

# IoT NOTARY: Attestable Sensor Data Capture in IoT Environments\*

NISHA PANWAR<sup>1</sup>, SHANTANU SHARMA<sup>2</sup>, GUOXI WANG<sup>3</sup>, SHARAD MEHROTRA<sup>3</sup>, NALINI VENKATASUBRAMANIAN<sup>3</sup>, MAMADOU H. DIALLO<sup>3</sup>, AND ARDALAN AMIRI SANI<sup>3</sup>, <sup>1</sup>Augusta University, USA. <sup>2</sup>New Jersey Institute of Technology, USA. <sup>3</sup>University of California, Irvine, USA.

Contemporary IoT environments, such as smart buildings, require end-users to trust data-capturing rules published by the systems. There are several reasons why such a trust is misplaced – IoT systems may violate the rules deliberately or IoT devices may transfer user data to a malicious third-party due to cyberattacks, leading to the loss of individuals’ privacy or service integrity. To address such concerns, we propose IoT NOTARY, a framework to ensure trust in IoT systems and applications. IoT NOTARY provides secure log sealing on live sensor data to produce a verifiable ‘proof-of-integrity,’ based on which a verifier can attest that captured sensor data adheres to the published data-capturing rules. IoT NOTARY is an integral part of TIPPERS, a smart space system that has been deployed at the University of California Irvine to provide various real-time location-based services on the campus. We present extensive experiments over realtime WiFi connectivity data to evaluate IoT NOTARY, and the results show that IoT NOTARY imposes nominal overheads. The secure logs only take 21% more storage, while users can verify their one day’s data in less than two seconds even using a resource-limited device.

CCS Concepts: • **Security and privacy** → **Security protocols**; *Mobile and wireless security*; *Domain-specific security and privacy architectures*; Social aspects of security and privacy;

Additional Key Words and Phrases: Internet of Things; smart homes; user privacy; channel and device activity; inference attacks.

## ACM Reference Format:

Nisha Panwar<sup>1</sup>, Shantanu Sharma<sup>2</sup>, Guoxi Wang<sup>3</sup>, Sharad Mehrotra<sup>3</sup>, Nalini Venkatasubramanian<sup>3</sup>, Mamadou H. Diallo<sup>3</sup>, and Ardalan Amiri Sani<sup>3</sup>. 2021. IoT NOTARY: Attestable Sensor Data Capture in IoT Environments. 1, 1 (August 2021), 24 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

Emerging IoT technologies [24, 30] promise to bring revolutionary changes to domains including healthcare, wearable devices, transportation, smart buildings/homes, smart infrastructure, and emergency response, among others. IoT interconnects (a potentially large number of) commodity devices (e.g., sensors, actuators, controllers) into an integrated network that empowers systems with new capabilities and/or brings in transformational improvements to the existing system [24, 28]. In IoT systems, sensors are used for fine-grained monitoring of the evolving state of the infrastructure and the environment. Our interest is in user-centric IoT spaces, wherein sensors of diverse types – cameras, cell phones, WiFi access points, beacons, bodyworn sensors, occupancy sensors, temperature sensors, light sensors, and acoustic sensors – are used to create awareness about subjects/end-users and their interactions among themselves and with the space. While fine-grained continuous monitoring offers numerous benefits, it also raises several privacy and security concerns (e.g., smoking habits, gender, and religious belief). To highlight the privacy concern, we first share our experience in building location-based services at UC Irvine using WiFi connectivity data.

### Use-Case: University WiFi Data Collection

In our ongoing project, entitled TIPPERS [31], we have developed a variety of location-based services based on a WiFi connectivity dataset. At UC Irvine, more than 2000 WiFi access-points and four WLAN controllers (managed by the university IT department) provide campus-wide wireless network coverage. Whenever a device connects to the campus WiFi network (through an access-point), the access-point generates a Simple Network Management Protocol

\*A preliminary version of this work was accepted in IEEE NCA 2019 under the title “IoT NOTARY: Sensor data attestation in smart environment.”

This version has been accepted in ACM Transactions on Internet Technology (TOIT), 2021.

Author’s address: Nisha Panwar<sup>1</sup>, Shantanu Sharma<sup>2</sup>, Guoxi Wang<sup>3</sup>, Sharad Mehrotra<sup>3</sup>, Nalini Venkatasubramanian<sup>3</sup>, Mamadou H. Diallo<sup>3</sup>, and Ardalan Amiri Sani<sup>3</sup>, <sup>1</sup>Augusta University, USA. <sup>2</sup>New Jersey Institute of Technology, USA. <sup>3</sup>University of California, Irvine, USA..

© 2021

XXXX-XXXX/2021/8-ART

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

(SNMP) trap for this association event. Each association event contains access-point-id,  $s_i$ , user device MAC address,  $d_j$ , and the time of the association,  $t_k$ . All SNMP traps  $\langle s_i, d_j, t_k \rangle$  are sent to access-point's controllers in realtime. The access-point controller anonymizes device MAC addresses (to preserve the privacy of users in the campus).

TIPPERS collects WiFi connectivity data from one of the controllers that manage 490 access-points and receives  $\langle s_i, d_j, t_k \rangle$  tuples for each connectivity event. However, before receiving any WiFi data, TIPPERS notifies all WiFi users about the data-capture rules by sending emails over a university mailing list. Subsequently, based on WiFi connectivity data  $\langle s_i, d_j, t_k \rangle$ , TIPPERS provides various realtime applications. Some of these services, e.g., computing occupancy levels of (regions in) buildings in the form of a live heatmap, require only anonymous data. Other services, e.g., providing location information (within buildings) or contextualized messaging (to provide messages to a user when he/she is in the vicinity of the desired location), require the user's original disambiguated data. To date, over one hundred users have registered into TIPPERS to utilize real-time services. A key requirement imposed by the university in sharing data with TIPPERS is that the system supports provable mechanisms to verify that individuals have been notified prior to their data (anonymized or not) being used for service provisioning. Also, an option for users to opt-out of sharing their WiFi connectivity data with TIPPERS must be supported. If users opt-out, the system must prove to the users that indeed their data was not shared with TIPPERS. TIPPERS uses immutable log-sealing to help all users to verify that the captured data is consistent with pre-notified data-capture rules.

Our experience in working with various groups in the campus is that (persistent) location data can be deemed quite sensitive by certain individuals with concerns about the spied upon by the administration or by others. Thus, mechanisms for notification of data-capture rules, secure log-sealing, and verification components made a sub-framework, entitled IoT NOTARY, which has become an integral part of TIPPERS. ■

Data-capture concerns in IoT environments are similar to that in mobile computing, where mobile applications may have continuous access to resident sensors on mobile devices. In the mobile setting, data-capture rules and permissions [22] are used to control data access, i.e., which applications have access to which data generated at the mobile device (e.g., location and contact list) for which purpose and in which context. At the abstract level, a data-capture rule informs the user about the nature of personally identifying information that an application captures, i.e., for what *purpose* and in what *context*. However, in IoT settings, the data-capture framework differs from that in the mobile settings, in two important ways:

- (1) Unlike the mobile setting, where applications can seek user's permission at the time of installation, in IoT settings, there are no obvious mechanisms/interfaces to seek users' preferences about the data being captured by sensors of the smart environment. Recent work [36] has begun to explore mechanisms using which environments can broadcast their data-capture rules to users and seek their explicit permissions.
- (2) Unlike the mobile setting, users cannot control sensors in IoT settings. While in mobile settings, a user can trust the device operating system not to violate the data-capture rules, in IoT settings, trust (in the environment controlling the sensors) may be misplaced. IoT systems may not be honest or may inadvertently capture sensor data, even if data-capture rules are not satisfied [4, 30].

We focus on the above-mentioned second scenario and determine ways to provide trustworthy sensing in an untrusted IoT environment. Thus, the users can verify their data captured by the IoT environment based on pre-notified data-capture rules. Particularly, we deal with three sub-problems, namely *secure notification* to the user about data-capture rules, *secure (sensor data) log-sealing* to retain *immutable* sensor data, as well as, data-capture rules, and *remote attestation* to verify the sensor data against pre-notified data-capture rules by a user, without being heavily involved in the attestation process.

## Our Contribution and Outline of the Paper

In this paper, we provide the following:

- A user-centric framework (§3.2) to ensure trustworthy data collection in untrusted IoT spaces, entitled IoT NOTARY that contains three entities (§2.1): (i) *infrastructure deployer* that installs sensors, (ii) a *service provider* (SP, e.g., TIPPERS) that establishes a list of data-capture rules that dictates the condition under which a sensor's data can/cannot be utilized to provide realtime services to the user, and (iii) *users* that use services provided by sensors, as well as, by SP, (if interested).
- Two models to inform the user about the data-capture rules (§4.1): notice-only model and notice-and-ACK model.

- A secure log-sealing mechanism (§4.2) implemented by secure hardware that cryptographically retains logs, data-capture rules, sensors' state, and contextual information to generate a *proof-of-integrity* in an immutable fashion.
- Optimized secure log-sealing mechanisms that regard the sensor state and user-device-id (§4.2.3 and §4.2.4), implemented by secure hardware to reduce the size of the secure log.
- A secure attestation mechanism (§4.3), mixed with SIGMA protocol [29], allowing a verifier (a user or a *non-mandatory* auditor) to securely attest the sealed logs as per the data-capture rules.
- Implementation results of IoT NOTARY on the university live WiFi data in §6.

## Advantages of IoT NOTARY

IoT NOTARY has the following distinct advantages:

- (1) **Privacy-preserving system.** IoT NOTARY is the first privacy-preserving system for the enforcement of IoT sensor data capturing policies and for the data attestation at the user to verify the presence/absence of their data against the data capturing policies. Moreover, during the policies' enforcement and attestation, any entity involved in IoT NOTARY cannot learn anything about other entities.
- (2) **Minimal overhead in proof generation.** IoT NOTARY does not incur computational and space overhead in secure log generation. In particular, IoT NOTARY takes 310ms for creating the secure log over 37K rows and 21% more space to store the log in a secure manner as compared to store the log in cleartext.
- (3) **Scalable.** IoT NOTARY is a highly scalable system in terms of data verification. In particular, a resource-constrained user (with a machine of 1-core 1GB RAM) can verify the presence or absence of their data over the data collected in a single day in most 30 seconds.
- (4) **A deployed system.** At the University of California, Irvine, IoT NOTARY has been deployed and is in use to capture WiFi connectivity events. The real-world testbed has been established to provide a real environment where IoT NOTARY's usability is tested and verified.

## 2 MODELING IOT DATA ATTESTATION

This section presents the entities involved in IoT NOTARY, threat model, and the desired security properties.

### 2.1 Entities

Our model has the following entities, as shown in Figure 1:

**Infrastructure Deployer (IFD).** IFD (which is the university IT department in our use-case; see §1) deploys and owns a network of  $p$  sensors devices (denoted by  $s_1, s_2, \dots, s_p$ ), which capture information related to users in a space. The sensor devices could be: (i) dedicated sensing devices, e.g., energy meters and occupancy detectors, or (ii) facility providing sensing devices, e.g., WiFi access-points and RFID readers. Our focus is on facility-providing sensing devices, especially WiFi access-points that also capture some user-related information in response to services. E.g., WiFi access-points capture the associated user-device-ids (MAC addresses), time of association, some other parameters (such as signal strength, signal-to-noise ratio); denoted by:  $\langle d_i, s_j, t_k, param \rangle$ , where  $d_i$  is the  $i^{th}$  user-device-id,  $s_j$  is the  $j^{th}$  sensor device,  $t_k$  is  $k^{th}$  time, and  $param$  is other parameters (we do not deal with  $param$  field and focus on only the first three fields). All sensor data is collected at a controller (server) owned by IFD. The controller may keep sensor data in cleartext or in encrypted form; however, it only sends encrypted sensor data to the service provider.

**Service Providers (SP).** SP (which is TIPPERS in our use-case; see §1) utilizes the sensor data of a given space to provide different *services*, e.g., monitoring a location and tracking a person. SP receives encrypted sensor data from the controller.

Data-capture rules. SP establishes data-capture rules (denoted by a list  $DC$  having different rules  $dc_1, dc_2, \dots, dc_q$ ). Data-capture rules are conditions on device-ids, time, and space. Each data-capture rule has an associated *validity* that indicates the time during which a rule is valid. Data-capture rules could be to capture user data by default (unless the user has explicitly opted out). Alternatively, default rules may be to opt-out, unless, users opt-in explicitly. Consider a default rule that individuals on the 6<sup>th</sup> floor of the building will be monitored from 9pm to 9am. Such a rule has an associated condition on the time and the id of the sensor used to generate the data. Now, consider a rule corresponding to a user with a device  $d_i$  opting-out of data capture based on the previously mentioned rule. Such an opt-out rule would have conditions on the user-id, as well as, on time and the sensor-id. For sensor data for




Fig. 1. Entities in IoT NOTARY.

which a default data-capture rule is opt-in, the captured data is forwarded to SP, if there do not exist any associated opt-out rules, whose associated conditions are satisfied by the sensor data. Likewise, for sensor data where the default is opt-out, the data is forwarded to SP only, if there exists an explicit opt-in condition. We refer to the sensor data to have a *sensor state* ( $s_i.state$  denotes the state of the sensor  $s_i$ ) of 1 (or active), if the data can be forwarded to SP; otherwise, 0 (or passive). In the remaining paper, unless explicitly noted, opt-out is considered as the default rule, for simplicity of discussion.

