Accountable Tracing Signatures

Markulf Kohlweiss and Ian Miers

Abstract. Demands for lawful access to encrypted data are a long standing obstacle to integrating cryptographic protections into communication systems. A common approach is to allow a trusted third party (TTP) to gain access to private data. However, there is no way to verify that this trust is well place as the TTP may open all messages indiscriminately. Moreover, existing approaches do not scale well when, in addition to the content of the conversation, one wishes to hide ones identity. Given the importance of metadata this is a major problem. We propose a new signature scheme as an accountable replacement for group signatures, *accountable forward and backward tracing signatures*.

Keywords: Accountability, traceable signatures, group signatures

> Доверяй, но проверяй. (Trust, but verify.)
> — (alleged) Russian proverb

1 Introduction

For digital communication there are few scaling issues to prevent mass surveillance and, indeed, no easy way to detect it. To ensure privacy, we rely on cryptographic guarantees which are typically—indeed hopefully—absolute. However, precisely because they are inviolable, widespread deployment of such systems, e.g., in cloud services, often raises governmental objections or requires mandated ways of providing access. This in part contributes to an interesting contradiction: robust end-to-end cryptography (e.g. PGP, OTR, and ZRTP) is available and can be used given sufficient motivation (e.g. the need to hide criminal activity), but is not widely deployed in the systems we use on a daily basis for mail, chat, or voice communication. Mass surveillance is easy, yet targeted surveillance where the targets suspect they might be targeted and take precautions is hard.

For cryptographic schemes, a commonly proposed solution for lawful access is to appoint a trusted party (or multiple partially-trusted parties) who escrow(s) a user's identity or information about her actions for future retrieval. This has been applied to anonymous signatures [CvH91,FV10], e-cash schemes [FPV09], and even saw very limited (and ultimately failed) real-world deployment in the from of the Clipper chip for encryption.

Even if we discount the philosophical, political, and moral objections to mandating lawful access, these approaches, while in many cases technologically novel, fail to provide adequate privacy protections in light of recent developments.

First, in the face of nation-states that are willing to compromise hardware, penetrate systems, and coerce individuals, such an approach seems foolhardy: eventually the trusted party's key or unfettered access to it will be extracted even if the party is itself trustworthy.

Second, escrow systems typically fail to hide metadata—something the last few years has proved quite valuable—and are not easily modified to do so. If a user's messages can easily be located (and subsequently decrypted with an escrowed key), then the privacy protections are severely lacking as such a system leaks metadata. If this is impossible, locating the messages to decrypt relies on expensive operations such as trial decryption with every escrowed key. This doesn't scale.

The question is, if we absolutely have to allow lawful access, how do we overcome these issues? We begin with three insights:

1. If no third party can be trusted, then each user must, solely on their own, hold the necessary key for opening their messages and the authority must be able to accountably gain access to this key.
2. Communication is increasingly via (cloud based) services. Even after raising suspicion, a user will likely continue to resubscribe to the service and this can be leveraged to gain access to his escrow key.
3. We gain accountability by forcing the authority to verifiably quantify—either by (i) directly identifying, or (ii) statistically describing (e.g. in a transparency report) —which (re)subscriptions involved a search.

Our approach is as follows: as in a group signature, users escrow their identity under an escrow public key as part of the signature. When a user resubscribes, the authority searches them by either (i) replacing that key with one the authority holds, or (ii) requiring the user to provide the escrow private key encrypted under an authority provided key. The former results in tracing going forward, while the latter allows backward tracing as well. In case a user is not traced, the authority does not replace keys and provides a random group element for a user to encrypt their key under.

The core of our approach is a key-oblivious encryption scheme where public keys are randomizable and randomized keys cannot be linked to each other or the original key. As a result, without the entropy used in key randomization, users cannot tell if their key was replaced—as in the case of forward tracing—or if they are encrypting to a random group element or the authority's public key—as in the case of backward tracing. Because users cannot tell if either search mechanism was invoked, users remain oblivious. However, once the randomness is revealed, the authority is held accountable.

Definitionally, this approach adds a new requirement to escrow systems: authentication. As such a scheme depends on user re-subscription, it is imperative that users demonstrate, for every potential escrow operation, that they are currently subscribed. Similarly, the service forwarding messages must verify users' identities for every message to ensure it can provide lawful access. Authenticating this in the clear erodes the sender's privacy and doing so completely anonymously makes it impossible to efficiently locate a user's messages once they are targeted.

At first glance, group signatures seem like a reasonable solution. They allow anonymity revocation by a group manager while still providing privacy to everyone else. Group managers, however, can open any message and are simply trusted not to. For accountable escrow, this is not the case and the group manager is fundamentally an adversary who seeks to actively violate user privacy. None the less, group signatures are a decent starting point. However, we need to add properties to ensure that
- the group manager can only open messages from users she selects when (re)enrolling them
- users cannot tell if they are enrolled with an opening key
- the group manager can prove that a given enrollment did not allow tracing.

By starting with the definitions for group signatures we inherit the strong guarantees that allow the group manager to prove that someone authored a message if-and-only-if they actually did so, i.e., she cannot frame users. While this property is understandably neglected in many escrow schemes, given the legal context in which escrowed message are likely to be used as evidence in court, attribution seems to be a fundamental property that should be dealt with.

Our Contribution. We provide four contributions. First, we propose accountability as a novel security requirement of escrow systems and formalize it in the definitions for accountable tracing

signatures. Second, we provide a practical construction for an accountable forward-tracing signature complete with proofs of security. This allows an authority to covertly and accountably tag and trace user messages once the user becomes suspect. Third, we extend our approach with an interactive subscribe protocol to build an accountable backward-tracing signature scheme where all of a user's messages can be identified even retroactively. Finally, we show how to augment either approach to create an accountable *tracing signcryption* scheme where messages are encrypted and opening a signature reveals both the author and the message content, giving us an efficient accountable wiretapping system.

Related work. Several cryptographic schemes, both from the academic literature and in practice, use a trusted party for escrow. For example, in a group signature scheme [CvH91], a group manager allows users to sign messages as coming from some member of the group while he alone maintains the ability to provably identify who signed the message. A related problem is given a suspicious user, identify all messages they have authored. Systems that support both functionalities are called traceable signatures [KTY04]. Variants of these properties have been defined for e-cash systems, with owner tracing, and coin tracing being defined analogously. Interestingly there are e-cash schemes that hold the authority accountable [KV01,KV03]. Unfortunately, these schemes require interaction with the tracing authority for every transaction and, as a result, the techniques do not necessarily scale.

Finally, key-escrow, where an encrypted message can be decrypted by both the recipient and an escrow authority, is both a well studied topic and one that has seen at least limited real world deployment in the (failed) Clipper chip. A key escrow mechanism in which multiple trustees need to collaborate to decrypt has been proposed by [LRC14].

2 Key-oblivious encryption

We call a public-key encryption scheme *key-oblivious* if (i) it allows for a large set of public keys all related to the same secret key, if (ii) existing public keys can be randomized to generate related keys, and if (iii) it cannot be discerned, without knowledge of the secret key and the randomness used in their generation, how public keys are related. The existence of such schemes is cryptographic folklore and we do not claim much novelty here. In order to allow for accountability, we insist on the ability to prove, for a given public key, which key was randomized to produce it. This leads to slight variants of the standard key-privacy and plaintext-indistinguishability games.

2.1 Syntax

We formalize a key-oblivious encryption scheme OE as a collection $\mathsf{OE}.(\mathsf{GroupGen}, \mathsf{KeyGen}, \mathsf{KeyRand}, \mathsf{Enc}, \mathsf{Dec})$ of five algorithms:

$\mathsf{GroupGen}(1^\lambda) \to \mathcal{G}$**:** generates parameters, usually a prime order group.
$\mathsf{KeyGen}(\mathcal{G}) \to (pk, sk)$**:** generates a key pair.
$\mathsf{KeyRand}(pk) \to pk'$**:** randomizes an existing public key into a public key for the same secret key.
$\mathsf{Enc}(pk, m) \to ct$**:** standard encryption functionality.
$\mathsf{Dec}(sk, ct) \to m$**:** standard decryption functionality.

In definitions and in protocols, we sometimes make the randomness of $\mathsf{KeyRand}$ explicit and write $(pk'; r) \leftarrow \mathsf{KeyRand}(pk)$ and $pk' = \mathsf{KeyRand}(pk; r)$. For two fixed random public keys $pk^{(0)}, pk^{(1)}$, $pk = \mathsf{KeyRand}(pk^{(b)}; r)$ acts as a *hiding* and *binding* bit commitment scheme, with r its opening. (Similarly, we write $(ct; s) \leftarrow \mathsf{Enc}(pk, m)$ and $ct = \mathsf{Enc}(pk, m; s)$ to make the randomness of Enc explicit in zero-knowledge proofs.)

2.2 Definitions

In addition to *key randomizability*—which corresponds to the commitment being hiding, we require that ciphertexts are plaintext indistinguishable even when the adversary can randomize the target keys. We term this *plaintext indistinguishability under key randomization*(INDr). We also require such schemes to be *key private* [BBDP01], again with the modification that this holds even for adversarially randomized keys. Even though stronger variants of these properties exist, security under chosen plaintext attacks suffices for our purposes. We note that any key randomizable encryption scheme can be made key private via the added step of randomizing the public key prior to encryption.

