

# Enhancing Cyber-Physical Security through Data Patterns

Richard Chow, Ersin Uzun  
Palo Alto Research Center  
Palo Alto, CA USA  
{rchow, euzun}@parc.com

Alvaro A. Cárdenas, Zhexuan Song, Sung Lee  
Fujitsu Laboratories of America  
Sunnyvale, CA US  
{alvaro.cardenas-mora, zhexuan.song, sung.lee}@us.fujitsu.com

**Abstract**—In this position paper, we propose a data-driven approach for security management in a network that interacts or receives inputs from physical systems – including human behavior. Our goal is to leverage the unique features of cyber-physical systems. In particular we propose: (1) the use of historical data from physical systems and human behaviors to enable anomaly detection, (2) the use of contextual data from multiple and diverse sensor readings to obtain a higher-level collective vision of the network for better event correlation and decision analysis, and (3) the use of physical sensor data and human behavior to enable fine-grained, dynamic access control and implicit authentication. We outline use cases describing how our ideas can be applied in the Home Area Network (HAN).

**Keywords**—home area networking; authentication; security; privacy; anomaly detection

## I. INTRODUCTION

As cyber-physical systems become more pervasive and important they will increasingly become the focus of attacks. One prominent recent example is Stuxnet. The Stuxnet worm targets industrial controllers and is believed to target the uranium enrichment infrastructure in Iran [6], [2], [3]. In addition, the rise of ubiquitous computing and the *Smart Grid* imply the deployment of billions of smart sensors and actuators embedded in our social infrastructures. By relying on information technology and networks, these new deployments will expose our social infrastructure to regular computer vulnerabilities available to an ever-growing set of motivated and highly-skilled attackers.

While many existing security mechanisms can be applied to cyber-physical systems, in this paper we explore some unique ways to enhance the security of these systems by leveraging the diverse physical and human behavior information collected by these systems. Examples include: (1) the use of historical data from physical systems and human behaviors to enable anomaly detection, (2) the use of contextual data from multiple and diverse sensors to obtain a higher-level collective vision of the network for better event correlation and decision analysis, and (3) the use of physical sensor data and human behavior to enable implicit authentication and fine-grained, dynamic access control.

For point (1), we propose the use of device and human profiles generated from a data driven approach. For example, a historical profile of the alarm system of a house that leaves

the alarm on every night might generate an incident report if one day the alarm is shut off at 3 AM in the morning. In general, previous work has shown that cyber-physical systems might benefit from anomaly detection techniques, as physical processes give off large amounts of data that are clearly non-deterministic in nature, and yet somewhat predictable [7].

While the use of physical process data has already been proposed as a way to detect computer attacks on cyber-physical systems [9], we believe that in order to obtain a better and more complete view of the system, the data of multiple and diverse sensors needs to be aggregated and combined to generate usable models with low false alarm rates. Therefore, we propose for point (2), the use of a master controller that collects multiple data sources and integrates them into the proper context. For example, while shutting off the home alarm system at 3 AM in the morning might be anomalous in itself, it might not be anomalous if the electric car associated with the home has also just pulled into the garage and was plugged into its charger. We argue in this position paper, especially in the case of the Home Area Network (HAN), that higher level patterns that involve collective behavior of devices and users should also be analyzed for anomalies. User behavioral patterns are also non-deterministic and yet somewhat predictable.

Finally, for point (3), we believe that the user and device profiles obtained by (1), and the contextual information obtained by (2) can be used to enhance user authentication, dynamic access control, revocation, and fine-grained access control policies. For example, a typical home area network currently provides full security or no security at all. It is very difficult for users at home to set the appropriate fine-grained access control to devices at home, and the proper revocation mechanisms. For example, if a neighbor brings their laptop to a home and is given the keys to access the wireless network, the neighbor will remain a valid (authenticated) user of the network even in his own home. We propose the use of device profiles for *implicit authentication*, allowing the revocation of devices that do not match our experience.

In the remainder of this paper we explore in more detail our arguments by providing additional use cases, a general system description, and a discussion of the security, usability, and performance of the system.

## II. ADDITIONAL USE CASES

### Collective Device Behavior

Higher-level components of cyber-physical systems may also exhibit anomalous behavior, and this behavior is best analyzed given appropriate context. A simple example is that intensive operation of an air-conditioning system is not anomalous given a heat wave, but may be under normal weather conditions. A more complex example is that shutting off the home alarm system at 3 AM might be very normal if the electric car has just pulled into the garage and is plugged into its charger.

Finding these sorts of rules may require sophisticated machine learning apparatus, but the advantage is more intelligent security decisions and fewer false positives. In large control systems this feature is usually called *situational awareness*.

