Public-Key Encryption Resistant to Parameter Subversion and Its Realization from Efficiently-Embeddable Groups

Benedikt Auerbach[1(✉)], Mihir Bellare[2], and Eike Kiltz[1]

[1] Horst-Görtz Institute for IT Security and Faculty of Mathematics,
Ruhr-University Bochum, Bochum, Germany
{benedikt.auerbach,eike.kiltz}@rub.de
[2] Department of Computer Science and Engineering,
University of California San Diego, 9500 Gilman Drive,
La Jolla, CA 92093, USA
mihir@eng.ucsd.edu

Abstract. We initiate the study of public-key encryption (PKE) schemes and key-encapsulation mechanisms (KEMs) that retain security even when public parameters (primes, curves) they use may be untrusted and subverted. We define a strong security goal that we call ciphertext pseudo-randomness under parameter subversion attack (CPR-PSA). We also define indistinguishability (of ciphertexts for PKE, and of encapsulated keys from random ones for KEMs) and public-key hiding (also called anonymity) under parameter subversion attack, and show they are implied by CPR-PSA, for both PKE and KEMs. We show that hybrid encryption continues to work in the parameter subversion setting to reduce the design of CPR-PSA PKE to CPR-PSA KEMs and an appropriate form of symmetric encryption. To obtain efficient, elliptic-curve-based KEMs achieving CPR-PSA, we introduce efficiently-embeddable group families and give several constructions from elliptic-curves.

1 Introduction

This paper initiates a study of public-key encryption (PKE) schemes, and key-encapsulation mechanisms (KEMs), resistant to subversion of public parameters. We give definitions, and efficient, elliptic-curve-based schemes. As a tool of independent interest, we define efficiently-embeddable group families and construct them from elliptic curves.

PARAMETER SUBVERSION. Many cryptographic schemes rely on some trusted, public parameters common to all users and implementations. Sometimes these are specified in standards. The Oakley primes [39], for example, are a small number of fixed prime numbers widely used for discrete-log-based systems. For ECC (Elliptic Curve Cryptography), the parameters are particular curves. Examples include the P-192, P-224, ... curves from the FIPS-186-4 [38] standard and Ed25519 [16].

There are many advantages to such broad use of public parameters. For example, it saves implementations from picking their own parameters, a task that can be error-prone and difficult to do securely. It also makes key-generation faster and allows concrete-security improvements in the multi-user setting [7]. Recent events indicate, however, that public parameters also bring a risk, namely that they can be *subverted*. The representative example is Dual-EC. We refer to [19] for a comprehensive telling of the story. Briefly, Dual EC was a PRG whose parameters consisted of a description of a cyclic group and two generators of the group. If the discrete logarithm of one generator to base the other were known, security would be compromised. The Snowden revelations indicate that NIST had adopted parameters provided by the NSA and many now believe these parameters had been subverted, allowing the NSA to compromise the security of Dual EC. Juniper's use of Dual EC further underscores the dangers [21].

SECURITY IN THE FACE OF PARAMETER SUBVERSION. DGGJR [26] and BFS [9] initiated the study of cryptography that retains security in the face of subverted parameters, the former treating PRGs and the latter treating NIZKs, where the parameter is the common reference string. In this paper we treat encryption. We define what it means for parameter-using PKE schemes and KEMs to retain security in the face of subversion of their parameters. With regard to schemes, ECC relies heavily on trusted parameters. Accordingly we focus here, providing various efficient elliptic-curve-based schemes that retain security in the face of parameter subversion.

CURRENT MITIGATIONS. In practice, parameters are sometimes specified in a verifiable way, for example derived deterministically (via a public algorithm) from publicly-verifiable coins. The coins could be obtained by applying a hash function like SHA1 to some specified constants (as is in fact done for the FIPS-186-4 curves [38] and in the ECC brainpool project), via the first digits of the irrational number π, or via lottery outcomes [5]. This appears to reduce the possibility of subversion, but BCCHLN [15] indicate that the potential of subverting elliptic curves still remains, so there is cause for caution even in this regard. Also, even if such mechanisms might "work" in some sense, we need definitions to understand what "work" means, and proofs to ensure definitions are met. Our work gives such definitions.

BACKGROUND. A PKE scheme specifies a parameter generation algorithm that returns parameters π, a key-generation algorithm that takes π and returns a public key pk and matching secret key sk, an encryption algorithm that given π, pk and message m returns a ciphertext c, and a decryption algorithm that given π, sk, c recovers m. We denote the classical notions of security by IND—indistinguishability of ciphertexts under chosen-ciphertext attack [8,22]—and PKH—public-key hiding, also called anonymity, this asks that ciphertexts not reveal the public key under which they were created [6]. For KEMs, parameter and key generation are the same, encryption is replaced by encapsulation—it takes π, pk to return an encapsulated key K and a ciphertext c that encapsulates K—and decryption is replaced by decapsulation—given π, sk, c it recov-

Fig. 1. Relations between notions of security. The notions are defined, and the relations hold, for both PKE schemes and KEMs. An arrow A → B is an implication: if a scheme meets A then it also meets B.

ers K. We continue to denote the classical goals by IND—this now asks for indistinguishability of encapsulated keys from random under chosen-ciphertext attack [23]—and PKH. We stress that these classical notions assume *honest parameter generation*, meaning *the parameters are trusted*.

We know that, in this setting, IND PKE is reduced, via hybrid encryption, to IND KEMs and ind-cpa symmetric encryption [23]. To the best of our knowledge, no analogous result exists for PKH.

Mass surveillance activities have made apparent the extent to which privacy can be violated purely by access to meta-data, including who is communicating with whom. PKE and KEMs providing PKH are tools towards systems that do more to hide identities of communicants. We will thus target this goal in the parameter subversion setting as well.

DEFINITIONS AND RELATIONS. For both PKE and KEMs, we formulate a goal called ciphertext pseudorandomness under parameter subversion attack, denoted CPR-PSA. It asks that ciphertexts be indistinguishable from strings drawn randomly from the ciphertext space, even under a chosen-ciphertext attack (CCA). We also extend the above-discussed classical goals to the parameter subversion setting, defining IND-PSA and PKH-PSA. For both PKE (Proposition 1) and KEMs (Proposition 2) we show that CPR-PSA implies both IND-PSA and PKH-PSA. We thus get the relations between the new and classical notions summarized in Fig. 1. (Here CPR is obtained by dropping the PSA in CPR-PSA, meaning it is our definition with honest parameter generation. This extends the notions of [26,37] to allow a CCA.)

We ask whether we can reduce the design of CPR-PSA PKE to the design of CPR-PSA KEMs via hybrid encryption. Proposition 3 says the answer is yes, but, interestingly, requires that the KEM has an extra property of well-distributed ciphertexts that we denote WDC-PSA. (The symmetric encryption scheme is required to have pseudo-random ciphertexts. Such symmetric schemes are easily obtained.) We now have a single, strong target for constructions, namely CPR-PSA+WDC-PSA KEMs. (By the above they imply CPR-PSA PKE, which in turn implies IND-PSA PKE and PKH-PSA PKE.) Our goal thus becomes to build efficient KEMs that are CPR-PSA+WDC-PSA.

Parameter-free schemes. We say that a scheme (PKE or KEM) is parameter free if there are no parameters. (Formally, the parameters are the empty string ε.) Note that a parameter-free scheme that is XXX secure is trivially also XXX-PSA secure. (XXX $\in$ {CPR, IND, PKH}.) This is an important observation, and some of our schemes will indeed be parameter-free, but, as we discuss next, this observation does not trivialize the problem.

Issues and challenges. In an attempt to achieve PSA security through the above observation, we could consider the following simple way to eliminate parameters. Given a XXX-secure parameter-using scheme, build a parameter-free version of it as follows: the new scheme sets its parameters to the empty string; key generation runs the old parameter generation algorithm to get π, then the old key generation algorithm to get pk and sk, setting the new public and secret keys to (π, pk) and (π, sk), respectively; encryption and decryption can then follow the old scheme. This trivial construction, however, has drawbacks along two dimensions that we expand on below: (1) security and (2) efficiency.

With regard to security, the question is, if the old scheme is XXX, is the new one too? (If so, it is also XXX-PSA, since it is parameter free, so we only need to consider the classical notions.) The answer to the question is yes if XXX = IND, but *no if* XXX $\in$ {PKH, CPR}. Imagine, as typical, that the parameters describe a group. Then in the new scheme, different users use different, independent groups. This will typically allow violation of PKH [6]. For example, in the El Gamal KEM, a ciphertext is a group element, so if two users have groups $\mathbb{G}_0$ and $\mathbb{G}_1$, respectively, one can determine which user generated a ciphertext by seeing to which of the two groups it belongs. The same is true for RSA where the group $\mathbb{G}_i = \mathbb{Z}_{N_i}$ is determined by the modulus N_i in the key of user i. Even when the moduli have the same bit length, attacks in [6] show how to violate PKH-security of the simple RSA KEM.

With regard to efficiency, the drawback is that we lose the benefits of parameter-using schemes noted above. In particular, key-generation is less efficient (because it involves parameter generation for the old scheme, which can be costly), and public keys are longer (because they contain the parameters of the old scheme). We aim to retain, as much as possible, the efficiency benefits of parameters while adding resistance to PSA.

BBDP [6] give (1) parameter-free IND+PKH RSA-based PKE schemes and (2) parameter-using discrete-log based IND+PKH PKE schemes. The former, since parameter-free, are IND-PSA+PKH-PSA, but they are not CPR-PSA and they are not as efficient as ECC-based schemes. The latter, while ECC-based and fast, are not secure against PSA.

The open question that emerges is thus to design efficient, ECC-based KEMs that are CPR-PSA+WDC-PSA. The technical challenge is to achieve CPR-PSA (and thus PKH-PSA) even though the groups of different users may be different.

Overview of the approach. We introduce and formalize *efficiently-embeddable group (eeg) families* and identify desirable security properties for them. We give a transform constructing CPR-PSA+WDC-PSA KEMs from secure eeg families. This reduces our task to finding secure eeg families. We

Table 1. Our elliptic curve based CPR-PSA+WDC-PSA KEMs. p denotes the modulus of the field. Efficiency of KE.G is dominated by the sampling time of the curves. Efficiency of KE.E (average, worst case) and KE.D (worst case) is given as the number of exponentiations on the curves. The key size is measured in bits, $k = \lceil |\mathbb{F}_p| \rceil$ being the bit length of the used modulus. For the rejection sampling based constructions, ℓ denotes the cut-off bound. For transform **eegToKE2** and the constructions based on Elligator curves (last two rows) see [4].

Eeg family	Transform	Parameter	Assumption	Efficiency			Key size
				KE.G	KE.E	KE.D	
EG_{twist}	**eegToKE1**	p	sCDH-PSA	t_{TGen}	2, 2	2	$10k$
EG_{twist}	**eegToKE2**	p	CDH-PSA	t_{TGen}	3, 3	3	$12k$
$EG_{twist\text{-}rs}^{\ell}$	**eegToKE1**	—	sCDH-PSA	t_{TGen}	3, $\ell+1$	1	$9k$
$EG_{twist\text{-}rs}^{\ell}$	**eegToKE2**	—	CDH-PSA	t_{TGen}	4, $\ell+2$	2	$11k$
$EG_{twist\text{-}re}$	**eegToKE1**	—	sCDH-PSA	t_{TGen}	3, 3	1	$9k$
$EG_{twist\text{-}re}$	**eegToKE2**	—	CDH-PSA	t_{TGen}	4, 4	2	$11k$
EG_{ell1}^{ℓ}, EG_{ell2}^{ℓ}	**eegToKE1**	p	sCDH-PSA	t_{EllGen}	3, $\ell+1$	1	$6k$
$EG_{ell1\text{-}rs}^{\ell}$, $EG_{ell2\text{-}rs}^{\ell}$	**eegToKE1**	—	sCDH-PSA	t_{EllGen}	5, $\ell+1$	1	$5k$

propose several instantiations of eeg families from elliptic curves with security based on different assumptions. An overview of the resulting KEMs is given in Table 1. We discuss our results in greater detail below.

