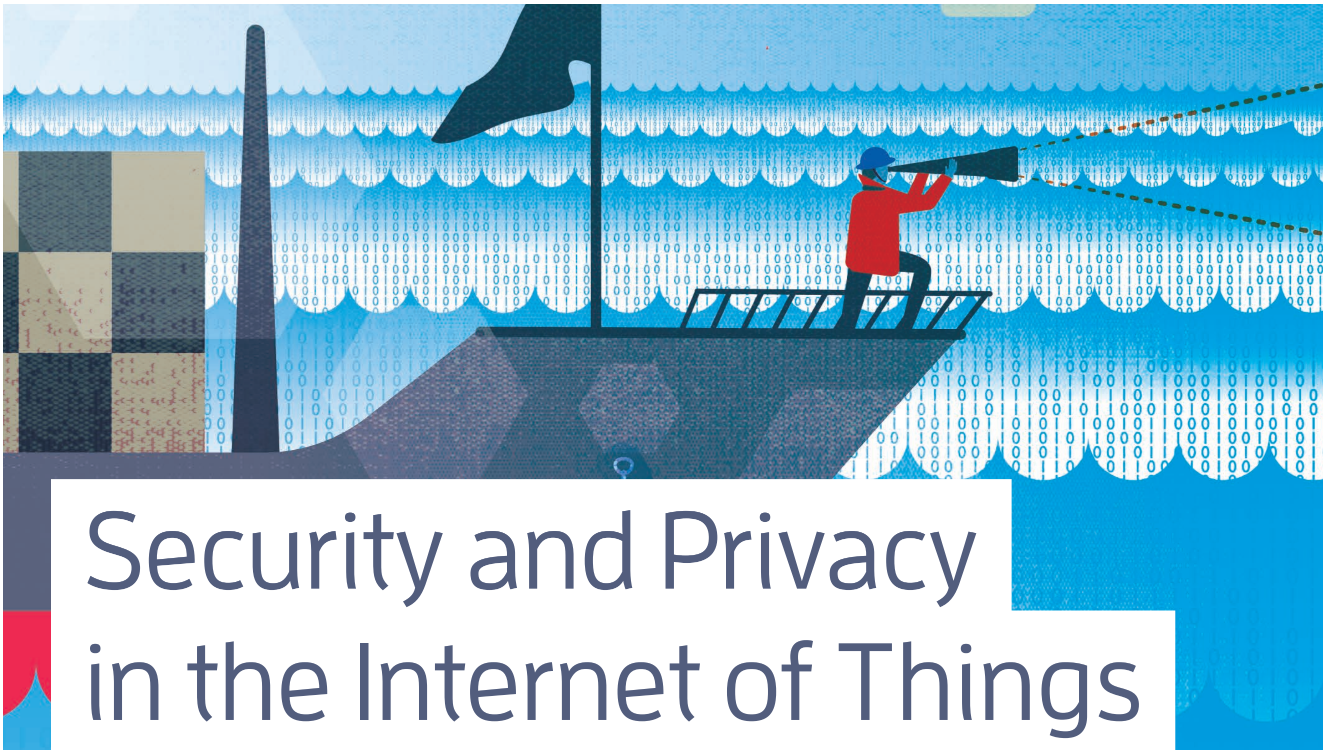

Security and Privacy in the Internet of Things

Kassem Fawaz, University of Wisconsin–Madison

Kang G. Shin, University of Michigan

Although the Internet of Things (IoT) computing paradigm is promising new applications, it introduces unprecedented security and privacy threats to individuals and their environments. The interactions within the IoT make it more challenging to protect users and their devices against these threats.

The Internet of Things (IoT) enables various new applications that will improve our quality of life. With tens of billions of connected devices, our fitness trackers, thermostats, door locks, heart pacemakers, cars, and appliances are becoming smarter and connected.[1] One significant obstacle to the broad adoption of the IoT is the associated security and privacy concerns.[2] Because these devices touch different aspects of our lives, they also bring unprecedented threats to users and their environments.

The IoT paradigm, however, brings along an attractive feature: the ability of users to interact with IoT devices, which enables them to issue commands and access information from the devices. In general, there are two user interaction surfaces in the IoT: indirect device-to-device (D2D) and direct human-to-device (H2D). The first relies on dedicated wireless network protocols, such as Bluetooth Low Energy (BLE), Zigbee, and ANT, to deliver user–device communication through a gateway. The second takes place through different user input (UI) mechanisms, ranging from traditional keypads and touchscreens to the most recent voice-based controls. These interaction surfaces are the perimeter that the attacker first breaches to inflict damage in the user's IoT environment. By exploiting the vulnerabilities of interaction surfaces, an attacker can gain unauthorized access to the user's IoT devices. Such unauthorized

Digital Object Identifier 10.1109/MC.2018.2888765
Date of publication: 16 April 2019

access can lead to an array of security and privacy threats.

