

Multi-Stage Group Key Distribution and PAKEs: Securing Zoom Groups against Malicious Servers without New Security Elements

Cas Cremers
CISPA Helmholtz Center for
Information Security
Saarbrücken, Germany
cremers@cispa.de

Eyal Ronen
Tel Aviv University
Tel Aviv, Israel
er@eyalro.net

Mang Zhao
CISPA Helmholtz Center for
Information Security
Saarbrücken, Germany
mang.zhao@cispa.de

Abstract—Video conferencing apps like Zoom have hundreds of millions of daily users, making them a high-value target for surveillance and subversion. While such apps claim to achieve some forms of end-to-end encryption, they usually assume an incorruptible server that is able to identify and authenticate all the parties in a meeting. Concretely this means that, e.g., even when using the “end-to-end encrypted” setting, malicious Zoom servers could eavesdrop or impersonate in arbitrary groups.

In this work, we show how security against malicious servers can be improved by changing the way in which such protocols use passwords (known as *passcodes* in Zoom) and integrating a password-authenticated key exchange (PAKE) protocol.

To formally prove that our approach achieves its goals, we formalize a class of cryptographic protocols suitable for this setting, and define a basic security notion for them, in which group security can be achieved assuming the server is trusted to correctly authorize the group members. We prove that Zoom indeed meets this notion. We then propose a stronger security notion that can provide security against malicious servers, and propose a transformation that can achieve this notion. We show how we can apply our transformation to Zoom to provably achieve stronger security against malicious servers, notably without introducing new security elements.

1. Introduction

Video conferencing apps such as Zoom are globally used by hundreds of millions of users on a daily basis [1], and aim to use cryptographic protocols to achieve some forms of end-to-end encryption. While there have been many recent advancements in highly-secure messaging protocols such as Signal, their core protocols are typically not suitable for real-time group applications, such as video conferencing, which have fundamentally different requirements involving real-time constraints, robustness, and usability.

In practice, real-time group protocols used in real-world widely deployed applications such as Zoom incorporate design choices based on real-time requirements that include assigning a group leader role to some participants, relying on key distribution instead of key agreement, and using simplified key evolution mechanisms. These design choices

improve features like robustness and usability, and enable real-time communication, at the cost of lower security guarantees compared to some state-of-the-art secure messaging protocols.

Nevertheless, many real-time group protocols explicitly claim that they can provide a form of end-to-end security. For example, in Zoom’s case, the ‘end-to-end security’ option in the app’s settings is explained as “Encryption key stored on your local device. No one else can obtain your encryption key, not even Zoom.” However, it was shown that a malicious Zoom server can eavesdrop or impersonate in groups [2], [3]. The underlying reason is that in practice, a Zoom server acts as the *sole root of trust* for the authenticity of users’ public keys and messages, and implicitly to those that are used to distribute group-specific public keys, which in turn are used by the leader to distribute the group key. Thus, if the Zoom server replaces some of these public keys, it can in fact learn your encryption key.

In this work, we propose a transformation to improve the security of a class of protocols against malicious servers, without introducing new security elements or even new message flows. We achieve this by reworking the way in which such protocols use passwords (known as *passcodes* in Zoom). In the Zoom protocol, the server inherently needs to know the group password and uses it to enforce access control for the group. We propose a modification in which the server no longer knows the password, which is distributed only to the group members and is used by them directly for access control. These passwords can be distributed as one would have done currently (e-mail, messaging app, phone, calendar appointment).

In our new threat model, we can no longer rely on the server providing a priori secure channels between group members. Instead, we employ password-authenticated key exchange to prevent offline guessing attacks on the protocol. We then formally prove that the transformed Zoom protocol achieves a strong form of security even in the presence of malicious servers.

At the technical level, we develop syntax and security notions for multi-stage group key distribution protocols, of which Zoom can be seen as an instance. We show under which assumptions the Zoom protocol (version 4.0) can

be proven secure in our multi-stage group key distribution model, and how the known attacks fit into this picture. We design a generic transformation that can turn any protocol in our class into a more secure protocol, for which we define a stronger security notion. We show that our transformation preserves the original security property, and provides security against malicious servers. We apply our transformation to the Zoom protocol, resulting in the ZoomPAKE protocol. We also show how the ZoomPAKE protocol prevents the attacks possible on the Zoom protocol.

Our main contributions are:

- We develop a solution to improve the security of Zoom-like apps against malicious servers, without introducing new security elements. The core observation is that Zoom already uses group-specific passwords, but they are by design known to the server. By leveraging techniques from password-authenticated key exchange, we can get rid of the reliance on the server for trusted channels.
- To formally prove the security of our solution, we need to develop substantial machinery. We propose a formal model and syntax of multi-stage group key distribution protocols, called mGKD, of which Zoom can be seen as an instance. For such protocols, we develop a basic security notion Sec-mGKD-pki, which assumes the server did not interfere with the public keys of a group’s participants, and prove that Zoom meets this notion. We show how real-world attacks manifest in this basic notion and notably how malicious zoom servers can manipulate groups.
- We formally prove that our transformation turns a protocol that is Sec-mGKD-pki secure into one that is also secure in a model that makes no assumptions on the server but only on the password, which we call Sec-mGKD-pw.
- We show how to efficiently apply our transformation to the Zoom version 4.0 protocol to obtain the ZoomPAKE protocol, in which the server no longer knows the password, and groups are protected against malicious servers.

Outline We discuss related work in Section 2 and notation in Section 3. In Section 4 we present our syntax for multi-stage group key distribution (mGKD) protocols and three security notions: basic Sec-mGKD-pki security, full end-to-end Sec-mGKD-pw security, and the combined Sec-mGKD-pw+ security. We show in Section 5 that the Zoom library can be modeled as a mGKD protocol and provably satisfies the basic Sec-mGKD-pki security, and show how impersonation attacks prevent it from satisfying the stronger Sec-mGKD-pw notion. In Section 6, we develop a generic transformation on any Sec-mGKD-pki secure mGKD protocol to achieve Sec-mGKD-pw and Sec-mGKD-pw+ security and apply it to Zoom.

We provide our customized notions w-PAKE security for PAKE and frob security for AEAD and proof sketches of all our theorems in appendix. We provide all details and full proofs in the long version [4].

2. Related Work

While there is a lot of adjacent related work that we will mention below, it turns out that there is surprisingly little directly related work in analysis of real-time group protocols. A number of surveys [5]–[10] examine numerous historical designs for secure group key establishment in different application scenarios. We identify three categories:

- centralized group key management/distribution protocols, where each group has a single trusted authority for group key generation and distribution.
- (continuous) group key agreement and distributed key management protocols, where every party in a group contributes to the group key generation and distribution.
- multi-factor key agreement and password-authenticated key exchange protocols, where the group key generation and distribution relies only on a secret that is used for authorization to a group.

We review each of these categories below.

2.1. Centralized Group Key Management Protocols

A centralized group key management (CGKM) protocol starts every group with a trusted authority, often referred to as the “Key Distribution Center” (KDC). The KDC is responsible for controlling for the whole group, e.g., member authentication, access control, and group key generation and distribution.

One of the first CGKM schemes is [11], [12]. In this approach, the KDC creates a “Group Key Packet” (GKP) for encrypting the communication payload with the help from the first group participant. The KDC sends the GKP to every party that wants to join the group and encrypts the new GKP to all group participants using the old one. To achieve forward secrecy, the KDC has to recreate the group whenever a participant leaves the group. After that, numerous tree-like CGKM constructions [13]–[18] were proposed to reduce computation cost. In these approaches, all trust resides in the KDC, which forms a single point of failure for compromise.

2.2. (Continuous) Group Key Agreement

Two important canonical group key agreement protocols are [19], [20]. Their constructions have a binary tree-like hierarchy, where the keying material of each party is a leaf node at the bottom of the tree and the shared group key is the top node of the tree. Every party can compute the key material on the path from its associated leaf node to the top using Diffie-Hellman Exchange (DHE). However, these designs are inherently synchronous: the initialization of the tree requires all parties to be online.

In the asynchronous secure messaging context, where participants might be offline and group keys need to be evolved, this approach does not work without modification. These drawbacks were lifted by the design in [21], leading to a line of papers in the continuous GKA (CGKA) domain, including [22]–[24] that focus on continuously evolving (ratcheting) group keys after the group establishment.

In general, ratcheting-like protocols such as CGKAs are impractical for real-time group applications, as they have to tolerate high amounts of packet loss and still being able to continue immediately when some packets arrive on time.

2.3. Multi-factor Key Agreement and Password-Authenticated Key Exchange

Multi-factor key agreement protocols often rely on three classes of human authentication factors: (1) something you know, e.g., passwords, (2) something you have, e.g., secure devices, and (3) something you are, e.g., biometric data. Among them, the password is possibly one of the most convenient means for sharing in practice, as it can be easily sent out-of-band, e.g., via email, in letters, or even in-person.

Human-chosen passwords are often low-entropy rather than uniformly at random. Password-Authenticated Key Exchange (PAKE) protocols are designed to allow some parties to establish a high-entropy session key with authentication based on a low-entropy shared password without being subject to offline guessing attacks. There are numerous modern and efficient 2-party PAKE constructions in the literature, such as CPace [25], [26] and SPAKE2 [27]–[29]. There are also several known group PAKE protocols [30]–[33]. However, the existing group PAKE protocols always require multiple rounds for the key agreement and is restricted to static groups. Thus, these group PAKE protocols are impractical for the real-time group applications, where the participants can freely join and leave the groups.

2.4. Existing Security Analysis for Zoom

In [2], [3], the authors describe several specific classes of impersonation attacks on end-to-end Zoom (version 2.3.1). First, a malicious meeting participant can impersonate any other participant inside this group, since there is no entity authentication in a group meeting. Second, the Zoom server can replay some messages and impersonate a legitimate user for a meeting. Third, if multiple users share a device, the Zoom server colluding with any user can impersonate any other users on the same device. Moreover, the authors also present a tampering attack based on potential implementation flaws and a Denial-of-Service attack. [2], [3] provide feasible countermeasures for each of above specific attacks. However, more general impersonation attacks by a malicious server are not considered in [2], [3].

The “accepted papers” page of Eurocrypt 2023 lists [34]. According to the title and abstract, [34] formally analyzes the security of the end-to-end Zoom protocol, focuses on various liveness properties, including dynamic group joining and leaving and the update of group roster and keys, but does not address assumptions on the server or PKI. However, at the time of this submission, [34] is not yet publicly available (neither on IACR eprint, arxiv, nor the author’s webpages), and thus we cannot comment on its details.

3. Preliminaries

We use the following notational conventions. We write ϵ to denote the empty string, and write $x \parallel y$ to denote the concatenation of strings x and y . We assume the existence of a reserved error symbol \perp . We write $y \leftarrow x$ for deterministic assignment or computation. We write $y \xleftarrow{\$} D$ for randomized sampling and $y \xleftarrow{\$} f(x)$ for non-deterministic computation. We write $\llbracket \cdot \rrbracket$ for a boolean statement that is either true (denoted by 1) or false (denoted by 0).

Within algorithms, we write “**require** C ” to denote that C is a requirement: if the condition C is not met, the algorithm undoes the execution and outputs \perp .

Moreover, we use several abbreviations. To avoid duplicating some longer variable names, we use some update-based variants of common operators: For a number x , we write $x++$ as a shorthand for $x \leftarrow x + 1$. For a set S and an element x , we write $S \overset{\pm}{\leftarrow} x$ for $S \leftarrow S \cup \{x\}$, and $S \bar{\leftarrow} x$ for $S \leftarrow S \setminus \{x\}$. We write PPT as an abbreviation of Probabilistic Polynomial Time.

4. multi-stage Group Key Distribution Protocols

In this section, we define our syntax for multi-stage Group Key Distribution mGKD protocols. Our class of mGKD protocols covers behaviors of a party for using real-time group services, such as long-term identity information generation, joining groups, group key rotation, and leaving groups. Then, we propose three security models that capture distinct security guarantees for real-time group services.

4.1. mGKD Definition

mGKD protocols are stateful interactive group communication protocols executed by a set of parties \mathcal{P} . Each party P must be uniquely identified by an identifier id_P . Each group is uniquely identified by an identifier gid . In each group, one party performs the leader role; all other group members perform the participant role. The role of each party in different groups might be distinct: a party can be leader in one group and participant in others. In practice, the leader is typically the so-called *host* that initiated the group. In this paper, we assume that each group gid is associated with a unique leader and that the group’s leader stays in the group for its entire duration. While one can in theory implement changing leaders by starting a new group, we leave the modeling and efficient implementation of multiple and changing leaders to future work.

Definition 1. A multi-stage group key distribution protocol $\text{mGKD} = (\text{SignUp}, \text{Schedule}, \text{Register}, \text{Join}, \text{Leave}, \text{KeyRotat})$ consists of the following algorithms:

Sign Up: $m_{\text{SignUp}} \xleftarrow{\$} \text{SignUp}(P)$ allows a (stateful) party P to initialize a long-term identity information for signing up. The private portion is locally stored. The public portion is output as an outgoing sign-up message m_{SignUp} .

Group Schedule: $m_{\text{GSch}}^{\text{gid}} \xleftarrow{\$} \text{Schedule}(P, \text{gid}, \text{gs})$ allows a (stateful) party P to take the role of the leader for

scheduling a group gid using a group secret gs . The output is an outgoing group schedule message m_{GSch}^{gid} for the server. The group secret gs is expected to be sent to authorized participants over secure out-of-band channels.

Register: Register = (Register-L, Register-P) consists of two sub-algorithms depending on the role of the caller:

- $m' \stackrel{s}{\leftarrow}$ Register-L(P, gid, gs, m) (resp. $m' \stackrel{s}{\leftarrow}$ Register-P(P, gid, gs, m)) allows a (stateful) party P to register for a group gid as leader (resp. participant) using a group secret gs and an incoming message m followed by group initialization. The output is an outgoing message m' .

Participant Join: Join = (Join-L, Join-P) allows a participant to interact with a leader for joining a group. This interactive phase consists of two sub-algorithms depending on the role of the caller:

- $m' \stackrel{s}{\leftarrow}$ Join-L($P, id_{P'}, gid, gs, m$) (resp. $m' \stackrel{s}{\leftarrow}$ Join-P($P, id_{P'}, gid, gs, m$)) allows a leader P (resp. a participant P) of the group gid to input a group secret gs and an incoming message m and to output an outgoing message m' for the participant (resp. the leader) P' .

Member Leave: Leave = (Leave-L, Leave-P) consists of two sub-algorithms depending on the role of the caller:

- Leave-L($P, gid, id_{P'}$) (resp. Leave-P($P, gid, id_{P'}$)) allows the leader (resp. a participant) P of the group gid to react to a party P' 's leaving. If $id_P = id_{P'}$, the leader (resp. the participant) P leaves the group gid and erases the corresponding state information.

Key Rotation: KeyRotat = (KeyRotat-L, KeyRotat-P) consists of two sub-algorithms depending on the role of the caller:

- $m_{KRot} \stackrel{s}{\leftarrow}$ KeyRotat-L(P, gid, m) allows the leader P of the group gid to input an incoming message m and to locally update the group key. The output is an outgoing message m_{KRot} that enables all participants of the same group gid to update group keys correspondingly.
- $m_{KRot} \stackrel{s}{\leftarrow}$ KeyRotat-P(P, gid, m) allows a participant P of the group gid to input an incoming message m and to locally update the group key. The output is an (optionally empty) outgoing message m_{KRot} .

