Memory-Tight Multi-Challenge Security of Public-Key Encryption

Joseph Jaeger and Akshaya Kumar

School of Cybersecurity and Privacy Georgia Institute of Technology, Atlanta, Georgia, US {josephjaeger, akshayakumar}@gatech.edu

Abstract. We give the first examples of public-key encryption schemes which can be proven to achieve multi-challenge, multi-user CCA security via reductions that are tight in time, advantage, *and* memory. Our constructions are obtained by applying the KEM-DEM paradigm to variants of Hashed ElGamal and the Fujisaki-Okamoto transformation that are augmented by adding uniformly random strings to their ciphertexts. The reductions carefully combine recent proof techniques introduced by Bhattacharyya'20 and Ghoshal-Ghosal-Jaeger-Tessaro'22. Our proofs for the augmented ECIES version of Hashed-ElGamal make use of a new computational Diffie-Hellman assumption wherein the adversary is given access to a pairing to a random group, which we believe may be of independent interest.

Keywords: Public-key cryptography, provably security, memory-tightness

1 Introduction

Secure deployment of cryptography requires concrete analysis of schemes to understand how the success probabilities of attackers grow with the amount of resources they employ to attack a system. The use of reduction-based cryptography enables such analysis by using an attacker with running time t and success probability ϵ to construct a related adversary with running time t' and success probability ϵ' against a computational problem whose security is better understood. A gold standard for concrete security reductions are tight reductions for which $t' \approx t$ and $\epsilon' \approx \epsilon$. We refer to such a reduction as TA-tight (time-advantage-tight) to distinguish it from other notions of tightness.

Auerbach, Cash, Fersch, and Kiltz [3] argued that the memory usage of an attacker can be crucial in determining its likelihood of success. This kicked of a line of works [31, 30, 24, 12, 11, 7, 19, 18, 9, 28, 10, 17] on *memory-aware* cryptography which accounts for the memory usage of attackers in security analyses. Auerbach, et al. focused in particular on incorporating memory considerations into the study of reductions. We refer to a reduction as TAM-tight if it is TA-tight and additionally $s \approx s'$ where these variables, respectively, denote the amount of memory used by the original adversary and the reduction adversary. In this work, we construct the first public-key encryption schemes with TAMtight proofs of multi-challenge (and multi-user) chosen-ciphertext attack (CCA) security. Our schemes are based on variants of the Hashed ElGamal and Fujisaki-Okamoto transformation key encapsulation mechanisms. These variants augment ciphertexts with random strings that are included in hash function calls.

MULTI-CHALLENGE SETTING. As mentioned, our focus in this work is on multichallenge and multi-user security. This is simply motivated by the fact that encryption schemes get deployed across many different users each of whom will encrypt many messages, so it is important to understand how the security of a scheme degrades as the number of encryptions increase. In particular, the goal of tight proofs is to show that security does not meaningfully degrade. Multiple papers [22, 25, 16] have looked at this in the non-memory-aware setting, providing schemes with TA-tight proofs of security. However, extending any of these proofs to the memory-aware setting is quite difficult.

Prior works on memory-tight CCA secure encryption have identified a primary difficulty in the multi-challenge setting which lies in how the decryption oracle handles challenge ciphertexts. Simply decrypting a challenge ciphertext would lead to trivial attacks against any scheme, so instead the decryption oracle has to recognize these ciphertexts and respond to them in a special manner.¹ This makes writing memory-tight security proofs difficult because the reduction adversary must emulate this differing behavior on decryption queries for challenge or non-challenge ciphertexts, but it is unclear how to go about identifying which are challenge ciphertexts other than remembering and checking against all ciphertexts that were previously returned to encryption queries. In the singlechallenge setting is a non-issue, because storing the single challenge ciphertext requires minimal memory.

MEMORY-TIGHTNESS OF HASHED ELGAMAL. In recent years, several papers have discussed the challenge of providing memory-tight security proofs for Hashed ElGamal. Auerbach, et al. [3] gave it at as an example of a proof they considered the memory complexty of, but were unable to improve. Follow-up work by Bhattacharyya [7] and Ghoshal and Tessaro [19] analyzed this further, giving what might seem at first to be contradictory results. Bhattacharyya gave a memorytight proof for Hashed ElGamal in the single-challenge setting while Ghoshal and Tessaro proved a lower-bound establishing that a memory-tight proof for Hashed ElGamal was not possible.

Resolving this contradiction requires more precisely understanding each result. The lower bound applies specifically for reductions to Strong Computational Diffie-Hellman (CDH) security [2] which are "black-box" in several ways, including that they do not depend on the particular group used. Ghoshal and Tessaro note that Bhattacharyya's result (for single-challenge security) avoids the lower bound by not being black-box in this manner; it depends on the group having

¹ An alternate definitional style would disallow the adversary from querying challenge ciphertexts to its decryption oracle, but prior the works argue this is an inappropriate restriction in the memory-bounded setting [18, 17].

Scheme/Transform	Assumption	Result
Hashed ElGamal (aECIES)	Pair CDH	\$CCA KEM
Cramer-Shoup ElGamal (aCS)	Strong CDH	\$CCA KEM
Twin ElGamal (aTWIN)	CDH	\$CCA KEM
Т	\$CPA PKE	OW-PCA PKE
aV	OW-PCA PKE	OW-PCVA PKE
aU [⊥]	OW-PCVA PKE	CCA KEM
KEM/DEM (KD)	\$CCA KEM/SKE	\$CCA PKE

Fig. 1. TAM-tight reductions we provide. Transformations T, aV , and aU^{\perp} are FO transforms discussed in Section 5. Results are multi-user, multi-challenge security.

an efficient pairing. However, for efficiency it is preferable to implement schemes using elliptical curves for which efficient pairing are believed not to exist.

Our result for Hashed ElGamal is black-box in the sense of Ghoshal and Tessaro. We avoid the lower-bound without requiring an efficient pairing by introducing and using an assumption (Pair CDH) which is stronger than Strong CDH, but is reasonable to assume holds in typical groups based on elliptic curves.² We discuss this assumption in more detail momentarily. Indeed, Ghoshal and Tessaro say in their paper [19, Sec. 3.1, p.42], "it appears much harder to extend our result to different types of oracles than [the Strong CDH oracle], as our proof is tailored at this oracle." Our new security notion gives an example of such an oracle to which their result cannot be extended.

1.1 Our results

We summarize our results in Fig. 1. Omitted proofs and results are provided in the full version of this paper [23].

HASHED ELGAMAL. Our first results consider the security of Hashed ElGamal. Following Bhattacharyya, we actually consider two variants which we refer to as the ECIES [1] and the Cramer-Shoup [8] variants. The negative results of Ghoshal and Tessaro apply only to the ECIES variant. In both, the decryption key is a value $x \in \mathbb{Z}_p^*$ and the encapsulation key is $X = \mathbf{g}^x$. Here \mathbf{g} is a generator of a group of prime order p. For encapsulation, one samples a fresh $y \leftarrow \mathbb{Z}_p^*$ and returns \mathbf{g}^y as the ciphertext. For ECIES the derived key is $H(X^y)$, while for Cramer-Shoup it is $H(\mathbf{g}^y, X^y)$. Our main results concern "augmented" versions of both of these schemes where the ciphertexts are instead (a, \mathbf{g}^y) where a is a uniformly random bitstring used as an additional input to the hash function.

² Technically, the lower bound also does not apply because we are considering an augmented scheme which differs from the one analyzed by Ghoshal and Tessaro. However, the augmentation is not important for this comparison, because in the single-challenge setting where the lower-bound was proven, Pair CDH TAM-tightly implies security of the non-augmented scheme as well.

To understand these results, let us discuss the high level idea of proving security for ECIES. A standard, single-challenge proof would work from the Strong CDH assumption in the random oracle model. In Strong CDH an adversary is given $X = \mathbf{g}^x$, $Y = \mathbf{g}^y$, and an oracle O which on input B, C tells whether $B^x = C$. Its goal is to return X^y . The only way to distinguish $H(X^y)$ from random is to query H on input X^y . So a reduction adversary will give X as the encryption key, Y as the challenge ciphertext, simulate random oracle and decapsulation queries, and checks if any of the random oracle queries are X^y in which case it returns that. The oracle O is used for checking whether random oracle queries are X^y (for a random oracle query Z, one queries O(Y,Z)) to check) and for maintaining consistency between random oracle and decapsulation queries for non-challenge ciphertexts. A decapsulation query for Y and a random oracle query for Y, Z should return the same result if $Y^x = Z$. The reduction can maintain this consistency by remembering all of the queries made to both oracles and then using O to check for this consistency. This is neither time- nor memory-tight.

Bhattacharyya was able to make this TAM-tight by introducing a new technique for this consistency aspect. They simulate the random oracle H(C) by $h(e(\mathbf{g}, C))$ where h is a random function and e is a pairing. Then the output of a non-challenge decapsulation query B can be simulated as h(e(X, B)). In our proof we use a similar technique, but replace the requirement for a pairingfriendly group by using a new variant of CDH we will discuss momentarily.

The first step in making the proof work in the multi-challenge setting is to use Diffie-Hellman rerandomization techniques so we can have multiple Diffie-Hellman challenges. We let the *u*-th user's public key be X^{x_u} and the *i*-th ciphertext by Y^{y_i} . For memory-tightness, we pick x_u and y_i using a (pseudo-)random function.³ Then if the adversary makes a random oracle query H(C) where $C = X^{x_u \cdot y \cdot y_u}$, we have $C^{1/(x_u \cdot y_u)} = X^y$. A challenge here is to know which *u* and *i* to use for such a random oracle query. A reduction could check each choice of *u*, *i*, but this would lose time-tightness. At the same time, for a decapsulation query *B* we must be able to identify if *B* was a prior challenge query. Storing all prior challenge queries loses memory-tightness.

Both of these issues are solved by our addition of an auxiliary string a to each ciphertext and hash query. The idea here is based on memory-tightness techniques of GGJT [17], in that we are going to hide the pertinent information we need in a. Rather than sampling a at random, if the *i*-th challenge query is made to user u then our reduction adversary picks a to be the "encryption" of (u, i). Now on future random oracle and decapsulation queries we can recover u, iby "decrypting" a. This allows us to properly simulate the view of the adversary.

PAIR CDH SECURITY. As we have been mentioning, we avoid the need for groups with pairing in our result for ECIES by making use of a new computational assumption. This assumption, we refer to as Pair CDH security, extends CDH security by giving the adversary access to an oracle which, on input A and B

³ In the body, this is separated out as a proof that single-challenge CDH tightly implies multi-challenge CDH.

(with discrete logarithms a and b) computes a and b then returns a random function applied to $a \cdot b$. This acts, in essence, as a pairing from the group under consideration to a randomly chosen group. Our use of this in our security proof for ECIES takes advantage of the fact that (i) the pairing is only needed for the proof, not in the construction itself and (ii) the proof does not require the ability to efficiently perform group operations with the output of the pairing. We think this notion may be of further interest if other proofs can be found where better tightness can be achieved using a pairing only in the reduction.

To justify our new assumption we analyze how it compares to existing assumptions. Pair CDH security is implied by CDH security if the group under consideration has an efficient pairing. This holds because we can emulate the random pairing by first applying the efficient pairing and then applying a random function (which may be pseudorandomly instantiated for efficiency). In turn, Pair CDH implies the Gap CDH assumption because a pairing can be used to check whether given group elements form a Diffie-Hellman triple.

These results do not justify the use of Pair CDH security for typical groups based on elliptic curves which do not have pairings. For this, we turn to nonstandard models (i.e. algebraic or generic group models [29, 26, 13]). In these models, we are able to show that CDH and Pair CDH are equivalent because learning anything from the oracle requires the ability to find non-trivial collisions in the pairing. The ability to find such collisions can in turn be used to solve the discrete logarithm problem.

FUJISAKI-OKAMOTO TRANSFORMATION. The other KEMs we consider are those derived from the Fujisaki-Okamoto Transformation which starts with a CPA secure public key encryption scheme and applies several random oracle based transformations to construct a CCA secure KEM. Hofheinz, Hövelmanns, and Kiltz [21] gave a nice modular approach for proving the security of several variants of this transformation. Bhattacharyya showed how to make these proofs memory-tight in the single-challenge setting (in some cases requiring one additional intermediate transformation). We extend these to the multi-challenge setting. For the final step of the transform, we need to consider an augmented transform of the existing scheme in which random strings are added to each ciphertext and incorporated into the hash queries. As before, our reduction samples these string as the encryption of the pertinent information it would need to identify challenge ciphertexts and respond to them appropriately.

For these results, we require that the starting CPA scheme have good multichallenge security. This is a significantly weaker starter point than multi-challenge CCA security because it avoids the issue of having to be able to identify challenge ciphertexts for the decryption oracle.

LIFTING TO PUBLIC-KEY ENCRYPTION. The approaches described above are for key encapsulation mechanisms. This raises the question of whether these tight reductions can be applied to public-key encryption via the KEM-DEM paradigm. It uses a KEM to generate a new symmetric key for a data encapsulation mechanism to encrypt the actual message with. This was previously looked at by GGJT [17], who gave a TAM-tight proof of security. However, because of their particular motivations, the proof assumed the KEM was constructed from a public-key encryption scheme. We show that (with some modifications) the proof works with generic KEMs as well.

2 Preliminaries

2.1 Notation

We recall basic notion and security definitions we will use in our paper.

PSEUDOCODE. For our proofs, we use the code based framework of [6]. If \mathcal{A} is an algorithm, then $x \leftarrow \mathcal{A}^{\mathcal{O}}(x_1, x_2, ...; r)$ denotes running \mathcal{A} on inputs $x_1, x_2, ...$ with coins r and having access to the set of oracles \mathcal{O} to produce output x. We use the notation \leftarrow ^s instead of \leftarrow when not explicitly specifying the coins r. If S is a set, |S| denotes its size and $x \leftarrow$ ^s S denotes sampling x uniformly from S. We use the symbol \perp to indicate rejection. When not specified, tables are initialized empty and integers are initialized to 0.

Security notions are defined with games such as the one in Fig. 3. The probability that the game G outputs true is $\Pr[G]$. We sometimes use a sequence of "hybrid" games in one figure for our proofs. We use comments of the form $//G_{[i,j)}$ to indicate that a line of code is included in games G_k for $i \leq k < j$. To identify changes made to the k^{th} hybrid, one looks for lines of code commented as $//G_{[i,k)}$ for code that is no longer included in the k^{th} hybrid and $//G_{[k,j)}$ for code that is new to the k^{th} hybrid.

COMPLEXITY MEASURES. Following ACFK [3], we measure the local complexities of algorithms and do not include the complexity of oracles that they interact with. We focus on the worst case runtime $\text{Time}(\mathcal{A})$ and memory used for local computation $\text{Mem}(\mathcal{A})$ of any algorithm \mathcal{A} .