PHYSICAL TRACKING

Wearable IoT devices, such as fitness trackers and medical devices, use wireless protocols to communicate with their gateways (i.e., access points or smartphones). By analyzing its wireless transmissions, a curious adversary can associate an IoT device with a unique identifier to continuously track it. The adversary can also deploy long-range scanners to track the mobility of users from their wearable devices over large areas. For example, researchers have been able to sense (short-range) Bluetooth devices more than a half-mile away.[3] The physical tracking of individuals has undesirable consequences, such as identifying their places of significance (home, work, ethnic stores, churches, or temples) and monitoring their behavior (leaving home or visits to hotels or hospitals).[4]

PERSONAL PROFILING

IoT devices collect sensory and usage data and communicate these data to their owners. A third-party entity with indirect access to the data, such as a wireless network sniffer, can draw sensitive inferences about the users and their behaviors. For example, network sniffers can identify the presence of the user's sensitive devices, such as a diabetic individual wearing a connected glucose monitor. Also, smart home devices leak the user's sleeping patterns, mobility, and activity through encrypted network traffic.[5]

UNAUTHORIZED CONTROL

The IoT paradigm brought users the ability to access their devices from a distance. Because many connected IoT devices are insecure, attackers can remotely access and control a user's devices via its wireless or voice interfaces. The repercussions to users vary from mere inconvenience to physical harm. For example, the unauthorized control of a smart bulb might pose an inconvenience to users, but an attacker controlling home-security systems and smart medical devices can pose grave dangers to their owners/bearers.

Optimally, secure and private access control must be part of the interaction protocols and, consequently, be built into the IoT device. A device must also be properly maintained throughout its lifetime by promptly incorporating interaction protocol updates. Nevertheless, with the thousands of manufacturers and developers, the IoT ecosystem is diverse, heterogeneous, and fragmented.[6] It is very challenging to ensure that all devices are properly secured at deployment time, let alone being maintained postproduction.[6]

Keeping the millions of deployed devices up to date requires patching by securely pushing firmware updates. Patch management is the leading security challenge in the emerging IoT[7,8] for many reasons. First, manufacturers might lack the ability to apply over-the-air updates[9] for some deployed IoT devices because they are neither programmable nor equipped with an Internet connection. Second, customers might not receive news about an update or be able to apply it even if available. Third, companies do not have enough financial incentives to maintain the devices after deployment.[10] There is, therefore, a need for a new approach to providing access control to the IoT devices at and after deployment.

By *access control*, we refer to the user's ability to enforce which entities are allowed to interact with his or her IoT devices. Developing such mechanisms is challenging, both theoretically and practically, for the following reasons. First, they must offer a provable security and privacy guarantee to prevent unauthorized access to the user's IoT devices. Second, they must be practical to deploy, requiring as few changes as possible to any IoT device. Third, they must allow for backward compatibility by not changing any underlying interaction protocol. Finally, these mechanisms should be implemented only with commercial, off-the-shelf hardware to ensure broad adoption.

In this article, we meet these challenges by designing, implementing, and evaluating a framework to enable external access control for both D2D and H2D interactions. We focus on two representative technologies: BLE-based input for D2D interactions and voice-based input for H2D interactions. Specifically, we present BLE-Guardian,[11] which provides external privacy and security protection to devices equipped with BLE, and VAuth,[12] which offers continuous authentication for voice-enabled devices.

BACKGROUND AND MOTIVATION

IoT interaction surfaces

Current IoT deployments enable individuals to interact with an IoT device either directly or indirectly, as shown in Figure 1. Through direct H2D interaction, a user can use a multitude of UI methods to communicate with the device. These include legacy interfaces, such as touchscreens and keypads, and more recent ones, such as voice-based control. Furthermore, a user can interact with IoT device indirectly via a gateway. The user employs a UI method to interact with an application on the gateway, which, in turn, uses the appropriate D2D interaction protocol, such as BLE, Zigbee, Z-Wave, and ANT, to interact with the IoT device. For example, an individual can use the Fitbit app on a smartphone to communicate with the Fitbit wearable device. As noted previously, in this article, we focus on two representative technologies for D2D and H2D interactions, i.e., BLE and voice inputs.

D2D interactions. The BLE protocol is a short-range wireless communication scheme that serves low-power devices. BLE powers a multitude of devices, such as sensors, fitness trackers, smart appliances and toys, and physical security devices, among others. Users interact with their BLE-powered devices through a BLE-equipped gateway, such as a smartphone.

Many manufacturers and developers are involved in producing BLE-equipped devices, with more than 75,000 unique BLE-equipped products available. The BLE standard defines how devices make their presence known to others via advertisement messages—e.g., wireless beacons containing information about the device, such as its address, name, and type. Also, it defines how more capable devices (e.g., smartphones) scan and connect to the BLE-equipped device.

These advertisements, however, can allow an unauthorized party to access or learn more about the BLE-equipped devices of a particular user or in a specific environment.[13] Apart from profiling the user's behavior and preferences, revealing the device's presence leads to critical privacy and security threats, especially for sensitive medical and home-security devices.

H2D interactions. With IoT devices becoming smaller and lacking traditional UIs, the IoT paradigm is offering voice as a newer interaction mechanism that does not require physical interaction. Voice is a desirable input interface in various scenarios where touch interfaces are dangerous or inappropriate, such as cooking, driving, and exercising. Many of our everyday appliances and devices are becoming voice activated, including alarm clocks, TVs, vacuum cleaners, home assistants, and thermostats.

Unfortunately, voice as an interaction mechanism brings serious security and privacy risks to users. Attackers can compromise voice-enabled devices by replaying user commands, impersonating the user's voice, and injecting hidden commands via audible or inaudible speech signals. Unauthorized access to voice-enabled IoT devices can have dangerous implications, including information theft, financial harm, and potential physical harm.