Game KR $\triangleq$
$b \leftarrow \{0, 1\}$
$\mathcal{G} \leftarrow \mathsf{GroupGen}(1^\lambda)$
$(pk, sk) \leftarrow \mathsf{KeyGen}(\mathcal{G})$
$pk_0 \leftarrow \mathsf{KeyRand}(pk)$
$(pk_1, sk_1) \leftarrow \mathsf{KeyGen}(\mathcal{G})$
$b' \leftarrow \mathcal{A}(pk, pk_b)$
return $(b' = b)$

Game KPr $\triangleq$
$b \leftarrow \{0, 1\}$
$\mathcal{G} \leftarrow \mathsf{GroupGen}(1^\lambda)$
$(pk_0, sk_0) \leftarrow \mathsf{KeyGen}(\mathcal{G})$
$(pk_1, sk_1) \leftarrow \mathsf{KeyGen}(\mathcal{G})$
$(m, pk'_0, r_0, pk'_1, r_1, st) \leftarrow \mathcal{A}_0(\mathcal{G}, pk_0, pk_1)$
if $\neg(pk'_i = \mathsf{KeyRand}(pk_i; r_i)$ for $i \in \{0, 1\})$ **then**
   return $\perp$
$ct \leftarrow \mathsf{Enc}(pk'_b, m)$
$b' \leftarrow \mathcal{A}_1(ct, st)$
return $(b' = b)$

Game INDr $\triangleq$
$b \leftarrow \{0, 1\}$
$\mathcal{G} \leftarrow \mathsf{GroupGen}(1^\lambda)$
$(pk, sk) \leftarrow \mathsf{KeyGen}(\mathcal{G})$
$(pk', r, m_0, m_1, st) \leftarrow \mathcal{A}_0(\mathcal{G}, pk)$
if $pk' \neq \mathsf{KeyRand}(pk; r)$ **then**
   return $\perp$
$ct \leftarrow \mathsf{Enc}(pk', m_b)$
$b' \leftarrow \mathcal{A}_1(ct, st)$
return $(b' = b)$

Definition 1 (Key randomizability). *Let* OE *be a key-oblivious encryption scheme. Consider Game* KR *played by adversary $\mathcal{A}$: The key-randomizability advantage of $\mathcal{A}$, $\mathbf{Adv}^{\mathsf{KR}}(\mathcal{A})$ is defined as $2 \cdot Pr[\mathsf{KR} : \mathsf{true}] - 1$. A scheme is key randomizable if for any polynomial time $\mathcal{A}$ this advantage is negligible.*

Definition 2 (Key privacy under key randomization). *Let* OE *be a key-oblivious encryption scheme. Consider Game* KPr *played by adversary $\mathcal{A}$: The key-privacy advantage of $\mathcal{A}$, $\mathbf{Adv}^{\mathsf{KPr}}(\mathcal{A})$ is defined as $2 \cdot Pr[\mathsf{KPr} : \mathsf{true}] - 1$. A scheme is key private if for any polynomial time $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ this advantage is negligible.*

Definition 3 (Plaintext indistinguishability under key randomization). *Let* OE *be a key-oblivious encryption scheme. Consider Game* INDr *played by adversary $\mathcal{A}$: The plaintext distinguishing advantage of $\mathcal{A}$, $\mathbf{Adv}^{\mathsf{INDr}}(\mathcal{A})$ is defined as $2 \cdot Pr[\mathsf{INDr} : \mathsf{true}] - 1$. A scheme is secure if for any polynomial time $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ this advantage is negligible.*

3 Accountable forward-traceable signatures

In a group signature scheme, a group manager can open any suspect message. In a tracing signature scheme, the manager can test if any message belongs to a suspect user. In an accountable tracing signature, the manager can do the same but must later reveal which users she deems suspect.

3.1 Syntax

We formalize an accountable tracing signature scheme ATS as a collection ATS.(Setup, GKg, UKg, Enroll, Sign, Verify, Open, Judge, Account) of 9 algorithms:

$\mathsf{Setup}(1^\lambda) \rightarrow gp$: Generates the public parameters for security level λ.
$\mathsf{GKg}(gp) \rightarrow (gpk, gsk)$: Generates the initial group key-pair.

$\mathsf{UKg}(gp) \rightarrow (upk, usk)$**:** Generates a user key-pair. Each user has a public key *upk* and a corresponding private key *usk*. Per [BMW03] this is necessary to provide any meaning to the assertion that a user actually did sign an opened message: without it, the group manager is free to simply assert that a key of their generation actually belongs to a user.

$\mathsf{Enroll}(gsk, upk, epoch, t) \rightarrow (cert, w^{escrw})$**:** The authority produces a certificate on a user's escrow public key. This certificate either provides full anonymity ($t = 0$) or allows for tracing ($t = 1$) depending on the bit t. The authority stores the witness w^{escrw} to t and returns the certificate to the user. The certificate *cert* contains the time range for which the user is enrolled. We call this counter the *epoch* of the certificate.

$\mathsf{Sign}(gpk, cert, usk, m) \rightarrow \sigma$ **:** Takes the group public key, a certificate, the user's private key, and a message to sign as input. It outputs a signature that may contain an escrow of the user's identity. The *epoch* of the signature corresponds to the *epoch* of the certificate.

$\mathsf{Verify}(gpk, m, \sigma, epoch) \rightarrow \{0, 1\}$**:** Given the group public key, a message, its signature, and an epoch, verifies the signature is valid for the specified message and epoch.

$\mathsf{Open}(gsk, m, \sigma) \rightarrow (upk, \psi)$**:** Given the group private key, a message and its signature; if possible (i.e. if signed using an escrow certificate), return the public key of the user and a proof that the user generated the message-signature pair. Otherwise returns $\bot$.

$\mathsf{Judge}(gpk, m, \sigma, epoch, upk, \psi) \rightarrow \{0, 1\}$**:** Given the group public key, a message, its signature, an epoch, a user's public key, and a proof that the user generated the signature of that epoch on that message, verifies the proof.

$\mathsf{Account}(gpk, cert, w^{escrw}, t) \rightarrow \{0, 1\}$**:** Given a certificate, a bit t saying whether it escrows the user's identity, and a witness w^{escrw} returns 1 if the witness confirms the choice of t.

We will expose some implementation details for the sake of simplifying definitions and avoiding excessive scaffolding. We use *cert.epoch* to denote the epoch of a certificate, *cert.upk* to denote the user public key being certified, *gpk.gp* the parameters of the group public key, and *gsk.csk* to denote the certificate signing key—this makes an adversary with access to this key strictly stronger than an adversary with access to an ENROLL oracle.

3.2 Definitions

The first three properties are only slight adaptations of the standard group signature properties. The last two, *anonymity with accountability* and *trace-obliviousness*, are fundamentally different from the guarantees that group signatures can provide. These stem from the requirement to hold the group manager accountable.

Definition 4 (Anonymity under tracing). *Let* ATS *be an accountable forward-tracing signature scheme. Consider the following game played by an adversary* $\mathcal{A}$*:*

Game $\mathsf{AuT} \triangleq$
$b \leftarrow \{0, 1\}$
$gp \leftarrow \mathsf{Setup}(1^\lambda)$
$(gpk, gsk) \leftarrow \mathsf{GKg}(gp)$
$b' \leftarrow \mathcal{A}^{\mathsf{Ch, OPEN}}(gpk, gsk.csk)$
return $(b' = b)$

Oracle $\mathsf{Ch}(sk_0, sk_1, cert_0, cert_1, m, w_0^{escrw}, w_1^{escrw}, t) \triangleq$
$\sigma_0 \leftarrow \mathsf{Sign}(gpk, sk_0, cert_0, m)$
$\sigma_1 \leftarrow \mathsf{Sign}(gpk, sk_1, cert_1, m)$
if $(\sigma_0 \neq \bot \wedge \sigma_1 \neq \bot \wedge cert_0.epoch = cert_1.epoch \wedge$
$\quad \mathsf{Account}(gpk, cert_0, w_0^{escrw}, t) \wedge \mathsf{Account}(gpk, cert_1, w_1^{escrw}, t))$
$\quad Q \leftarrow Q \cup \{\sigma_b\}$
$\quad$ **return** σ_b
else
$\quad$ **return** $\bot$

Oracle $\mathsf{OPEN}(m, \sigma) \triangleq$
if $(\sigma \in Q)$ **then return** $\bot$
return $\mathsf{Open}(gsk, m, \sigma)$

The anonymity under tracing advantage of $\mathcal{A}$*,* $\mathbf{Adv}^{\mathsf{AuT}}(\mathcal{A})$ *is defined as* $2 \cdot Pr[\mathsf{AuT} : \mathsf{true}] - 1$*.* ATS *is anonymous under tracing if for any polynomial time* $\mathcal{A}$ *this advantage is negligible.*

Anonymity under tracing corresponds to the standard anonymity property of group signatures which guarantees anonymity toward everyone except the group manager. It ensures that even when being traced users are anonymous to the general public.