FUNCTIONS AND IDEAL MODELS. We define $\mathsf{Fcs}(D, R)$ (resp. $\mathsf{lnj}(D, R)$) to be the set of all functions (resp. injections) mapping from D to R. For $f \in \mathsf{lnj}(D, R)$, we define f^{-1} to be its inverse (with $f^{-1}(y) = \bot$ if y has no preimage). If D_t and R_t are sets for each $t \in T$, then we define $\mathsf{Fcs}(T, D, R)$ (resp. $\mathsf{lnj}(T, D, R)$) to be the set of functions f so that $f(t, \cdot) \in \mathsf{Fcs}(D_t, R_t)$ (resp. $f(t, \cdot) \in \mathsf{lnj}(D_t, R_t)$). We let $f_t(\cdot) = f(t, \cdot)$.

For $f \in \text{Inj}(D, R)$ we let f^{\pm} denote the function defined by $f^{\pm}(+, x) = f(x)$ and $f^{\pm}(-, x) = f^{-1}(x)$. We often write f(x) or $f^{-1}(x)$ in place of $f^{\pm}(+, x)$ or $f^{\pm}(-, x)$. We let $\text{Inj}^{\pm}(D, R) = \{f^{\pm} : f \in \text{Inj}(D, R)\}$ and extend this to define $\text{Inj}^{\pm}(T, D, R)$ analogously.

Ideal models (e.g. the random oracle or ideal cipher model) are captured by having a scheme S specify a set of functions S.IM. Then, at the beginning of a security game for S, a random $\mathcal{H} \in S.IM$ is sampled. The adversary and some algorithms of the scheme S are then given oracle access to \mathcal{H} . The standard model is captured by S.IM being a singleton set containing the identity function.

If F and G are sets of functions, then we define $(F,G) = F \times G = \{ f \times g : f \in F, g \in G \}$. Here, $f \times g$ is the function defined by $f \times g(0,x) = f(x)$ and $f \times g(1,x) = g(x)$. In the code of an algorithm expecting oracle access to

 $f \times g \in \mathsf{F} \times \mathsf{G}$, we write f(x) or g(x) with the natural meaning. We extend this notation to more than two sets of functions as well.

SWITCHING LEMMA. Our proofs make use of the indistinguishability of random functions and injections, as captured by the following standard result.

Lemma 1 (Switching Lemma). Fix T, D, R and $N = \min_{t \in T} |R_t|$. For any adversary \mathcal{A} making at most q queries, it holds that $|\Pr[\mathcal{A}^f \Rightarrow 1] - \Pr[\mathcal{A}^g \Rightarrow 1]| \leq 0.5 \cdot q^2/N$, where the probability is taken over the randomness of \mathcal{A} , sampling $f \leftarrow \operatorname{s} \operatorname{Fcs}(T, D, R)$, and sampling $g \leftarrow \operatorname{s} \operatorname{Inj}(T, D, R)$.

2.2 Memory-tightness background

 \mathcal{F} -ORACLE ADVERSARIES. We adopt GGJT's [17] oracle adversary formulation for our proofs in the memory-aware setting, i.e., we allow reductions to access uniformly random functions or invertible random injections. Our reductions are of the form shown below for some set of functions \mathcal{F} and algorithm \mathcal{B} . We call such an adversary \mathcal{A} an \mathcal{F} -oracle adversary.

$$\frac{\text{Adversary } \mathcal{A}^{O}(\text{in})}{f \leftarrow \mathfrak{F}}$$

out $\leftarrow \mathcal{B}^{O,f}(\text{in})$
Return out

The complexity of adversary \mathcal{A} would include the (large) complexity of sampling, storing, and computing f. However, as proposed in [17], we present theorems in terms of the *reduced complexity* of an oracle aided adversary which is defined as $\mathbf{Time}^*(\mathcal{A}) = \mathbf{Time}(\mathcal{B})$ and $\mathbf{Mem}^*(\mathcal{A}) = \mathbf{Mem}(\mathcal{B})$.

We refer readers to Lemma 2 of [17] which bounds how much an adversary may be aided by a random object by replacing it with a pseudorandom version of the object. Pseudorandom injections can typically be instantiated by appropriately chosen encryption schemes.

There is a small issue when pseudorandomly instantiating a random function if the game \mathcal{A} plays is inefficient. This is the case for some of our reduction adversaries playing CDH variants wherein they have access to some inefficient oracle based on the group. Then the pseudorandomness reduction adversary from [17] will be inefficient because it simulates the game that \mathcal{A} is playing. However, we can simply use pseudorandom schemes believed to be secure even against adversaries with access to the inefficient oracle. This seems reasonable as we can choose a pseudorandom scheme which seems unrelated to the group.

MESSAGE ENCODING TECHNIQUES. The message encoding technique proposed by GGJT in [17] programs randomness that a reduction provides to an adversary in a special way that stores retrievable state information. This is achieved by generating randomness as the output of random injections. The reduction may then invert randomness generated thusly to retrieve state information. For example, consider a key encapsulation mechanism that outputs ciphertexts of the

PKE Syntax	KEM Syntax	SKE Syntax
$\overline{(ek,dk)} \leftarrow PKE.K$	$(ek, dk) \leftarrow KEM.K$	$\overline{K \leftarrow \text{SKE.K}}$
$c \leftarrow *PKE.E^{\mathcal{H}}(ek, m)$	$(c, K) \leftarrow KEM.E^{\mathcal{H}}(ek)$	$c \leftarrow SKE.E^{\mathcal{H}}(K,m)$
$m \leftarrow PKE.D^{\mathcal{H}}(dk,c)$	$K \leftarrow KEM.D^{\mathcal{H}}(dk,c)$	$m \leftarrow SKE.D^{\mathcal{H}}(K,c)$

Fig. 2. Syntax of a public key encryption scheme PKE , key encapsulation mechanism KEM, and symmetric key encryption scheme SKE . The ideal model oracle is \mathcal{H} .

form (a, c) where a is uniformly random. Then a reduction can simulate challenge ciphertexts by setting a = f(i) where f is a random injection and i is some pertinent information the reduction would want to know if the adversary later makes oracle queries for the same ciphertext. Then the reduction can recover this information during future queries as $i \leftarrow f^{-1}(a)$.

MAP-THEN-RANDOM-FUNCTION. We describe the main proof technique of Bhattacharyya [7], namely "map-then-rf".⁴ This technique allows the reduction to use the composition of an injection and a random function to replace a random function. This relies on the simple fact that if $h \in \text{Inj}(D, S)$, then sampling f according to $f \leftarrow \text{s} \text{Fcs}(D, R)$ or $g \leftarrow \text{s} \text{Fcs}(S, R); f \leftarrow g \circ h$ are equivalent, meaning, if g is a random function, and h is any injection, then $f \leftarrow g \circ h$ is a random function. This allows a reduction to compute the output f(x) given h(x), even if it does not know x.

2.3 Public Key Encryption

SYNTAX. A public key encryption scheme, PKE, specifies three algorithms - the key generation algorithm (PKE.K) that returns a pair of keys (ek, dk) where ek is the encryption key and dk is the corresponding decryption key, the encryption algorithm (PKE.E) that takes the encryption key ek and a message m and returns ciphertext c, and the decryption algorithm PKE.D that takes the decryption key dk and a ciphertext c and returns message m (or the special symbol \perp to indicate rejection). The syntax of these algorithms is given in Fig. 2.

Perfect correctness requires $\mathsf{PKE}.\mathsf{D}^{\mathcal{H}}(dk,c) = m$ for all $(ek,dk) \in [\mathsf{PKE}.\mathsf{K}]$, all m, all $\mathcal{H} \in \mathsf{PKE}.\mathsf{IM}$, and all $c \in [\mathsf{PKE}.\mathsf{E}^{\mathcal{H}}(ek,m)]$. The weaker notion of δ -correctness requires that for all (not necessarily efficient) \mathcal{D} ,

$$\Pr[\mathsf{PKE}.\mathsf{D}^{\mathcal{H}}(dk,\mathsf{PKE}.\mathsf{E}^{\mathcal{H}}(ek,m)) \neq m : m \leftarrow \mathfrak{D}^{\mathcal{H}}(ek,dk)] \leq \delta(q)$$

where q upper bounds the number of \mathcal{H} queries \mathcal{D} makes. The probability is over $(ek, dk) \leftarrow \mathsf{PKE.K}, \mathcal{H} \leftarrow \mathsf{PKE.IM}$, and the coins of \mathcal{D} and PKE.E. When not stated otherwise, schemes are assumed to be perfectly correct.

⁴ Bhattacharyya actually uses "map-then-prf", as they were not using the oracle adversary formulation.

Game $G_{PKE,b}^{mu-\$cca}(\mathcal{A})$	$ \operatorname{Enc}_b(u,m) $	$\operatorname{Dec}(u,c)$
<i>H</i> ← ^{\$} PKE.IM	$\overline{c_1 \leftarrow \text{s PKE.E}^{\mathcal{H}}(ek_u, m)}$	If $M[u,c] \neq \bot$
$(ek_{(\cdot)}, dk_{(\cdot)}) \leftarrow *PKE.K$	$c_0 \leftarrow PKE.\mathcal{C}(ek_u, m)$	Return $M[u, c]$
$b' \leftarrow \mathcal{A}^{\operatorname{New}, \operatorname{Enc}_b, \operatorname{Dec}, \mathcal{H}}$	$M[u, c_b] \leftarrow m$	$m \leftarrow PKE.D^{\mathcal{H}}(dk,c)$
Return $b' = 1$	Return c_b	Return m
New(u)		
$\overline{\text{Return }}ek_u$		

Fig. 3. Game defining mu-\$cca security of PKE.

We define the encryption keyspace as $\mathsf{PKE.Ek} = \{ek : (ek, dk) \in [\mathsf{PKE.K}]\}$ and assume that for each $ek \in \mathsf{PKE.Ek}$ and allowed message length n, there exists a set $\mathsf{PKE.C}(ek, n)$ such that $\mathsf{PKE.E}^{\mathcal{H}}(ek, m) \in \mathsf{PKE.C}(ek, |m|)$ always holds. We assume this set is disjoint for distinct message lengths and let $\mathsf{PKE.C}^{-1}(ek, c)$ return n such that $c \in \mathsf{PKE.C}(ek, n)$. We let $\mathsf{PKE.R}$ denote the set from which $\mathsf{PKE.E}$ draws its randomness. Sometimes we assume that all messages to be encrypted are drawn from a set $\mathsf{PKE.M}$ of equal length messages and then let $\mathsf{PKE.C}$ simply denote the set of all possible ciphertexts.

INDISTINGUISHABLE FROM RANDOM SECURITY. We consider indistinguishable from random, chosen ciphertext attack (\$CCA) security as captured by Fig. 3. The definition multi-user and multi-challenge (allowing multiple challenges per user). It requires ciphertexts output by the encryption scheme be indistinguishable from random, even when given access to a decryption oracle. In this game, the adversary obtains the encryption key ek_u for user u by querying NEW(u). It makes an encryption query ENC(u, m) to receive a challenge encryption of m by u and a decryption query DEC(u, c) to have u decrypt c. The adversary needs to distinguish between the real world (b = 1) in which a query to ENC(u, m) returns a real encryption of m and the ideal world (b = 0) in which the same query returns a uniformly random element of PKE. $C(ek_u, |m|)$.

Table entry M[u, c] stores the message encrypted in user u's challenge ciphertext c. If the adversary queries DEC with a challenge ciphertext it returns M[u, c] rather than performing the decryption. Prior works on memory-aware cryptography [19, 17] considered other ways a decryption oracle might respond to challenge ciphertexts and argued that this is the "correct" convention. The advantage of an adversary \mathcal{A} is $\operatorname{Adv}_{\mathsf{PKE}}^{\mathsf{mu-\$cca}}(\mathcal{A}) = \Pr[\mathsf{G}_{\mathsf{PKE},1}^{\mathsf{mu-\$cca}}(\mathcal{A})] - \Pr[\mathsf{G}_{\mathsf{PKE},0}^{\mathsf{mu-\$cca}}(\mathcal{A})]$. In this and future definition we let \mathcal{U} denote the set of allowed user identifiers u.

The general framework of capturing multi-user security by allowing the attacker to access separate instances of oracles for each user with shared secret bit across them is originally due to Bellare, Boldyreva, and Micali [5] who provided a definition for IND-CPA secure public-key encryption.

ONE-WAYNESS SECURITY. Following the one-wayness security definitions in [21], we define variants of one-wayness security of PKE schemes in the multi-user,

Game $G_{PKE}^{mu-ow-w}(\mathcal{A})$	CHAL(u,i)	PCO(u, m, c)
<i>H</i> ← ^{\$} PKE.IM	$\overline{\mathrm{If}\ C[u,i]} \neq \bot$	$\overline{m'} \leftarrow PKE.D^{\mathcal{H}}(dk_u, c)$
$(ek_{(\cdot)}, dk_{(\cdot)}) \leftarrow PKE.K$	Return $C[u, i]$	Return $m = m'$
$O \leftarrow \perp //w = cpa$	$m \leftarrow SPKE.\mathcal{M}$	CVO(u, c)
$O \leftarrow PCO //w = pca$	$c \leftarrow PKE.E^{\mathcal{H}}(ek_u, m)$	$\left \frac{O V O(u,c)}{U}\right $
$O \leftarrow (PCO, CVO) / w = pcva$	$C[u,i] \leftarrow c$	$m' \leftarrow PKE.D^{\mathcal{H}}(dk_u, c)$
$(m', u, i) \leftarrow \mathcal{A}^{\text{New, CHAL}, O, \mathcal{H}}$	Return c	Return $(m' \in PKE.\mathcal{M})$
Return $PCO(u, m', CHAL(u, i))$	NEW(u)	
	$\boxed{\operatorname{Return} ek_u}$	

Fig. 4. Game defining mu-ow-*w* security of PKE for $w \in \{cpa, pca, pcva\}$.

multi-challenge setting in Fig. 4. We define three variants - One-Wayness under Chosen Plaintext Attacks (OW-CPA), One-Wayness under Plaintext Checking Attacks (OW-PCA) and One-Wayness under Plaintext and Validity Checking Attacks (OW-PCVA). The difference between each variant $w \in \{cpa, pca, pcva\}$ is in the auxilliary oracle(s) O that the adversary is given access to.

In each variant, the adversary is tasked with finding the decryption of a challenge ciphertext which encrypt a message randomly sampled from PKE. \mathcal{M} . In the game $G_{\mathsf{PKE}}^{\mathsf{mu-ow-cpa}}$, the adversary does not have access to any auxilliary oracle as indicated by $O \leftarrow \bot$. In the game $G_{\mathsf{PKE}}^{\mathsf{mu-ow-pca}}$, the adversary has access to the Plaintext Checking Oracle PCO which takes as input a valid message-ciphertext pair, and returns true if the message is a valid decryption of the ciphertext and false otherwise. The adversary has access to both oracles, PCO and CVO, in $G_{\mathsf{PKE}}^{\mathsf{mu-ow-pcva}}$ where CVO takes as input a ciphertext, and returns true if the ciphertext decrypts to a valid message. For each variant, we define $\mathsf{Adv}_{\mathsf{PKE}}^{\mathsf{mu-ow-w}}(\mathcal{A}) = \Pr[G_{\mathsf{PKE}}^{\mathsf{mu-ow-w}}]$. Note that an adversary may re-query CHAL(u,i) to get back the same ciphertext. This makes it hard to prove one-wayness, but easier write proofs starting from one-wayness. We sometime assume challenge identifiers, i, are drawn from a fixed set \mathcal{I} .

