# Privacy-Preserving Protocols for Smart Cameras and Other IoT Devices

YOHAN BEUGIN, QUINN BURKE, BLAINE HOAK, RYAN SHEATSLEY, and ERIC PAULEY, University of Wisconsin-Madison, USA

GANG TAN and SYED RAFIUL HUSSAIN, The Pennsylvania State University, USA

PATRICK MCDANIEL, University of Wisconsin-Madison, USA

Millions of consumers depend on smart camera systems to remotely monitor their homes and businesses. However, the architecture and design of popular commercial systems require users to relinquish control of their data to untrusted third parties, such as service providers (e.g., the cloud). Third parties therefore can (and in some instances have) access the video footage without the users' knowledge or consent—violating the core tenet of user privacy. In this paper, we introduce CaCTUs, a privacy-preserving smart camera system that *returns control to the user*; the root of trust begins with the user and is maintained through a series of cryptographic protocols designed to support popular features, such as sharing, deleting, and viewing videos live. In so doing, we demonstrate that it is feasible to implement a performant smart-camera system that leverages the convenience of a cloud-based model while retaining the ability to control access to (private) data. We then discuss how our techniques and protocols can also be extended to privacy-preserving designs of other IoT devices recording time series data.

## 1 INTRODUCTION

Smart camera systems are changing the way consumers secure their homes and businesses. Commercial camera systems have been remarkably successful; they have become the *de facto* monitoring system, as they offer the following essential services with plug-and-play support: (1) watch live and recorded video feeds, (2) share videos with others, (3) delete recorded videos, (4) recover access to the system, and (5) perform a full factory reset. Yet, while the market demand for smart camera systems continues to grow rapidly as reported by *Ring* [12, 26, 27, 59], *Wyze* [60], and *Arlo* [4], consumers have come to realize that the costs of owning a smart camera system are not exclusively monetary.

Commercially available smart camera systems follow a threat model that mandates undue trust *by design*; the service provider is granted unfettered access to the video content of any consumer who uses their system. Ring has come under legal scrutiny [21, 24] for allowing more than 2,000 government agencies to directly request videos from users without formal due process [9, 34, 46, 49]. Perhaps even more troubling, employees are viewing and annotating live user streams

for research [6, 15] while others are abusing their access to view and share users' videos online [29, 53]. Moreover, research has posited that these systems can be transformed into mass surveillance systems, given their widespread adoption [10–12, 32, 34, 38, 39, 48]. The message behind the underlying design of modern smart camera systems is clear: users do not have control over their own videos and system, compromising user privacy.

To this end, we answer the following question: *can users afford all of the features present in commercial smart camera systems, without compromising their privacy?* A cryptographic approach is a plausible way to protect users' privacy in that it enables users to solely assume control over videos stored in the cloud. However, practical realizations of such systems face several key challenges: cryptographic protections (e.g., encryption) incur computational overheads, affecting system performance, since stored videos are encrypted; a fine-grained sharing scheme (i.e., sharing specific video fragments) requires a user-controlled key management system; and generating, storing, rotating, and re-negotiating cryptographic keys poses further challenges on performance and usability. Troublingly, even if such challenges could be addressed, events have demonstrated that encryption alone is not sufficient to protect users from abuses by governments through coercion [23, 33]. Thus, meeting performance, security, privacy, and usability goals demands a novel approach that is sensitive to the unique requirements of this domain.

In this paper, we present CaCTUs, a privacy-preserving smart **Ca**mera system **C**ontrolled **T**otally by **Us**ers. Inspired by information privacy laws, CaCTUs is designed to enforce three privacy goals through known security properties: (1) *the right to not be seen*: the user is assured *confidentiality* of stored videos and live video streams; (2) *the right of sole ownership*: the user (and only the user) is trusted, and has *complete mediation* over access to their data by others; (3) *the right to be forgotten*: deleted videos are not recoverable, even in cases of coercion. Moreover, we show that the protocols involved in CaCTUs can be generalized to other IoT devices.

To meet the required feature set of commercial systems and address the stringent technical challenges and privacy goals of smart camera systems, we design CaCTUs as follows: it allows the user to solely assume control of the smart camera system through a direct and physical pairing process (that is, without relying on or trusting third parties); isolates and protects access to video footage through encryption, key rotation, and key management; enables viewing live and stored videos through performance-aware cryptographic algorithms; supports video deletion and factory reset via key rotation and management; and provides fine-grained (i.e., on the scale of seconds) peer-to-peer delegation of video footage through a binary key tree. We make the following contributions:

(1) We present CaCTUs, a privacy-preserving smart **Ca**mera system **C**ontrolled **T**otally by **Us**ers, that returns controls of the system to users without compromising features found in commercial smart camera systems.

(2) We perform a performance and functional user evaluations of our system and find that CaCTUs can serve a live video stream at a resolution of 480 p, at a frame rate of 10 fps, with a latency of 2 s and that it is natural and easy to use, all while meeting our privacy goals.

(3) We demonstrate how the protocols used in CaCTUs can be adapted to a broader class of devices.

We release CaCTUs as an open-source artifact for the development of future privacy-preserving IoT devices[1].

## 2 BACKGROUND

### 2.1 Smart Camera Systems

A smart camera system is a collection of cameras that are connected to the Internet, allowing *owners* (i.e., those who purchase and configure the system) to view live and recorded videos of their homes from anywhere. Most companies

---

[1]Available at https://github.com/siis/CaCTUs.

sell their systems as an integrated ecosystem: cameras work with a purpose-built smartphone application that allows the owner to view footage, delegate access, and administer their smart camera system. At the core of these systems are five functions: (1) recording and streaming, (2) sharing (delegation), (3) deleting, (4) access recovery, and (5) factory reset. Each function places requirements and motivates the architecture of the ecosystems available to consumers.

*Recording and Streaming.* Camera systems allow owners to view live and recorded footage from all cameras they own using a smartphone application, allowing them to monitor the current status of their property. As the most fundamental function provided by camera systems, this is expected to work reliably and globally: users want to be able to view footage anywhere, and recover footage even in the case of physical failures of the camera or home Internet connection. To facilitate this, consumer smart camera systems currently entrust the data to a cloud provider, streaming camera data to cloud storage as it is captured and making it available to the owner's device. As a result, access to the footage is managed by the cloud provider, who must be trusted to prevent unauthorized access.

*Delegation.* Owners want to share access to their camera systems with others. We refer to this capability as *delegation*. Whether used to provide a house-sitter with access to live footage during a vacation, or sharing video of an incident after the fact, this delegation is expected to be *fine-grained*, meaning it applies to specific users (*delegatees*) for only the portion of time that they need access. Consumer smart camera systems allow policy enforcement as a means of delegation: each user has an attached policy for the time range of live or recorded footage they may access, and cloud storage mediates this to prevent delegatees from exceeding their policies.

*Deletion.* Owners expect the ability to fully delete their data to prevent further access by any party. As recorded data from camera systems are saved to the cloud, owners must trust cloud providers to delete their data when requested, including copies stored elsewhere in the cloud.

*Access Recovery.* Since access to smart camera systems is mediated by a set of credentials (e.g., a username and password), these systems must account for the possibility of a user losing them. When authentication is performed by a cloud service, this is relatively straightforward: user's identity is verified via other means, such as a password reset through email. As we will discuss, however, such recovery is only trivial because of the trust assumptions of these systems, and we will see that this critical function requires careful thought under other trust models.

*Reset.* Finally, owners may wish to stop use of the smart camera system. They will expect all of their stored footage to be deleted, access to live footage revoked, and the device returned to a condition where it may be set up by another user. This is equivalent to delete all stored data and then reset the physical camera itself.

## 2.2 Privacy in Smart Camera Systems

Smart camera systems have been shown to have both privacy and security risks [6, 9, 14, 15, 29, 34, 46, 49, 53]. Motivated by legal frameworks (e.g., the California Consumer Privacy Act (CCPA), California Privacy Rights Act (CPRA), and the European General Data Protection Regulation (GDPR), among others [8, 35, 36, 42–44, 47]), we identify the following concrete rights that a *privacy-preserving system* should afford to the system owner, in practice, they imply that device owners must have exclusive control over the collection of data, its uses, and the access delegations to it:

(1) **Right to not be seen**: no unauthorized user can view stored videos or live video streams.
(2) **Right of sole ownership**: the owner retains full control of their data and who they trust.
(3) **Right to be forgotten**: deleted videos are not recoverable, even in cases of coercion.
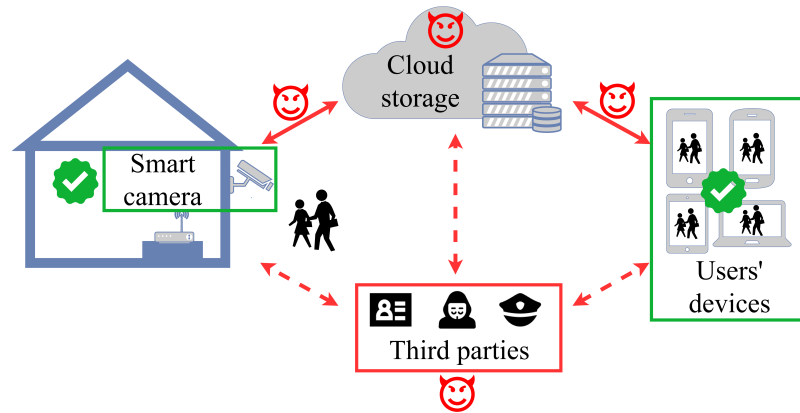
Fig. 1. Overview of the components and actors in the privacy-preserving smart camera system CaCTUs: the camera devices and the users' devices are trusted while the cloud storage provider, the networks, and any other third party are untrusted.

## 2.3 Threat Model

Our goal in this work is to demonstrate a smart camera system that provides feature parity with commercial systems while placing no trust in a cloud provider or other third party. As such, we work under a threat model wherein edge devices (i.e., the smart camera and end-user devices) are trusted, but the cloud storage provider, network, and any other third-party service are untrusted (see Figure 1 for an overview of CaCTUs's threat model).

We only trust the devices owned by the users (cameras, smartphones, laptops, or tablets) to securely handle the encryption and decryption keys used in the system, and we trust the device manufacturer to provide us with a camera device that correctly executes its functionality. This additionally implies that supply-chain exploits against the camera manufacturer are out of scope. We trust the other applications running on the users' devices (or that the operating system sufficiently isolates these applications) and we assume that the cryptographic algorithms used provide the advertised guarantees (e.g., Diffie-Hellman assumption [16] and RSA public-key cryptosystem [52]).

We also acknowledge that access under our system may be universally delegated: once granted access to a video, a party is not prevented from sharing with others or downloading and storing the videos somewhere else. Partial mitigations to this may be considered, but as such sharing can occur outside the purview of our system complete prevention is not possible. Encrypted frames are assumed to be publicly accessible (as we do not trust the cloud to do any access control mediation), thus we acknowledge that access pattern to the cloud storage may be leaked in CaCTUs. Ongoing research in Private Information Retrieval (PIR) or Oblivious RAM (ORAM) could provide potential mitigations through the use for example of random accesses and dummy writes to the cloud storage. Finally, physically tampering with the devices (modifying the hardware, chip-level changes to edit the software execution, etc.) and denial of service attacks are outside the scope of our work.