We assume all incoming and outgoing messages of an mGKD protocol are publicly accessible; we leave this implicit as the concrete mechanisms can differ substantially between protocols, but could for example be a PKI or a “bulletin board” on the server. In contrast, the input group secret gs of the Schedule algorithm is expected to be chosen by the leader and be sent to authorized parties for joining the group over secure out-of-band channels. Before a party P joins a group gid , P has to register for this group, no matter whether P has previously joined the group gid and left. The Member Leave phase enables every party P to react to a participant P' leaving the group, notably, without any additional incoming message. This captures the scenario where the server might notify group members that a participant has left the group without sending any leave request due to unexpected network disconnection. The KeyRotat algorithm enables every party

to update their group key, the execution frequency of which can be decided in advance, in a regular schedule and/or when a party joins or leaves the group.

To model concurrent or sequential groups of a party P , let π_P^{gid} denote party P 's session with respect to the short-term group gid . In addition, each party P has an associated long-term state st_P that is shared by all of P 's sessions.

Definition 2. In a mGKD protocol, each party P has the following state variables. The long-term state variables are initialized during the Sign Up phase:

- $st_P.id$: the associated and unique identifier of the party P . In this paper, we assume it equal to id_P .
- $st_P.isk$: the identity secret key of the party P .
- $st_P.ipk$: the identity public key of the party P .

The short-term per-group state variables are initialized during the Register phase:

- $\pi_P^{gid}.sk$: the group-specific secret key of the party P for joining the group gid .
- $\pi_P^{gid}.pk$: the group-specific public key of the party P for joining the group gid .
- $\pi_P^{gid}.gk$: the current group key used by party P , which is supposed to be shared by all parties in the group gid . This variable is initialized with \perp .
- $\pi_P^{gid}.gkid$: the index of the current group key of the party P in the group gid . This variable is initialized with 0.
- $\pi_P^{gid}.GP$: The set of all parties in the group gid . This variable is initialized with the empty set \emptyset .
- $\pi_P^{gid}.status \in \{\perp, registered, joined\}$: the status that indicates whether the party P has initialized the state for (but not yet registered for), or registered for (but not yet joined), or joined the group gid . This variable is \perp by default.

Due to the page limit, we omit the correctness analysis and refer the reader to the full version of the paper.

4.2. A Generic Security Model

We next define a generic Sec-mGKD-X security model. By presenting different instantiations of the freshness conditions that the attacker must obey and of the winning conditions that the attacker must pursue, we then introduce three distinct concrete models for $X \in \{pki, pw, pw+\}$ in Section 4.3, Section 4.4, and Section 4.5.

Trust Model. We assume that all parties' sampled randomness is independent, uniform, and unpredictable. For simplicity, we assume every leader samples group secret from a same distribution \mathcal{D} (according to the underlying protocol). We assume that every party joins every group at most once. We assume that every leader stays in the corresponding group for the entire duration and leaves only when all other participants have left. This means, once the leader leaves the corresponding group, this group is immediately marked as “invalid” and no party is allowed to register for or join this group anymore. We assume that every party can send register request for every group at most once. Our model assumes a single shared group key for communication. Thus,

impersonation attacks between parties inside the group is out of scope of this work.

Threat Model. We allow the attacker to have full control over the network and can eavesdrop, drop, and insert messages during all phases. We allow the attacker to corrupt the long-term state st_P for any participant P to capture the real-world scenarios where the hardware devices might be stolen. Moreover, we allow attackers to compromise the short-term per-group state π_P^{gid} for any participant P after joining any group gid to capture attacks during the ongoing communications. We also allow attackers to leak arbitrary group keys (to analyze the impact on the security of previous and future group keys). We allow attackers to reveal the group secret for any group to capture the real-world scenarios where the the out-of-band channels might be vulnerable.

Security Experiment. The security experiment is conducted between a challenger \mathcal{C} and an attacker \mathcal{A} against a mGKD protocol Π . At the beginning of the experiment, the challenger \mathcal{C} samples a random challenge bit $b \in \{0, 1\}$. During the experiment, \mathcal{C} produces two sequences of variables $\{GP^{\text{gid}, \text{gkid}}\}_{\text{gid}, \text{gkid}}$ and $\{gk_P^{\text{gid}, \text{gkid}}\}_{\text{gid}, \text{gkid}, P}$. The variable $GP^{\text{gid}, \text{gkid}}$ aims to record the identifier of every party that is expected to know the gkid -th group key in the group gid from the leader's view. The challenger \mathcal{C} monitors the states of the leader P of any group gid . Whenever $\pi_P^{\text{gid}} \cdot GP$ changes, \mathcal{C} records $GP^{\text{gid}, \pi_P^{\text{gid}} \cdot \text{gkid}} \leftarrow GP^{\text{gid}, \pi_P^{\text{gid}} \cdot \text{gkid}} \cup \pi_P^{\text{gid}} \cdot GP$. The variable $gk_P^{\text{gid}, \text{gkid}}$ records the gkid -th group key derived by any party P in the group gid . Whenever $\pi_P^{\text{gid}} \cdot \text{gkid}$ changes for any party P and group gid during the experiment, \mathcal{C} stores the new group key $gk_P^{\text{gid}, \pi_P^{\text{gid}} \cdot \text{gkid}} \leftarrow \pi_P^{\text{gid}} \cdot gk$. The attacker \mathcal{A} can interact with \mathcal{C} by querying the following oracles, where \mathcal{C} responds according to Π . To simplify the explanation, we partition the oracles into categories.

Oracle Category 1: Setup of groups and parties. This category includes a NEWPARTY oracle that simulates the Sign Up phase of a party, a NEWGROUP that simulates the Group Schedule phase of a group, and a AUTHORIZE oracle that simulates the group secret transmission from the leader to the authorized participants over out-of-band channels.

- NEWPARTY(id_P): This oracle can be queried at most once on each input. The challenger \mathcal{C} initializes a state st_P by setting $st_P.\text{id} \leftarrow \text{id}_P$. Then, \mathcal{C} runs $m_{\text{SignUp}}^P \xleftarrow{\$} \text{SignUp}(st_P)$ and followed by forwarding the sign-up message m_{SignUp}^P to \mathcal{A} . The party P is marked as “created”.
- NEWGROUP(id_P, gid): This oracle can be invoked at most once for each gid . The input party P must be marked as “created”. The challenger \mathcal{C} samples the secret of the group gid by $gs^{\text{gid}} \xleftarrow{\$} \mathcal{D}$ and runs $m_{\text{GSch}}^{\text{gid}} \xleftarrow{\$} \text{Schedule}(P, \text{gid}, gs^{\text{gid}})$ for an associated outgoing message $m_{\text{GSch}}^{\text{gid}}$. Then, \mathcal{C} marks the group gid as “created” and “valid” and marks P as the leader of the group gid and “authorized”. Finally, \mathcal{C} returns $m_{\text{GSch}}^{\text{gid}}$ to \mathcal{A} .
- AUTHORIZE(gid, id_P): The group gid must be marked as both created and valid, and the party P must be marked as created and have not registered for the group gid , i.e.,

$\pi_P^{\text{gid}}.\text{status} = \perp$. The challenger \mathcal{C} marks P as authorized for the group gid .

Oracle Category 2: Register phase. This category includes a REGISTERAUTH oracle, which simulates that an authorized party registers for a group using honest group secret, and a REGISTERINJECT oracle, which simulates that an unauthorized (malicious) party registers for a group with an input using some chosen group secret.

- REGISTERAUTH($\text{id}_P, \text{gid}, m$): This oracle can be queried at most once for each tuple $(\text{id}_P, \text{gid})$. The group gid must be marked as both created and valid and the party P must be authorized for the group gid . The challenger \mathcal{C} runs Register-L($P, \text{gid}, gs^{\text{gid}}, m$) if P is the leader of the group gid and Register-P($P, \text{gid}, gs^{\text{gid}}, m$) otherwise. In both cases, \mathcal{C} forwards the output message m' to \mathcal{A} .
- REGISTERINJECT($\text{id}_P, \text{gid}, gs, m$): This oracle can be queried at most once for each tuple $(\text{id}_P, \text{gid})$. The group gid must be marked as both created and valid and the party P must be unauthorized for the group gid . The challenger \mathcal{C} runs Register-P(P, gid, gs, m) and forwards the output message m' to \mathcal{A} .

Oracle Category 3: Participant Join phase. This category includes a SENDJOINAUTH oracle, which simulates that an authorized party (as either leader or participant) sends messages to another party (either authorized or unauthorized) during the Participant Join phase in the group, and a SENDJOININJECT oracle, which simulates that an unauthorized (malicious) participant sends messages (to the leader) during the Participants Join phase in the group.

- SENDJOINAUTH($\text{id}_P, \text{id}_{P'}, \text{gid}, m$): The challenger \mathcal{C} first checks
 - whether gid is marked as created and valid,
 - whether P is authorized for this group gid ,
 - whether both parties P and P' have been created and registered for this group,
 - whether either P or P' is the leader of the group gid ,
 - whether the leader of the group, either P or P' , has joined the group, and that the other party hasn't joined the group yet.

If any of the check fails, \mathcal{C} directly returns \perp to \mathcal{A} . Otherwise, \mathcal{C} runs Join-L($P, \text{id}_{P'}, \text{gid}, gs^{\text{gid}}, m$) if P is the leader of the group gid or Join-P($P, \text{id}_{P'}, \text{gid}, gs^{\text{gid}}, m$) if P is a participant. Then, the output message m' of either Join-L or Join-P is returned to \mathcal{A} .

- SENDJOININJECT($\text{id}_P, \text{id}_{P'}, \text{gid}, gs, m$): The challenger \mathcal{C} first checks
 - whether gid is marked as created and valid,
 - whether P is unauthorized for the group gid ,
 - whether P' is the leader of the group gid ,
 - whether the participant P has been created and registered for this group gid but has not joined the group yet, and
 - whether the leader P' has joined the group gid .

If any of the check fails, \mathcal{C} directly returns \perp to \mathcal{A} . Otherwise, \mathcal{C} runs Join-P($P, \text{id}_{P'}, \text{gid}, gs, m$) and returns the output message m' of Join-P to \mathcal{A} .

Oracle Category 4: Member Leave phase. This category includes a SENDLEAVE oracle, which simulates that a party notices another party leaving the group gid , and an ENDDGROUP oracle, which simulates that a leader leaves and ends the group.

- SENDLEAVE($id_P, gid, id_{P'}$): The challenger \mathcal{C} aborts if
 - the gid is not marked as both created and valid, or
 - the leaving party P' is the leader of the group gid , or
 - the party P has not joined the group.
 Otherwise, \mathcal{C} runs Leave-L($P, gid, id_{P'}$) if P is the leader of the group gid and Leave-P($P, gid, id_{P'}$) otherwise.
- ENDDGROUP(gid): The challenger \mathcal{C} first checks
 - whether the group gid is created and valid, and
 - whether the leader P of the group gid is the unique party in his local party list, i.e., $\pi_P^{gid}.GP = \{id_P\}$.
 If either check fails, \mathcal{C} aborts. Otherwise, \mathcal{C} runs Leave-L(P, gid, id_P) and marks the group gid as “invalid”.

Oracle Category 5: Key Rotation phase. This category includes only one SENDKEYROTAT oracle that simulates the process where a party updates their local group key.

- SENDKEYROTAT(id_P, gid, m): The party P must have joined the group gid . The challenger \mathcal{C} runs KeyRotat-L(P, gid, m) if P is the leader of the group gid and KeyRotat-P(P, gid, m) otherwise. The output message of either KeyRotat-L or KeyRotat-P is returned to \mathcal{A} .

Oracle Category 6: Secret information leakage. This category includes four oracles CORRUPT, COMPROMISE, LEAK, and REVEAL, that respectively simulates that the attacker \mathcal{A} knows the long-term state of a party, the short-term per-group state of a party for a group, a group key of a party in a group, and the group secret of a group.

- CORRUPT(id_P): The challenger \mathcal{C} first checks whether the party P is created. If the check fails, \mathcal{C} simply returns \perp . Otherwise, \mathcal{C} returns st_P to the attacker \mathcal{A} and marks st_P as “corrupted”.
- COMPROMISE(id_P, gid): The challenger \mathcal{C} checks whether the party P has joined the group gid , i.e., $\pi_P^{gid}.status = joined$. If the check fails, \mathcal{C} simply returns \perp . Otherwise, \mathcal{C} returns π_P^{gid} to the attacker \mathcal{A} , followed by marking π_P^{gid} as “compromised” and $gk_P^{(gid, \pi_P^{gid}.gkid)}$ as “leaked”.
- LEAK($id_P, gid, gkid$): The challenger \mathcal{C} checks whether $gk_P^{(gid, gkid)}$ has been set. If $gk_P^{(gid, gkid)} = \perp$, then \mathcal{C} simply returns \perp . Otherwise, \mathcal{C} marks $gk_P^{(gid, gkid)}$ as “leaked” and returns $gk_P^{(gid, gkid)}$ to \mathcal{A} .
- REVEAL(gid): If the group gid is not created, then the challenger \mathcal{C} simply returns \perp . Otherwise, \mathcal{C} marks gid as “revealed” and returns gs^{gid} to \mathcal{A} .

Oracle Category 7: Test challenge bit. This category includes only one TEST oracle that returns either a real group key if the challenge bit $b = 0$, or a random key if $b = 1$.

- TEST($id_P, gid, gkid$): This oracle can be queried at most once. If the party P is authorized for the group gid and the party P has produced $gkid$ -th group key, i.e., $gk_P^{(gid, gkid)} \neq$

\perp , then the challenger \mathcal{C} returns $gk_P^{(gid, gkid)}$ to \mathcal{A} if the challenge bit $b = 0$, or a random key from the same space if $b = 1$. Then, \mathcal{C} further marks the party P , the group identifier gid , and the group key index $gkid$ as “tested”. Otherwise, \mathcal{C} immediately returns \perp .

Advantage Measures. In the end, the attacker \mathcal{A} outputs a bit $b' \in \{0, 1\}$. Under two freshness conditions $\text{fresh}_{KAuth}^{\text{Sec-mGKD-X}}$ and $\text{fresh}_{KPriv}^{\text{Sec-mGKD-X}}$ that prevent the attacker \mathcal{A} from trivially winning the experiment, we say the attacker \mathcal{A} wins the experiment Sec-mGKD-X against a mGKD protocol Π , if either of the following events is triggered:

- 1) [Event E_{KAuth}] there exists any group gid with the leader P^{gid} , any party P' that is authorized for the group gid , and any group key index $gkid$, such that $gk_{P'}^{(gid, gkid)} \neq \perp$ but $gk_{P^{gid}}^{(gid, gkid)} \neq gk_{P'}^{(gid, gkid)}$, without violating the freshness condition $\text{fresh}_{KAuth}^{\text{Sec-mGKD-X}}(id_{P'}, gid, gkid)$.
- 2) [Event E_{KPriv}] $b = b'$ without violating the freshness condition $\text{fresh}_{KPriv}^{\text{Sec-mGKD-X}}(id_{P'}, gid, gkid)$, where P' , gid , and $gkid$ are respectively the tested party, group identifier, and group key index.

