Continuous Authentication and Authorization for the Internet of Things

Muhammad Shahzad and Munindar P. Singh • *North Carolina State University*

How can users be authenticated and authorized continuously for the Internet of Things, when most small smart devices lack the conventional interfaces used for authentication (such as keyboards, mice, and touchscreens)? Here, the authors explore potential solutions along with a related case study.

We're venturing into the era of the Internet of Things (IoT). As computing devices become smaller, smarter, and ubiquitous, computing has begun to embed into our environments by attaching to physical objects or things. IoT is bringing computing both onto our bodies and into our daily surroundings. Examples of on-body computing devices include human activity trackers, smart watches, and semi-permanent insulin pumps. Examples of in-environment computing devices include intelligent thermostats, smart appliances, remotely controllable household equipment, and weather-based automated lawn irrigation systems.

Although IoT devices often have compute power close to those of conventional computing devices from a few years ago, one of the ways in which typical IoT devices differ is that they lack conventional user interfaces in the form of keyboards, mice, and touchscreens. Examples of such computing devices include the Fitbit activity tracker, sewable computing devices such as the Arduino Lilypad, and smart fabrics. A motivation for eliminating such user interfaces isn't so much to reduce the cost as that the conventional interfaces often aren't appropriate for the intended applications. For example, a smart fabric can have embedded antennas and communicate information about the person wearing the fabric to devices such as smartphones, but it wouldn't quite make sense to attach a touchscreen to a shirt or a keypad to a Fitbit.

This lack of a user interface gives rise to a fundamentally challenging question: How do we authenticate and authorize users for the IoT, where we lack conventional user interfaces? For example, one of the latest features of the Apple Watch is that if a user owns a (sufficiently new version of) Macbook Pro, an iPhone, and an Apple Watch, the user can set up the Macbook Pro to automatically unlock without entering a password. More specifically, as soon as the user opens the lid of her Macbook Pro, the laptop automatically unlocks if the following four conditions are satisfied:

- The user is wearing the Apple Watch.
- The Apple Watch is connected to the user's iPhone via Bluetooth.
- The watch is in close proximity to the Macbook Pro.
- Either the iPhone or the watch has been unlocked at least once since the user last put on the watch.

This convenient feature carries a security threat, however. Suppose an attacker, possibly posing as a friend, gets hold of the user's watch and has physical access to her computer. Such a scenario might occur if the two are in a lunch meeting and the user steps away from the table to pick up something from the buffet, but leaves behind her watch and computer. If the attacker wears that watch and the user happens to unlock her phone while away from the table — but within the Bluetooth range of the watch — the watch unlocks as well. Then the attacker can use the watch to unlock the user's Macbook Pro without having to guess her password.

Although this technology employs an Apple Watch, which does have a conventional user interface in the form of a touchscreen, in principle, it can be extended to any wearable device with a Bluetooth interface, such as a Fitbit. These kinds of examples highlight the present challenge: How can we continuously authenticate a person using a device without a conventional interface? Here, we consider various solutions, including a case study for a Wi-Fi-based human authentication system (Wi-Fi uses radio frequencies near 2.4 and 5 GHz).

Prospective Solutions

Because of the diversity of devices and applications, a universal solution to the problem of continuous authentication of users on devices without conventional interfaces might not exist. However, we can make progress by dividing IoT devices with which humans interact into two categories and studying solution directions for these categories separately. The first category consists of devices that maintain permanent physical contact with the user during usage, such as activity trackers, smart watches, and insulin pumps. The second category consists of devices that don't maintain permanent physical contact with humans, such as intelligent thermostats, occupancy sensors, and smart household appliances.

Authentication on Devices That Maintain Continuous Physical Contact

Devices that maintain contact with the user can support new forms of biometric authentication. Most of these devices fall into two categories. Devices of the first category either contain an inertial measurement unit (IMU), which is comprised of an accelerometer and a gyroscope, or can have an IMU embedded quite easily. Devices of the second category contain a photoplethysmogram (PPG) sensor, which is comprised of a few (often two or three) LEDs and a few (again, often two or three) light sensors. Both IMUs and PPG sensors can enable user authentication.