Whenever SP creates a new data-capture rule, SP must send a *notice message* to user devices about the current usage of sensor data (this phase is entitled *notification phase*). SP uses Intel Software Guard eXtension (SGX) [14], which works as a trusted agent of IFD, for securely storing sensor data corresponding to data-capture rules. SGX keeps all valid data-capture rules in the secure memory and only allows to keep such data that qualifies pre-notified valid data-capture rules; otherwise, it discards other sensor data. Further, SGX creates immutable and verifiable logs of the sensor data (this phase is entitled *log-sealing phase*). The assumption of secure hardware at a machine is rational with the emerging system architectures, e.g., Intel machines are equipped with SGX [1]. However, existing SGX architectures suffer from side-channel attacks, e.g., cache-line, branch shadow, page-fault attacks [44], which are outside the scope of this paper.

**Users.** Let  $d_1, d_2, \dots, d_m$  be  $m$  (user) devices carried by  $u_1, u_2, \dots, u_{m'}$  users, where  $m' \leq m$ . Using these devices, users enjoy services provided by SP. We define a term, entitled *user-associated data*. Let  $\langle d_i, s_j, t_k \rangle$  be a sensor reading. Let  $d_i$  be the  $i^{th}$  device-id owned by a user  $u_i$ . We refer to  $\langle d_i, s_j, t_k \rangle$  as user-associated data with the user  $u_i$ . Users worry about their privacy, since SP may capture user data without informing them, or in violation of their preference (e.g., when the opt-out was a default rule or when a user opted-out from an opt-in default). Users may also require SP to prove service integrity by storing all sensor data associated with the user (when users have opted-in into services), while minimally being involved in the attestation process and storing records at their sides (this phase is entitled *attestation phase*).

**Auditor.** An auditor is a *non-mandatory* trusted third-party that can (periodically) verify entire sensor data against data-capture rules. Note that a user can only verify his/her data, not the entire sensor data or sensor data related to other users, since it may reveal the privacy of other users.

## 2.2 Threat Model

We assume that SP and users may behave like adversaries. The adversarial SP may *store* sensor data without informing data-capture rules to the user. The adversarial SP may *tamper* with the sensor data by inserting, deleting, modifying, and truncating sensor readings and secured-logs in the database. By tampering with the sensor data, SP may *simulate* the sealing function over the sensor data to produce secured-logs that are identical to real secured-logs. Thus, the adversary may hinder the attestation process and make it impossible to detect any tampering with the sensor data by the verifier (that may be an auditor or a user). Further, as mentioned before that SP utilizes sensor data to provide services to the user. However, an adversarial SP may provide *false answers* in response to user queries. We assume that the adversarial SP cannot obtain the secret key of the enclave (by any means of side-channel attacks on SGX).

Since we assumed that sensors are trusted and cannot be spoofed, we do not need to consider a case when sensors would collude with SP to fabricate the logs.

An adversarial user may *repudiate* the reception of notice messages about data-capture rules. Further, an adversarial user may *impersonate* a real user, and then, may retrieve the sensor data and secured-log during the verification phase. This way an adversarial user may reveal the privacy of the users by observing sensor data. Further, a user may infer the identity of other users associated with sensor data by potentially launching *frequency-count attacks* (e.g., by determining which device-ids are prominent).

### 2.3 Security Properties

In the above-mentioned adversarial model, an adversary wishes to learn the (entire/partial) data about the user, without notifying or by mis-notifying about data-capture rules, such that the user/auditor cannot detect any inconsistency between data-capture rules and stored sensor data at SP. Hence, a secure attestation algorithm must make it detectable, if the adversary stores sensor data in violation of the data-capture rules notified to the user. To achieve a secure attestation algorithm, we need to satisfy the following properties:

**Authentication.** Authentication is required: (i) between SP and users, during notification phase; thus, the user can detect a rogue SP, as well as, SP can detect rogue users, and (ii) between SP and the verifier (auditor/user), before sending sensor data to the verifier to prevent any rogue verifier to obtain sensor data. Thus, authentication prevents threats such as impersonation and repudiation. Further, a periodic mutual authentication is required between IFD and SP, thereby discarding rogue sensor data by SP, as well as, preventing any rogue SP to obtain real sensor data.

**Immutability and non-identical outputs.** We need to maintain the immutability of notice messages, sensor data, and the sealing function. Note that if the adversary can alter notice messages after transmission, it can do anything with the sensor data, in which case, sensor data may be completely stored or deleted without respecting notice messages. Further, if the adversary can alter the sealing function, the adversary can generate a proof-of-integrity, as desired, which makes flawless attestation impossible. The output of the sealing function should not be identical for each sensor reading to prevent an adversary to forge the sealing function (and to prevent the execution of a frequency-count attack by the user). Thus, immutability and non-identical outputs properties prevent threats, e.g., inserting, deleting, modifying, and truncating the sensor data, as well as, simulating the sealing function.

**Minimality, non-refutability and privacy-preserving verification.** The verification method must find any misbehavior of SP, during storing sensor data inconsistent with pre-notified data-capture rules. However, if the verifiers wish to verify a subset of the sensor data, then they should not verify the entire sensor data. Thus, SP should send a minimal amount of sensor data to the verifier, enabling them to attest to what they wish to attest. Further, the verification method: (i) cannot be refuted by SP, and (ii) should not reveal any additional information to the user about all the other users during the verification process. These properties prevent SP to store only sensor data that is consistent with the data-capture rules notified to the user. Further, these properties preserve the privacy of other users during attestation and impose minimal work on the verifier.

### 2.4 Assumptions

This section presents assumptions, we made, as follows:

- (1) The sensor devices are assumed to be computationally inefficient to locally generate a verifiable log for the continuous data stream as per the data-capture rules.
- (2) Sensor devices are tamper-proof, and they cannot be replicated/spoofed (i.e., two devices cannot have an identical id). In short, we assume a correct identification of sensors, before accepting any sensor-generated data at the controller at IFD, and it ensures that no rogue sensor device can generate the data on behalf of an authentic sensor. Further, we assume that an adversary cannot deduce any information from the dataflow between a sensor and the controller. (Recall that in our setting the university IT department collects the entire sensor data from their owned and deployed sensors, before sending it to TIPPERS.)
- (3) We assume the existence of an authentication protocol between the controller and SP, so that SP receives sensor data only from authenticated and desired controller.
- (4) The communication channels between SP and users, as well as, between SP and auditor are insecure. Thus, our solution incorporates an authenticated key exchange based on SIGMA protocol (which protects the sender's identity). When the verifier's identity is proved, the cryptographically sealed logs are sent to the verifier.

- (5) By any side-channel attacks on SGX, one cannot tamper with SGX and retrieve the secret key of SGX. (Otherwise, the adversary can simulate the sealing process.)
- (6) A hash function  $H$  is known to the service provider, the auditor, and the user.

**Definition: Attestation Protocol.** Now, we define an attestation protocol that consists of the following components:

- $Setup()$ : Given a security parameter  $1^k$ ,  $Setup()$  produces a public key of the enclave ( $PK_E$ ) and a corresponding private key ( $PR_E$ ), used by the enclave to securely write sensor logs.
- $Sealing(PR_E, \langle d_i, s_j, s_j.state, t_k \rangle, dc_l)$ : Given the key  $PR_E$ ,  $Sealing()$  (which executes inside the enclave) produces a verifiable *proof-of-(log)-integrity* ( $\mathcal{PI}$ ) and *proof-of-integrity for user/service (query) verification* ( $\mathcal{PU}$ ), based on the received sensor readings and the current data-capture-rule,  $dc_l$ .
- $Verify(PK_E, \langle *, s_j, s_j.state, t_k, \mathcal{PI}, dc_l \rangle, Sealing(PR_E, \langle d_i, s_j, s_j.state, t_k \rangle, dc_l))$ : Given the public key  $PK_E$ , sensor data, proof, and data-capture rule,  $Verify()$  is executed at the verifier, where  $*$  denotes the presence/absence of a user-device-id, based on  $dc_l$ .  $Verify()$  produces 1, iff  $\mathcal{PI} = Verify(PK_E, \langle *, s_j, s_j.state, t_k, dc_l \rangle, Sealing(PR_E, \langle d_i, s_j, s_j.state, t_k \rangle, dc_l))$ ; otherwise, 0. Similarly,  $Verify()$  can attest  $\mathcal{PU}$ .

Note that the functions  $Sealing()$  and  $Verify()$  are known to the user, auditor, and SP. However, the private key  $PR_E$  is only known to the enclave.

### 3 IOT NOTARY: CHALLENGES AND APPROACH

This section provides challenges we faced during IoT NOTARY development and an overview of IoT NOTARY.

#### 3.1 Challenges and Solutions

We classify the problem of data attestation in IoT into three categories: (i) secure notification, (ii) log-sealing, and (iii) log verification. Before going into details, we provide the challenges that we faced during the development and the way we overcome those challenges (Figure 2), as follows:

**C1. Secured notification for data-capture rules.** The declaration of data-capture rules requires a reliable and secure notification scheme, thereby users can verify the sender of notice messages. Trivially, this can be done through a unique key establishment between each user and SP. However, this incurs a major overhead at SP, as well as, SP can send different messages to different users.

**Solution.** To address the above challenge, we develop two solutions: One is based on a secure notifier that delivers a cryptographically encrypted notice message, which is signed by the enclave, to all the users (see §4.1). Another solution uses an acknowledgment from the user and does not need any trusted notifier (see §4.1).

**C2. Tamper-proof cryptographically secured-log sealing.** The verification process depends on immutable sensor data that is stored at SP. A trivial way is to store the entire data using a strong encryption technique, mixed with access-pattern hiding mechanisms. While it will prevent tampering with the data (except deletion), SP cannot develop real-time services on this data. Thus, the first challenge is how to store sensor data at SP according to data-capture rules; hence, the verifier can attest the data. The second challenge arises due to the computational cost at the verifier and communication cost between the verifier and SP to send the verifiable sensor data; e.g., if the verifier wishes to attest only one-day old sensor data over the sensor data of many years, then sending many years of sensor data to the verifier is impractical. Finally, the last challenge is how to securely store data when data-capture rules are set to be false (i.e., not to store data). In this case, not storing any data would not provide a way for verification, since SP may store data and can inform the verifier that there is no data as per the existing data-capture rules.

**Solution.** To address the first challenge, we develop a cryptographic sealing method based on hash-chains and XOR-linked-lists to ensure immutable sensor logs, after the sealed logs leave the enclave. Thus, any log addition/deletion/update is detectable (see §4.2.1). To address the second challenge, we execute sealing on small-sized chunks, which each maintains its hash-chain and XOR-links. The XOR-links ensure the log completeness, i.e., a chunk before and after the desired chunk has not been altered (see §4.2.1). To address the third challenge, we store the *device state* of the first sensor reading for which the data-capture rule is set to false. Further, we discard all subsequent sensor readings, unless finding a sensor reading for which the data-capture rule is to *store data* (§4.2.3).

**C3. Privacy-preserving log verification.** In case of log-integrity verification, SP can provide the entire sensor data with a cryptographically sealed log to the trusted auditor. But, the challenge arises, if a user asks to verify her




Fig. 2. Phases in IoT NOTARY.

user-associated data/query results. Here, SP cannot provide the entire sensor data to the user, since it will reveal other users' privacy.

**Solution.** To address this challenge, we develop a technique to non-deterministically encrypt the user-id before cryptographically-sealing the sensor data. However, only non-deterministic encryption is also not enough to verify the log completeness, we compute XOR operation on each sensor reading, and then, create XOR-linked-list (see §4.2.2).

### 3.2 IoT NOTARY: An Overview

This section presents an overview of the three phases and dataflow among different entities and devices, see Figure 3.

**Notification phase: SP to Users messages.** This is the first phase that notifies users about data-capture rules for the IoT space using notice messages (in a verifiable manner for later stages). Such messages can be of two types: (i) notice messages, and (ii) notice-and-acknowledgment messages. SP establishes (the default) data-capture rules and informs trusted hardware (1). Trusted hardware securely stores data-capture rules (2, 5) and informs the *trusted notifier* (3) that transmits the message to all users (4). Only notice messages need a trusted notifier to transmit the message (see §4.1).

**Log-sealing phase: Sensor devices to SP messages.** Each sensor sends data to the controller (1). The controller receives the correct data, generated by the actual sensor, as per our assumptions (and settings of the university IT department). The controller sends encrypted data to SP (6) that authenticates the controller using any existing authentication protocol, before accepting data. Trusted hardware (Intel SGX) at SP reads the encrypted data in the enclave (7).