Threat model

Our framework aims to protect a user-facing IoT device at the access stage, i.e., before an adversary can establish an unauthorized connection to the device. In the case of BLE, the adversary (or unauthorized/unwanted device) can sniff the device's advertisements, issue scan requests, and attempt to connect to the device. In the voice-input case, the attacker can interfere with the audio channel to hijack the device voice interface and deceive it into executing malintended voice commands. The adversary can have different passive and active capabilities, from curious individuals scanning nearby devices (e.g., using a mobile app), to those with moderate technical knowledge using commercial sniffers, and all the way to sophisticated adversaries with software-defined radios.

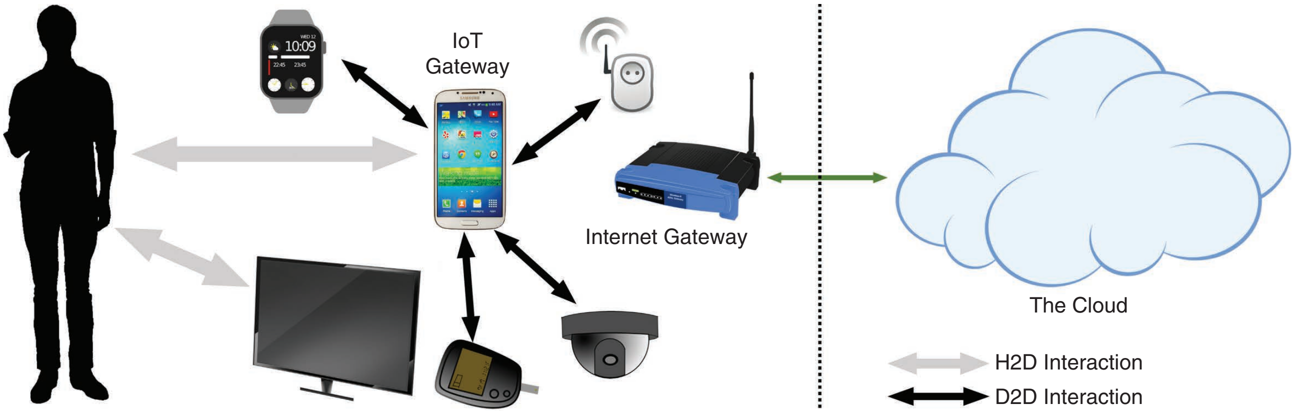

FIGURE 1. An IoT system model depicting D2D and H2D interactions. (Source: Creative Commons and Pixabay; used with permission.)

Our framework achieves security and privacy protection at the device level so that, if it authorizes a client to access the IoT device, all applications running on that device will have the same access privileges. Moreover, attacks—including jamming the channel entirely, masquerading as fake devices to trick users into connecting to them, or attacking the bonding process—are orthogonal to our framework. Finally, once our framework enables an authorized client to connect to the IoT device, it will not have any control over what follows later.

How are the security threats handled in practice?

The BLE standard offers security and privacy provisions to protect the owners/bearers of BLE devices, including whitelisting, address randomization, and direct advertisements. With whitelisting, only a set of authorized devices, those with a preexisting trust relationship, can connect to the BLE device.

This trust relationship can be built through pairing and bonding procedures between the two BLE devices. Randomizing the device address in each advertisement message prevents scanners from tracking the bearer of a BLE device over time. Direct advertisements prevent sniffers from inferring information about the device (e.g., name or type). In theory, these privacy and security provisions should protect the BLE device against tracking, profiling, and unauthorized access.

In our measurement study (presented later in this article), we discovered that, because of poor design and/or implementation, BLE advertisements leak an alarming amount of information, allowing the tracking, profiling, and fingerprinting of users. Almost all existing approaches addressing some of these threats rely on mechanisms that necessarily include changes to the protocol itself or to the way the BLE-equipped devices function.[14,15] These mechanisms are impractical to use in current and future IoT deployments.

Regarding voice interfaces, there is no standard on how to secure access to voice-enabled devices. The state-of-the-art solutions proposed to thwart unauthorized access to voice-enabled devices are based on voice biometric technologies. These mechanisms require training a signature of the user's voice in advance and then matching it in real time to authenticate the user's command. Voice biometrics, similar to other biometrics-based authentications, rely on static signatures. An adversary can overcome their protection by synthesizing speech commands that match the voice signature of the authorized user. In addition, existing IoT devices are resource constrained and may not be able to use sophisticated voice authentication mechanisms.

The general problem of securing IoT interactions

These issues have implications well beyond current IoT deployments or BLE and voice interaction technologies. The diversity and heterogeneity of IoT devices' manufacturers and developers will be a common feature of future IoT deployments. One could envision significant updates to interaction protocols every two to three years, but there is no guarantee that such updates will be disseminated to all existing devices at the same point in time. For instance, the BLE protocol underwent an overhaul in 2013 from version 4.1 to version 4.2, which carries enhanced security and privacy features. From our measurement study, we found that the vast majority of existing BLE devices still implement version 4.1 of the BLE protocol. The same applies to research proposals that require introducing major changes to IoT devices.

In rest of this article, we describe two systems (BLE-Guardian[11] and VAuth[12]) that we designed to provide external access control for BLE-equipped devices (representing D2D) and voice-enabled devices (representing H2D).