We define $\text{Adv}_{\Pi}^{\text{Sec-mGKD-X}}(\mathcal{A})$ as the advantage that \mathcal{A} can win the Sec-mGKD-X experiment against a mGKD protocol Π , namely,

$$\text{Adv}_{\Pi}^{\text{Sec-mGKD-X}}(\mathcal{A}) := \max \left(\left| \text{Pr}[E_{KPriv}] - \frac{1}{2} \right|, \text{Pr}[E_{KAuth}] \right).$$

Definition 3 (Sec-mGKD-X). *We say that a mGKD protocol Π is Sec-mGKD-X-secure for $X \in \{pki, pw, pw+\}$, if the advantage $\text{Adv}_{\Pi}^{\text{Sec-mGKD-X}}(\mathcal{A})$ is negligible for any PPT adversary \mathcal{A} .*

4.3. The Sec-mGKD-pki Security Model

Our basic security model Sec-mGKD-pki captures the following security guarantees for an authorized party in a group assuming the honest distribution of long-term sign-up message of all parties within this group. Note that this basic model guarantees that only group members can learn the key, and where group membership is determined by the *server*. In reality, and in our model, the server may insert group members that are not authorized by the leader and do not know the passcode.

- 1) **(Implicit) Group Key Authentication:** If a group member accepts a group key, then the leader must have produced the same group with the same group key index.
- 2) **Group Key Secrecy:** If a group member accepts a group key, then an attacker cannot derive this key, even if it knows other group keys.
- 3) **Perfect Forward Secrecy:** An attacker that compromises a party’s long-term keys, can not learn the group keys of any group the party was previously in.

Definition 4 (Sec-mGKD-pki Freshness Conditions). *We say the freshness condition $\text{fresh}_{KAuth}^{\text{Sec-mGKD-pki}}(id_{P'}, gid, gkid)$, where gid has a unique leader P^{gid} , holds if and only if*

- 1) *the per-group states $\pi_{P^{gid}}^{gid}$ and $\pi_{P'}^{gid}$ are not compromised,*

- 2) the long-term states $st_{P^{gid}}$ and $st_{P'}$ are not corrupted before P^{gid} and P' joined the group gid , and
- 3) the sign-up messages $m_{SignUp}^{P^{gid}}$ and $m_{SignUp}^{P'}$ of P^{gid} and P' are honestly delivered to the other.

We say the freshness condition $\text{fresh}_{KPriv}^{\text{Sec-mGKD-pki}}(id_{P'}, gid, gkid)$ holds if and only if

- 1) the group key $gk_P^{(gid, gkid)}$ is not leaked for all parties P in the group gid with $id_P \in GP^{(gid, gkid)}$,
- 2) the short-term state π_P^{gid} is not compromised for all parties P in the group gid with $id_P \in GP^{(gid, gkid)}$,
- 3) the long-term state $st_{P^{gid}}$ of the leader P^{gid} in the group gid is not corrupted before all parties P with $id_P \in GP^{(gid, gkid)}$ joined the group gid ,
- 4) the long-term state st_P of all participants P in the group gid with $id_P \in GP^{(gid, gkid)}$ is not corrupted before P joined the group gid , and
- 5) the sign-up messages m_{SignUp}^P of all parties P with $id_P \in GP^{(gid, gkid)}$ are honestly distributed within the group gid .

Conclusion. Our Sec-mGKD-pki security model captures all guarantees listed at the beginning of this subsection:

- 1) **(Implicit) Group Key Authentication:** If authentication does not hold, the attacker \mathcal{A} can win via E_{KAuth} .
- 2) **Group Key Secrecy:** The attacker is allowed to leak arbitrarily many group keys except for the tested one. If group key secrecy does not hold, \mathcal{A} can win via E_{KPriv} .
- 3) **Perfect Forward Secrecy:** The attacker is allowed to corrupt the long-term state of the tested party. If perfect forward secrecy does not hold, \mathcal{A} can win via E_{KPriv} .

4.4. The Sec-mGKD-pw Security Model

The basic Sec-mGKD-pki model has two restrictions:

- First, both freshness conditions $\text{fresh}_{KPriv}^{\text{Sec-mGKD-pki}}$ and $\text{fresh}_{KAuth}^{\text{Sec-mGKD-pki}}$ require the honest distribution of the sign-up messages. Since the sign-up messages are distributed by servers or by PKI in practice, this restriction is also known as “trusted PKI” assumption in the related literature. In the full end-to-end setting, i.e., no trusted PKI or server exists, a Sec-mGKD-pki secure mGKD protocol might still suffer from machine-in-the-middle attacks such that the attacker can easily impersonate any participant towards the group leader and impersonate any group leader towards any participant.
- Second, both freshness conditions $\text{fresh}_{KPriv}^{\text{Sec-mGKD-pki}}$ and $\text{fresh}_{KAuth}^{\text{Sec-mGKD-pki}}$ allow the attacker to reveal all group secrets but do not prevent unauthorized parties from knowing the group keys. Thus, this Sec-mGKD-pki model does not capture the security benefit provided by the group secret transmitted over secure out-of-band channels.

The goal of the Sec-mGKD-pw security model is to preserve the guarantees achieved in Sec-mGKD-pki model and to capture the following additional security guarantee, while getting rid of the “trusted PKI” assumption.

- 4) **(Implicit) Group Member Authentication:** If any party produces the same group key as the leader of a group, then

this party must be authorized for this group, as long as the group secret is not revealed.

We replace the “trusted PKI” assumption with the arguably simpler assumption of a shared “secure (short) group secret”. We assume that (short) group secrets (such as passwords or pin codes) can be distributed over out-of-band secure channels, e.g., by email, encrypted messaging application (e.g., Signal), or even in person. In fact, a similar passcode mechanism has been widely deployed in real life by many service providers such as Zoom for access control management (see Section 5.1). The only difference is that the current Zoom passcode mechanism hands over passcode and the rule of verifying it to the (possibly) untrusted server, while in our model the group secret is known only to the participants.

Definition 5 (Sec-mGKD-pw Freshness Conditions). We say the freshness condition $\text{fresh}_{KAuth}^{\text{Sec-mGKD-pw}}(id_{P'}, gid, gkid)$, where gid has a unique leader P^{gid} , holds if and only if

- 1) the per-group states $\pi_{P^{gid}}^{gid}$ and $\pi_{P'}^{gid}$ are not compromised,
- 2) the long-term states $st_{P^{gid}}$ and $st_{P'}$ are not corrupted before P^{gid} and P' joined the group gid , and
- 3) the group gid is not revealed.

We say the freshness condition $\text{fresh}_{KPriv}^{\text{Sec-mGKD-pw}}(id_{P'}, gid, gkid)$ holds if and only if

- 1) the group key $gk_P^{(gid, gkid)}$ is not leaked for all authorized parties P in the group gid with $id_P \in GP^{(gid, gkid)}$,
- 2) the short-term state π_P^{gid} is not compromised for all authorized parties P in the group gid with $id_P \in GP^{(gid, gkid)}$,
- 3) the long-term state $st_{P^{gid}}$ of the leader P^{gid} in the group gid is not corrupted before all authorized parties P with $id_P \in GP^{(gid, gkid)}$ joined the group gid ,
- 4) the long-term state st_P of all authorized participants P in the group gid with $id_P \in GP^{(gid, gkid)}$ is not corrupted before P joined the group gid , and
- 5) the group gid is not revealed.

Conclusion. Note that Sec-mGKD-pki and Sec-mGKD-pw models share the same oracles and similar freshness conditions. Thus, it is straightforward that our Sec-mGKD-pw model also satisfies all guarantees listed in Section 4.3. However, we stress two main distinctions of the freshness conditions in Sec-mGKD-pw and Sec-mGKD-pki models that make the security guarantees provided by these two models very different.

- 1) **[Full End-to-End Security]** Both $\text{fresh}_{KAuth}^{\text{Sec-mGKD-pw}}$ and $\text{fresh}_{KPriv}^{\text{Sec-mGKD-pw}}$ allow attackers to manipulate the transmission of all messages, including the sign-up messages of all parties in any group, which are forbidden by $\text{fresh}_{KAuth}^{\text{Sec-mGKD-pki}}$ sub-point 3) and $\text{fresh}_{KPriv}^{\text{Sec-mGKD-pki}}$ sub-point 5). Instead, $\text{fresh}_{KAuth}^{\text{Sec-mGKD-pw}}$ sub-point 3) and $\text{fresh}_{KPriv}^{\text{Sec-mGKD-pw}}$ sub-point 5) require no leakage of the group secrets. This indicates that our Sec-mGKD-pw security model captures the full end-to-end security, i.e., against a malicious server. Consequently, this new Sec-mGKD-pw model solves the first restriction of Sec-mGKD-pki model.

2) While the $\text{fresh}_{\text{KPriv}}^{\text{Sec-mGKD-pki}}(\text{id}_{P'}, \text{gid}, \text{gkid})$ condition sub-points 1) - 4) require no group key leakage, no short-term per-group state compromise, and no long-term state corruption for all parties P in the group with $\text{id}_P \in GP^{(\text{gid}, \text{gkid})}$, our new $\text{fresh}_{\text{KPriv}}^{\text{Sec-mGKD-pw}}(\text{id}_{P'}, \text{gid}, \text{gkid})$ condition has the same requirements in sub-points 1) - 4) but only for the authorized parties in every group. By this, our Sec-mGKD-pw model captures the following property:

- **(Implicit) Group Member Authentication:** The attacker can leak arbitrarily many group keys of any unauthorized party in the tested group. If an unauthorized party can successfully produce a same group key as a leader, then the attacker can test this leader, leak the group key of this unauthorized party, and win via the event E_{KPriv} .

4.5. The Sec-mGKD-pw+ Security Model

Note that the Sec-mGKD-pki and Sec-mGKD-pw models rely on different assumptions: trusted PKI and secure group secret. We then define a third Sec-mGKD-pw+ model that incorporates the above two models. The goal of the Sec-mGKD-pw+ model is to capture the security of a mGKD protocol if either the “trusted PKI” or the “secure group secret” assumption holds.

Definition 6 (Sec-mGKD-pw+ Freshness Conditions). *We say the freshness condition $\text{fresh}_{\text{KAuth}}^{\text{Sec-mGKD-pw+}}(\text{id}_{P'}, \text{gid}, \text{gkid})$ holds if and only if $\text{fresh}_{\text{KAuth}}^{\text{Sec-mGKD-pki}}(\text{id}_{P'}, \text{gid}, \text{gkid})$ or $\text{fresh}_{\text{KAuth}}^{\text{Sec-mGKD-pw}}(\text{id}_{P'}, \text{gid}, \text{gkid})$ holds.*

We say the freshness condition $\text{fresh}_{\text{KPriv}}^{\text{Sec-mGKD-pw+}}(\text{id}_{P'}, \text{gid}, \text{gkid})$ holds if and only if $\text{fresh}_{\text{KPriv}}^{\text{Sec-mGKD-pki}}(\text{id}_{P'}, \text{gid}, \text{gkid})$ or $\text{fresh}_{\text{KPriv}}^{\text{Sec-mGKD-pw}}(\text{id}_{P'}, \text{gid}, \text{gkid})$ holds.

The following corollary is straightforward by definition:

Corollary 1. *Let Π denote a mGKD protocol. If Π is Sec-mGKD-pki and Sec-mGKD-pw secure, then Π is also Sec-mGKD-pw+ secure, and vice versa.*

5. Zoom’s protocol is a mGKD protocol

The Zoom library allows parties to establish end-to-end group meeting communication. In this section, we first introduce the Zoom overview [35] (version 4.0) in Section 5.1. Then, we show that Zoom library is Sec-mGKD-pki secure in Section 5.2 but not Sec-mGKD-pw secure in Section 5.3.

5.1. The Zoom End-to-End Connection Overview

For end-to-end encrypted meetings, Zoom only supports connecting from installed clients: Browser-based connections are not supported. Zoom distinguishes between users and devices by non-cryptographic user identifiers uid and hardware identifiers hid . We model the identifier of each party P by a pair of user and hardware identifiers, i.e., $\text{id}_P = (\text{uid}_P, \text{hid}_P)$ and assume that party identifiers are unique¹.

¹The Zoom white-paper [35] states that the user identifiers uid are assigned by servers and the hardware identifiers hid are randomly sampled. Based on this, we assume that they are unique in practice.

Zoom deploys two infrastructures for transmitting cryptographic primitives: an identity management system and a multimedia router (MMR). While the identity management system distributes cryptographic public keys generated by individual clients, the MMR distributes cryptographic messages between parties in a meeting. The connection between parties and servers are established on TLS-tunnels over TCP and are encrypted with AES in GCM mode. In this paper, we assume the existence of Zoom servers but do not explicitly model them, because our goal is to consider them adversary-controlled. Zoom allows every party to set up a group meeting. Groups are uniquely identified by their group identifiers gid . Each group meeting is equipped with a specific “bulletin board”, where all parties can post (their own) and retrieve (others’) cryptographic messages. The server is able to control and tamper with the bulletin boards.

The end-to-end secure Zoom library consists of six phases following the mGKD protocol syntax², of which we give an overview in Figure 1. We recall the cryptographic algorithms ZSign and ZBox underlying the Zoom library in Appendix B. For domain separation, Zoom uses hardcoded context strings ($\text{ctx}_1, \text{ctx}_2, \text{ctx}_3$).

Sign Up $\text{SignUp}(P)$: During the Sign Up phase, every party P samples an identity public-private ZSign key pair and stores them into the long-term state st_P . The party P outputs the identity public key to the server as the sign-up message. This algorithm is executed only once for each party, i.e., each user on each hardware device.³

Group Schedule $\text{Schedule}(P, \text{gid}, \text{gs})$: During the Group Schedule phase, the leader P parses a passcode pc^{gid} from the input gs . The leader P sends pc^{gid} to not only the server as the group schedule message $m_{\text{GSch}}^{\text{gid}}$ for the access control management, but also to the authorized participants for joining the group over out-of-band channels, e.g., email.

Register $\text{Register} = (\text{Register-L}, \text{Register-P})$: The Register phase enables every party P to register for joining the meeting gid . We separate the description for Register-P($P', \text{gid}, \text{gs}, m$), where the P' is a participant of the group gid , and for Register-L($P, \text{gid}, \text{gs}, m$), where P the leader of the group gid .

- Register-P($P', \text{gid}, \text{gs}, m$): The input message m is given by the server and should be correctly parsed as a special mUUID string. The mUUID string is a server-generated per-group-instance random string that the individual parties cannot control. Moreover, the participant P' also inputs a group secret gs that can be correctly parsed as a passcode pc^{gid} . This algorithm first samples a public-private per-group ZBox key pair and stores them into

²The official Zoom white-paper [35] only sketches the re-joining mechanism informally, and does not specify any mechanism to change leaders. We leave the analysis both functionalities to future work.

³The Zoom library also supports anonymous log-in: people without a Zoom account can also join a group meeting as a “guest participant” (note that the guest cannot play the role of leader). Before a guest joins a group, the Sign Up algorithm is always executed. This prevents other parties from tracing them across meetings by noticing when a long-term key is reused [35].