Specifically, using the IMU, we can develop authentication techniques that are based on the principle that users frequently move their limbs in unique patterns throughout the time they use the device. An example of a well-known trait that differs across users is gait. If we can extract the patterns in the output of an IMU sensor due to the user's unique gait, we can use simple machine learning techniques to learn these patterns and apply them to continuously authenticate the user based on gait. If such a technique was developed and put into practice, a device of the first category (worn by the user) could monitor the user continuously and frequently authenticate the user's legitimacy before allowing the user to perform appropriate operations. For example, the watch in the Macbook Pro setting wouldn't authenticate the attacker, which would prevent the Macbook Pro from being spuriously unlocked. Several behavioral biometrics solutions have been proposed that employ the IMU to authenticate users.[1,2]

Similarly, the PPG sensor provides an opportunity to study the PPG signal for unique patterns in blood flow rhythm. Researchers have shown that due to slight variations in every human's heartbeat rhythm, echocardiogram (ECG) signals contain small information — but this is enough to indicate what's unique to an individual user.[3,4] Consequently, just by using the ECG signal, we can design user schemes to authenticate users. Although several ECG-based user authentication systems have been proposed, this technology has yet to achieve sufficient effectiveness to see widespread deployment. Because the PPG signal is generated based on the amount of blood flow in the user's veins, which depends on how the user's heart pumps blood, the PPG signal could contain enough information to enable user authentication. Using the PPG signal is particularly challenging, however, because this signal is sensitive to the motion of a person's limbs: that is, the PPG measurement depends on the person's speed of movement. Fortunately, most devices these days that come with a PPG sensor also come with an IMU. Therefore, we can potentially use information from the IMU sensors to measure the amount of motion of the limb and combine the two signals — or correct the measurement of the PPG sensor — to enable authentication.

Authentication on Devices That Don't Maintain Permanent Physical Contact

A more challenging problem is to design an authentication scheme that can identify users for devices that don't maintain permanent contact with users. Such devices include those embedded into our environment. For example, consider an application that integrates a user's calendar with the user's home lighting and is controlled with speech. Whenever a calendar generates a notification for the user, the user's location is automatically determined through proximity or movement sensors, and lights of the appropriate room are flashed to alert the user. Suppose this application is enabled or disabled through voice commands from a specific user. Then, an attacker could replay previously recorded voice commands of the original user. That is, voice-based authentication is insufficient. We need effective methods to continuously and unintrusively authenticate users without, for example, requiring the user to wear sensors such as IMUs.

A potential approach for developing such an authentication system is to employ pervasive modalities such as radio frequency (RF) signals, ambient light, and sound, which are present all around us. The intuition behind RF-based authentication is

that the wireless channel metrics — such as channel state information (CSI) and received signal strength (RSS) — change based on a user's presence and movement. The patterns of change in these metrics depend on the way the user moves. Because different users have different gaits, they produce different patterns of change in wireless channel metrics. An RF-based user authentication system can apply machine learning techniques to associate each user with his or her patterns of change and identify the user at runtime based on the learned associations.

Specifically, with a human walking around, because a human is mostly made of water, the Wi-Fi signal reflected by the human body generates unique, although small, variations in CSI measurements on the receiver due to the well-known multipath effect of wireless signals. These variations in CSI enable signal processing techniques to obtain gait information such as walking speed, gait cycle time, footstep length, and movement speeds of legs and torso. Because each human has a unique gait, the gait patterns that the Wi-Fi receiver obtains can be used to recognize a walking human subject.

Similarly, the intuition behind light-based authentication is that as a user moves in an indoor environment, the amount of light he or she reflects and blocks depends on his or her patterns of movement. As different users have different gaits, the patterns of change in intensity of light, as measured by light sensors deployed on the floor, are also different. A light-based user authentication system can learn these patterns and apply them to identify users at runtime. A similar intuition holds for audio-based user authentication.

Authorization in the IoT

So far we've talked of authentication because it provides concrete use cases. But authentication by itself is usually meaningless. The point of authentication is to provide a basis for making a decision — about which resources to provide access to which person for what purpose and when. In broad terms, we aren't so much interested in seeing who specifically is around but what information or device to share with that person under what circumstances.

Consider a situation where a user is wearing various health-monitoring devices, including an ECG reader. Ordinarily, the data gathered by such devices would be confidential. Now, suppose the user is having a serious medical problem, such as a heart attack. In such a case, it might be acceptable behavior for the relevant software application to reveal the data from his ECG reader as well as data about recent physical activity to anyone who is present nearby and might be willing and able to help. But if the user is already in a hospital, then perhaps the application doesn't need to be quite so forthcoming in revealing its user's information to strangers. That is, here the decision changes from a focus on authentication of the counterparty to determining some attributes of the information resource and of the current context. Indeed, modern approaches[5] to authorization express policies in terms of attributes of principals, resources, and contexts instead of specific identifiers or roles. The IoT can readily accommodate such approaches by accumulating a rich variety of attribute values from the available devices.