D2D INTERACTIONS

As a case study of D2D interactions, we consider BLE-equipped devices.[16] We present the results of a measurement study that highlights the security and privacy issues with BLE devices in the wild. Then, we introduce BLE-Guardian, which addresses those issues.

Measurement

We conducted a measurement campaign to investigate whether BLE devices implement the required security and privacy provisions of the BLE standard. We collected the advertisements of 214 unique types of BLE devices in the vicinity of 100 individuals during 2016.

Table 1 summarizes part of our measurement findings. First, almost all of the devices that we observed used *indirect advertisements*, the nonprivate type of advertisement. As is evident from Table 1, BLE devices typically advertise their names in the clear, which leads to identifying the user's type of device, such as Dexcom RX, a BLE-equipped glucose monitor. Also, some devices advertise unique identifiers as part of the device's name. Although the device might randomize its address, those unique identifiers could still lead to device tracking. We also found that many popular devices (e.g., Fitbit

TABLE 1. A sample of devices with revealing names.

Name	Type
LG LAS751M (27:5D)	Music streaming
JS00002074	Digital pen
ihere	Key finder
spacestation	Battery/storage extension
Jabra Pulse Smart	Smart bulb
Dexcom RX	Glucose monitor
Clover Printer 0467	Printer
Frances's Band ea:9d LE	Smart band
Gear Fit (60ED)	Activity tracker
Lyve Home-00228	Photo storage
Matthias-FUSE	Headset
Richelle's Band b2:6a LE	Smart band
vivosmart #3891203273	Activity tracker
KFDNX	Key fob
OTbeat	Heart-rate monitor
Thermos-4653	Smart thermos
POWERDRIVER-L10C3	Smart power inverter

products) used their advertising address consistently for weeks, i.e., without randomizing it. Apart from the information leaked about the device from advertisements, our lab experiments revealed that various devices accept connections directly from untrusted devices or use default personal identification numbers (PINs), or no PINs, to pair. Once connected, a client can access data from the device that leads to user tracking and profiling.

Our results suggest a significant discrepancy between the protocol specifications and the devices' operations. To bridge this gap, we present BLE-Guardian, which acts as an external access-control system to minimize the exposure of BLE devices.

BLE-Guardian

BLE-Guardian is an external protection system that prevents unauthorized entities from scanning and connecting to the user's BLE devices. Conceptually, BLE-Guardian consists of device hiding and access-control modules. The device-hiding module ensures that the BLE device is invisible to scanners in the area, and the access-control module ensures that only authorized clients are allowed to discover, scan, and connect to the BLE device.

Device hiding. BLE-Guardian uses an external hardware to hide a BLE device from adversaries by jamming its advertisements. A BLE device is supposed to advertise its presence periodically, with the period (advertising interval) being a preset value between 20 ms and 10.24 s. A BLE device sleeps for the period of the advertising interval and for an additional random time between 0 and 10 ms before advertising. After advertising, the device waits 10 ms for incoming connections before sleeping again. The additional random delay (after sleeping for the advertising interval) serves to reduce the probability of advertisements from different devices colliding.

BLE-Guardian overcomes two challenges to jam a BLE device. First, it has to avoid completely jamming the advertisement channels so as not to harm innocuous BLE devices. Second, BLE-Guardian must jam the BLE device exactly when it is advertising; missing the advertisement will leak the device's presence to potentially unauthorized entities. BLE-Guardian overcomes both challenges by predicting the BLE device's next advertisement. Through a brief learning phase, BLE-Guardian estimates the advertising period of the BLE device. Using this estimate, BLE-Guardian focuses only on the 10-ms interval in which the BLE device is expected to advertise. During that interval, it keeps sensing the channel until it detects a transmission, after which it immediately jams the channel for 10 ms to cover the advertisement and the subsequent listening period. BLE-Guardian then sleeps until the BLE device's next expected advertising interval.

Access control. Because hiding the device breaks the BLE protocol, BLE-Guardian packs an access-control protocol to allow a legitimate client to connect to the BLE device. Typically, the clients, such as smartphones, laptops, and gateways, are more powerful, are programmable, and have dual Bluetooth radios. BLE-Guardian utilizes Bluetooth Classic as an out-of-band channel to authorize the connections of legitimate clients to the BLE device. When an authorized client is in the vicinity of the BLE device, BLE-Guardian engages the following protocol:

1. Communicate over the out-of-band channel a set of random connection parameters to be used in the next connection request.
2. Lift the jamming immediately after the BLE device finishes advertising (after 350 μs). At this point, the BLE device should be listening for incoming connections.
3. Advertise on behalf of the BLE device with an advertisement message that includes less information about the device.
4. Monitor the medium for unauthorized connection requests, i.e., those with connection parameters not matching the precommunicated set to the authorized clients.
5. Jam the unauthorized connections, and alert the user about possible attacks to its BLE devices.

Evaluation. We implemented BLE-Guardian using an external Ubertooth radio that connects to a smartphone. Ubertooth is an off-the-shelf programmable radio that allows wireless transmission and reception at each Bluetooth channel. We also implemented a mobile app to control the settings of BLE-Guardian.