## 3 CACTUS

### 3.1 Overview

In the following sections, we will describe how CaCTUs meets the privacy goals described in Section 2 by providing the following three security properties: (1) *confidentiality*: stored videos and live video streams cannot be viewed by
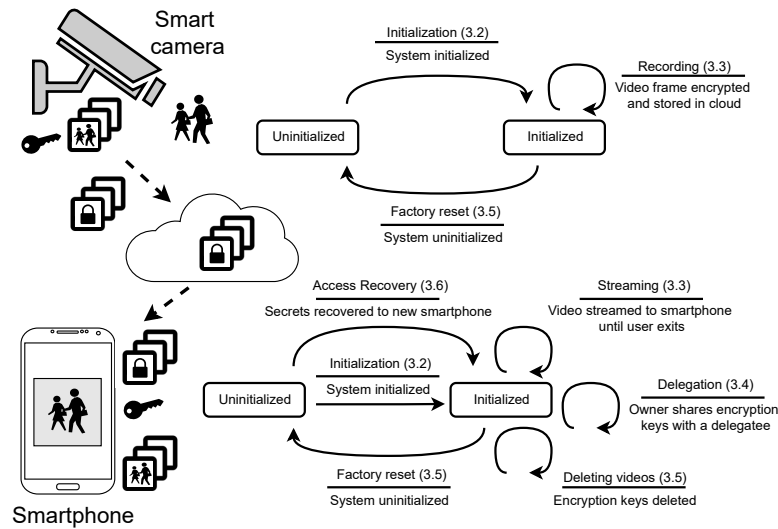
Fig. 2. Overview of the recording and streaming operations; the camera records video frames, encrypts them locally, sends the encrypted frames to the cloud, from where the smartphone application downloads them, decrypts them locally, and plays the video. The right part is the overview of the states and actions in the CaCTUs system for the camera at the top right and the owner's smartphone at the bottom right (references indicated between parentheses are the subsections to refer to for more details).

unauthorized parties, (2) *complete mediation*: the user fully controls access to their data by others, and (3) *deletion*: deleted videos are not recoverable, even in cases of coercion. Each subsection motivates the privacy goals before providing technical details of the feature. We refer to Appendix A for the notation used in our cryptographic constructions and algorithms and to Appendix D for the storyboard of the CaCTUs's smartphone application. For clarity, we consider only a single camera, though our approach readily generalizes to multi-camera systems by applying the described protocols to each camera individually.

Figure 2 shows an overview and state diagram of CaCTUs. The camera locally encrypts recorded video frames before uploading them to cloud storage. The smartphone application performs the reverse operations: it downloads the frames, decrypts them locally, and plays the video. Regular key rotation and secure key management allow the system to support secure streaming, delegation, deletion, recovery, and reset.

## 3.2 Initialization

During initialization, the user's smartphone and camera establish a trust association used for all other steps, so the security of this step is critical to that of the system as a whole. In CaCTUs, we adapt the *Seeing-Is-Believing* (SiB) technique introduced by McCune, Perrig, and Reiter [37] to establish an authenticated communication channel between devices that share no prior context (and without having to trust any third party). Specifically, we use the visual channel as an out-of-band means to verify the authenticity of each end of a Bluetooth channel.

This secure pairing bootstraps the system to allow the owner of the device to communicate directly with the camera (for system initialization) while ensuring confidentiality and integrity. Thus, negotiation of encryption keys to establish the root of trust can be done without any other party involved. Moreover, the required proximity and physical interaction

Table 1. Protocol followed by the camera and the owner's smartphone during the initialization, 👁 and ⌁ respectively correspond to what is obtained through the visual and Bluetooth channels.

| | Camera | 👁 | ⌁ | Smartphone |
|---|---|---|---|---|
| 1 | (QR code) | $\xrightarrow{h_{PK_f}}$ | | |
| 2 | | | $\xrightarrow{PK_f}$ | $h'_f = hash(PK_f)$ <br> $h'_f \overset{?}{=} h_{PK_f}$ |
| 3 | $h'_o = hash(PK_o)$ | | $\xleftarrow{PK_o}$ | $gen(SK_o, PK_o)$ |
| 4 | $h'_o \overset{?}{=} h_{PK_o}$ | $\xleftarrow{h_{PK_o}}$ | | (QR code) |
| 5 | $DH(SK_f, PK_o)$ | | $\xleftrightarrow{check}$ | $DH(SK_o, PK_f)$ |
| 6 | $gen(SK_c, PK_c)$ | | $\xrightarrow[(RSA)]{PK_c}$ | |
| 7 | $init(secrets)$ | | $\xleftarrow[(RSA)]{secrets}$ | $gen(escrow)$ <br> $passphrase$ |

between the devices would render attack attempts easily detectable by the system owners. The initialization protocol between the camera and the owner's smartphone application is as follows (see Table 1):

(1) A pair of factory-generated asymmetric keys $(SK_f, PK_f)$ is present on the camera device, and the hash of its public key $h_{PK_f}$ is embedded into a QR code on the back of the device (recall from Section 2.3 that the supply chain is trusted). The owner's smartphone scans this QR code and stores its content.

(2) The camera and the owner's smartphone connect through Bluetooth and the camera sends its public key $PK_f$ to the owner, who computes the hash of the camera's public key $h'_f$ and checks that it matches the hash retrieved from the QR code.

(3) If they match, the owner's smartphone generates its own asymmetric pair of keys $(SK_o, PK_o)$ and sends its public key $PK_o$ through Bluetooth, the camera computes the hash $h'_o$ of the owner's public key.

(4) On the screen of the owner's smartphone is displayed the QR code with the hash of the owner's public key. The camera retrieves the content from the QR code and checks that it matches $h'_o$.

(5) If the key hashes match, both devices now verify that the other device knows the secret key corresponding to the public key that they advertised earlier. This is can be done for instance by applying the Diffie-Hellman key exchange to compute their shared secret and then by exchanging a series of encrypted messages where both parties prove their knowledge.

(6) The camera then generates a new asymmetric key pair $(SK_c, PK_c)$ that it will use for future communication (to avoid further reliance on the factory-generated key). The camera then shares its public key $PK_c$ with the owner through Bluetooth in an authenticated way using RSA and the factory-generated asymmetric key pair

$(SK_f, PK_f)$. Note that we could have used the shared secret computed at the previous step through Diffie-Hellman, however, we chose to use RSA to align this step of our initialization protocol with the future protocols presented in CaCTUs as well as because we want authentication of the messages.

(7) Next, the owner sends their *secrets* to the camera device to complete system setup, this is done in a secure and authenticated way using RSA, but this time the new asymmetric key pair of the camera is used $(SK_c, PK_c)$. First, they send *wifi credentials* of the wifi network the camera should connect to. Then, they generate and send a *seed key* that will be used to derive the keys to encrypt video frames. Lastly, they send *escrow material* (protected by a non-recoverable *passphrase*) that may be used by the owner to recover access to the system (see Section 3.6 for details about the escrow material). If necessary, during this step the owner could also configure the cloud storage option they want to use if different than the default one.

At this point, the system is initialized, the camera begins recording, and the owner can execute other functions.

## 3.3 Recording and Streaming Videos

Here we describe how to ensure the *confidentiality*, *integrity*, *authenticity*, and *freshness* of the recorded video footage. The video frames and metadata are encrypted locally at the camera and asymmetrically signed in blocks of $N$ frames, before being uploaded to cloud storage. At each key rotation, the camera device securely erases the encryption keys previously used as it does not need them anymore. To view a video, users of the system download the encrypted data from the cloud storage, derive the decryption keys locally (if they have access to them), and decrypt the frames to rebuild the video. Thus, only users with access to the appropriate decryption keys can view the footage. As the camera is recording, it performs the following:

(1) Consider a block of $N$ frames. Each frame $F_i$ is recorded at timestamp $t_i$.

(2) A key rotation scheme is used to derive encryption keys for a given frame (this will prove useful for delegation, discussed in Section 3.4). The rotation scheme $\mathcal{K}$ provides a key $k_i$, the key used to encrypt frame $F_i$. An initialization vector $IV_i$ is randomly generated.

$$\{k_i = Extract(\mathcal{K}, i) | i \in [\![1, N]\!]\}$$

$$\{IV_i = RandBytes(16) | i \in [\![1, N]\!]\}$$

(3) Authenticated encryption is used to symmetrically encrypt each frame (**confidentiality)** into the corresponding ciphertext $C_i$ using the AES algorithm in Galois/Counter Mode (GCM, chosen for its performance) with a 256-bit key. Additional authenticated data $t_i$ is passed along and a tag $\tau_i$ is generated (**freshness and integrity)**.

$$\{C_i || t_i, \tau_i = AES256Enc(IV_i, k_i, F_i, t_i) | i \in [\![1, N]\!]\}$$

(4) A signature $\sigma$ of the block is computed using the private key $SK_c$ of the camera and the $N$ tags of the frames in this block (**authenticity)**.

$$\sigma = Sign(SK_c, \tau_1 || \tau_2 || ... || \tau_N)$$

(5) The encrypted and authenticated frames $\langle \{C_i || t_i, IV_i, \tau_i | i \in [\![1, N]\!]\}, \sigma \rangle$ are uploaded to the cloud along with their corresponding metadata (initialization vector used for encryption, tag, and timestamp).

Each user who has access to the correct decryption keys can download these encrypted and authenticated frames $\langle \{C_i || t_i, IV_i, \tau_i | i \in [\![1, N]\!]\}, \sigma \rangle$. To view the video, the user performs the following:

(1) We consider a block of $N$ frames of signature $\sigma$ downloaded on demand. Each ciphertext $C_i$ along with its tag $\tau_i$, encrypted using the initialization vector $IV_i$, corresponds to a frame recorded at timestamp $t_i$ .

(2) The signature $\sigma$ of the block is verified with the public key $PK_c$ of the camera and the $N$ tags of the frames in this block (**authenticity**).

$$1 \stackrel{?}{=} Verify(PK_c, \sigma, \tau_1||\tau_2||...||\tau_N)$$

(3) If the signature is correct, the corresponding symmetric key $k_i$ is extracted from key rotation scheme $\mathcal{K}$. Then, symmetric authenticated decryption turns each ciphertext $C_i$ into the corresponding frame $F_i$ (**confidentiality**) while checking data integrity (**integrity and freshness**)

$$\{k_i = Extract(\mathcal{K}, i)|i \in [\![1, N]\!]\}$$

$$\{F_i = AES256Dec(IV_i, k_i, C_i||t_i, \tau_i)|i \in [\![1, N]\!]\}$$

Encryption is performed *end-to-end*: data is encrypted locally at the camera before being stored in the cloud and decrypted locally at the smartphone after being retrieved. Furthermore, integrity is ensured using an authenticated encryption scheme, so that the video footage cannot be tampered with during transmission or storage. Lastly, the identity of the camera is embedded into the video frames (and signed) so that users can attest the authenticity of the video footage. This scheme allows the user to verify the integrity, authenticity, and freshness of an arbitrary set of video frames, so that they then decrypt the frames and rebuild the video for playback.
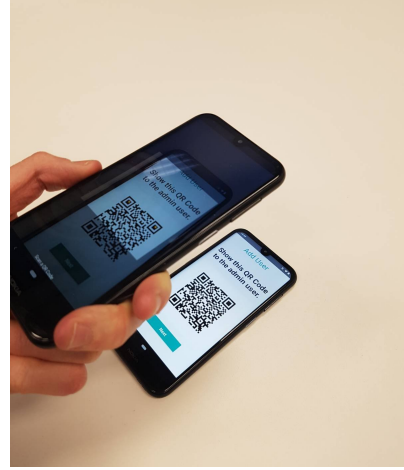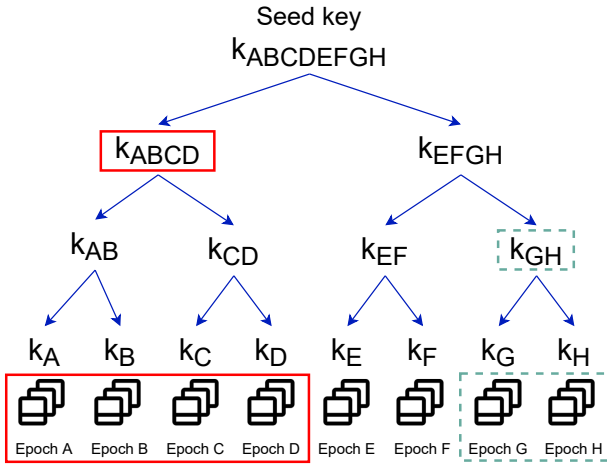
## 3.4 Delegation

Users must have complete mediation over access to their videos. CaCTUs achieves this by ensuring that the owner has control of the keys used to encrypt the video footage. However, they may also want to delegate access to their videos (e.g., to friends or family) for different periods of time. Achieving fine-grained sharing capabilities for delegatees is nontrivial: we want to support delegation without knowing beforehand to whom the owners will delegate access or for how long. To enable this, we rotate the keys used to encrypt the video frames at the end of every *epoch* (a fixed-size time interval). We use a binary key tree construction to facilitate the management of all the keys for the camera device, owners, and delegatees. A peer-to-peer pairing is adopted for sharing keys so that there is no reliance on a third party.