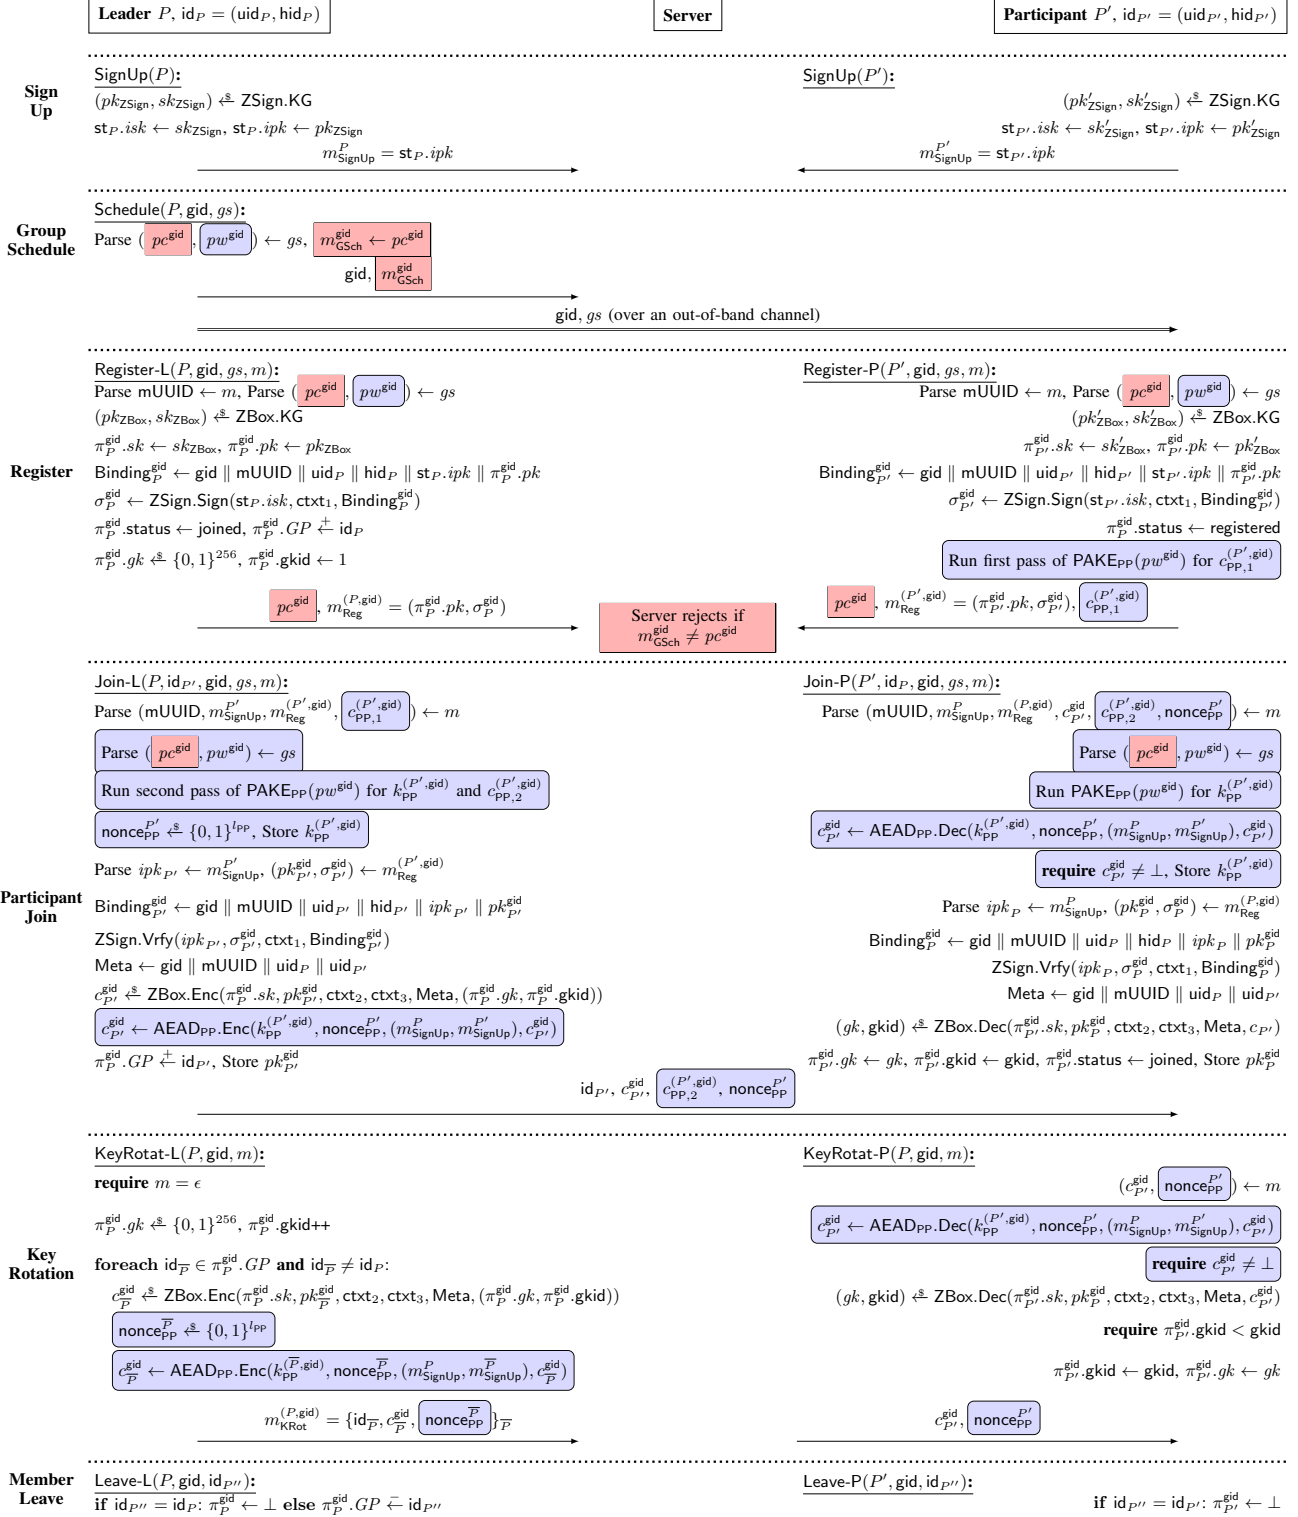


Figure 1. Overview of the **Zoom** protocol and our modified **ZoomPAKE** protocol. The boxes of the form added denote the additions from our transformation for ZoomPAKE, i.e., which are not in current Zoom. Note we do require any additional message flows. The boxes of the form redundant denote the elements that become redundant in our PAKE-based design: thus, for ZoomPAKE, we can essentially set pc^{gid} to the empty string, and still obtain the same guarantees, effectively replacing the old passcode by the new PAKE password. We recall the cryptographic algorithms ZSign and ZBox from the Zoom library in Appendix B. All communications from and to the server are performed over unilaterally secured TLS connections, authenticating the server. Possible out-of-band channels in the Group Schedule include, e.g., e-mail, messaging app, phone, calendar invite, or in-person.

the state $\pi_{P'}^{\text{gid}}$. Next, it computes $\text{Binding}_{P'}^{\text{gid}}$, which is the concatenation of the group identifier gid , server-generated random string mUUID , as well as the party P' 's identifier $\text{id}_{P'} = (\text{uid}, \text{hid})$, identity public key $\text{st}_{P'}.ipk$, and per-group public key $\pi_{P'}^{\text{gid}}.pk$. Then, it signs the binding information $\text{Binding}_{P'}^{\text{gid}}$ for a signature $\sigma_{P'}^{\text{gid}}$ using ZSign.Sign algorithm, the identity secret key $\text{st}_{P'}.isk$, and a context string ctx_1 . The passcode pc^{gid} and the output register message $m_{\text{Reg}}^{(P', \text{gid})}$ consisting of the per-group public key $\pi_{P'}^{\text{gid}}.pk$ and the signature $\sigma_{P'}^{\text{gid}}$ is sent to the server. The server adds $m_{\text{Reg}}^{(P', \text{gid})}$ to the “bulletin board” of the group gid , if the passcode pc^{gid} matches the group schedule message $m_{\text{GSch}}^{\text{gid}}$ received from the group leader, and rejects it otherwise. The status $\pi_{P'}^{\text{gid}}.\text{status}$ of the participant P' in the group gid is set to registered.

- **Register-L(P, gid, gs, m):** If the party P is the leader of the group gid , P runs the same execution as a participant except for setting the status $\pi_P^{\text{gid}}.\text{status}$ to joined rather than registered. Moreover, the leader P initializes the group by sampling the first group key $\pi_P^{\text{gid}}.gk$ of bit length 256 and sets the group key index $\pi_P^{\text{gid}}.gkid$ to 1. The identifier id_P is added into the local party set $\pi_P^{\text{gid}}.GP$.

Participants Join $\text{Join} = (\text{Join-L}, \text{Join-P})$: The Zoom library executes this interactive sub-protocol between the leader P and a participant P' only one-pass:

- **Join-L($P, \text{id}_{P'}, \text{gid}, gs, m$):** When the leader P notices the joining request of a new participant P' , P retrieves an incoming message m from the server and the group gid 's “bulletin board” followed by parsing it into: (1) a server-generated randomness mUUID , (2) the participant P' 's sign-up message $m_{\text{SignUp}}^{P'}$, and (3) the participant P' 's register message $m_{\text{Reg}}^{(P', \text{gid})}$. Next, the leader P parses the participant P' 's identity public key $ipk_{P'}$, per-group public key $pk_{P'}^{\text{gid}}$, and per-group signature $\sigma_{P'}^{\text{gid}}$ from the incoming message, followed by using them to produce the participant's binding information $\text{Binding}_{P'}^{\text{gid}}$. If the binding information cannot pass the verification ZSign.Vrfy upon the participant's identity public key $ipk_{P'}$, signature $\sigma_{P'}^{\text{gid}}$, and the context ctx_1 , then the leader aborts and undoes the previous executions. Otherwise, the leader creates a metadata Meta by concatenating the group identifier gid , server-generated randomness mUUID , the leader's user identifier id_P , and the participant's user identifier $\text{id}_{P'}$. Finally, the leader P encrypts the current group key $\pi_P^{\text{gid}}.gk$ as well as the index $\pi_P^{\text{gid}}.gkid$ using the ZBox.Enc encryption algorithm and the leader P 's per-group secret key $\pi_P^{\text{gid}}.sk$, the participant P' 's per-group public key $\pi_{P'}^{\text{gid}}.pk$, and auxiliary information $\text{ctx}_2, \text{ctx}_3$, and Meta . The identifier of the participant P' and the ZBox ciphertext $c_{P'}^{\text{gid}}$ are sent to P' via the server. The leader P stores the identifier $\text{id}_{P'}$ of the participant P' into the local party set $\pi_P^{\text{gid}}.GP$ and stores the participant's per-group public key $pk_{P'}^{\text{gid}}$.
- **Join-P($P', \text{id}_P, \text{gid}, gs, m$):** When a participant P' , who registered for a group gid , receives an incoming message

m that includes (1) a server-generated randomness mUUID , (2) the leader P 's sign-up message m_{SignUp}^P , (3) the leader P 's register message $m_{\text{Reg}}^{(P, \text{gid})}$, and (4) the leader's reply $c_{P'}^{\text{gid}}$, P' first parses the incoming message, creates the leader's binding information, and verifies it using ZSign.Vrfy algorithm, similar to the leader's execution. Then, P' also generates the meta data Meta and uses it together with its own per-group secret key $\pi_{P'}^{\text{gid}}.sk$, the leader's per-group public key pk_P^{gid} , and the contexts ctx_2 and ctx_3 , to decrypt the ciphertext $c_{P'}^{\text{gid}}$. If any error occurs during above steps, then the participant P' aborts and undoes the previous executions. Otherwise, the participant P' stores the decrypted group key as well as the associated index into the per-group state $\pi_{P'}^{\text{gid}}.gk$ and $\pi_{P'}^{\text{gid}}.gkid$, followed by setting the status $\pi_{P'}^{\text{gid}}.\text{status}$ to joined. Moreover, P' stores the leader P 's per-group public key into the state.⁴

Key Rotation $\text{KeyRotat} = (\text{KeyRotat-L}, \text{KeyRotat-P})$: The execution of this algorithm is distinct according to the caller P 's role: a leader or a participant. We separate the description for $\text{KeyRotat-L}(P, \text{gid}, m)$, where P is the leader of the group gid , and for $\text{KeyRotat-P}(P', \text{gid}, m)$, where P' is a participant of the group gid .

- **KeyRotat-L(P, gid, m):** The leader P of the group gid executes this algorithm without any auxiliary incoming input, i.e., $m = \epsilon$. The leader P samples a new group key $\pi_P^{\text{gid}}.gk$ of bit length 256 and increments the corresponding index $\pi_P^{\text{gid}}.gkid$ by 1. Similar to the encryption during the Participant Join phase, the leader encrypts the new group key and index for each party in its local party set $\pi_P^{\text{gid}}.GP$ except for himself. The output is a ciphertext bundle that includes the identifiers of each participant and the customized ciphertexts.

The server is expected to split the ciphertext bundles and to send each ciphertext to the specified participant.

- **KeyRotat-P(P', gid, m):** The participant P' of the group gid first parses the incoming message m from the server to an ZBox ciphertext $ct_{P'}^{\text{gid}}$. Then, P' decrypts the new group key gk and index $gkid$ as during the Participant Join phase. If any error occurs in the above steps or the decrypted group key index is smaller than or equal to the local one, then the participant P' aborts and undoes the previous executions. Otherwise, P' simply overwrites the local group key as well as the index by the new ones.

Member Leave $\text{Leave} = (\text{Leave-L}, \text{Leave-P})$: The execution of this algorithm is distinct according to the caller P 's role: a leader or a participant. We separate the description for $\text{Leave-L}(P, \text{gid}, \text{id}_{P''})$, where the P is the leader of the group gid , and for $\text{Leave-P}(P', \text{gid}, \text{id}_{P''})$, where P' is a participant of the group gid .

- **Leave-L($P, \text{gid}, \text{id}_{P''}$):** If $\text{id}_P = \text{id}_{P''}$, i.e., the leader P wants to leave the group gid , then P erases the per-group

⁴In practice, Zoom has an independent mechanism for leader P to synchronize the party set $\pi_P^{\text{gid}}.GP$ with every participant P' . We omit it here since, this does not impact Zoom security analysis in our models.

- instance π_P^{gid} . Otherwise, the leader P notices a party P'' leaving the group gid , P simply removes the identifier of the participant $\text{id}_{P''}$ from the local party set $\pi_P^{\text{gid}}.GP$.
- **Leave-P(P' , gid , $\text{id}_{P''}$):** If $\text{id}_{P'} = \text{id}_{P''}$, i.e., the participant P' wants to leave the group gid , then P' erases the per-group instance $\pi_{P'}^{\text{gid}}$. Otherwise, P' performs no action.

Instantiations. Underlying the ZSign and ZBox algorithms, the key derivation function H_1 is SHA256 and H_2 is HKDF algorithm (using an empty salt parameter). The length l underlying ZBox algorithm is 192. The elliptic curve ECDH underlying Diffie-Hellman key exchange is Curve25519. The ZSign algorithm relies DS on EdDSA over Ed25519. The AEAD algorithm is xchacha20poly1305.

5.2. Zoom is Sec-mGKD-pki secure

Below, we investigate the provable security of the Zoom library. We give the proof sketch in Appendix C.