2.4 Key Encapsulation Mechanisms

SYNTAX. A key encapsulation mechanism, KEM, consists of three algorithms - the key generation algorithm (KEM.K) that returns a pair of keys (ek, dk) where ek is the encapsulation key and dk is the corresponding decapsulation key, the encapsulation algorithm (KEM.E) that takes the encapsulation key ek and returns a ciphertext-key pair (c, K) where $K \in \text{KEM}.\mathcal{K}$ and the decapsulation algorithm KEM.D that takes the decapsulation key dk and a ciphertext c and returns a key K (or \perp to indicate rejection). The syntax of these algorithms is shown in Fig. 2. Perfect correctness requires that KEM.D^{\mathcal{H}}(dk, c) = K for all $(ek, dk) \in [\text{KEM.K}]$, all $\mathcal{H} \in \text{KEM.IM}$, and all $(c, K) \in [\text{KEM.E}^{\mathcal{H}}(ek)]$.

$ \begin{array}{ c c c c c c c c c c c c c c c c c c c$	$\frac{\text{DECAP}(u,c)}{\text{If } T[u,c] \neq \bot}$ Return $T[u,c]$ $K \leftarrow KEM.D^{\mathcal{H}}(dk_u,c)$ Return K
---	--

Fig. 5. Game defining mu-\$cca security of KEM.

We define encryption keyspace $\mathsf{KEM}.\mathsf{Ek} = \{ek : (ek, dk) \in [\mathsf{KEM}.\mathsf{K}]\}$. For $ek \in \mathsf{KEM}.\mathsf{Ek}$ we let $\mathsf{KEM}.\mathcal{C}(ek)$ denote the ciphertext set $\{c : (c, K) \in [\mathsf{KEM}.\mathsf{E}(ek)]\}$ and define $|\mathsf{KEM}.\mathcal{C}| = \min_{ek \in \mathsf{KEM}.\mathsf{Ek}} |\mathsf{KEM}.\mathcal{C}(ek)|$. We let $\mathsf{KEM}.\mathcal{R}$ denote the set from which $\mathsf{KEM}.\mathsf{K}$ draws it randomness. We say that KEM is ϵ -uniform if for all $ek \in \mathsf{KEM}.\mathsf{Ek}, \mathcal{H} \in \mathsf{KEM}.\mathsf{IM}$, and (not necessarily efficient) \mathcal{D} it holds that

$$\Pr[\mathcal{D}(c) = 1 : c \leftarrow \mathsf{KEM}.\mathcal{C}(ek)] - \Pr[\mathcal{D}(c) = 1 : (c, \cdot) \leftarrow \mathsf{KEM}.\mathsf{E}^{\mathcal{H}}(ek)] \leq \epsilon.$$

INDISTINGUISHABLE FROM RANDOM SECURITY. Our notion of \$CCA security for KEMs is presented in Fig. 5, which requires that keys and ciphertexts output by the scheme be indistinguishable from random. The adversary is given a user instantiation oracle NEW, encapsulation oracle ENCAP, and a decapsulation oracle DECAP. Its goal is to distinguish between the real world (b = 1) where ENCAP returns true outputs from KEM.E and the ideal world (b = 0) where it returns a pair (c, K) chosen uniformly at random from KEM. $C(ek) \times KEM.K$.

The table T stores the keys corresponding to challenge ciphertexts output by the encapsulation oracle. The decapsulation oracle uses T to respond to challenge queries. The advantage of an adversary \mathcal{A} is defined as $\mathsf{Adv}_{\mathsf{KEM}}^{\mathsf{mu-\$cca}}(\mathcal{A}) = \Pr[\mathsf{G}_{\mathsf{KEM},0}^{\mathsf{mu-\$cca}}(\mathcal{A})] - \Pr[\mathsf{G}_{\mathsf{KEM},0}^{\mathsf{mu-\$cca}}(\mathcal{A})]$. We also define CCA security (via $\mathsf{Adv}_{\mathsf{KEM}}^{\mathsf{mu-\$cca}}$ and $\mathsf{G}_{\mathsf{KEM},b}^{\mathsf{mu-cca}}$) analogously to \$CCA security, except in the ENCAP oracle c_0 is set to equal c_1 rather than being sampled at random.

2.5 Symmetric Key Encryption

SYNTAX. A symmetric key encryption scheme, SKE, consists of three algorithms - the key generation algorithm (SKE.K) that returns a key K, the encryption algorithm (SKE.E) that takes the key K and a message m and returns ciphertext c, and the decryption algorithm SKE.D that takes the key K and a ciphertext c and returns message m (or \bot to indicate rejection). The syntax of these algorithms is given in Fig. 2. Perfect correctness requires that SKE.D^{\mathcal{H}}(K, c) = m for $K \in [SKE.K]$, all m, all $\mathcal{H} \in SKE.IM$, and all $c \in [SKE.E^{\mathcal{H}}(K,m)]$. We define the ciphertext, message, and expansion lengths of SKE by SKE.cl(|m|) = $|SKE.E^{\mathcal{H}}(K,m)|$ (requiring this to hold for all \mathcal{H}, K, m), SKE.ml(SKE.cl(l)) = l, and SKE.xl = min_l SKE.cl(l) - l respectively.

Game $G^{mu-\$cca}_{SKE,b}(\mathcal{A})$	$\operatorname{Enc}_b(u,m)$	DEC(u, c)
<i>H</i> ← s SKE.IM	$\overline{c_1} \leftarrow SKE.E^{\mathcal{H}}(K_u, m)$	$\overline{\text{If } M[u,c]} \neq \bot:$
<i>K</i> (.) ←\$ SKE.K	$c_0 \leftarrow \{0,1\}^{SKE.cl(m)}$	Return $M[u, c]$
$b' \leftarrow \mathcal{A}^{\mathrm{Enc}_b, \mathrm{Dec}, \mathcal{H}}$	$M[u, c_b] \leftarrow m$	$m \leftarrow SKE.D^{\mathcal{H}}(K_u, c)$
Return $b' = 1$	Return c_b	Return m

Fig. 6. Game defining mu-\$cca security of SKE.

INDISTINGUISHABLE FROM RANDOM CCA SECURITY. Our notion of \$CCA security for SKE schemes is captured by Fig. 6, which requires that ciphertexts output by the encryption scheme be indistinguishable from ciphertexts chosen at random. In this game, the adversary is given access to an encryption oracle ENC and a decryption oracle DEC. The adversary needs to distinguish between the real world (b = 1), where ENC returns an encryption of m under K_u and the ideal world (b = 0) where the output of ENC is sampled uniformly at random. The advantage of an adversary \mathcal{A} is defined as $\operatorname{Adv}_{\mathsf{SKE}}^{\mathsf{mu}-\$cca}(\mathcal{A}) = \Pr[\mathsf{G}_{\mathsf{SKE},0}^{\mathsf{mu}-\$cca}(\mathcal{A})] - \Pr[\mathsf{G}_{\mathsf{SKE},0}^{\mathsf{mu}-\$cca}(\mathcal{A})]$. We will only need "one-time" security in which the adversary only makes one encryption query per user.

3 Diffie-Hellman Definitions

In this section, we introduce the Computational Diffie-Hellman (CDH) assumptions we need for our later proofs. The first is a multi-user, multi-challenge variant of Strong CDH (which we need for one of our coming KEM proofs). We verify this is TAM-tightly implied by single-challenge variants. The second is a new definition we introduce, Pair CDH, which gives the adversary oracle access to a pairing from the group under consideration to a random group. We provide several results to understand the plausibility of Pair CDH security. We show that it always implies Gap CDH security and is {AM,TM}-tightly equivalent to CDH in algebraic/generic group models [29, 26, 13] or if the group has a pairing.

3.1 Group Syntax

A prime order group \mathbb{G} is a tuple (\mathbf{g}, p, \circ) where g is a group generator of prime order p under the group operation \circ . In our definitions we will treat the group as a priori fixed. We typically omit writing the group operation \circ explicitly and instead write group operations using multiplicative notation. We let $\langle \mathbf{g} \rangle = \{\mathbf{g}^a : a \in \mathbb{N}\}$. The discrete log(arithm) of an element $X \in \langle \mathbf{g} \rangle$ is the value $\operatorname{dlog}(X) \in \mathbb{Z}_p$ such that $\mathbf{g}^{\operatorname{dlog}(X)} = X$. We let $1_{\mathbb{G}} = \mathbf{g}^0$ denote the identity element. A pairing from $\mathbb{G} = (\mathbf{g}, p, \circ)$ to $\mathbb{G}_2 = (\mathbf{g}_2, p_2, \circ_2)$ is a map $e : \langle \mathbf{g} \rangle \times \langle \mathbf{g} \rangle \to \langle \mathbf{g}_2 \rangle$ satisfying $e(\mathbf{g}^x, \mathbf{g}^y) = \mathbf{g}_2^{xy}$. We let $\operatorname{Time}(\mathbb{G})$ and $\operatorname{Mem}(\mathbb{G})$ denote the time and memory complexity of computing exponentiations or multiplications in $\langle \mathbf{g} \rangle$.

$ \frac{\text{Game } G^{x}_{\mathbb{G}}(\mathcal{A})}{(g, p, \circ) \leftarrow \mathbb{G}} \\ x_{(\cdot)} \leftarrow \mathbb{Z}^{*}_{p}; y_{(\cdot, \cdot)} \leftarrow \mathbb{Z}^{*}_{p} \\ O \leftarrow \bot //G^{cdh} \\ O \leftarrow \text{STRONG} //G^{scdh} \\ O \leftarrow \text{GAP} //G^{gcdh} \\ O \leftarrow \text{PAIR} //G^{pcdh} \\ f \leftarrow \mathbb{Inj}(\mathbb{Z}_{p}, \mathbb{Z}_{p}) //G^{pcc} \\ (u, v, \mathbb{Z}_{p}) \leftarrow \mathfrak{g} \wedge \mathbb{N}^{New,CHALC} $	h $ \frac{ \operatorname{NEW}(u) }{X_u \leftarrow \mathbf{g}^{x_u}; \operatorname{Return} X_u} \\ \frac{ \operatorname{CHAL}(u,i) }{Y_{u,i} \leftarrow \mathbf{g}^{y_{u,i}}; \operatorname{Return} Y_{u,i}} $	$\frac{ GAP(A, B, C) }{a \leftarrow \operatorname{dlog}(A); b \leftarrow \operatorname{dlog}(B)}$ Return $(C = \mathbf{g}^{ab})$ STRONG (u, B, C) Return GAP (\mathbf{g}^{xu}, B, C) $\frac{ PAIR(A, B) }{a \leftarrow \operatorname{dlog}(A); b \leftarrow \operatorname{dlog}(B)}$ Return $f(ab)$
$f \leftarrow \sup(\mathbb{Z}_p, \mathbb{Z}_p) //G^{pcc}$	h	$\boxed{a \leftarrow \operatorname{dlog}(A); b \leftarrow \operatorname{dlog}(B)}$ Beturn $f(ab)$
$\operatorname{Return} \left(Z = \mathbf{g}^{x_u y_{u,i}} \right)$		

Fig. 7. Security games capturing several variants of the computational Diffie-Hellman problem, namely, CDH, Gap CDH, Strong CDH, and Pair CDH. The last of these is a new notion we introduce which gives the attacker access to a pairing from \mathbb{G} to a random group.

3.2 Computational Diffie-Hellman variants

In this paper we will make use of several variants of the Computational Diffie-Hellman assumption. These security notions are defined by the game shown in Fig. 7. In each, the adversary is given access to a \mathbf{g}^x and \mathbf{g}^y with the goal of producing \mathbf{g}^{xy} . For our later security proofs, it was useful to write "multiuser" and "multi-challenge" version of these games. Thus rather than giving the adversary a single \mathbf{g}^x , we give it access to an oracle NEW which on input a string u (which we think of as identifying a user) returns a fresh \mathbf{g}^{x_u} . Similarly, the adversary is given access to an oracle CHAL which on inputs string u and i (which we think of as identifying a challenge) returns a fresh $\mathbf{g}^{y_{u,i}}$. For the memory-tightness of future proofs, it is important that the attacker can repeat queries, obtaining the same result as before. The goal of the attacker is to return $\mathbf{g}^{x_u y_{u,i}}$ for any choice of u and i.

The different variants of CDH are captured by the games differing in what (if any) auxiliary oracle O the adversary is given. The standard notion of CDH security is captured by the game G^{cdh} in which the adversary is not given any auxiliary oracle, as expressed by the code O $\leftarrow \bot$. Gap CDH security [27] is captured by G^{gcdh} in which the adversary's oracle GAP takes as input a tuple (A, B, C) and outputs a boolean indicating whether this is a valid Diffie-Hellman tuple (i.e. $C = \mathbf{g}^{\operatorname{dlog}(A) \operatorname{dlog}(B)}$). The Strong CDH game G^{scdh} [2] is a weakened version of Gap CDH in which the oracle only allows tuples of the form (\mathbf{g}^{x_u}, B, C) .

The final variant is a new security notion we introduce called Pair CDH. In this game $\mathbf{G}^{\mathsf{pcdh}}$, the adversary is given access to the oracle PAIR which on input (A, B) returns $f(\mathbf{g}^{ab})$ where a, b are the discrete logs of A, B and f is a random injection. This oracle can be thought of being a pairing to a random group $\mathbb{G}_2 = (\mathbf{g}_2, p, \circ_2)$ where $\mathbf{g}_2 = \mathsf{PAIR}(\mathbf{g}, \mathbf{g})$ and $h \circ_2 h' = f(f^{-1}(h) \circ f^{-1}(h'))$. Note that \mathcal{A} is not able to efficiently compute the operation \circ_2 .

Adversary $\mathcal{B}_{x}^{\mathrm{New,Chal},\mathrm{O}}$	SIMNEW(u)
$(\mathbf{g}, p, \circ) \leftarrow \mathbb{G}$	$\overline{x'_u \leftarrow g(u)}$
$g \leftarrow Fcs(\mathcal{U}, \mathbb{Z}_p^*); h \leftarrow Fcs(\mathcal{U} \times \mathcal{I}, \mathbb{Z}_p^*)$	Return $X^{x'_u}$
$X \leftarrow \text{NEW}(1); Y \leftarrow \text{CHAL}(1,1)$ If $x = \text{scdh}$ then SIMO \leftarrow SIMSTRONG Else SIMO $\leftarrow O$	$\frac{\text{SIMCHAL}(u,i)}{y'_{u,i} \leftarrow h(u,i)}$
$(u, i, Z) \leftarrow \mathcal{A}^{\text{SIMNew,SIMCHAL,SIMO}}$	Return $Y^{y'_{u,i}}$
$\begin{aligned} x'_u \leftarrow g(u); y'_{u,i} \leftarrow h(u,i) \\ \text{Return } (1,1,Z^{1/(x'_u y'_{u,i})}) \end{aligned}$	$\frac{\text{SIMSTRONG}(u, B, C)}{x'_u \leftarrow g(u)}$ Return O(1, B, C ^{1/x'_u})

Fig. 8. Adversary used for Lemma 2.