Recall that frames are encrypted using a symmetric key derived from a key rotation scheme $\mathcal{K}$. In practice, to support delegation, this rotation scheme is a binary key tree, inspired by the key tree introduced by Kocher [31]. The tree is of a fixed depth $d_{\mathcal{K}}$. In the tree, each leaf node holds some cryptographic key $k$ and covers a specific epoch: a time interval $[t_j, t_{j+1})$ of fixed-size $\delta_{\mathcal{K}}$. The root node of the tree is initialized with a seed key that is negotiated during the initialization of CaCTUs (see Section 3.2). The timestamp of this negotiation is used for $t_0$, with $t_{j+1} = t_j + \delta_{\mathcal{K}}$. The leaf nodes in the tree hold the encryption keys for every epoch. Each node in the tree can be derived from the root node knowing the derivation equations and relations between the parent node and its two children. Within each epoch, the symmetric key used for each frame is identical. The derivation of keys is based on a Hash-Based Key Derivation function (HKDF), which is a one-way process [28]. If we have $k_{parent}$, then $k_{left} = HKDF(k_{parent})$ and $k_{right} = HKDF(k_{parent} \oplus 1)$. Therefore, for a given key in the tree, a user can only derive the keys below it but not the ones above (see Figure 3a for an illustration).

The binary key tree is useful for several reasons. First, it decreases the amount of keys that need to be shared with the delegatees, as the derivation algorithm is publicly known by the Kerckhoffs's principle and a specific part of the key tree can be reconstructed on-demand by a delegatee that is given access to a node of the tree. Moreover, it is storage

(a) Key Tree Construction used by CaCTUs, the keys to be shared to give access to the corresponding footage are framed.

(b) Establishment of a secure and authenticated Bluetooth pairing with a delegatee.

Fig. 3. Delegation in CaCTUs

space-efficient, as the key rotation mechanism generates a large number of encryption keys. When key rotation happens and the camera device securely erases the previous encryption keys, the camera device only needs to save at most $d_{\mathcal{K}}$ nodes to keep functioning. Further, it is simple to find which node is responsible for the encryption or decryption of a specific timestamped frame (through a binary search). Lastly, it facilitates the recovery process for the owners as only a subset of keys needs to be recovered to rebuild the tree from the escrow material.

Delegation is supported by giving some keys of the key tree to each delegatee (an example is given in Figure 3a where keys to share and their corresponding epochs are framed). From there, each delegatee can derive the associated keys to correctly decrypt and view only video footage captured within the time window they were authorized for. The depth $d_{\mathcal{K}}$ and the epoch size $\delta_{\mathcal{K}}$ are configurable parameters of the system. The size of an epoch $\delta_{\mathcal{K}}$ corresponds to the lowest delegation granularity achievable in CaCTUs.

Delegation uses a peer-to-peer approach similar to the physical pairing of the initialization with the camera. Table 2 shows the details of the operations performed during the delegation protocol between the owner and a delegatee:

(1) The pair of asymmetric keys $(SK_o, PK_o)$ is present on the smartphone of the owner. When the delegation process starts, the smartphone application of the owner displays a QR code in which the hash of the owner's public key $h_{PK_o}$ is embedded. The delegatee's smartphone scans this QR code and stores its content.

(2) The delegatee and the owner's smartphone connect through Bluetooth and the owner sends its public key $PK_o$ to the delegatee, who computes the hash of the owner's public key $h'_o$ and checks that it matches the hash retrieved from the QR code.

(3) If they match, the delegatee's smartphone generates its own asymmetric pair of keys $(SK_d, PK_d)$ and sends its public key $PK_d$ through Bluetooth, the owner computes the hash $h'_d$ of the delegatee's public key.

(4) On the screen of the delegatee's smartphone is displayed the QR code with the hash of the delegatee's public key. The owner retrieves the content from the QR code and checks that it matches $h'_d$.

(5) If the key hashes match, both devices now verify that the other device knows the secret key corresponding to the public key that they advertised earlier. This is can be done for instance by applying the Diffie-Hellman key exchange to compute their shared secret and then by exchanging a series of encrypted messages where both parties prove their knowledge.

(6) Then, the owner extracts the keys $\{k_i\}$ to share and send them in an encrypted and authenticated way using RSA to the delegatee to give them access to the corresponding videos. Note that this last step does not necessary need to be done through Bluetooth and could be done over the Internet too without undermining the security of the delegation protocol.

Table 2. Protocol followed by the smartphone application of the owner and the delegatee during delegation, 👁 and ⌘ respectively correspond to what is obtained through the visual and Bluetooth channels.

| | Owner | 👁 | ⌘ | Delegatee |
|---|---|---|---|---|
| 1 | [QR code] | $\xrightarrow{h_{PK_o}}$ | | |
| 2 | | | $\xrightarrow{PK_o}$ | $h'_o = hash(PK_o)$ <br> $h'_o \overset{?}{=} h_{PK_o}$ |
| 3 | $h'_d = hash(PK_d)$ | | $\xleftarrow{PK_d}$ | $gen(SK_d, PK_d)$ |
| 4 | $h'_d \overset{?}{=} h_{PK_d}$ | $\xleftarrow{h_{PK_d}}$ | | [QR code] |
| 5 | $DH(SK_o, PK_d)$ | | $\overset{check}{\longleftrightarrow}$ | $DH(SK_d, PK_o)$ |
| 6 | keys $\{k_i\}$ | | $\overset{\{k_i\}}{\underset{(RSA)}{\longrightarrow}}$ | $\mathcal{K} = init(\{k_i\})$ |

## 3.5 Deleting Videos and Factory Reset

CaCTUs users must be able to delete their videos and factory reset their system. As video frames are uploaded to the cloud, deleting them would require trusting the cloud storage provider to do so. However, as video frames can not be decrypted without the decryption keys, deleting these keys is equivalent to deleting the corresponding video. Owners can achieve this by deleting select decryption keys (i.e., a subset of nodes in the key tree) and by updating the *escrow material*, so that the keys below them in the tree cannot be recomputed, even in cases of coercion. Thus, with a tree depth of $d_\mathcal{K}$, for each portion of video content composed of $n_e$ epochs to delete, the upper bound of the number of nodes that must be deleted from the tree is $O(d_\mathcal{K} n_e)$. To provide an example: in Figure 3a, to delete $k_A$ (key for epoch A), keys $\{k_{AB}, k_{ABCD}, k_{ABCDEFGH}\}$ must also be deleted. Thus, $\{k_B, k_{CD}, k_{EFGH}\}$ must be saved in this sparser key tree so that the corresponding videos can still be decrypted. Recall also that the camera device securely deletes the encryption keys as soon as it does not need them anymore due to key rotation. Note that with this mechanism, delegatees may still know some decryption keys that were deleted from the owner's device. As discussed in our threat model (see Section 2.3), once granted access to a video, a party is not prevented anyway from having already shared or downloaded the video.

To factory reset, the camera and the owner must forget about all the encryption keys (symmetric and asymmetric), as well as the *secrets* exchanged during initialization, and the *escrow material*. Thus, for both cases (see Table 3a):

(1) The owner sends the timestamped and authenticated request encoding the operation as well as any updated *key material* encrypted and authenticated with RSA (recall that the owner and the camera have shared their asymmetric public keys during initialization). This can be done through Bluetooth or remotely over the Internet.

(2) At reception, the camera verifies the authenticity of the deletion or reset operation; if valid it sends an ack and performs it. When receiving the ack, the owner's smartphone applications completes the delete or reset operation locally.

Table 3. Deletion, factory reset, and access recovery protocols.

(a) Protocol followed by the owner and the camera during deletion or factory reset operations. ✳ and 🌐 respectively corresponds to the Bluetooth or any other channel such as the Internet.

| | Owner | ✳ or 🌐 | Camera |
|---|---|---|---|
| 1 | *update key material* | $\overset{operation\ \&\ key\ material}{\longrightarrow}$ $(RSA)$ | *verify operation?* |
| 2 | *complete operation* | $\overset{ack}{\longleftarrow}$ $(RSA)$ | *perform operation* |

(b) Protocol followed by the camera and the smartphone application of the user trying to recover access. ✳ correspond to what is obtained through the Bluetooth channel.

| | New smartphone | ✳ | Camera |
|---|---|---|---|
| 1 | | $\overset{request}{\longrightarrow}$ | |
| 2 | *passphrase known?* | $\overset{escrow}{\underset{material}{\longleftarrow}}$ | |
| 3 | $(SK_o, PK_o), PK_c,$ and $\mathcal{K}$ retrieved | | |

## 3.6 Access Recovery

In case owners lose access to their smartphone, they must be able to recover access to the system. However, in a privacy-preserving system where no third party is trusted, achieving this is nontrivial. To solve this, we generate an *escrow material* during the initialization step that contains all of the information needed by the owners to recover access to their system. The escrow material is encrypted by an unrecoverable passphrase only known by the owner and stored on the camera. The escrow material gives access to the following secrets :

- The *owner's asymmetric key pair* $(SK_o, PK_o)$ encrypted with a randomly generated key of size 128 bits and corresponding to the passphrase displayed to the owner during initialization.
- The *key material* necessary to build the key tree $\mathcal{K}$ (for details about this key tree see Section 3.4) asymmetrically encrypted with the owner's key.
- The *asymmetric public key* $PK_c$ of the camera (does not need to be encrypted).

Table 3b details the steps executed during access recovery to the system:

(1) The owner uses their new smartphone to request through Bluetooth to the camera the escrow material.

(2) The camera sends back the escrow material. Note that no assumption regarding the status of the owner has been made so far; anyone who is in physical range from the camera can request the escrow material. However, only the owners have knowledge of the recovery passphrase and can use it to decrypt the escrow material.

(3) If the passphrase is known, they are able to recover access to the asymmetric key pair $(SK_o, PK_o)$ of the owner, the public key $PK_c$ of the camera, and the *key material* necessary to build the key tree $\mathcal{K}$.

## 4  EVALUATION

The goal of this section is to evaluate the effectiveness and efficiency of CaCTUs with respect to four metrics: security, privacy, usability, and performance. To this end, we evaluate three research questions:

- **RQ1:** Does CaCTUs enforce our privacy requirements, while offering the same feature set of commercial systems?
- **RQ2:** Is CaCTUs easy to use for end users?
- **RQ3:** Can CaCTUs operate at sufficient resolution, frame rate, and latency to meet the needs of users?

*Experimental Setup.* Experiments were performed using a *Raspberry Pi 4 Model B* with 2GB of RAM, equipped with a camera module. Our implementation of the CaCTUs camera system is written in C, while the paired mobile companion application was written in Java and installed on a *Nokia 4.2* smartphone, running Android 10. Further evaluation setup details are described in Appendix B.

### 4.1  Privacy and Security Analysis (RQ1)

We begin by analyzing how CaCTUs achieves the security and privacy properties discussed in Section 2.2.

*Confidentiality.* CaCTUs aims to protect user data from being read by unauthorized third parties. As such, *end-to-end encryption* is used to enforce that only the owners have access to their videos; data is encrypted locally at the camera before being stored in the cloud and decrypted locally at the smartphone after being retrieved. The encryption keys are exclusively located on the camera device and the owner's smartphone. Owners manage access to the decryption keys and decide with whom to share them. As a result, parties who have not been delegated access through key sharing are unable to view the encrypted frames, protecting confidentiality of the footage, which ensures the right to not be seen. It is also worth noting that the key-tree construction of CaCTUs allows to delegate access at a fine granularity; users exactly control which portion of their video footage they are sharing with others.

*Complete Mediation.* CaCTUs uses end-to-end encryption which means that the owners of the systems are the only ones to have access to the decryption keys. Thus, delegation can only be performed by the owner (who holds the material necessary to build the key tree); and no other party is able to grant access to the footage, this provides complete access mediation to the videos. The key tree ensures this mediation cryptographically; a user without the seed or delegated keys is unable to decrypt footage regardless of access control on the encrypted data, assuming the cryptographic primitives are secure. This enforces the right to sole ownership of the videos recorded by the system. Additionally, during initialization a root of trust is established between the camera device and the owner's smartphone, specifically they have exchanged their public keys. The camera device will perform some administrative actions such as updating the escrow material or factory resetting if and only if the request is authenticated and signed by the owner's asymmetric key. This is another aspect of CaCTUs that enforces that owners have sole control over their systems.