Definition 5 (Traceability). *Let* ATS *be an accountable forward-tracing signature scheme. Consider the following game played by adversary* $\mathcal{A}$*:*

Game Trace $\triangleq$
$gp \leftarrow \mathsf{Setup}(1^\lambda)$
$(gpk, gsk) \leftarrow \mathsf{GKg}(gp)$
$(m, \sigma) \leftarrow \mathcal{A}^{\mathsf{UKG,ENROLL,SIGN,OPEN}}(gpk)$
$(upk, \psi) \leftarrow \mathsf{Open}(gsk, m, \sigma)$
return ($\mathsf{Verify}(gpk, m, \sigma) = 1 \wedge$
$(m, \sigma) \notin Q \wedge$
$(\mathsf{Judge}(gpk, m, \sigma, upk, \psi) = 0 \vee$
$upk = \perp$))

Oracle OPEN$(m, \sigma) \triangleq$
$(upk, \psi) \leftarrow \mathsf{Open}(gsk, m, \sigma)$
return (upk, ψ)

Oracle UKG$(upk, epoch) \triangleq$
$(upk, usk) \leftarrow \mathsf{UKg}(gp)$
$S[upk] \leftarrow usk$
return upk

Oracle ENROLL$(upk, epoch, t) \triangleq$
$(cert, w^{escrw}) \leftarrow \mathsf{Enroll}(gsk, upk, epoch, (t \vee upk \notin \mathsf{dom}(S)))$
return $cert$

Oracle SIGN$(cert, m) \triangleq$
$usk = S[cert.upk]$
if $(usk = \perp)$ **then return** $\perp$
$\sigma \leftarrow \mathsf{Sign}(gpk, cert, usk, m)$
$Q \leftarrow Q \cup \{(m, \sigma)\}$
return σ

The traceability advantage of $\mathcal{A}$, $\mathbf{Adv}^{\mathsf{Trace}}(\mathcal{A})$ is defined as $Pr[\mathsf{Trace} : \mathsf{true}]$. ATS is traceable if for any polynomial time $\mathcal{A}$ this advantage is negligible.

Informally, traceability requires that every valid message will trace to someone as long as the adversary does not have a certificate and private key for a non-tracing certificate.

Although the attacker is free to choose the type of certificate for honest users whose keys are generated by the UKG oracle, he is not allowed to get untraceable certificates on user keys of his choosing. In the definition this is ensured by $(t \vee upk \notin \mathsf{dom}(S))$ always evaluating to 1 on such occasions. This is a slight departure from the standard tracing game [BMW03], where ENROLL always produces traceable certificates.

Definition 6 (Non-frameability). *Let* ATS *be an accountable forward-tracing signature scheme. Consider the following game played by adversary* $\mathcal{A}$*:*

Game NF $\triangleq$
$gp \leftarrow \mathsf{Setup}(1^\lambda)$
$(gpk, st) \leftarrow \mathcal{A}_0(gp)$
if $gpk.gp \neq gp$ **then**
   return $\perp$
$(m, \sigma, upk, \psi) \leftarrow \mathcal{A}_1^{\mathsf{UKG,SIGN}}(st)$
return $(\mathsf{Judge}(gpk, m, \sigma, upk, \psi) = 1$
$\wedge (m, \sigma) \notin Q \wedge upk \in S$

Oracle UKG$(upk, epoch) \triangleq$
$(upk, usk) \leftarrow \mathsf{UKg}(gp)$
$S[upk] \leftarrow usk$
return upk

Oracle SIGN$(cert, m) \triangleq$
$usk = S[cert.upk]$
if $(usk = \perp)$ **then return** $\perp$
$\sigma \leftarrow \mathsf{Sign}(gpk, cert, usk, m)$
$Q \leftarrow Q \cup \{(m, \sigma)\}$
return σ

The non-frameability advantage of $\mathcal{A}$, $\mathbf{Adv}^{\mathsf{NF}}(\mathcal{A})$ is defined as $Pr[\mathsf{NF} : \mathsf{true}]$. ATS is non-frameable if for any polynomial time $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ this advantage is negligible.

Non-frameability requires that no one, not even the group manager, can sign messages on the user's behalf. At the same time, it guarantees that users, if they are being traced, have to take responsibility for the messages they sign, i.e., traced signatures ensure non-repudiation.

Definition 7 (Anonymity with accountability). *Let* ATS *be an accountable forward-tracing signature scheme. Consider the following game played by an adversary $\mathcal{A}$:*

Game AwA $\triangleq$
$b \leftarrow \{0,1\}$
$gp \leftarrow \mathsf{Setup}(1^\lambda)$
$(gpk, st) \leftarrow \mathcal{A}_0(gp)$
if $gpk.gp \neq gp$
  return $\bot$
$b' \leftarrow \mathcal{A}_1^{\mathsf{Ch}}(st)$
return $(b' = b)$

Oracle $\mathsf{Ch}(sk_0, sk_1, cert_0, cert_1, m, w_0^{escrw}, w_1^{escrw})) \triangleq$
$\sigma_0 = \mathsf{Sign}(gpk, sk_0, cert_0, m)$
$\sigma_1 = \mathsf{Sign}(gpk, sk_1, cert_1, m)$
if $(\sigma_0 \neq \bot \wedge \sigma_1 \neq \bot \wedge cert_0.epoch = cert_1.epoch \wedge$
  $\mathsf{Account}(gpk, cert_0, w_0^{escrw}, 0) \wedge \mathsf{Account}(gpk, cert_1, w_1^{escrw}, 0))$
  return σ_b
else
  return $\bot$

The anonymity advantage of $\mathcal{A}$, $\mathbf{Adv}^{\mathsf{AwA}}(\mathcal{A})$ is defined as $2 \cdot Pr[\mathsf{AwA} : \mathsf{true}] - 1$. ATS *is anonymous if for any polynomial time* $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ *this advantage is negligible.*

Intuitively, this captures the notion that the user is anonymous even from a corrupt authority that has full control of the system. It requires that even if every single parameter in the system is adversarially controlled, a user is anonymous as long as the escrow key in his certificate is accountably private.

Definition 8 (Trace-obliviousness). *Let* ATS *be an accountable forward-tracing signature scheme. Consider the following game played by an adversary $\mathcal{A}$:*

Game TO $\triangleq$
$gp \leftarrow \mathsf{Setup}(1^\lambda)$
$(gpk, gsk) \leftarrow \mathsf{GKg}(gp)$
$b \leftarrow \{0,1\}$
$(b') \leftarrow \mathcal{A}^{\mathsf{Ch,ENROLL,OPEN}}(gpk)$
return $(b' = b)$

Oracle $\mathsf{Ch}(upk, epoch) \triangleq$
$(cert, w^{escrw}) \leftarrow \mathsf{Enroll}(gsk, upk, epoch, b)$
$U = U \cup \{upk\}$
return $cert$

Oracle $\mathsf{OPEN}(m, \sigma) \triangleq$
$(upk, \psi) \leftarrow \mathsf{Open}(gsk, m, \sigma)$
if $upk \in U$ **then**
  return $\bot$
else
  return (upk, ψ)

Oracle $\mathsf{ENROLL}(upk, epoch, t) \triangleq$
$(cert, w^{escrw}) \leftarrow \mathsf{Enroll}(gsk, upk, epoch, t)$
return $cert$

The trace-obliviousness advantage of $\mathcal{A}$, $\mathbf{Adv}^{\mathsf{TO}}(\mathcal{A})$ is defined as $2 \cdot Pr[\mathsf{TO} : \mathsf{true}] - 1$. ATS *is trace oblivious if for any polynomial time* $\mathcal{A}$ *this advantage is negligible.*

Trace-obliviousness requires that users cannot tell whether they are being traced.

Remarks. Note two things. First, if all keys are tracing keys, then (i) the trace game would be standard, (ii) the trace-oblivious game is moot, and (iii) there is no accountable anonymity. As a result, we would have a standard group signature scheme. Perhaps more interestingly, if we dropped the requirement for anonymity under tracing, we would not require simulation extractability for the proof system.

Note, moreover, that the anonymity under tracing adversary $\mathcal{A}$ provides all input to Ch. A simple hybrid argument thus shows that anonymity under tracing with a single challenge query implies security with multiple challenge queries. This means that even one-time simulation extraction is sufficient to prove security.

4 Accountable forward-tracing signature from key-oblivious encryption

Assume we have an unforgeable signature scheme $\mathsf{SIG}.(\mathsf{GroupGen}, \mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify})$, a one-time signature scheme $\mathsf{OTS}.(\mathsf{GroupGen}, \mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify})$, a key-oblivious encryption scheme

OE.(GroupGen, KeyGen, KeyRand, Enc, Dec), and a simulation-extractable non-interactive zero-knowledge proof system Π.(GroupGen, Setup, Prove, VfyProof, SimExtSetup, Sim, Ext). (For efficiency we require that SIG.GroupGen = OTS.GroupGen = OE.GroupGen = Π.GroupGen and refer to this algorithm as GroupGen.)

We construct an accountable tracing scheme ATS:

ATS.Setup(1^λ): Runs $\mathcal{G} \leftarrow$ GroupGen(1^λ), $(pk^{(0)}, sk^{(0)}) \leftarrow$ OE.KeyGen($\mathcal{G}$), $crs \leftarrow$ Π.Setup($\mathcal{G}$). The secret key $sk^{(0)}$ is discarded, in fact $pk^{(0)}$ should be generated in such a way that $sk^{(0)}$ is not known to any party. Outputs $gp = (\mathcal{G}, pk^{(0)}, crs)$.

ATS.GKg(gp): Runs $(cpk, csk) \leftarrow$ SIG.Keygen($gp.\mathcal{G}$) and $(pk^{(1)}, sk^{(1)}) \leftarrow$ OE.KeyGen($gp.\mathcal{G}$). Returns $gpk = (gp, cpk, opk = pk^{(1)})$ and $gsk = (gpk, csk, osk = sk^{(1)})$.

ATS.UKg(gp): Returns $(upk, usk) \leftarrow$ SIG.Keygen($gp.\mathcal{G}$).