Consider for example a home network that controls appliances, heating and cooling, and lighting. In most networks available to consumers, there is typically a lack of fine-grained access control: access to the central controller implies total control of the system. However, adding such fine-grained access control runs the danger of making the system unusable.

Instead, we propose adding a behavioral analysis system on top of the controller security system. We amass a collective baseline of historical settings, readings, and network traffic. The baselines can be used to detect intrusion, both cyber and physical (as well as broken or misconfigured equipment). Detected anomalies in a device may result in revocation from the network, or a request for re-authentication or two-factor authentication. Revocation in home area networks is not an easy task to accomplish in current mass market deployments.

Collective contextual information can also be used to design more flexible access control mechanisms. For example, with the proliferation of surveillance cameras in everyday life and webcams at home computers, the number of unsecured cameras on the Internet has become an increasing cause of concern. Bloggers have reported the ability to tap into thousands of raw webcam feeds with a few simple Google searches, and the Spanish police arrested a suspect on charges of developing a computer virus that can activate a webcam without the owner's permission [5]. We propose that the context for camera networks might be used in the access control policy. For instance, the context contains properties such as the content of the video stream (e.g., people or the type of events happening), when the access request is made, and location of the subject. These properties add the flexibility of describing rules such as "only people in a room should be able to control the camera located in that room." Camera networks might also be part of complex networks including other sensors such as audio, temperature, humidity etc. The readings of these sensors might give additional

context relevant to the access control policy. For example, emergency response teams (firefighters or paramedics) may be allowed to access any web camera if the fire alarm is on.

### User behavior

The cyber-physical systems of tomorrow are all envisioned to be connected and remotely controllable. With the advent of the smart grid, smart appliances within a consumer's HAN can be controlled remotely by consumers and potentially by utility companies. See, for example, [1] for the latest iPhone applications for remotely controlling the home and [4] for the PG&E SmartAC program. Given the ubiquity of Internet-connected smartphones, it is clear that cyberattacks against such devices is a way to attack the HAN. In fact, as argued by Neuman [10], creating a smartphone botnet may be a relatively easy way to generate traffic affecting the large scale power grid.

Given the possibility of direct intrusion into the home, the proper authentication of users becomes even more critical. One approach is to protect each remote control instance, for example requiring traditional authentication mechanisms (such as a password, one-time-password, or user certificate) for each remote control directive. This strategy may provide a reasonable level of security when used properly (for example, unpredictable passwords, secure storage of certificates, etc.), but does not scale well when the number of devices increases. Imagine, for a moment, the following common scenario: you have many personal Internet-connected devices such as a laptop, desktop, smartphone and even an Internet-connected TV that you use to control all the appliances in your HAN, such as the air-conditioning unit, alarm, electric car, lights and other smart appliances. Having a separate password for all these devices and authenticating to each every time you need access is very unusable while having the same password for all of them would significantly reduce the security.

Another approach is to authenticate the user once and provide general remote control capabilities after this authentication, at least for some limited period of time. These technologies reduce the burden on the user by requiring, say, only one password. However, these technologies not only ease access for legitimate users but ease it for attackers as well. A natural answer is to add an additional authentication factor, but traditional second factors such as hardware tokens or biometrics reduce usability, especially on a mobile platform.

We propose that patterns of user behavior can be used as a second factor, in line with the theme of this paper of using behavioral patterns to increase security. For mobile phones in particular, the behavioral patterns can be based from rich data, such as location, calls, SMS, and web site visits. See [8] and [13] for some work in this direction, called *implicit authentication*. This approach does not directly address malware on the phone, but reduces the risk of a

lost phone providing a gateway into the home network. One of the themes in this paper is to generalize the techniques of implicit authentication beyond humans to a collection of devices.

We remark also that improving authentication on mobile phones is critical because smartphones might become the de facto remote control device. We believe the phone will also be involved in various enrollment protocols in the HAN. For instance, in [12] the phone is the party that mediates the pairing between two devices.

### III. SYSTEM DESCRIPTION

We outline here one possible system architecture based on the above ideas. See Figure 1 for a representation in the home network. Appliances are enrolled into a star-shaped network centered at a Master Controller. Commands to each appliance go through and are vetted by the Master Controller. The Master Controller contains an Authentication Service to authenticate commands, and also a Pattern Analyzer to build data models and evaluate recent data based on the models.

Appliances periodically upload data to an authentication service. These include settings, sensor readings, status, etc. In the case of a device with a user interface, user-related data may be recorded as well. For instance, data such as location, call logs, SMS logs, and web sites visited (all suitably anonymized and obfuscated) might be collected for a cell phone. This data is stored in the Database and used by the Pattern Analyzer to generate a model for the data. For instance, a simple model might say that the air conditioner is usually on during the afternoon but off at night.