EFFICIENTLY-EMBEDDABLE GROUP FAMILIES. As described above, having users utilize different groups typically enables linking ciphertexts to the intended receiver and hence violating CPR-PSA. However, certain families of groups allow to efficiently map group elements to a space, which is independent of the particular group of the family. Building on these types of group families it is possible to achieve CPR-PSA secure encryption while still allowing each user to choose his own group.

We formalize the required properties via *efficiently embeddable group families*, a novel abstraction that we believe is of independent interest. An eeg family EG specifies a parameter generation algorithm EG.P sampling parameters to be used by the other algorithms, and a group generation algorithm EG.G sampling a group from the family. Embedding algorithm EG.E embeds elements of the group into some embedding space EG.ES. The group element can be recovered using inversion algorithm EG.I. An important property is that the embedding space only depends on the parameters and in particular not on the used group. Looking ahead, the KEM's public key will contain a group sampled with EG.S and ciphertexts will be embeddings. We require two security properties for EG in order to achieve CPR-PSA+WDC-PSA KEMs. Both assume parameter subversion attacks and are defined with respect to a sampling algorithm EG.S, which samples (not necessarily uniformly distributed) group elements. The first, embedding pseudorandomness (EPR-PSA), is that embeddings of group elements sampled

with EG.S are indistinguishable from uniform. Further we give a definition the strong computational Diffie-Hellman assumption (sCDH-PSA) with respect to EG—an adaption of the interactive assumption introduced in [2] to our setting. It differs from the usual strong computational Diffie-Hellman assumption in two points. The group used for the challenge is sampled using EG.G on a parameter of the adversary's choice and additionally one of the exponents used in the challenge is sampled with sampling algorithm EG.S.

KEY ECAPSULATION MECHANISMS FROM EEG FAMILIES. We provide a transform **eegToKE1** of eeg families to secure KEMs. If the eeg family is both EPR-PSA and sCDH-PSA the resulting KEM is CPR-PSA and WDC-PSA.

KEY ENCAPSULATION FROM WEAKER ASSUMPTIONS. In the full version of this paper [4] we give a second transform **eegToKE2** from eeg families to secure KEMs. It is applicable to eeg families consisting of groups, which order has no small prime factors. Its security is based on the weaker computational Diffie-Hellman assumption (CDH-PSA), i.e. it achieves a CPR-PSA and WDC-PSA KEM under the weaker assumption that EG is both EPR-PSA and CDH-PSA. However, this comes at the cost of larger key size and slower encryption and decryption.

INSTANTIATIONS FROM ELLIPTIC CURVES. We propose several instantiations of eeg families from elliptic curves. It is well known that elliptic curves are not all equal in security. We target elliptic-curve groups over the field $\mathbb{F}_p$ for a large odd prime p since they are less vulnerable to discrete-log-finding attacks than groups over fields of characteristic two [28,40]. While the usage of standardized primes allows for more efficient implementations, several cryptanalysts further suggest that p should be as random as possible for maximal security, see for example Brainpool's RFC on ECC [36]. These constraints make building eeg families more challenging. We offer solutions for both cases. We first identify an eeg family implicitly given in prior work [34,37]. The family consists of curve-twist pairs of elliptic curves. Its embedding space depends on the modulus p of the underlying field, which serves as parameter of the construction.

Building on eeg family $\mathsf{EG}_{\text{twist}}$ we also provide alternatives, which no longer rely on a fixed modulus. The constructions have empty parameters and p is sampled at random in the group generation algorithm. The technical challenge is to still achieve pseudorandom embeddings in an embedding space independent of the group. Our solution $\mathsf{EG}^{\ell}_{\text{twist-rs}}$ achieves this by using rejection sampling with cut-off parameter ℓ. Its embedding space consists of bit strings of length only dependent on the security parameter. The sampling algorithm has a worst-case running time of ℓ exponentiations, but the average cost is two exponentiations independently of ℓ. Eeg family $\mathsf{EG}_{\text{twist-re}}$ uses a range expansion technique from [33] and improves on $\mathsf{EG}^{\ell}_{\text{twist-rs}}$ both in terms of efficiency and security. As in the other construction embeddings are bit strings, but sampling only requires a single exponentiation.

SECURITY OF THE INSTANTIATIONS. We now discuss the security properties of our instantiations in greater detail. An overview is given in Table 2. All of

Table 2. Security of our eeg families. The modulus of the used field is denoted by p. $\Delta_{\text{EPR-PSA}}$ denotes the maximal advantage of an (unbounded) adversary in breaking EPR-PSA. ℓ denotes the cut-off bound used in the construction based on rejection sampling.

Eeg family	Curve type	Parameter	$\Delta_{\text{EPR-PSA}}$	See
$\mathsf{EG}_{\text{twist}}$	Twist	p	0	Sect. 5.2
$\mathsf{EG}_{\text{twist-rs}}^{\ell}$	Twist	—	$(1/2)^{\ell}$	Sect. 5.3
$\mathsf{EG}_{\text{twist-re}}$	Twist	—	0	Sect. 5.4
$\mathsf{EG}_{\text{ell1}}^{\ell}$, $\mathsf{EG}_{\text{ell2}}^{\ell}$	Elligator	p	$(2/3)^{\ell}$	[4]
$\mathsf{EG}_{\text{ell1-rs}}^{\ell}$, $\mathsf{EG}_{\text{ell2-rs}}^{\ell}$	Elligator	—	$(4/5)^{\ell}$	[4]

our constructions achieve EPR-PSA statistically. Embeddings in eeg families $\mathsf{EG}_{\text{twist}}$, and $\mathsf{EG}_{\text{twist-re}}$ are perfectly random, i.e. any (unbounded) adversary has advantage 0 in breaking EPR-PSA. For family $\mathsf{EG}_{\text{twist-rs}}^{\ell}$ the advantage decays exponentially in the cut-off bound ℓ.

Diffie-Hellman problem sCDH-PSA is non standard. It is defined with respect to the eeg family's sampling algorithm and assumes parameter subversion attacks. However, for all of our proposed instantiations we are able to show that sCDH-PSA can be reduced to assumptions, which no longer depend on the sampling algorithms, but use uniformly sampled exponents instead. Considering the parameters of our constructions, they belong to one of two classes. Eeg familiy $\mathsf{EG}_{\text{twist}}$ uses the modulus p as parameter, which might be subject to subversion. Accordingly sCDH-PSA in this case corresponds to the assumption that the adversary's possibility to choose p does not improve its capacities in solving Diffie-Hellman instances on either the curve or its twist for a curve-twist pair sampled from the family. Eeg families $\mathsf{EG}_{\text{twist-rs}}^{\ell}$ and $\mathsf{EG}_{\text{twist-re}}$ serve as more conservative alternatives. They are parameter-free and each user choses his own modulus at random, resulting in the weaker assumption that solving Diffie-Hellman instances over curves sampled with respect to a randomly chosen modulus is hard.

INSTANTIATIONS FROM ELLIGATOR CURVES. In the full version of this paper [4] we provide alternatives to our curve-twist pair based constructions. Eeg families $\mathsf{EG}_{\text{ell1}}^{\ell}$, $\mathsf{EG}_{\text{ell2}}^{\ell}$, $\mathsf{EG}_{\text{ell1-rs}}^{\ell}$ and $\mathsf{EG}_{\text{ell2-rs}}^{\ell}$ make use of the Elligator1 and Elligator2 curves of [17]. $\mathsf{EG}_{\text{ell1}}^{\ell}$ and $\mathsf{EG}_{\text{ell2}}^{\ell}$ were implicitly given in [17] and use the modulus of the underlying field as parameter. Constructions $\mathsf{EG}_{\text{ell1-rs}}^{\ell}$ and $\mathsf{EG}_{\text{ell2-rs}}^{\ell}$ serve as parameter-free alternatives.

RELATED WORK. One might consider generating parameters via a multi-party computation protocol so that no particular party controls the outcome. It is unclear however what parties would perform this task and why one might trust any of them. PKE resistant to parameter subversion provides greater security.

Parameter subversion as we consider it allows the adversary full control of the parameters. This was first considered for NIZKs [9] and (under the term back-

doored) for PRGs [25,26]. Various prior works, in various contexts, considered relaxing the assumptions on parameters in some way [20,30,32,35], but these do not allow the adversary full control of the parameters and thus do not provide security against what we call parameter subversion.

Algorithm-substitution attacks, studied in [3,10–12,24], are another form of subversion, going back to the broader framework of kleptography [43,44]. The cliptography framework of RTYZ [41] aims to capture many forms of subversion. In [42] the same authors consider PKE that retains security in the face of substitution of any of its algorithms, but do not consider parameter subversion.

2   Preliminaries

NOTATION. We let ε denote the empty string. If X is a finite set, we let $x \leftarrow\!\!\text{\$} X$ denote picking an element of X uniformly at random and assigning it to x. All our algorithms are randomized and polynomial time (PT) unless stated otherwise. An adversary is an algorithm. Running time is worst case. If A is an algorithm, we let $y \leftarrow A(x_1, \ldots; r)$ denote running A with random coins r on inputs $x_1, \ldots$ and assigning the output to y. We let $y \leftarrow\!\!\text{\$} A(x_1, \ldots)$ be the result of picking r at random and letting $y \leftarrow A(x_1, \ldots; r)$. We let $[A(x_1, \ldots)]$ denote the set of all possible outputs of A when invoked with inputs $x_1, \ldots$. We use the code based game playing framework of [14]. (See Fig. 3 for an example.) By $\Pr[\text{G}]$ we denote the probability that the execution of game G results in the game returning true. We also adopt the convention that the running time of an adversary refers to the worst case execution time of the game with the adversary. This means that the time taken for oracles to compute replies to queries is included. The random oracle model [13] is captured by a game procedure RO that implements a variable output length random oracle. It takes a string x and an integer m and returns a random m-bit string. We denote by $\mathcal{P}_k$ the set of primes of bit length k and by $[d]$ the set $\{0, \ldots, d-1\}$. Furthermore, the uniform distribution on M is denoted by U_M. If two random variables X and Y are equal in distribution we write $X \sim Y$. The statistical distance between X and Y is denoted by $\Delta(X; Y)$. If $\Delta(X; Y) \leq \delta$ we say X is δ-close to Y.

3   Public-Key Encryption Resistant to Parameter Subversion

In this section we recall public-key encryption schemes and key encapsulation mechanisms. For both primitives we define the strong security notion of pseudorandomness of ciphertexts in the setting of parameter subversion and show that it implies both indistinguishability of encryptions and public-key hiding. We further define the security notion of well-distributedness of ciphertexts for key encapsulation mechanisms. Finally, we recall symmetric encryption schemes and revisit the hybrid encryption paradigm in the setting of ciphertext pseudorandomness under parameter subversion attacks.

3.1 Public-Key Encryption Schemes

Below we give a syntax for public-key encryption schemes. It follows [23], but uses slightly different notation and includes an additional algorithm setting up global parameters to be utilized by all users. We then formalize a novel security requirement of pseudorandomness of ciphertexts under parameter subversion attacks (CPR-PSA), which says that even if the parameters of the scheme are controlled by the adversary, ciphertexts obtained under any public key are indistinguishable from random elements of the ciphertext space, which depends only on the security parameter, the message length and the global parameters. We then recall two existing requirements of public-key encryption schemes adapting them to the setting of parameter subversion attacks. The first is the well-known notion of indistinguishability of encryptions [31], the second, from [1,6], is that ciphertexts under different public keys are indistinguishable, which they called anonymity or key hiding and we call public-key hiding. In Proposition 1 we show that the first requirement implies the other two, allowing us to focus on it subsequently. We model the possibility of subverted parameters by having the adversary provide the parameters, which are used in the security games.