*Data Deletion.* An emerging property of the previous two ones is that owners have sole ownership and control over the encryption keys. As the recorded footage can not be decrypted without having these keys, deleting the keys is equivalent to deleting the corresponding video files that are by themselves just meaningless blobs of data (as they are encrypted). As a result, if owners decide they want to delete some or all of their videos, they can simply delete the corresponding encryption keys on their smartphone as well as in the escrow material on the camera device, which makes it impossible for them to view the corresponding encrypted footage stored in the cloud, even in case of coercion. Note that no third party needs to be trusted to perform this operation and that this ensures owners' right to be forgotten.
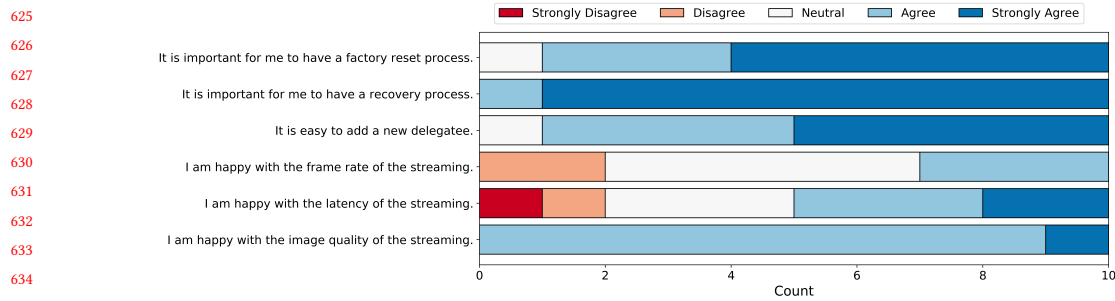
Fig. 4.  Likert scale evaluation of different statements made by the 10 participants during the functional user evaluation of CaCTUs.

## 4.2   Functional User Evaluation (RQ2)

In order to assess the functional ease of use of CaCTUs and identify what aspects of the system could be improved, we performed a functional user evaluation of our implementation with ten participants. Studies have shown that this group size is sufficient to identify most of the issues within a system design [40, 41]. To this end, participants were asked to initialize the system and use the different functionalities as if they had just bought it: view live and recorded streams, share videos with a delegatee, recover access after simulating the loss of their smartphone, and reset the system. The detailed protocol that was followed is described in Appendix C. We obtained approval from the Institutional Review Board (IRB) of our university and all other institutional requirements were met for this evaluation. Out of the ten participants, four were female and six male, seven of them were between 20-25 years old and the three between 45-55 years old.

*Goals and Limitations.* Our modest goals through this functional user evaluation of CaCTUs were to assess whether guided users can perform the functions of the system that we laid out and identify where our proof-of-concept of the implementation fell short (with respect to performance or feature design) and how it could be improved. As a consequence, our participants knew they were evaluating CaCTUs. Thus, our functional evaluation differs from traditional usability studies that may compare different system designs to identify which is the most usable, has the best interface, or what UX options are optimal for ensuring that users understand a specific privacy concept of the system. We defer a more comprehensive usability study to future work.

*User Interface.* Overall, the participants thought that the interface of the implemented smart camera system was simple to understand and navigate through. They liked that for every process there were step-by-step instructions displayed to them. They noted that these directions were clear, straightforward, and self-explanatory.

*Secure and Authenticated Pairing.* Despite that pairing requires both setting up a Bluetooth connection and scanning QR codes, the participants found the process simple and easy. They expressed that it was more straightforward than steps they had to perform with other systems. They also felt that using both Bluetooth and a visual channel was more secure even if they did not explicitly always know why.

*Quality of the Footage.* As shown in Figure 4, the participants agreed that the image quality of the video stream (480 p) and the frame rate (10 fps) were sufficient for security and surveillance purposes. However, some believed that the latency of the system was a bit of a downside as it tended to defeat the initial purpose of being able to monitor what

13

was recorded by the smart camera system in real-time. We took this comment into account and further improved our implementation to achieve the results presented in Section 4.3, we also discuss in Section 5 other potential optimizations to CaCTUs.

*Granularity Options.* The participants were impressed by the granularity to which they could view specific segments of video footage (up to the second), but believed that the same granularity option for delegation was not necessary; they stated that the primary use case would be to share access for several hours or days. This concern can easily be addressed by modifying the granularity options displayed in the application. The majority of the participants also expressed desire to have an option to quickly delegate unlimited access, but then figured out that they would not be able to revoke such access without having to reset the system.

*Access Recovery and Factory Reset.* Regarding the recovery process, the participants were divided on whether the recovery passphrase should be randomly generated by the system or if it should give the owner the opportunity to choose it. They felt that they could lose the recovery passphrase or forget about it if they did not choose it. However, they agreed that as this passphrase allows to recover full access to the system, it might be less secure to let the owners pick their own. One participant remarked that the ability to recover access without needing to trust a third party (i.e., without using a recovery email address for instance) was interesting. As shown on Figure 4, most participants found it important to be able to recover access to and factory reset their system, as oftentimes owners may want to recover access (to retrieve videos) before factory resetting it. The participants found the recovery and factory reset processes easy to perform.

*Missing Features.* Participants expressed that they would like to see the following features implemented in a commercial system: motion detection-triggered recording, remote pairing and delegation, two-way audio support to listen in on and remotely speak through devices, password-locked smartphone application for additional security, and application availability across different platforms (e.g., Android, iOS, and web interface). See Section 5 for details about such extensions to CaCTUs.

## 4.3 Performance Evaluation (RQ3)

We now evaluate the streaming performance of CaCTUs, focusing our efforts on three key metrics: latency (delay from time of recording), stream image resolution, and frame rate. As a baseline, commercial systems achieve frame rates of 30 fps in 1080 p (1920 x 1080 pixels), with a latency on the order of milliseconds. However, we note that these systems have been largely optimized for commercial use and lack privacy-preserving features as CaCTUs. We discuss additional potential optimizations to CaCTUs than the ones presented in this section in Section 5.

*System Latency.* Recall that users pointed out in our functional evaluation that the latency of the system was too high. Taking that into account, we have since introduced a new optimization in our implementation to improve the streaming experience and reduce the latency without sacrificing security; uniform frame dropping based on the delay that the live stream is at. Thus, for the parameters considered in the functional user evaluation, we obtained a latency of 2 s at a frame rate of 10 fps for a video resolution of 480 p.

Figure 5a displays the evolution of the latency of our implementation without and with frame dropping while live streaming at different frame rates and video qualities. Frame dropping clearly improves the latency of CaCTUs as the implementation is now able to keep the delay, between the recording and rendering times, to a lower value than previously by uniformly dropping frames. Specifically, at high frame rates our unoptimized implementation was not
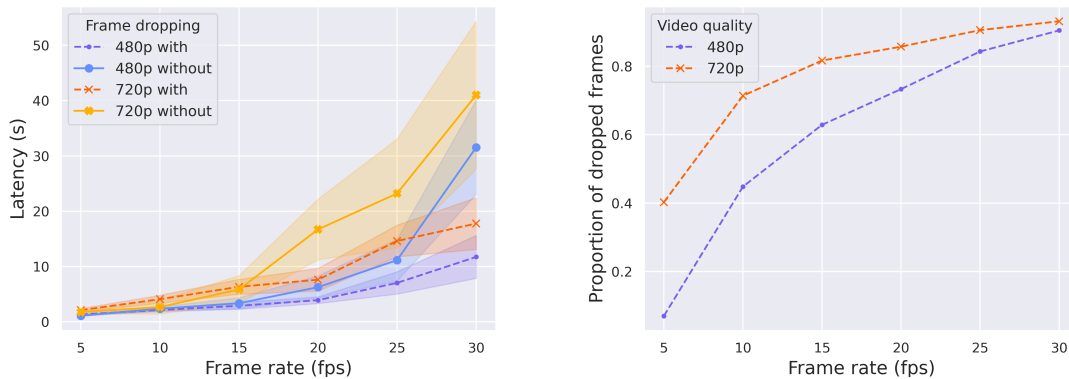
(a) Latency (± standard deviation represented by the shaded areas) for different frame rates of live streaming at a quality of 480 p and 720 p using frame dropping or no, averaged over 1,000 frames.

(b) Proportion of frames dropped by our implementation while live streaming at different frame rates for a video quality of 480 p and 720 p.

Fig. 5. Frame dropping

able to keep up with the live stream as too many frames needed to be processed. As we use uniform frame droping based on the current delay where the stream is at, the transition between the frames of the output video on the smartphone application still looks smooth and does not appear as if it was lagging or buffering. Figure 5b shows the proportion of frames being dropped for different frame rates and video qualities by our system; the higher the frame rate, the more frames are dropped in proportion. These two figures show that our implementation of CaCTUs behaves the best for a lower video quality and a lower frame rate, confirming our choice of using a frame rate of 10 fps and a video resolution of 480 p for the rest of our evaluation. As expressed by the users of our study, these parameters are sufficient for streams recorded by a security camera (see Section 5 for other optimizations).

*System Bottlenecks.* Next, we measure the effect that each phase of streaming has on the stream latency. The values have been averaged over 1,000 frames. For this evaluation, we picked the same parameters as for a worst-case scenario (i.e., $d_{\mathcal{K}} = 32$ and $\delta_{\mathcal{K}} = 10$ s), to show the performance baseline that can be expected from CaCTUs. Table 4b shows the results when the camera device is recording for a video quality of 480 p. As shown, the largest contributor to the latency is the upload/download of the encrypted frames to/from the cloud storage server during live stream, this would have the same contribution to latency in a system not using encryption techniques. We leave to future work improving that aspect and discuss means to do so in Section 5. Note that within the same epoch the same key is used, but between epochs, we need to derive the new key, that is why the standard deviation is larger than the average for the key extraction.

*Lifespan and storage space.* The depth $d_{\mathcal{K}}$ and the epoch size $\delta_{\mathcal{K}}$ are configurable parameters of the system. For our implementation, we selected $d_{\mathcal{K}} = 32$ and $\delta_{\mathcal{K}} = 10$ s. Importantly, this choice of parameters demonstrates the worst-case performance users can expect from CaCTUs. Specifically, these parameters result in a key tree that covers a lifespan of 1362 years at a 10 s level of granularity. We provide more reasonable parameter choices (and, thus, expected performance at deployment) in Table 4a, where we also present the storage space required in the worst-case scenario where every other epoch videos have been deleted by the user. In practice, such a scenario is very unlikely to happen as

15

it will render the system unusable as well as because owners are unlikely to delete beforehand keys that would have been used to encrypt future videos. Thus, far less amount of storage space is required; moreover, our binary key tree structure enables dynamic derivation of lower keys.

Table 4

(a) Lifespan of the key tree for an epoch time $\delta_{\mathcal{K}}$ of 10 s and 60 s and storage space needed to save to disk the encryption keys in the worst-case scenario for different depth sizes $d_{\mathcal{K}}$.

| $d_{\mathcal{K}}$ | Lifespan for | | Storage space (worst-case scenario) |
|---|---|---|---|
| | $\delta_{\mathcal{K}} = 10$ s | $\delta_{\mathcal{K}} = 60$ s | |
| 24 | 5 years | 32 years | 256 MB |
| 26 | 21 years | 128 years | 1 GB |
| 28 | 85 years | 511 years | 4 GB |
| 30 | 340 years | 2043 years | 16 GB |
| 32 | 1362 years | 8172 years | 64 GB |

(b) Time delay averaged over 1,000 frames and standard deviation $\sigma$ of each operation, while recording at a resolution of 480 p.

| Device | Operation | Delay | $\sigma$ |
|---|---|---|---|
| Camera | Key Extraction | 0.05 ms | 0.2 ms |
| | Frame Encryption | 4.0 ms | 1.3 ms |
| | Signature | 8.8 ms | 2.8 ms |
| | Upload | 662 ms | 536 ms |
| Smartphone | Download | 204 ms | 358 ms |
| | Signature Verification | 0.7 ms | 0.5 ms |
| | Key Extraction | 0.03 ms | 0.3 ms |
| | Frame Decryption | 0.6 ms | 0.6 ms |

## 5 DISCUSSION

In the following, we discuss potential optimizations for improving the performance of CaCTUs and practical considerations for deploying it as a commercial system.

### 5.1 Improving Latency

We identify several components that could be improved to reduce the overall latency of the system, namely; (1) cryptographic accelerators, (2) network relays, (3) streaming libraries, and (4) video compression techniques. Such improvements are described below.

As discussed in Section 4, our implementation used a Raspberry Pi, wherein the (relatively weak) CPU was responsible for handling all processing, including encryption. Since cryptographic operations dominate many of the features in CaCTUs, we expect substantial gains in latency by leveraging dedicated cryptographic accelerators as seen in many other crypto-dominated applications, such as in IoT [30].