We conducted several evaluations of BLE-Guardian to assess its effectiveness and overhead on the advertisement channels. We used BCM20702A0 and Nordic nRF51822 chips to emulate the presence of up to 10 BLE devices to be protected (with varying advertising intervals). Figure 2 shows the impact of BLE-Guardian on an innocuous BLE device advertising at four different advertising intervals. As is evident from the figure, BLE-Guardian has little impact on the innocuous device while it is protecting up to six BLE devices. As BLE-Guardian protects more devices, the impact on the innocuous device will increase, jamming up to 50% of its advertisements. When advertising at the highest frequency possible (advertising interval = 20 ms), the innocuous device will experience the highest overhead. This overhead will have little impact on the user experience, because jamming 50% of the device's advertisements effectively doubles its advertising interval to 40 ms, which is still acceptable. Moreover, our measurement study revealed that real-world devices advertise with lower frequency. BLE-Guardian does not affect BLE data exchange because it takes place on different channels. The rest of our evaluation revealed that BLE-Guardian is effective in protecting the privacy of a BLE device and has limited energy overhead.

Discussion. BLE-Guardian performs its functionality using a radio that provides the basic functionalities of reception and transmission over the BLE advertisement channels. It requires no modification of the BLE device. BLE-Guardian is effective against various passive and active attackers. It prevents a single-antenna attacker from receiving the device's advertising, thereby preventing an array of tracking and profiling privacy threats. It also prevents an active attacker from issuing fraudulent connection requests to the BLE device. Nevertheless, BLE-Guardian is less effective for a multiantenna attacker that can extract the advertisement signal as well as for a high-power attacker that can overwhelm the BLE device. In the latter case, BLE-Guardian can detect the existence of the attacker and alert the user.

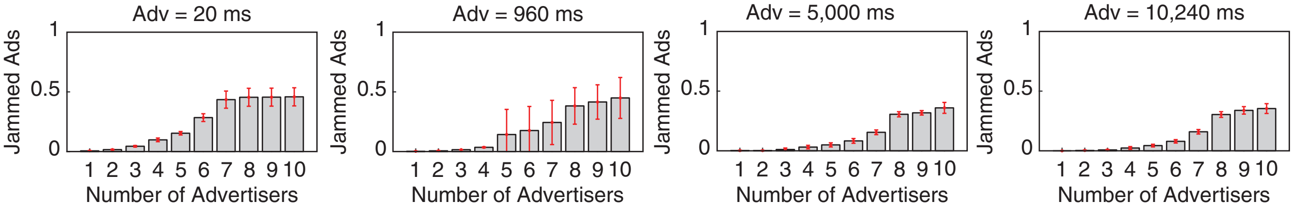

FIGURE 2. The effect of BLE-Guardian on innocuous devices in an environment. Each bar represents the portion of the jammed advertisements, averaged over the number of devices. The whiskers indicate the standard deviations. Ads: advertisements; Adv: advertising interval.

H2D INTERACTIONS

For H2D interactions, we present VAuth as a new method to provide continuous authentication for voice-enabled devices and assistants.

VAuth

VAuth relies on 1) a wearable component that captures the on-body vibrations of the speaker issuing a command and 2) a device component that matches the on-body vibrations with the speech received by the microphone of the voice-enabled device. Voice is a pressure wave that travels through the human body and over the air. The microphone of the voice assistant captures the voice command off the air, and VAuth's wearable component captures the vibrations resulting from the voice wave traveling through the speaker's body. Because both the vibration and speech signals are products of the same speech process (the user issuing a command), they should match in the temporal domain. VAuth leverages this observation to match both signals to decide whether or not to release the voice command to the voice assistant. The device component communicates with the wearable component over Bluetooth Classic, which provides a secure medium to authenticate the wearable and ensure data integrity. The cornerstone of VAuth is the matching algorithm, which decides whether the microphone and vibration signals match.

Matching algorithm. The matching algorithm takes as inputs two signals:

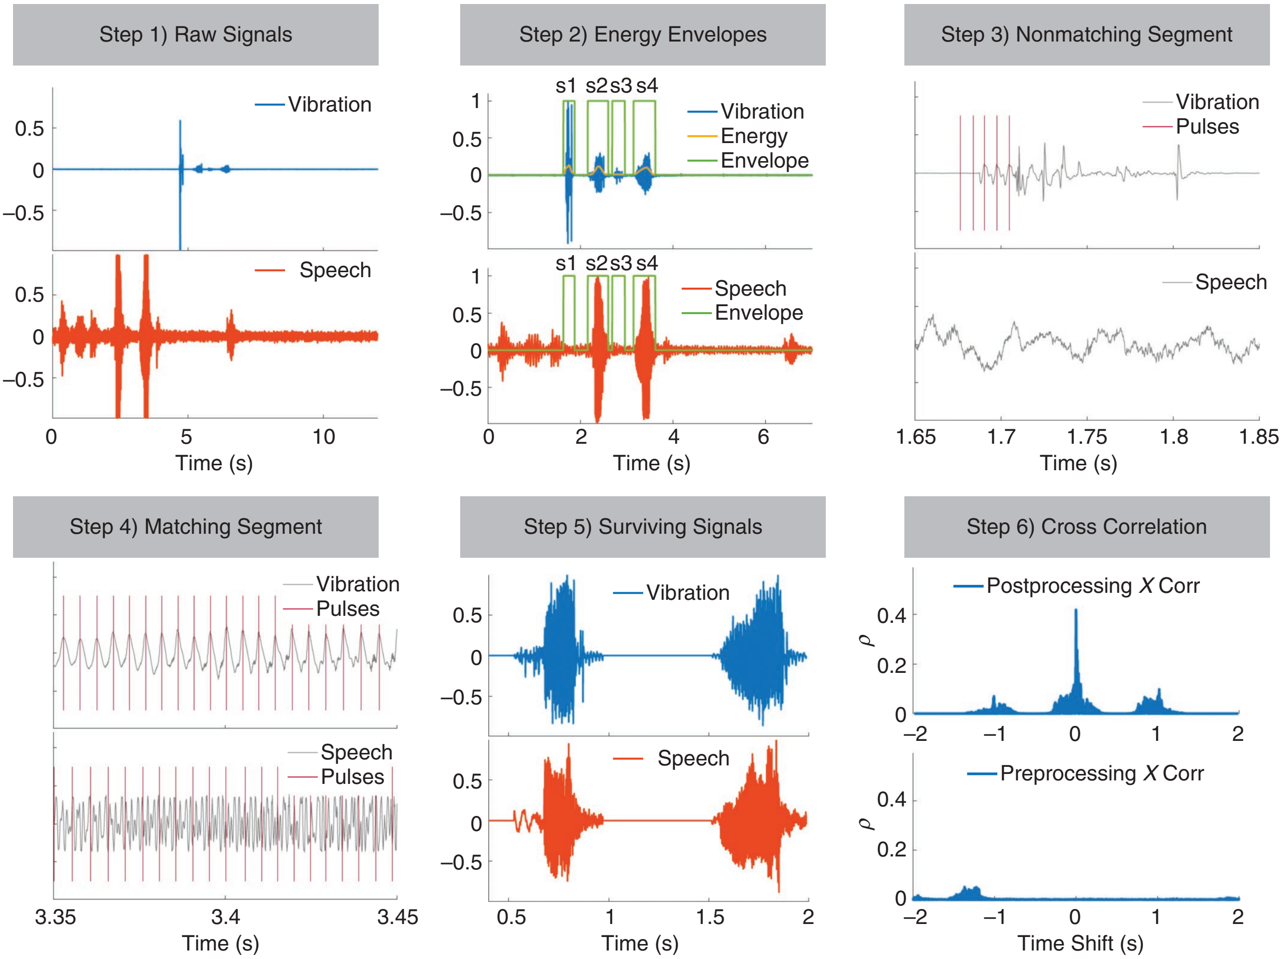

FIGURE 3. The matching algorithm of VAuth showing the progression from the raw signals (step 1) to the decision making (steps 5 and 6). In this case, the signals match as is evident from the top plot of step 6.

the vibration and speech (step 1 of Figure 3). First, VAuth identifies the energy envelope of the vibration signal, which corresponds to time intervals where the signal energy exceeds the average noise. The energy envelope denotes the instances where VAuth's wearable component recorded vibrations that could result from speech. Then, the matching algorithm overlays the energy envelope of the vibration signal on top of the speech signal (step 2 of Figure 3). This step of the matching algorithm provides the first security property of VAuth. It nullifies all parts of the speech signal that did not originate from the speaker (i.e., those that do not have a corresponding vibration signal).

Then, VAuth partitions both vibration and speech signals into a sequence of segments. A *segment* refers to a continuous interval of time during which the vibration signal has energy content. VAuth compares each segment from the vibration signal to its counterpart from the speech signal. In particular, it evaluates the percent difference in the sequence of the glottal periods of both segments. If the percent difference is high, the matching algorithm nullifies both the vibration and speech segments (steps 3 and 4 of Figure 3). This step provides the second security property of VAuth: it prevents an adversary from injecting speech into the medium when the user is actively communicating with the voice assistant.

Finally, the matching algorithm performs a cross-correlation operation on the surviving portions of the vibration and speech signals (steps 5 and 6 of Figure 3). The bottom plot from step 6 of Figure 3 shows the cross-correlation profile of the raw vibration and speech signals (before any processing). It is clear that the preprocessing cross correlation holds no information, indicating that both signals actually match, thus justifying the initial processing steps that result in a cleaner cross-correlation profile.

VAuth finally passes the cross-correlation profile (the top plot from step 6 of Figure 3) into a pretrained support vector machine (SVM) classifier that decides whether both signals match. The SVM classifier is user independent and trained offline. Only when both signals match does VAuth pass the surviving voice command to the voice assistant for additional processing.

Evaluation. We built the wearable component of VAuth using a wideband accelerometer and an off-the-shelf Bluetooth transmitter. We implemented the device component (including the matching component) as an Android patch to secure access to its Google Now smart assistant. We evaluated the performance of VAuth over a set of 18 users while issuing 30 English commands in six different scenarios. These scenarios corresponded to three different placements of VAuth on the speaker's body and two different movement conditions. Table 2 summarizes VAuth's true-positive and false-positive rates.

As is evident from the table, the 10th percentile of the true-positive rate is above 0.9 in most cases. This result indicates that VAuth correctly matched more than 90% of the voice commands to their vibration counterparts successfully. In one case, two of the test participants did not wear the VAuth device properly, which resulted in negative matches because it did not properly capture the vibration signal. The false-positive rate was lower than 0.0034 for 90% of the voice commands. This result indicates that, in a handful of cases, VAuth generated apositive match when the input signals do not match. Further analyzing the false-positive cases revealed that they correspond to nonintelligible commands, which result from the matching algorithm's nullifying nonmatching segments.