PUBLIC-KEY ENCRYPTION. A *public-key encryption scheme* (PKE) PE specifies the following. Parameter generation algorithm PE.P takes input 1^k, where $k \in \mathbb{N}$ is the security parameter, and returns global parameters π. Key-generation algorithm PE.G takes input $1^k, \pi$ and returns a tuple (pk, sk) consisting of the public (encryption) key pk and matching secret (decryption) key sk. PE.CS associates to k, π and message length $m \in \mathbb{N}$ a finite set $\mathsf{PE.CS}(k, \pi, m)$ that is the *ciphertext space* of PE. Encryption algorithm PE.E takes $1^k, \pi, pk$ and a message $M \in \{0,1\}^*$ and returns a ciphertext $c \in \mathsf{PE.CS}(k, \pi, |M|)$. Deterministic decryption algorithm PE.D takes $1^k, \pi, sk$ and a ciphertext c and returns either a message $M \in \{0,1\}^*$ or the special symbol $\bot$ indicating failure. The correctness condition requires that for all $k \in \mathbb{N}$, all $\pi \in [\mathsf{PE.P}(1^k)]$, all $(pk, sk) \in [\mathsf{PE.G}(1^k, \pi)]$ and all $M \in \{0,1\}^*$ we have $\Pr\big[\mathsf{PE.D}(1^k, \pi, sk, c) = M\big] \geq 1 - \mathsf{PE.de}(k)$, where the probability is over $c \leftarrow_\$ \mathsf{PE.E}(1^k, \pi, pk, M)$ and $\mathsf{PE.de} : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ is the *decryption error* of PE. Our PKEs will be in the ROM [13], which means the encryption and decryption algorithms have access to a random oracle specified in the security games. Correctness must then hold for all choices of the random oracle. We say a PKE is *parameter-free* if PE.P returns ε on every input 1^k.

CIPHERTEXT PSEUDORANDOMNESS. Consider game $\mathbf{G}^{\text{cpr-psa}}_{\mathsf{PE},\mathcal{A}}(k)$ of Fig. 2 associated to PKE PE, adversary $\mathcal{A}$ and security parameter k, and let

$$\mathbf{Adv}^{\text{cpr-psa}}_{\mathsf{PE},\mathcal{A}}(k) = 2\Pr[\mathbf{G}^{\text{cpr-psa}}_{\mathsf{PE},\mathcal{A}}(k)] - 1.$$

We say that PE has pseudorandom ciphertexts under parameter subversion attacks (also called CPR-PSA) if the function $\mathbf{Adv}^{\text{cpr-psa}}_{\mathsf{PE},\mathcal{A}}(\cdot)$ is negligible for every $\mathcal{A}$. In the game, b is a challenge bit. When $b = 1$, the challenge ciphertext c^* is an encryption of a message of the adversary's choice, but if $b = 0$ it is chosen at random from the ciphertext space. Given the public key and challenge ciphertext, the adversary outputs a guess b' and wins if b' equals b, the game

Games $\mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{cpr\text{-}psa}}(k), \mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{ind\text{-}psa}}(k), \mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}(k)$

$c^* \leftarrow \bot$
$b \leftarrow_\$ \{0,1\}$
$b' \leftarrow_\$ \mathcal{A}^{\mathrm{INIT,ENC,DEC,RO}}(1^k)$
Return $(b = b')$

$\underline{\mathrm{RO}(x,m)} \,//\, \mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{cpr\text{-}psa}}, \mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{ind\text{-}psa}}, \mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}$
If $(T[x,m] = \bot)$
    then $T[x,m] \leftarrow_\$ \{0,1\}^m$
Return $T[x,m]$

$\underline{\mathrm{ENC}(M)} \,//\, \mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{cpr\text{-}psa}}$
If $(pk = \bot)$ then return $\bot$
If $(b = 0)$ then $c^* \leftarrow_\$ \mathsf{PE.CS}(k,\pi,|M|)$
Else $c^* \leftarrow_\$ \mathsf{PE.E}^{\mathrm{RO}}(1^k,\pi,pk,M)$
Return c^*

$\underline{\mathrm{ENC}(M_0,M_1)} \,//\, \mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{ind\text{-}psa}}$
If $(pk = \bot)$ then return $\bot$
If $(|M_0| \neq |M_1|)$ then return $\bot$
$c^* \leftarrow_\$ \mathsf{PE.E}^{\mathrm{RO}}(1^k,\pi,pk,M_b)$
Return c^*

$\underline{\mathrm{ENC}(M)} \,//\, \mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}$
If $(pk_0 = \bot \vee pk_1 = \bot)$
    return $\bot$
$c^* \leftarrow_\$ \mathsf{PE.E}^{\mathrm{RO}}(1^k,\pi,pk_b,M)$
Return c^*

$\underline{\mathrm{INIT}(\pi)} \,//\, \mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{cpr\text{-}psa}}, \mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{ind\text{-}psa}}$
$(pk, sk) \leftarrow_\$ \mathsf{PE.G}(1^k,\pi)$
Return pk

$\underline{\mathrm{INIT}(\pi)} \,//\, \mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}$
$(pk_0, sk_0) \leftarrow_\$ \mathsf{PE.G}(1^k,\pi)$
$(pk_1, sk_1) \leftarrow_\$ \mathsf{PE.G}(1^k,\pi)$
If $(pk_0 = \bot \vee pk_1 = \bot)$
    return $\bot$
Return (pk_0, pk_1)

$\underline{\mathrm{DEC}(c)} \,//\, \mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{cpr\text{-}psa}}, \mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{ind\text{-}psa}}$
If $(c = c^*)$ then return $\bot$
Else return $\mathsf{PE.D}^{\mathrm{RO}}(1^k,\pi,sk,c)$

$\underline{\mathrm{DEC}(c)} \,//\, \mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}$
If $(c = c^*)$ then return $\bot$
$M_0 \leftarrow \mathsf{PE.D}^{\mathrm{RO}}(1^k,\pi,sk_0,c)$
$M_1 \leftarrow \mathsf{PE.D}^{\mathrm{RO}}(1^k,\pi,sk_1,c)$
Return (M_0, M_1)

Fig. 2. Games defining security of PKEs. In each game the adversary is given access to oracles. The game, to which an oracle belongs, is indicated behind the oracle's name. In each game oracles INIT and ENC may be queried only once. Further INIT has to be queried before using any of the other oracles.

returning true in this case and false otherwise. The adversary has access to an oracle INIT, which sets up the public key using parameters of the adversary's choice, and an oracle ENC to generate the challenge ciphertext. Furthermore it has access to the random oracle and a decryption oracle crippled to not work on the challenge ciphertext. We require that the adversary queries the oracles INIT and ENC only once. Furthermore INIT has to be queried before using any of the other oracles.

INDISTINGUISHABILITY OF ENCRYPTIONS. Consider game $\mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{ind\text{-}psa}}(k)$ of Fig. 2 associated to PKE PE, adversary $\mathcal{A}$ and security parameter k, and let

$$\mathbf{Adv}_{\mathsf{PE},\mathcal{A}}^{\mathrm{ind\text{-}psa}}(k) = 2\Pr[\mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{ind\text{-}psa}}(k)] - 1.$$

We say that PE has indistinguishable encryptions under parameter subversion attacks (also called IND-PSA) if the function $\mathbf{Adv}_{\mathsf{PE},\mathcal{A}}^{\mathrm{ind\text{-}psa}}(\cdot)$ is negligible for every

$\mathcal{A}$. In the game, b is a challenge bit. The adversary has access to an oracle INIT, which sets up the public key using parameters of the adversary's choice, and an oracle ENC, which receives as input two messages M_0, M_1 of the same length and outputs the challenge ciphertext c^*. When $b = 0$, the challenge ciphertext is an encryption of M_0, if $b = 1$ an encryption of M_1. Given the public key and challenge ciphertext, the adversary outputs a guess b' and wins if b' equals b, the game returning true in this case and false otherwise. Again, the adversary has access to the random oracle and a decryption oracle crippled to not work on the challenge ciphertext. We require that the adversary queries the oracles INIT and ENC only once. Furthermore INIT has to be queried before using any of the other oracles.

PUBLIC-KEY HIDING. Consider game $\mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}(k)$ of Fig. 2 associated to PKE PE, adversary $\mathcal{A}$ and security parameter k, and let

$$\mathbf{Adv}_{\mathsf{PE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}(k) = 2\Pr[\mathbf{G}_{\mathsf{PE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}(k)] - 1.$$

We say that PE is public-key hiding under parameter subversion attacks (also called PKH-PSA) if the function $\mathbf{Adv}_{\mathsf{PE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}(\cdot)$ is negligible for every $\mathcal{A}$. In the game, b is a challenge bit. Unlike the prior games, two key pairs are generated, not one. The challenge ciphertext c^* is an encryption of a message of the adversary's choice under pk_b. Given the public keys and the challenge ciphertext, the adversary outputs a guess b' and wins if b' equals b. This time the crippled decryption oracle returns decryptions under both secret keys. The adversary sets up the public keys with its call to oracle INIT, and an uses oracle ENC to generate the challenge ciphertext. Again we require that the adversary queries the oracles INIT and ENC only once. Furthermore INIT has to be queried before using any of the other oracles.

RELATIONS. The following says that pseudorandomness of ciphertexts implies both indistinguishable encryptions and anonymity. We give both asymptotic and concrete statements of the results.

Proposition 1. *Let* PE *be a PKE that has pseudorandom ciphertexts under parameter subversion attacks. Then:*

1. PE *is IND-PSA. Concretely, given an adversary $\mathcal{A}$ the proof specifies an adversary $\mathcal{B}_0$ such that $\mathbf{Adv}_{\mathsf{PE},\mathcal{A}}^{\mathrm{ind\text{-}psa}}(k) \leq 2 \cdot \mathbf{Adv}_{\mathsf{PE},\mathcal{B}_0}^{\mathrm{cpr\text{-}psa}}(k)$ for every $k \in \mathbb{N}$, and $\mathcal{B}_0$ has the same running time and query counts as $\mathcal{A}$.*
2. PE *is PKH-PSA. Concretely, given an adversary $\mathcal{A}$ the proof specifies an adversary $\mathcal{B}_1$ such that $\mathbf{Adv}_{\mathsf{PE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}(k) \leq 2 \cdot \mathbf{Adv}_{\mathsf{PE},\mathcal{B}_1}^{\mathrm{cpr\text{-}psa}}(k)$ for every $k \in \mathbb{N}$, and $\mathcal{B}_0$ has the same running time and query counts as $\mathcal{A}$.*

The proof of the proposition can be found in the full version of this paper [4].