Note that the smart meter may also upload data to the Master Controller. In this way, patterns of household electric data may be incorporated into the Pattern Analyzer. As detailed in [11], these patterns can identify the use of even non-smart appliances, as well as reconstruct daily routines, including sleep habits of inhabitants.

When the user wishes to change a setting on a device, say, through an application on his phone, the application connects with the Master Controller with the usual password. Depending on the policy enforced for the device, the Master Controller may measure how well recent behavior on the phone matches historical behavior. If the behavior is very different, another credential may be requested. The policy may also request another credential if a change is requested that is unusual for that device, in conjunction with all other available data. For example, turning on the air conditioner in the middle of the night may be unusual, but especially if nobody is in the house.

We allow the possibility of a device directly interacting with another device, not going through the Master Controller. For instance, in Figure 1, a cell phone might be used to control a car without going through the Master Controller. In this case, the user has chosen to drop back to unaided

traditional measures, without the added assurance provided by the Master Controller. The user may not want to rely on the Master Controller being operational, for instance.

A novel aspect of the system is that the data uploaded to the Master Controller must be treated as sensitive. Knowledge of the data would be equivalent to knowing a security key in a traditional system, since the algorithms used by the Master Controller must be considered public. Hence, communication channels to the Master Controller should be secured, and we expect that data will not reside in the long term on individual devices, limiting the risk of compromises of individual devices. We recognize, however, that for some devices historical data facilitates daily operation of the device (for example, the recent calls on a cell phone).

#### A. Anomaly Detection

One key question is how well anomaly detection works for this kind of data. Experiments described in [13] indicate that user modeling on cell phones is promising. With training data of around one or two weeks and four types of data collected (web sites, call logs, SMS, location), one can set thresholds such that an adversary would be detected within about 10 phone usages, while the legitimate user is identified as illegitimate about every 130 usages. It is reasonable to believe (but have no hard evidence) that most devices would have less variance and less richness in their data and so would be more predictable.

In [13], the approach is to treat each type of data collected separately. For each type of data, the software builds a model based on training data and then evaluates recent data against the model, deriving a score. See Figure 2. The scores for each data type are combined, assuming independence of each data type. To enable the discovery of the complex rules and interactions between HAN appliances, it would be necessary to detect and model the actual correlations between data types, going beyond the work in [13].




Figure 2. Scoring Data

### IV. SUMMARY AND DISCUSSION

In this position paper, we propose data-driven security management for cyber-physical systems. We argue that in the cyber-physical systems of the near future, such as HANs, better and more usable anomaly/attack detection can only be built if the collective information from all devices as

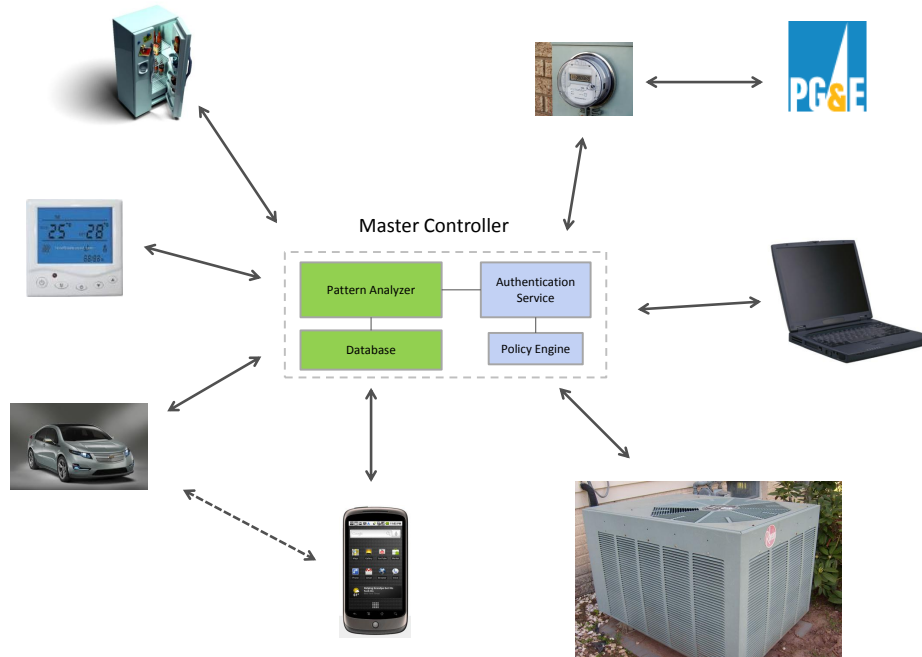


Figure 1. System Architecture

well as the behavior of their user and the inferred context are integrated into the decision making process. From the examples given in previous sections, we indicate how this collective high-level vision can detect anomalies for devices that operate within usual limits when considered individually, or how false alarms can be prevented with a collaborative and contextual view.