Streaming live video can be demanding on the network. Specifically, popular video streaming platforms (such as Netflix, Hulu, and Disney Plus) employ a variety of techniques to bring video data as close to the user as possible. These techniques often take the form of caches, content delivery networks, or dedicated network infrastructure designed to serve high-bandwidth content quickly [1, 20, 54]. Naturally, these techniques could substantially improve the performance of CaCTUs to be even closer to commercial-grade systems.

The smart camera systems available today use optimized streaming protocols to deliver video content quickly [1]. Given that it was necessary for us to implement our video streaming protocol from scratch (to support our encryption and key rotation schemes), we could see further improvements by augmenting current protocols to support our design. In a similar vein, our current implementation operates at the frame-level, while commercial systems operate at the block-level (and thus exploit compression algorithms commonly used in video streaming applications [3]). Moreover, popular techniques such as adaptive video playback or more advanced frame dropping could also be used to improve the throughput of CaCTUs. We defer such improvements to future work.

## 5.2 Deploying CaCTUS as a Commercial System

Here, we highlight some challenges (motivated by notable features in commercial systems and suggestions from our functional user evaluation of CaCTUs) that should be addressed to realize commercial implementations of CaCTUs without compromising any of our privacy goals.

*Relaxing Proximity.* To uphold our privacy goals, delegatees need to be within local proximity of the owners (e.g., to perform delegation and pairing through the QR codes). While this was not an area of concern in the functional user evaluation of CaCTUs since participants were always close to the camera system, this can be challenging if users wish to delegate access remotely. We did conceptually derive a scheme to support this capability, while upholding our privacy goals. Specifically, CaCTUs could be extended to asymmetrically[2] encrypt the keys needed by delegatees and upload them to the cloud storage or use another third party service like email, from where the remote delegatee could retrieve, decrypt, and then use these keys to access the video stream for some predetermined period of time.

*Revoking Access to Unrecorded Videos.* Note that a similar scheme to the one described in the previous paragraph could also be used to allow owners to revoke access to delegatees to unrecorded videos without needing to factory reset the system. The sharing of the keys only needs to be done periodically at a certain interval that could be configured by the owner for each delegatee.

*Supporting Audio.* In our current implementation, we focused on supporting video data only, and not audio. However, the system design does not prohibit extensions to support audio-video recording and streaming. The extension is straightforward: we can treat some audio sample as a "frame", as it is done for video, and apply similar encryption operations. For sampling, a Linear Predictive Coding (LPC), which is widely deployed by telephone companies for speech encoding and processing, could be used [55]. Moreover, LPC can even be leveraged to transmit audio in reverse; that is, be applied so the users could speak into their smartphones and have the audio replayed through the camera, as commonly seen in commercial systems.

*Supporting Motion Detection.* Many commercial systems support forms of motion detection to notify users of events they may be interested in viewing. We find this feature valuable for both users as well as the camera system in that video recording need only be saved if motion was detected[3]. CaCTUs could support this capability through addition of a physical hardware sensor or a software solution. This addition would save significant space in cloud storage systems and reduce the cost related to storage. In a similar vein older data could be overwritten at a fixed interval that could be configured by the owner.

*Physical Security.* Finally, as explained in our threat model, we did not take into consideration physical tampering attacks against the camera device. Techniques such as package design so that the device is tamper-proof should be explored. Specifically, local storage, battery, reset buttons, serial ports, firmware, etc., should be protected to prevent physical access by untrusted parties.

---

[2]Clearly, how the public and private keys are computed and transferred between the owner and the remote delegatee needs to be done in a security-preserving manner.

[3]To mitigate side-channel attacks and to not leak behavior patterns, meaningless data could be randomly uploaded to the cloud.

## 6 EXTENDING TO OTHER DEVICES

The protocols used in CaCTUs could be easily extended to other IoT devices recording time series data or monitoring some events such as: temperature, humidity level, power usage, sensors, etc. Instead of video stream, which is bandwidth-intensive, any other blob of data recorded by the considered IoT device can be encrypted by following the same exact protocol and by using the same key tree construction. Depending on the application (room temperature, humidity condition, current power consumption of some devices around a house, or status of different sensors in a factory, etc.), it could be more convenient to use blocks of $N = 1$ elements for the encryption. As video stream is most likely the most bandwidth-intensive of these applications, performance for such devices is expected to be even better. In the next section, we discuss how the protocols designed in CaCTUs can be readily extended to other type of IoT devices, we then demonstrate it by adapting our open source artifact to be able to send any arbitrary timestamped data.

### 6.1 Discussion

*Privacy Concerns.* As for CaCTUs, owners should also be concerned that other devices usually located inside their homes can record and transmit valuable data about them or their behaviors. Thus, users of such systems should have similar rights than the ones enforced in CaCTUs; owners must have exclusive control over the collection of data, how that data is used, and by whom it is accessed. As a result, the threat and trust models considered in this work readily apply to other devices as well: the device manufacturer and the service or cloud providers do not have to be trusted more than for what they are needed (i.e., manufacture the product, maintain the software, and provide the infrastructure) which can be performed as demonstrated by CaCTUs without having access to the recorded data.

*Functionality.* Additionally, if we consider the life cycle of these IoT devices, similarities with the functionality of CaCTUs start to emerge; they also need to be initialized, record data, support access recovery, etc. As other IoT devices require an initialization phase during their setup, a physical pairing similar to the one we presented for CaCTUs can be followed to establish a root of trust between the owner and their IoT device without having to trust any third party. The protocols presented in this work were designed to build a privacy-preserving smart camera system, however, in practice our techniques do not require the recorded data to be video frames or the device to be a smart camera, making this extension to other IoT devices possible. Sometimes, IoT devices are connected to a hub or a base station in users' homes; in that case, the hub could be considered as a shared user to whom access was delegated.

*Potential Adaptations.* The physical pairing presented in CaCTUs was leveraging the fact that both devices had a camera sensors, however, this may not be the case for other IoT devices. In that case, the step can either be skipped and we defer to the *Seeing-Is-Believing* (SiB) technique introduced by McCune, Perrig, and Reiter [37] for a full discussion of the guarantees provided in such a case, or replaced with another mechanism as long as it still requires the owner to perform an action demonstrating to the device that they are in its physical proximity. For that, we should normally easily be able to come up with some physical task for the owner to perform that involves the capability of the sensor embedded in the device. For instance during pairing, a thermostat could ask to be placed closer to a warmer or colder source (e.g., chimney, freezer, etc.), a smart outlet could require the owner to plug a device and use it for a moment, a proximity or physically-activated sensor could ask the owner to trigger it in a simple pattern, etc. Similarly, any wireless protocol other than Bluetooth can be used for the pairing.

Table 5. Comparison of privacy guarantees of smart camera systems proposed in the literature (ROI stands for Region Of Interest) as well as some of the most popular commercial systems.

| | Camera System | Right to not be seen<br>End-to-End<br>Encryption | Right of sole ownership<br>Only the Owner is Trusted<br>& Controls Access | Right to be forgotten<br>Video<br>Deletion |
|---|---|---|---|---|
| Research | PrivacyCam [13] | ✗ (ROI only) | ✗ | ✗ |
| | TrustCAM [57] | ✗ (ROI only) | ✗ (trusts a central station) | ✗ |
| | TrustEYE.M4 [58] | ✓ | ✗ | ✗ |
| | SoC-based [25] | ✓ | ✗ (relies on a Trusted Authority) | ✗ |
| | Signcryption [56] | ✓ | ✗ (uses a Key Distribution Center) | ✗ |
| | Pinto [62] | ✗ (ROI only) | ✗ | ✗ |
| | CaCTUs | ✓ | ✓ | ✓ |
| Commercial | Arlo [5] | ✗ | ✗ | ✓ |
| | Blink [7] | ✗ | ✗ | ✓ |
| | Eufy [18] | ✓ | ✗ | ✓ |
| | Nest [22] | ✗ | ✗ | ✓ |
| | Ring [51] | ✓ (opt-in) | ✗ | ✓ |
| | Wyze [61] | ✗ | ✗ | ✓ |

## 6.2 Artifact: Adding Support for Timestamped Data

## 7 RELATED WORK

### 7.1 Research Systems

While this work is the first to examine a smart camera system that is privacy-preserving, there is already a large body of research studying the security of smart camera systems. Here, we detail the gaps in prior work and how CaCTUs addresses them. Refer to Table 5 for a comparison of the privacy guarantees of these systems.

Alharbi and Aspinall introduced a security analysis framework for IoT smart cameras that weighs the threat of significant risks (e.g., unencrypted video streaming) across various platforms [2]. However, it focuses on the security of the camera device and only partially addresses some vulnerabilities of other components of the system. For instance, the authors do not discuss the use of cloud storage to remotely access recordings, nor procedures for securely pairing a smartphone and camera. CaCTUs provides an end-to-end secure solution starting from camera initialization to protect the secrecy and integrity of both communications and data.

Haider and Rinner proposed a SoC-based smart camera that uses physically unclonable functions (PUFs) to generate encryption keys [25]. The keys are used to encrypt video frames at the camera before storing them in remote cloud storage, thus removing the requirement that the owner trust the cloud service provider. However, this approach has several limitations. First, a trusted authority is required to create camera device fingerprints during key generation; in a commercial system, the trusted authority will likely be the manufacturer, whom in general is not trusted by the camera owner. Moreover, the key extraction procedure using PUFs only produces a fixed number of encryption keys, which does not align with the feature-set typically desired by smart-camera owners—e.g., being able to delegate camera access to other people with fine granularity. Finally, as the fingerprint is physically embedded into the hardware, camera owners will likely need to get a brand new device if the encryption keys are leaked. CaCTUs addresses these shortcomings by making the owner (i.e., their smartphone) the root of trust: they store the secrets and share expendable keys with both

the camera and delegatees. This simultaneously gives control back to the owner while enabling delegation and system reset without specialized (or new) hardware.

Winkler and Rinner introduced TrustEYE.MP4 [58], a monitoring framework that provides similar secrecy and integrity guarantees envisioned by CaCTUs. However, as the same key is used to encrypt all the videos, they do not address the unique challenges in delegating access (e.g., how to share and revoke access to video data at particular timescales) or deleting videos. Similarly, Ullah, Rinner, and Marcenaro proposed using signencryption [56] to encrypt and sign video frames at once, but they also did not consider the challenges related to video deletion and delegation. CaCTUs addresses these issues by storing the system secrets at the owner's smartphone and using a secure key rotation scheme to enable fine-grained delegation.

Finally, PrivacyCam [13], TrustCAM [57], and Pinto [62] do not fully address the privacy challenges in a smart camera system as they attempt to solve a problem of a different nature and setting. These systems try to protect the anonymity of people or of vehicle license plates recorded in public spaces by detecting privacy sensitive regions and selectively encrypting or blurring them while leaving the rest of the video frame unperturbed. However, even this approach is limited in protecting anonymity: people can still be identified through their clothes or actions, and blurred videos still disclose behavior patterns such as when users are at home. Moreover, the system does not address data privacy of the video recordings.

### 7.2 Commercial Systems

As of May 2022, among the most popular smart camera systems commercially available, we find mainly doorbells, indoor and outdoor security cameras, as well as some baby monitors from the following companies; Arlo [5], Blink (Amazon company) [7], Eufy [18], Nest (Google company) [22], Ring (Amazon company) [51], or Wyze [61]. All these systems offer a cloud-based solution for video storage and streaming, with the exception of Eufy that primarily uses a base station installed in users' homes to store videos (Eufy also has a cloud storage available in option).

We analyzed these commercial systems with the public information sourced from product data sheets and technical reports or other documents published by the company. Our analysis revealed that some of these systems while they provide security, fail to provide *end-to-end encryption*. End-to-end encryption is an encryption mechanism through which two parties (can be extended to more than two) are able to communicate between each other with the guarantee that they are the only ones able to see the content of their exchange, i.e., no third party is able to access the data that is being communicated. End-to-end encryption is generally achieved by encrypting the data with encryption keys only known by the parties that should be able to access and see that data, essentially transforming the issue into managing properly the encryption keys. In the context of smart camera systems, we are mainly concerned by the fact that videos recorded by the system are end-to-end encrypted between the camera device and the users, i.e., the service and cloud providers as well as any other third party are unable to access and see these videos in clear.