*Working of the enclave.* The enclave decrypts the data and checks against the pre-notified data-capture rules. Recall that the decrypted data is of the format:  $\langle d_i, s_j, t_k \rangle$ , where  $d_i$  is  $i^{th}$  user-device-id,  $s_j$  is the  $j^{th}$  sensor device, and  $t_k$  is  $k^{th}$  time. After checking each sensor reading, the enclave adds a new field, entitled *sensor (device) states*. The sensor state of a sensor  $s_j$  is denoted by  $s_j.state$ , which can be *active* or *passive*, based on capturing user data. For example,  $s_j.state = active$  or (1), if data captured by the sensor  $s_j$  satisfies the data-capture rules; otherwise,  $s_j.state = passive$  or (0). For all the sensors whose  $state = 0$ , the enclave deletes the data. Then, the enclave cryptographically seals sensor data, regardless of the sensor state, and provides cleartext sensor data of the format:  $\langle d_i, s_j, s_j.state = 1, t_k \rangle$  to SP (8) that provides services using this data (9). Note that the cryptographically sealed logs and cleartext sensor data are kept at untrusted storage of SP (8, 10).

**Verification phase: SP to verifier messages.** In our model, an auditor and a user can verify the sensor data. The auditor can verify the entire/partial sensor data against data-capture rules by asking SP to provide cleartext sensor data and cryptographically sealed logs (8, 10). The users can also verify their own data against pre-notified messages or can verify the results of the services provided by SP using only cryptographically sealed logs (12). Note that using an underlying authentication technique (as per our assumptions), auditor/users and SP authenticate each other before transmitting data from SP to auditor/users.

## 4 ATTESTATION PROTOCOL

This section presents three phases of attestation protocol. The important notations used in the presentation of the protocol are listed in Table 1.

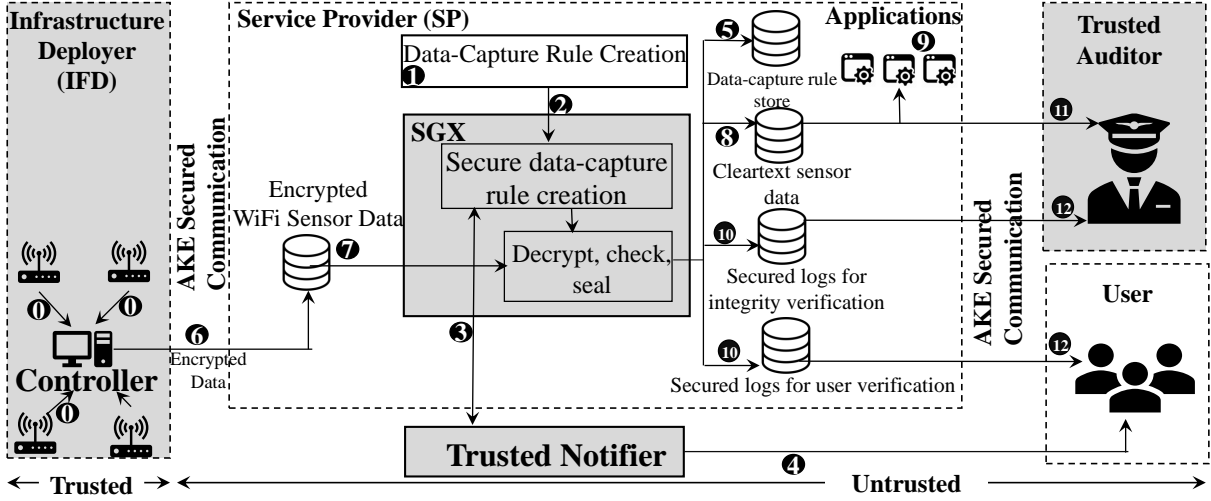


Fig. 3. Dataflow and computation in the protocol. Trusted parts are shown in shaded box.

**Preliminary Setup Phase.** We assume a preliminary setup phase that distributes public keys ( $PK$ ) and private keys ( $PR$ ), as well as, registers user devices into the system. The trusted authority (which is the university IT department in our setup of TIPPERS) generates/renews/revokes keys used by the secure hardware enclave (denoted by  $\langle PK_E, PR_E \rangle$ ) and the notifier (denoted by  $\langle PK_N, PR_N \rangle$ ). The keys are provided to the enclave during the secure hardware registration process. Also,  $\langle PK_{di}, PR_{di} \rangle$  denotes keys of the  $i^{th}$  user device. *Usages of keys:* The controller uses  $PK_E$  to encrypt sensor readings before sending them to SP.  $PR_E$  is also used by the enclave to write encrypted sensor logs and decrypt sensor readings.  $PK_N$  is used during the notification phase by SGX to send an encrypted message to the notifier. User device keys are used during device registration, as given below.

We assume a registration process during which a user identifies herself to the underlying system. For instance, in a WiFi network, users are identified by their mobile devices, and the registration process consists of users providing the MAC addresses of their devices (and other personally identifiable information, *e.g.*, email and a public key). During registration, users also specify their preferred modality through which the system can communicate with the user (*e.g.*, email and/or push messages to the user device). Such communication is used during the notification phase.

#### 4.1 Notification Phase

The notification phase informs data-capture rules established by SP to the (registered) users by explicitly sending *notice messages*. We consider two models for notification, differing based on acknowledgment from users.

In the *notice-only model (NoM)*, SP informs users of data-capture rules, but users may not acknowledge receipt of the message. Such a model is used to implement policies, when data capture is mandatory, and the user cannot exercise control, over data capture. Since there is no acknowledgment, SP is only required to ensure that it sends a notice, but is not required to guarantee that the user received the notice. In contrast, a *notice-and-ACK model (NaM)* is intended for discretionary data-capture rules that require explicit permission from users prior to data capture. Such rules may be associated, for instance, with fine-grained location services that require users' location. A user can choose not to let SP track his location, but will likely not be able to avail of some services.

Implementation of notification differs based on the model used. Interestingly, since NaM requires acknowledgment, the notification phase is easier as compared to NoM that uses a trusted notifier to deliver the message to users. Below we discuss the implementation of both models:

*Notification implementation in NoM.* NoM assumes that, by default, data-capture rules are set not to retain any user data, unless SP, first, informs SGX about a data-capture rule, (*i.e.*, SP cannot use the encrypted sensor data for building any application, see 9 in Figure 3). When SP creates a new data-capture rule, SP must inform SGX. Then, the enclave encrypts the data-capture rule using the public key (*i.e.*,  $PK_N$ ) of the notifier and informs the trusted notifier (via SP) about the encrypted data-capture rule by writing it outside of the enclave (in our user-case §1, the university IT department works as a trusted notifier). Data-capture rules are maintained by SP on stable storage, which is read by SGX into the enclave to check, if the sensor data should be forwarded to SP. SGX can retain a cache of rules in



the enclave, if such rules are still valid (and hence used for enforcement).<sup>1</sup> Finally, the trusted notifier acknowledges SP about receiving the encrypted data-capture rule, and then, informs users of the encrypted data-capture rule via signed notice messages. On receiving the notice message, the users may decrypt it and obtain the data-capture rule.

To see the role of *trusted hardware* above, suppose that SP was responsible for informing users about data-capture rules directly. Data-capture rules are also required by SGX during log-sealing (PHASE 2). An adversarial SP may inform SGX, not to users, or may inform non-identical rules to users and to SGX. Hence, SP cannot inform the rule to users directly.

To see the role of the *trusted notifier* above, suppose that SP can directly inform users about encrypted data-capture rules obtained from SGX. An adversarial SP may not deliver the data-capture rule to all or some of the users; thus, an encrypted data-capture rule is not helpful. Hence, a trusted notifier ensures that the notice message is sent to all the registered users. Note that the trusted notifier might be a trusted website that lists all the data capture rules which users can access.

*Implementation of notification in NaM.* Unlike NoM, the notification phase of NaM does not require the trusted notifier. In NaM, by default, SP cannot utilize all those sensor readings having device-ids for which the users have not acknowledged. Likewise NoM, in NaM, SP informs data-capture rules to SGX that encrypts the rule and writes outside of the enclave. The encrypted rules are delivered by SP to users, unlike NoM. On receiving the message, a user may securely acknowledge the enclave about her consent. The enclave retains all those device-ids that acknowledge the notice message for the log-sealing phase and considers those device-ids during the log-sealing phase to retain their data while discarding data of others.

## 4.2 Log Sealing Phase

The second phase consists of cryptographically sealing the sensor data for future verification against pre-notified data-capture rules. The sensor data is sealed into secured logs using authenticated data structures, *e.g.*, hash-chains and XOR-linked lists (as shown in Figures 4, 5), by the sealing function,  $Sealing(PR_E, \langle d_i, s_j, s_j.state, t_k \rangle)$ , executed in the enclave at SP. Let us explain log-sealing in the context of WiFi connectivity data. The enclave reads the encrypted sensor data (7 in Figure 3) and executes the three steps: (i) decrypts the data, (ii) checks the data against pre-notified valid data-capture rules, and (iii) cryptographically seals the data and store *appropriate secured logs*.

Below we explain our log sealing approach. To simplify the discussion, we first consider the case when all the sensor data satisfies some data-capture rule (*i.e.*, the state of all the sensor data is one), and hence, data is forwarded to and stored at SP §4.2.1. Then, we adapt the protocol to deal with the case when some sensor data satisfies some data-capture rule (*i.e.*, the state of some sensor data is one, and hence, data is forwarded to and stored at SP), while remaining sensor data does not satisfy any rule (*i.e.*, the state of the remaining sensor data is zero, and hence, data is forwarded to SP) 4.2.3.

**4.2.1 Sealing Entire Sensor Data.** Informally, the sealing function executes a hash function on each sensor reading (or value), whose output is used to create a chain of hash digests. At the end of the sensor readings/values, the sealing function generates an authenticated proof-of-integrity by mixing a computationally-hard secure string. For example, consider four values:  $v_1, v_2, v_3$ , and  $v_4$ . The sealing function works as follows:

Value	Hash output
$v_1$	$h_1 \leftarrow H(v_1    H(0))$
$v_2$	$h_2 \leftarrow H(v_2    h_1)$
$v_3$	$h_3 \leftarrow H(v_3    h_2)$
$v_4$	$h_4 \leftarrow H(v_4    h_3)$
Proof-of-integrity	$\langle sign_{PR_E}(SS \oplus h_4) \rangle$

In this example, all the values are hashed while including the hash digest of the previous value, except the first value. Note that only the first value is hashed with the hash digest of zero. In the end, the proof-of-integrity is prepared by signing XORed-value of the hash digest of the last value and a secret random string,  $SS$ . Note that the secret random string is generated for each chunk in a specific manner, which will be clear in the detailed description of the protocol; please see below.

<sup>1</sup>Since the enclave has limited memory, the enclave cannot retain all the valid and non-valid data-capture rules after a certain size. Thus, the enclave writes all the non-valid data-capture rules on the disk after computing a secured hash digest over all the rules. Taking a hash over the rules is required to maintain the integrity of all the rules, since any rule written outside of the enclave can be tampered by SP. Recall that altering a rule will make it impossible to verify partial/entire sensor data.

Sensor data after checking data-capture rule in the enclave	
$d_i$	$\triangleq$ $i^{th}$ user-device-id
$t_k$	$\triangleq$ $k^{th}$ time
$s_j$	$\triangleq$ $j^{th}$ sensor device
$s_j.state$	$= \begin{cases} 1, & \text{if the sensor } s_j \text{'s state is active; the data can be forwarded to SP} \\ 0, & \text{otherwise} \end{cases}$
Data in log sealing phase	
$C_x$	$\triangleq$ $x^{th}$ data chunk
$g^b$	$\triangleq$ A random string for chunk $C_x$
$S_{eoc}^x$	$\triangleq$ XORed value of random string $g^a, g^b, g^c$ from chunk $C_v, C_x, C_y$
$h_i^x$	$\triangleq$ The hash digest of $i^{th}$ sensor reading in chunk $C_x$ (for log integrity)
$H$	$\triangleq$ A hash function known to the service provider, the auditor, and the user
$hu_i^x$	$\triangleq$ The hash digest of $i^{th}$ sensor reading in chunk $C_x$ (for user verification)
$\mathcal{PI}$	$\triangleq$ Proof-of-(log)-integrity of a chunk, used by the trusted auditor
$\mathcal{PU}$	$\triangleq$ Proof-of-integrity of a chunk, for the user/service (query) verification

Table 1. Table of Notation for the Attestation Protocol

$\langle d_1, s_1, 1, t_1 \rangle$ $\langle d_2, s_2, 1, t_2 \rangle$ $\langle d_3, s_1, 1, t_3 \rangle$ $\langle d_4, s_1, 1, t_4 \rangle$	$h_1 \leftarrow H(d_1 \  s_1 \  1 \  t_1 \  H(0))$ $h_2 \leftarrow H(d_2 \  s_2 \  1 \  t_2 \  h_1)$ $h_3 \leftarrow H(d_3 \  s_1 \  1 \  t_3 \  h_2)$ $h_4 \leftarrow H(d_4 \  s_1 \  1 \  t_4 \  h_3)$ $\mathcal{PI}_{C_x} \leftarrow (g^b, \text{Sign}_{PRE}(h_4 \oplus S_{eoc}^x))$	$o_1 \leftarrow H(d_1 \  t_1)$ $o_2 \leftarrow H(d_2 \  t_2)$ $o_3 \leftarrow H(d_3 \  t_3)$ $o_4 \leftarrow H(d_4 \  t_4)$ $hu_{end} \leftarrow hu_1 \oplus hu_2 \oplus hu_3 \oplus hu_4$ $\mathcal{PU}_{C_x} \leftarrow (g^b, \text{Sign}_{PRE}(h_q^{end} \oplus S_{eoc}^x))$
<b>Sensor data after passing the enclave</b>	<b>Sealing function execution for log-integrity</b>	<b>Sealing function execution for user's data/ query verification</b>

Fig. 4. Cryptographically sealing procedure executed on a chunk,  $C_x$ . Gray-shaded data is not stored on the disk. White-shaded data is stored on the disk and accessible by SP. The figure shows proof-of-integrity only for one chunk,  $C_x$ ; hence, some notations are abusively used. Note that we used  $h_i$  to denote a hash digest computed for verifying the entire sensor data, while  $hu_i$  denotes a hash digest computed for verifying the user data or query results.