ATS.Enroll($gsk, upk, epoch, b$): Given a user public key, computes $(epk; w^{escrw}) \leftarrow$ OE.KeyRand($pk^{(b)}$) and a signature $\sigma_{cert} \leftarrow$ SIG.Sign($csk, (upk, epk, epoch)$). Returns certificate $cert = ((upk, epk, epoch), \sigma_{cert})$ and witness w^{escrw}.

ATS.Sign($gpk, cert, usk, m$): Parses gpk as (gp, cpk, opk) and $cert$ as $((upk, epk, epoch), \sigma_{cert})$. Runs $(pk_{ots}, sk_{ots}) \leftarrow$ OTS.Keygen(gp), $\sigma_u \leftarrow$ SIG.Sign(usk, pk_{ots}), and $(escrw; s) \leftarrow$ OE.Enc($epk, (upk, \sigma_u)$). Computes a proof π using crs for the following relation to prove knowledge of $(upk, epk, \sigma_{cert}, \sigma_u)$:

$$((escrw, pk_{ots}, cpk, epoch), (upk, epk, \sigma_{cert}, \sigma_u, s)) \in R_{sig}$$
$$\text{iff } (\text{SIG.Verify}(cpk, \sigma_{cert}, (upk, epk, epoch)) = 1 \wedge$$
$$escrw = \text{OE.Enc}(epk, (upk, \sigma_u); s) \wedge$$
$$\text{SIG.Verify}(upk, \sigma_u, pk_{ots}) = 1) .$$

Computes $\sigma_{ots} \leftarrow$ OTS.Sign($sk_{ots}, (m, (escrw, pk_{ots}, cpk, epoch), \pi)$) and sets $\sigma = (\pi, \sigma_{ots}, pk_{ots}, escrw)$. If ATS.Verify($gpk, m, \sigma$) = 0, returns $\perp$, otherwise returns σ. This check is needed as we guarantee anonymity even for maliciously formed inputs, as long as the signature verifies and the user is not being traced.

ATS.Verify($gpk, m, \sigma, epoch$): Parses gpk as (gp, cpk, opk) and σ as $(\pi, \sigma_{ots}, pk_{ots}, escrw)$. First, verifies the one-time signature: OTS.Verify($pk_{ots}, \sigma_{ots}, (m, (escrw, pk_{ots}, cpk, epoch), \pi)$). Second, verifies the proof Π.VfyProof($crs, (escrw, pk_{ots}, cpk, epoch), \pi$). If all checks succeed, returns 1, otherwise returns 0.

ATS.Open($gpk, gsk, m, \sigma, epoch$): Parses gpk as (gp, cpk, opk) and σ as $(\pi, \sigma_{ots}, pk_{ots}, escrw)$. Runs ATS.Verify($gpk, m, \sigma, escrw$) and abort if it fails. Extracts osk from gsk and runs (upk, σ_u) = OE.Dec($osk, escrw$). Returns $(upk, \psi = \sigma_u)$.

ATS.Judge($gpk, m, \sigma, upk, \psi$): Parses gpk as (gp, cpk, opk), σ as $(\pi, \sigma_{ots}, pk_{ots}, escrw)$, and ψ as σ_u. Runs ATS.Verify($gpk, m, \sigma, epoch$) and then runs SIG.Verify(upk, σ_u, pk_{ots}). If both checks succeed, returns 1, otherwise 0.

ATS.Account($gpk, cert, w^{escrw}, b$): Returns 1, if $cert.epk =$ OE.KeyRand($pk^{(b)}; w^{escrw}$), otherwise 0.

Theorem 1. *If SIG is unforgeable, OTS is strongly unforgeable, OE is plaintext indistinguishable, key private and key randomizable, and Π is zero-knowledge and simulation-extractable then ATS is anonymous under tracing, traceable, non-frameable, accountably anonymous, and trace oblivious as defined in Section 3.2. See the proofs of Theorems 2-6 in Appendix A.1.*

5 Instantiation

We now detail how to instantiate our scheme.

EIKO. Intuitively key-oblivious encryption can be obtained from key-private homomorphic public-key encryption. For a given key pair, let $C(m; r)$ be the ciphertext of message m with randomness r. Consider the homomorphic property $C(m; r) \otimes C(m'; r') = C(m \cdot m'; r + r')$. We construct a key-oblivious encryption scheme by letting each oblivious public key corresponds to an encryption $C(1; r)$. We randomize a key by computing $C(1; r * r')$ for a random r' (this can be computed using $\otimes$ by square-and-multiply). Encryption corresponds to computing $C(m; r)$ by multiplying the public key $C(1; r)$ with the ciphertext $C(m; 0)$, i.e. the message encrypted with randomness zero—see also [GJJS04].

We will use such a simple construction of key-oblivious encryption based on ElGamal (EIKO):

EIKO.GroupGen(1^λ): Pick a group $\mathbb{G}$ of prime order q with generator g, and return $\mathcal{G} = (\mathbb{G}, q, g)$.
EIKO.KeyGen($\mathcal{G}$): Sample $r, \gamma \leftarrow \mathbb{Z}_q$ and return $pk = (g^r, (g^\gamma)^r)$, $sk = (pk, \gamma)$.
EIKO.KeyRand(pk): Let $pk = (A, B)$, sample $r \leftarrow \mathbb{Z}_q$, return $pk' = (A^r, B^r)$ and randomness r.
EIKO.Enc(pk, m): Let $pk = (A, B)$, sample $s \leftarrow \mathbb{Z}_q$, and return $ct = (A^s, B^s \cdot m)$.
EIKO.Dec(sk, ct): Let $sk = (pk, \gamma)$ and $ct = (C, D)$. Return $m = D/C^\gamma$.

Lemma 1 (Key privacy of EIKO). *If the decisional Diffie-Hellman problem holds in $\mathbb{G}$ then EIKO is a key-private key-oblivious encryption scheme.*

Lemma 2 (Key randomizability of EIKO). *If the decisional Diffie-Hellman problem holds in $\mathbb{G}$ then EIKO is a key-randomizable encryption scheme.*

Lemma 3 (Plaintext indistinguishability of EIKO).
If the decisional Diffie-Hellman problem holds in $\mathbb{G}$ then EIKO satisfies INDr.

SIG and OTS. To ensure efficiency, we need structure-preserving signatures (SPS) [AHO10]. SPS can either be proved secure in the generic group model [AGHO11], using q-type assumptions [AHO10], or using static assumptions [ACD$^+$12]. Usually the stronger the assumption, the better the performance. This gives us several options. Moreover, if we only use the signature scheme to sign freshly created one-time signature (OTS) public keys, we do not require full adaptive unforgeability, but it suffices if the signature scheme is secure for random messages. We can thus use the xSIG scheme from [ACD$^+$12] secure under DDH_2 and $xDLIN_1$ with $|\sigma_{sig}| = 6$. For the OTS scheme, which does not need to be structure preserving, on the other hand, we will always make use of the scheme of [HJ12] with $|pk_{ots}| = 2$ and $|\sigma_{ots}| = 2$. Note that the xSIG scheme requires random messages of a specific form, and thus requires extended pk_{ots} of size $2 \times 3 = 6$.

Groth-Sahai proofs and simulation extraction. We will use Groth-Sahai (GS) proofs to efficiently instantiate the simulation-extractable proof system Π for the relation R_{sig} used in the signing algorithm of our accountable forward-tracing signature scheme. GS proofs operate over bilinear groups and we thus pick suitable primitives for SIG, OTS, and OE. In our instantiation, we let GroupGen on input security parameter 1^λ output the bilinear group $\mathcal{G} := (q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_t, e, G, \hat{G})$ used by Π.Setup. We let of SIG, OTS, which will be pairing based, use the same groups. While OE will use $(\mathbb{G}_1, q, G)$ as its $(\mathbb{G}, q, g)$.

Recall the properties of bilinear groups:

- q is a λ-bit prime,
- $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_t$ are groups of prime order q with efficiently computable group operations, membership tests, and bilinear mapping $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_t$,
- G and $\hat{G}$ are random generators of $\mathbb{G}_1$, $\mathbb{G}_2$, and $e(G, \hat{G})$ generates $\mathbb{G}_t$, and
- $\forall A \in \mathbb{G}_1, \forall B \in \mathbb{G}_2, \forall x, y \in \mathbb{Z} : e(A^x, B^y) = e(A, B)^{xy}$.

- an equation of the form $\prod_i \prod_j e(A_i, B_j)^{\gamma_{ij}} = 1$ for constants $\gamma_{ij} \in \mathbb{Z}_p$, and constants or variables $A_i \in \mathbb{G}_1$, $B_j \in \mathbb{G}_2$ is called a pairing product equation (PPE for short).

These are exactly the type of equations that can be proved using the GS proof system. Pairing product equations that are linear, that is equations in which variables appear only on one side of the pairing are more efficient to prove.