For $x \in {cdh, scdh, gcdh, pcdh}$ we define $Adv_{\mathbb{G}}^{\times}(\mathcal{A}) = \Pr[G_{\mathbb{G}}^{\times}(\mathcal{A})]$. We sometimes need to restrict user identifiers, u, to be from some fixed set \mathcal{U} and challenge identifiers, i, to be a from a fixed set \mathcal{I} .

MULTI-CHALLENGE SECURITY. Standard proofs use Diffie-Hellman rerandomization techniques to show that single-challenge security TA-tightly implies multichallenge security for most variants of Diffie-Hellman-based security notions. The following lemma extends this to TAM-tightness for the notions considered in this paper. The proof is an extension of standard Diffie-Hellman rerandomization techniques that picks the values used for rerandomization as the output of a random function, rather than picking them randomly and storing them.

Lemma 2 (Single-challenge \Rightarrow multi-challenge). Let \mathbb{G} be a group and $x \in \{cdh, scdh, gcdh, pcdh\}$. Let \mathcal{A} be an adversary for $G_{\mathbb{G}}^{\times}$ with $(q_{NEW}, q_{CHAL}, q_O) =$ **Query**(\mathcal{A}). Then we can construct a $(Fcs(\mathcal{U}, \mathbb{Z}_p^*), Fcs(\mathcal{U} \times \mathcal{I}, \mathbb{Z}_p^*))$ -oracle adversary \mathcal{B}_{\times} (given in the proof) such that

 $\begin{aligned} \mathsf{Adv}_{\mathbb{G}}^{\mathsf{x}}(\mathcal{A}) &= \mathsf{Adv}_{\mathbb{G}}^{\mathsf{x}}(\mathcal{B}_{\mathsf{x}}) \\ \mathbf{Query}(\mathcal{B}_{\mathsf{x}}) &= (1, 1, q_{\mathrm{O}}) \\ \mathbf{Time}^{*}(\mathcal{B}_{\mathsf{x}}) &= O(\mathbf{Time}(\mathcal{A}) + (q_{\mathrm{NEW}} + q_{\mathrm{CHAL}} + q_{\mathrm{O}} + 1)\mathbf{Time}(\mathbb{G})) \\ \mathbf{Mem}^{*}(\mathcal{B}_{\mathsf{x}}) &= O(\mathbf{Mem}(\mathcal{A}) + 2\mathbf{Mem}(\mathbb{G})). \end{aligned}$

Proof (of Lemma 2). Consider the adversary \mathcal{B}_{X} shown in Fig. 8. It makes a single query NEW(1) to obtain a group element X and a single query CHAL(1, 1) to obtain a group element Y. Then it runs \mathcal{A} . Let $x = \operatorname{dlog}(X)$ and $y = \operatorname{dlog}(Y)$.

It responds to SIMNEW(u) queries with $X^{x'_u}$ for x'_u the output of its random function. Letting $x_u = xx'_u$, note that $X^{x'_u} = \mathbf{g}^{x_u}$ and x_u is uniformly random because x'_u is. So this oracle has the correct distribution. It responds to SIMCHAL queries with $Y^{y'_{u,i}}$ for $y'_{u,i}$ output by its random function. Letting $y_{u,i} = yy'_{u,i}$, note that $Y^{y'_{u,i}} = \mathbf{g}^{y_{u,i}}$ and that $y_{u,i}$ is uniformly random because $y'_{u,i}$ is. For CDH, Gap CDH, or Pair CDH security \mathcal{B}_{x} gives \mathcal{A} direct access to O. For Strong CDH security ($\mathsf{x} = \mathsf{scdh}$), \mathcal{B}_{x} simulates the oracle by replacing a query (u, B, C) with a query $(1, B, C^{1/x'_u})$ which has the same behavior.

When \mathcal{A} finally halts and outputs (u, i, Z) the adversary \mathcal{B}_{x} halts and outputs $(1, 1, Z^{1/(x'_u y'_{u,i})})$. We claim that \mathcal{B}_{x} wins whenever \mathcal{A} would. To see this, note that if \mathcal{A} wins then $Z = \mathbf{g}^{x_u y_{u,i}} = \mathbf{g}^{(xx'_u)(yy'_{u,i})}$ and so $Z^{1/(x'_u y'_{u,i})} = \mathbf{g}^{xy}$. \Box

3.3 Studying Pair CDH

Pair CDH is a new computational assumption that we've introduced for this work. In this section we provide a few results to give a sense of its difficulty. Namely, we note that Pair CDH security implies Gap CDH security and that it is equivalent to CDH security in certain settings (if \mathbb{G} has an efficient pairing to some \mathbb{G}_2 , in the algebraic model, and in the generic group model). Given Lemma 2, we focus on the case that the adversary makes only single query each to NEW and CHAL. For brevity, we sketch the relationships here.

PAIR CDH \Rightarrow GAP CDH. To see that Pair CDH security implies Gap CDH security we need only note that GAP(A, B, C) = true if and only if PAIR(A, B) = PAIR (\mathbf{g}, C) . Hence a Pair CDH adversary can efficiently simulate the view of a Gap CDH adversary.

CDH + PAIRING \Rightarrow PAIR CDH. We claim that CDH security implies Pair CDH security if \mathbb{G} has an efficient pairing e to some group \mathbb{G}_2 with generator \mathbf{g}_2 . We can achieve TAM-tightness in this implication by using a $\operatorname{lnj}(\langle \mathbf{g}_2 \rangle, \mathbb{Z}_p)$ oracle adversary. Letting f' denote the random injection, our CDH adversary can simulate PAIR by responding to queries for (A, B) with f'(e(A, B)). If a = $\operatorname{dlog}(A)$ and $b = \operatorname{dlog}(B)$, then $f'(e(A, B)) = f'(e(\mathbf{g}, \mathbf{g})^{ab})$. Note that $F(\cdot) =$ $e(\mathbf{g}, \mathbf{g})^{(\cdot)}$ is an injection and so $f(\cdot) = f'(e(\mathbf{g}, \mathbf{g})^{(\cdot)})$ is a random injection. Hence this perfectly emulates PAIR.

 $\text{CDH} + \text{AGM}/\text{GGM} \Rightarrow \text{PAIR CDH}$. We claimed that CDH security implies Pair CDH security in the algebraic group model. More precisely, we {AM,TM}tightly show that CDH security implies Pair CDH security using a $\text{Fcs}(\mathbb{Z}_p^6, \mathbb{Z}_p)$ oracle adversary. We show a way to imperfectly simulate PAIR for algebraic adversaries such that distinguishing this from the real oracle requires the ability to solve the discrete log problem (given \mathbf{g}^c for a random c, find c). Noting that CDH security implies discrete log security gives our claim.

Let X and Y denote the challenge group elements and let $x = \operatorname{dlog}(X)$ and $y = \operatorname{dlog}(Y)$. An algebraic adversary, when making an oracle query (A, B) to PAIR is required to additionally provide "explanations" (a_1, a_2, a_3) and (b_1, b_2, b_3) such that $A = g^{a_1}X^{a_2}Y^{a_3}$ and $B = g^{b_1}X^{b_2}Y^{b_3}$. Then the true PAIR would respond with $f((a_1 + a_2x + a_3y) \cdot (b_1 + b_2x + b_3y))$. Our CDH adversary will think of this input to f as a degree-two polynomial $P_{A,B}(\mathbf{x}, \mathbf{y}) \in \mathbb{Z}_p[\mathbf{x}, \mathbf{y}]$ whose coefficients it can compute given the explanations for A and B. Letting (c_1, c_2, \ldots, c_6) denote these coefficients and $f' \in \operatorname{Fcs}(\mathbb{Z}_p^6, \mathbb{Z}_p)$, we simulate the output of PAIR as $f'(c_1, c_2, \ldots, c_6)$. Distinguishing this from the true oracle requires finding (A, B) and (A', B') such that $P_{A,B} \neq P_{A',B'}$ (as polynomials),

but $P_{A,B}(x,y) = P_{A',B'}(x,y)$. Using analysis techniques from [4], we can use the ability to find such "colliding" polynomials to solve the discrete log problem. We provide details of this analysis in the full version [23].

To achieve TM-tightness, the discrete log reduction picks two of the PAIR oracle queries at random and assumes that they give colliding polynomials. To achieve AM-tightness, we can check every pair of queries for collisions using the memory-tight rewinding technique of Auerbach, et al [3]. Namely, each time we reach a new PAIR oracle query while running the Pair CDH adversary, we pause and run an extra copy of that adversary from the start using the same coins. While running this extra copy, each time it makes a PAIR oracle query we check if this gives a colliding polynomial with the query we paused at in the first adversary. Ignoring memory tightness, \mathcal{A} could simply remember all of the PAIR oracle queries and check them at the end of execution, but then it is not clear how to achieve better time efficiency than checking each pair of queries.

When working in a generic group model [29, 26] we can use the same line of reasoning and then information theoretically bound the probability that an adversary finds colliding polynomials by $O(q^2/p)$ where q is the number of queries the Pair CDH adversary makes.

4 Hashed ElGamal KEMs

In this section we present the first example of KEMs with TAM-tight proofs in the multi-challenge setting. The KEMs we consider are variants of the ECIES and Cramer-Shoup Hashed ElGamal KEMs. These variants augment the existing schemes by adding random strings to the ciphertexts and random oracle queries. Our reductions make use of these strings to store pertinent information that will be needed to answer later oracle queries.

4.1 Augmented ECIES

AUGMENTED VERSION. We start with the ECIES [1] variant of Hashed ElGamal. Our augmented version of ECIES includes a random string *a* in the ciphertext. The augmented ECIES key encapsulation mechanism $\mathsf{aECIES}[\mathbb{G}, \mathcal{K}, l]$ is parameterized by a group $\mathbb{G} = (\mathbf{g}, p, \circ)$, key space \mathcal{K} , and length of the random string, *l*. The parameters \mathbb{G}, \mathcal{K} , and *l* are fixed for an instantiation of ECIES, so we use aECIES and $\mathsf{aECIES}[\mathbb{G}, \mathcal{K}, l]$ interchangeably. We define the scheme as follows with $\mathsf{aECIES}.\mathcal{K} = \mathcal{K}$ and $\mathsf{aECIES}.IM = \mathsf{Fcs}(\{0, 1\}^l \times \mathbb{G}, \mathcal{K})$.

<u>aECIES.K</u>	$aECIES.E^{\mathcal{H}}(ek)$	$aECIES.D^{\mathcal{H}}(dk,(a,Y))$
$x \leftarrow \mathbb{Z}_p^*$	$a \leftarrow {\$} \{0,1\}^l$	$Z \leftarrow Y^{dk}$
$ek \leftarrow \mathbf{g}^x$	$y \leftarrow \mathbb{Z}_p^*$	$K \leftarrow \mathcal{H}(a, Z)$
$dk \leftarrow x$	$Y \leftarrow \hat{\mathbf{g}^y}$	Return K
Return (ek, dk)	$Z \leftarrow ek^y$	
	$K \leftarrow \mathcal{H}(a, Z)$	
	Return $((a, Y), K)$	

OVERVIEW OF EXISTING TECHNIQUES AND ASSOCIATED CHALLENGES. Bhattacharyya [7] studied ECIES in the memory-aware setting. They pointed out the technique of simulating random oracles with PRFs introduced in [3] cannot be used for this family of KEMs, as in general, decapsulation queries cannot be simulated by the reduction. For example, if a PRF F is used to compute hashes as F(k, Z) instead of the random oracle H, for a decapsulation query Y the reduction would need to return $F(k, Y^{dk})$ which it cannot compute.⁵

Bhattacharyya used the map-then-prf technique as a workaround for groups with pairings. In this technique, the input Z to the random oracle is first operated on by a bilinear map e(g, Z), and then by the PRF F. Hence, the query H(Z)is simulated as $F(k, e(\mathbf{g}, Z))$ and a decapsulation query for Y can be simulated as F(k, e(ek, Y)) for all non-challenge ciphertexts. The reduction remembers the challenge ciphertext and returns appropriately when it is queried to DECAP.

This does not scale to the multi-user, multi-challenge setting since it requires that the reduction remembers all the challenge ciphertexts, incurring a memory overhead. Our solution for augmented ECIES combines Ghoshal et al.'s message encoding technique [17] with the map-then-rf technique. We encode the identifying information of challenge ciphertexts in a using a random injection so that this information can be recovered when an appropriate oracle query is made. To avoid the need for an efficiently computable pairing we make use of our new Pair CDH assumption. Our result is captured in Theorem 1.

Theorem 1 (Pair CDH \Rightarrow **\$CCA).** Let $\mathsf{aECIES} = \mathsf{aECIES}[\mathbb{G}, \mathcal{K}, l]$ where $\mathbb{G} = (\mathbf{g}, p, \circ)$ is a prime order group. Define $D_1 = \{0, 1\}^l \times \mathbb{G}$ and $D_2 = \mathcal{U} \times [q_{\text{ENCAP}}]$. Let \mathcal{A} be an adversary with $\mathbf{Query}(\mathcal{A}) = (q_{\text{NEW}}, q_{\text{ENCAP}}, q_{\text{DECAP}}, q_{\text{H}})$ and assume $2^l > |\mathcal{U}| \cdot q_{\text{ENCAP}}$. Then Fig. 12 gives a $(\mathsf{Fcs}(D_1, \mathcal{K}), \mathsf{lnj}(D_2, \{0, 1\}^l))$ -oracle adversary \mathcal{B} such that

$$\mathsf{Adv}_{\mathsf{aECIES}}^{\mathsf{mu-\$cca}}(\mathcal{A}) \leqslant 2\mathsf{Adv}_{\mathbb{G}}^{\mathsf{pcdh}}(\mathcal{B}) + \frac{q_{\mathrm{ENCAP}}^2}{2^l} + \frac{2q_{\mathrm{ENCAP}}(2q_{\mathrm{DECAP}} + |\mathcal{U}| \cdot q_{\mathrm{H}})}{2^l(p-1)}$$

 $\mathbf{Query}(\mathcal{B}) = ((q_{\text{NEW}} + q_{\text{ENCAP}}), (q_{\text{ENCAP}} + q_{\text{DECAP}} + q_{\text{H}}), (q_{\text{ENCAP}} + q_{\text{DECAP}} + 2q_{\text{H}}))$ $\mathbf{Time}^{*}(\mathcal{B}) = O(\mathbf{Time}(\mathcal{A})) \text{ and } \mathbf{Mem}^{*}(\mathcal{B}) = O(\mathbf{Mem}(\mathcal{A})).$

CHOOSING THE AUXILIARY STRING LENGTH. When instantiating this scheme one must choose the parameter l which determines the length of a. Larger l incurs a communication cost, while too small of a l can harm the concrete security results. We can expect the $q_{\text{ENCAP}}^2/2^l$ term to dominate the information theoretic part of the bound. With cautious choice of l = 256, the size of the ciphertext is not too significantly increased, but even with $q_{\text{ENCAP}} = 2^{64}$ encapsulations (an intentional overestimate of what seems likely) we get $q_{\text{ENCAP}}^2/2^l = 2^{-128}$.