The rest of our evaluation found that these results are consistent across four other languages (Korean, Arabic, Persian, and Chinese). Furthermore, our VAuth prototype can match the commands with a delay of less than 0.5 s and can last for a week on a 500-mAh battery.

TABLE 2. VAuth true-positive and false-positive rates.

Placement	Movement	True positive, % in 10th percentile	False positive, % in 90th percentile
Next to the ear (as an earbud)	Stationary	100	0.26
	Jogging	91.33	0.33
On top of the nose (as eyeglasses)	Stationary	96.66	0.26
	Jogging	96.66	0.34
Back of the neck (as a necklace)	Stationary	81.66	0.15
	Jogging	93.33	0.15

ABOUT THE AUTHORS

KASSEM FAWAZ is an assistant professor in the Department of Electrical and Computer Engineering at the University of Wisconsin–Madison. His research interests include the security and privacy of the interactions between users and connected systems. Fawaz received a Ph.D. in computer science and engineering from the University of Michigan. He is a Member of the IEEE. Contact him at kfawaz@wisc.edu.

KANG G. SHIN is the Kevin and Nancy O'Connor Professor of Computer Science in the Department of Electrical Engineering and Computer Science at the University of Michigan, Ann Arbor. His research interests include quality of service-sensitive computing and networking as well as embedded real-time and cyberphysical systems. Shin received a Ph.D. from Cornell University. He is a Fellow of the IEEE. Contact him at kgshin@umich.edu.

Discussion. VAuth relies on the principle of security by possession; if an adversary gains access to the user's wearable component, it can issue unauthorized commands. The user can regain access by unpairing the compromised wearable, after which it will become unusable. The problem of authenticating wearables (and other devices) to an actual person is still open. The proposed solutions to this problem are prone to replay attacks from compromised signatures. More research is needed to properly and securely identify an authorized device user.

From a broader perspective, our framework offers general design directions for securing future IoT deployments. As more companies and manufacturers join the IoT sphere, fragmentation and heterogeneity issues are likely to remain and grow. In such a fragmented ecosystem, an IoT deployment is as secure as its weakest link. In this article, we presented an alternative paradigm in which the trust base shifts from the various manufacturers and developers to a framework that secures the interaction surfaces of the deployed IoT devices. Eventually, the administrators and owners of IoT devices can control their own security by deciding which entities are allowed access to devices in an IoT environment. The same design philosophy could extend to other D2D interaction protocols, such as Zigbee, or H2D interfaces, such as gesture control.

REFERENCES

1. A. Nordrum, "Popular Internet of Things forecast of 50 billion devices by 2020 is outdated," *IEEE Spect.*, 2016. [Online]. Available: https://spectrum.ieee.org/tech-talk/telecom/internet/popular-internet-of-things-forecast-of-50-billion-devices-by-2020-is-outdated
2. K. Rose, S. Eldridge, and L. Chapin, "The Internet of Things: An overview understanding the issues and challenges of a more connected world," Internet Society, Oct. 2015. Accessed on: June 1, 2016. [Online]. Available: https://www.internetsociety.org/sites/default/files/ISOC-IoT-Overview-20151014 0.pdf
3. S. Lester, "The emergence of Bluetooth low energy," Context Information Security. Accessed on: Nov. 28, 2017. [Online]. Available: https://www.contextis.com/blog/the-emergence-of-bluetooth-low-energy
4. A. J. Blumberg and P. Eckersley, "On locational privacy, and how to avoid losing it forever," Electronic Frontier Foundation, 2009. Accessed on: Nov. 28, 2017. [Online]. Available: https://www.eff.org/files/eff-locational-privacy.pdf
5. N. Apthorpe, D. Reisman, and N. Feamster, "A smart home is no castle: Privacy vulnerabilities of encrypted IoT traffic," in *Proc. First Workshop Data and Algorithmic Transparency*, 2016, pp. 1–6.
6. Federal Trade Commission, "Internet of Things, privacy & security in a connected world," Jan. 2015. [Online]. Available: https://www.ftc.gov/system/files/documents/reports/federal-trade-commission-staff-report-november-2013-workshop-entitled-internet- things-privacy/150127iotrpt.pdf
7. John Pescatore, "A SANS analyst survey: Securing the 'Internet of Things' survey," SANS, Jan. 2014. Accessed on: Jan. 18, 2016. [Online]. Available: https://www.sans.org/reading-room/whitepapers/analyst/securing-internet-things-survey-34785
8. Federal Bureau of Investigation, "Internet of Things poses oppor tunities for cyber crime," Sept. 2015. Accessed on: Jan. 18, 2016. [Online]. https://www.ic3.gov./media/2015/150910.aspx
9. European Commission, "Article 29 Data Protection Working Party. Opinion 8/2014 on the on recent developments on the Internet of Things," Justice and Consumers, Sept. 2014. Accessed on: Jan. 18, 2018. [Online].