Theorem 1. *Let Π denote the end-to-end Zoom protocol in Section 5.1. Assume the $\epsilon_{H_1}^{\text{coll-res}}$ -collision resistance of the underlying H_1 , the $\epsilon_{\text{DS}}^{\text{euf-cma}}$ -euf-cma security of DS, the $\epsilon_{\text{AEAD}}^{(n,m)\text{-frob}}$ - (n, m) -FROB security, $\epsilon_{\text{AEAD}}^{\text{ind\$-cca}}$ -IND\\$-CCA security, and $\epsilon_{\text{AEAD}}^{\text{cti-cpa}}$ -cti-cpa security of the AEAD. Assume the $\epsilon_{\text{ECDH}, H_2}^{\text{mn-prf-ODH}}$ hardness of the mn-prf-ODH problem over ECDH and function H_2 . The advantage of any PPT attacker \mathcal{A} that breaks the Sec-mGKD-pki security of Π is bounded by,*

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{Sec-mGKD-pki}}(\mathcal{A}) &\leq \epsilon_{H_1}^{\text{coll-res}} + q_{\text{NEWPARTY}} \epsilon_{\text{DS}}^{\text{euf-cma}} \\ &+ c_{\text{maxReg}} q_{\text{NEWGROUP}} \left(\epsilon_{\text{AEAD}}^{\text{cti-cpa}} + \epsilon_{\text{AEAD}}^{(n,m)\text{-frob}} \right) \\ &+ c_{\text{maxReg}}^{(n_{\text{party}}-1)} (n_{\text{party}} - 1) \left(\epsilon_{\text{ECDH}, H_2}^{\text{mn-prf-ODH}} + \epsilon_{\text{AEAD}}^{\text{ind\$-cca}} \right) \end{aligned}$$

where c_{maxParty} denotes the maximal number of parties in every group, c_{maxReg} denotes the maximal number of register requests for every group, $n_{\text{party}} \leq c_{\text{maxParty}}$ denotes the number of parties in the set $GP^{(\text{gid}, \text{gkid})}$ for tested group identifier gid and group key index gkid , and $q_{\mathcal{O}}$ denote the maximal number of the queries to any oracle \mathcal{O} .

Note that $H_1 = \text{SHA256}$ provides an expected collision resistance of 128 bits [36]. The EUF-CMA security of Ed25519 was proven in [37]. The security of xchacha20poly1305 was discussed in [38]. The maximal number of parties per meeting is $c_{\text{maxParty}} = 1000$ [35]. The above theorem shows that the end-to-end Zoom library provably provides Sec-mGKD-pki security and satisfies all properties listed in Section 4.3.

Remark 1. *Theorem 1 shows that Zoom achieves Sec-mGKD-pki independent of the passcode: the passcode is only used implicitly for access control by honest servers.*

5.3. Zoom is not Sec-mGKD-pw secure

Recall that the Sec-mGKD-pki model has two restrictions in Section 4.3. A natural question arises whether these restrictions apply to the Sec-mGKD-pki secure Zoom library.

Does Zoom Provide Trusted PKI? The PKI is expected to “enable users of an insecure public network such as the Internet to securely and privately exchange data through the use of a public and a private cryptographic key pair that is obtained and shared through a trusted authority” [39, Chapter 1]. As we mentioned in Section 5.1, all public keys of all parties in the Zoom library are uploaded to an infrastructure, called “identity management system”, that is fully controlled by Zoom. The identity management system distributes the identity public keys. While Zoom claims the end-to-end security, the goal of which is to protect the secrecy and integrity of the exchanged content between every two parties against all third parties including the service providers, assuming Zoom-controlled PKI trusted is controversial and doubtful. Considering a malicious server, the server can easily perform the “machine-in-the-middle” attack by forging the sign-up messages and impersonate any party towards others.

Although there do exist other group meeting providers, such as Cisco WebEx and Skype, that employ a third-party PKI, such as Microsoft Certificate Authority (CA), we stress that the reliability of PKI is still imperfect. Eckersley and Burns [40] revealed that 14 CAs had been compromised. Moreover, a number of attacks that successfully break several CAs, including DigiNotar, Comodo, GlobalSign, StartSSL, and TurkTrust, have been publicly noticed [41]. It is prudent to consider the potential PKI compromise.

Does Zoom Provide (Implicit) Group Member Authentication? Unfortunately, this does not hold for Zoom. Although the end-to-end Zoom library asks every leader to create every new group together with an associated passcode, the leaders hand over the power of passcode verification to the untrusted server. By colluding with the untrusted server, an unauthorized (and malicious) party can join every group without the knowledge of any passcode. The consequence is that nobody in the group (including the leader) can distinguish authorized participants from the others, in particular, in the end-to-end setting.

6. A Generic Approach to Sec-mGKD-pw Security: Password-Protected Transformation

If we could assume that each group gid has a unique high-entropy group secret g_s^{gid} only shared by the leader and authorized participants of the group gid , we could design a trivial construction that meets the Sec-mGKD-pw security. We can simply use a message authentication code MAC with the group secret g_s^{gid} as key to sign and verify all outgoing and incoming messages.

However, in practice we use low-entropy passwords for usability, allowing the passwords to be shared over various out-of-band channels. For instance, real-world service providers often support only short passwords⁵. This restricts the upper bound of the password entropy and enables attackers to perform dictionary attacks on the password, e.g., by brute force guessing.

⁵For instance, meeting passcodes in Zoom are 1-16 digit numeric lock codes; the default meeting password in Cisco WebEx has ≥ 11 characters.

In this section, we introduce a generic Password-Protected (PP) transformation that provably transforms any Sec-mGKD-pki secure mGKD protocol Π to another Sec-mGKD-pw secure $\Pi' = \text{PP}[\Pi, \text{PAKE}_{\text{PP}}, \text{AEAD}_{\text{PP}}]$ protocol by using a password-authenticated key exchange PAKE_{PP} and an authenticated encryption with associated data AEAD_{PP} . We also prove that the PP transformation preserves Sec-mGKD-pki security, i.e., if Π is Sec-mGKD-pki secure, so is Π' . In this sense, Π' satisfies stronger security Sec-mGKD-pw+, due to Corollary 1. Finally, we illustrate how to apply our PP transformation to the Zoom library, and provide efficient instantiations for PAKE_{PP} and AEAD_{PP} , without causing additional message flows.

6.1. The Generic Transformation

The goal of our PP transformation is to ensure that only the authorized parties that know the group secret can recover any group key, even if the server is malicious. The high-level overview of our PP transformation is to (1) let the leader and every participant run a PAKE_{PP} protocol upon a new password (included in the group secret) to produce a symmetric key k_{PP} during the Participant Join phase, and (2) use the key k_{PP} and an AEAD_{PP} scheme to encrypt/decrypt the original transcript of the mGKD protocol Π during the Participant Join and Key Rotation phases. Note that every participant has to first register for a group before joining it. To avoid introducing additional message flows, we design our PP transformation to shift the first pass of PAKE_{PP} to the participant's Register phase. We give the formal definition of our PP transformation below.

Definition 7. Let $\Pi = (\text{SignUp}, \text{Schedule}, \text{Register}, \text{Join}, \text{Leave}, \text{KeyRotat})$ denote a multi-stage group key distribution protocol. Let PAKE_{PP} denote a password-authenticated key exchange scheme. Let AEAD_{PP} denote an authenticated encryption with associated data. We define the password-protected (PP) transformation $\text{PP}[\Pi, \text{PAKE}_{\text{PP}}, \text{AEAD}_{\text{PP}}]$ that outputs $\Pi' = (\text{SignUp}', \text{Schedule}', \text{Register}', \text{Join}', \text{Leave}', \text{KeyRotat}')$ as follows:

Sign Up $\text{SignUp}'(P)$: Run $m_{\text{SignUp}}^P \xleftarrow{\$} \text{SignUp}(P)$ and stores m_{SignUp}^P locally into the long-term state st_P .

Group Schedule $\text{Schedule}'(P, \text{gid}, gs)$: Parse $(gs_{\Pi}^{\text{gid}}, pw^{\text{gid}}) \leftarrow gs$, run $m_{\text{GSch}}^{\text{gid}} \xleftarrow{\$} \text{Schedule}(P, \text{gid}, gs_{\Pi}^{\text{gid}})$, and output $m_{\text{GSch}}^{\text{gid}}$. The full group secret gs is sent to authorized parties over out-of-band channels.

Register $\text{Register}' = (\text{Register-L}', \text{Register-P}')$: We define the sub-algorithms as follows:

- **Register-L'** (P, gid, gs, m) : Parse $(gs_{\Pi}^{\text{gid}}, pw^{\text{gid}}) \leftarrow gs$, run $m' \xleftarrow{\$} \text{Register-L}(P, \text{gid}, gs_{\Pi}^{\text{gid}}, m)$, and output m' .
- **Register-P'** (P, gid, gs, m) : First, parse $(gs_{\Pi}^{\text{gid}}, pw^{\text{gid}}) \leftarrow gs$. Next, run $m' \xleftarrow{\$} \text{Register-P}(P, \text{gid}, gs_{\Pi}^{\text{gid}}, m)$. Then, run the first pass of PAKE_{PP} upon the password pw^{gid} for a ciphertext $c_{\text{PP}}^{(P, \text{gid})}$. Finally, output $(m', c_{\text{PP}}^{(P, \text{gid})})$.

Participant Join $\text{Join}' = (\text{Join-L}', \text{Join-P}')$: This phase consists of two steps. In either step, if any error occurs during this algorithm, the caller P aborts and undoes the

executions. In the first step, the leader and the participant run PAKE_{PP} until PAKE_{PP} outputs a key k_{PP} .

- **Join-L'** $(P, \text{id}_{P'}, \text{gid}, gs, m)$ or **Join-P'** $(P, \text{id}_{P'}, \text{gid}, gs, m)$: The caller P first parses $(gs_{\Pi}^{\text{gid}}, pw^{\text{gid}}) \leftarrow gs$ from the group secret and other necessary information for running PAKE_{PP} from the input message m . Then, P runs the next pass of PAKE_{PP} on pw^{gid} . If the key k_{PP} is still unavailable, P directly outputs the outgoing message of PAKE_{PP} . Otherwise, the key k_{PP} is stored into the per-group state π_P^{gid} .

If the leader and the participant have computed the key k_{PP} before this algorithm invocation or in the first step in this invocation, they execute the following second step.

- **Join-L'** $(P, \text{id}_{P'}, \text{gid}, gs, m)$: The leader P first parses $(gs_{\Pi}^{\text{gid}}, pw^{\text{gid}}) \leftarrow gs$ from the group secret. If the original Join-L algorithm needs any incoming information from the participant P' , then the leader P extracts an AEAD_{PP} ciphertext and an AEAD_{PP} nonce from the input m . Next, the leader P decrypts the AEAD_{PP} ciphertext using the key k_{PP} , the AEAD_{PP} nonce, and an associated data consisting of both parties' sign-up messages, and obtains a message m_1 . Then, the leader P extracts other necessary information m_2 from the input m for running $m' \xleftarrow{\$} \text{Join-L}(P, \text{id}_{P'}, \text{gid}, gs_{\Pi}^{\text{gid}}, m_1 \parallel m_2)$. After that, P encrypts m' using the AEAD_{PP} key k_{PP} , a random nonce, and an associated data consisting of both parties' sign-up messages. Finally, the leader P outputs the AEAD_{PP} ciphertext and nonce.
- **Join-P'** $(P, \text{id}_{P'}, \text{gid}, gs, m)$: The participant P first parses $(gs_{\Pi}^{\text{gid}}, pw^{\text{gid}}) \leftarrow gs$ from the group secret. If the original Join-P algorithm needs any incoming information from the leader P' , then the participant P extracts an AEAD_{PP} ciphertext and an AEAD_{PP} nonce from the input m . Next, the participant P decrypts the AEAD_{PP} ciphertext using the key k_{PP} , the AEAD_{PP} nonce, and an associated data consisting of both parties' sign-up messages, and obtains a message m_1 . Then, the participant P extracts other necessary information m_2 from the input m for running $m' \xleftarrow{\$} \text{Join-P}(P, \text{id}_{P'}, \text{gid}, gs_{\Pi}^{\text{gid}}, m_1 \parallel m_2)$. After that, P encrypts m' using the AEAD_{PP} key k_{PP} , a random nonce, and an associated data consisting of both parties' sign-up messages. Finally, the participant P outputs the AEAD_{PP} ciphertext and nonce.

Member Leave $\text{Leave}' = (\text{Leave-L}', \text{Leave-P}')$: These algorithms are identical to the original $\text{Leave} = (\text{Leave-L}, \text{Leave-P})$. Note that if a per-group state is erased, then the stored key k_{PP} must also be erased.

Key Rotation $\text{KeyRotat}' = (\text{KeyRotat-L}', \text{KeyRotat-P}')$: We define the sub-algorithms as follows. If any error occurs during the above execution, then the caller aborts and undoes the executions in this invocation.

- **KeyRotat-L'** (P, gid, m) : The leader P first runs the original $m_{\text{KRot}} \xleftarrow{\$} \text{KeyRotat-L}(P, \text{gid}, m)$. Then, the leader P extracts the portion $c_{\bar{P}}$ in m_{KRot} that is specific to every participant \bar{P} in the group gid , followed by encrypting it using the stored corresponding AEAD_{PP}

key k_{PP} , a random nonce, and an associated data consisting of both parties' sign-up messages as in the Participant Join phase. Finally, the leader P outputs the AEAD_{PP} ciphertext and nonce for every participant \bar{P} in the group gid .

- **KeyRotat- $P'(P, \text{gid}, m)$:** The participant P first extracts an AEAD_{PP} ciphertext and an AEAD_{PP} nonce from the input m . Next, P recovers a message m_1 from the AEAD_{PP} ciphertext using the stored corresponding AEAD_{PP} key k_{PP} , the AEAD_{PP} nonce, and an associated data consisting of both parties' sign-up messages as in the Participant Join phase. Then, the participant P extracts other necessary information m_2 from the input m for running $m_{\text{KRot}} \stackrel{\$}{\leftarrow} \text{KeyRotat-P}(P, \text{gid}, m_1 \parallel m_2)$ and outputting m_{KRot} .

The theorem below shows that our PP transformation provably turns a Sec-mGKD-pki secure Π into a Sec-mGKD-pw secure $\Pi' = \text{PP}[\Pi, \text{PAKE}_{PP}, \text{AEAD}_{PP}]$ protocol. We provide a proof sketch in Appendix D.