The sealing operation consists of the following three phases: (i) chunk creation, (ii) hash-chain creation, and (iii) proof-of-integrity creation; described below.

**PHASE 1: Chunk creation.** The first phase of the sealing operation finds an appropriate size of a chunk (to speed up the attestation process). Note that the incoming encrypted sensor data may be large, and it may create problems during verification, due to increased communication between SP and the verifier. Also, the verifier needs to verify the entire data, which has been collected over a large period of time (e.g., months/years). Further, creating cryptographic sealing over the entire sensor data may also degrade the performance of *Sealing()* function, due to the limited size of the SGX enclave. Thus, we first determine an appropriate chunk size, for each of which the sealing function is executed.

The chunk size depends on time epochs, the enclave size, the computational overhead of executing sealing on the chunk, and the communication overhead for providing the chunk to the verifier. A small chunk size reduces the

communication overhead and maintains the log minimality property, thereby during the log verification phase, a verifier retrieves only the desired log chunks, instead of retrieving the entire sensor data. Consequently, SP stores many chunks.

**PHASE 2: Hash-chain creation.** Consider a chunk,  $C_x$ , that can store at most  $n$  sensor readings, each of them of the format:  $\langle d_i, s_j, t_k \rangle$ . The sealing function checks each sensor reading against data-capture rules and adds sensor state to each reading, as:  $\langle d_i, s_j, s_j.state, t_k \rangle$ . Since in this section we assumed that all sensor data will be stored, the sensor state of each sensor reading is set to 1. The sealing function starts with the first sensor reading of the chunk  $C_x$ , as follows:

First sensor reading. For the first sensor reading of the chunk, the sealing function computes a hash function on value zero, i.e.,  $H(0)$ . Then, the sealing function mixes  $H(0)$  with the remaining values of the sensor reading, i.e., sensor-id, device-id, sensor state, and time, at which it computes the hash function, denoted by  $H(d_1||s_j||s_j.state||t_k||H(0))$  that results in a hash digest, denoted by  $h_1^x$ . After processing the complete first sensor reading of the chunk  $C_x$ , the enclave writes cleartext first sensor reading of  $C_x$ , i.e.,  $\langle d_1, s_j, s_j.state, t_k \rangle$  on the disk, which can be accessed by SP.

Second sensor reading. Let  $\langle d_2, s_j, s_j.state, t_{k+1} \rangle$  be the second sensor reading. For this, the sealing function works identically to the processing of the first sensor reading. It computes a hash function on the second sensor values, while mixing it with the hash digest of the first sensor reading, i.e.,  $H(d_2||s_j||s_j.state||t_{k+1}||h_1^x)$  that results in a hash digest, say  $h_2^x$ . Finally, the enclave writes the second sensor reading in cleartext on the disk.

Processing the remaining sensor readings. Likewise, for the second sensor reading processing, the sealing function computes the hash function on all the remaining sensor readings of the chunk  $C_x$ . After processing the last sensor reading of the chunk  $C_x$ , the hash digest  $h_n^x$  is obtained.

**PHASE 3: Proof-of-integrity creation.** Since each sensor reading is written on disk, SP can alter sensor readings, to make it impossible to verify log integrity by an auditor. Thus, to show that all the sensor readings are kept according to the pre-notified data-capture rules, the sealing function prepares an immutable proof-of-integrity for each chunk, as follows:

For each chunk  $C_i$ , the sealing function generates a random string, denoted by  $g^j$ , where  $i \neq j$ . Let  $C_v$ ,  $C_x$ , and  $C_y$  be three consecutive chunks (see Figure 5), based on consecutive sensor readings. Let  $g^a$ ,  $g^b$ , and  $g^c$  be random strings for chunks  $C_v$ ,  $C_x$ , and  $C_y$ , respectively. The use of random strings will ensure that any of the consecutive chunks have not been deleted by SP (will be clear in §4.3). Now, for producing the proof-of-integrity for the chunk  $C_x$ , the sealing function: (i) executes XOR operation on  $g^a$ ,  $g^b$ ,  $g^c$ , whose output is denoted by  $S_{eoc}^x$ , where  $eoc$  denotes the end-of-chunk; (ii) signs  $h_n^x$  XORed with  $S_{eoc}^x$  with the private key of the enclave; and (iii) writes the proof-of-integrity for log verification of the chunk  $C_x$  with the random string  $g^b$ , as follows:

$$\mathcal{PI}_{C_x} = (g^b, \text{Sign}_{PRE}(h_n^x \oplus S_{eoc}^x))$$

**Note.** We do not generate the proof for each sensor reading. The enclave writes only the proof and the random string for each chunk to the disk, which is accessible by SP. Further, the sensor readings having the state one are written on the disk, based on which SP develops services.

**Example.** Please see Figure 4, where the **middle box** shows PHASE 2 execution on four sensor readings. Note that the hash digest of each reading is passed to the next sensor reading on which a hash function is computed with the sensor reading. After computing  $h_4$ , the proof-of-integrity,  $\mathcal{PI}$ , is created that includes signed  $h_4 \oplus S_{eoc}^x$  and a random string,  $g^b$ .

*Note.  $g^*$  for the first chunk.* The initialization of the log sealing function requires an initial seed value, say  $g^*$ , due to the absence of  $0^{th}$  chunk. Thus, in order to initialize the secure binding for the first chunk, the seed value is used as a substitute random string.

**4.2.2 Sealing Data for User Data/Service Verification.** While capturing *user-associated data*, users may wish to verify their user-associated data against notified messages. Note that *the protocol presented so far requires entire cleartext data to be sent to the verifier to attest log integrity* (it will be clear soon in §4.3). However, such cleartext data transmission is not possible in the case of user-associated data verification, since it may reveal other users' privacy. Thus, to allow verification of user-associated data (or service/query result<sup>2</sup> verification), we develop a new sealing

<sup>2</sup>The users, who access services developed by SP (as mentioned in §1), may also wish to verify the query results, since SP may tamper with the data to show the wrong results.

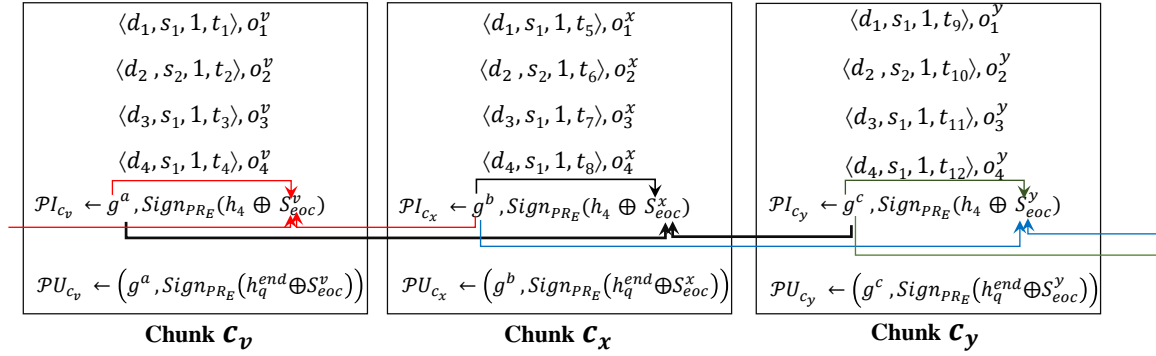


Fig. 5. PHASE 3: end of chunk,  $S_{eoc}$ , creation for three chunks. Observe that  $S_{eoc}^x = g^a \oplus g^b \oplus g^c$ .

method, consists of the three phases: (i) chunk creation, (ii) hash-generation, and (iii) proof-of-integrity creation. The chunk creation phase of this new sealing method is identical to the above-mentioned chunk creation phase 1; see §4.2.1. Below, we only describe PHASE 2 and PHASE 3.

**PHASE 2: Hash-generation.** Consider a chunk,  $C_x$ , that can have at most  $n$  sensor readings, each of them of the format:  $\langle d_i, s_j, s_j.state, t_k \rangle$ . Our objective is to hide users' device-id and its frequency-count (*i.e.*, which device-id is prominent in the given chunk). Thus, on the  $i^{\text{th}}$  sensor reading, the sealing function mixes  $d_j$  with  $t_k$ , and then, computes a hash function over them, denoted by  $H(d_j || t_k)$  that results in a digest value, say  $o_i$ . Note that hash on device-ids mixed with time results in two different digests for more than one occurrence of the same device-id. Note that  $o_i$  helps the user to know his presence/absence in the data during attestation, but it will not prove that tampering has not happened with the data. Then, the sealing function mixes  $o_i$  with the sensor state (to produce a proof of sensor state) of the  $i^{\text{th}}$  sensor reading, and on which it computes the hash function, denoted by  $H(o_i || s_j.state)$  that results in a hash digest, denoted by  $hu_i^x$ . After processing the  $i^{\text{th}}$  sensor reading of the chunk  $C_x$ , the enclave writes  $o_i$  on the disk. After processing all the  $n$  sensor readings of the chunk  $C_x$ , the sealing function computes XOR operation on all hash digests,  $hu_i^x$ , where  $1 \leq i \leq n$ :  $hu_1^x \oplus hu_2^x \oplus \dots \oplus hu_n^x$ , whose output is denoted by  $hu_{end}^x$ . (Reason of computing  $hu_{end}^x$  will be clear in §4.3).

**PHASE 3: Proof-of-integrity creation for the user.** The sealing function prepares an immutable proof-of-integrity for users, denoted by  $\mathcal{P}U$ , for each chunk and writes on the disk. Likewise, proof-of-integrity for entire log verification,  $\mathcal{P}I$  (§4.2.1), for each chunk, the sealing function obtains  $S_{eoc}$ ; refer to PHASE 3 in §4.2.1. Now, for producing  $\mathcal{P}U$  for the chunk  $C_x$ , the sealing function: (i) signs  $hu_{end}^x$  XORed with  $S_{eoc}^x$  with the private key of the enclave, and (ii) writes the signed output with the random string of the chunk,  $g^b$ , as  $\mathcal{P}U_{C_x}$ .

$$\mathcal{P}U_{C_x} = (g^b, \text{Sign}_{\text{PRE}}(hu_{end}^x \oplus S_{eoc}^x))$$

**Note.** The enclave writes hash digests,  $o_i$  for each sensor reading, the proof for user verification, and the random string for each chunk on the disk. Of course, the sensor readings having the state one are written on the disk.

**Example.** Please see Figure 4, where the **last box** shows PHASE 2 execution on four sensor readings to obtain the proof-of-integrity for the user,  $\mathcal{P}U$ . ■

**4.2.3 Sealing Mixed State Sensor Data.** The protocol so far has assumed that all data has the sensor state of one. We next consider how it can be generalized to a situation when some sensor readings may not satisfy the data capture rules (and hence, have the sensor state of zero). Recall that (as mentioned in §3.2), the enclave decrypts the sensor data received from the IFD and checks against the pre-notified data-capture rules, and if a sensor reading captured by a sensor  $s_i$  does not satisfy the data-capture rules, then the sensor state of the sensor reading becomes zero. Please note that the sensor state of zero does not indicate that the sensor is turned off.

**Example 4.2.3.1.** In a chunk, assume the following six sensor readings, produced by two sensors  $s_1$  and  $s_2$ .

- $\langle d_1, s_1, 1, t_1 \rangle$
- $\langle d_2, s_2, 0, t_2 \rangle$
- $\langle d_2, s_2, 0, t_3 \rangle$
- $\langle d_3, s_2, 0, t_4 \rangle$
- $\langle d_3, s_2, 1, t_5 \rangle$

$$\langle d_1, s_1, 1, t_6 \rangle$$

In this case, there is no need to store all sensor readings having a state zero. However, doing it carelessly may provide an opportunity to the adversary to delete all the sensor readings and prove that such readings are deleted due to the sensor state being zero, due to not satisfying data-capturing rules. Thus, to avoid storing sensor readings having sensor state zero, we provide a method below. ■

Informally, the sealing function checks each sensor reading and deletes all those sensor readings for which sensor state is passive, due to not satisfying data-capture rules. In this case, it is not mandatory to seal each sensor readings, as mentioned in §4.2.1. Thus, the sealing function provides a *filter operation* that removes sensor readings whose device state is passive, while storing sensor readings whose device states are active. But, the sealing function cryptographically stores the data with minimal information of sensor readings whose device states are passive. Below we describe PHASE 2 for sealing mixed state sensor data. Note that PHASE 1 (chunk creation) and PHASE 3 (proof-of-integrity creation) are identical to the method described in §4.2.1.