Available: https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp223_en.pdf

10. B. Schneier, "The Internet of Things is wildly insecure and often unpatchable," *Wired*, Jan. 2014. Accessed on: Jan. 18, 2016. [Online]. https://www.wired.com/2014/01/theres-no-good-way-to-patch-the-internet-of-things-and-thats-a-huge-problem/
11. K. Fawaz, K.-H. Kim, and K. G. Shin, "Protecting privacy of BLE device users," in *Proc. 25th USENIX Security Symp. (USENIX Security 16)*, 2016, pp. 1205–1221.
12. H. Feng, K. Fawaz, and K. G. Shin, "Continuous authentication for voice assistants," in *Proc. 23rd Annu. Int. Conf. Mobile Computing and Networking (MobiCom '17)*. New York, 2017, pp. 343–355. [Online]. Available: http://doi.acm.org/10.1145/3117811.3117823
13. S. Lester, "The emergence of Bluetooth Low Energy," Context Information Security, May 2015. [Online]. Available: http://www.contextis.com/resources/blog/emergence-bluetooth-low-energy/
14. A. Leonard, "Wearable honeypot: A major qualifying project report," Worcester Polytechnic Inst., MA, Apr. 30, 2015. [Online]. Available: https://web.wpi.edu/Pubs/E-project/Available/E-project-042915-123749/unrestricted/Wearable_Honeypot.pdf
15. P. Wang, "Bluetooth Low Energy-privacy enhancement for advertisement," DiVA Portal, 2014. [Online]. Available: http://www.diva-portal.org/smash/get/diva2:750267/FULLTEXT01.pdf
16. Bluetooth SIG, "Specification of the Bluetooth system version 4.2," Dec. 2014. [Online]. Available: https://www.bluetooth.org/en-us/specification/adopted-specifications

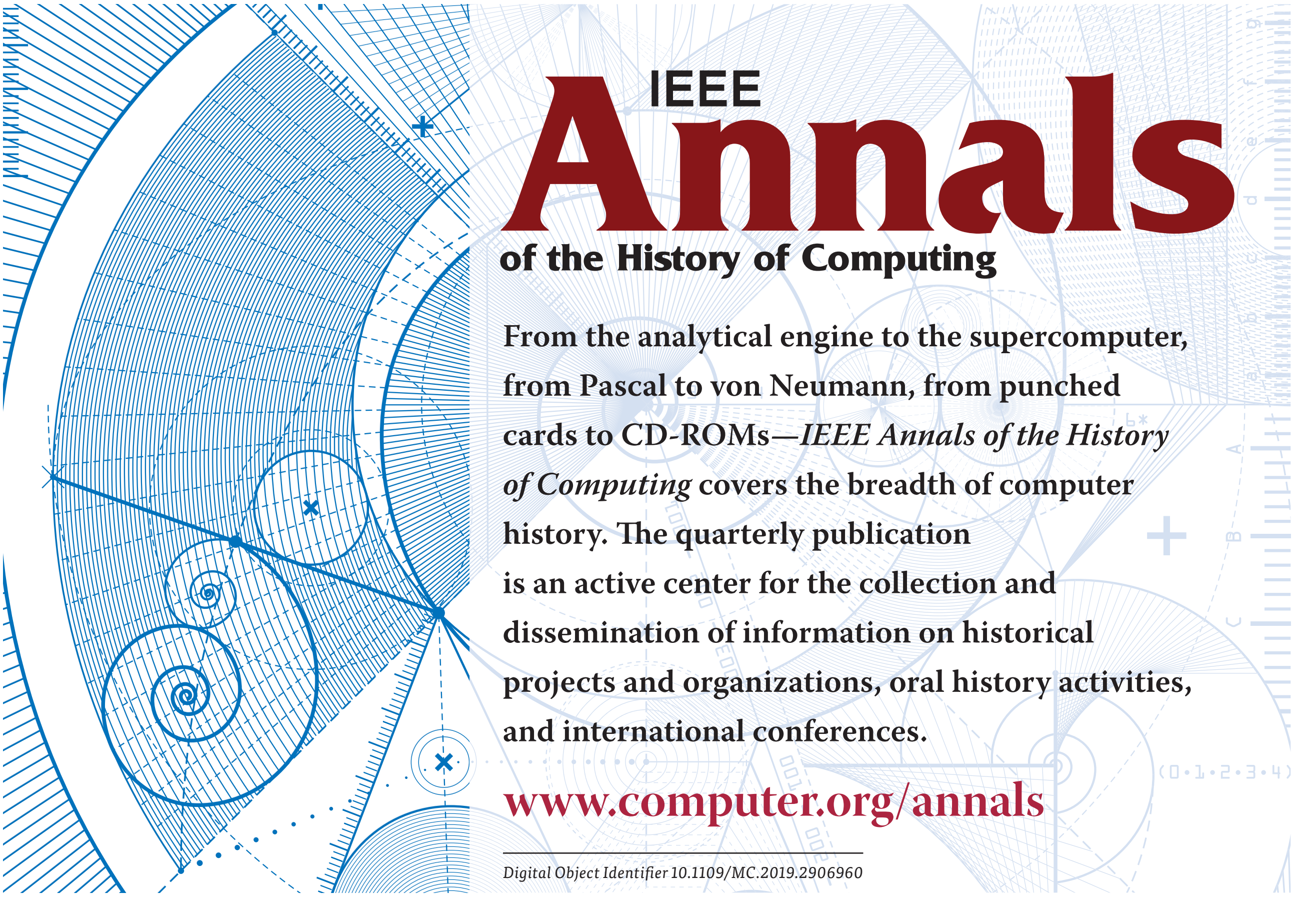

Digital Object Identifier 10.1109/MC.2019.2906960