Out of the six companies considered in this section (see Table 5), only Eufy and Ring offer some commercial systems that let users have access to end-to-end encryption. The other commercial systems require users to trust the service or cloud providers to handle at some point their data in clear; videos are usually encrypted *in transit* (TLS) between the camera and the cloud. Then, in the cloud, videos are encrypted *at rest* with some encryption mechanism put in place by the service provider. This encryption sometimes leverages an underlying mechanism offered by the cloud provider (e.g., S3 bucket encryption on AWS). The bottom point of this approach (encryption in transit and then at rest) is that users are not in control of the encryption keys; they need to trust the service or cloud providers with their videos. This is why in these systems access mediation is enforced by the service provider; users need to authenticate

themselves, proving who they are in order to decrypt their videos. This is also the technical explanation on how these companies are able to serve a valid search warrant and provide videos to authorities; they are the ones mediating the access control mechanism—users are not solely in control here, which does not align with the threat model considered in our work. Moreover, similar trust in the platform is needed when users want to delete their recorded footage.

Eufy sells a smart camera system that stores and encrypts the videos on a base station directly located in users' homes. While it seems that this system takes into account different privacy aspects, limitations might appear. For instance, it is unclear how the keys used to encrypt and decrypt the videos are generated, where and how they are stored, who can access them, how they are linked to an account and obtained by a user, or if they are rotated or not, specifically in light of the incident that occurred in 2021; when Eufy deployed an update on their servers that caused users to see other users' feeds in their apps [45]. Moreover, it is unclear if they are using end-to-end encryption—as defined earlier—, how they can disclose video recordings when being served with a valid search warrant [19].

Ring offers end-to-end encryption on an opt-in basis on some of their devices, the enrollment protocol to establish the root of trust between the owner's smartphone and the camera device has some similarities with the one presented in CaCTUs. However, by enabling end-to-end encryption, Ring users lose access to some features of their system (shared users unable to see videos, impossibility to access live stream from different devices at the same time, etc.), specifically, the end-to-end protocols of Ring were not designed to enable fine grained delegation as opposed to CaCTUs [50].

To summarize, the main differences with CaCTUs are that these commercial systems have a different trust model that the one we considered and entrust at some point in their life cycle the service or cloud providers with users data and access control to it. With CaCTUs, we proposed techniques that do not require users to operate in such threat models or trust the service or cloud providers, and still be offered the same feature set as of commercial systems.

## 8 CONCLUSION

This paper presented privacy-preserving protocols to design CaCTUs; a smart camera system that provides users with full control over their system, isolation and protection of the access to their video footage, deletion and factory reset, as well as peer-to-peer and fine-grained delegation of their videos. However, these protocols are not only restricted to video streaming but are readily extensible to IoT devices recording other time series data. Specifically, we showed how to leverage physical and direct pairing between devices to initialize systems without having to trust third parties, use performance-aware cryptographic algorithms for real time data streaming, and efficiently manage encryption keys through a binary key tree to support key rotation and enable data deletion and fine-grained (i.e., on the order of seconds) peer-to-peer delegation. CaCTUs serves as an existence proof that smart camera systems and other IoT devices need not compromise the privacy of users to be afforded the modern capabilities that commercial systems offer today.

National Science Foundation, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for government purposes notwithstanding any copyright notation hereon.

## REFERENCES

[1] Vijay K Adhikari, Yang Guo, Fang Hao, Volker Hilt, Zhi-Li Zhang, Matteo Varvello, and Moritz Steiner. 2014. Measurement study of Netflix, Hulu, and a tale of three CDNs. *IEEE/ACM Transactions on Networking* 23, 6 (2014), 1984–1997.

[2] Rana Alharbi and David Aspinall. 2018. An IoT analysis framework: An investigation of IoT smart cameras' vulnerabilities. In *Living in the Internet of Things: Cybersecurity of the IoT - 2018*. Institution of Engineering and Technology, London, 1–10. https://doi.org/10.1049/cp.2018.0047

[3] John G Apostolopoulos, Wai-tian Tan, and Susie J Wee. 2002. Video streaming: Concepts, algorithms, and systems. *HP Laboratories, report HPL-2002-260* (2002).

[4] Arlo. 2021. Arlo - Investor Relations. https://investor.arlo.com/ir-home/default.aspx Last Accessed: 2021-04-06.

[5] Arlo. 2022. Wireless Smart Home HD Security Cameras, Lights and Doorbells | Arlo. https://www.arlo.com/en-us

[6] Sam Biddle. 2019. For Owners of Amazon's Ring Security Cameras, Strangers May Have Been Watching Too. *The Intercept* (Jan. 2019). https://theintercept.com/2019/01/10/amazon-ring-security-camera/ Last Accessed: 2020-09-08.

[7] Blink. 2022. Blink for Home. https://blinkforhome.com

[8] European Data Protection Board. 2019. Guidelines 3/2019 on processing of personal data through video devices. https://edpb.europa.eu/sites/edpb/files/consultation/edpb_guidelines_201903_videosurveillance.pdf Last Accessed: 2021-04-19.

[9] Thomas Brewster. 2018. Smart Home Surveillance: Governments Tell Google's Nest To Hand Over Data 300 Times. *Forbes* (2018). https://www.forbes.com/sites/thomasbrewster/2018/10/13/smart-home-surveillance-governments-tell-googles-nest-to-hand-over-data-300-times/ Last Accessed: 2020-09-08.

[10] Lauren Bridges. 2021. Amazon's Ring is the largest civilian surveillance network the US has ever seen | Lauren Bridges. *The Guardian* (May 2021). http://www.theguardian.com/commentisfree/2021/may/18/amazon-ring-largest-civilian-surveillance-network-us Last Accessed: 2021-05-18.

[11] Dell Cameron. 2019. Amazon Is Marketing Face Recognition to Police Departments Partnered With Ring: Report. *Gizmodo* (2019). https://gizmodo.com/amazon-is-marketing-face-recognition-to-police-departme-1839073749 Last Accessed: 2020-09-08.

[12] Dell Cameron and Dhruv Mehrotra. 2019. Ring's Hidden Data Let Us Map Amazon's Sprawling Home Surveillance Network. *Gizmodo* (2019). https://gizmodo.com/ring-s-hidden-data-let-us-map-amazons-sprawling-home-su-1840312279 Last Accessed: 2020-01-13.

[13] A. Chattopadhyay and T. E. Boult. 2007. PrivacyCam: a Privacy Preserving Camera Using uCLinux on the Blackfin DSP. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*. 1–8. https://doi.org/10.1109/CVPR.2007.383413

[14] Catalin Cimpanu. 2019. Hackers keep dumping Ring credentials online 'for the giggles'. *ZDNet* (2019). https://www.zdnet.com/article/hackers-keep-dumping-ring-credentials-online-for-the-giggles/ Last Accessed: 2020-01-13.

[15] Dani Deahl. 2019. Ring let employees watch customer videos, claim reports. *The Verge* (Jan. 2019). https://www.theverge.com/2019/1/10/18177305/ring-employees-unencrypted-customer-video-amazon Last Accessed: 2020-09-08.

[16] W. Diffie and M. Hellman. 1976. New directions in cryptography. *IEEE Transactions on Information Theory* 22, 6 (Nov. 1976), 644–654. https://doi.org/10.1109/TIT.1976.1055638

[17] Bill Dirks, Micheal H. Schimek, Hans Verkuil, Martin Rubli, Andy Walls, and Mauro Carvalho Chehab. 2009. Video for Linux Two API Specification. https://www.linuxtv.org/downloads/legacy/video4linux/API/V4L2_API/spec-single/v4l2.html Last Accessed: 2021-03-21.

[18] eufy. 2022. eufy, Smart Home Simplified. https://us.eufylife.com

[19] eufy. 2022. Privacy Commitment. https://us.eufylife.com/pages/privacy-commitment

[20] Ken Florance. 2016. About Netflix - How Netflix Works With ISPs Around the Globe to Deliver a Great Viewing Experience. *About Netflix* (March 2016). https://about.netflix.com/en/news/how-netflix-works-with-isps-around-the-globe-to-deliver-a-great-viewing-experience Last Accessed: 2021-05-28.

[21] Yesenia Flores. 2020. Bad Neighbors? How Amazon's Ring Video Surveillance Could be Undermining Fourth Amendment Protections. https://www.californialawreview.org/amazon-ring-undermining-fourth-amendment/ Last Accessed: 2021-05-18.

[22] Nest & Google. 2022. Nest & Google - The best of Google. The best of Nest. https://store.google.com/us/category/google_nest?hl=en-US

[23] Andy Greenberg. 2021. Two Cases' Lessons: If Cops Don't Know What You Encrypted, They Can't Make You Decrypt It. https://www.forbes.com/sites/andygreenberg/2012/02/24/two-cases-lessons-if-cops-dont-know-what-you-encrypted-they-cant-make-you-decrypt-it/ Last Accessed: 2021-08-02.

[24] Matthew Guariglia and Matthew Maas. 2021. LAPD Requested Ring Footage of Black Lives Matter Protests. *Electronic Frontier Foundation* (Feb. 2021). https://www.eff.org/deeplinks/2021/02/lapd-requested-ring-footage-black-lives-matter-protests Last Accessed: 2021-05-11.

[25] I. Haider and B. Rinner. 2017. Private Space Monitoring with SoC-Based Smart Cameras. In *2017 IEEE 14th International Conference on Mobile Ad Hoc and Sensor Systems (MASS)*. IEEE, Orlando, FL, USA, 19–27. https://doi.org/10.1109/MASS.2017.15

[26] John Herrman. 2020. Who's Watching Your Porch? *The New York Times* (Jan. 2020). https://www.nytimes.com/2020/01/19/style/ring-video-doorbell-home-security.html Last Accessed: 2021-04-06.

[27] Brian Huseman. 2020. Huseman reply to Wyden, Markey, Van Hollen, Coons, Peters letter about Ring's Data Security Practices. https://regmedia.co.uk/2020/01/08/ringsenateresponse.pdf Last Accessed: 2021-04-06.

[28] Jonathan Katz and Yehuda Lindell. 2014. Introduction to Modern Cryptography, Second Edition. Chapman & Hall/CRC.

[29] Cathie Keck. 2019. Amazon's Ring Security Cameras May Have Let Employees Spy on Customers: Report. *Gizmodo* (2019). https://gizmodo.com/amazons-ring-security-cameras-may-have-let-employees-sp-1831658669 Last Accessed: 2020-09-08.

[30] Peter Kietzmann, Lena Boeckmann, Leandro Lanzieri, Thomas C. Schmidt, and Matthias Wählisch. 2021. A Performance Study of Crypto-Hardware in the Low-end IoT. In *Proceedings of the 2021 International Conference on Embedded Wireless Systems and Networks*. Junction Publishing, USA, 79–90. https://doi.org/10.5555/3451271.3451279

[31] Paul Kocher. 2011. Complexity and the challenges of securing SoCs. In *2011 48th ACM/EDAC/IEEE Design Automation Conference (DAC)*. Association for Computing Machinery, New York, NY, USA, 328–331.

[32] Rachel Kraus. 2019. Ring watched your kids trick or treat and then bragged about it. *Mashable* (Nov. 2019). https://mashable.com/article/ring-halloween-surveillance/ Last Accessed: 2021-05-11.

[33] David Kravets. 2016. Indefinite prison for suspect who won't decrypt hard drives, feds say. *Ars Technica* (May 2016). https://arstechnica.com/tech-policy/2016/05/feds-say-suspect-should-rot-in-prison-for-refusing-to-decrypt-drives/ Last Accessed: 2021-05-31.

[34] Colin Lecher. 2019. Ring reportedly outed camera owners to police with a heat map. *The Verge* (Dec. 2019). https://www.theverge.com/2019/12/3/20993814/ring-user-location-heat-map-police-privacy-tool-camera-owners Last Accessed: 2020-09-08.

[35] California State Legislature. 2018. TITLE 1.81.5. California Consumer Privacy Act of 2018 [1798.100 - 1798.199.100].

[36] California State Legislature. 2020. The California Privacy Rights Act of 2020.

[37] J.M. McCune, A. Perrig, and M.K. Reiter. 2005. Seeing-is-believing: using camera phones for human-verifiable authentication. In *2005 IEEE Symposium on Security and Privacy (S P'05)*. 110–124. https://doi.org/10.1109/SP.2005.19

[38] Alfred Ng. 2019. Amazon's Ring wanted to use 911 calls to activate its video doorbells. *CNET* (Sept. 2019). https://www.cnet.com/home/smart-home/amazons-ring-wanted-to-use-911-calls-to-activate-its-video-doorbells/ Last Accessed: 2021-05-18.