**PHASE 2: Sealing operation.** Formally, to create hash-chain for this case, the sealing function will do the following: For the first  $i^{th}$  sensor reading (for example,  $\langle d_j, s_k, 0, t_i \rangle$ ) whose *state* = 0, the enclave executes two operations: (i) sealing function that computes a hash function, as:  $H(s_j || s_j.state || t_k || h_{i-1})$ , whose output is denoted by  $h_i$ , and (ii) filter operation that deletes the device-id  $d_j$  and stores  $\langle s_k, 0, t_i \rangle$  on the disk. Now for all the successive sensor readings until encountering a sensor reading, say  $l$ , with *state* = 1, all the sensor readings are discarded (not stored on the disk), as well as, the hash function is not executed. However, the sealing function computes the hash function on the  $l^{th}$  sensor reading ( $\langle d_x, s_k, 1, t_l \rangle$ ), as:  $H(d_x || s_j || 1 || t_l || h_i)$  and stores  $l^{th}$  sensor reading on the disk.

**Example 4.2.3.2.** Now, we apply the above-mentioned PHASE 2 on the sensor readings given in Example 4.2.3.1. Here, the enclave does not store the second, third, and fourth sensor readings on disk and delete them, while the remaining sensor readings will be stored on the disk. In addition, the enclave will store only the sensor device and its state of the first reading (*i.e.*, the second sensor reading) for which state was zero and generate verifiable logs, such that the verifier can verify that the three entries has been deleted due to not satisfying data-capturing rules. The sealing function computes the hash-chain and proof-of-integrity as follows:

$$\begin{aligned} h_1 &\leftarrow H(d_1 || s_1 || 1 || t_1 || H(0)) \\ h_2 &\leftarrow H(s_2 || 0 || t_2 || h_1) \\ h_3 &\leftarrow H(d_3 || s_2 || 1 || t_5 || h_2) \\ h_4 &\leftarrow H(d_1 || s_1 || 1 || t_6 || h_3) \\ \mathcal{PI} &\leftarrow \langle secret\ string, sign_{PRE}(secret\ string \oplus h_4) \rangle \end{aligned}$$

Note that here we compute a hash function on the first, second, fifth, and last sensor readings, while do not seal the third and fourth sensor readings, due to their passive (or zero) sensor states.<sup>3</sup> ■

**4.2.4 Log-size Optimization.** The above-mentioned procedure given in §4.2.3 regards the sensor states and stores less amount of cryptographically sealed data on the disk, by not storing consecutive sensor readings having sensor state of zero. However, such improvement is limited, if the sensor readings have state zero and one in an alternative sequence.

**Example 4.2.4.1.** Consider the following seven sensor readings obtained from two sensors  $s_1$  and  $s_2$ . Here, the state of sensor  $s_2$  is zero, due to not satisfying the data-capturing rules; however, the state of sensor  $s_1$  is one.

$$\begin{aligned} &\langle d_1, s_1, 1, t_1 \rangle \\ &\langle d_2, s_2, 0, t_2 \rangle \\ &\langle d_2, s_1, 1, t_3 \rangle \\ &\langle d_3, s_2, 0, t_4 \rangle \\ &\langle d_3, s_1, 1, t_5 \rangle \\ &\langle d_1, s_2, 0, t_6 \rangle \\ &\langle d_1, s_1, 1, t_7 \rangle \end{aligned}$$

In such scenario, the method given in §4.2.3 is not efficient, since it will store all the sensor readings and produce hash digest for each sensor reading, as follows:

<sup>3</sup>We can only compress  $x > 1$  continuous sensor readings, say  $\langle *d, *s, *s.state = 0, *t \rangle$  (where  $*d$ ,  $*s$ , and  $*t$  denote any device-id, sensor-id, and time, respectively) to produce a proof that such  $x$  readings have been deleted. However, we cannot compress  $x$  sensor readings having  $*s.state = 1$ , since it disallows to verify service integrity (*e.g.*, a user query, how many times a user has visited a location, cannot be verified, if  $x$  readings with  $*s.state = 1$  have been deleted).

$$\begin{aligned}
h_1 &\leftarrow H(d_1||s_1||1||t_1||H(0)) \\
h_2 &\leftarrow H(s_2||0||t_2||h_1) \\
h_3 &\leftarrow H(d_2||s_1||1||t_3||h_2) \\
h_4 &\leftarrow H(s_2||0||t_4||h_3) \\
h_5 &\leftarrow H(d_3||s_1||1||t_5||h_4) \\
h_6 &\leftarrow H(s_2||0||t_6||h_5) \\
h_7 &\leftarrow H(d_1||s_1||1||t_7||h_6)
\end{aligned}$$

Thus, in order to reduce the size of the secured log for the case when sensor readings have state zero and one in an alternative sequence, below, we propose a method that stores logs for each sensor or each user device. ■

**Per sensor-based logging.** We implement a log-sealing procedure (given in §4.2.1 and §4.2.3) on per sensor, *i.e.*, we group the sensor readings produced by the same sensor before producing cryptographically sealed logs. This optimization method works as follows, for  $n$  sensor readings and  $x$  sensors:

- (1) The enclave creates  $x$  buffers, one buffer for each of the  $x$  sensors.
- (2) The enclave reads the  $n$  sensor readings from the disk and decrypts them.
- (3) Each of the  $n$  sensor readings is allocated to one of the buffers based on the sensor.
- (4) On each sensor reading allocated to a buffer, the enclave either executes the method given in §4.2.1 if all the sensor readings have state one, or executes the method given in §4.2.3 if the sensor readings have state one and zero.
- (5) At the end, the enclave writes cryptographically secured logs and proof-of-integrity for each buffer's sensor log, on the disk.

**Example 4.2.4.2.** Now, we show how one can reduce the size of secured logs for the seven sensor readings given in Example 4.2.4.1 by keeping logs per sensor-based. Here, the enclave maintains two buffers: one for  $s_1$  and another for  $s_2$ . All the sensor readings of  $s_1$  are sealed using the method of §4.2.1, and the sensor readings of  $s_2$  are sealed using the method of §4.2.3. Thus, the enclave executes the following computation:

The first buffer dealing with the sensor readings of sensor  $s_1$ :

$$\begin{aligned}
h_1 &\leftarrow H(d_1||s_1||1||t_1||H(0)) \\
h_2 &\leftarrow H(d_2||s_1||1||t_3||h_1) \\
h_3 &\leftarrow H(d_3||s_1||1||t_5||h_2) \\
h_4 &\leftarrow H(d_1||s_1||1||t_7||h_3)
\end{aligned}$$

$$\mathcal{PI}_{s_1} \leftarrow \langle \text{secret string}, \text{sign}_{PR_E}(\text{secret string} \oplus h_4) \rangle$$

The second buffer dealing with the sensor readings of sensor  $s_2$ :

$$h_5 \leftarrow H(s_2||0||t_2||H(0))$$

$$\mathcal{PI}_{s_2} \leftarrow \langle \text{secret string}, \text{sign}_{PR_E}(\text{secret string} \oplus h_5) \rangle$$

Note that, here, we used the notations  $\mathcal{PI}_{s_1}$  and  $\mathcal{PI}_{s_2}$  to indicate the proof-of-integrity generated for sensors  $s_1$  and  $s_2$ , respectively. The enclave writes cleartext sensor readings corresponding to the sensor  $s_1$ , the tuple  $\langle s_2, 0, t_2 \rangle$ , hash digests  $h_1, h_2, \dots, h_5$ , and proof-of-integrity for both sensors  $\mathcal{PI}_{s_1}$  and  $\mathcal{PI}_{s_2}$ , on the disk. ■

**Issues.** The above-mentioned method, while reduces the size of cryptographically sealed logs, it faces two issues, as follows:

- (1) *More sensors.* It may happen that the number of sensors is significantly more, thereby it is not easy to maintain buffers for each sensor, due to a limited memory of the secure hardware. In this case, we can still use the above method; however, one buffer is responsible for more than one sensor. For example, if there are four sensors,  $s_1, s_2, \dots, s_4$ , and the secure hardware can hold only two buffers, then the first buffer might be responsible for sensor readings corresponding to two sensors  $s_1$  and  $s_2$ , and another buffer might be responsible for sensor readings corresponding to the remaining two sensors. To allocate sensor readings to buffers, one can use a hash function on the sensor-id to know the buffer id.
- (2) *Different-sized log.* Note that when we create multiple buffers for sensors, the enclave writes multiple chunks and proof-of-integrity corresponding to multiple buffers. The size of each chunk may not be identical, due to having a different number of hash digests. It may reveal information to the adversary that which sensor's state is zero. To avoid such leakage, the enclave may pad output produced for each buffer with some fake values and may include this information in proof-of-integrity to show how many fake values are added. For instance, in Example 4.2.4.2, the output corresponding to the second buffer has only one hash digest (*i.e.*,  $h_5$ ) while the output corresponding to the first buffer will have four hash digests (*i.e.*,  $h_1, h_2, h_3, h_4$ ). Thus, to write the same size data for each buffer,

the enclave may pad the output corresponding to the second buffer with three fake hash digests and write this information in the proof-of-integrity, as follows:  $\mathcal{PI}_{s_2} \leftarrow \langle \text{secret string}, \text{sign}_{PR_E}(\text{secret string} \oplus h_5, 3) \rangle$ . Note that while adding fake hash digests, the user does not need to verify any fake digest, and thus, verifying only desired data will reduce the verification time, as we will see in Experiment 7 in §6.

**Note: Per user-based logging.** We can also apply the same procedure for each user device-id to reduce the size of the cryptographically sealed data, by creating buffers for each user device id. This may reduce the verification time for user-related data.

### 4.3 Attestation Phase

The attestation phase contains two sub-phases: (i) key establishment between the verifier and service provider to retrieve logs (§4.3.1), and (ii) verification of the retrieved logs (§4.3.2).

**4.3.1 Key Establishment.** The secure log retrieval is crucial for proof validation and non-trivial in the presence of a de-centralized verifier model, where anyone can execute a remote request to attest the secured logs against the data-capture rules. Thus, before the log transmission from the service provider to the verifier (*i.e.*, a user or an auditor), for each attestation request, the verifier’s identity must be authenticated prior to the session key establishment.

Our log retrieval scheme is based on Authenticated Key Exchange (AKE) protocol [29], where a verifier and the service provider (*i.e.*, prover) dynamically establish a session key (shown in Figure 6). This dynamic key establishment provides forward secrecy for all future sessions, such that, for any compromised session in the future, all sessions in the past remain secure. To achieve these properties, we use SIGMA [29] protocol, which is the cryptographic base for the Internet Key Exchange (IKE) protocol. The family of SIGMA protocols is based on Diffie-Hellman key exchange [19]. We only show a 3-round version of SIGMA, as it provides verifier/sender-identity protection during the key establishment process. Recall that in our solution, the prover is a centralized service provider, but the verifier can be anyone, and therefore, we use this identity protection method to achieve the verifier’s identity privacy during the session.

Without the loss of generality, we, first, define the Computational Diffie Hellman (CDH) [19]. During the session, the verifier and the prover must execute the computations within a cyclic group  $G = \langle g \rangle$  that has a generator  $g$  of prime order  $q$ . According to the CDH assumption, the computation of a discrete logarithm function on public values  $(g, g^x, g^y)$  is hard within the cyclic group  $G$ . In particular, given the publicly known value  $g$ , it is hard to distinguish  $g^{xy}$  from  $g^x$  and  $g^y$  without knowing the ephemeral secrets  $x$  and  $y$ .

Figure 6 depicts SIGMA-based communication flow between a verifier and the service provider/prover. Initially, a verifier selects a secret value  $x$ , computes  $g^x$  as a public exponent, and sends it to the prover. Similarly, the prover selects a secret value  $y$ , computes  $g^y$ , and receives the public exponent  $g^x$  from the verifier. Next, the prover computes a joint secret value as:  $e = g^{xy}$ , and also uses it to derive a message authentication code (MAC) key  $MAC_k$ .

The prover composes a message structure as:  $[g^y, SP_{id}, MAC_k(SP_{id}||g^x||g^y)]$ , where  $g^y$  is the prover’s public exponent for session key generation,  $SP_{id}$  is the identity of sender/prover, and  $MAC_k(SP_{id}||g^x||g^y)$  is the message authentication code generated on sender’s identity and all public exponents used for session key derivation. Note that in order to generate this MAC, a separate key  $MAC_k$  is derived from the session key  $e$ .

Next, the verifier receives this message and retrieves the public exponent  $g^y$  to generate a local copy of session key  $e$ , as well as, the message authentication code generating key  $MAC_k$  derived from  $e$ . At this stage, both parties have locally computed a secure session key  $e$ ; however, there is still an identity disclosure required from the verifier to prove that the freshly generated key  $e$  binds to an authentic identity holder  $v_{id}$ , where  $v_{id}$  is the verifier’s identity.