We use the standard R-or-Sign technique to construct a simulation-extractable proof system Π from GS proofs, SIG and OTS. The left side of the 'or' proves the original statement (this is R_{sig} which we denote S_0), while the right side proves knowledge of a signature (σ_{sim}) under a public key in the crs (we denote this S_1). This signature certifies a one-time signature public key, which in turn is used to sign the instance and the Groth-Sahai proof. For details on the construction we refer to the construction of [CCS09]. For the performance analysis we fall back on [ADK+13]:

The number of group elements in a proof of SE-NIZK is counted as follows. Let $S_0 = (R_{sig}(x, w) = 1)$ and $S_1 = (\mathsf{Verify}_{sim}(pk_{sim}, \sigma_{sim}, pk_{ots\_sim}) = 1)$ be the statements of the left and right side respectively, both represented by pairing product equations. The proof size of the SE-NIZK is as follows:

$$\text{(size of } S_0) + \text{(size of switcher)} + \text{(size of } S_1) + \text{(size of } OTS) \tag{1}$$
$$= \text{(size of } S_0) \tag{1}$$
$$+ (|com| \times 1 + |\pi_{NL}| \times 1) \tag{2}$$
$$+ (|com| \times (|\sigma_{sim}| + S_1(C)) + |\pi_L| \times (S_1(L) + S_1(C)) + |\pi_{NL}| \times S_1(NL) \tag{3}$$
$$+ (|pk_{ots\_sim}| + |\sigma_{ots\_sim}|) \tag{4}$$

$|\pi_L|$ denotes the cost of proving a single linear PPE (e.g, $e(A, X) = 1$ where A is a constant),
$S_1(L)$ denotes the number of linear relations needed to prove $\mathsf{Verify}_{sim}(\ldots) = 1$
$|\pi_{NL}|$ denotes the cost of proving a non-linear PPE (e.g., $e(X, Y) = 1$)
$S_1(L)$ denotes the number of linear PPEs.
$S_1(C)$ denotes the number of constant pairings, e.g. $e(A, B)$, in PPEs.
pk_{ots_sim} is the size of the one-time signature key signed by the key in the crs (not part of a PPE)
σ_{ots_sim} is the size of the actual signature (again not part of a PPE).

For GS proofs over DLIN in asymmetric groups, we have $(|com|, |\pi_L|, |\pi_{NL}|) = (3, 3, 18)$; and for GS proofs over XDH, we have $(|com|, |\pi_L|, |\pi_{NL}|) = (2, 2, 8)$.

We consider the one-time simulation-extractable scheme SE-NIZK4 from [ADK+13] and an instantiation based on the xSIG scheme of [ACD+12] with unbounded simulation extraction. For these constructions, we summarize the overhead for achieving simulation extractability on top of simply proving the original statement, computed as the sum of equations (2)+(3)+(4), for both DLIN and XDH-based GS proofs in Table 1.

SE-NIZK overhead	simulatability	$|\sigma_{sim}|$	$S_1(C)$	$S_1(L)$	$S_1(NL)$	$|pk_{ots\_sim}|$	$|\sigma_{ots\_sim}|$	DLIN	XDH
SE-NIZK4 [ADK+13]	one-time	3	2	3	0	2	2	55	34
SE-NIZK+xSIG [ACD+12]	unbounded	6	2	1	1	2×3	2	80	48

Table 1. Overhead of GS proof.

Putting it all together. We now examine the actual cost of our ATS scheme. It is dominated by proving the statement for the core relation R_{sig} of our accountable tracing signature scheme (S_0 above).

The equation for the proof size for R_{sig} is given below and the results for various instantiations is summarized in Table 2.

$$\text{(size of } S_0) \tag{5}$$
$$= |com| \times (|upk| + |epk| + |\sigma_{cert}| + |\sigma_u| + 1 + S_0(C)) \tag{6}$$
$$+ |\pi_L| \times (S_0(L) + S_0(C)) + |\pi_{NL}| \times S_0(NL) \tag{7}$$

Where $S_0(\{L, C, NL\})$ are defined as in the same way as $S_1(\{L, C, NL\})$

To reduce the size of *epk* we use a key-oblivious randomness-reusing variant of EIKO similar to [Kur02] to allow us to efficiently encrypt multiple group elements. Effectively, each group element is encrypted under a distinct key using the same randomness. This results in a ciphertext overhead of only a single group elements and thus shorter oblivious keys. (The PPE for verifiable encryption is $\bigwedge_{i=1}^{n} e(D_i, \hat{G}) = e(B_i, \hat{G}^s)e(m_i, \hat{G}) \wedge e(C, \hat{G}) = e(A, \hat{G}^s)$). Note that keys now correspond to encryptions of vectors of 1.

For performance, we instantiate SIG with two different schemes. As discussed above, we can for instance use a general purpose structure-preserving signature scheme (SPS) with secret key *csk* for the group manager and use an XRMA-secure scheme with secret keys *usk* for users.

As a further optimization, instead of the complete *upk*, we encrypt a single unique group element that serves as an identifier of *upk*. Instead of adding an additional group element to *cert*, the group manager can reuse one of the random elements of *upk* by ensuring that it is unique. The user encrypts this element together with σ_u, so OE needs to encrypt only $|\sigma_u| + 1$ group elements, thus $|epk| = |\sigma_u| + 2$ because of the ciphertext overhead.

Relation R_{sig}	$|upk|$	$|epk|$	$|\sigma_{cert}|$	$|\sigma_u|$	$S_0(C)$	$S_0(L)$	$S_0(NL)$	DLIN	XDH
SPS [AGHO11]+SPS [AGHO11]	1	$2 + |\sigma_u|$	3	3	$8 + |\sigma_u|$	$4 + |\sigma_u|$	4	198	116
SPS [AGHO11]+xSIG [ACD$^+$12]	10	$2 + |\sigma_u|$	3	6	$6 + |\sigma_u|$	$3 + |\sigma_u|$	3	237	144
SIG2 [ACD$^+$12]+xSIG [ACD$^+$12]	10	$2 + |\sigma_u|$	14	6	$4 + |\sigma_u|$	$4 + |\sigma_u|$	4	279	170

Table 2. Cost of ATS relation R_{sig}

An ATS signature σ consists of the elements $(\pi, \sigma_{ots}, pk_{ots}, escrw)$. To reduce the overhead, we generate only a single pk_{ots} (and one-time signature σ_{ots}) and use it both in S_0 and S_1. In both cases σ_{ots}, respectively σ_{ots_sim}, signed the instance, the proof, and the message. They can thus be merged.

The overall signature size of an ATS signature is thus the sum of the number of group elements needed to prove R_{sig}, the simulation-extraction overhead, and the ciphertext size $|escrw| = |epk|$. In our example instantiations it ranges from 155 to 367 group elements.

Because of the availability of a large number of structure-preserving primitives, there is plenty of room for optimization, especially when one is aiming for both efficiency and weak cryptographic assumptions. We stress, moreover, that our instantiation does not make use of random oracles in its security proof.

6 Backward-tracing and message-escrow extensions

We only formally describe and analyze a base scheme and our approach can be extended in several directions to fit specific application requirements. We discuss two such extensions for applications that require backward-tracing and encryption respectively.

6.1 Accountable backward-tracing signatures

So far we have considered monitoring a suspect's future actions. In the case of recovering past actions, we cannot retroactively tag a message and must, instead, extract something from the user to identify her messages.

With some applications (e.g. cloud based email), where users may maintain an encrypted inbox/outbox of their messages merely (accountably) extracting the necessary decryption key is sufficient. We can decrypt the inbox rather than resort to trial message decryption. For other applications, it seems search costs are on the order of $m \times t$ where m is the number of messages in the system and t is the number of targeted users.

In either case, retrieving the user's key introduces a second issue: restoring privacy. For forward tracing, the authority merely needs to replace the escrow key with a randomization of $pk^{(0)}$ when a warrant expires. For backward tracing, things are more complicated as we need the user to replace her key with a new one without realizing she did so. This requires more than just key obliviousness: the user must only hold a share of her private key. If not, she can simply test if she can decrypt messages encrypted under the new key.

We augment EIKO with a basic distributed key generation functionality to form DEIKO. We model distributed key generation using the following algorithms: $\mathsf{ShareGen}(\mathcal{G})$ generates public and private key shares (ps, ss), while $\mathsf{CombinePS}$ and $\mathsf{CombineSS}$ combine a vector of public and private key shares into a public and a private key respectively. We extend EIKO with algorithms $\mathsf{ShareGen}$, $\mathsf{CombinePS}$ and $\mathsf{CombineSS}$ for generating keys from public and private shares:

$\mathsf{DEIKO.ShareGen}(\mathcal{G})$: $\alpha \leftarrow \mathbb{Z}_q$, $ps = g^{\alpha}$, $ss \leftarrow \alpha$. Returns (ps, ss).
$\mathsf{DEIKO.CombinePS}(\boldsymbol{ps})$: $r \leftarrow \mathbb{Z}_q$, Returns $pk = (g^r, (\prod ps_i)^r)$.
$\mathsf{DEIKO.CombineSS}(\boldsymbol{ss})$: Returns $sk = \sum ss_i$.

Adding backward tracing. An Accountable backward-tracing signature scheme is a set of 8 algorithms: $(\mathsf{Setup}, \mathsf{GKg}, \mathsf{UKg}, \mathsf{Sign}, \mathsf{Verify}, \mathsf{Open}, \mathsf{Judge}, \mathsf{Account})$ and one protocol $\mathsf{Subscribe}(gpk, usk) \leftrightarrow \mathsf{Enroll}(gsk, upk, b)$.

Our construction is based on our forward-tracing scheme from Section 4 with two modifications (i) epk instead of being a key for EIKO is now for DEIKO and (ii) we replace Enroll with an interactive protocol $(\mathsf{Subscribe}, \mathsf{Enroll})$, that handles key generation, key retrieval in the case of a warrant, and key renewal on warrant expiration.