Theorem 2. Let Π denote a mGKD protocol. Let PAKE_{PP} denote a password-authenticated key exchange scheme. Let AEAD_{PP} denote an authenticated encryption with associated data scheme. Let $\Pi' = \text{PP}[\Pi, \text{PAKE}_{PP}, \text{AEAD}_{PP}]$. Let $\mathcal{D} = \mathcal{D}_{\Pi} \times \mathcal{D}_{pw}$ denote the distribution of the group secrets. Assume the $\epsilon_{\text{PAKE}_{PP}, \mathcal{D}_{pw}}^{\text{w-PAKE}}$ -w-PAKE security of the underlying PAKE_{PP} , the $\epsilon_{\text{AEAD}_{PP}}^{\text{d-frob}}$ -d-FROB security, $\epsilon_{\text{AEAD}_{PP}}^{\text{ind}\$-cca}$ -IND $\$$ -CCA security, and $\epsilon_{\text{AEAD}_{PP}}^{\text{cti-cpa}}$ -cti-cpa security of AEAD_{PP} . If there exists any PPT attacker \mathcal{A} that breaks the Sec-mGKD-pw security of Π' , then there must exist a PPT attacker \mathcal{B} that breaks the Sec-mGKD-pki security of Π such that

$$\text{Adv}_{\Pi'}^{\text{Sec-mGKD-pw}}(\mathcal{A}) \leq q_{\text{NewGroup}} \left(\epsilon_{\text{PAKE}_{PP}, \mathcal{D}_{pw}}^{\text{w-PAKE}} + \epsilon_{\text{AEAD}_{PP}}^{\text{d-frob}} + c_{\text{maxReg}} (\epsilon_{\text{AEAD}_{PP}}^{\text{cti-cpa}} + \epsilon_{\text{AEAD}_{PP}}^{\text{ind}\$-cca}) + \text{Adv}_{\Pi}^{\text{Sec-mGKD-pki}}(\mathcal{B}) \right)$$

where c_{maxReg} denotes the maximal number of register requests for every group and $q_{\mathcal{O}}$ denotes the maximal number of queries to any oracle \mathcal{O} .

Below, we further show that our PP transformation preserves the Sec-mGKD-pki security of the original mGKD protocol Π . We provide a proof sketch in Appendix E.

Theorem 3. Let Π denote a mGKD protocol. Let PAKE_{PP} denote a password-authenticated key exchange scheme. Let AEAD_{PP} denote an authenticated encryption with associated data. Let $\Pi' = \text{PP}[\Pi, \text{PAKE}_{PP}, \text{AEAD}_{PP}]$. If there exists any PPT attacker \mathcal{A} that breaks the Sec-mGKD-pki security of Π' , then there exists another PPT attacker \mathcal{B} that breaks the Sec-mGKD-pki security of Π such that

$$\text{Adv}_{\Pi'}^{\text{Sec-mGKD-pki}}(\mathcal{A}) \leq \text{Adv}_{\Pi}^{\text{Sec-mGKD-pki}}(\mathcal{B})$$

Combining these two theorems, our PP transformation provably endows a Sec-mGKD-pki secure mGKD protocol Π with Sec-mGKD-pw security while preserving the original one, i.e., Sec-mGKD-pw+ security due to Corollary 1.

6.2. Application to the Zoom Library

Then, we illustrate how to apply our PP transformation to the Zoom library in Section 5.1 by using a 2-pass PAKE_{PP} scheme and an AEAD_{PP} scheme. We call the transformed version ZoomPAKE and show the resulting protocol in Figure 1, where we use boxes to indicate the modifications. Note that Zoom achieves Sec-mGKD-pki security without relying on passcodes, as stated in Remark 1. The passcodes for the server's access control underlying Zoom are redundant in the stronger threat model, and can therefore be set to the empty string without impacting Sec-mGKD-pw security. In the following we assume that the passcode is set to the empty string, which amounts to replacing the passcode with the new PAKE password computations.

The Sign Up and Member Leave phases are unchanged.

ZoomPAKE Group Schedule Phase: This algorithm is nearly identical to the original one except that the leader sends the new password to authorized parties over an out-of-band channel instead of sending the passcode to the server.

ZoomPAKE Register Phase: Similar to the previous, parties no longer need to send the passcode to the server. Then, each participant P' runs the first pass of PAKE_{PP} on the password pw^{gid} for a ciphertext $c_{PP,1}^{(P', \text{gid})}$ and outputs both the original outgoing messages and $c_{PP,1}^{(P', \text{gid})}$.

ZoomPAKE Participant Join Phase: Our PP transformation modifies both leaders' and participants' execution.

From the leader P 's side, P first parses an additional PAKE_{PP} ciphertext $c_{PP,1}^{(P', \text{gid})}$ from the input, and uses the group secret as the password pw^{gid} . Next, the leader P runs the second pass of PAKE_{PP} with necessary input for a key $k_{PP}^{(P', \text{gid})}$ and a ciphertext $c_{PP,2}^{(P', \text{gid})}$. Then, P stores $k_{PP}^{(P', \text{gid})}$ and samples a random nonce of length l_{PP} uniformly at random. After that, P executes the original computation. When the ZBox encryption $c_{P'}^{\text{gid}}$ is derived, the leader P further re-encrypts it using the key $k_{PP}^{(P', \text{gid})}$, the random nonce $\text{nonce}_{PP}^{P'}$, and the associated data consisting both parties' sign-up messages. The PAKE_{PP} ciphertext $c_{PP,2}^{(P', \text{gid})}$ and the AEAD_{PP} ciphertext are output.

From the participant P' 's side, P' first parses two more components from the input messages: a PAKE_{PP} ciphertext $c_{PP,2}^{(P', \text{gid})}$ and a nonce $\text{nonce}_{PP}^{P'}$. The participant P' also uses the group secret gs as the password pw^{gid} . Note that the original ciphertext $c_{P'}^{\text{gid}}$ is not the one of ZBox anymore but the one of AEAD_{PP} . Next, P' runs PAKE_{PP} for a key $k_{PP}^{(P', \text{gid})}$ and decrypts the AEAD_{PP} ciphertext $c_{P'}^{\text{gid}}$ using the key $k_{PP}^{(P', \text{gid})}$, the nonce $\text{nonce}_{PP}^{P'}$, and the associated data consisting of the leader P 's and the participant P' 's sign-up messages, for the original ZBox ciphertext. If any error occurs during this step, the participant P' simply aborts. Otherwise, the key $k_{PP}^{(P', \text{gid})}$ is stored locally. The remaining computation of P' remains the same.

ZoomPAKE Key Rotation Join Phase: The key rotation phase is very similar to the original one. The only difference

from the leader P 's side is that P has to encrypt the ZBox ciphertext using AEAD_{PP} for every participant \bar{P} in his local party set, i.e., $\text{id}_{\bar{P}} \in \pi_P^{\text{gid}}.GP$ and $\text{id}_{\bar{P}} \neq \text{id}_P$, using the stored key $k_{\text{PP}}^{(\bar{P}, \text{gid})}$, output by PAKE_{PP} , a independently random nonce $\text{nonce}_{\text{PP}}^{\bar{P}}$, the associated data consisting of the sign-up messages of P and \bar{P} . The output is a ciphertext bundle that includes AEAD_{PP} ciphertexts rather than the original ZBox ciphertexts.

When receiving the AEAD_{PP} ciphertext, a participant P' first decrypts it by using the stored key $k_{\text{PP}}^{(P', \text{gid})}$ for a ZBox ciphertext $c_{\text{PP}}^{\text{gid}}$. Then, P' simply runs the original KeyRotat-P algorithm using the new ciphertext $c_{\text{PP}}^{\text{gid}}$.

Instantiation Suggestions. We suggest to instantiate the underlying PAKE_{PP} with CPace [25] or SPAKE2 [27] for the w-PAKE security, see Section A.3. The AEAD_{PP} can be instantiated with CAU-C4 or CAU-SIV-C4 [42] for the d-FROB security, see Section A.1.

7. Technical Summary

In this paper, we propose a new mGKD protocol that captures the behaviors of the Zoom library and three associated security models: the basic Sec-mGKD-pki model considers restricted end-to-end encrypted security assuming the existence of a trusted PKI; the Sec-mGKD-pw model captures full end-to-end encrypted security without any trusted PKI; and the Sec-mGKD-pw+ that combines the Sec-mGKD-pki and Sec-mGKD-pw models. We proved that the Zoom library version 4.0 satisfies the basic Sec-mGKD-pki, but does not provide end-to-end Sec-mGKD-pw security.

To improve the Sec-mGKD-pki security of any mGKD protocol (including the Zoom library) to the Sec-mGKD-pw+ security, we propose a novel PP transformation that makes use of the group secret transmitted over out-of-band channels and cryptographic primitives PAKE and AEAD. Intuitively, to get the group keys that encrypts the real messages in the group chat, every participant must first additionally execute PAKE with the group leader for a shared key. This shared key is peer-wise independent: the group leader knows all shared keys and the participant only knows the one that it produces. Whenever the leader needs to rotate and distribute a new group key, the leader must additionally “wrap” the original ciphertext, i.e., encrypt it using the shared keys and AEAD, and every participant needs to first unwrap the original ciphertext in order to recover the real group keys. In particular, the application of our PP transform to the Zoom library is very efficient in terms of the communication rounds, as it does not cause any additional round trip time.

References

- [1] “How many people use Zoom?” 2022, <https://www.zippia.com/advice/zoom-meeting-statistics/>, (Accessed Feb 2023).
- [2] T. Isobe and R. Ito, “Security analysis of end-to-end encryption for Zoom meetings,” *IEEE Access*, vol. 9, pp. 90 677–90 689, 2021.
- [3] —, “Security Analysis of End-to-End Encryption for Zoom Meetings,” Cryptology ePrint Archive, Paper 2021/486, 2021, <https://eprint.iacr.org/2021/486>.

- [4] C. Cremers, E. Ronen, and M. Zhao. (2023) Multi-Stage Group Key Distribution and PAKEs: Securing Zoom Groups against Malicious Servers without New Security Elements (Full version with detailed proofs). <https://cispa.saarland/group/cremers/publications/extended/msgkd-pakes.html>.
- [5] M. J. Moyer, J. R. Rao, and P. Rohatgi, “A survey of security issues in multicast communications,” *IEEE network*, vol. 13, no. 6, 1999.
- [6] S. Rafaeeli and D. Hutchison, “A Survey of Key Management for Secure Group Communication,” *ACM Comput. Surv.*, vol. 35, no. 3, p. 309–329, sep 2003, <https://doi.org/10.1145/937503.937506>.
- [7] P. S. Kruus, “A survey of multicast security issues and architectures,” NAVAL RESEARCH LAB WASHINGTON DC, Tech. Rep., 1998.
- [8] C. Boyd, A. Mathuria, and D. Stebila, *Protocols for Authentication and Key Establishment, Second Edition*, ser. Information Security and Cryptography. Springer, 2020.
- [9] A. J. Menezes, P. C. Van Oorschot, and S. A. Vanstone, *Handbook of applied cryptography*. CRC press, 2018.
- [10] B. Poettering, P. Rösler, J. Schwenk, and D. Stebila, “SoK: Game-based Security Models for Group Key Exchange,” Cryptology ePrint Archive, Paper 2021/305, 2021, <https://eprint.iacr.org/2021/305>.
- [11] H. Harney and C. Muckenhirn. Group Key Management Protocol (GKMP) Specification. <https://www.rfc-editor.org/rfc/rfc2093>.
- [12] —. Group Key Management Protocol (GKMP) Architecture. <https://www.rfc-editor.org/rfc/rfc2094>.
- [13] C. K. Wong, M. Gouda, and S. S. Lam, “Secure group communications using key graphs,” *IEEE/ACM transactions on networking*, vol. 8, no. 1, pp. 16–30, 2000.
- [14] A. T. Sherman and D. A. McGrew, “Key establishment in large dynamic groups using one-way function trees,” *IEEE transactions on Software Engineering*, vol. 29, no. 5, pp. 444–458, 2003.
- [15] J. Goshi and R. E. Ladner, “Algorithms for dynamic multicast key distribution trees,” in *Proceedings of the twenty-second annual symposium on Principles of distributed computing*, 2003, pp. 243–251.
- [16] L. Xu and C. Huang, “Computation-efficient multicast key distribution,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 19, no. 5, pp. 577–587, 2008.
- [17] H. Lu, “A novel high-order tree for secure multicast key management,” *IEEE Transactions on Computers*, vol. 54, no. 2, pp. 214–224, 2005.
- [18] Z. Liu, Y. Lai, X. Ren, and S. Bu, “An efficient LKH tree balancing algorithm for group key management,” in *2012 International Conference on Control Engineering and Communication Technology*. IEEE, 2012, pp. 1003–1005.
- [19] A. Perrig, “Efficient collaborative key management protocols for secure autonomous group communication,” in *International Workshop on Cryptographic Techniques and E-Commerce (CrypTEC’99)*, 1999.
- [20] Y. Kim, A. Perrig, and G. Tsudik, “Simple and fault-tolerant key agreement for dynamic collaborative groups,” in *Proceedings of the 7th ACM Conference on Computer and Communications Security*, 2000, pp. 235–244.
- [21] K. Cohn-Gordon, C. Cremers, L. Garratt, J. Millican, and K. Milner, “On Ends-to-Ends Encryption: Asynchronous group messaging with strong security guarantees,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, 2018.
- [22] K. Klein, G. Pascual-Perez, M. Walter, C. Kamath, M. Capretto, M. Cueto, I. Markov, M. Yeo, J. Alwen, and K. Pietrzak, “Keep the dirt: tainted treekem, adaptively and actively secure continuous group key agreement,” in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021, pp. 268–284.
- [23] J. Alwen, S. Coretti, Y. Dodis, and Y. Tselekounis, “Security analysis and improvements for the IETF MLS standard for group messaging,” in *Advances in Cryptology—CRYPTO 2020: 40th Annual International Cryptology Conference, Santa Barbara, CA, USA, August, 2020, Proceedings, Part I*. Springer, 2020, pp. 248–277.