Thus, the verifier sends a message, consisting of  $[v_{id}, MAC_k(v_{id}||SP_{id}||g^x||g^y), \langle \text{log query} \rangle_e]$ , where  $MAC_k(v_{id}||SP_{id}||g^x||g^y)$  is the message authentication code on verifier’s identity  $v_{id}$ , responder’s identity  $SP_{id}$ , and all public exponents exchanged so far, and  $\langle \text{log query} \rangle_e$  is the log query request protected by the freshly generated session key  $e$ . This step binds all public key exponents exchanged with the claimed identity holders together and marks the end of the authenticated session key exchange process. The service provider, then, retrieves the corresponding logs according to the log query request in the last message and sends the log encrypted using the session key  $e$  to the verifier.

**4.3.2 Verification of Logs.** This section presents procedures for log verification at the auditor and a user.

**Verification process at the auditor.** Recall that the auditor can verify any part of the sensor data. Suppose the auditor wishes to verify a chunk  $C_x$ ; see Figure 5. Hence, entire sensor data (the data written in the first box of

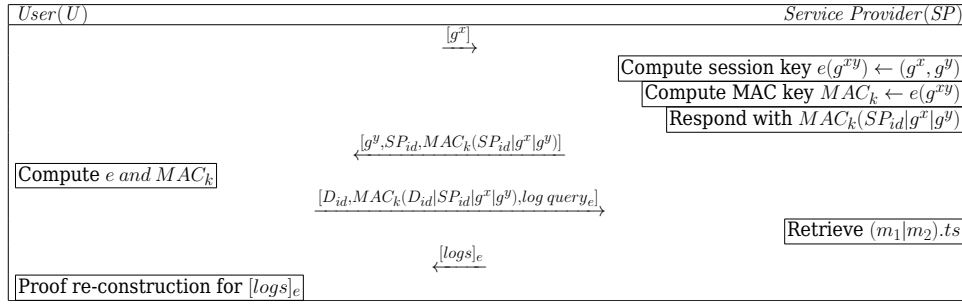


Fig. 6. Secure log retrieval by the verifier.

Figure 4) of the chunk  $C_x$ , random strings  $g^a$ ,  $g^b$ , and  $g^c$  (corresponding to the previous and next chunks of  $C_x$ ; see Figure 5), and proof-of-integrity  $\mathcal{PI}_{C_x}$  are provided to the auditor. The auditor performs the same operation as in PHASE 2 of §4.2.2. Also, the auditor computes the end-of-chunk string  $S_{eoc}^x = g^a \oplus g^b \oplus g^c$ . In the end, the auditor matches the results of  $h_n^x \oplus S_{eoc}^x$  against the decrypted value of received  $\mathcal{PI}_{C_x}$ , and if both the values are identical, then it shows that the entire chunk is unchanged.

Note that since SP transfers sensor readings of the chunk  $C_x$ , random strings ( $g^a$ ,  $g^b$ , and  $g^c$ ) and  $\mathcal{PI}_{C_x}$  to the user, SP can alter any transmitted data. However, SP cannot alter the signed  $Sign_{PR_E}(h_n^x \oplus S_{eoc}^x)$ , due to the unavailability of the private key of the enclave,  $PR_E$ , which was generated and provided by the trusted authority to the enclave. Thus, by following the above-mentioned procedure on the sensor readings of  $C_x$ , any inconsistency created by SP will be detected by the auditor.

**Verification process at the user.** If the user wishes to verify his data in a chunk, say  $C_x$ , the user is provided all hash digests computed over device-id and time ( $o_i$ , see the last box in Figure 4), time, sensor state, random strings  $g^a$ ,  $g^b$ , and  $g^c$  (see Figure 5), and the proof  $\mathcal{PU}$  by SP. Since, the user knows her device-id, first, the user verifies her occurrences in the data by computing the hash function on her device-id mixed with provided time values and compares against received hash digests. This confirms the user's presence/absence in the data. Also, to verify that no hash-digest is modified/deleted by SP, the user computes the hash function on the sensor state mixed with the received  $o_i$  ( $1 \leq i \leq n$ , where  $n$  is the number of sensor readings in  $C_x$ ) and computes  $hu_{end}^x = h_1^x \oplus h_2^x \oplus \dots \oplus h_n^x$ . Finally, the user computes  $hu_{end}^x \oplus S_{eoc}^x$  and compares against the decrypted value of  $\mathcal{PU}$ . The correctness of this method can be argued in a similar manner to the correctness of the verification at the auditor.

## 5 SYSTEM IMPLEMENTATION AND DEPLOYMENT

IoT NOTARY has been implemented and deployed as an integral part of TIPPERS that is an IoT testbed at the University of California Irvine, providing various real-time location-based services on the campus. The university IT department is a trusted infrastructure deployer in this testbed. At UCI, users need to register their devices' MAC addresses in a system shown in Figure 7a to get access to the campus-wide WiFi network service. When a user's device connects to the campus WiFi network infrastructure, an association event, including the user device MAC address, access point ID, and timestamp of the connection, is generated by the access point and collected by the IT department as part of the infrastructure maintenance tasks. The IoT NOTARY protocol is used when a service provider, e.g., a location-based service in TIPPERS, asks for the WiFi association data for service provisioning.

Before sharing data with a service provider, i.e., TIPPERS, the university IT department sends the notifications to the public through emails, website messages (Figure 7c), and signage (Figure 7b) attached to the buildings where the data is collected. The IT department generates a secret key and corresponding SGX enclave program at the setup phase for each service provider. The SGX program is given to the service provider and runs on its SGX-enabled system. In operation, the IT department encrypts the WiFi association data using the secret key and sends the encrypted data to the service provider's SGX enclave program. The IT department is also responsible for policy management in our testbed as the university's authority requires. A consent management website shown in Figure 7 is provided to users on campus to opt-out or opt-in their devices for data sharing with TIPPERS.

When the encrypted data arrives at TIPPERS, the SGX program decrypts the data in the secure enclave. It, further, enforces the policies by removing the data that contains opted-out devices. Then, tamper-proof logs are generated in the process and stored in a database. Before giving the data to TIPPERS' applications, the SGX program further



Home Services Support Information Security About OIT

OIT > UCInet Mobile Access > Registration

UCInet Mobile Access Registration

Summary: To use the UCInet Mobile Access network, you will need to register the hardware (MAC) address of your ethernet (for wired connections) or wireless card. This is a

**Manage Your MAC Addresses**

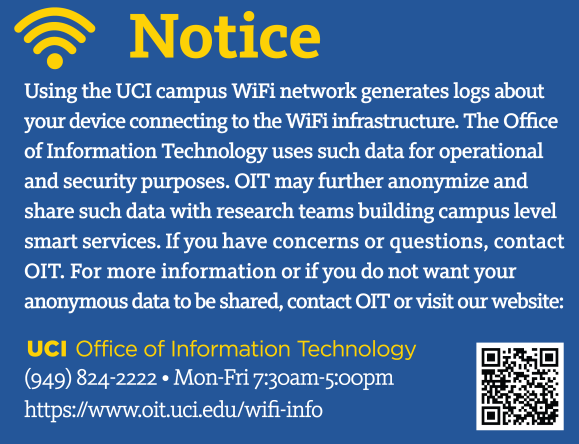
Add MAC Address  
Add Guest Registrations You are authorized to add time-limited guest registrations.

Register  
Connect  
Locations  
Help  
Security and More Info

MAC Addresses that are registered to your UCInetID

Registered at	Last seen	MAC Address
2019-10-07 12:29:35	2021-03-22 16:46:42	08:00:27:12:34:56
2019-10-07 14:35:19	2021-03-23 17:03:16	08:00:27:12:34:56
2019-10-07 14:35:37	2021-03-05 18:01:35	08:00:27:12:34:56
2019-10-07 14:35:46	2020-06-30 14:29:15	08:00:27:12:34:56
2019-10-07 14:35:54	2020-02-02 18:23:11	08:00:27:12:34:56


(a) System for registering device's MAC address to get WiFi access on campus



**Notice**

Using the UCI campus WiFi network generates logs about your device connecting to the WiFi infrastructure. The Office of Information Technology uses such data for operational and security purposes. OIT may further anonymize and share such data with research teams building campus level smart services. If you have concerns or questions, contact OIT. For more information or if you do not want your anonymous data to be shared, contact OIT or visit our website:

UCI Office of Information Technology  
(949) 824-2222 • Mon-Fri 7:30am-5:00pm  
<https://www.oit.uci.edu/wifi-info>



(b) The data-capturing notification signage

Undergraduates Graduate Students Staff Faculty Researchers IT Profession

UCI Office of Information Technology

ENHANCED BY Google

Home Services Support Information Security About OIT

For the latest information regarding technology support for the Coronavirus situation, please visit [UCI TechPrsp](#).

You are here: Home / UCInet Mobile Access (WiFi) / WiFi Connection Data Sharing

**WiFi Connection Data Sharing**

UCI's Office of Information Technology (OIT) provides WiFi network service to the campus. When a device (smartphone, laptop, etc.) connects to WiFi, OIT collects data about the device including its identity, WiFi access point, data, and time. This data is used for operational purposes. Only connection data is saved; no information transmitted over the network is collected or shared. The University of California Electronic Communications Policy protects your privacy.

**Data Sharing**

Beginning in May 2019 anonymized WiFi connection data is being shared with researchers. The anonymized data is not personally identifiable. You can opt out of this sharing. No information transmitted over the connection is collected or shared.

(c) The data-capturing notification on IT department website

UCI University of California, Irvine

Hello [User Name] Logout

Using the UCI campus WiFi network generates logs about your device connecting to the WiFi infrastructure. OIT uses such data for operational and security purposes. OIT may further anonymize and share such data with research teams building campus level smart services. You may use this application to opt-out all or some of your devices from sharing the anonymized WiFi connection information. Please visit the [OIT WiFi Security Page](#) for more information.

**Registered Devices for UCInet Mobile Access:**

In order for you to get Internet access from UCInet Mobile Access, you will need to register the hardware (MAC) addresses of your mobile devices. Please visit the [UCInet Mobile Access - Manual Registration](#) page to manage your registered devices

**Opt-out Devices:**

No WiFi connection information generated by the devices in opt-out devices list will be shared. You can add/remove the devices to/from opt-out list using this application

**Your Registered Devices for UCInet Mobile Access**

MAC Address: [MAC Address]	Opted Out
MAC Address: [MAC Address]	Add

(d) System for users to specify their policies (opt-out)

Fig. 7. The campus deployment of IoT NOTARY

de-identifies the MAC address. Then, TIPPERS uses the data and provides valuable location-based services, such as occupancy analysis. IoT NOTARY via TIPPERS provides APIs for the auditor and the users to check the presence/absence of their data by retrieving the desired logs for a given duration and an application for executing the attestation module at their-ends.

## 6 EXPERIMENTAL EVALUATION

This section presents our experimental results on live WiFi data. We execute IoT NOTARY on a 4-core 16GB RAM machine equipped with SGX at Microsoft Azure cloud.

**Setup.** In our setup, the IT department at UCI is the trusted infrastructure deployer. It also plays the role of the trusted notifier (notifying users over emailing lists). At UCI, 490 WiFi sensors, installed over 30 buildings, send data to a controller that forwards data to the cloud server, where IoT NOTARY is installed. The cloud keeps cryptographic log digests that are transmitted to the verifier, while sensor data, qualifies data-capture rules, is ingested into real-time applications supported by TIPPERS. We use SHA-256 as the hashing algorithm and 256-bit length random strings in IoT NOTARY. We allow users to verify the data collected over the last 30minutes (min).

**Dataset size.** Although IoT NOTARY deals with live WiFi data, we report results for data processed by the system over 180 days during which time IoT NOTARY processed 13GB of WiFi data having 110 million WiFi events.

**Data-capture rules.** We set the following four data-capture rules: (i) *Time-based*: always retain data, except from  $t_i$  to  $t_j$  each day; (ii) *User-location-based*: do not store data about specified devices if they are in a specific building;

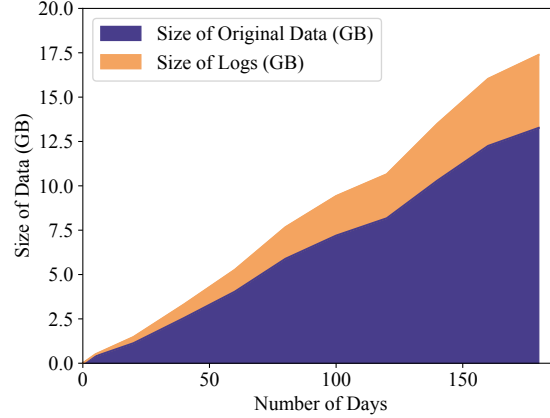


Fig. 8. Exp 1: Storage overhead.

(iii) *User-time-based*: do not capture data having a specific device-id from  $t_x$  to  $t_y$  ( $x \neq i, y \neq j$ ) each day; and (iv) *Time-location-based*: do not store any data from a specific building from time  $t_x$  to  $t_y$  each day. The validity of these rules was 40 days. After each 40-days, variables  $i, j, x, y$  were changed.