Several of coming theorems use of an auxiliary string a of length l. Similar reasoning applies and l = 256 seems like a sufficient choice for all of them.

⁵ We discuss the use of PRFs to match prior work, but in the oracle adversary framework, we use random function oracles instead.

INTUITION. For each ENCAP query, our Pair CDH adversary programs the random string a as the output of a random injection applied to user identity u and counter i and simulates the random oracle $\mathcal{H}(a, Z)$ as $\mathcal{H}(a, \text{PAIR}(\mathbf{g}, Z))$. This allows us to simulate decapsulations because $\text{PAIR}(\mathbf{g}^x, \mathbf{g}^y) = \text{PAIR}(\mathbf{g}, \mathbf{g}^{xy})$.

Our adversary simulates *j*-th challenge ciphertexts $Y_{u,i}$ as CHAL(u, i). To determine if a decapsulation query (u, a, Y) is for a challenge ciphertext, the reduction first inverts *a* to obtain (v, j). If v = u, it re-queries CHAL(u, j) obtain the corresponding ciphertext $Y_{u,j}$. If $Y = Y_{u,j}$, the reduction assumes this was a challenge ciphertext. Finally, when the adversary \mathcal{A} queries the oracle H with (a, Z) such that $a^{-1} = (u, j)$ and $PAIR(\mathbf{g}, Z) = PAIR(NEW(u), CHAL(u, j))$, the reduction outputs Z and wins the Pair CDH game.

Proof (of Theorem 1). We use a sequence of hybrids H_0^1 through H_1^1 , H_0^2 through H_2^2 , and H_0^3 through H_3^3 presented in Figures 9, 10, and 11 where we establish the following claims that upper bound the advantage of adversary \mathcal{A} .

$$\begin{array}{ll} 1. \ \mathsf{Adv}_{\mathsf{a}\mathsf{ECleS}}^{\mathsf{mu}-\$\mathsf{cca}}(\mathcal{A}) &= 2 \Pr[\mathsf{H}_{0}^{1}] - 1 & \quad 6. \ \Pr[\mathsf{H}_{1}^{3}] \leqslant \Pr[\mathsf{H}_{2}^{3}] + \frac{q_{\mathsf{ENCAP}}(2q_{\mathsf{D}\mathsf{ECAP}} + |\mathcal{U}| \cdot q_{\mathsf{H}})}{2^{l}(p-1)} \\ 2. \ \Pr[\mathsf{H}_{0}^{1}] \leqslant \Pr[\mathsf{H}_{1}^{1}] + \frac{q_{\mathsf{ENCAP}}^{2}}{2 \cdot 2^{l}} & \quad 7. \ \Pr[\mathsf{H}_{2}^{3}] \leqslant \Pr[\mathsf{H}_{3}^{3}] + \Pr[\mathsf{bad}] \\ 3. \ \Pr[\mathsf{H}_{1}^{1}] &= \Pr[\mathsf{H}_{0}^{2}] = \Pr[\mathsf{H}_{1}^{2}] & \quad 8. \ \Pr[\mathsf{H}_{3}^{3}] \leqslant \frac{1}{2} \\ 4. \ \Pr[\mathsf{H}_{1}^{2}] &= \Pr[\mathsf{H}_{2}^{2}] = \Pr[\mathsf{H}_{0}^{3}] & \quad 9. \ \Pr[\mathsf{bad}] \leqslant \mathsf{Adv}_{\mathbb{G}}^{\mathsf{pcdh}}(\mathcal{B}) \\ 5. \ \Pr[\mathsf{H}_{0}^{3}] &= \Pr[\mathsf{H}_{1}^{3}] \end{array}$$

TRANSITION TO H_0^1 . We claim that the view of \mathcal{A} in H_0^1 is identical to its view in $G_{aECIE5,b}^{mu-\$cca}$ (Fig. 5) if b is chosen uniformly. In the latter, ENCAP₁ returns a ciphertext-key pair (c_1, K_1) such that c is the encapsulation of K_1 , and ENCAP₀ returns a ciphertext-key pair (c_0, K_0) where c_0 and K_0 are uniformly random elements of the ciphertext space and key space respectively. The same holds for the ENCAP oracle in H_0^1 . The table T in $G_{aECIE5,b}^{mu-\$cca}$ is indexed by (u, c) and stores the key that was returned by the ENCAP_b oracle for ciphertext c. The table Tin H_0^1 behaves analogously. The DECAP oracle in $G_{aECIE5,b}^{mu-\$cca}$ returns key K_b that was output by the ENCAP_b oracle when queried on a challenge ciphertext, and returns honest decapsulations otherwise. The same is true for H_0^1 . Note that H_0^1 's final output is whether b' = b, so standard conditional probability calculations give that $Adv_{aECIE5}^{mu-\$cca}(\mathcal{A}) = 2 \Pr[H_0^1] - 1$.

TRANSITION H_0^1 to H_1^1 . In H_1^1 , we make the following changes.

- 1. In the ENCAP oracle we switch from
 - (a) sampling the values $a_{(.,.)}$ uniformly to assigning them as the output of a random injection g evaluated on (u, i).
 - (b) sampling the values $y_{(.,.)}$ uniformly to assigning them the output of a random function h evaluated on (u, i). Switching to the random function h does not change the view of the adversary because the ordered pair (u, i) never repeats.
- 2. In the DECAP oracle we switch the If condition from checking whether $T[u, a, Y] \neq \bot$ to evaluating the boolean $u = v \land j \in I[u] \land Y = \mathbf{g}^{h(u,i)}$. These conditions are equivalent since $T[u, a, Y] \neq \bot$ iff a = g(u, j) and $Y = \mathbf{g}^{h(u,j)}$ for some $j \in I[u]$.

Hybrids H^1_{κ} for $0 \leq \kappa \leq 1$	New(u)
$b \leftarrow \$ \{0, 1\}$	$\overline{\text{Return }} X_u$
$i \leftarrow 0$	H(a,Z)
$x_{(\cdot)} \leftarrow \mathbb{Z}_p^*; X_{(\cdot)} \leftarrow \mathbf{g}^{x_{(\cdot)}}$	$\overline{\operatorname{Return}} \mathcal{H}(a, Z)$
$\mathcal{H} \leftarrow $ aECIES.IM	DECAP(u, a, Y)
$g \leftarrow \lim_{L \to \infty} (D_2, \{0, 1\}) / /\Pi_{[1,\infty)}$	$\frac{1}{(v, i)} \leftarrow a^{-1}(a) / H^{1}_{1}$
$n \leftarrow \text{FCS}(D_2, \mathbb{Z}_p) / / \Pi_{[1,\infty)}$	$u_{n,i} \leftarrow h(u, i) / / H_{1,i}^{1}$
$b \leftarrow \mathcal{A}$ Boturn $b' = b$	$[f T[u, a, Y] \neq \perp //H_{[0, 1)}^{1}$
Return $b = b$	$If u = v \land i \in I[u] \land Y = \mathbf{g}^{y_{u,j}} / / H_{i}^{1}$
$\frac{\text{ENCAP}_b(u)}{(u+1)}$	Return $T[u, a, Y]$
$i \leftarrow i + 1$ $I[a_i] \leftarrow I[a_i] + \{i\}$	$Z \leftarrow Y^{x_u}$
$a \leftarrow \{0, 1\}^l / (H_1^1,)$	$K \leftarrow \mathcal{H}(a, Z)$
$u_{u,i} \leftarrow \{0,1\} / /\Pi_{[0,1)}$	Return K
$g_{u,i} \leftarrow a(u,i) / (H_{1,i})$	
$u_{u,i} \leftarrow g(u,i) / H^{1}_{[1,\infty)}$	
$Y_{u,i} \leftarrow \boldsymbol{\sigma}^{y_{u,i}}$	
$Z_{u,i} \leftarrow X_{u}^{y_{u,i}}$	
$K_{u,i}^{1} \leftarrow \mathcal{H}(a_{u,i}, Z_{u,i})$	
$K_{u,i}^0 \leftarrow \mathcal{K}$	
$T[u, a_{u,i}, Y_{u,i}] \leftarrow K_{u,i}^b$	
Return $((a_{u,i}, Y_{u,i}), K^b_{u,i})$	

Fig. 9. First set of hybrids H_0^1 through H_1^1 used for proof of Theorem 1.

The only change in the adversary's view comes from 1(a). The switching lemma (Lemma 1) gives us $\Pr[\mathsf{H}_0^1] \leq \Pr[\mathsf{H}_1^1] + \frac{q_{\text{Excap}}^2}{2\cdot 2^l}$.

TRANSITION H_1^1 TO H_0^2 . The transition to hybrid H_0^2 is shown in Fig. 10. We have highlighted the ways in which H_0^2 differs from H_2^1 . Our changes are the following.

- 1. In the ENCAP oracle, the table \tilde{T} has been added to record the responses to the \mathcal{H} queries made within ENCAP.
- 2. In the H oracle, an If block is added to check if the input tuple (a, Z) was previously queried to \mathcal{H} from within the ENCAP oracle. In essence, the If block returns $\mathcal{H}(a, Z)$ when queried on a challenge ciphertext. The H oracle would behave the same way without the If block, as in H¹₁.

Hence, $\Pr[\mathsf{H}_1^1] = \Pr[\mathsf{H}_0^2]$.

TRANSITION H_0^2 TO H_1^2 . In game H_1^2 , we introduce the pairing oracle PAIR. It is only used by oracles within the game; the adversary does not have direct access to it. Note that $PAIR(\mathbf{g}, Z) = PAIR(\mathbf{g}^{x_u}, Y)$ iff $Z = Y^{x_u}$.

In the H oracle, we switch the condition from checking $\tilde{T}[a, Z] \neq \bot$ to evaluating the boolean $j \in I[u] \wedge \text{PAIR}(\mathbf{g}, Z) = \text{PAIR}(X_u, Y)$ where $(u, j) = g^{-1}(a), Y = \mathbf{g}^{h(u,j)}$. These two conditions are equivalent. Note that $\tilde{T}[a, Z] \neq \bot$

Hybrids H^2_{κ} for $0 \leq \kappa \leq 2$	NEW(u)
$b \leftarrow \$ \{0, 1\}$	$\overline{\text{Return }} X_u$
$i \leftarrow 0$	H(a, Z)
$x_{(\cdot)} \leftarrow \mathbb{Z}_p^*; X_{(\cdot)} \leftarrow \mathbf{g}^{x_{(\cdot)}}$	$\overline{(u, i)} \leftarrow a^{-1}(a) / H_{\text{EL}}^2$
$\mathcal{H} \leftarrow $ aECIES.IM $//H^2_{[0,2)}$	$\begin{array}{ccc} (\omega, j) & g & (\omega) / / \cdot \cdot [1, \infty) \\ y & \leftarrow h(y, j) \cdot Y & \leftarrow \sigma^{y_{u,j}} / / H_{c}^{2} \end{array}$
$\mathcal{H} \leftarrow Fcs(D_1, \mathcal{K}) //H^2_{[2,\infty)}$	$\begin{array}{c} y_{u,j} \leftarrow n(w,j), r_{u,j} \leftarrow g \qquad //\Pi_{[1,\infty)} \\ \text{If } \tilde{T}[a, 7] \neq 1 //H^2 \end{array}$
$q \leftarrow s \ln(\mathbb{Z}_p, \mathbb{Z}_p) //H^2_{[1,\infty)}$	If $I[a, Z] \neq \bot //\Pi_{[0,1]}$ If $i \in I[a] : \text{Data}(m, Z) = \text{Data}(X, Y, L) //H^2$
$g \leftarrow \operatorname{slnj}(D_2, \{0, 1\}^l); h \leftarrow \operatorname{slnj}(D_2, \mathbb{Z}_p^*)$	$\prod j \in I[u] \land \operatorname{PAIR}(\mathbf{g}, \mathbb{Z}) = \operatorname{PAIR}(X_u, Y_{u,j}) / / \Pi_{[1,\infty)}$
$b' \leftarrow \mathcal{A}^{\operatorname{New}, \operatorname{Encap}_b, \operatorname{Decap}, \operatorname{H}}$	Return $T[a, Z]$
Return $b' = b$	Return $\mathcal{H}(a, Z) / / H_{[0,2)}$
$ENCAP_b(u)$	Return $\mathcal{H}(a, \operatorname{PAIR}(\mathbf{g}, Z)) //H_{[2,\infty)}^{\mathbb{Z}}$
$\overline{i \leftarrow i+1}$	$\operatorname{PAIR}(X,Y) //\operatorname{H}^{2}_{[1,\infty)}, \text{ internal}$
$I[u] \leftarrow I[u] \cup \{i\}$	$\overline{x \leftarrow \operatorname{dlog}(X)}; y \leftarrow \operatorname{dlog}(Y)$
$a_{u,i} \leftarrow g(u,i)$	Return $q(xy)$
$y_{u,i} \leftarrow h(u,i)$	Decap(u, a, Y)
$Y_{u,i} \leftarrow \mathbf{g}^{y_{u,i}}$	$\overline{(v, i)} \leftarrow a^{-1}(a)$
$Z_{u,i} \leftarrow X_u^{g_{u,i}}$	(z, j) = g = (z) $u_{u, j} \leftarrow h(u, j)$
$K_{u,i}^{1} \leftarrow \mathcal{H}(a_{u,i}, Z_{u,i}) //H^{2}_{[0,2)}$	If $u = v \land j \in I[u] \land Y = \mathbf{g}^{y_{u,j}}$
$K_{u,i}^1 \leftarrow \mathcal{H}(a_{u,i}, \operatorname{PAIR}(\mathbf{g}, Z_{u,i})) //H^2_{[2,\infty)}$	Return $T[u, a, Y]$
$K^0_{u,i} \leftarrow \mathcal{K}$	$Z \leftarrow Y^{x_u}$
$\tilde{T}[a_{u,i}, Z_{u,i}] \leftarrow K^1_{u,i}$	$K \leftarrow \mathcal{H}(a, Z) //H^2_{[0,2)}$
$T[u, a_{u,i}, Y_{u,i}] \leftarrow K^b_{u,i}$	$K \leftarrow \tilde{\mathcal{H}}(a, \operatorname{PAIR}(\mathbf{g}, Z)) / / H^2_{[2,\infty)}$
Return $((a_{u,i}, Y_{u,i}), K^b_{u,i})$	Return K

Fig. 10. Second set of hybrids H_0^2 through H_2^2 used for proof of Theorem 1. Grey highlighting is used to show the difference between H_1^1 and H_0^2 . Note that PAIR is used internally by other oracles and is not directly accessible to A.

means it was filled in an encapsulation query. Suppose this was the *j*-th query and was to *u*. Then a = g(u, j) must hold, *j* would have been added to I[u], and $Z = X_u^{h(u,j)}$ (so PAIR(\mathbf{g}, Z) = PAIR(X_u, Y)). Thus, $\Pr[\mathsf{H}_0^2] = \Pr[\mathsf{H}_1^2]$.