In this example, the decision about whether to share some data is based upon the data values themselves (for example, sharing with anyone if the ECG indicates distress), but in general the decision might be based on the totality of available information. In particular, we can have situations where an application grants access to wearable devices based on environmental devices or the other way around. For example, in a home eldercare setting, if the environmental sensors (whether Wi-Fi or light-based) indicate a lack of movement for a prolonged period, the eldercare application might disclose data from wearable devices that capture the resident elder's health condition. Conversely, the data from a wearable sensor being anomalous might lead the application to verify whether a qualified caregiver was currently in the same room as the resident.

Case Study

To validate the effectiveness of such an approach to leverage variations in pervasive modalities, working with colleagues, we developed a Wi-Fi-based human authentication system, called WifiU, which recognizes users based on their gait.[6] We developed WifiU entirely using COTS Wi-Fi devices to capture fine-grained gait patterns. WiFiU consists of two Wi-Fi devices: one for continuously sending signals, which can be a router, and one for continuously receiving signals, which can be a laptop. In WifiU, the receiver measures channel state information (CSI) of each received Wi-Fi frame. Fundamentally, WifiU recognizes humans based on who they are, because WifiU extracts unique biometrics information from Wi-Fi signals and uses it to perform human authentication.

Compared with traditional gait-recognition systems, which use cameras, floor sensors, or wearable sensors to capture gait information, WifiU is easier to deploy and has better coverage. From the deployment perspective, WifiU doesn't require any special hardware (such as floor sensors) and doesn't require the human subject to wear any hardware (such as an IMU). Wi-Fi devices are ubiquitous and most homes and offices are covered by Wi-Fi signals. The hardware that we experimented with — namely a Net-Gear JR6100 Wi-Fi router and Think-Pad X200 laptop (with an Intel 5300 WiFi NIC) — required no modifications.

Furthermore, unlike cameras, WifiU doesn't require lighting and works in the dark just as well as in bright light.

In designing WifiU, we faced many technical challenges. For example, it's nontrivial to profile gait patterns using CSI dynamics. Extracting gait information from CSI signals is difficult, because the signal reflections of different body parts are mixed together in the CSI waveform. As different human body parts move at different speeds while walking, the radio signal reflections from different body parts have different frequencies. To separate the radio signal reflections from the different body parts, we convert CSI waveforms (of two dimensions: time and amplitude) into spectrograms in the time-frequency domain (of three dimensions: time, frequency, and amplitude). We apply spectrogram enhancement techniques to reduce signal noise. The resulting spectrograms yield detailed human gait information similar to those generated by Doppler radars.

We conducted experiments on WiFiU using our gait database that contains more than 2,800 gait instances collected from 50 human subjects walking in a typical laboratory with an area of 50 square meters (see Figure 1). We anonymized all collected data to protect participants' privacy. Over the 50 subjects, WifiU achieves recognition accuracies of 79.3, 89.5, and 93.0 percent for the Top-1, Top-2, and Top-3 candidates, respectively. (Here, Top-N means that one of the selected N persons is the person who truly generated that gait observation.)

With the current implementation using a single wireless link, WiFiU has three limitations. First, the distance between the human subject and the Wi-Fi devices is limited to six meters. To address this limitation in future work, we can deploy multiple Wi-Fi sender-receiver pairs in the target area. Second, the recognition accuracy is limited to 92.3 percent for Top-1 candidates. Whereas
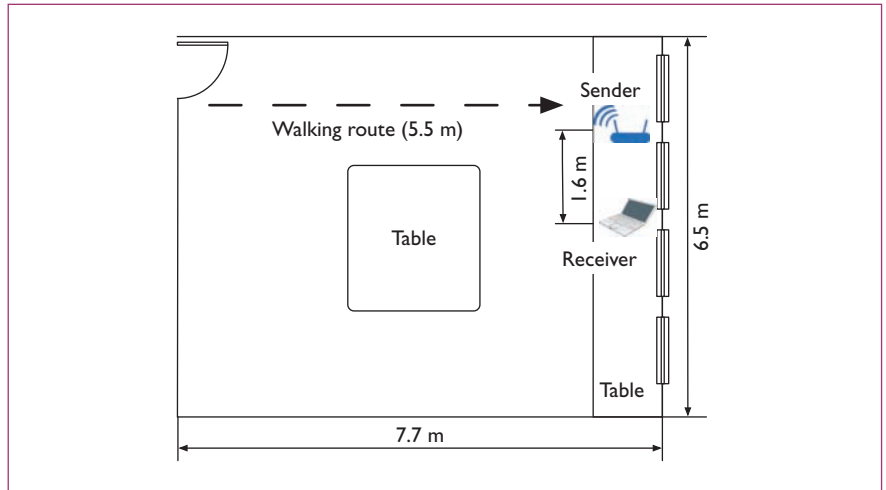