3.2 Key Encapsulation Mechanisms

Below we first give a syntax for key encapsulation mechanisms. It follows [23] but with notation a bit different and including an additional algorithm setting

Game $\mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\text{cpr-psa}}(k)$

$b \leftarrow_\$ \{0,1\}$
$b' \leftarrow_\$ \mathcal{A}^{\text{INIT,DEC,RO}}(1^k)$
Return $(b = b')$

Game $\mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\text{ind-psa}}(k)$

$b \leftarrow_\$ \{0,1\}$
$b' \leftarrow_\$ \mathcal{A}^{\text{INIT,DEC,RO}}(1^k)$
Return $(b = b')$

Game $\mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\text{pkh-psa}}(k)$

$b \leftarrow_\$ \{0,1\}$
$b' \leftarrow_\$ \mathcal{A}^{\text{INIT,DEC,RO}}(1^k)$
Return $(b = b')$

$\underline{\text{DEC}(c)} \;//\; \mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\text{cpr-psa}}, \mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\text{ind-psa}}$

If $(c = c^*)$ then return $\perp$
$K \leftarrow \mathsf{KE.D}^{\text{RO}}(1^k, \pi, sk, c)$
Return K

$\underline{\text{DEC}(c)} \;//\; \mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\text{pkh-psa}}$

If $(c = c^*)$ then return $\perp$
$K_0 \leftarrow \mathsf{KE.D}^{\text{RO}}(1^k, \pi, sk_0, c)$
$K_1 \leftarrow \mathsf{KE.D}^{\text{RO}}(1^k, \pi, sk_1, c)$
Return (K_0, K_1)

$\underline{\text{RO}(x, m)} \;//\; \mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\text{cpr-psa}}, \mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\text{ind-psa}}, \mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\text{pkh-psa}}$

If $(T[x,m] = \perp)$
    then $T[x,m] \leftarrow_\$ \{0,1\}^m$
Return $T[x,m]$

$\underline{\text{INIT}(\pi)} \;//\; \mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\text{cpr-psa}}$

$(pk, sk) \leftarrow_\$ \mathsf{KE.G}(1^k, \pi)$
If $(pk = \perp)$ then return $\perp$
If $(b = 1)$ then $(K^*, c^*) \leftarrow_\$ \mathsf{KE.E}^{\text{RO}}(1^k, \pi, pk)$
Else $K^* \leftarrow_\$ \mathsf{KE.KS}(k)$
$c^* \leftarrow_\$ \mathsf{KE.CS}(k, \pi)$
Return (pk, K^*, c^*)

$\underline{\text{INIT}(\pi)} \;//\; \mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\text{ind-psa}}$

$(pk, sk) \leftarrow_\$ \mathsf{KE.G}(1^k, \pi)$
If $(pk = \perp)$ then return $\perp$
$(K^*, c^*) \leftarrow_\$ \mathsf{KE.E}^{\text{RO}}(1^k, \pi, pk)$
If $(b = 0)$ then $K^* \leftarrow_\$ \mathsf{KE.KS}(k)$
Return (pk, K^*, c^*)

$\underline{\text{INIT}(\pi)} \;//\; \mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\text{pkh-psa}}$

$(pk_0, sk_0) \leftarrow_\$ \mathsf{KE.G}(1^k, \pi)$
$(pk_1, sk_1) \leftarrow_\$ \mathsf{KE.G}(1^k, \pi)$
If $(pk_0 = \perp \vee pk_1 = \perp)$ then return $\perp$
$(K^*, c^*) \leftarrow_\$ \mathsf{KE.E}^{\text{RO}}(1^k, \pi, pk_b)$
Return (pk_0, pk_1, K^*, c^*)

Fig. 3. Games defining security of key encapsulation mechanism KE. In each game the adversary is given access to oracles. The game, to which an oracle belongs, is indicated behind the oracle's name. In each game oracle INIT must be queried only once, which has to be done before using any of the other oracles.

up global parameters to be utilized by all users. As for public-key encryption schemes we formalize the security requirement of pseudorandomness of cipher-texts under parameter subversion attacks (CPR-PSA). We then adapt the two existing KEM requirements of indistinguishability of encryptions [23] and public-key hiding [1,6] to the setting of parameter subversion attacks. In Proposition 2 we show that—as in the case of public-key encryption—the first requirement implies the other two. We furthermore define a new security requirement called well-distributedness of ciphertexts, which is necessary to achieve CPR-PSA in the hybrid PKE construction. It states that key-ciphertext pairs generated using the KEM's encapsulation algorithm are indistinguishable from choosing a cipher-text at random and then computing its decapsulation.

KEMs. A *key encapsulation mechanism* (KEM) KE specifies the following. Parameter generation algorithm $\mathsf{KE.P}$ takes input 1^k, where $k \in \mathbb{N}$ is the security parameter, and returns global parameters π. Key-generation algorithm $\mathsf{KE.G}$

takes input $1^k, \pi$ and returns a tuple (pk, sk) consisting of the public (encryption) key pk and matching secret (decryption) key sk. KE.KS associates to k a finite set KE.KS(k) only depending on the security parameter that is the *key space* of KE. KE.CS associates to k and parameters π a finite set KE.CS(k, π) that is the *ciphertext space* of KE. Encapsulation algorithm KE.E takes $1^k, \pi, pk$ and returns (K, c) where $K \in$ KE.KS(k) is the *encapsulated key* and $c \in$ KE.CS(k, π) is a ciphertext encapsulating K. Deterministic decapsulation algorithm KE.D takes $1^k, \pi, sk$ and a ciphertext c and returns either a key $K \in$ KE.KS(k) or the special symbol $\bot$ indicating failure. The correctness condition requires that for all $k \in \mathbb{N}$, all $\pi \in$ [KE.P(1^k)] and all $(pk, sk) \in$ [KE.G$(1^k, \pi)$] we have $\Pr\left[\text{KE.D}(1^k, \pi, sk, c) = K\right] \geq 1 - \text{KE.de}(k)$, where the probability is over $(K, c) \leftarrow_\$ \text{KE.E}(1^k, \pi, pk)$ and KE.de $: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ is the *decryption error* of KE. Our KEMs will be in the ROM [13], which means the encapsulation and decapsulation algorithms have access to a random oracle specified in the security games. Correctness must then hold for all choices of the random oracle. We say a KEM is *parameter-free* if KE.P returns ε on every input 1^k.

CIPHERTEXT PSEUDORANDOMNESS. Consider game $\mathbf{G}_{\text{KE},\mathcal{A}}^{\text{cpr-psa}}(k)$ of Fig. 3 associated to KEM KE, adversary $\mathcal{A}$ and security parameter k, and let

$$\mathbf{Adv}_{\text{KE},\mathcal{A}}^{\text{cpr-psa}}(k) = 2\Pr[\mathbf{G}_{\text{KE},\mathcal{A}}^{\text{cpr-psa}}(k)] - 1.$$

We say that KE has pseudorandom ciphertexts under parameter subversion attacks (also called CPR-PSA) if the function $\mathbf{Adv}_{\text{KE},\mathcal{A}}^{\text{cpr-psa}}(\cdot)$ is negligible for every $\mathcal{A}$. In the game, b is a challenge bit. When $b = 1$, the challenge key K^* and ciphertext c^* are generated via the encapsulation algorithm, but if $b = 0$ they are chosen at random, from the key space and ciphertext space, respectively. Given the public key, challenge key and challenge ciphertext, the adversary outputs a guess b' and wins if b' equals b, the game returning true in this case and false otherwise. The adversary has access to an oracle INIT, which sets up the challenge. We require that the adversary queries INIT before using any of the other oracles and that it queries INIT only once. Further the adversary has access to an oracle for decapsulation under sk, crippled to not work when invoked on the challenge ciphertext. It, and the encapsulation and decapsulation algorithms, have access to the random oracle RO. The parameters used in the game are provided by the adversary via its call to INIT.

INDISTINGUISHABILITY OF ENCAPSULATED KEYS FROM RANDOM. Consider game $\mathbf{G}_{\text{KE},\mathcal{A}}^{\text{ind-psa}}(k)$ of Fig. 3 associated to KEM KE, adversary $\mathcal{A}$ and security parameter k, and let

$$\mathbf{Adv}_{\text{KE},\mathcal{A}}^{\text{ind-psa}}(k) = 2\Pr[\mathbf{G}_{\text{KE},\mathcal{A}}^{\text{ind-psa}}(k)] - 1.$$

We say that KE has encapsulated keys indistinguishable from random under parameter subversion attacks (also called IND-PSA) if the function $\mathbf{Adv}_{\text{KE},\mathcal{A}}^{\text{ind-psa}}(\cdot)$ is negligible for every $\mathcal{A}$. In the game, b is a challenge bit. When $b = 1$, the challenge key K^* and ciphertext c^* are generated via the encapsulation algorithm, while if $b = 0$ the key is switched to one drawn randomly from the key space,

the ciphertext remaining real. Given the public key, challenge key and challenge ciphertext, the adversary outputs a guess b' and wins if b' equals b. Again the adversary has access to a crippled decapsulation oracle and the random oracle and provides the parameters used in the game via his call to the oracle INIT, which has to be queried before using any of the other oracles.

PUBLIC-KEY HIDING. Consider game $\mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}(k)$ of Fig. 3 associated to KEM KE, adversary $\mathcal{A}$ and security parameter k, and let

$$\mathbf{Adv}_{\mathsf{KE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}(k) = 2\Pr[\mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}(k)] - 1.$$

We say that KE is public-key hiding under parameter subversion attacks (also called PKH-PSA) if the function $\mathbf{Adv}_{\mathsf{KE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}(\cdot)$ is negligible for every $\mathcal{A}$. In the game, b is a challenge bit. Unlike the prior games, two key pairs are generated, not one. The challenge key K^* and ciphertext c^* are generated via the encapsulation algorithm under pk_b. Given the public keys, challenge key and challenge ciphertext, the adversary outputs a guess b' and wins if b' equals b. This time the crippled decapsulation oracle returns decapsulations under both secret keys. Again the adversary provides the parameters to be used in the game via his single call to the oracle INIT, which has to be queried before using any of the other oracles.

RELATIONS. The following says that in the parameter subversion setting CPR-PSA implies both IND-PSA and PKH-PSA. We give both the asymptotic and concrete statements of the results.

Proposition 2. *Let* KE *be a KEM that has pseudorandom ciphertexts under parameter subversion attacks. Then:*

1. KE *is* IND-PSA*. Concretely, given an adversary $\mathcal{A}$ the proof specifies an adversary $\mathcal{B}$ such that $\mathbf{Adv}_{\mathsf{KE},\mathcal{A}}^{\mathrm{ind\text{-}psa}}(k) \le 2 \cdot \mathbf{Adv}_{\mathsf{KE},\mathcal{B}}^{\mathrm{cpr\text{-}psa}}(k)$ for every $k \in \mathbb{N}$, and $\mathcal{B}$ has the same running time and query counts as $\mathcal{A}$.*
2. KE *is* PKH-PSA*. Concretely, given an adversary $\mathcal{A}$ the proof specifies an adversary $\mathcal{B}$ such that $\mathbf{Adv}_{\mathsf{KE},\mathcal{A}}^{\mathrm{pkh\text{-}psa}}(k) \le 2 \cdot \mathbf{Adv}_{\mathsf{KE},\mathcal{B}}^{\mathrm{cpr\text{-}psa}}(k)$ for every $k \in \mathbb{N}$, and $\mathcal{B}$ has the same running time and query counts as $\mathcal{A}$.*

The proof of the proposition can be found in the full version of this paper [4].

WELL-DISTRIBUTED CIPHERTEXTS. Consider game $\mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\mathrm{wdc\text{-}psa}}(k)$ of Fig. 4 associated to KEM KE, adversary $\mathcal{A}$ and security parameter k, and let

$$\mathbf{Adv}_{\mathsf{KE},\mathcal{A}}^{\mathrm{wdc\text{-}psa}}(k) = 2\Pr[\mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\mathrm{wdc\text{-}psa}}(k)] - 1.$$

We say KE has well distributed ciphertexts under parameter subversion attacks (also called WDC-PSA), if the function $\mathbf{Adv}_{\mathsf{KE},\mathcal{A}}^{\mathrm{wdc\text{-}psa}}(\cdot)$ is negligible for every adversary $\mathcal{A}$. In the game b is a challenge bit. If b equals 1 the adversary as response to querying the initialization procedure, which may be done at most once, receives a key-ciphertext pair generated using KE.E. If b equals 0 it receives a pair (c^*, K^*) generated by choosing c^* at random and then setting K^* to be the decapsulation of c^*. The adversary has access to a decryption oracle. We require that the adversary queries INIT before querying any of the other oracles. Looking ahead, all of our instantiations achieve this notion statistically.