- [24] R. Barnes, B. Beurdouche, R. Robert, J. Millican, E. Omara, and K. Cohn-Gordon. The Messaging Layer Security (MLS) Protocol. <https://datatracker.ietf.org/doc/html/draft-ietf-mls-protocol-20>.
- [25] B. Haase and B. Labrique, “AuCPace: Efficient verifier-based PAKE protocol tailored for the IIoT,” *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2019, no. 2, p. 1–48, Feb. 2019, <https://tches.iacr.org/index.php/TCHES/article/view/7384>.
- [26] M. Abdalla, B. Haase, and J. Hesse, “Security analysis of CPace,” in *Advances in Cryptology—ASIACRYPT 2021: 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6–10, 2021, Proceedings, Part IV*. Springer, 2021, pp. 711–741.
- [27] M. Abdalla and D. Pointcheval, “Simple password-based encrypted key exchange protocols,” in *Topics in Cryptology—CT-RSA 2005: The Cryptographers’ Track at the RSA Conference 2005, San Francisco, CA, USA, February, 2005. Proceedings*. Springer, 2005, pp. 191–208.
- [28] M. Abdalla and M. Barbosa, “Perfect Forward Security of SPAKE2,” Cryptology ePrint Archive, Paper 2019/1194, 2019, <https://eprint.iacr.org/2019/1194>.
- [29] M. Abdalla, M. Barbosa, T. Bradley, S. Jarecki, J. Katz, and J. Xu, “Universally Composable Relaxed Password Authenticated Key Exchange,” Cryptology ePrint Archive, Paper 2020/320, 2020, <https://eprint.iacr.org/2020/320>.
- [30] E. Bresson, O. Chevassut, and D. Pointcheval, “Group Diffie-Hellman key exchange secure against dictionary attacks,” in *Advances in Cryptology—ASIACRYPT 2002: 8th International Conference on the Theory and Application of Cryptology and Information Security Queenstown, New Zealand, December 1–5, 2002 Proceedings 8*. Springer, 2002, pp. 497–514.
- [31] M. Abdalla, E. Bresson, O. Chevassut, and D. Pointcheval, “Password-based group key exchange in a constant number of rounds,” in *Public Key Cryptography—PKC 2006: 9th International Conference on Theory and Practice in Public-Key Cryptography, New York, NY, USA, April 24–26, 2006. Proceedings 9*. Springer, 2006, pp. 427–442.
- [32] M. Abdalla and D. Pointcheval, “A scalable password-based group key exchange protocol in the standard model,” in *ASIACRYPT*, vol. 4284. Springer, 2006, pp. 332–347.
- [33] M. Abdalla, J.-M. Bohli, M. I. G. Vasco, and R. Steinwandt, “(Password) authenticated key establishment: From 2-party to group,” in *Theory of Cryptography: 4th Theory of Cryptography Conference, TCC 2007, Amsterdam, The Netherlands, February 21–24, 2007. Proceedings 4*. Springer, 2007, pp. 499–514.
- [34] Y. Dodis, D. Jost, B. Kesavan, and A. Marcedone, “End-to-End Encrypted Zoom Meetings: Proving Security and Strengthening Liveness,” in *Advances in Cryptology—EUROCRYPT 2023: 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, Proceedings*. Springer-Verlag, 2023.
- [35] J. Blum, S. Booth, B. Chen, O. Gal, M. Krohn, J. Len, K. Lyons, A. Marcedone, M. Maxim, M. E. Mou, A. Namavari, J. O’Connor, S. Rien, M. Steele, M. Green, L. Kissner, and A. Stamos. Zoom end-to-end encryption whitepaper. <https://github.com/zoom/zoom-e2e-whitepaper> Version 4.0 (Released on 18.11.2022).
- [36] Q. Dang, *Recommendation for Applications Using Approved Hash Algorithms*. Special Publication (NIST SP), National Institute of Standards and Technology, Gaithersburg, MD, 2012-08-24 2012, https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=911479 (Accessed Feb 2023).
- [37] J. Brendel, C. Cremers, D. Jackson, and M. Zhao, “The provable security of Ed25519: theory and practice,” in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021, pp. 1659–1676.
- [38] S. Arciszewski, “XChaCha: eXtended-nonce ChaCha and AEAD_XChaCha20_Poly1305,” Internet Engineering Task Force, Internet-Draft draft-irtf-cfrg-xchacha-03, Jan. 2020, work in Progress: <https://datatracker.ietf.org/doc/draft-irtf-cfrg-xchacha/03/>.
- [39] J. R. Vacca, *Public key infrastructure: building trusted applications and Web services*. Auerbach Publications, 2004.
- [40] P. Eckersley and J. Burns, “The (Decentralized) SSL Observatory,” in *20th USENIX Security Symposium (USENIX Security 11)*. San Francisco, CA: USENIX Association, Aug. 2011, <https://www.usenix.org/conference/usenix-security-11/decentralized-ssl-observatory>.
- [41] J. Stapleton, “PKI Under Attack,” *ISSA*, vol. 11, no. 3, 2013, https://cdn.ymaws.com/www.members.issa.org/resource/resmgr/JournalPDFs/PKI_Under_Attack_ISSA0313.pdf.
- [42] M. Bellare and V. T. Hoang, “Efficient schemes for committing authenticated encryption,” in *Advances in Cryptology—EUROCRYPT 2022: 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30–June 3, 2022, Proceedings, Part II*. Springer, 2022, pp. 845–875.
- [43] P. Rogaway, “Authenticated-Encryption with Associated-Data,” in *Proceedings of the 9th ACM Conference on Computer and Communications Security*, ser. CCS ’02, 2002, p. 98–107.
- [44] G. Barwell, D. Page, and M. Stam, “Rogue decryption failures: Reconciling AE robustness notions,” in *Cryptography and Coding: 15th IMA International Conference, IMACC 2015, Oxford, UK, December 15–17, 2015. Proceedings 15*. Springer, 2015, pp. 94–111.
- [45] J. Brendel, M. Fischlin, F. Günther, and C. Janson, “PRF-ODH: Relations, instantiations, and impossibility results,” in *Advances in Cryptology—CRYPTO 2017: 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20–24, 2017, Proceedings, Part III 37*. Springer, 2017, pp. 651–681.
- [46] D. J. Bernstein, T. Lange, and P. Schwabe. NaCl: Networking and Cryptography library. <https://nacl.cr.yp.to/>, (Accessed Jan 2023).
- [47] F. Denis. The Sodium cryptography library. <https://doc.libsodium.org/>, (Accessed Jan 2023).

Appendix A. Preliminaries

Due to the page limit, we only introduce novel and customized security notions here and recall other notions in our full version [4].

A.1. Authenticated Encryption with Associated Data

Definition 8 ([43]). Let Key , Nonce , Data , Message , Ciphertext respectively denote the space of keys, nonces, associated data, messages, and ciphertexts. An authenticated encryption with associated data scheme $\text{AEAD} = (\text{AEAD.Enc}, \text{AEAD.Dec})$ is a tuple of algorithms where

- AEAD.Enc the encryption algorithm inputs a key $k \in \text{Key}$, a nonce $\text{nonce} \in \text{Nonce}$, an associated data $D \in \text{Data}$, and a message m and (deterministically) outputs a ciphertext c , i.e., $c \leftarrow \text{AEAD.Enc}(k, \text{nonce}, D, m)$.
- AEAD.Dec the decryption algorithm inputs a key $k \in \text{Key}$, a nonce $\text{nonce} \in \text{Nonce}$, an associated data $D \in \text{Data}$, and a ciphertext $c \in \text{Ciphertext}$ and deterministically outputs a message $m \in \text{Message} \cup \{\perp\}$, i.e., $m \leftarrow \text{AEAD.Dec}(k, \text{nonce}, D, c)$.

Definition 9. We say an AEAD is ϵ -ind $\$$ -cca (resp. cti-cpa) secure, if the below defined advantage of any PPT attacker \mathcal{A} against $\text{Exp}_{\text{AEAD}}^{\text{IND}\$-\text{CCA}}$ (resp. $\text{Exp}_{\text{AEAD}}^{\text{CTI-CPA}}$) experiment in Figure 2 is bounded by,

$$\text{Adv}_{\text{AEAD}}^{\text{IND}\$-\text{CCA}} := \left| \Pr[\text{Exp}_{\text{AEAD}}^{\text{IND}\$-\text{CCA}}(\mathcal{A}) = 1] - \frac{1}{2} \right| \leq \epsilon$$

$$\text{Adv}_{\text{AEAD}}^{\text{CTI-CPA}} := \Pr[\text{Expr}_{\text{AEAD}}^{\text{CTI-CPA}}(\mathcal{A}) = 1] \leq \epsilon$$

Below, we define two customized notions. The (n, m) -frob notion ensures that each ciphertext cannot be decrypted two distinct messages upon different nonces. The d-frob security ensures that each ciphertext cannot be decrypted to two valid messages upon different associated data. They both are implied by the CMT-4 security in [42].

Definition 10. We say an AEAD has ϵ - (n, m) -frob (resp. d-frob) secure, if the below defined advantage of any PPT attacker \mathcal{A} against $\text{Expr}_{\text{AEAD}}^{(n, m)\text{-FROB}}$ (resp. $\text{Expr}_{\text{AEAD}}^{\text{d-FROB}}$) experiment in Figure 3 is bounded by,

$$\text{Adv}_{\text{AEAD}}^{(n, m)\text{-FROB}} := \Pr[\text{Expr}_{\text{AEAD}}^{(n, m)\text{-FROB}}(\mathcal{A}) = 1] \leq \epsilon$$

$$\text{Adv}_{\text{AEAD}}^{\text{d-FROB}} := \Pr[\text{Expr}_{\text{AEAD}}^{\text{d-FROB}}(\mathcal{A}) = 1] \leq \epsilon$$

A.2. lr-prf-ODH Assumption

We recall the generic lr-prf-ODH definition in [45].

Definition 11. Let ECDH denote a cyclic group G with order q over an elliptic curve EC with a generator g . Let $H : \mathbb{G} \times \{0, 1\}^* \rightarrow \{0, 1\}^{lh}$ denote a function. We define a generic security notion lr-prf-ODH which is parameterized by $l, r \in \{n, s, m\}$ indicating how often the attacker is allowed to query a certain “left” resp. “right” oracle (ODH_u resp. ODH_v) where n indicates that no query is allowed, s that a single query is allowed, and m that multiple (polynomially many) queries are allowed to the respective side. Consider the following security game $\text{Expr}_{\text{ECDH}, H}^{\text{lr-prf-ODH}}$ between a challenger \mathcal{C} and a PPT attacker \mathcal{A} .

- 1) The challenger \mathcal{C} samples $u \xleftarrow{\$} \mathbb{Z}_q$ and provides \mathbb{G} , g , and g^u to the attacker \mathcal{A} .
- 2) If $l = m$, \mathcal{A} can issue arbitrarily many queries to the following oracle ODH_u .
 ODH_u **oracle.** On a query of the form (S, x) , \mathcal{C} first checks if $S \notin \mathbb{G}$ and returns \perp if this is the case. Otherwise, it computes $y \leftarrow H(S^u, x)$ and returns y .
- 3) Eventually, \mathcal{A} issues a challenge query x^* . On this query, \mathcal{C} samples $v \xleftarrow{\$} \mathbb{Z}_q$ and a bit $b \xleftarrow{\$} \{0, 1\}$ uniformly at random. It then computes $y_0^* = H(g^{uv}, x^*)$ and samples $y_q^* \xleftarrow{\$} \{0, 1\}^{lh}$ uniformly at random. The challenger returns (g^v, y_b^*) to \mathcal{A} .
- 4) Next, \mathcal{A} may issue (arbitrarily interleaved) queries to the following oracles ODH_u and ODH_v (depending on l and r).

ODH_u **oracle.** The attacker \mathcal{A} may ask no ($l = n$), a single ($l = s$), or arbitrarily many ($l = m$) queries to this oracle. On a query of the form (S, x) , the challenger first checks if $S \notin G$ or $(S, x) = (g^v, x^*)$ and returns \perp if this is the case. Otherwise, it computes $y \leftarrow H(S^u, x)$ and returns y .

ODH_v **oracle.** The attacker \mathcal{A} may ask no ($r = n$), a single ($r = s$), or arbitrarily many ($r = m$) queries to this oracle. On a query of the form (T, x) , the challenger first checks if $T \notin G$ or $(T, x) = (g^u, x^*)$ and returns \perp

if this is the case. Otherwise, it computes $y \leftarrow H(T^v, x)$ and returns y .

- 5) At some point, \mathcal{A} stops and outputs a guess $b' \in \{0, 1\}$.

We say that the attacker wins the lr-prf-ODH game if $b' = b$. We say the lr-prf-ODH problem is ϵ -hard over ECDH and H , if the advantage of any PPT attacker \mathcal{A} that wins above $\text{Expr}_{\text{ECDH}, H}^{\text{lr-prf-ODH}}$ experiment is bounded by ϵ .

In this paper, we only need the mn-prf-ODH assumption.

A.3. Password-Authenticated Key Exchange

The password-authenticated key exchange (PAKE) protocols allow two parties to establish a high-entropy key over an insecure channel using a shared low-entropy password. Below, we first define a weak security (w-PAKE) security model against a PAKE protocol. This model is weaker than and therefore implied by the security model defined in [26], [28], [29]. Thus, some modern and widely used PAKE schemes, including CPace [25] or SPAKE2 [27] that are respectively proven secure in [26] and [28], [29], are provably secure in this w-PAKE model.

Protocol Members. This weak semantic security model only considers the two-party setting. I.e., the PAKE protocol has only two members: either an initiator I or a responder R .

The initiator I indeed captures the behaviors of (all) participants in each group gid in our mGKD protocol. The responder R indeed captures the behaviors of the (unique) leader in each group gid in our mGKD protocol.

Long-Lived Keys / Passwords. The initiator I and the responder R hold the same password pw , which is sampled from a distribution \mathcal{D} . In many literature, the password is also called the “long-lived key”.

Protocol Execution. The interaction between an attacker \mathcal{A} and the protocol members occurs only via oracle queries, which model the attacker capabilities in a real attack. During the execution, the attacker may create several instances of a member. We consider the concurrent model, i.e., several instances may be active at any given time. Let U^{id} denote the instance with identifier id of a member U . Let $b \in \{0, 1\}$ be a bit chosen uniformly at random. The attacker \mathcal{A} can query following three oracles:

- $\text{SENDPAKE}(U^{\text{id}}, m)$: This query models an active attack, in which the adversary may tamper with the message being sent over the public channel. The output of this query is the message that the member instance U^{id} would generate upon receipt of message m .
- $\text{COMPROMISEPAKE}(U^{\text{id}})$: This query models the misuse of session keys by a member. If a session key is not defined for instance U^{id} or if a TESTPAKE query was asked to either U^{id} or to its partner, then return \perp . Otherwise, return the session key held by the instance U^{id} .
- $\text{TESTPAKE}(U^{\text{id}})$: This query tries to capture the adversary’s ability to tell apart a real session key from a random one. If no session key for instance U^{id} is defined or U^{id} is not fresh (which is defined below), then return the undefined

$\text{Exp}_{\text{AEAD}}^{\text{IND}\text{-}\text{CCA}}$: 1 $b \leftarrow_{\mathcal{R}} \{0, 1\}$ 2 $\mathcal{L}_{\text{ENC}}, \mathcal{L}_{\text{DEC}} \leftarrow \emptyset$ 3 $k \leftarrow_{\mathcal{R}} \mathcal{K}$ 4 $b' \leftarrow_{\mathcal{R}} \mathcal{A}^{\text{ENC}, \text{DEC}}()$ 5 return $\llbracket b = b' \rrbracket$	$\text{ENC}(\text{nonce}, D, m)$: 6 require $(\text{nonce}, -, -, -) \notin \mathcal{L}_{\text{ENC}}$ 7 require $(\text{nonce}, D, m, -) \notin \mathcal{L}_{\text{DEC}}$ 8 if $b = 0$ 9 $c \leftarrow \text{AEAD.Enc}(k, \text{nonce}, D, m)$ 10 else $c \leftarrow_{\mathcal{R}} \{0, 1\}^{\ell(m)}$ 11 $\mathcal{L}_{\text{ENC}} \stackrel{\perp}{\leftarrow} (\text{nonce}, D, m, c)$ 12 return c	$\text{DEC}(\text{nonce}, D, c)$: 13 require $(\text{nonce}, D, -, c) \notin \mathcal{L}_{\text{ENC}}$ 14 $m \leftarrow \text{AEAD.Dec}(k, \text{nonce}, D, c)$ 15 if $m \neq \perp$ 16 $\mathcal{L}_{\text{DEC}} \stackrel{\perp}{\leftarrow} (\text{nonce}, D, m, c)$ 17 return m	$\text{Exp}_{\text{AEAD}}^{\text{CTI-CPA}}$: 1 $\mathcal{L}_c \leftarrow \emptyset, k \leftarrow_{\mathcal{R}} \mathcal{K}$ 2 $(\text{nonce}, D, c) \leftarrow_{\mathcal{R}} \mathcal{A}^{\text{ENC}}()$ 3 require $c \notin \mathcal{L}_c$ 4 return $\llbracket \text{AEAD.Dec}(k, \text{nonce}, D, c) \neq \perp \rrbracket$ $\text{ENC}(\text{nonce}, D, m)$: 5 $c \leftarrow \text{AEAD.Enc}(k, \text{nonce}, D, m)$ 6 $\mathcal{L}_c \leftarrow \mathcal{L}_c \cup \{c\}$ 7 return c
---	--	--	--

Figure 2. IND $\text{-}\text{CCA}$ security [44] and CTI-CPA security [43] for an AEAD scheme. The functi on ℓ is a function that inputs the message length and outputs the corresponding ciphertext length.