TRANSITION H_1^2 TO H_2^2 (MAP-THEN-RF). In H_2^2 , the random function \mathcal{H} is replaced by a random function $\tilde{\mathcal{H}}$ from $\mathsf{Fcs}(D_1, \mathcal{K})$ where $D_1 = \{0, 1\}^l \times \mathbb{Z}_p$. We replaced the function \mathcal{H} as $\mathcal{H}(a, Z) = \tilde{\mathcal{H}}(a, \mathsf{PAIR}(\mathbf{g}, Z))$. Then \mathcal{H} is a random function if $\tilde{\mathcal{H}}$ is, $\mathsf{PAIR}(g, .)$ is an injection. Hence, $\Pr[H_1^2] = \Pr[H_2^2]$.

TRANSITION H_2^2 TO H_0^3 . Game H_0^3 is shown in Fig. 11. We have highlighted the ways in which H_0^3 differs from H_2^2 . In H_0^3 , the key $K_{(.,.)}^0$ is assigned the output of a random function \mathcal{E} from $\mathsf{Fcs}(D_1, \mathcal{K})$, instead of sampling it at random. This does not change the adversary's view as *a* never repeats. Hence, $\Pr[H_2^2] = \Pr[H_0^3]$.

TRANSITION H_0^3 TO H_1^3 . In H_1^3 , the table *T* is no longer used. The change is in the DECAP oracle where the If condition evaluates whether the bit *b* is 0. Note that the DECAP oracle always returns K^b . When b = 0, it returns $\mathcal{E}(a, \text{PAIR}(\mathbf{g}, Z))$

Hybrids H^3_{κ} for $0 \leq \kappa \leq 3$	NEW(u)
$b \leftarrow \$ \{0, 1\}$	$\overline{\text{Return }} X_u$
$i \leftarrow 0$	H(a, Z)
$x_{(\cdot)} \leftarrow \mathbb{Z}_p^*; X_{(\cdot)} \leftarrow \mathbf{g}^{x_{(\cdot)}}$	$\frac{1}{(u,i)} \leftarrow a^{-1}(a)$
$\mathcal{H} \leftarrow Fcs(D_1, \mathcal{K})$	$\begin{array}{c} (\alpha, j) & g & (\alpha) \\ u_{u,j} \leftarrow h(u, j); Y_{u,j} \leftarrow g^{y_{u,j}} \end{array}$
$\mathcal{E} \leftarrow \mathbb{F}cs(D_1,\mathcal{K})$	If $j \in I[u] \land \text{PAIR}(\mathbf{g}, Z) = \text{PAIR}(X_u, Y_{u,i}) //\text{H}_{[0,2)}^3$
$q \leftarrow * lnj(\mathbb{Z}_p, \mathbb{Z}_p)$	If PAIR(\mathbf{g}, Z) = PAIR($X_{u}, Y_{u, i}$) //H ³ _{10,00}
$g \leftarrow s \ln j(D_2, \{0, 1\}^\iota);$	Beturn $\tilde{T}[a, Z] / H_{12}^3$
$h \leftarrow s \vdash cs(D_2, \mathbb{Z}_p^*)$	had \leftarrow true $//H^3$
$Dad \leftarrow false / / H_{[3,\infty)}^{m}$	$ABORT //H^3$
$b' \leftarrow \$ \mathcal{A}' \land b' \land $	$\operatorname{Potum}_{\mathcal{I}(\alpha)} \mathcal{I}(\alpha, \operatorname{Pot}_{(\alpha)}, \mathbb{Z}))$
Return $b = b$	Return $\mathcal{H}(a, \operatorname{FAIR}(\mathbf{g}, \mathbb{Z}))$
$\overline{\mathrm{Encap}_b(u)}$	$\underline{\operatorname{PAIR}(X,Y)}$ //Internal
$i \leftarrow i + 1$	$x \leftarrow \operatorname{dlog}(X); y \leftarrow \operatorname{dlog}(Y)$
$I[u] \leftarrow I[u] \cup \{i\} //H^{\mathfrak{s}}_{[0,2)}$	Return $q(xy)$
$a_{u,i} \leftarrow g(u,i)$	$\operatorname{DeCAP}(u, a, Y)$
$y_{u,i} \leftarrow h(u,i)$	$(v,j) \leftarrow g^{-1}(a)$
$Y_{u,i} \leftarrow \mathbf{g}^{g_{u,i}}$	$y_{u,j} \leftarrow h(u,j)$
$\begin{array}{ccc} Z_{u,i} \leftarrow X_{u} \\ K^{1} & \tilde{\mathcal{I}}(z & \text{Datp}(z, Z)) \end{array}$	$Z \leftarrow Y^{x_u}$
$K_{u,i} \leftarrow \mathcal{H}(a_{u,i}, \text{PAIR}(\mathbf{g}, Z_{u,i}))$	If $u = v \land j \in I[u] \land Y = \mathbf{g}^{y_{u,j}} //H^3_{[0,2)}$
$\tilde{\mathbf{X}}_{u,i} \leftarrow \mathcal{E}(\boldsymbol{d}_{u,i}, \operatorname{PAIR}(\mathbf{g}, \boldsymbol{\Delta}_{u,i}))$	If $u = v \land j \in I[u] \land Y = \mathbf{g}^{y_{u,j}} \land b = 0 //H^3_{[1,2)}$
$I [a_{u,i}, Z_{u,i}] \leftarrow K_{\overline{u},i} / / H^{\bullet}_{[0,3)}$	If $u = v \wedge Y = \mathbf{g}^{y_{u,j}} \wedge b = 0 //H^3_{[2,\infty)}$
$T[u, a_{u,i}, Y_{u,i}] \leftarrow K_{u,i}^{\circ} //H_{[0,1)}^{\circ}$	Return $T[u, a, Y] //H^3_{[0,1]}$
Return $((a_{u,i}Y_{u,i}), K_{u,i}^o)$	Return $\mathcal{E}(a, \operatorname{PAIR}(\mathbf{g}, Z)) / / H_{1}^{3}$
	$K \leftarrow \tilde{\mathcal{H}}(a \text{ PAIR}(\boldsymbol{\sigma} Z))$
	Return K

Fig. 11. Third set of hybrids H_0^3 through H_3^3 used for proof of Theorem 1.

which is what was used to compute K^0 in ENCAP. The If condition evaluates to false under two cases

- 1. (a,Y) is a challenge ciphertext and b = 1
- 2. (a,Y) is not a challenge ciphertext.

In both these cases, the DECAP oracle returns $\tilde{\mathcal{H}}(a, \text{PAIR}(\mathbf{g}, Z))$ which is the same as K^1 . Therefore, this modification does not change the view of the adversary and $\Pr[\mathsf{H}_0^3] = \Pr[\mathsf{H}_1^3]$.

TRANSITION H_1^3 TO H_2^3 . In H_2^3 we change the If statements in the H and DECAP oracles. In both places, we remove the check $j \in I[u]$. H_1^3 and H_2^3 differ only under the following events:

1. The adversary makes a DECAP query for (u, a, Y) such that $(u, j) = g^{-1}(a)$ and $Y = \mathbf{g}^{h(u,j)}$ but $j \notin I[u]$. 2. The adversary makes a H query for (a, Z) such that $\text{PAIR}(\mathbf{g}, Z) = \text{PAIR}(X_u, Y)$ and $Y = \mathbf{g}^{h(u,j)}$ (where $(u, j) = g^{-1}(a)$) but $j \notin I[u]$.

To cause either of these events, the adversary must "guess" a point a in the image of g other than the at most q_{ENCAP} such points for which it was given the corresponding ciphertexts by ENCAP. We analyze the probability of this event in H_1^3 where the behaviors of DECAP and H can only depend on the values of $h(u, \cdot)$ and $g(u, \cdot)$ for inputs in I[u]. Hence, \mathcal{A} only learns about $h(u, \cdot)$ and $g(u, \cdot)$ through its queries to ENCAP. At some fixed point in time, let n denote the total number of ENCAP queries that \mathcal{A} has made so far and n_u denote the number of ENCAP queries it has made to user u.

First consider a query DECAP(u, a, Y). For this to differ between the games it must hold that a is in the image of $g(u, \cdot)$. There are q_{ENCAP} total such values, of which the adversary has already seen n_u from ENCAP. This is out of the 2^l values in the codomain, of which the adversary has already seen n from ENCAP. Thus there is a $(q_{\text{ENCAP}} - n_u)/(2^l - n) \leq (q_{\text{ENCAP}} - n_u + n)/2^l \leq 2q_{\text{ENCAP}}/2^l$ chance that the adversary picks such an a. The adversary must additionally have guessed the correct $\mathbf{g}^{h(u,j)}$, which it has an at most 1/(p-1) chance of having done.

Now consider a query H(a, Z). For this to differ between the games it must hold that a is in the image of $g(\cdot, \cdot)$. There are $|\mathcal{U}| \cdot q_{\text{ENCAP}}$ total such values, of which the adversary has already seen n from ENCAP. This is out of the 2^l values in the codomain, of which the adversary has already seen n from ENCAP. Thus there is a $(|\mathcal{U}| \cdot q_{\text{ENCAP}} - n)/(2^l - n) \leq |\mathcal{U}| \cdot q_{\text{ENCAP}}/2^l$ chance that the adversary picks such an a. The adversary must additionally have guessed the correct $X_u^{h(u,j)}$, which it has an at most 1/(p-1) chance of having done.

Applying a union bound over all DECAP and H queries gives the claimed bound $\Pr[\mathsf{H}_1^3] \leq \Pr[\mathsf{H}_2^3] + (2q_{\text{ENCAP}}q_{\text{DECAP}} + |\mathcal{U}| \cdot q_{\text{ENCAP}}q_{\text{H}})/(2^l \cdot (p-1)).$

TRANSITION H_2^3 TO H_3^3 . In H_3^3 , the table \tilde{T} is removed. The change is in the H oracle, wherein if the adversary queries the H oracle with a challenge ciphertext, the H oracle aborts. Using the fundamental lemma of game playing, this probability is bounded as $\Pr[H_3^3] \leq \Pr[H_3^3] + \Pr[\mathsf{bad}]$.

In H_3^3 , the adversary is unable to compute H(a, Z) for challenge ciphertexts using the H oracle or the DECAP oracle. The DECAP oracle returns K_b on all challenge ciphertexts (which is the same value that was returned by the ENCAP oracle) and the H oracle aborts on challenge ciphertexts. Therefore, the adversary's view in H_3^3 is independent of the bit b. Hence, $\Pr[H_3^3] \leq 1/2$.

To bound $\Pr[\mathsf{bad}]$, we construct an adversary \mathcal{B} given in Fig. 12 against the Pair CDH security of \mathbb{G} . We claim that \mathcal{B} perfectly simulates H_3^3 for \mathcal{A} . Note that the challenege ciphertext is computed using \mathcal{B} 's challenge oracle, and the If condition in the H has been replaced with a call to \mathcal{B} 's PAIR oracle. Whenever the flag **bad** is set, \mathcal{B} outputs the corresponding (u, j, Z) and wins the Pair CDH game. Therefore, $\Pr[\mathsf{bad}] \leq \mathsf{Adv}_{\mathbb{G}}^{\mathsf{pcdh}}(\mathcal{B})$.

$\text{Adversary } \mathcal{B}_{pcdh}^{\text{New, Chal}, \text{Pair}}$	SIMENCAP _b (u)
$b \leftarrow \{0, 1\}$	$i \leftarrow i + 1$
$i \leftarrow 0$	$X_u \leftarrow \operatorname{NEW}(u)$
$\tilde{\mathcal{H}} \leftarrow sFcs(D_1, \mathcal{K})$	$a_{u,i} \leftarrow g(u,i)$
$\mathcal{E} \leftarrow Fcs(D_1, \mathcal{K})$	$Y_{u,i} \leftarrow CHAL(u,i)$
$g \leftarrow \operatorname{slnj}(D_2, \{0, 1\}^l)$	$w \leftarrow \operatorname{PAIR}(X_u, Y_{u,i})$
bad \leftarrow false	$K_{u,i}^1 \leftarrow \tilde{\mathcal{H}}(a, w)$
$b' \leftarrow \mathfrak{A}^{\mathrm{O}}$	$K_{u,i}^0 \leftarrow \mathcal{E}(a, w)$
Return \perp	Return $((a_{u,i}Y_{u,i}), K_{u,i}^b)$
$\operatorname{Sim}(a, Z)$	$\operatorname{SIMDECAP}(u, a, Y)$
$\overline{(u,j)} \leftarrow g^{-1}(a)$	$(v,j) \leftarrow g^{-1}(a)$
$X_u \leftarrow \text{NEW}(u)$	$X_u \leftarrow \operatorname{NEW}(u)$
$Y_{u,j} \leftarrow \text{CHAL}(u,j)$	$Y_{u,j} \leftarrow \operatorname{CHAL}(u,j)$
If $\operatorname{PAIR}(X_u, Y_{u,i}) \leftarrow \operatorname{PAIR}(\mathbf{g}, Z)$	If $u = v \wedge Y = Y_{u,j} \wedge b = 0$
bad \leftarrow true	Return $\mathcal{E}(a, \operatorname{PAIR}(X_u, Y_{u,j}))$
OUTPUT(u, j, Z)	$K \leftarrow \tilde{\mathcal{H}}(a, \operatorname{PAIR}(X_u, Y_{u,j}))$
Return $\tilde{\mathcal{H}}(a, \operatorname{PAIR}(\mathbf{g}, Z))$	Return K

Fig. 12. Adversary \mathcal{B} for Theorem 1; $O = \{NEW, SIMENCAP_b, SIMDECAP, SIMH\}.$

4.2 Augmented Cramer-Shoup KEM

AUGMENTED VERSION. In this section we present a memory-tight reduction for an augmented version of the Cramer-Shoup KEM [8]. The augmented Cramer-Shoup key encapsulation mechanism $aCS[\mathbb{G}, \mathcal{K}, l]$ is parameterized by a group $\mathbb{G} = (\mathbf{g}, p, \circ)$, key space \mathcal{K} , and length of the random string, l. The parameters \mathbb{G}, \mathcal{K} , and l are constants for any instantiation of the Cramer-Shoup KEM and hence, we override notation and use aCS and $aCS[\mathbb{G}, \mathcal{K}, l]$ interchangably. We let $aCS.IM = Fcs(\{0, 1\}^l \times \mathbb{G}^2, \mathcal{K})$ and define the scheme as follows.