Game $\mathbf{G}_{\mathsf{KE},\mathcal{A}}^{\mathrm{wdc\text{-}psa}}(k)$

$b \leftarrow\!\!\$ \{0,1\}$
$b' \leftarrow\!\!\$ \mathcal{A}^{\mathrm{INIT,DEC,RO}}(1^k)$
Return $(b = b')$

$\underline{\mathrm{INIT}(\pi)}$

$(pk, sk) \leftarrow\!\!\$ \mathsf{KE.G}(1^k, \pi)$
If $(pk = \bot)$ then return $\bot$
If $(b = 1)$ then $(K^*, c^*) \leftarrow\!\!\$ \mathsf{KE.E}^{\mathrm{RO}}(1^k, \pi, pk)$
Else $c^* \leftarrow\!\!\$ \mathsf{KE.CS}(k, \pi)$
$K^* \leftarrow \mathsf{KE.D}^{\mathrm{RO}}(1^k, \pi, sk, c^*)$
Return (pk, K^*, c^*)

$\underline{\mathrm{RO}(x, m)}$

If $(T[x, m] = \bot)$
    then $T[x, m] \leftarrow\!\!\$ \{0,1\}^m$
Return $T[x, m]$

$\underline{\mathrm{DEC}(c)}$

If $(c = c^*)$ then return $\bot$
$K \leftarrow \mathsf{KE.D}^{\mathrm{RO}}(1^k, \pi, sk, c)$
Return K

Fig. 4. Game defining well-distributedness of ciphertexts of KEs.

3.3   Symmetric Encryption

Below, we recall symmetric encryption. Our definition follows [23] but uses different notation. We further define the security notion of ciphertext pseudorandomness for symmetric key encryption.

ONE-TIME SYMMETRIC-KEY ENCRYPTION. A symmetric-key encryption scheme (SKE) specifies the following. $\mathsf{SE.KS}$ associates to security parameter k key space $\mathsf{SE.KS}(k)$. $\mathsf{SE.CS}$ associates to security parameter k and message length $m \in \mathbb{N}$ the ciphertext space $\mathsf{SE.CS}(k, m)$. Deterministic encryption algorithm $\mathsf{SE.E}$ takes as input 1^k, key $K \in \mathsf{SE.KS}(k)$ and a message $M \in \{0,1\}^*$ and returns ciphertext $c \in \mathsf{SE.CS}(k, |M|)$. Deterministic decryption algorithm $\mathsf{SE.D}$ on input $1^k, K \in \mathsf{SE.KS}(k), c \in \mathsf{SE.CS}(k, m)$ returns either a message $M \in \{0,1\}^m$ or the special symbol $\bot$ indicating failure. For correctness we require that $M = \mathsf{SE.D}(1^k, K, c)$ for all k, all $K \in \mathsf{SE.KS}(k)$ and all $M \in \{0,1\}^*$, where $c \leftarrow \mathsf{SE.E}(1^k, K, M)$.

ONE-TIME SECURITY. Consider game $\mathbf{G}_{\mathsf{SE},\mathcal{A}}^{\mathrm{cpr}}(k)$ of Fig. 5 associated to SKE SE, adversary $\mathcal{A}$ and security parameter k, and let

$$\mathbf{Adv}_{\mathsf{SE},\mathcal{A}}^{\mathrm{cpr}}(k) = 2\Pr[\mathbf{G}_{\mathsf{SE},\mathcal{A}}^{\mathrm{cpr}}(k)] - 1.$$

We say that SE has pseudorandom ciphertexts (also called CPR) if the function $\mathbf{Adv}_{\mathsf{SE},\mathcal{A}}^{\mathrm{cpr}}(\cdot)$ is negligible for every $\mathcal{A}$. We require that ENC is queried at most once.

3.4   PKE from Key Encapsulation and Symmetric-Key Encryption

Below, we analyze hybrid encryption in the setting of parameter subversion. Formally we give a transform **KEMToPE** that associates to KEM KE and symmetric-key encryption scheme SE a public-key encryption scheme PE. The construction essentially is the hybrid encryption scheme of [23] including an additional parameter generation algorithm. The scheme's parameter generation, key

$$
\begin{array}{ll}
\underline{\text{Game } \mathbf{G}^{\text{cpr}}_{\mathsf{SE},\mathcal{A}}(k)} & \underline{\text{Enc}(M)} \\
b \leftarrow\!\!\$\ \{0,1\} & \text{If } (b=0) \text{ then } c^* \leftarrow\!\!\$\ \mathsf{SE.CS}(k,|M|) \\
K \leftarrow\!\!\$\ \mathsf{SE.KS}(k) & \text{Else } c^* \leftarrow \mathsf{SE.E}(1^k, K, M) \\
b' \leftarrow \mathcal{A}^{\text{Enc},\text{Dec}}(1^k) & \text{Return } c^* \\
\text{Return } (b=b') & \underline{\text{Dec}(c)} \\
& \text{If } (c=c^*) \text{ then return } \bot \\
& \text{Else return } \mathsf{SE.D}(1^k, K, c)
\end{array}
$$

Fig. 5. Game defining one-time security notions of SKEs.

$$
\begin{array}{ll}
\underline{\mathsf{PE.P}(1^k)} & \underline{\mathsf{PE.E}(1^k, \pi, pk, M)} \\
\pi \leftarrow\!\!\$\ \mathsf{KE.P}(1^k) & (K, c_1) \leftarrow\!\!\$\ \mathsf{KE.E}^{\text{RO}}(1^k, \pi, pk) \\
\text{Return } \pi & c_2 \leftarrow \mathsf{SE.E}(1^k, K, M) \\
\underline{\mathsf{PE.G}(1^k, \pi)} & \text{Return } (c_1, c_2) \\
(pk, sk) \leftarrow\!\!\$\ \mathsf{KE.G}(1^k, \pi) & \underline{\mathsf{PE.D}(1^k, \pi, sk, c)} \\
\text{Return } (pk, sk) & (c_1, c_2) \leftarrow c \\
& K \leftarrow \mathsf{KE.D}^{\text{RO}}(1^k, \pi, sk, c_1) \\
& M \leftarrow \mathsf{SE.D}(1^k, K, c_2) \\
& \text{Return } M
\end{array}
$$

Fig. 6. PKE **KEMToPE**[KE, SE] associated to KEM KE and SE SE.

generation encryption and decryption algorithms are in Fig. 6. PE's ciphertext space is given by $\mathsf{PE.CS}(k, \pi, m) = \mathsf{KE.CS}(k, \pi) \times \mathsf{SE.CS}(k, m)$. It is easy to verify that PE has decryption error $\mathsf{PE.de}(k) = \mathsf{KE.de}(k)$. The following essentially states that hybrid encryption also works in setting of ciphertext pseudorandomness under parameter subversion attacks, i.e., combining a KEM that is both CPR-PSA and WDC-PSA with a SKE that is CPR yields a CPR-PSA PKE, where the well-distributedness of the KEM's ciphertext is necessary to correctly simulate the decryption oracle in the CPR-PSA game with respect to PE.

Proposition 3. *Let* KE *a KEM and* SE *a SE such that* $\mathsf{KE.KS}(k) = \mathsf{SE.KS}(k)$ *for all* $k \in \mathbb{N}$. *Let* $\mathsf{PE} = \mathbf{KEMToPE}[\mathsf{KE}, \mathsf{SE}]$ *be the PKE associated to* KE *and* SE. *If* KE *is* CPR-PSA *and* WDC-PSA *and if* SE *is* CPR *then* PE *is* CPR-PSA *Concretely, given adversary* $\mathcal{A}$ *against* $\mathbf{G}^{\text{cpr-psa}}_{\mathsf{PE},\mathcal{A}}(k)$, *there exist adversaries* $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ *having the same running time and query count as* $\mathcal{A}$, *which satisfy*

$$
\mathbf{Adv}^{\text{cpr-psa}}_{\mathsf{PE},\mathcal{A}}(k) \leq 2\,\mathbf{Adv}^{\text{cpr-psa}}_{\mathsf{KE},\mathcal{B}_1}(k) + \mathbf{Adv}^{\text{wdc-psa}}_{\mathsf{KE},\mathcal{B}_2}(k) + \mathbf{Adv}^{\text{cpr}}_{\mathsf{SE},\mathcal{B}_3}(k) + \mathsf{KE.de}(k).
$$

The proof of the proposition can be found in the full version of this paper [4].

$$
\begin{array}{|ll|}
\hline
\text{Game } \mathbf{G}^{\text{epr-psa}}_{\text{EG},\mathcal{A}}(k) & \text{INIT}(\pi) \\
\hline
b \leftarrow\!\!\$ \{0,1\} & G \leftarrow\!\!\$ \text{ EG.G}(1^k,\pi) \\
b' \leftarrow\!\!\$ \mathcal{A}^{\text{INIT}}(1^k) & \text{If } (G = \bot) \text{ then return } \bot \\
\text{Return } (b = b') & (\langle\mathbb{G}\rangle, n, g) \leftarrow G \\
& \text{If } (b = 1) \text{ then} \\
& \quad y \leftarrow\!\!\$ \text{ EG.S}(1^k,\pi,G) \\
& \quad c \leftarrow\!\!\$ \text{ EG.E}(1^k,\pi,G,g^y) \\
& \text{Else } c \leftarrow\!\!\$ \text{ EG.ES}(k,\pi) \\
& \text{Return } (G,c) \\
\hline
\end{array}
$$

Fig. 7. Game defining embedding pseudorandomness of eeg family EG.

4 KEMs from Efficiently Embeddable Group Families

In this section we define efficiently embeddable group families (eeg). We define the security notion of pseudorandom embeddings under parameter subversion attacks (EPR-PSA) and adapt the strong computational Diffie-Hellman problem (sCDH-PSA) to the setting of efficiently embeddable group families and parameter subversion. Further we give a generic constructions of key encapsulation mechanisms from eeg families. It achieves security assuming the eeg family is sCDH-PSA and EPR-PSA.

4.1 Efficiently Embeddable Group Families

EFFICIENTLY EMBEDDABLE GROUP FAMILIES. An embeddable group family EG specifies the following. Parameter generation algorithm EG.P takes as input 1^k, where $k \in \mathbb{N}$ is the security parameter, and returns parameters π. Group generation algorithm EG.G on input $1^k, \pi$ returns a tuple $G = (\langle\mathbb{G}\rangle, n, g)$, where $\langle\mathbb{G}\rangle$ is a description of a cyclic group $\mathbb{G}$ of order n, and g is a generator of $\mathbb{G}$. EG.ES associates to k a finite set EG.ES(k,π) called the embedding space that is only dependent on k and π. Sampling algorithm EG.S on input of $1^k, \pi$ and $G \in [\text{EG.G}(1^k,\pi)]$ outputs $y \in \mathbb{Z}_n$. (Not necessarily uniformly distributed.) Embedding algorithm EG.E receives as input $1^k, \pi, G \in [\text{EG.G}(1^k,\pi)]$ and $h \in \mathbb{G}$ and returns an element $c \in \text{EG.ES}(k,\pi)$. Deterministic inversion algorithm EG.I on input of $1^k, \pi, G \in [\text{EG.G}(1^k,\pi)]$ and $c \in \text{EG.ES}(k,\pi)$ returns an element of $\mathbb{G}$. The correctness condition requires that for all $k \in \mathbb{N}$, all $\pi \in \text{EG.P}(1^k)$ and all $G \in [\text{EG.G}(1^k,\pi)]$ we have $\Pr\left[\text{EG.I}(1^k,\pi,G,h) = g^y\right] \geq 1 - \text{EG.ie}(k)$, where the probability is over $y \leftarrow\!\!\$ \text{EG.S}(1^k,\pi,G)$ and $h \leftarrow\!\!\$ \text{EG.E}(1^k,\pi,G,g^y)$, and EG.ie : $\mathbb{N} \to \mathbb{R}_{\geq 0}$ is the *inversion error* of EG. If EG.P returns ε on every input 1^k, i.e. if no parameters are used, we say that EG is *parameter-free*.