Figure 1. Data collection environment. Walking in an area of 50 square meters, we gathered more than 2,800 gait instances from 50 human subjects.

this accuracy might be acceptable for many personalized services such as adjusting room temperature and background music in smart buildings, it might not be high enough for settings that require high accuracy, such as accessing your email. Third, the number of walking human subjects is limited to one. In practice, the targets should walk along a given path one by one to ensure good recognition performance, as is the case for an airport security check. To address this limitation in future work, we plan to use multiple Wi-Fi receivers to separate signals of multiple humans using the differences in received signals at multiple receivers.

The IoT is in a nascent stage, but the arc of technology and the potential benefits it offers suggest that the IoT's presence will only increase. By placing people in information-rich environments, especially those that are natural and feel natural, the IoT exposes users to new security and privacy threats. It simultaneously demands stronger (that is, continuous) authentication and authorization and takes away conventional information modalities. Fortunately, the IoT provides new ways to address these challenges through innovative uses of technology. Therefore, the prospects of unintrusive authentication and authorization leading to context-sensitive policies are encouraging.

However, these unintrusive authentication technologies create potential privacy threats through the infrastructure in that an attacker who can obtain access to the infrastructure might apply these techniques without the user being aware of having been identified. For example, an attacker can potentially read Wi-Fi signals to identify victims without being detected. Consider a scenario where a burglar attempts to figure out who is at home by eavesdropping on the Wi-Fi signal emitted by the Wi-Fi router in the victim's house. As Wi-Fi signals can penetrate through obstacles such as furniture, wooden doors, and walls, the burglar needs to only passively measure the CSI of the signal outside the house without needing to decode the Wi-Fi packets' content. Therefore, it would be difficult for the victim to prevent certain breaches of privacy. Although avoiding privacy breach isn't the focus of this article, we hope this work highlights this privacy risk to the research community and encourages future work. A previous column[7] addresses the privacy risks in intelligent user interfaces, some of which might be exacerbated in combination with the IoT.

References

1. R. Mayrhofer and H. Gellersen, "Shake Well before Use: Authentication Based on Accelerometer Data," *Proc. Int'l Conf. Pervasive Computing*, 2007, pp. 144–161.

2. T.T. Ngo et al., "The Largest Inertial Sensor-Based Gait Database and Performance Evaluation of Gait-Based Personal Authentication," *Pattern Recognition*, vol. 47, no. 1, 2014, pp. 228–237.

3. S.I. Safie, J.J. Soraghan, and L. Petropoulakis, "Electrocardiogram (ECG) Biometric Authentication Using Pulse Active Ratio (PAR)," *IEEE Trans. Information Forensics and Security*, vol. 6, no. 4, 2011, pp. 1315–1322.

4. Z. Zhang et al., "ECG-Cryptography and Authentication in Body Area Networks," *IEEE Trans. Information Technology in Biomedicine*, vol. 16, no. 6, 2012, pp. 1070–1078.

5. D.R. Kuhn, E.J. Coyne, and T.R. Weil, "Adding Attributes to Role-Based Access Control," *Computer*, vol. 43, no. 6, 2010, pp. 79–81.

6. W. Wang, A.X. Liu, and M. Shahzad, "Gait Recognition Using WiFi Signals," *Proc. Int'l Conf. Pervasive and Ubiquitous Computing*, 2016, pp. 363–373.

7. C.J. Hazard and M.P. Singh, "Privacy Risks in Intelligent User Interfaces," *IEEE Internet Computing*, vol. 20, no. 6, 2016, pp. 57–61.

Muhammad Shahzad is an assistant professor in the Department of Computer Science at North Carolina State University. His research interests include the measurement, networking, and user interface aspects of the IoT as well as measurements, design, and modeling of computer networks. Shahzad has a PhD in computer science from Michigan State University. Contact him at mshahza@ncsu.edu.

Munindar P. Singh is a computer science professor at North Carolina State University. His research interests include the conception, engineering, and governance of sociotechnical systems as a way to tackle concerns such as security and privacy. Singh is a Fellow of IEEE and the American Association for Artificial Intelligence (AAAI), a former Editor in Chief of *IEEE Internet Computing*, and the current Editor in Chief of *ACM Transactions on Internet Technology*. Contact him at singh@ncsu.edu.