$\mathsf{Subscribe}$, detailed in Figure 1, uses a blind decryption scheme [Gre11] with algorithms $\mathsf{BE.}(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Blind}, \mathsf{BlindDec}, \mathsf{UnBlind})$. Intuitively, in $\mathsf{Subscribe}$, the user provides the authority with an encryption of her share of the escrow key. The authority can gain oblivious access to this share via a blind decryption query. To maintain trace-obliviousness, the authority normally issues blind decryption queries on an encryption of 0. Again, revealing the randomness—in this case for blinding the ciphertext—renders this accountable.

The process for key renewal is best understood via Figure 1. It leverages the fact that a user, since she knows only shares of escrow keys, cannot tell at the end of subscription whether she holds an escrow key generated from an old share, a new share, or a randomized key shared with all traced users.

6.2 Extending to encryption and message escrow

Both signature schemes can be readily adapted to form an escrowed encryption scheme by having the message be a ciphertext and including in $escrw$ a copy of the plaintext, and modifying the proof in the signature accordingly. Formally, such as scheme has 10 algorithms: $(\mathsf{Setup}, \mathsf{GKg},$

Subscribe(gpk, usk) **Enroll**$(gsk, upk, epoch, b)$

$\mathsf{read}(upk, (epk_{old}, gps_{old}, gss_{old}, E_{old},$
$\qquad\qquad b_{old}, witnesses, keys))$
$starttracing = b \wedge \neg b_{old};$
$stoptracing = b_{old} \wedge \neg b$
$w_0 \leftarrow \$$
if $(starttracing)$ **then**
$\quad B = \mathsf{BE.Blind}(E_{old}, w_0)$
else
$\quad B = \mathsf{BE.Blind}(E_0, w_0)$

$\xleftarrow{\quad B \quad}$

$(ps, ss) \leftarrow \mathsf{OE.ShareGen}(gp)$
$E \leftarrow \mathsf{BE.Enc}(bpk, ss)$
$D = \mathsf{BE.Dec}(bsk, B)$
π is a proof of correct decryption

$\xrightarrow{\quad ps, E, D, \pi \quad}$

$w_1 \leftarrow \$$
verify the proof of correct decryption
if $(starttracing)$
$\quad ss_{old} = \mathsf{BE.Unblind}(D, w_0)$
$\quad k = \mathsf{OE.CombineSS}(gss_{old}, ss_{old})$
$\quad epk = \mathsf{OE.KeyRand}(pk^{(1)}; w_1)$
$\quad \mathsf{store}(upk, (epk, \bot, \bot, \bot, 1, witnesses, keys :: k))$
else if $(stoptracing)$ **then**
$\quad (gps, gss) \leftarrow \mathsf{OE.ShareGen}(gp)$
$\quad epk = \mathsf{OE.CombinePS}(gps, ps; w_1)$
$\quad \mathsf{store}(upk, (epk, gps, gss, E,$
$\qquad\qquad 0, witnesses :: (w_0, w_1), keys))$
else
$\quad epk = \mathsf{OE.KeyRand}(epk_{old}; w_1)$
$\quad \mathsf{store}(upk, (epk, gps_{old}, gss_{old}, E_{old},$
$\qquad\qquad b, witnesses :: (w_0, w_1), keys)$
$sig \leftarrow \mathsf{Sign}(gsk, (epk, upk, epoch))$
$cert = (upk, epk, epoch, sig)$

$\xleftarrow{\quad cert \quad}$

Fig. 1. The $\mathsf{Subscribe} \leftrightarrow \mathsf{Enroll}$ protocol

$\mathsf{UKg}, \mathsf{Enroll}, \mathsf{Signcrypt}, \mathsf{Decrypt}, \mathsf{Verify}, \mathsf{Open}, \mathsf{Judge}, \mathsf{Account})$. Sign is replaced by $\mathsf{Signcrypt}$ and augmented to additionally take a public key under which the message is encrypted. These keys can either be from an external source or produced by UKg. The resulting "signature" can be decrypted by $\mathsf{Decrypt}$ only with the corresponding private key. Verify can still be run by anyone of course.

Concretely, $\mathsf{ATS.Sign}$ becomes

$(\sigma, escrw) \leftarrow \mathsf{Signcrypt}(gpk, cert, usk, m, pk)$**:**
    Parse gpk as (cpk, crs, opk) and $cert$ as $(upk, epk, epoch)$. Compute $(ct; s_{pk}) \leftarrow \mathsf{PKEnc}(pk, m)$.
    Run $(pk_{ots}, sk_{ots}) \leftarrow \mathsf{OTS.Keygen}(gp)$, $\sigma_u \leftarrow \mathsf{SIG.Sign}(usk, pk_{ots})$, and $(escrw, s_{oe}) \leftarrow \mathsf{OE.Enc}(epk,$
    $(upk, \sigma_u, m))$. Compute a proof π for the relation

$$((escrw, pk_{ots}, cpk, epoch, ct), (upk, \sigma_{cert}, \sigma_u, m, s_{oe}, s_{pk})) \in R_{sig}$$
$$\text{iff } (\mathsf{SIG.Verify}(cpk, \sigma_{cert}, (upk, epk, epoch)) = 1 \wedge$$
$$escrw = \mathsf{OE.Enc}(epk, (upk, \sigma_u, m); s_{oe}) \wedge$$
$$ct = \mathsf{PKEnc}(pk, m; s_{pk}) \wedge$$
$$\mathsf{SIG.Verify}(upk, \sigma_u, pk_{ots}) = 1) \,.$$

Run $\sigma_{ots} \leftarrow$ OTS.Sign$(sk_{ots}, (escrw, pk_{ots}, cpk, epoch, ct, \pi))$, set $\sigma = (\pi, \sigma_{ots}, pk_{ots}, escrw, ct)$ and if ATS.Verify$(gpk, m, \sigma) = 0$, return $\perp$ else, return σ).

7 Transparency reports and conclusions

Accountable tracing signatures hold those who demand "lawful" access to encrypted messages accountable for what they access. With it, under some circumstances at least, demands for lawful access to cryptographic systems can be dealt with without allowing mass surveillance of message content and, crucially, metadata.

For most ordinary criminal cases, the existence of a search warrant is already revealed when data obtained from it is used in court. Thus the requirement to reveal searches after the fact is innocuous. However, many of the more troubling issues stem from orders which demand both access and silence. Currently, many companies issue transparency reports purportedly giving statistical data about the volume of such requests. These, again, are not accountable. Using an ATS scheme or its message escrow variant, however, these transparency reports become provable. If every epk is stored in a public ledger, then the authority can easily issue proofs attesting that less than some fraction of its transactions use tracing keys for warrants accompanied by gag orders.

Accountable tracing signatures do not, of course, preclude the use of backdoors, software vulnerabilities, or other non-cryptographic attack vectors. However, given that they provide a vector for lawful access (and arguably bounded "unlawful" access to whatever extent the transparency report allows), they eliminate part of the motivation. Moreover, by potentially eliminating the lawful access objection to strong cryptography and allowing its deployment, they make mass surveillance far more difficult.

While accountable tracing signatures hold the authority accountable, they obviously only do so after the fact. An attacker who compromises the authority's keys gets access to all data until they are detected. Thus, if the goal is to maximize security, such systems should be avoided unless the alternatives are a non-accountable escrow system or no cryptographic protections at all.

References

[ACD⁺12] Masayuki Abe, Melissa Chase, Bernardo David, Markulf Kohlweiss, Ryo Nishimaki, and Miyako Ohkubo. Constant-size structure-preserving signatures: Generic constructions and simple assumptions. In *Advances in Cryptology - ASIACRYPT 2012 - 18th International Conference on the Theory and Application of Cryptology and Information Security, Beijing, China, December 2-6, 2012. Proceedings*, pages 4–24, 2012.

[ADK⁺13] Masayuki Abe, Bernardo David, Markulf Kohlweiss, Ryo Nishimaki, and Miyako Ohkubo. Tagged one-time signatures: Tight security and optimal tag size. In *Public-Key Cryptography - PKC 2013 - 16th International Conference on Practice and Theory in Public-Key Cryptography, Nara, Japan, February 26 - March 1, 2013. Proceedings*, pages 312–331, 2013.

[AGHO11] Masayuki Abe, Jens Groth, Kristiyan Haralambiev, and Miyako Ohkubo. Optimal structure-preserving signatures in asymmetric bilinear groups. In *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, pages 649–666, 2011.

[AHO10] Masayuki Abe, Kristiyan Haralambiev, and Miyako Ohkubo. Signing on elements in bilinear groups for modular protocol design. Cryptology ePrint Archive, Report 2010/133, 2010.

[BBDP01] Mihir Bellare, Alexandra Boldyreva, Anand Desai, and David Pointcheval. Key-privacy in public-key encryption. In *Advances in Cryptology–ASIACRYPT 2001*, pages 566–582. Springer, 2001.

[BMW03] Mihir Bellare, Daniele Micciancio, and Bogdan Warinschi. Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions. In *Advances in Cryptology–Eurocrypt 2003*, pages 614–629. Springer, 2003.

[CCS09] Jan Camenisch, Nishanth Chandran, and Victor Shoup. A public key encryption scheme secure against key dependent chosen plaintext and adaptive chosen ciphertext attacks. In *Advances in Cryptology - EUROCRYPT 2009*, volume 5479, pages 351–368, 2009.

[CvH91] David Chaum and Eugène van Heyst. Group signatures. In *EUROCRYPT*, volume 547 of *Lecture Notes in Computer Science*, pages 257–265, 1991.

[FPV09]   Georg Fuchsbauer, David Pointcheval, and Damien Vergnaud. Transferable constant-size fair e-cash. In Juan A. Garay, Atsuko Miyaji, and Akira Otsuka, editors, *CANS*, volume 5888 of *Lecture Notes in Computer Science*, pages 226–247. Springer, 2009.