EMBEDDING PSEUDORANDOMNESS. Consider game $\mathbf{G}^{\text{epr-psa}}_{\text{EG},\mathcal{A}}(k)$ of Fig. 7 associated to eeg family EG, adversary $\mathcal{A}$ and security parameter k. Let

$$\mathbf{Adv}^{\text{epr-psa}}_{\text{EG},\mathcal{A}}(k) = 2\Pr[\mathbf{G}^{\text{epr-psa}}_{\text{EG},\mathcal{A}}(k)] - 1.$$

$$
\begin{array}{ll}
\underline{\text{Game } \mathbf{G}_{\mathsf{EG},\mathcal{A}}^{\text{scdh-psa}}(k)} & \underline{\text{INIT}(\pi)} \\[2pt]
Z \leftarrow_{\$} \mathcal{A}^{\text{INIT,DDH}}(1^k) & G \leftarrow_{\$} \mathsf{EG}.\mathsf{G}(1^k, \pi) \\
\text{Return } (Z = g^{xy} \wedge G \neq \bot) & \text{If } (G = \bot) \text{ then return } \bot \\[4pt]
\underline{\text{DDH}(\tilde{Y}, \tilde{Z})} & (\langle \mathbb{G} \rangle, n, g) \leftarrow G \\[2pt]
\text{Return } (\tilde{Y}^x = \tilde{Z}) & x \leftarrow_{\$} \mathbb{Z}_n \\
& y \leftarrow_{\$} \mathsf{EG}.\mathsf{S}(1^k, \pi, G) \\
& \text{Return } (G, g^x, g^y)
\end{array}
$$

Fig. 8. Experiment for the strong computational Diffie-Hellman problem with respect to eeg family EG. Oracle INIT may be queried only once and has to be queried before using oracle DDH.

We say that EG has pseudorandom embeddings under parameter subversion attacks (also called EPR-PSA) if the function $\mathbf{Adv}_{\mathsf{EG},\mathcal{A},.}^{\text{epr-psa}}$ is negligible for every $\mathcal{A}$. In the game, b is a challenge bit. When $b = 1$, the challenge embedding c^* is generated by sampling an exponent using EG.S and embedding the group generator raised to the exponent with EG.E. If $b = 0$ the adversary is given an embedding sampled uniformly from the embedding space. Given the group and the embedding, the adversary outputs a guess b' and wins if b' equals b. The parameters used in the game are provided by the adversary making a single call to the oracle INIT. All of our instantiations sample exponents such that the resulting embeddings are statistically close to uniform on $\mathsf{EG}.\mathsf{ES}(k, \pi)$, and hence achieve this notion statistically.

DIFFIE-HELLMAN PROBLEM WITH RESPECT TO EG. The computational Diffie-Hellman problem for a cyclic group $\mathbb{G}$ of order n, which is generated by g, asks to compute g^{xy} given g^x and g^y, where $x, y \leftarrow_{\$} \mathbb{Z}_n$. In the strong computational Diffie-Hellman problem introduced by Abdalla et al. in [2] the adversary additionally has access to an oracle, which may be used to check whether $Y^x = Z$ for group elements $Y, Z \in \mathbb{G}$. We provide a definition for the strong computational Diffie-Hellman problem with respect to eeg families EG, which allows parameter subversion. An additional difference is that y is not chosen uniformly from $\mathbb{Z}_n$ but instead sampled using EG.S.

Thus, consider game $\mathbf{G}_{\mathsf{EG},\mathcal{A}}^{\text{scdh-psa}}(k)$ of Fig. 8. The game is associated to eeg family EG, adversary $\mathcal{A}$ and security parameter k. The adversary has access to an oracle INIT setting up a problem instance according to the parameters it is provided. Let

$$\mathbf{Adv}_{\mathsf{EG},\mathcal{A}}^{\text{scdh-psa}}(k) := \Pr\left[\mathbf{G}_{\mathsf{EG},\mathcal{A}}^{\text{scdh-psa}}(k)\right].$$

We say that the strong computational Diffie-Hellman problem under parameter subversion (also called sCDH-PSA) is hard with respect to EG if $\mathbf{Adv}_{\mathsf{EG},\mathcal{A}}^{\text{scdh-psa}}(\cdot)$ is negligible for every adversary $\mathcal{A}$.

$$
\begin{array}{l|l|l}
\underline{\mathsf{KE}.\mathsf{G}_1(1^k, \pi)} & \underline{\mathsf{KE}.\mathsf{E}_1^{\mathrm{RO}}(1^k, \pi, pk)} & \underline{\mathsf{KE}.\mathsf{D}_1^{\mathrm{RO}}(1^k, \pi, sk, c)} \\
G \leftarrow\!\!\$\ \mathsf{EG}.\mathsf{G}(1^k, \pi) & (G, X) \leftarrow pk & (G, x, pk) \leftarrow sk \\
\text{If } (G = \bot) \text{ return } \bot & y \leftarrow\!\!\$\ \mathsf{EG}.\mathsf{S}(1^k, \pi, G) & Y \leftarrow \mathsf{EG}.\mathsf{I}(1^k, \pi, G, c) \\
(\langle \mathbb{G} \rangle, n, g) \leftarrow G & Y \leftarrow g^y & K \leftarrow \mathrm{RO}((pk, c, Y^x), m(k)) \\
x \leftarrow\!\!\$\ \mathbb{Z}_n;\ X \leftarrow g^x & c \leftarrow\!\!\$\ \mathsf{EG}.\mathsf{E}(1^k, \pi, G, Y) & \text{Return } K \\
pk \leftarrow (G, X) & K \leftarrow \mathrm{RO}((pk, c, X^y), m(k)) & \\
sk \leftarrow (G, x, pk) & \text{Return } (K, c) & \underline{\mathsf{KE}.\mathsf{P}_1(1^k)} \\
\text{Return } (pk, sk) & & \pi \leftarrow\!\!\$\ \mathsf{EG}.\mathsf{P}(1^k) \\
& & \text{Return } \pi
\end{array}
$$

Fig. 9. KEM $\mathsf{KE}_1 = \mathbf{eegToKE1}[\mathsf{EG}, m]$ built from eeg family EG and polynomial m via our transform. The KE has key space $\mathsf{KE}.\mathsf{KS}(k) = \{0,1\}^{m(k)}$ and ciphertext space $\mathsf{KE}.\mathsf{CS}(k, \pi) = \mathsf{EG}.\mathsf{ES}(k, \pi)$.

4.2 Key Encapsulation from Efficiently Embeddable Group Families

In this section we give a generic construction of a key encapsulation mechanism from an eeg family EG. Its security is based on the strong Diffie-Hellman problem, i.e. if sCDH-PSA is hard with respect to EG, the KEM is IND-PSA. If additionally EG has pseudorandom embeddings, the KEM has pseudorandom and well-distributed ciphertexts. The construction is similar to the standard El Gamal based key encapsulation mechanism as for example used in [2,23]. As an intermediate step in the proof that the construction is CPR-PSA we obtain that it is IND-PSA. The proof of this property follows the outlines of the proofs given in [2,23]. Afterwards we use the pseudorandomness of the eeg family's embeddings to show, that our construction achieves pseudorandom and well-distributed ciphertexts.

Formally, we define a transform $\mathbf{eegToKE1}$ that associates to an eeg family EG and a polynomial $m : \mathbb{N} \to \mathbb{N}$ a KEM $\mathsf{KE} = \mathbf{eegToKE1}[\mathsf{EG}, m]$. The parameter generation, key generation, encryption and decryption algorithms of KE are in Fig. 9. The construction is in the ROM, so that encryption and decryption invoke the RO oracle. The key space is $\mathsf{KE}.\mathsf{KS}(k) = \{0,1\}^{m(k)}$. The ciphertext space $\mathsf{KE}.\mathsf{CS}(k, \pi) = \mathsf{EG}.\mathsf{ES}(k, \pi)$ is the embedding space of EG. It is easy to verify that $\mathsf{KE}.\mathsf{de} = \mathsf{EG}.\mathsf{ie}$, meaning the decryption error of the KEM equals the inversion error of the eeg family.

SECURITY OF THE CONSTRUCTION. The following says that if sCDH-PSA is hard with respect to eeg family EG then $\mathbf{eegToKE1}[\mathsf{EG}, m]$ has desirable security properties.

Theorem 4. *Let* $\mathsf{KE} = \mathbf{eegToKE1}[\mathsf{EG}, m]$ *be the KEM associated to eeg family* EG *and polynomial* $m : \mathbb{N} \to \mathbb{N}$ *as defined in Fig. 9. Assume that* EG *is* EPR-PSA *and that* sCDH-PSA *is hard with respect to* EG. *Then*

(i) KE *has pseudorandom ciphertexts under parameter subversion attacks.*
(ii) KE *has well-distributed ciphertexts under parameter subversion attacks.*

Moreover, if EG *is parameter-free so is* KE. *Concretely, given an adversary* $\mathcal{A}$ *making at most* $q(k)$ *queries to* RO *the proof specifies adversaries* $\mathcal{B}_1$ *and* $\mathcal{B}_2$ *having the same running time as* $\mathcal{A}$ *satisfying*

$$\mathbf{Adv}_{\mathsf{KE}}^{\mathrm{cpr\text{-}psa}}(\mathcal{A})(k) \leq \mathbf{Adv}_{\mathsf{EG},\mathcal{B}_1}^{\mathrm{scdh\text{-}psa}}(k) + \mathbf{Adv}_{\mathsf{EG},\mathcal{B}_2}^{\mathrm{epr\text{-}psa}}(k),$$

where $\mathcal{B}_2$ *makes at most* $q(k)$ *queries to* DDH. *Furthermore given an adversary* $\mathcal{A}'$ *the proof specifies an adversary* $\mathcal{B}'$ *having the same running time as* $\mathcal{A}'$ *such that,*

$$\mathbf{Adv}_{\mathsf{KE},\mathcal{A}'}^{\mathrm{wdc\text{-}psa}}(k) \leq \mathbf{Adv}_{\mathsf{EG},\mathcal{B}'}^{\mathrm{epr\text{-}psa}}(k) + \mathsf{EG.ie}(k).$$

The proof of the theorem can be found in the full version of this paper [4]. In the full version of this paper [4] we also provide a transform **eegToKE2**, which achieves security under the weaker CDH-PSA assumption with respect to EG.

5 Efficiently Embeddable Group Families from Curve-Twist Pairs

In this section we give instantiations of eeg families based on elliptic curves. The main tool of the constructions is a bijection of [34] mapping points of an elliptic curve and its quadratic twist to an interval of integers. We first give a construction using parameters, the parameter being a prime p of length k serving as the modulus of the prime field the curves are defined over. The construction has embedding space $[2p+1]$. Since we assume, that the parameter shared by all users might be subject to subversion, security of this construction corresponds to the assumption that there exist no inherently bad choices for p, i.e. that for *any* sufficiently large prime p it is possible to find elliptic curves defined over $\mathbb{F}_p$ on which the strong computational Diffie-Hellman assumption holds.

As an alternative we also give parameter-free eeg-families whose security is based on the weaker assumption that for *random* k-bit prime p it is possible to find elliptic curves defined over $\mathbb{F}_p$, such that the strong computational Diffie-Hellman assumption holds. Since in this construction the modulus p is sampled along with the curve, it is no longer possible to use $[2p + 1]$ as the embedding space of the eeg family. We propose two solutions to overcome this, one using rejection sampling to restrict the embedding space to the set $[2^k]$, the other one is based on a technique from [33] and expands the embedding space to $[2^{k+1}]$.