**Exp 1. Storage overhead at the cloud.** We fix the size of each chunk to 5MB, and on average, each of them contains  $\approx 37K$  sensor readings, covering around 30min data of 30 buildings in peak hours. Based on the 5MB chunk size, we got 3291 chunks for 180 days. For each chunk, the sealing function generates two types of logs: (i) for auditor verification that produced proof-of-integrity  $\mathcal{PI}$  of size 512bytes, and (ii) for user verification that produces hashed values (see Figure 4) and proof-of-integrity for users  $\mathcal{PU}$  of size 1.05MB. Figure 8 shows 180-days WiFi data size without having sealed logs (dark blue color) and with sealed logs (light brown color).<sup>4</sup>

**Exp 2. Performance at the cloud.** For each 5MB chunk, the sealing function took around 310ms to seal each chunk. This includes time to compute  $\mathcal{PI}$ ,  $\mathcal{PU}$  and encrypt them.

**Exp 3. Auditor verification time.** The auditor at our campus has a 7th-Gen quad-core i7CPU and 16GB RAM machine. It downloads the chunks from the cloud and executes auditor verification. The auditor varied the number of chunks from 1 to 3000; see Table 2. Note that to attest one-day data across 30 buildings, the auditor needs to download at most 50 chunks, which took less than 1min to verify. Observe that as the number of chunks increases, the time also increases, due to executing the hash function on more data.

Number of Chunks	1	50	100	500	1000	3000
$\approx$ duration (day)	30-60min	1-2	2-5	8-18	35-55	175
Verification time (seconds)	1	49	102	544	1160	4400

Table 2. Exp 3: The auditor verification time. Duration varies due to different class schedules in buildings and working hours.

**Exp 4: Verification at a resource-constrained user.** To show the practicality of IoT NOTARY for resource-constrained users, we considered four types of users, differing on computational capabilities (*e.g.*, available main memory (1GB/2GB) and the number of cores (1 or 2 cores)). Each user verified 1/10/20-days data; see Figure 9. Note that verifying 1-day data, which is  $\approx 50$  blocks, at resource-constrained users took at most 30s. As the number of blocks increases, the computational time also increases, where the maximum computational time to verify 20-days data was  $< 10$ min. As the days increase, so does data transmitted to the user, which spills over to disk causing an increased latency. Also, we execute the same experiment on a powerful user having 4-core and 16GB machine. Note that as the number of core and memory increases, it results in parallel processing and the absence of disk data read. Thus, the computation time decreases (see user 5 in Figure 9).

<sup>4</sup>The reason for getting more chunks is that during non-peak hours a 5MB chunk can store sensor readings for more than one hour. However, as per our assumption, we allow the user to verify the data collected over the last 30min. Hence, regardless of the chunk is full or not, we compute the sealing function on each chunk after 30min.




Fig. 9. Exp 4: Verification time.

Chunk Size (KB)	5	100	1,000	5,000	10,000	25,000	40,000
Time (millisecond)	30	38	120	246	579	1,789	3,389

Table 3. Exp 5: Impact of the chunk size on log sealing execution.




Fig. 10. Exp 5: A tradeoff between log size and latency.

**Exp 5. Impact of the chunk size.** We have selected chunk size to be 5MB that can hold at most 30min WiFi data. Now, we investigate the impact of chunk size on the sealing function execution time, the verification time, and latency in obtaining the most recent data. Table 3 shows that as the chunk size increases, the chunk holds more data, and hence, executing the sealing function on large-sized data takes more time. Similarly, a large-sized chunk also increases the latency in obtaining the most recent data (see Figure 10), since unless filling the chunk, the enclave cannot produce the proof-of-integrity. We can also create a chunk having only a single row; however, it will increase the size of secured logs (see Figure 10). Observe that (in Figure 10), when the chunk size is 5KB, the secured log size is  $\approx 7$ GB, while when the chunk size is 5MB, the secured log size is  $\approx 4$ GB. Figure 11 shows verification time also increases as the chunk size increases.

**Exp 6: Impact of communication.** We measured the communication impact when a verifier downloaded the sensor data and/or sealed log for attestation. Consider a case when the verifier attests only one-hour/one-day data. The average size of one-hour (one-day) data in a peak hour was 14MB (250MB) having 103K (1.2M) connectivity events, while in a non-peak hour, it was 2MB (50MB) having 13.5K (320K) connectivity events. When using slow (100MB/s),

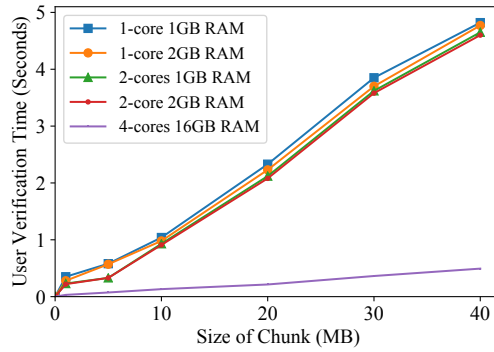


Fig. 11. Exp 5: User verification time with different log chunk sizes.




Fig. 12. Exp 8: Size of the secured log when using different log sealing methods.

medium (500MB/s), and fast (1GB/s) speed of data transmission, the data transmission time in the case of 1-hour/1-day data was negligible.

**Exp 7: Impact of parallelism.** The processing time at each server can be reduced by parallelizing the computation. We investigated the impact of parallel processing to seal a 5MB chunk when having the number of threads 2 or 4 that took 322ms and 310ms, respectively. Increasing more threads did not provide speed-up, since the execution time increases due to thread maintenance overheads. Note that we only parallelized the hash function computation for  $\mathcal{PU}$ , (while  $\mathcal{PI}$  cannot be computed in parallel, due to the formation of hash chains).

**Exp 8. Impact of log optimization.** We compare the impact of per sensor- and per-user-device-id-based optimization methods (denoted by OPT-SENSOR and OPT-USER in Figure 11, given in §4.2.4) against the non-optimized method (denoted by NON-OPT in Figure 11, given in §4.2.1). Since there were 490 WiFi access points in our experiments, we pre-allocated 490 buffers in SGX memory for 490 sensors, one buffer for each sensor. As a result, each buffer size was  $\approx 85$ KB. We also created buffers for groups of devices to implement OPT-USER. Here, we again created 490 buffers, and an identical group of user devices is allocated to a buffer. Particularly, the enclave extracted the user device’s MAC address, hashed to get the buffer identity, and placed the sensor reading to the corresponding buffer. When the buffer got full, the data in the buffer is cryptographically sealed and written on the disk. Here, we compare storage overhead and verification time.

Figure 11 shows the storage overhead to store only cryptographically-secured logs, when using OPT-SENSOR, OPT-USER, and NON-OPT methods. Since OPT-SENSOR and OPT-USER methods produce chunks of size 85KB, for a fair comparison, we set the log chunk size to be 85KB in NON-OPT method. Figure 11 shows that OPT-SENSOR and OPT-USER save  $\approx 6.5\%$  and  $\approx 2.1\%$  space, respectively, compared to NON-OPT method.

We also investigated the benefit in the performance improvement of user data verification, when using OPT-SENSOR, OPT-USER, and NON-OPT methods. Note that when using OPT-SENSOR and NON-OPT methods, a chunk may store data

associated with other users. Thus, it needs to verify additional data, which does not belong to the user. In contrast, when using OPT-USER methods, a user has to verify only the desired data that belongs to him/her. It is clear that OPT-USER method requires to verify less amount of data, and hence, less verification time, compared to OPT-SENSOR and NON-OPT methods. Table 4 shows verification time and number of chunks required to verify one-day data at a resource-constrained user (1-core 1GB RAM). Observe that OPT-USER method takes significantly less time compared to NON-OPT method.

Method	# chunks	Verification time
NON-OPT	3012	31.2s
OPT-SENSOR	86	0.89s
OPT-USER	57	0.71s

Table 4. Exp 8: User verification performance when using different methods.

**Summary of the experiments.** We conduct extensive experiments on a real-world setting to evaluate IoT NOTARY’s overhead, performance, and impact of critical parameters. The results can be interpreted and presented from two principal stakeholders’ perspectives in the proposed framework: the service provider and users. On the service provider side, deploying IoT NOTARY incurs a little overhead. In a default setting, only 21% of additional storage is needed to retain the tamper-proof logs (Exp 1), and Enforcing policy and generating logs processes for 30-minute data take sub-second computation time (Exp 2). The service provider needs to select the appropriate parameters for an efficient deployment according to its requirements. For example, the chunk size represents the trade-off between the latency for verification and log storage size (Exp 5). Besides, according to different system configurations and data patterns, the service provider could further improve the performance and reduce the overhead by increasing parallelism or incorporating optimized logging algorithms (Exp 7 and Exp 8). On the auditor or users’ side, their interactions with IoT NOTARY is for log verification, and IoT NOTARY’s verification algorithms allow both the auditor and users to verify data efficiently on different device configurations without incurring computational overhead (Exp 3 and Exp 4).

## 7 COMPARISON WITH EXISTING WORK

We classify the related work in the scope of IoT attestation into the following four categories: (i) the work that verifies the internal memory of sensors to verify that sensors have not tampered, (ii) the work that verifies the data integrity, (iii) the work that focuses on privacy-preserving data capturing and data release, and (iv) the work that performs auditing in decentralized storage networks.

**Attestation of the state of the device in IoT.** In IoT settings, several techniques have been developed to verify the internal memory state of untrusted devices through a trusted remote verifier. (Note that such a type of verification is different from the data verification done by IoT NOTARY.) For example, AID [26] attests the internal state of neighboring devices through key exchange and proofs-of-non-absence. In AID, the adversary can compromise communication channels, the internal state of the device, and the cryptographic keys. SEDA [6] attests low-end embedded devices in a swarm and provides the number of devices that pass attestation. However, SEDA attests neighboring peer devices only. Similarly, DARPA [25] and SANA [5] allow detection of physical attacks by using heartbeat messages and provide aggregate network attestation, with high computation and communication costs, which are quadratic in the network size. SMARM [11] protects against malware by scanning memory in a secret randomized order. However, it may require many iterations to eventually detect malware. RADIS [13] assumes that a compromised IoT service can influence genuine operations of other services without changing the software. Thus, RADIS provides control-flow attestation for distributed services among interconnected IoT devices. [43] provided SGX-based attestation method for physical attacks on the sensor, e.g., modifying memory and changing I/O signals. In short, all such work only deals with attestation of sensor devices, and their methods cannot be used to verify sensor data against the data-capture rules.

In contrast, IoT NOTARY does not deal with the verification of the internal state of sensor devices, since in our case, (WiFi access-point) sensors are assumed to be deployed by a trusted entity (e.g., the university IT department). Of course, cyberattacks are possible on sensors to maliciously record data and that can also be detected by IoT NOTARY, while not verifying the sensors.

**Data integrity verification.** To verify data integrity, several protocols based on zero-knowledge proofs, multi-party computation (MPC), and Merkle trees have been proposed. For example, [21, 23] proposed a privacy-preserving scheme based on zero-knowledge proofs to detect log-exclusion attacks. [23] provided solutions for accountability and auditing through hierarchical multi-party computation (MPC) and succinct zero-knowledge proof statements. [27] considered verification process for MPC using a trusted third-verifier. [21] provided a privacy-preserving certificate transparency service, which signs a message four times, where an auditor can trace the certified logs. Other techniques have proposed the concept of *excerpts* and *snapshots* for log integrity verification. For example, in [45] used hash-chains for integrity protection and identity-based searchable encryption. [16] proposed a Merkle tree-based history tree to prove the sequence of logs over time. [49] proposed a Bloom tree that stores proof of logs at an untrusted cloud. Further, access-pattern-hiding cryptographic techniques [10, 50] may be used to verify any stored log, since an adversary cannot skip the computation on some parts of the data, due to executing an identical computation on each sensor reading. Techniques, *e.g.*, function secret-sharing [10] or vSQL [50], may be used to verify query results on cleartext. However, these techniques cannot detect log deletion. Also, all such techniques incur significant time. For example, vSQL took more than 4000 seconds to verify a SQL query. In addition, any end-to-end encryption model is not sufficient for verification.

In contrast, IoT NOTARY provides complete security to sensor data and a real-time data attestation approach. Unlike [21], IoT NOTARY requires only two signatures per file, where one is used to validate log completeness, and another is used to validate the log ordering with respect to adjacent logs.

**IoT data privacy policy enforcement.** Several works propose solutions for data privacy policy enforcement before the data is released for any service provisioning. From an architectural perspective, these solutions could be coarsely classified into two categories: *middle-box methods* and *on-device methods*, depending on where the policy enforcement functions are placed. The *middle-box methods* import a new trusted and commonly centralized service that aggregates and processes the data according to the privacy policies before the data is released to any service provider. Personal data stores (PDS) [18, 32] provide a general rule-based access management framework for sensitive personal data, including mobile sensor data. Besides access control functions, the model that [15] introduced also helps the data controller and service providers demonstrating compliance with privacy regulations. [17] combines policy enforcement with edge computing and proposes a solution where the policy enforcement functions could be deployed flexibly at a location in a user's trusted domain. The middle-box methods have been adopted for processing health data [9, 40, 47], location data [8, 33], and video data [42].