 $\begin{array}{c|c} \underline{\mathsf{aCS.K}} \\ x \leftarrow \mathbb{s} \mathbb{Z}_p^* \\ ek \leftarrow \mathbf{g}^x \\ dk \leftarrow x \\ \operatorname{Return}(ek, dk) \end{array} \begin{vmatrix} \mathsf{aCS.E}^{\mathcal{H}}(ek) \\ a \leftarrow \mathbb{s} \{0, 1\}^l \\ y \leftarrow \mathbb{s} \mathbb{Z}_p^* \\ Y \leftarrow \mathbf{g}^y \\ Z \leftarrow ek^y \\ K \leftarrow \mathcal{H}(a, Y, Z) \\ \operatorname{Return}((a, Y), K) \end{vmatrix} \begin{vmatrix} \mathsf{aCS.D}^{\mathcal{H}}(dk, (a, Y)) \\ \overline{Z} \leftarrow Y^{dk} \\ K \leftarrow \mathcal{H}(a, Y, Z) \\ \operatorname{Return} K \end{vmatrix}$

OVERVIEW OF EXISTING TECHNIQUES AND ASSOCIATED CHALLENGES. A traditional security reduction for the Cramer-Shoup KEM from the Strong CDH problem in the single-user, single-challenge setting would use l = 0 and the lazy sampling technique to simulate \mathcal{H} as a random oracle. The reduction would maintain a table T to store \mathcal{H} queries and corresponding responses. When the adversary makes a decapsulation query on Y, the reduction would check the table to see if an entry T[Y, Z] exists such that $\operatorname{GAP}(X, Y, Z) = \operatorname{true}$ where $X = \mathbf{g}^x$ is the public key. If the entry exists, it would return the corresponding value, and if it does not exist, the reduction would sample a new uniformly random element K from the key set \mathcal{K} . The reduction would then store $T[Y, _] \leftarrow K$ and return K. The second entry would be filed in the table T when the adversary makes a hash query for (Y, Z) such that $\operatorname{GAP}(X, Y, Z) = \operatorname{true}$. The reduction with the Strong CDH game if it outputs a Z such that $Z = \mathbf{g}^{xy}$, which it does by waiting for the Cramer-Shoup adversary to query its hash oracle on inputs (Y, Z) such that $\operatorname{GAP}(X, Y, Z) = \operatorname{true}$.

Like with ECIES, the random oracle simulation using PRF technique cannot be used here as it is not possible for the reduction to simulate decapsulation queries using the PRF. Bhattacharya avoided this issue using the map-thenprf technique, defining $\mathcal{H}(Y, Z)$ so that when $Z = Y^{dk}$, $\mathcal{H}(Y, Z)$ is computable from ek, Y. This allows properly responding to decapsulation queries when the reduction only has access to Y and cannot compute $Z = Y^{dk}$.

This proof breaks in the multi-challenge setting because it is not clear how to identify and respond to challenge ciphertexts without simply storing them all. Once again, augmentation with a random string a allows encoding the information needed to respond to queries appropriately. Our result is captured by the following theorem. The proof of Theorem 2 is given in the full version [23].

Theorem 2 (Strong CDH \Rightarrow **\$CCA).** Let $\mathbb{G} = (\mathbf{g}, p, \circ)$ be a group of prime order p. Let \mathcal{K} and l be fixed. Define $\mathsf{aCS} = \mathsf{aCS}[\mathbb{G}, \mathcal{K}, l]$.

Let $D_1 = \{0,1\}^l \times \mathbb{G} \times \mathbb{G} \cup \{\star\}$ and $D_2 = \mathcal{U} \times [q_{\text{ENCAP}}]$. Let \mathcal{A} be a mu-\$cca adversary with $\mathbf{Query}(\mathcal{A}) = (q_{\text{NEW}}, q_{\text{ENCAP}}, q_{\text{DECAP}}, q_{\text{H}})$ and assume $2^l > 2q_{\text{ENCAP}}$. We construct a $(\mathsf{Fcs}(D_1, \mathcal{K}) \times \mathsf{Inj}(D_2, \{0,1\}^l))$ -oracle adversary \mathcal{B} such that

$$\mathsf{Adv}_{\mathsf{aCS}}^{\mathsf{mu-\$cca}}(\mathcal{A}) \leqslant 2\mathsf{Adv}_{\mathbb{G}}^{\mathsf{scdh}}(\mathcal{B}) + \frac{q_{\mathrm{ENCAP}}^2}{2^l} + \frac{2q_{\mathrm{ENCAP}}(2q_{\mathrm{DECAP}} + |\mathcal{U}| \cdot q_{\mathrm{H}})}{2^l(p-1)}$$

 $\mathbf{Query}(\mathcal{B}) = (q_{\text{New}}, (q_{\text{ENCAP}} + q_{\text{DECAP}} + q_{\text{H}}), q_{\text{H}})$ $\mathbf{Time}^*(\mathcal{B}) = O(\mathbf{Time}(\mathcal{A})) \text{ and } \mathbf{Mem}^*(\mathcal{B}) = O(\mathbf{Mem}(\mathcal{A})).$

INTUITION. For this proof, we program the random string a as the output of a random injection g applied to a user identity u and a counter i and simulate the random oracle H using the random function $\tilde{\mathcal{H}}(\alpha(a, Y, Z))$ where α is an injection from $\{0, 1\}^l \times \mathbb{G} \times \mathbb{G} \cup \{\star\}$ to $\{0, 1\}^l \times \{0, 1\} \times \mathbb{G}^2$ that can be computed with knowing Z when $Z = Y^{ek}$.

Our reduction adversary plays the Strong CDH game where it gets its challenge ciphertexts Y from the oracle CHAL(u, i). To determine if a decapsulation query (u, a, Y) is for a challenge ciphertext, the reduction first inverts a to obtain (v, j). If u = v, it queries CHAL(u, j) to obtain the ciphertext $Y_{u,j}$. Then it can simply check if $Y = Y_{u,j}$. Finally, when the adversary \mathcal{A} queries the oracle H with (a, Y, Z) such that $g^{-1}(a) = (u, j)$ and STRONG(u, Y, Z) =true and Y = CHAL(u, j), the reduction outputs (u, j, Z) and wins the Strong CDH game.

5 Fujisaki-Okamoto Transformation

The Fujisaki-Okamoto [14, 15] transformations use a random oracle to construct an IND-CCA secure KEM from a weakly (IND-CPA) secure PKE scheme. Bhattacharyya presented memory-tight reductions for the modules analyzed in [21] in the single-user, single-challenge setting [7]. In our work, we use the message encoding technique along with map-then-rf technique to prove memory-tight reductions for one version of the Fujisaki-Okamoto transformations in the multi-user, multi-challenge setting. In the following subsections, we present the definitions and memory-tight reductions for the transformations T, aV, and aU^{\perp} .

5.1 Transformation T [IND-CPA \rightarrow OW-PCA]

The transformation T constructs a deterministic OW-PCA secure public key encryption scheme TKE = T[PKE] from an IND-CPA secure public key encryption scheme PKE. We define TKE as follows with TKE.IM = Fcs(PKE.M, PKE.R) × PKE.IM and TKE.M = PKE.M.

$$\frac{\mathsf{TKE}.\mathsf{E}^{\mathcal{H}\times\mathcal{H}'}(ek,m)}{c \leftarrow \mathsf{PKE}.\mathsf{E}^{\mathcal{H}'}(ek,m;\mathcal{H}(m))} \frac{\mathsf{TKE}.\mathsf{D}^{\mathcal{H}\times\mathcal{H}'}((ek,dk),c)}{m' \leftarrow \mathsf{PKE}.\mathsf{D}^{\mathcal{H}'}(dk,c)}$$

Return c
If $m' = \bot$ or $\mathsf{PKE}.\mathsf{E}^{\mathcal{H}'}(ek,m';\mathcal{H}(m')) \neq c$
Return \bot
Return m'

For key generation, TKE.K samples $(ek, dk) \leftarrow$ PKE.K and outputs (ek, (ek, dk)). Note that TKE is *tidy*, meaning that if m =TKE.D $^{\mathcal{H} \times \mathcal{H}'}((ek, dk), c)$, then c =TKE.E $^{\mathcal{H} \times \mathcal{H}'}(ek, m)$.

We present a memory-tight reduction for T in the multi-user, multi-challenge setting using the randomness programming technique. Our result is captured in Theorem 3, whose proof is given in the full version [23].

Theorem 3 (IND-CPA \Rightarrow **OW-PCA).** Let TKE = T[PKE]. If PKE is δ -correct, then TKE is δ' -correct for $\delta'(q) = (q + 1)\delta(q)$. Let \mathcal{A} be an adversary against TKE with **Query**(\mathcal{A}) = $(q_{\text{NEW}}, q_{\text{CHAL}}, q_{\text{PCO}}, q_{\text{H}})$. Assume PKE's algorithms make at most q_{PKE} oracle queries and define $q^* = q_{\text{H}} + q_{\text{CHAL}}(q_{\text{PKE}} + 1) + q_{\text{PCO}}(2q_{\text{PKE}} + 1) + q_{\text{PKE}}$. We construct an (Fcs(PKE. \mathcal{M} , PKE. \mathcal{R}), $\text{Inj}^{\pm}(\mathcal{U} \times \mathcal{I}, \text{PKE}.\mathcal{M})$)-oracle adversary \mathcal{B} such that

$$\mathsf{Adv}_{\mathsf{TKE}}^{\mathsf{mu-ow-pca}}(\mathcal{A}) \leqslant \mathsf{Adv}_{\mathsf{PKE}}^{\mathsf{mu-cpa}}(\mathcal{B}) + |\mathcal{U}| \cdot \delta'(q^*) + \frac{0.5q_{\mathsf{CHAL}}^2 + |\mathcal{U}||\mathcal{I}|(q_{\mathrm{H}} + q_{\mathrm{PCO}} + 1)}{|\mathsf{PKE}.\mathcal{M}|}$$

 $\mathbf{Query}(\mathcal{B}) = (q_{\text{NEW}}, q_{\text{CHAL}})$ $\mathbf{Time}^*(\mathcal{B}) = O(\mathbf{Time}(\mathcal{A}))$ $\mathbf{Mem}^*(\mathcal{B}) = O(\mathbf{Mem}(\mathcal{A})).$

The notion of CPA security we require is interesting in that the adversary's encryption queries are of the form (u, i, m) where *i* is a "challenge identifier" and if it exactly repeats a query (u, i, m) it is given back the ciphertext from the earlier query.

INTUITION. For this proof, we program the random messages m to be the output of a random injection g applied to user identifier u and counter i. Our reduction adversary simulates the CHAL oracle for \mathcal{A} by its own encryption oracle on g(u, i). If any message that \mathcal{A} queries to its random oracle or outputs at the end of execution is in the image of g, then out reduction assumes it is in the real world and outputs 1. In the ideal world, the view of \mathcal{A} is independent of g so we can information theoretically bound the probability it finds such a message.

5.2 Transformation aV [OW-PCA \rightarrow OW-PCVA]

The augmented transformation aV constructs a deterministic OW-PCVA secure public key encryption scheme VKE = aV[TKE] from a deterministic OW-PCA secure scheme TKE. The unaugmented V transformation was given (with a singleuser, single-challenge memory-tight reduction) in [7]. Our augmentation adds a random string to the encryption key which is included with every hash function query. We define VKE as follows with VKE.IM = Fcs({0,1}^l × TKE. \mathcal{M} , {0,1}^{γ}) × TKE.IM and VKE. \mathcal{M} = TKE. \mathcal{M} , where *l* and γ are fixed.

 $\begin{array}{c|c} \frac{\mathsf{VKE}.\mathsf{E}^{\mathcal{H}\times\mathcal{H}'}((a,ek),m)}{c_1\leftarrow\mathsf{TKE}.\mathsf{E}^{\mathcal{H}'}(ek,m)} & \mathsf{VKE}.\mathsf{D}^{\mathcal{H}\times\mathcal{H}'}((a,ek,dk),c) \\ c_2\leftarrow\mathsf{TKE}.\mathsf{E}^{\mathcal{H}'}(ek,m) \\ c\leftarrow(c_1,c_2) \\ \mathsf{Return}\ c \\ \end{array} \\ \begin{array}{c|c} \mathsf{Keturn}\ \mathcal{L} \\ \mathsf{Return}\ \mathcal{L} \\ \mathsf{Return}\ \mathcal{L} \\ \mathsf{Return}\ \mathcal{L} \end{array} \\ \begin{array}{c|c} \mathsf{VKE}.\mathsf{D}^{\mathcal{H}\times\mathcal{H}'}((a,ek,dk),c) \\ (c_1,c_2)\leftarrow c \\ m'\leftarrow\mathsf{TKE}.\mathsf{D}^{\mathcal{H}'}(dk,c_1) \\ \mathsf{If}\ m'=\bot\ \mathsf{or}\ \mathcal{H}(a,m')\neq c_2\ \mathsf{or}\ \mathsf{TKE}.\mathsf{E}^{\mathcal{H}'}(ek,m')\neq c_1 \\ \mathsf{Return}\ \mathcal{L} \\ \mathsf{Return}\ m' \end{array}$

For key generation, VKE.K samples $a \leftarrow \{0,1\}^l$ and $(ek, dk) \leftarrow TKE.K$, then returns ((a, ek), (a, ek, dk)). Note that aV is tidy and if T is tidy and δ' -correct, then aV is δ' -correct.

We present a memory-tight reduction for aV in the multi-user, multi-challenge setting using the randomness programming technique. Our result is captured in Theorem 4 whose proof is given in the full version of the paper [23].

Theorem 4 (OW-PCA \Rightarrow **OW-PCVA).** Let VKE = aV[TKE] and suppose TKE is δ' -correct. Let \mathcal{A} be an adversary against VKE with $\mathbf{Query}(\mathcal{A}) = (q_{\mathrm{U}}, q_{\mathrm{CHAL}}, q_{\mathrm{PCO}}, q_{\mathrm{CVO}}, q_{\mathrm{H}})$. Assume TKE's algorithms make at most q_{TKE} oracle queries and define $q^* = q_{\mathrm{TKE}}(q_{\mathrm{H}} + q_{\mathrm{CHAL}} + 2q_{\mathrm{PCO}} + 2q_{\mathrm{CVO}})$. We construct an (Fcs($\{0, 1\}^l \times \mathsf{TKE}.\mathcal{M}, \{0, 1\}^\gamma$), Fcs(TKE. $\mathcal{C}, \{0, 1\}^\gamma$), $\mathsf{Inj}^{\pm}(\mathcal{U}, \{0, 1\}^l)$)-oracle adversary \mathcal{B} against TKE such that

 $\begin{aligned} \mathsf{Adv}_{\mathsf{VKE}}^{\mathsf{mu}\text{-}\mathsf{ow}\text{-}\mathsf{pcva}}(\mathcal{A}) &\leq \mathsf{Adv}_{\mathsf{TKE}}^{\mathsf{mu}\text{-}\mathsf{ow}\text{-}\mathsf{pca}}(\mathcal{B}) + 4|\mathcal{U}| \cdot \delta'(q^*) + 0.5|\mathcal{U}|^2/2^l + q_{\mathrm{CVO}}/2^{\gamma} \\ \mathbf{Query}(\mathcal{B}) &= (q_{\mathrm{NEW}}, q_{\mathrm{CHAL}}, q_{\mathrm{PCO}}, q_{\mathrm{H}} \cdot q_{\mathsf{TKE}}) \\ \mathbf{Time}^*(\mathcal{B}) &= O(\mathbf{Time}(\mathcal{A})) \text{ and } \mathbf{Mem}^*(\mathcal{B}) = O(\mathbf{Mem}(\mathcal{A})). \end{aligned}$

5.3 Augmented Transformation aU^{\perp} [OW-PCVA \rightarrow \$IND-CCA]

The transformation aU^{\perp} constructs an IND-CCA secure key encapsulation mechanism $aUEM = aU^{\perp}[VKE]$ from a deterministic, OW-PCVA secure public key

encryption scheme VKE. We define aUEM as follows where aUEM.K = VKE.K and aUEM.IM = $Fcs(\{0,1\}^l \times VKE.\mathcal{M} \times VKE.\mathcal{C},\mathcal{K}) \times VKE.IM$ (\mathcal{K} is an arbitrary set used as the key set of aUEM).

$$\begin{array}{l} \underbrace{ \mathsf{aUEM}.\mathsf{E}^{\mathcal{H}\times\mathcal{H}'}(ek) }{m \leftarrow \mathsf{s}\;\mathsf{VKE}.\mathcal{M}} \\ a \leftarrow \mathsf{s}\;\{0,1\}^l \\ c \leftarrow \mathsf{VKE}.\mathsf{E}^{\mathcal{H}'}(ek,m) \\ K \leftarrow \mathcal{H}(a,m,c) \\ \mathrm{Return}\;((a,c),K) \end{array} \begin{array}{l} \underbrace{ \mathsf{aUEM}.\mathsf{D}^{\mathcal{H}\times\mathcal{H}'}(dk,(a,c)) }{m \leftarrow \mathsf{VKE}.\mathsf{D}^{\mathcal{H}'}(dk,c) \\ \mathrm{If}\;m = \bot \\ \mathrm{Return}\;\bot \\ K \leftarrow \mathcal{H}(a,m,c) \\ \mathrm{Return}\;K \end{array}$$

The following theorem gives our security result.