5.1 Elliptic Curves

Let $p \geq 5$ be prime and $\mathbb{F}_p$ a field of order p. An elliptic curve over $\mathbb{F}_p$ can be expressed in short Weierstrass form, that is as the set of projective solutions of an equation of the form

$$YZ^2 = X^3 + aXZ^2 + bZ^3,$$

where $a, b \in \mathbb{F}_p$ with $4a^3 + 27b^2 \neq 0$. We denote the elliptic curve generated by p, a, b by $E(p, a, b)$. $E(p, a, b)$ possesses exactly one point with Z-coordinate 0,

the so called point at infinity $\mathcal{O} = (0 : 1 : 0)$. After normalizing by $Z = 1$ the curve's other points can be interpreted as the solutions $(x, y) \in \mathbb{F}_p^2$ of the affine equation $y^2 = x^3 + ax + b$. It is possible to establish an efficiently computable group law on $E(p, a, b)$ with $\mathcal{O}$ serving as the neutral element of the group. We use multiplicative notation for the group law to be consistent with the rest of the paper.

TWISTS OF ELLIPTIC CURVES. In [34, Sect. 4] Kaliski establishes the following one-to-one correspondence between two elliptic curves defined over $\mathbb{F}_p$ which are related by twisting and a set of integers.

Lemma 5. *Let $p \in \mathbb{N}_{\geq 5}$ be prime. Let $u \in \mathbb{Z}_p$ be a quadratic nonresidue modulo p and $a, b \in \mathbb{Z}_p$ such that $4a^3 + 27b^2 \neq 0$. Consider the elliptic curves $E_0 := E(p, a, b)$ and $E_1 := E(p, au^2, bu^3)$. Then $|E_0| + |E_1| = 2p + 2$. Furthermore, the functions $l_0 : E_0 \longrightarrow [2p + 2]$ and $l_1 : E_1 \longrightarrow [2p + 2]$ defined as*

$$
l_0(P) = \begin{cases} p & \text{if } P = \mathcal{O}_0 \\ (ux \bmod p) & \text{if } (P = (x, y)) \wedge (0 \leq y \leq (p-1)/2), \\ (ux \bmod p) + p + 1 & \text{if } (P = (x, y)) \wedge ((p-1)/2 < y) \end{cases}
$$

$$
l_1(P) = \begin{cases} 2p + 1 & \text{if } P = \mathcal{O}_1 \\ x & \text{if } (P = (x, y)) \wedge (0 < y \leq (p-1)/2) \\ x + p + 1 & \text{if } (P = (x, y)) \wedge ((y = 0) \vee ((p-1)/2 < y)) \end{cases}
$$

are injective with nonintersecting ranges, where $\mathcal{O}_0$ and $\mathcal{O}_1$ denote the neutral elements of E_0 and E_1 respectively.

Lemma 6. *The functions l_0 and l_1 can be efficiently inverted. That is, given $z \in [2p + 1]$, one can efficiently compute the unique $(P, \delta) \in E_0 \cup E_1 \times \{0, 1\}$ such that $l_\delta(P) = z$.*

The proof of the lemma can be found in the full version of this paper [4].

Definition 7. *A curve-twist generator TGen on input of security parameter 1^k and a k-bit prime p returns (G_0, G_1), where $G_0 = (\langle E_0 \rangle, n_0, g_0)$ and $G_1 = (\langle E_1 \rangle, n_1, g_1)$ are secure cyclic elliptic curves defined over the field $\mathbb{F}_p$. More precisely we require $E_0 := E(p, a, b)$ and $E_1 := E(p, au^2, bu^3)$ for $a, b \in \mathbb{F}_p$ such that $(4a^3 + 27b^2) \neq 0$ and quadratic nonresidue u. Furthermore we require that g_0 generates E_0 and g_1 generates E_1 as well as $|E_0| = n_0$, $|E_1| = n_1$ and $\gcd(n_0, n_1) = 1$.*

GENERATION OF SECURE TWISTED ELLIPTIC CURVES. There exist several proposals for properties an elliptic curve over a prime field $\mathbb{F}_p$ should have to be considered secure (e.g., [18, 27]). Firstly, the elliptic curve's order is required to be either the product of a big prime and a small cofactor—or preferably prime. Secondly, several conditions preventing the transfer of discrete logarithm problems on the curve to groups, where faster algorithms to compute discrete

logarithms may be applied, should be fulfilled. Finally, for our applications we need both the elliptic curve and its quadratic twist to be secure, a property usually called twist security. For concreteness, we suggest to implement $\mathsf{TGen}(1^k, p)$ by sampling the necessary parameters a, b, u with rejection sampling such that the resulting curve $E(p, a, b)$ fulfills the three security requirement mentioned above. This way, TGen can be implemented quite efficiently[1] and furthermore, with overwhelming probability, the resulting curve fulfills all relevant security requirements from [18,27] that are not covered by the three security properties explicitly mentioned above.

COMPUTATIONAL PROBLEMS ASSOCIATED TO TGen. Let TGen a curve-twist generator. We give two versions of the strong computational Diffie-Hellman assumption with respect to TGen. In the first version the prime p on which TGen is invoked is chosen by the adversary, while in the second version p is sampled uniformly at random from all k-bit primes. For $d \in \{0, 1\}$ consider games $\mathbf{G}_{\mathsf{TGen},\mathcal{A}}^{\mathrm{twist}_d\text{-cp-scdh}}(\cdot)$ and $\mathbf{G}_{\mathsf{TGen},\mathcal{A}}^{\mathrm{twist}_d\text{-up-scdh}}(\cdot)$ of Fig. 10. We define advantage functions

$$\mathbf{Adv}_{\mathsf{TGen},\mathcal{A}}^{\mathrm{twist}_d\text{-cp-scdh}}(k) = \Pr\Big[\mathbf{G}_{\mathsf{TGen},\mathcal{A}}^{\mathrm{twist}_d\text{-cp-scdh}}(k)\Big],$$

$$\mathbf{Adv}_{\mathsf{TGen},\mathcal{A}}^{\mathrm{twist}_d\text{-up-scdh}}(k) = \Pr\Big[\mathbf{G}_{\mathsf{TGen},\mathcal{A}}^{\mathrm{twist}_d\text{-up-scdh}}(k)\Big].$$

Definition 8. *Let* TGen *be a curve-twist generator. We say the strong computational Diffie-Hellman assumption for chosen (uniform) primes holds with respect to curve-twist generator* TGen, *if both* $\mathbf{Adv}_{\mathsf{TGen},\mathcal{A}}^{\mathrm{twist}_0\text{-cp-scdh}}(\cdot)$ *and* $\mathbf{Adv}_{\mathsf{TGen},\mathcal{A}}^{\mathrm{twist}_1\text{-cp-scdh}}(\cdot)$ *(or* $\mathbf{Adv}_{\mathsf{TGen},(P_k)_k,\mathcal{A}}^{\mathrm{twist}_0\text{-up-scdh}}(\cdot)$ *and* $\mathbf{Adv}_{\mathsf{TGen},(P_k)_k,\mathcal{A}}^{\mathrm{twist}_1\text{-up-scdh}}(\cdot)$ *respectively) are negligible for all adversaries* $\mathcal{A}$.

5.2 An Eeg Family from Elliptic Curves

In [34] Kaliski implicitly gives an eeg family based on elliptic curves. The family is parameter-using, the parameter being a prime p serving as the modulus of the field the elliptic curves are defined over. The definition of eeg family $\mathsf{EG}_{\mathrm{twist}}$ may be found in Fig. 11. Parameter generation algorithm $\mathsf{EG}_{\mathrm{twist}}.\mathsf{P}$ on input of security parameter 1^k returns a randomly sampled k-bit prime[2] p. Group generation algorithm $\mathsf{EG}_{\mathrm{twist}}.\mathsf{G}$ on input of parameter $\pi = p$ checks, whether p is

---

[1] In [29] Galbraith and McKee consider elliptic curves E chosen uniformly from the set of elliptic curves over a fixed prime field $\mathbb{F}_p$. They give a conjecture (together with some experimental evidence) for a lower bound on the probability of $|E|$ being prime. Using a similar technique [27] argue, that the probability of a uniformly chosen elliptic curve over a fixed prime field $\mathbb{F}_p$ to be both secure and twist secure is bounded from below by $0.5/\log^2(p)$. Since their definition of security of an elliptic curve includes primality of the curve order and since due to Lemma 5 the orders of curve and twist sum up to $2p + 2$, this in particular implies that the curve and its twist are cyclic and have coprime group order.

[2] In practice one would preferably instantiate $\mathsf{EG}_{\mathrm{twist}}$ with a standardized prime.

$$\begin{array}{ll}
\underline{\text{Game } \mathbf{G}^{\text{twist}_d\text{-cp-scdh}}_{\mathsf{TGen},\mathcal{A}}(k)} & \underline{\text{INIT}(\pi)} \\
\end{array}$$

Game $\mathbf{G}^{\text{twist}_d\text{-cp-scdh}}_{\mathsf{TGen},\mathcal{A}}(k)$	$\text{INIT}(\pi)$
$Z \leftarrow_{\$} \mathcal{A}^{\text{INIT},\text{DDH}}(1^k)$	$p \leftarrow \pi$
Return $(Z = g_d^{xy})$	If $(p \notin \mathcal{P}_k)$ then return $\bot$
	$(G_0, G_1) \leftarrow_{\$} \mathsf{TGen}(1^k, p)$
$\underline{\text{DDH}(\tilde{Y}_d, \tilde{Z}_d)}$	$(\langle E_d \rangle, n_d, g_d) \leftarrow G_d$
If $\tilde{Y}_d \notin E_d \vee \tilde{Z}_d \notin E_d$	$x \leftarrow_{\$} \mathbb{Z}_{n_d}; \ y \leftarrow_{\$} \mathbb{Z}_{n_d}$
$\quad$ return $\bot$	$X \leftarrow g_d^x, \ Y \leftarrow g_d^y$
Return $(\tilde{Y}_d^{\ x} = \tilde{Z}_d)$	Return (G_0, G_1, X, Y)

Game $\mathbf{G}^{\text{twist}_d\text{-up-scdh}}_{\mathsf{TGen},\mathcal{A}}(k)$	INIT
$Z \leftarrow_{\$} \mathcal{A}^{\text{INIT},\text{DDH}}(1^k)$	$p \leftarrow_{\$} \mathcal{P}_k$
Return $(Z = g_d^{xy})$	$(G_0, G_1) \leftarrow_{\$} \mathsf{TGen}(1^k, p)$
	$(\langle E_d \rangle, n_d, g_d) \leftarrow G_d$
$\underline{\text{DDH}(\tilde{Y}_d, \tilde{Z}_d)}$	$x \leftarrow_{\$} \mathbb{Z}_{n_d}; \ y \leftarrow_{\$} \mathbb{Z}_{n_d}$
If $(\tilde{Y}_d \notin E_d \vee \tilde{Z}_d \notin E_d)$	$X \leftarrow g_d^x, \ Y \leftarrow g_d^y$
$\quad$ return $\bot$	Return (G_0, G_1, p, X, Y)
Return $(\tilde{Y}_d^{\ x} = \tilde{Z}_d)$	

Fig. 10. Experiments for the sCDH problem for chosen (uniform) primes with respect to $d \in \{0, 1\}$, adversary $\mathcal{A}$ and curve-twist generator TGen.