[39] Alfred Ng. 2019. Ring let police view map of video doorbell installations for over a year. *CNET* (2019). https://www.cnet.com/news/ring-gave-police-a-street-level-view-of-where-video-doorbells-were-for-over-a-year/ Last Accessed: 2020-01-13.

[40] Jakob Nielsen. 2000. Why You Only Need to Test with 5 Users. https://www.nngroup.com/articles/why-you-only-need-to-test-with-5-users/ Last Accessed: 2021-08-30.

[41] Jakob Nielsen and Thomas K. Landauer. 1993. A mathematical model of the finding of usability problems. , 206–213 pages. https://doi.org/10.1145/169059.169166

[42] House of Commons of Canada. 2020. Bill C-11 (First Reading).

[43] European Parliament and Council of the European Union. 2009. ePrivacy Directive - Directive 2009/136/EC.

[44] European Parliament and Council of the European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance).

[45] Ben Patterson. 2021. Anker's Eufy division pledges to bolster security after privacy snafu. https://www.techhive.com/article/579437/ankers-eufy-division-pledges-to-bolster-security-following-privacy-snafu.html

[46] Kari Paul. 2019. Amazon's doorbell camera Ring is working with police – and controlling what they say. *The Guardian* (Aug. 2019). https://www.theguardian.com/technology/2019/aug/29/ring-amazon-police-partnership-social-media-neighbor Last Accessed: 2020-01-13.

[47] The Constitution Project. 2007. Guidelines for public video surveillance - A guide to protecting communities and preserving civil liberties. https://archive.constitutionproject.org/pdf/Video_Surveillance_Guidelines_Report_w_Model_Legislation4.pdf Last Accessed: 2021-04-19.

[48] Ring. 2019. Ring Video Doorbells Get 15+ Million Dings This Halloween and Capture Cute Costumes and Fun Pranks. *The Ring Blog* (Nov. 2019). https://blog.ring.com/neighborhood-stories/ring-video-doorbells-get-15-million-dings-this-halloween-and-capture-cute-costumes-and-fun-pranks/ Last Accessed: 2021-05-18.

[49] Ring. 2021. Active Agency Map. https://www.google.com/maps/d/viewer?mid=1eYVDPh5itXq5acDT9b0BVeQwmESBa4cB Last Accessed: 2021-05-18.

[50] Ring. 2021. Ring End-to-End Encryption Whitepaper. *Ring Whitepaper* (July 2021). ring.com/end-to-end-encryption

[51] Ring. 2022. Home Security Systems | Smart Home Automation. https://ring.com

[52] Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. 1983. Cryptographic communications system and method. https://patents.google.com/patent/US4405829/en Last Accessed: 2021-08-18.

[53] Lucas Ropek. 2021. A Home Security Worker Hacked Into Surveillance Systems to Watch People Have Sex. *Gizmodo* (Jan. 2021). https://gizmodo.com/a-home-security-worker-hacked-into-surveillance-systems-1846111569 Last Accessed: 2021-01-23.

[54] Hani Salah, Sandra Zimmermann, and Juan A. Cabrera G. 2020. Chapter 5 - Content distribution. , 93-103 pages. https://doi.org/10.1016/B978-0-12-820488-7.00016-5

[55] Gérard Scorletti. 2016. Traitement du Signal. (Sept. 2016), 203 pages. https://cel.archives-ouvertes.fr/cel-00673929 Last Accessed: 2021-05-11 (Lecture material in French).

[56] Subhan Ullah, Bernhard Rinner, and Lucio Marcenaro. 2017. Smart cameras with onboard signcryption for securing IoT applications. In *2017 Global Internet of Things Summit (GIoTS)*. IEEE, Geneva, Switzerland, 1–6. https://doi.org/10.1109/GIOTS.2017.8016279

[57]  T. Winkler and B. Rinner. 2010. TrustCAM: Security and Privacy-Protection for an Embedded Smart Camera Based on Trusted Computing. In *2010 7th IEEE International Conference on Advanced Video and Signal Based Surveillance*. IEEE, Boston, MA, USA, 593–600. https://doi.org/10.1109/AVSS.2010.38

[58]  Thomas Winkler and Bernhard Rinner. 2015. Secure embedded visual sensing in end-user applications with TrustEYE.M4. In *2015 IEEE Tenth International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP)*. Institute of Electrical and Electronics Engineers, Singapore, 1–6. https://doi.org/10.1109/ISSNIP.2015.7106934

[59]  Ron Wyden, Chris Van Hollen, Edward Markey, Christopher Coons, and Gary Peters. 2019. Wyden, Markey, Van Hollen, Coons, Peters Question Ring's Data Security Practices. https://www.wyden.senate.gov/news/press-releases/wyden-markey-van-hollen-coons-peters-question-rings-data-security-practices Last Accessed: 2021-04-06.

[60]  Wyze. 2018. Wyze Cam - Our Story. https://wyze.com/our-story Last Accessed: 2021-04-06.

[61]  Wyze. 2022. Wyze. https://www.wyze.com/

[62]  Hyunwoo Yu, Jaemin Lim, Kiyeon Kim, and Suk-Bok Lee. 2018. Pinto: Enabling Video Privacy for Commodity IoT Cameras. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. Association for Computing Machinery, New York, NY, USA, 1089–1101. https://doi.org/10.1145/3243734.3243830

# A   NOTATION

# B   EXPERIMENTAL SETUP

*Camera Device.* On a *Raspberry Pi 4 Model B Rev 1.1* (Broadcom BCM2711, 1.5 GHz quad-core Cortex-A72 ARM v7 64-bit, 2GB RAM), we used the Video4Linux2 driver [17] to interface with the camera sensor and capture frames that are then encrypted using OpenSSL3.0[4]. The Raspberry Pi Camera Module v2 that we used has a still resolution of 8 Megapixels, a sensor resolution of 3280 × 2464 pixels, and supports the three following video modes 1080 p/30 fps, 720 p/60 fps, and 480 p/90 fps (respectively video quality and maximum frame rate).

*Android Smartphone.* We used a *Nokia 4.2* smartphone with Android 10 on which we have installed the implemented application. In this application, we use C native libraries that we have cross-compiled, and C code to download and decrypt the frames. We leveraged the MediaCodec class[5] to perform the encoding and decoding of video files, as well as the Quirc[6] and Bluetooth libraries to perform the pairing.

*Cloud Storage.* An *AWS EC2 t3.small* instance was used to deploy a Nginx web server. Upon request, we serve the list of encrypted frames that were recorded during the time frame specified in the request.

# C   FUNCTIONAL USER EVALUATION PROTOCOL

All institutional requirements were met for this functional user evaluation of CaCTUs. We obtained approval from the Institutional Review Board (IRB) of our university and a consent form was signed by the participants at the beginning of their session. We also tested the protocol with coworkers and collaborators beforehand to identify possible limitations.

All the material was provided to the participants that were guided by a researcher through the different tasks to perform. We introduced each task with a real-life scenario to help the participants behave as if they were using the system in their real life. To collect feedback, we observed a talk aloud process asking the participants to express aloud what they are doing or looking for while performing the task, allowing us to better identify potential issues in the system. Between each task, we also asked specific questions about the process that had just been completed.

## C.1   Before Each Session

Before each session, we verified that all the material needed for the session was provided, working, and its initial state:

---

[4] https://www.openssl.org/
[5] https://developer.android.com/reference/android/media/MediaCodec
[6] https://github.com/dlbeer/quirc

Table 6. Notation used by CaCTUs cryptographic constructions.

| Symbol | Name |
|---|---|
| $\|\|$ | Concatenate function |
| $\oplus$ | XOR function |
| $AES256Dec$ | Symmetric authenticated decryption function |
| $AES256Enc$ | Symmetric authenticated encryption function |
| $C_i$ | $i$th cipher frame |
| $check$ | Check if a party knows the secret key to the corresponding public key |
| $d_\mathcal{K}$ | Depth key tree |
| $\delta_\mathcal{K}$ | Epoch size |
| $escrow\ material$ | Escrow material encrypted with a passphrase |
| $Extract$ | Key extraction function from $\mathcal{K}$ |
| $F_i$ | $i$th frame |
| $gen$ | Generate function |
| $h_{PK_d}$ | Hash of the public key of a delegatee |
| $h_{PK_f}$ | Hash of the factory-generated public key of the camera |
| $h_{PK_o}$ | Hash of the public key of the owner |
| $hash$ | Hash function |
| $HKDF$ | Hash-based key derivation function |
| $init$ | Initialization function |
| $IV_i$ | $i$th initialization vector |
| $\mathcal{K}$ | Binary key tree |
| $k_i$ | Symmetric encryption key for $F_i$ |
| $passphrase$ | Passphrase encrypting the escrow material |
| $RandBytes$ | Random bytes generation function |
| $RSA$ | Rivest–Shamir–Adleman encryption and signature scheme |
| $secrets$ | Secrets sent by the owner to the camera during the initialization |
| $seed\ key$ | Encryption material owned by the root node of $\mathcal{K}$ |
| $\sigma$ | Asymmetric signature of a block of frames |
| $Sign$ | Asymmetric signature function |
| $(SK_c, PK_c)$ | Asymmetric key pair of the camera |
| $(SK_d, PK_d)$ | Asymmetric key pair of a delegatee |
| $(SK_f, PK_f)$ | Factory-generated asymmetric key pair of the camera |
| $(SK_o, PK_o)$ | Asymmetric key pair of the owner |
| $\tau_i$ | $i$th tag |
| $t_i$ | Timestamp of $i$th frame |
| $[t_j, t_{j+1})$ | Time interval of the $j$th epoch |
| $Verify$ | Signature verification function |
| $wifi\ credentials$ | Wifi credentials of the owner's network |

(1) **Camera Device:** 1 *Raspberry Pi 4 Model B 2GB*, with camera sensor and our software installed.

(2) **Android Smartphone Devices:** 3 *Nokia 4.2* with application pre-installed, wifi configured, location enabled (for Bluetooth discovery of nearby devices), and Bluetooth disabled with no prior device paired.

(3) **Other:** Piece of paper and pen provided (to write down the recovery passphrase).

## C.2 During Each Session

We first started by asking some preliminary questions about their background and their experience with smart camera systems. Next, we explained that the objective of the session was to evaluate the functional usability of the smart camera system that we designed. Details about privacy violations in current available systems were briefly described to help the participants understand the motivation of the project. Then, we presented at very high level how the system we implemented was giving back full control to the users by using end-to-end encryption and strict access mediation to the decryption keys. We also verified that the participants were either familiar with Android or we showed them how to navigate between applications on Android.

*C.2.1 Initialization.* You want to secure your home, so you just bought this new smart camera system online. You received the package with the camera and just downloaded the application on your smartphone. Go ahead with the rest of the configuration.

(1) What did you like about the initialization?
(2) What did you dislike about the initialization?
(3) Any further comments?

*C.2.2 System Usage.* You have your new smart camera all set up, so now you want to be able to see what is happening inside/outside of your home. For that you open the application to view the live streaming and access the different functionalities of the application.

(1) How do you feel about the quality of the video streaming?
(2) What is your opinion about the following statements (Likert scale: strongly disagree, disagree, neutral, agree, strongly agree)?
   (a) I am happy with the image quality of the streaming.
   (b) I am happy with the latency of the streaming.
   (c) I am happy with the frame rate of the streaming.
(3) What did you like about the usage of the system/application?
(4) What did you dislike about the usage of the system/application?
(5) Any further comments?

*C.2.3 Delegation.* You want to give access to someone else to the live streaming of your camera, as you are going on vacation abroad. They have downloaded the application on their phone, you need to add them as a new delegatee on your application.

(1) What did you like about the delegation process?
(2) What did you dislike about the delegation process?
(3) What is your opinion about the following statement: It is easy to add a new delegatee (Likert scale: strongly disagree, disagree, neutral, agree, strongly agree)?
(4) Would you like to see any change in the delegation process or the options for the access control?
(5) Who would you typically add as a delegatee and for how long?
(6) What do you think of the granularity of the delegation control?
(7) Any further comments?

*C.2.4 Access Recovery.* Unfortunately, on your way back home during the layover, you lost your smartphone, which was the device you used to access your camera system, and when you come back home, you figured out that someone has broken in during your holidays and robbed you. You buy a new smartphone, install back the application on it, and want to recover access to your system to see what happened. Luckily, you wrote down your recovery passphrase.