[FV10]   Georg Fuchsbauer and Damien Vergnaud. Fair blind signatures without random oracles. In *AFRICACRYPT*, volume 6055 of *Lecture Notes in Computer Science*, pages 16–33, 2010.

[GJJS04] Philippe Golle, Markus Jakobsson, Ari Juels, and Paul F. Syverson. Universal re-encryption for mixnets. In *CT-RSA*, volume 2964 of *Lecture Notes in Computer Science*, pages 163–178, 2004.

[Gre11]   Matthew Green. Secure blind decryption. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *Public Key Cryptography*, volume 6571 of *Lecture Notes in Computer Science*, pages 265–282. Springer, 2011.

[HJ12]   Dennis Hofheinz and Tibor Jager. Tightly secure signatures and public-key encryption. In *CRYPTO*. Springer, 2012.

[KTY04]   Aggelos Kiayias, Yiannis Tsiounis, and Moti Yung. Traceable signatures. In *Advances in Cryptology-EUROCRYPT 2004*, pages 571–589. Springer, 2004.

[Kur02]   Kaoru Kurosawa. Multi-recipient public-key encryption with shortened ciphertext. In *Public Key Cryptography*, pages 48–63. Springer, 2002.

[KV01]   Dennis Kügler and Holger Vogt. Auditable tracing with unconditional anonymity. 2001.

[KV03]   Dennis Kügler and Holger Vogt. Offline payments with auditable tracing. In *Financial Cryptography*, pages 269–281. Springer, 2003.

[LRC14]   Jia Liu, Mark D Ryan, and Liqun Chen. Balancing societal security and individual privacy: Accountable escrow system. In *27th IEEE Computer Security Foundations Symposium (CSF)*, 2014.

A   Proofs

A.1   Security for ATS.

Theorem 2 (Anonymity under tracing). *If* Π *is zero-knowledge and simulation-extractable and* OE *is plaintext indistinguishable and key private, then the construction described in Section 4 is anonymous under tracing as defined in Section 3.2.*

Proof. We now proceed to describe a sequence of hybrid experiments.

– ∂_0. The original AuT game.
– ∂_1. Same as ∂_0, except that in the Ch oracle we use Sim to simulate π_b in σ_b and store the simulated proof along with its inputs in a log LS. By the zero-knowledge property of the proof system, the attacker has a negligible advantage in distinguishing between this and the previous game.
– ∂_2. Same as ∂_1, except that for proofs $\pi \in LS$ we answer with the stored data from the simulator. For $\pi \notin LS$ we use extraction to answer OPEN oracle queries without using the decryption key in gsk. If this fails we abort. Because Π is simulation extractable, by definition the probability of failing to extract on a proof that was not directly produced by the simulator (i.e. is not in LS) is negligible, and hence so is the probability of abort and the attacker's advantage in distinguishing between this and the previous game.
– ∂_3 and ∂_4 correspond to ∂_2 and ∂_3 of the anonymity with accountability proof.

In ∂_4, all inputs to $\mathcal{A}$ are independent of b and its advantage is therefore zero.

Theorem 3 (Traceability). *If* Π *is extractable,* SIG *is unforgeable, and* OTS *is strongly unforgeable, then the construction described in Section 4 is traceable as defined in Section 3.2.*

Proof. We proceed through a sequence of hybrids.

– ∂_0. The original traceability game.
– ∂_1. The same as ∂_0, except that we attempt to extract $(upk, epk, \sigma_{cert}, \sigma_u)$ from the proof π in σ and abort without $\mathcal{A}$ winning if extraction fails. The difference in the advantage of $\mathcal{A}$ winning compared with ∂_0 is negligible by the extractability of the proof system.

- $\beth_2$. Same as $\beth_1$, except that we abort if the attacker uses a signature σ_{cert} on a fresh upk, epk.

    Lemma 4. *The difference in the advantage of $\mathcal{A}$ winning compared with $\beth_1$ is negligible by the unforgeability of* SIG.

- $\beth_3$. Same as $\beth_2$, except that we abort if the attacker uses a signature σ_u on a fresh OTS public key pk_{ots}, i.e., one that did not result from a signing query.

    Lemma 5. *The difference in the advantage of $\mathcal{A}$ winning compared with $\beth_2$ is negligible by the unforgeability of* SIG.

- $\beth_4$. Same as $\beth_3$, except that we abort if the attacker reuses a OTS public key to sign a different message (i.e not the one stored in Q). The probability of this happening, and hence aborting, is negligible by a slight variant of Lemma 4 and 5 for one-time signatures.

In $\beth_4$ the advantage of $\mathcal{A}$ is zero.

Theorem 4 (Non-frameability). *Let* SIG *be an unforgeable signature scheme and* OTS *a strongly unforgeable one-time signature scheme, then the construction described in Section 4 is non-frameable as defined in Section 3.2.*

Proof. Informally causing Judge to validate is the same as forging one of the two signature in ψ. This is assumed impossible for secure signature schemes.

The proof is thus nearly identical to that of traceability, except that $\mathcal{A}$ already provides σ_u as part of his output and it is thus does not need to extracted. For the hybrids see $\beth_3$ and $\beth_4$ of the traceability proof.

Theorem 5 (Anonymity with accountability). *If* Π *is zero-knowledge and* OE *is plaintext indistinguishable and key private, then the construction described in Section 4 is accountably anonymous as defined in Section 3.2.*

Proof. We proceed through a sequence of hybrids.

- $\beth_0$. The original AwA game.
- $\beth_1$. Same as $\beth_0$, except that we replace the zero-knowledge proofs by simulated proofs. By the zero-knowledge property of the proof system, the attacker has a negligible advantage in distinguishing between this and the previous game.
- $\beth_2$. Same as $\beth_1$. except we modify ATS.Sign in game Ch to produce an escrow ciphertext of the encryption of 0 for σ_b.

    Lemma 6. *By the INDr-CPA property of* EIKO*, the new ciphertext is indistinguishable from the old one and so $\beth_1$ and $\beth_2$ are indistinguishable.*

- $\beth_3$. Same as $\beth_2$. except we modify ATS.Sign in game Ch to use a fresh random key as the escrow key.

    Lemma 7. *By the key privacy property of* EIKO*, the new ciphertext is indistinguishable from the old one and so $\beth_3$ and $\beth_4$ are indistinguishable.*

In $\beth_3$, Ch returns a simulated proof and an encryption of zero under a new key. Hence all inputs to $\mathcal{A}$ are independent of b and its advantage is therefore zero.

Theorem 6 (Trace-obliviousness). *If* OE *is key randomizable, and* Π *is extractable, then the construction described in Section 4 is trace-oblivious as defined in Section 3.2.*

Proof. We proceed through a sequence of hybrids.

- ∂_0. The original trace-obliviousness game.
- ∂_1. Same as ∂_0, except that we change OPEN to invoke the extractor for Π to decrypt *escrw* instead of using the private key. The attacker has a negligible advantage in distinguishing between this game and ∂_0 by the extractability of the proof system.
- ∂_2. Same as ∂_1, except that we change Enroll in ENROLL to return freshly generated public keys instead of randomized keys.

  Lemma 8. *By the key randomizability property of* EIKO*, the new key is indistinguishable from the old one.*

  In ∂_2 the output of all oracles is independent of b, and therefor the adversary's advantage is zero.

A.2   Proofs of supporting lemmas

Proof (Lemma 1). Given an attacker $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ who breaks KPr we construct a distinguisher $\mathcal{D}$ for DDH.

$$
\begin{aligned}
&\textbf{Distinguisher } \mathcal{D}(\mathbb{G}, X, Y, T) \triangleq \\
&b \leftarrow \{0,1\} \\
&\lambda, \mu, \xi, r_0, r_1, \leftarrow \mathbb{Z}_q \\
&X_0 \leftarrow X; Y_0 \leftarrow Y; T_0 \leftarrow T \\
&X_1 \leftarrow X \cdot g^\lambda; Y_1 \leftarrow Y^\xi \cdot g^\mu; T_1 \leftarrow T^\xi \cdot X^\mu \cdot Y^{\lambda\xi} \cdot g^{\lambda\xi} \\
&pk_0 \leftarrow (g^{r_0}, Y_0^{r_0}) \\
&pk_0 \leftarrow (g^{r_1}, Y_1^{r_1}) \\
&(m, pk_0', w_0, pk_1', w_1, st) \leftarrow \mathcal{A}_0(\mathcal{G}, pk_0, pk_1) \\
&\textbf{if } pk_0' \neq \mathsf{KeyRand}(pk_0; w_0) \vee pk_1' \neq \mathsf{KeyRand}(pk_1; w_1) \\
&\qquad \textbf{return } \bot \\
&d \leftarrow \mathcal{A}_1(X^{r_b w_b}, T_b^{r_b w_b} \cdot M, st) \\
&\qquad \textbf{return } (d = b)
\end{aligned}
$$

Given Diffie-Hellman's random self-reducibility, if X, Y, T is a Diffie-Hellman triple, then so is X_1, Y_1, T_1. Moreover, both triples are distributed identically regardless and produce the proper distributions of keys in EIKO. In the case that T is a random group element, then the challenge ciphertext given to A contains no information, as Y_b and $T_b^{w_b}$ are identically distributed regardless of b. Hence, $\mathcal{A}$'s advantage at guessing the bit is negligible. On the other hand, in the case where the challenge is a valid Diffie-Hellman triple, $\mathcal{A}$'s inputs are the same as in **Game** KO, since $T_b^{r_b w_b} = (g^{x_b \cdot y_b})^{r_b \cdot w_b} = g^{x_b \cdot y_b \cdot r_b \cdot w_b} = (X^{r_b \cdot w_b})^{y_b} = C^{y_b}$.