Theorem 5 (OW-PCVA \Rightarrow **CCA).** Let $\mathsf{aUEM} = \mathsf{aU}^{\perp}[\mathsf{VKE}]$ where VKE is tidy and δ' -correct. Let \mathcal{A} be an adversary against aUEM with $\mathbf{Query}(\mathcal{A}) = (q_{\text{NEW}}, q_{\text{ENCAP}}, q_{\text{DECAP}}, q_{\text{H}})$. Assume VKE 's algorithms make at most q_{VKE} oracle queries and define $q^* = q_{\mathsf{VKE}}(2q_{\mathrm{H}} + q_{\text{ENCAP}} + 2q_{\text{DECAP}})$. Let $D_1 = \{0, 1\}^l \times \mathsf{VKE}.\mathcal{M} \cup \{\star\} \times \mathsf{VKE}.\mathcal{C}$ and $D_2 = \mathcal{U} \times \mathcal{I}$. We construct an $(\mathsf{Inj}^{\pm}(D_2, \{0, 1\}^l), \mathsf{Fcs}(D_1, \mathcal{K}), \mathsf{Fcs}(D_1, \mathcal{K}))$ -oracle adversary \mathcal{B} such that

 $\begin{aligned} \mathsf{Adv}_{\mathsf{aUEM}}^{\mathsf{mu-cca}}(\mathcal{A}) &\leqslant 2\mathsf{Adv}_{\mathsf{VKE}}^{\mathsf{mu-ow-pcva}}(\mathcal{B}) + 2|\mathcal{U}| \cdot \delta'(q^*) + 6|\mathcal{U}| \cdot |\mathcal{I}|/2^l\\ \mathbf{Query}(\mathcal{B}) &= (q_{\mathsf{NEW}}, (q_{\mathsf{ENCAP}} + q_{\mathsf{DECAP}} + q_{\mathsf{H}}), q_{\mathsf{H}}, q_{\mathsf{DECAP}})\\ \mathbf{Time}^*(\mathcal{B}) &= O(\mathbf{Time}(\mathcal{A})) \text{ and } \mathbf{Mem}^*(\mathcal{B}) = O(\mathbf{Mem}(\mathcal{A})). \end{aligned}$

INTUITION. This proof is very similar to the proof of Theorem 2. Once again, we program the random string a to be the output of the injective function g(u, i), and the message m to be the ouput of the random function h(u, i). We use the map-then-rf technique to simulate the oracle H using the random function $\tilde{\mathcal{H}}(\alpha(a, m, c))$ where α is an injective function.

The reduction adversary gets challenge ciphertexts c from the CHAL oracle. To simulate the H oracle it uses the PCO oracle and to simulate the DECAP oracle, it uses its CVO oracle. When the aUEM adversary queries the H oracle with a tuple (a, m, c) such that PCO(u, m, c) = 1 and c = CHAL(u, j) where $(u, j) = g^{-1}(a)$, the reduction outputs (u, j, m).

6 Memory-Tight Reduction for PKE Schemes

In this section, we provide a modified version of the TAM-tight security proof from [17] to show the security of the KEM/DEM construction of public key encryption. Thus, combining one of the KEMs studied in the rest of the paper with an appropriate symmetric encryption scheme gives a PKE scheme with a TAM-tight reduction in the multi-user, multi-challenge setting.

KEM/DEM SCHEME. Let SKE be a symmetric key encryption scheme and KEM be a key encapsulation mechanism. Then the KEM/DEM encryption scheme KD = KD[KEM, SKE] is defined as follows, with $KD.IM = KEM.IM \times SKE.IM$.

The following theorem TAM-tightly proves the security of KD.

Theorem 6. Let SKE be a symmetric key encryption scheme, KEM be a ϵ uniform key encapsulation mechanism, and KD = KD[SKE, KEM]. Let $T' = \mathcal{U} \times \bigcup_{ek \in \mathsf{KEM}.\mathsf{Ek}} \mathsf{KEM}.\mathcal{C}(ek) \times \mathbb{N}$, Let $T = \mathcal{U} \times \mathsf{KEM}.\mathsf{Ek}$, $D_{(u,ek)} = [q_{\mathrm{ENC}}]$ and $R_{(u,ek)} = \mathsf{KEM}.\mathcal{C}(ek)$. $D'_{(u,c^k,l)} = \{0,1\}^l$, and $R'_{(u,c^k,l)} = \{0,1\}^{\mathsf{SKE.cl}(l)}$. Let \mathcal{A} be a mu-\$cca adversary against KD with $\mathbf{Query}(\mathcal{A}) = (q_{\mathrm{NEW}}, q_{\mathrm{ENC}}, q_{\mathrm{DEC}}, q_{\mathrm{H}})$. Then we can construct a (SKE.IM, $\mathsf{Inj}^{\pm}(T', D', R')$)-oracle adversary $\mathcal{B}_{\mathsf{KEM}}$ and (KEM.IM, $\mathsf{Fcs}(\mathcal{U}, \mathsf{KEM}.\mathcal{R}), \mathsf{Inj}^{\pm}(T, D, R)$)-oracle adversary against $\mathcal{B}_{\mathsf{SKE}}$ such that

$$\begin{split} \mathsf{Adv}_{\mathsf{KD}}^{\mathsf{mu}-\$\mathsf{cca}}(\mathcal{A}) &\leqslant 2\mathsf{Adv}_{\mathsf{KEM}}^{\mathsf{mu}-\$\mathsf{cca}}(\mathcal{B}_{\mathsf{KEM}}) + \mathsf{Adv}_{\mathsf{SKE}}^{\mathsf{mu}-\$\mathsf{cca}}(\mathcal{B}_{\mathsf{SKE}}) + q_{\mathrm{ENC}} \cdot \epsilon \\ &+ \frac{1.5q_{\mathrm{ENC}}^2 + 2q_{\mathrm{ENC}}q_{\mathrm{DEC}}}{|\mathsf{KEM}.\mathcal{C}|} + \frac{q_{\mathrm{ENC}}^2 + 2q_{\mathrm{DEC}}}{2^{\mathsf{SKE}.\mathsf{xl}}}. \end{split}$$

The complexities of \mathcal{B}_{KEM} and \mathcal{B}_{SKE} basically match those of \mathcal{A} . Moreover, \mathcal{B}_{SKE} makes at most one encryption query per user.

INSTANTIATING KD. This result proves the multi-challenge, multi-user security of KD, but requires appropriate choices of KEM and SKE. Naturally, one could choose any of the KEMs studied earlier in this work for the first component.⁶ For the symmetric encryption we need a scheme which achieves single-challenge, multi-user security. We are not aware of any TAM-tight multi-user analysis of symmetric encryption scheme, so one instead needs to pick a scheme whose multiuser, single challenge security is sufficiently strong against memory-unbounded adversaries. One reasonable option could be GCM with a random nonce. In the ideal cipher model, Hoang, Tessaro, and Thiruvengadam [20] showed a strong bound for this setting which is essentially independent of the number of users.

References

- Michel Abdalla, Mihir Bellare, and Phillip Rogaway. DHIES: An encryption scheme based on the Diffie-Hellman problem. Contributions to IEEE P1363a, September 1998.
- Michel Abdalla, Mihir Bellare, and Phillip Rogaway. The oracle Diffie-Hellman assumptions and an analysis of DHIES. In David Naccache, editor, CT-RSA 2001, volume 2020 of LNCS, pages 143–158. Springer, Heidelberg, April 2001.
- Benedikt Auerbach, David Cash, Manuel Fersch, and Eike Kiltz. Memory-tight reductions. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 101–132. Springer, Heidelberg, August 2017.

⁶ For using the Fujisaki-Okamoto Transformation, which we proved CCA secure, note that $|\mathsf{Adv}_{\mathsf{KEM}}^{\mathsf{mu-Scca}}(\mathcal{B}_{\mathsf{KEM}}) - \mathsf{Adv}_{\mathsf{KEM}}^{\mathsf{mu-Scca}}(\mathcal{B}_{\mathsf{KEM}})| \leq q_{\mathrm{ENC}} \cdot \epsilon.$

- 4. Balthazar Bauer, Georg Fuchsbauer, and Julian Loss. A classification of computational assumptions in the algebraic group model. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 121–151. Springer, Heidelberg, August 2020.
- Mihir Bellare, Alexandra Boldyreva, and Silvio Micali. Public-key encryption in a multi-user setting: Security proofs and improvements. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 259–274. Springer, Heidelberg, May 2000.
- Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *EU-ROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006.
- Rishiraj Bhattacharyya. Memory-tight reductions for practical key encapsulation mechanisms. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 249–278. Springer, Heidelberg, May 2020.
- Ronald Cramer and Victor Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. Cryptology ePrint Archive, Report 2001/108, 2001. https://eprint.iacr.org/2001/108.
- Wei Dai, Stefano Tessaro, and Xihu Zhang. Super-linear time-memory trade-offs for symmetric encryption. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part III*, volume 12552 of *LNCS*, pages 335–365. Springer, Heidelberg, November 2020.
- Denis Diemert, Kai Gellert, Tibor Jager, and Lin Lyu. Digital signatures with memory-tight security in the multi-challenge setting. In Mehdi Tibouchi and Huaxiong Wang, editors, ASIACRYPT 2021, Part IV, volume 13093 of LNCS, pages 403–433. Springer, Cham., December 2021.
- Itai Dinur. On the streaming indistinguishability of a random permutation and a random function. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020*, *Part II*, volume 12106 of *LNCS*, pages 433–460. Springer, Heidelberg, May 2020.
- Itai Dinur. Tight time-space lower bounds for finding multiple collision pairs and their applications. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020*, *Part I*, volume 12105 of *LNCS*, pages 405–434. Springer, Heidelberg, May 2020.
- Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 33–62. Springer, Heidelberg, August 2018.
- 14. Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In Michael J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 537–554. Springer, Heidelberg, August 1999.
- 15. Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. *Journal of Cryptology*, 26(1):80–101, January 2013.
- Romain Gay, Dennis Hofheinz, Eike Kiltz, and Hoeteck Wee. Tightly CCA-secure encryption without pairings. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 1–27. Springer, Heidelberg, May 2016.
- 17. Ashrujit Ghoshal, Riddhi Ghosal, Joseph Jaeger, and Stefano Tessaro. Hiding in plain sight: Memory-tight proofs via randomness programming. In *EUROCRYPT* 2022, 2022. https://eprint.iacr.org/2021/1409.

- Ashrujit Ghoshal, Joseph Jaeger, and Stefano Tessaro. The memory-tightness of authenticated encryption. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 127–156. Springer, Heidelberg, August 2020.
- Ashrujit Ghoshal and Stefano Tessaro. On the memory-tightness of hashed El-Gamal. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 33–62. Springer, Heidelberg, May 2020.
- Viet Tung Hoang, Stefano Tessaro, and Aishwarya Thiruvengadam. The multi-user security of GCM, revisited: Tight bounds for nonce randomization. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, ACM CCS 2018, pages 1429–1440. ACM Press, October 2018.
- Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 341–371. Springer, Heidelberg, November 2017.
- Dennis Hofheinz and Tibor Jager. Tightly secure signatures and public-key encryption. In Reihaneh Safavi-Naini and Ran Canetti, editors, CRYPTO 2012, volume 7417 of LNCS, pages 590–607. Springer, Heidelberg, August 2012.
- Joseph Jaeger and Akshaya Kumar. Memory-tight multi-challenge security of public-key encryption. Cryptology ePrint Archive (To Appear), 2022. https://eprint.iacr.org/2022/.
- Joseph Jaeger and Stefano Tessaro. Tight time-memory trade-offs for symmetric encryption. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019*, *Part I*, volume 11476 of *LNCS*, pages 467–497. Springer, Heidelberg, May 2019.
- Benoît Libert, Marc Joye, Moti Yung, and Thomas Peters. Concise multi-challenge CCA-secure encryption and signatures with almost tight security. In Palash Sarkar and Tetsu Iwata, editors, ASIACRYPT 2014, Part II, volume 8874 of LNCS, pages 1–21. Springer, Heidelberg, December 2014.
- Ueli M. Maurer. Abstract models of computation in cryptography (invited paper). In Nigel P. Smart, editor, 10th IMA International Conference on Cryptography and Coding, volume 3796 of LNCS, pages 1–12. Springer, Heidelberg, December 2005.
- Tatsuaki Okamoto and David Pointcheval. The gap-problems: A new class of problems for the security of cryptographic schemes. In Kwangjo Kim, editor, *PKC 2001*, volume 1992 of *LNCS*, pages 104–118. Springer, Heidelberg, February 2001.
- Ido Shahaf, Or Ordentlich, and Gil Segev. An information-theoretic proof of the streaming switching lemma for symmetric encryption. In 2020 IEEE International Symposium on Information Theory (ISIT), pages 858–863, 2020.
- Victor Shoup. Lower bounds for discrete logarithms and related problems. In Walter Fumy, editor, *EUROCRYPT'97*, volume 1233 of *LNCS*, pages 256–266. Springer, Heidelberg, May 1997.
- 30. Stefano Tessaro and Aishwarya Thiruvengadam. Provable time-memory trade-offs: Symmetric cryptography against memory-bounded adversaries. In Amos Beimel and Stefan Dziembowski, editors, TCC 2018, Part I, volume 11239 of LNCS, pages 3–32. Springer, Heidelberg, November 2018.
- Yuyu Wang, Takahiro Matsuda, Goichiro Hanaoka, and Keisuke Tanaka. Memory lower bounds of reductions revisited. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 61–90. Springer, Heidelberg, April / May 2018.