(1) What did you like about the recovery process?
(2) What did you dislike about the recovery process?
(3) What is your opinion about the following statement: It is important for me to have a recovery process. (Likert scale: strongly disagree, disagree, neutral, agree, strongly agree)?
(4) Any further comments?

*C.2.5 Factory Reset.* Finally, you want to reconfigure your system as you lost your smartphone, but before you want to make sure to factory reset the system.

(1) What did you like about the factory reset process?
(2) What did you dislike about the factory reset process?
(3) What is your opinion about the following statement: It is important for me to have a factory reset process. (Likert scale: strongly disagree, disagree, neutral, agree, strongly agree)?
(4) Any further comments?

## C.3 After Each Session

After each session, we made sure that every device was reset and back into its initial state, as if the session had not occurred.

## D STORYBOARD

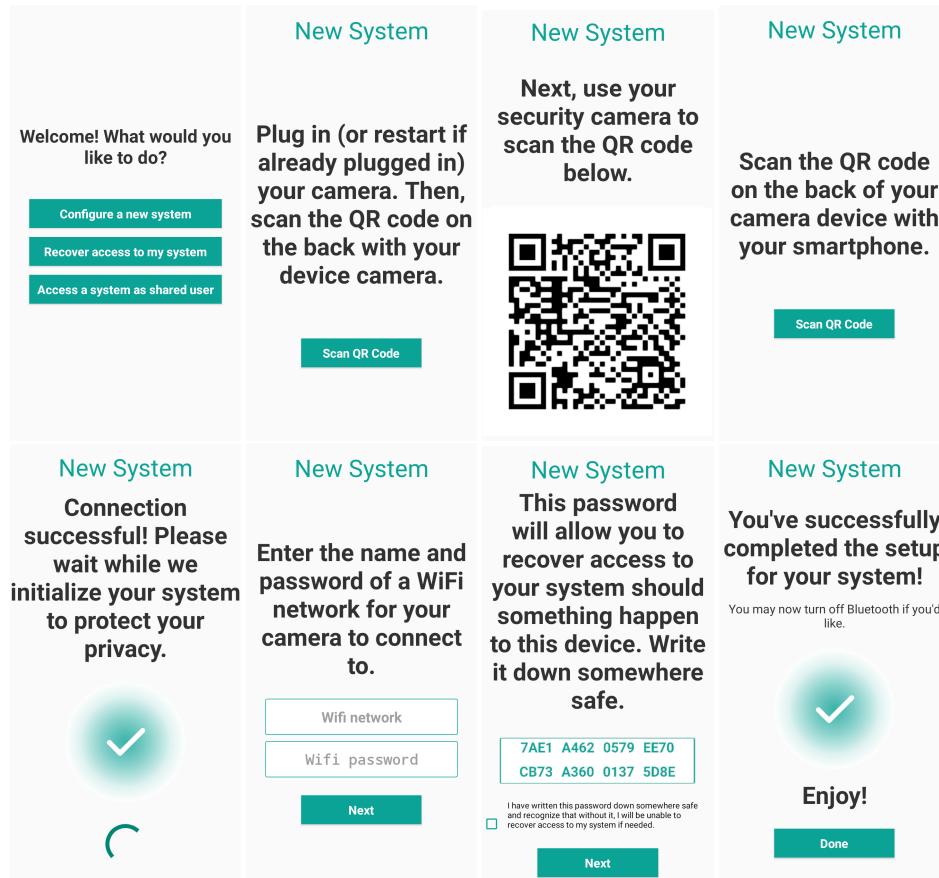Following is the storyboard of the CaCTUs's smartphone application.

Fig. 6. Screenshots of the CaCTUs smartphone application of the owner during initialization of the system.



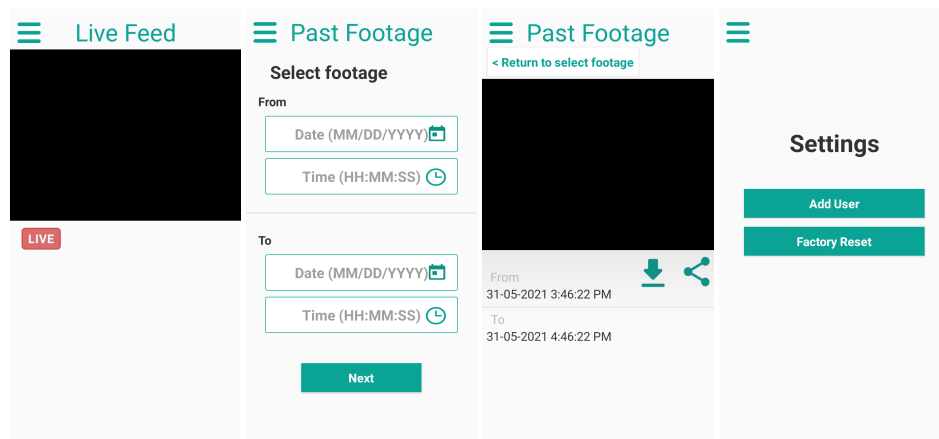Fig. 7. Screenshots of the CaCTUs smartphone application: home page (live feed), access to past footage, and settings.
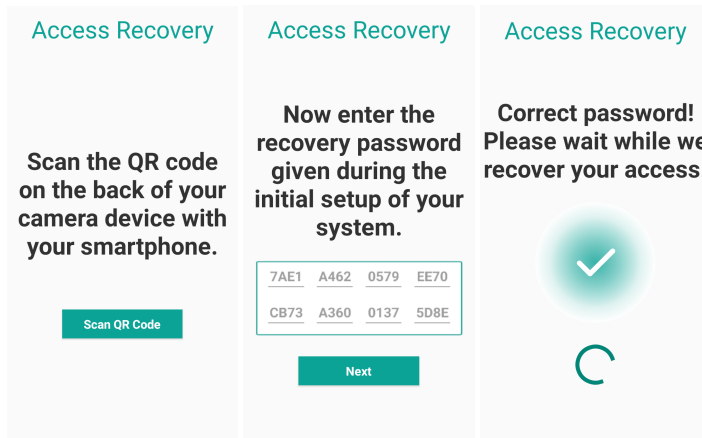
Fig. 8. Screenshots of the CaCTUs smartphone application of the owner during access recovery.
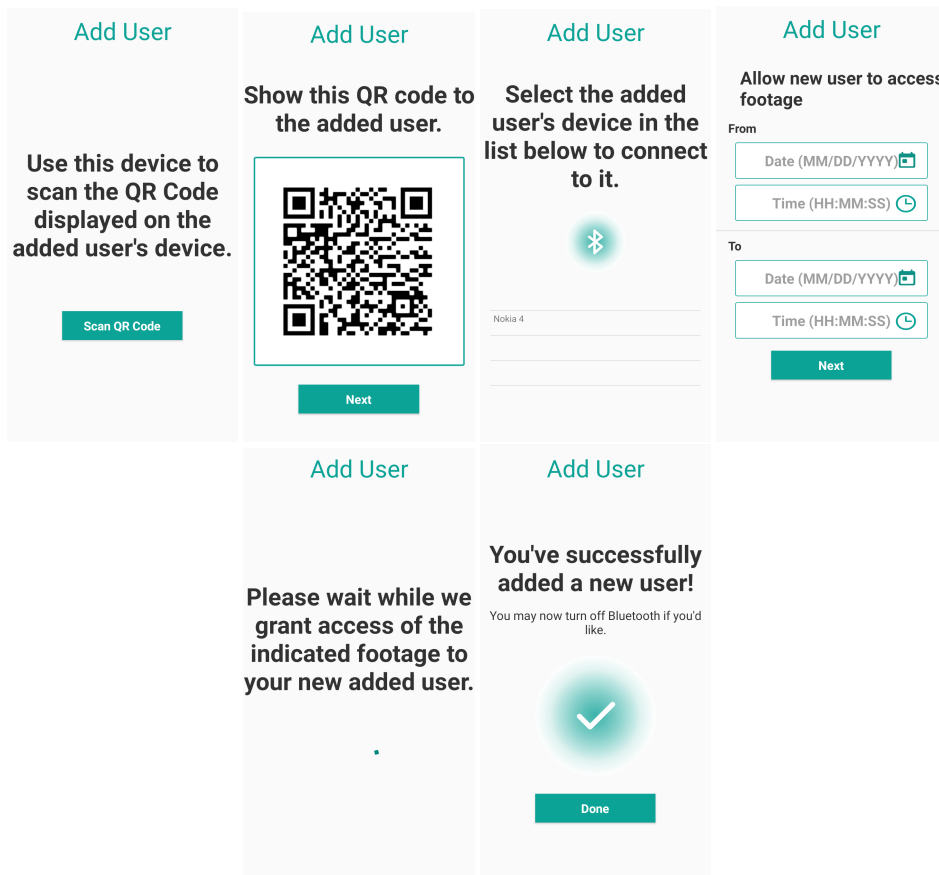


Fig. 9. Screenshots of the CaCTUs smartphone application of the owner during delegation.

Welcome! What would you like to do?

Configure a new system

Recover access to my system

Access a system as shared user

## Add User

**Show this QR Code to the admin user.**

**Next**

## Add User

**Now use this device to scan the QR Code displayed on the admin's device.**

Scan QR Code

## Add User

**Enable bluetooth in order to connect with the admin's device.**

Enable Bluetooth

## Add User

**Please wait while the admin user finishes to delegate you access to the system.**

⟩

## Add User

**You've successfully been added as a new user!**

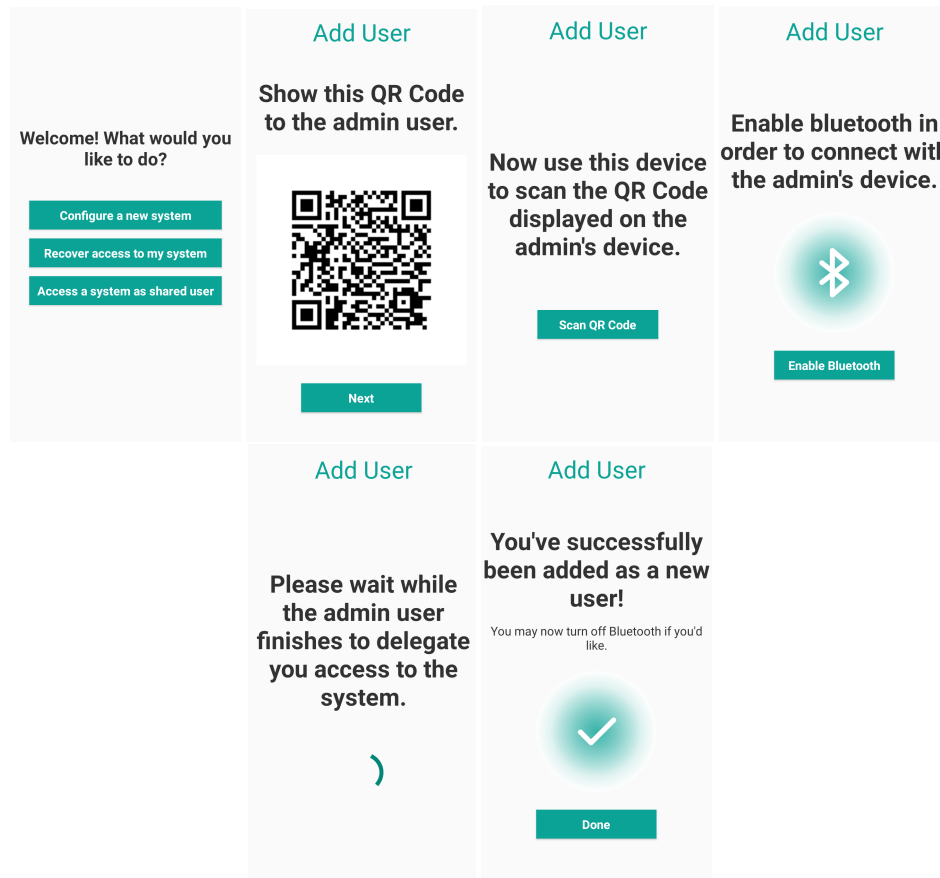You may now turn off Bluetooth if you'd like.

Done

Fig. 10. Screenshots of the CaCTUs smartphone application of the delegatee during delegation.