Thus if $\mathcal{A}$ has a non-negligible advantage at breaking **Game** KO, then $\mathcal{D}$ breaks DDH.

Proof (Lemma 2). Given a DDH challenge $(\mathbb{G}, X, Y, T)$, the reduction is immediate.

$$
\begin{aligned}
&\textbf{Distinguisher } \mathcal{D}(\mathbb{G}, X, Y, T) \triangleq \\
&r, \leftarrow \mathbb{Z}_q \\
&d \leftarrow \mathcal{A}((g^r, X^r), (Y^r, T^r)) \\
&\qquad \textbf{return } d
\end{aligned}
$$

In the case where $(X = g^x, Y = g^y, T = g^{xy})$, then our original key is (g^r, g^{xr}), and the second one is $((g^r)^y, (g^{xr})^y))$, i.e. the original key re-randomized by y. On the other hand, where $(X = g^x, Y = g^y, T = g^z)$, the second key is unrelated. Thus if $\mathcal{A}$ distinguishes between a real or random key with a non-negligible advantage, then the distinguisher above breaks DDH with the same advantage.

Proof (Lemma 3). Given an attacker $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ who wins against INDr we construct a DDH distinguisher $\mathcal{D}$.

$$
\begin{aligned}
&\textbf{Distinguisher } \mathcal{D}(\mathbb{G}, X, Y, T) \triangleq \\
&\quad b \leftarrow \{0, 1\} \\
&\quad r \leftarrow \mathbb{Z}_q \\
&\quad pk \leftarrow (g^r, X^r) \\
&\quad pk', r', m_0, m_1 \leftarrow \mathcal{A}_0(\mathcal{G}, pk) \\
&\quad \textbf{if } pk' \neq \mathsf{KeyRand}(pk; r') \textbf{ then} \\
&\quad\quad \textbf{return } \bot \\
&\quad ct \leftarrow ((Y^r)^{r'}, (T^r)^{r'} \cdot m_b) \\
&\quad b' \leftarrow \mathcal{A}_1(ct) \\
&\quad \textbf{return } (b = b')
\end{aligned}
$$

In the case where (X, Y, T) is a DDH triple, $\mathcal{A}$'s view is identical to INDr game: she receives a public key g^{rx} and a base g^r and then a properly formed ciphertext $\left(((g^r)^{r'})^y, ((g^r)^{r'})^{xy} \cdot m\right)$. On the other hand, when T is a random group element, $(T^r)^{r'} \cdot m_b$ reveals no information about m_b. Hence her advantage is negligible. Thus if A has an advantage in winning INDr, then the above breaks DDH with the same advantage.

Proof (Lemma 4). Given an attacker $\mathcal{A}_{\supset}$ who forgers a signature in Trace, we construct a reduction to the standard EU-CMA signature game where an attacker A is given access to a signing oracle and a public key and produces a forgery on a message not previously signed by the oracle.

$\;\;\;\mathcal{A}^{\mathsf{SIG}}(pk)$ works as follows. Given pk it simulates the standard Trace game using the verification key it received as input as the authorities key in the parameters. We modify SIGN to return signatures generated by the oracle SIG oracle. When $\mathcal{A}_{\supset}$ triggers abort by producing a signature on a key not in S, $\mathcal{A}$ return it as a signature forgery.

Proof (Lemma 5). This proof is nearly identical to that of Lemma 4. Instead of inserting the challenge public key into the parameters, however, we must have ENROLL embed this key in some generated user key. Unfortunately, we can only do so for one query and must simply blindly guess upon which one to do so. Having done so, the game continues as in the proof above and, if we guessed correctly, we get the appropriate forgery. Because $\mathcal{A}_{\supset}$ makes at most $poly()$ queries to ENROLL, there is a $\frac{1}{poly()}$ chance $\mathcal{A}_{\supset}$ forges a signature under the target key (i.e. that the forgery resulted from the query we picked), thus $\mathcal{A}_{\supset}$'s probability of abort is negligible.

Proof (Lemma 6). $\mathcal{A}_{\supset}$ makes at most $poly()$ queries to the Ch oracle. We construct a series of hybrids—$\supset_2^0, \ldots, \supset_2^i, \ldots, \supset_2^{poly()}$ where $\supset_2^{poly()}$ is equivalent to $\supset_2$—replacing the ciphertext in each successive query with 0. Given an adversary $\mathcal{A}_{\supset}^i$ who detects the hybrid that modifies the ith query to Ch, we construct an adversary $\mathcal{A}_{IND} = (\mathcal{A}_0, \mathcal{A}_1)$ who breaks INDr as follows.

$\;\;\;\mathcal{A}_0$ runs the standard AwA game with $\mathcal{A}_{\supset}^i$ using the provided parameters as the (non)escrow key in gp. On the ith query to the Ch oracle, $\mathcal{A}_0$ returns $(cert_b.epk, w_b^{escrw}), ctm, 0)$ as (pk', r, m_0, m_1) where ctm is the correct content of an escrow ciphertext. Upon receiving the challenge ciphertext, $\mathcal{A}_1$ constructs σ using ct as the escrow ciphertext and allows the AwA game to continue. Finally, it returns the resulting bit.

$\;\;\;$In the case where m_0 is chosen in the $INDr$ game, $A_{\supset}^i$'s view is identical to that of $\supset_2^{i-1}$ (i.e. where the ith query is untampered with and results in a proper ciphertext). On the other hand, in the case where m_1 is chosen, her view is identical to the case of $\supset_2^i$ (i.e. where the ith ciphertext is an encryption of the all zero string). Thus $\mathcal{A}_{\supset}^i$'s advantage is the same as $\mathbf{Adv}^{\mathsf{INDr}}(\mathcal{A})$ which is negligible. Thus $\mathcal{A}_{\supset}$'s advantage for the whole set of substitutions is $\mathbf{Adv}^{\mathsf{INDr}}(\mathcal{A}) \cdot poly()$ which is still negligible.

Proof (Lemma 7). This proof proceeds similarly to lemma 6. Again, $\mathcal{A}_\supset$, this time distinguishing between ∂_2 and ∂_3, makes at most polynomially many queries to Ch. We construct a series of hybrids —$\partial_3^0, \ldots \partial_3^i, \ldots \partial_3^{poly()}$ where $\partial_3^{poly()}$ is equivalent to ∂_3—in which we swap the key in the ith query. Given an adversary $\mathcal{A}_\supset^i$ who detects the hybrid that modifies the ith query to Ch, we construct an adversary $\mathcal{A}_{KP} = (\mathcal{A}_0, \mathcal{A}_1)$ who breaks KPr as follows.

$\mathcal{A}_0(pk_0, pk_1)$ runs the standard AwA game with $\mathcal{A}_\supset$ using pk_0 as the key in gp. On the ith query to the Ch oracle, $\mathcal{A}_0$ samples fresh randomness r and returns $(0, cert_b.epk, w_0^{escrw}, \mathsf{KeyRand}(pk_1; r),$ $r, 0)$ as $(m, pk_0', w_0, pk_1', w_1, st)$. Upon receiving the challenge ciphertext, $\mathcal{A}_1$ constructs σ using ct as the escrow ciphertext and allows the AwA game to continue. Finally, it returns the resulting bit.

In the case where pk_0 is chosen as the encryption key in the KPr game, $A_\supset^i$'s view is identical to ∂_3^{i-1}, where the provided key is used. On the other hand, in the case where pk_2 is chosen, her view is identical to ∂_3^i where a fresh key is used. Thus her advantage in distinguishing between any two hybrids is $\mathbf{Adv}^{\mathsf{KPr}}(\mathcal{A})$ which is negligible. This implies $\mathcal{A}_\supset$'s advantage $\mathbf{Adv}^{\mathsf{KPr}}(\mathcal{A}) \cdot poly()$ which is still negligible.

Proof (Lemma 8). This proof proceeds similarly to the others above. We construct a series of hybrids, one for each query to ENROLL where we sequentially replace the returned key with a random one. We denote these $\partial_2^0, \ldots \partial_2^i, \ldots \partial_2^{poly()}$ where $\partial_2^{poly()}$ is equivalent to ∂_2. Given an adversary $\mathcal{A}_\supset^i$ who detects the hybrid that tampers with the ith query, We construct $\mathcal{A}_{KR}(pk, pk_b)$ as follows.

$\mathcal{A}_{KR}$ embeds pk as the escrow key in the parameters and runs $A_\supset^i$. On the query to ENROLL, it returns pk_b. The game continues as normal. Finally, it returns the resulting bit.

In the case of $\mathcal{A}_{KR}(pk, pk_0)$, $A_\supset^i$ receives a randomized key and her view is identical to that of ∂_2^{i-1}. On the other hand, in the case of $\mathcal{A}_{KR}(pk, pk_1)$, her view is identical to that in ∂_2^i. Thus her advantage is $\mathbf{Adv}^{\mathsf{KR}}(\mathcal{A})$ when detecting tampering with any one query. $\mathcal{A}_\supset$'s advantage in distinguishing between ∂_1 and ∂_2 is $\mathbf{Adv}^{\mathsf{KR}}(\mathcal{A}) \cdot poly()$ which is negligible.