On the other hand, there are some obvious limitations to such a system. For example, one real advantage of a collaborative view over a range of smart devices is the ability to capture the behavioral patterns of a human user and make contextual decisions based on these patterns. However, humans do not always follow established behavioral patterns, and in the case of multiple users interacting with the same system, the system may be slow or ineffective in detecting attacks/anomalies due to the lack of one clear behavioral pattern. Having one master controller is another weakness of the system as it may be a single point of failure, but this weakness can be easily overcome by employing fault-tolerance practices such as replication.

Note that decision making in the proposed system is based on automated modeling which requires sufficient initial training data and also dynamically evolves with time to be more effective. However, the ongoing training process makes the system vulnerable at the beginning and can possibly be exploited by smart attackers to slowly evolve the system to an insecure state if it is not designed carefully. In the case of user authentication with behavioral data, one important

limitation is the reaction time of the system to a user change. In other words, it is not possible for the system to detect changes in the user behavior before certain number interactions between the system and an attacker. Thus, we only recommend this kind of implicit authentication as a second factor or as a low-security authentication mechanism.

One main motivation is usability. Since security is not the main objective when users interact with a system, they get irritated when they are bothered with security mechanisms (for example, requests for an additional credential) while trying to achieve their goal (for example, turning on the air-conditioner). With this system, the overall authentication can be made more usable by reducing the requests for additional credentials. False alarms by security systems are one of the biggest issues of current security systems, affecting both usability and security. Too many false alarms not only make the system less usable, but also seriously damages security (for example, nobody pays attention to car alarms). Hence, usability of our system is intimately tied to the rates of both false positives and false negatives of our anomaly detection algorithms; acceptable rates would need to be determined with user studies.

#### REFERENCES

- [1] Apps for Remote Control. On the Web at <http://www.apple.com/iphone/apps-for-everything/remote.html>.
- [2] Israeli Test on Worm Called Crucial in Iran Nuclear Delay. On the Web at <http://www.nytimes.com/2011/01/16/world/middleeast/16stuxnet.html>.

- [3] Malware Aimed at Iran Hit Five Sites, Report Says. On the Web at <http://www.nytimes.com/2011/02/13/science/13stuxnet.html>.
- [4] PG&E SmartAC Program. On the Web at <http://www.pge.com/myhome/saveenergymoney/energysavingprograms/smartac/>.
- [5] Spanish police nab suspected creator of webcam Trojan. On the Web at [http://www.computerworld.com/s/article/99034/Spanish\\_police\\_nab\\_suspected\\_creator\\_of\\_webcam\\_Trojan](http://www.computerworld.com/s/article/99034/Spanish_police_nab_suspected_creator_of_webcam_Trojan).
- [6] Stuxnet. On the Web at <http://en.wikipedia.org/wiki/Stuxnet>.
- [7] A. A. Cardenas, S. Amin, Z.-S. Lin, Y.-L. Huang, C.-Y. Huang, and S. Sastry. Attacks against process control systems: Risk assessment, detection, and response. In *ASIACCS*, 2011.
- [8] M. Jakobsson, E. Shi, P. Golle, and R. Chow. Implicit Authentication for Mobile Devices. In *HotSec '09: Proceedings of the 4th USENIX Workshop on Hot Topics in Security*, 2009.
- [9] Y. Mo and B. Sinopoli. Secure control against replay attacks. In *Proceedings of the 47th annual Allerton conference on Communication, control, and computing*, Allerton'09, pages 911–918, Piscataway, NJ, USA, 2009. IEEE Press.
- [10] C. Neuman. Challenges in Security for Cyber-Physical Systems. On the Web at <http://cimic.rutgers.edu/positionPapers/CPS-Neuman.pdf>.
- [11] E. Quinn. Privacy and the New Energy Infrastructure. On the Web at <http://www.townsend.com/Templates/media/files/media%20coverage/Elias%20Quinn%20-%20SmartGridPriv.pdf>.
- [12] N. Saxena, M. B. Uddin, and J. Voris. Universal device pairing using an auxiliary device. In *Proceedings of the 4th symposium on Usable privacy and security*, SOUPS '08, pages 56–67, New York, NY, USA, 2008. ACM.
- [13] E. Shi, Y. Niu, M. Jakobsson, and R. Chow. Implicit authentication through learning user behavior. In *ISC*, pages 99–113, 2010.