In contrast to introducing an additional trusted entity in the system, the *on-device methods* aim to enforce the privacy policies at the data source directly. Achieving this goal requires more computational power on the sensing device. Thus, on-device methods are generally designed for a mobile computing setting, where a smartphone works as a sensing device. [12] and [46] propose general frameworks for privacy and context-aware sensor management on smartphones for users to define their sensor data access policy in a fine-grain manner. In addition, the pervasive photograph-capturing capability on mobile devices causes privacy concerns nowadays [37]. Several solutions have been proposed to ensure data capturing policies for video feeds on mobile devices, including smartphones [3, 39], wearable devices [48], and augmented reality (AR) devices [38].

Unlike existing middle-box solutions, IoT NOTARY minimizes infrastructure deployer (IFD)'s overhead. It does not require the IFD to implement the data capturing policy enforcement and logging functions in its existing infrastructure. Instead, the service provider (SP) who eventually uses the data will pay overhead. In our experience, this model is critical for the agreement of collecting data from IFDs, because infrastructure and organizations, *e.g.*, network infrastructure, water infrastructure, and building management, may not have extra resources for additional services other than their regular workload. Additionally, IoT NOTARY's model also leads to better scalability of supporting more SPs. When a new SP requires data from IFD, IFD only needs to spawn the trusted agent program for the new SP initially. In operation, the communication overhead of sending the encrypted data to the new SP is the only additional overhead. On the other hand, comparing with on-device solutions, IoT NOTARY supports existing commodity devices without modifications. This advantage is vital for supporting heterogeneous devices in an IoT scenario.

**Auditing in decentralized storage networks.** Decentralized storage networks (DSN) refer to a group of untrusted nodes (or people) at multiple locations who are incentivized to join the network, store the data of other users, and keep data accessible. Since the data is stored at multiple untrusted nodes in DSN, several data auditing schemes such as centralized private auditing, Merkle tree-based auditing, succinct proof framework, SNARK frameworks, data

tagging, and Pedersen commitment-based zkLedger [2, 7, 20, 34, 35, 41] have been proposed. In contrast, IoT NOTARY provides a central data repository and mechanisms for data attestation by several users.

## 8 CONCLUSION

This paper presented a framework, IoT NOTARY for sensor data attestation that embodies cryptographically enforced log-sealing mechanisms to produce immutable proofs, used for log verification. Our solution improves the naïve end-to-end encryption model, where retroactive verification is not provable. The service verification mechanism on failing at users allows them to revoke services of the concerned IoT space. Therefore, a user is not required to blindly trust in the IoT space, and we empower the users with the right-to-audit instead of right-to-own the data captured by sensors. IoT NOTARY is a part of a real IoT system (TIPPERS) and provides verification on live WiFi data with almost no overheads on users.

At the high level, IoT NOTARY works in three phases, as follows: (i) Notification, in which the system informs the users about the data capturing policies established by the organizations; (ii) Log sealing, in which the system stores the received sensor data in a secure manner using hash chains and produces the proof; and (iii) Verification, in which the user verifies the presence/absence of their data against the data capturing policies. We, future, provided two techniques to optimize the secure logging sealing process, by creating per sensor-based logging and per user-based logging. The experimental results show the effectiveness of IoT NOTARY in a real IoT environment, consisting of 490 WiFi access-points across 30 buildings of the University of California, Irvine. The experiments show that the overheads of IoT NOTARY in creating the secure log, storing the log, as well, in verifying secure logs at the resource-constrained users are affordable.

**Future directions.** Now, we provide the future directions that one can explore and improve IoT NOTARY, as follows:

- (1) **Contextual policies.** While we presented IoT NOTARY for the case of monotonous data capturing policies, which are applicable to all the users, an interesting direction would be in exploring the contextual data capturing policies, where the context is determined based on data captured by other sensors. Extending the system to handle such data-capture rules, e.g., “Do not capture my WiFi connectivity data if I am the only person connected to the access point,” is a non-trivial challenge.
- (2) **User registration.** IoT NOTARY currently works in an organization that maintains a user-device registry for their authentication. However, such a registry may not exist in spaces such as shopping malls or stadiums. A potential improvement in IoT NOTARY addressing this issue could be authorizing verification query based on device identity and only allowing a device to verify its own data. However, ensuring the integrity of the device identifier remains a challenge.
- (3) **Device spoofing.** The current IoT NOTARY implementation for wireless network data utilizes the MAC address as the device identifier to enforce data capturing policies. However, the method may be vulnerable to identifier spoofing attacks. It is important to design or utilize an appropriate device authentication mechanism in the network and use more secure device identifiers.
- (4) **Sensor spoofing.** IoT NOTARY’s enforcement of data-capturing rules relies on the assumption that the service provider cannot tamper with the deployment of sensors. For example, swapping sensors in two rooms may cause the data capturing inconsistent with the published rules. Thus, another interesting direction of work would be in enhancing IoT NOTARY for ensuring trust on sensor deployment. Potential solutions, including building sensor registry and verifying sensors at the network level, pose new challenges.
- (5) **A very-large-scale deployment.** The current IoT NOTARY system works in a campus-scale deployment, where hundreds to thousands of sensors exist. It is interesting to explore how the system could be extensively deployed in a city-scale scenario, where a massive amount of sensors generate data. It will require IoT NOTARY to processing data obtained from multiple trusted agents, in parallel. This improvement proposes new challenges like load balancing and data-capturing rules consistency at a wide coverage area.
- (6) **Extending IoT NOTARY for decentralized storage networks.** The current IoT NOTARY system works in a centralized storage system, and extending IoT NOTARY for DSN comes with challenges such as: how to inform users where their data has been stored, how the users will communicate with different nodes in DSN, what will happen if data is stored according to per sensor or per user, and what will be the overhead of using DSN compared to the centralized system.

## REFERENCES

- [1] <https://newsroom.intel.com/newsroom/wp-content/uploads/sites/11/2017/09/8th-gen-intel-core-product-brief.pdf>.
- [2] Filecoin: A Decentralized Storage Network by Protocol Labs, available at: <https://filecoin.io/filecoin.pdf>.
- [3] P. Aditya et al. I-Pic: A platform for privacy-compliant image capture. In *MobiSys*, page 235–248, 2016.
- [4] S. K. Aikins. Connectivity of smart devices: Addressing the security challenges of the internet of things. In *Connectivity Frameworks for Smart Devices: The Internet of Things from a Distributed Computing Perspective*, 2016.
- [5] M. Ambrosin et al. SANA: secure and scalable aggregate network attestation. In *CCS*, pages 731–742, 2016.
- [6] N. Asokan et al. Seda: Scalable embedded device attestation. In *CCS*, pages 964–975, 2015.
- [7] E. Ben-Sasson et al. Snarks for C: verifying program executions succinctly and in zero knowledge. In *CRYPTO*, volume 8043, 2013.
- [8] A. R. Beresford et al. Mix zones: User privacy in location-aware services. In *IEEE Annual Conference on Pervasive Computing and Communications Workshops*, pages 127–131, 2004.
- [9] R. Boussada et al. Privacy-preserving aware data transmission for iot-based e-health. *Comput. Networks*, 162, 2019.
- [10] E. Boyle et al. Function secret sharing. In *EUROCRYPT*, 2015.
- [11] X. Carpent et al. Remote attestation of iot devices via smarm: Shuffled measurements against roving malware. In *HOST*, pages 9–16, 2018.
- [12] S. Chakraborty et al. ipShield: A framework for enforcing context-aware privacy. In *NSDI*, pages 143–156, 2014.
- [13] M. Conti et al. Radis: Remote attestation of distributed iot services. In *Sixth International Conference on Software Defined Systems (SDS)*, pages 25–32, 2019.
- [14] V. Costan et al. Intel SGX explained. *IACR Cryptology ePrint Archive*, 2016:86, 2016.
- [15] A. Crabtree et al. Building accountability into the Internet of Things: the IoT databox model. *J. Reliab. Intell. Environ.*, 4(1):39–55, 2018.
- [16] S. A. Crosby et al. Efficient data structures for tamper-evident logging. In *USENIX*, pages 317–334, 2009.
- [17] N. Davies et al. Privacy mediators: Helping iot cross the chasm. In *HotMobile*, page 39–44, 2016.
- [18] Y.-A. De Montjoye et al. openpds: Protecting the privacy of metadata through safeanswers. *PLoS one*, 9(7):e98790, 2014.
- [19] W. Diffie et al. New directions in cryptography. *IEEE Trans. Information Theory*, 22(6):644–654, 1976.
- [20] Y. Du et al. Enabling secure and efficient decentralized storage auditing with blockchain. *IEEE Transactions on Dependable and Secure Computing*, 2021.
- [21] S. Eskandarian et al. Certificate transparency with privacy. *PoPETs*, 2017(4):329–344, 2017.
- [22] E. Fernandes et al. Security implications of permission models in smart-home application frameworks. *IEEE Security & Privacy*, 15(2):24–30, 2017.
- [23] J. Frankle et al. Practical accountability of secret processes. In *USENIX*, pages 657–674, 2018.
- [24] J. Gubbi et al. Internet of things (iot): A vision, architectural elements, and future directions. *FGCS*, (7):1645–1660, 2013.
- [25] A. Ibrahim et al. Darpa: Device attestation resilient to physical attacks. In *WiSec*, pages 171–182, 2016.
- [26] A. Ibrahim et al. AID: autonomous attestation of IoT devices. In *SRDS*, 2018.
- [27] W. Jiang et al. Transforming semi-honest protocols to ensure accountability. *Data Knowl. Eng.*, 65(1):57–74, 2008.
- [28] S. L. Keoh et al. Securing the internet of things: A standardization perspective. *Internet of Things Journal*, 1(3):265–275, 2014.
- [29] H. Krawczyk. SIGMA: The ‘SIGn-and-MAC’ approach to authenticated diffie-hellman and its use in the IKE protocols. In *CRYPTO*, 2003.
- [30] S. Madakam et al. Security mechanisms for connectivity of smart devices in the internet of things. 2016.
- [31] S. Mehrotra et al. TIPPERS: A privacy cognizant IoT environment. In *PerCom Workshops*, pages 1–6, 2016. <http://tippersweb.ics.uci.edu/>.
- [32] M. Y. Mun et al. Personal data vaults: a locus of control for personal data streams. In *CoNEXT*, page 17, 2010.
- [33] G. Myles et al. Preserving privacy in environments with location-based applications. *IEEE Pervasive Computing*, 2(1):56–64, 2003.
- [34] N. Narula et al. zkledger: Privacy-preserving auditing for distributed ledgers. In *NSDI*, pages 65–80, 2018.
- [35] B. Parno et al. Pinocchio: Nearly practical verifiable computation. In *IEEE SP*, pages 238–252.
- [36] A. Rao et al. Expecting the unexpected: Understanding mismatched privacy expectations online. In *SOUPS*, pages 77–96, 2016.
- [37] Y. Rashidi et al. "you don't want to be the next meme": College students' workarounds to manage privacy in the era of pervasive photography. In *SOUPS*, pages 143–157.
- [38] F. Roesner et al. World-driven access control for continuous sensing. In *CCS*, pages 1169–1181, 2014.
- [39] J. Shu et al. Cardea: context-aware visual privacy protection for photo taking and sharing. In *MMSys*, pages 304–315, 2018.
- [40] W. Tang et al. Secure data aggregation of lightweight e-healthcare IoT devices with fair incentives. *IEEE Internet Things J.*, 6(5):8714–8726, 2019.
- [41] D. Vorick et al. Sia: Simple decentralized storage. 2014.
- [42] J. Wang et al. A scalable and privacy-aware iot service for live video analytics. In *MMSys*, pages 38–49. ACM, 2017.
- [43] J. Wang et al. Enabling security-enhanced attestation with Intel SGX for remote terminal and iot. *TCDICS*, 37(1):88–96, 2018.
- [44] W. Wang et al. Leaky cauldron on the dark land: Understanding memory side-channel hazards in SGX. In *CCS*, pages 2421–2434, 2017.
- [45] B. R. Waters et al. Building an encrypted and searchable audit log. In *NDSS*, volume 4, pages 5–6, 2004.
- [46] Z. Xu et al. SemaDroid: A privacy-aware sensor management framework for smartphones. In *CODASPY*, page 61–72, 2015.
- [47] Y. Yang et al. Privacy-preserving fusion of IoT and big data for e-health. *Future Generation Computer Systems*, 86:1437–1455, 2018.
- [48] E. Zarepour et al. A context-based privacy preserving framework for wearable visual lifeloggers. In *PerCom Workshops*, 2016.
- [49] S. Zawoad et al. Towards building forensics enabled cloud through secure logging-as-a-service. *IEEE TDSC*, 13:148–162, 2016.
- [50] Y. Zhang et al. vSQL: Verifying arbitrary SQL queries over dynamic outsourced databases. In *IEEE SP*, pages 863–880, 2017.