-Program-Closer-to-Autonomous-Aerial-Refueling.html.
12. U. Drolia et al., *Autoplug: An Automotive Test-Bed for Electronic Controller Unit Testing and Verification,* tech. report, School of Eng. and Applied Science, Univ. of Pennsylvania, 2011.
13. D. Lemp, S. Köhl, and M. Plöger, *ECU Network Testing by Hardware-in-the-Loop Simulation,* tech. report, dSPACE GmbH, 2004.
14. M. Aoyama, ''Computing for the Next Generation Automobile,'' *Computer,* June 2012, pp. 32-37.
15. *CAN Specification: Version 2.0,* tech. report, Robert Bosch GmbH, 1991.
16. S. Lorenz, ''The FlexRay Electrical Physical Layer Evolution,'' *Automotive,* 2010, pp. 1-6.
17. K. Akdemir et al., ''An Emerging Threat: Eve Meets a Robot,'' *Proc. 2nd Int'l Conf. Trusted Systems,* LNCS 6802, Springer, 2011, pp. 271-289.
18. S. Chen et al., ''Feasibility Analysis of Vehicular Dynamic Spectrum Access via Queueing Theory Model,'' *IEEE Comm. Magazine,* Nov. 2011, pp. 156-163.
19. D. Agrawal et al., ''Trojan Detection Using IC Fingerprinting,'' *Proc. IEEE Symp. Security and Privacy,* IEEE, 2007, pp. 296-310.
20. S. Kinney, *Trusted Platform Module Basics: Using TPM in Embedded Systems,* Elsevier, 2006.
21. C. Hu, *Solving Today's Design Security Concerns,* tech. report, Xilinx, 2010.
22. C. Gorog, ''Protect Firmware from Counterfeiting,'' 2011; http://www.embeddedintel.com/specialfeatures.php?article=126.
23. L.K. Shade, ''Implementing Secure Remote Firmware Updates,'' *Proc. Embedded Systems Conf.,* UBM Electronics, 2011.

Alexander M. Wyglinski is an associate professor in the Electrical and Computer Engineering Department at Worcester Polytechnic Institute. His research interests include wireless communications, cognitive radio, software-defined radio, cyber-physical systems and security, and wireless system optimization and adaptation. Wyglinski has a PhD in electrical engineering from McGill University. He is a senior member of IEEE.

Xinming Huang is an associate professor in the Electrical and Computer Engineering Department at Worcester Polytechnic Institute. His research interests include circuits and systems design for reconfigurable computing, wireless communications, signal processing, and cyber security. Huang has a PhD in electrical and computer engineering from Virginia Tech. He is a senior member of IEEE.

Taskin Padir is an assistant professor in the Electrical and Computer Engineering and Robotics Departments at Worcester Polytechnic Institute. His research interests include the design, modeling, and control of robotic systems and intelligent vehicles. Padir has a PhD in electrical and computer engineering from Purdue University. He is a member of IEEE.

Lifeng Lai is an assistant professor in the Electrical and Computer Engineering Department at Worcester Polytechnic Institute. His research interests include information theory, stochastic signal processing, and wireless network security. Lai has a PhD in electrical and computer engineering from Ohio State University. He is a member of IEEE.

Thomas R. Eisenbarth is an assistant professor in the Electrical and Computer Engineering Department at Worcester Polytechnic Institute. His research interests focus on embedded systems security. Eisenbarth has a DrIng in electrical and computer engineering from Ruhr University Bochum, Germany. He is a member of IEEE.

Krishna Venkatasubramanian is an assistant professor in the Computer Science Department at Worcester Polytechnic Institute. His research interests focus on cyber-physical systems and security. Venkatasubramanian has a PhD in computer science from Arizona State University. He is a member of IEEE.