$\text{Exp}_{\text{AEAD}}^{(n, m)\text{-FROB}}$: 1 $(c, (k_1, \text{nonce}_1, D_1), (k_2, \text{nonce}_2, D_2)) \leftarrow_{\mathcal{R}} \mathcal{A}()$ 2 require $\perp \notin \{c, k_1, \text{nonce}_1, D_1, k_2, \text{nonce}_2, D_2\}$ 3 $m_1 \leftarrow \text{AEAD.Dec}(k, \text{nonce}_1, D, c)$ 4 $m_2 \leftarrow \text{AEAD.Dec}(k, \text{nonce}_2, D, c)$ 5 require $(\text{nonce}_1, m_1) \neq (\text{nonce}_2, m_2)$ 6 return $\llbracket m_1 \neq \perp \rrbracket$ and $\llbracket m_2 \neq \perp \rrbracket$	$\text{Exp}_{\text{AEAD}}^{\text{d-FROB}}$: 1 $(c, (k_1, \text{nonce}_1, D_1), (k_2, \text{nonce}_2, D_2)) \leftarrow_{\mathcal{R}} \mathcal{A}()$ 2 require $\perp \notin \{c, k_1, \text{nonce}_1, D_1, k_2, \text{nonce}_2, D_2\}$ 3 $m_1 \leftarrow \text{AEAD.Dec}(k, \text{nonce}_1, D, c)$ 4 $m_2 \leftarrow \text{AEAD.Dec}(k, \text{nonce}_2, D, c)$ 5 require $D_1 \neq D_2$ 6 return $\llbracket m_1 \neq \perp \rrbracket$ and $\llbracket m_2 \neq \perp \rrbracket$
---	--

Figure 3. (n, m)-FROB security and d-FROB security for an AEAD scheme.

symbol \perp . Otherwise, return the session key for instance U^{id} if $b = 1$ or a random key of the same size if $b = 0$.

Notation. We say an instance U^{id} *opened* if a query $\text{COMPROMISEPAKE}(U^{\text{id}})$ has been made by the attacker. We say an instance U^{id} is *unopened* if it is not opened. We say an instance U^{id} *tested* if a query $\text{TESTPAKE}(U^{\text{id}})$ has been made by the attacker. We say an instance U^{id} is *untested* if it is not tested. We say an instance U^{id} has *accepted* if it goes into an accept mode after receiving the last expected protocol message.

Session Identifiers and Partnering. We define the *session identifiers* (sid) as the transcript of the conversation between the initiator and the responder instances before acceptance.

We say two instances I^{id_1} and R^{id_2} to be *partners* if the following conditions are met:

- Both I^{id_1} and R^{id_2} accept, and
- Both I^{id_1} and R^{id_2} share the same identifiers sid .

Freshness. The notion of freshness is defined to avoid cases in which adversary can trivially break the security of the scheme. The goal is to only allow the attacker to ask TESTPAKE queries to fresh oracle instances. More specifically, we say an instance U^{id} is fresh if it has accepted and if both U^{id} and its partner are unopened and untested.

Semantic Security. Consider an execution of the above experiment for an attacker \mathcal{A} against a PAKE protocol Π . The attacker \mathcal{A} wins the experiment if and only if \mathcal{A} guesses $b' = b$, where b is the hidden bit used by the TESTPAKE oracle. The advantage of \mathcal{A} breaking the weak security of Π is defined as

$$\text{Adv}_{\Pi, \mathcal{D}}^{\text{w-PAKE}}(\mathcal{A}) := |\Pr[\mathcal{A} \text{ wins}] - \frac{1}{2}|$$

We say that Π is $\epsilon_{\text{PAKE}, \mathcal{D}}^{\text{w-PAKE}}$ -w-PAKE secure if for any PPT attackers \mathcal{A} it always holds that

$$\text{Adv}_{\Pi, \mathcal{D}}^{\text{w-PAKE}}(\mathcal{A}) \leq \epsilon_{\Pi, \mathcal{D}}^{\text{w-PAKE}}$$

Appendix B. Cryptographic Algorithms

The Zoom library makes use of the interface and implementation of two building blocks in the NaCl [46]-inspired libsodium library [47]: *Signing* and *Authenticated Public-Key Encryption* (aka. Box).

ZSign.KG : 7 $(pk_{\text{ZSign}}, sk_{\text{ZSign}}) \leftarrow_{\mathcal{R}} \text{DS.KG}$ 8 return $(pk_{\text{ZSign}}, sk_{\text{ZSign}})$	$\text{ZSign.Sign}(sk_{\text{ZSign}}, \text{ctxt}, m)$: 11 $m' \leftarrow H_1(\text{ctxt}) \parallel H_1(m)$ 12 $\sigma \leftarrow \text{DS.Sign}(sk_{\text{ZSign}}, m')$ 13 return σ
$\text{ZSign.Vrfy}(pk_{\text{ZSign}}, \sigma, \text{ctxt}, m)$: 9 $m' \leftarrow H_1(\text{ctxt}) \parallel H_1(m)$ 10 return $\text{DS.Vrfy}(pk_{\text{ZSign}}, \sigma, m')$	

Figure 4. The Zoom-Signing algorithm ZSign. Zoom instantiates H_1 with SHA256 and DS with EdDSA over Ed25519.

Zoom Signing Algorithm: The construction of the Zoom Signing algorithm $\text{ZSign} = (\text{ZSign.KG}, \text{ZSign.Sign}, \text{ZSign.Vrfy})$ is depicted in Figure 4.

- The key generation algorithm ZSign.KG simply generates and outputs a DS key pair.
- The signing algorithm ZSign.Sign inputs a secret key sk_{ZSign} , a context ctxt , and a message m . The ZSign.Sign first computes the hash function H_1 over respective context ctxt and message m , followed by concatenating them. Then, the ZSign.Sign computes and outputs the signature of the concatenation using DS upon the input secret key sk_{ZSign} .
- The verification ZSign.Vrfy algorithm inputs a public key pk_{ZSign} , a signature σ , a context ctxt , and a message m . This ZSign.Vrfy algorithm simply computes the concatenation as in the signing algorithm and outputs the DS verification result DS.Vrfy upon the public key pk_{ZSign} , the signature σ , and the concatenation.

Zoom Authenticated Public-Key Encryption (aka. Box)

ZBox.KG:	ZBox.Enc($sk_{ZBox}^S, pk_{ZBox}^R, \text{ctxt}_{H_2}, \text{ctxt}_{Cipher}, \text{Meta}, m$):	ZBox.Dec($sk_{ZBox}^R, pk_{ZBox}^S, \text{ctxt}_{H_2}, \text{ctxt}_{Cipher}, \text{Meta}, c$):
1 $(pk_{ZBox}, sk_{ZBox}) \xleftarrow{\$} \text{ECDH}$	3 $\text{nonce} \xleftarrow{\$} \{0, 1\}^l$	10 $\text{Parse}(c', \text{nonce}) \leftarrow c$
2 return (pk_{ZBox}, sk_{ZBox})	4 $K' \leftarrow sk_{ZBox}^S \cdot pk_{ZBox}^R$	11 $K' \leftarrow sk_{ZBox}^R \cdot pk_{ZBox}^S$
	5 $K \leftarrow H_2(K', \text{ctxt}_{H_2})$	12 $K \leftarrow H_2(K', \text{ctxt}_{H_2})$
	6 $D \leftarrow H_1(\text{ctxt}_{Cipher}) \parallel H_1(\text{Meta})$	13 $D \leftarrow H_1(\text{ctxt}_{Cipher}) \parallel H_1(\text{Meta})$
	7 $c' \leftarrow \text{AEAD.Enc}(K, \text{nonce}, D, m)$	14 $m \leftarrow \text{AEAD.Dec}(K, \text{nonce}, D, c')$
	8 $c \leftarrow (c', \text{nonce})$	15 require $m \neq \perp$
	9 return c	16 return m

Figure 5. The Zoom-Box algorithm ZBox. We have that $l = 192$, and the underlying function H_1 denotes SHA256. The function H_2 denotes HKDF (using an empty salt parameter). ECDH is performed on Curve25519. “ \cdot ” denotes scalar multiplication. The AEAD is instantiated with xchacha20poly1305.

Algorithm: The Zoom Box algorithm $\text{ZBox} = (\text{ZBox.KG}, \text{ZBox.Enc}, \text{ZBox.Dec})$ is depicted in Figure 5.

- The key generation algorithm ZBox.KG samples and outputs a Diffie-Hellman key pair over an elliptic curve ECDH.
- The encryption algorithm ZBox.Enc takes as inputs a sender’s secret key sk_{ZBox}^S , a receiver’s public key pk_{ZBox}^R , two contexts ctxt_{H_2} and ctxt_{Cipher} , a meta data Meta, and a message m . It first samples a random nonce of bit length l . Next, it computes the Diffie-Hellman exchange of sk_{ZBox}^S and pk_{ZBox}^R , which is combined with the context ctxt_{H_2} and used as input to the hash function H_2 for a key K . Then, it computes the associated data D by concatenating $H_1(\text{ctxt}_{Cipher})$ and $H_1(\text{Meta})$. The output is the nonce nonce and an AEAD ciphertext produced from the key K , nonce nonce, data D , and message m .
- The decryption algorithm ZBox.Dec takes as input a receiver’s secret key sk_{ZBox}^R , a sender’s public key pk_{ZBox}^S , two contexts ctxt_{H_2} and ctxt_{Cipher} , a meta data Meta, and a ciphertext c . This algorithm parses the nonce nonce from the ciphertext c , computes the key K and the associated data D as in the encryption algorithm, and executes the AEAD decryption. If the AEAD outputs a message $m \neq \perp$, this algorithm simply outputs this m , and aborts otherwise.

Appendix C. Proof Sketch of Theorem 1

Proof Sketch. By the collision resistance of H_1 , the DS scheme underlying ZSign algorithm never signs two identical input. The euf-cma security of DS then ensures that all parties obtains other parties’ honest sign-up messages in every group. Then, we consider two cases. If \mathcal{A} can win via the event E_{KAuth} , then we can easily guess the group identifier $\widetilde{\text{gid}}$ with some leader $P^{\widetilde{\text{gid}}}$ and party \widetilde{P} that trigger \mathcal{A} to win with probability at least $1/c_{\text{maxReg}} q_{\text{NEWGROUP}}$. By mn-prf-ODH security of ECDH and H_2 underlying ZBox algorithm, the key of AEAD between $P^{\widetilde{\text{gid}}}$ and \widetilde{P} is indistinguishable from random. If \mathcal{A} can win via the event E_{KAuth} , then the attacker must be able to either forge a ciphertext or a nonce of AEAD, which breaks either the CTI-CPA security of (n, m) -FROB security of AEAD.

If \mathcal{A} win via the event E_{KPriv} , then we can easily guess the group identifier $\widetilde{\text{gid}}$ with some leader $P^{\widetilde{\text{gid}}}$ and the identifier of all parties in the set $GP^{(\widetilde{\text{gid}}, \widetilde{\text{gkid}})}$ for the tested group key index $\widetilde{\text{gkid}}$ with probability at least

$1/c_{\text{maxParty}} c_{\text{maxReg}}^{(n_{\text{party}}-1)} q_{\text{NEWGROUP}}$. By a sequence of $(n_{\text{party}} - 1)$ hybrid games, we know that all keys and ciphertext of AEAD between every participant in $\widetilde{\text{gid}}$ and the leader $P^{\widetilde{\text{gid}}}$ should be indistinguishable from random due to the mn-prf-ODH security of ECDH and H_2 and ind $\$$ -cca security of AEAD. Thus, \mathcal{A} cannot win via E_{KPriv} . \square

Appendix D. Proof Sketch of Theorem 2

Proof Sketch. We can easily guess the group $\widetilde{\text{gid}}$ that enables \mathcal{A} to win with probability at least q_{NEWGROUP} . Due to the w -PAKE security of the PAKE_{PP} scheme with \mathcal{D}_{pw} , we can ensure that the keys $k_{PP}^{(P', \widetilde{\text{gid}})}$ of all authorized participants P' in the group $\widetilde{\text{gid}}$, which are output by PAKE_{PP} of and will be used for AEAD_{PP}, are random. Moreover, the key $k_{PP}^{(P', \widetilde{\text{gid}})}$ produced by the leader $P^{\widetilde{\text{gid}}}$ of the group $\widetilde{\text{gid}}$ is either same as the one produced by the corresponding authorized participant P' or independently random. By c_{maxReg} hybrid games on the cti-cpa security of AEAD_{PP}, all authorized parties P' in the group $\widetilde{\text{gid}}$ must agree on all AEAD_{PP} ciphertexts with the leader $P^{\widetilde{\text{gid}}}$. By d-frob security of AEAD_{PP}, agreeing ciphertexts indicates that agreeing on the associated data, i.e., sign-up messages. Moreover, by the ind $\$$ -cca security of AEAD AEAD_{PP}, all AEAD_{PP} ciphertexts produced by $P^{\widetilde{\text{gid}}}$ for any unauthorized party are indistinguishable from random and therefore leaks no information about any group keys. Finally, if \mathcal{A} win via the event E_{KAuth} or E_{KPriv} against Π' , then \mathcal{A} can also win the same event against Π ’s Sec-mGKD-pki security. \square

Appendix E. Proof Sketch of Theorem 3

Proof Sketch. The proof can be easily given by a reduction. The attacker \mathcal{B} can easily simulates the real Sec-mGKD-pki game against Π' to \mathcal{A} by sampling all passwords for all groups and runs the PP transformation by himself. All information that \mathcal{B} needs can be obtained by querying its challenger. Finally, \mathcal{B} forwards the bit output by \mathcal{A} and wins whenever \mathcal{A} wins. \square

Appendix F. Meta-Review

F.1. Summary

The paper presents a protocol version of the audio-video chat application Zoom that aims to achieve end-to-end security, in a group chat setting, even in the face of a malicious server. They rely on PAKEs in order to construct said protocol. They also present a detailed formal security analysis of the protocol currently used in Zoom.

F.2. Scientific Contributions

- Provides a Valuable Step Forward in an Established Field
- Independent Confirmation of Important Results with Limited Prior Research

F.3. Reasons for Acceptance

- 1) The paper provides a valuable step forward in an established field, as it is one of the first formal studies of the protocol Zoom uses (which is a widely used service) specially in regards to its claims on end-to-end encryption. They formalise the security of the protocol and state that it is insecure in the face of malicious servers (which was previously know but not formalised).
- 2) The paper proposes a new protocol that achieves strong security/privacy notions in the face of a malicious server. This protocol can be used by Zoom.

F.4. Noteworthy Concerns

The paper is rigorous in their technical definitions, but an in-depth analysis of performance/efficiency considerations of the new protocol is lacking. This is very much needed given the constrains of working with video/audio, and to further emphasise the practicality of their protocol.

Appendix G. Response to the Meta-Review

Our proposed changes do not change the actual message transmission layer, only the key exchange parts. Moreover, our new solution is tailored to the current Zoom message flow, thus does not add extra communication rounds to the key exchange, which are the main performance constraint for this type of applications.