indeed a prime of appropriate length, and—if so—runs a curve-twist generator $\mathsf{TGen}(1^k, \pi)$ to obtain the description of two cyclic secure cyclic elliptic curves $G_0 = (\langle E_0 \rangle, n_0, g_0)$ and $G_1 = (\langle E_1 \rangle, n_1, g_1)$. Its output is $(\langle \mathbb{G} \rangle, n, g)$, where $\mathbb{G} \leftarrow E_0 \times E_1$ is the direct product of the two elliptic curves, $n \leftarrow n_0 \cdot n_1$ and $g \leftarrow (g_0, g_1)$. Here we assume that the description $\langle \mathbb{G} \rangle$ of $\mathbb{G}$ includes the values n_0 and n_1, which are used by $\mathsf{EG}_{\text{twist}}$'s other algorithms. Note that $|\mathbb{G}| = n$ and since n_0 and n_1 are coprime, g generates $\mathbb{G}$. Furthermore, if we regard E_0 and E_1 as subgroups of $\mathbb{G} = E_0 \times E_1$ in the natural way, we may rewrite the set $E_0 \cup E_1 \subseteq \mathbb{G}$ as

$$\begin{aligned}
E_0 \cup E_1 &= \{(h_0, \mathcal{O}_1) \mid h_0 \in E_0\} \cup \{(\mathcal{O}_0, h_1) \mid h_1 \in E_1\} \\
&= \{(g_0, g_1)^y \mid y \in \mathbb{Z}_n : y \equiv 0 \mod n_0 \text{ or } y \equiv 0 \mod n_1\}
\end{aligned}$$

Algorithm $\mathsf{EG}_{\text{twist}}.\mathsf{S}$ uses this property to efficiently sample $y \in \mathbb{Z}_n$ such that $g^y \sim U_{E_0 \cup E_1}$. It first samples $z \leftarrow_{\$} \mathbb{Z}_{2p+1}$. If $z < n_0$ it returns $\varphi_{\text{crt}}(z, 0)$. Else it returns $\varphi_{\text{crt}}(0, z - n_0 - 1)$. Here φ_{crt} denotes the canonical isomorphism $\varphi_{\text{crt}} : \mathbb{Z}_{n_0} \times \mathbb{Z}_{n_1} \to \mathbb{Z}_n$. As a result $y \leftarrow_{\$} \mathsf{EG}_{\text{twist}}.\mathsf{S}(1^k, G)$ satisfies $y \sim U_M$, where $M := \{y \in \mathbb{Z}_n \mid y \equiv 0 \mod n_0 \text{ or } y \equiv 0 \mod n_1\}$. Embedding algorithm $\mathsf{EG}_{\text{twist}}.\mathsf{E}$ receives as input $1^k, \pi, G$ and $h = (h_0, h_1) \in \mathbb{G}$. It first checks, whether h lies outside of the support $[\mathsf{EG}_{\text{twist}}.\mathsf{S}(1^k, \pi, G)]$ of the sampling algorithm, i.e. whether both $h_0 \neq \mathcal{O}_0$ and $h_1 \neq \mathcal{O}_1$. In this case the element is mapped to 0. If h is an element of $[\mathsf{EG}_{\text{twist}}.\mathsf{S}(1^k, \pi, G)]$, algorithm $\mathsf{EG}_{\text{twist}}.\mathsf{E}$ returns $l_0(h_0)$ if $h_1 = \mathcal{O}_1$, and $l_1(h_1)$ if $h_1 \neq \mathcal{O}_1$. Here $l_0 : E_0 \to [2p + 2]$ and $l_1 : E_1 \to [2p + 2]$ denote the maps of Lemma 5. By Lemma 5 the map $\mathsf{EG}_{\text{twist}}.\mathsf{E}(1^k, G, \cdot)|_{E_0 \cup E_1}$ is

$$
\begin{array}{ll}
\underline{\mathsf{EG}_{\text{twist}}.\mathsf{P}(1^k)} & \underline{\mathsf{EG}_{\text{twist}}.\mathsf{S}(1^k, \pi, G)} \\
p \leftarrow_\$ \mathcal{P}_k & p \leftarrow \pi \\
\pi \leftarrow p & z \leftarrow_\$ \mathbb{Z}_{2p+1} \\
\text{Return } \pi & \text{If } (z < n_0) \text{ return } \varphi_{\text{crt}}(z, 0) \\
& \text{Else return } \varphi_{\text{crt}}(0, z - n_0 - 1) \\
\underline{\mathsf{EG}_{\text{twist}}.\mathsf{G}(1^k, \pi)} & \\
p \leftarrow \pi & \underline{\mathsf{EG}_{\text{twist}}.\mathsf{E}(1^k, \pi, G, (h_0, h_1))} \\
\text{If } (p \notin \mathcal{P}_k) \text{ return } \bot & \text{If } (h_0 \neq \mathcal{O}_0 \wedge h_1 \neq \mathcal{O}_1) \text{ return } 0 \\
(G_0, G_1) \leftarrow_\$ \mathsf{TGen}(1^k, p) & \text{Elseif } h_1 = \mathcal{O}_1 \text{ return } l_0(h_0) \\
(\langle E_0\rangle, g_0, n_0) \leftarrow G_0; \ (\langle E_1\rangle, g_1, n_1) \leftarrow G_1 & \text{Else return } l_1(h_1) \\
\mathbb{G} \leftarrow E_0 \times E_1; \ g \leftarrow (g_0, g_1); \ n \leftarrow n_0 \cdot n_1 & \\
G \leftarrow (\langle\mathbb{G}\rangle, n, g) & \underline{\mathsf{EG}_{\text{twist}}.\mathsf{I}(1^k, \pi, G, z)} \\
\text{Return } G & \text{If } (z \in \operatorname{im}(l_0)) \text{ return } l_0^{-1}(z) \\
& \text{Else return } l_1^{-1}(z)
\end{array}
$$

Fig. 11. Definition of eeg family $\mathsf{EG}_{\text{twist}}$ with embedding space $\mathsf{EG}_{\text{twist}}.\mathsf{ES}(k, \pi) = [2p + 1]$. l_0 and l_1 denote the maps from Lemma 5, φ_{crt} the canonical isomorphism $\mathbb{Z}_{n_0} \times \mathbb{Z}_{n_1} \to \mathbb{Z}_n$.

a bijection between $E_0 \cup E_1$ and $[2p + 1]$ and we obtain $\mathsf{EG}_{\text{twist}}.\mathsf{E}(1^k, G, g^y) \sim U_{[2p+1]}$ for y sampled with $\mathsf{EG}_{\text{twist}}.\mathsf{S}(1^k, G)$. We obtain the following.

Lemma 9. EG_{twist} *as defined in Fig. 11 is an eeg family with embedding space* $\mathsf{EG}_{twist}.\mathsf{ES}(k, G) = [2p + 1]$ *and inversion error* $\mathsf{EG}_{twist}.\mathsf{ie}(k) = 0$. *Furthermore* EG_{twist} *has pseudorandom embeddings. More precisely, for every (potentially unbounded) adversary $\mathcal{A}$ we have*

$$
\mathbf{Adv}^{\text{epr-psa}}_{\mathsf{EG}_{twist}, \mathcal{A}}(k) = 0.
$$

A proof of the lemma can be found in the full version of the paper [4]. Concerning the hardness of sCDH-PSA with respect to $\mathsf{EG}_{\text{twist}}$ we obtain the following.

Lemma 10. *Let EG_{twist} be the embeddable group generator constructed with respect to twisted elliptic curve generator TGen as described above. If the strong Diffie-Hellman assumption for chosen primes holds with respect to TGen, then the strong Diffie-Hellman assumption holds with respect to EG_{twist}.*

Concretely for every adversary $\mathcal{A}$ against game $\mathbf{G}^{\text{scdh-psa}}_{\mathsf{EG}_{twist}, \mathcal{A}}(\cdot)$, which makes at most Q queries to its DDH-oracle, there exist adversaries $\mathcal{B}_0, \mathcal{B}_1$ against games $\mathbf{G}^{\text{twist}_0\text{-cp-scdh}}_{\mathsf{TGen}, \mathcal{B}_0}(\cdot)$ or $\mathbf{G}^{\text{twist}_1\text{-cp-scdh}}_{\mathsf{TGen}, \mathcal{B}_1}(\cdot)$ respectively making at most Q queries to their DDH-oracles, satisfying

$$
\mathbf{Adv}^{\text{scdh-psa}}_{\mathsf{EG}_{twist}, \mathcal{A}}(k) \leq \mathbf{Adv}^{\text{twist}_0\text{-cp-scdh}}_{\mathsf{TGen}, \mathcal{B}_0}(k) + \mathbf{Adv}^{\text{twist}_1\text{-cp-scdh}}_{\mathsf{TGen}, \mathcal{B}_1}(k).
$$

The proof of the lemma can be found in the full version of this paper [4].

5.3 A Parameter-Free Eeg Family Using Rejection Sampling

Eeg family $\mathsf{EG}_{\text{twist}}$ of Sect. 5.2 is parameter-using, the parameter being the size p of the field $\mathbb{F}_p$. Correspondingly, hardness of sCDH-PSA with respect to $\mathsf{EG}_{\text{twist}}$

$$
\begin{array}{lll}
\underline{\mathsf{EG}^\ell_{\text{twist-rs}}.\mathsf{P}(1^k)} & \underline{\mathsf{EG}^\ell_{\text{twist-rs}}.\mathsf{S}(1^k, \pi, G)} & \underline{\mathsf{EG}^\ell_{\text{twist-rs}}.\mathsf{E}(1^k, \pi, G, h)} \\
\text{Return } \varepsilon & (G', p) \leftarrow G & (G', p) \leftarrow G' \\
& \text{For } \ell^* = 1 \text{ to } \ell & z \leftarrow_\$ \mathsf{EG}_{\text{twist}}.\mathsf{E}(1^k, p, G', h) \\
\underline{\mathsf{EG}^\ell_{\text{twist-rs}}.\mathsf{G}(1^k, \pi)} & \quad \text{Do } y \leftarrow \mathsf{EG}_{\text{twist}}.\mathsf{S}(1^k, p, G') & \text{Return } z \\
p \leftarrow_\$ \mathcal{P}_k & \quad \text{If } (\mathsf{EG}_{\text{twist}}.\mathsf{E}(1^k, p, G, g^y) < 2^k) & \underline{\mathsf{EG}^\ell_{\text{twist-rs}}.\mathsf{I}(1^k, \pi, G, z)} \\
G' \leftarrow_\$ \mathsf{EG}_{\text{twist}}.\mathsf{G}(1^k, p) & \qquad \text{return } y & (G', p) \leftarrow G' \\
G \leftarrow (G', p) & \text{Return } \bot & h \leftarrow \mathsf{EG}_{\text{twist}}.\mathsf{I}(1^k, p, G', z) \\
\text{Return } G & & \text{Return } h
\end{array}
$$

Fig. 12. Parameter-free eeg family $\mathsf{EG}^\ell_{\text{twist-rs}}$.

follows from the assumption, that the elliptic curves output by curve-twist generator TGen are secure, independently of the prime p the curve-twist generator TGen is instantiated with. In this section we show how $\mathsf{EG}_{\text{twist}}$ can be used to construct an eeg family $\mathsf{EG}^\ell_{\text{twist-rs}}$ for which hardness of sCDH-PSA follows from the weaker assumption that TGen instantiated with a *randomly* chosen prime is able to sample secure elliptic curves. The construction is parameter-free and has embedding space $[2^k]$. The size p of the field over which the elliptic curves are defined is now sampled as part of the group generation. The embedding algorithm uses rejection sampling to ensure that embeddings of group elements g^y for y sampled with $\mathsf{EG}^\ell_{\text{twist-rs}}.\mathsf{S}$ are elements of $[2^k]$. The specification of $\mathsf{EG}^\ell_{\text{twist-rs}}$'s algorithms may be found in Fig. 12.

Theorem 11. *Let* $\ell : \mathbb{N} \to \mathbb{N}$ *be a polynomial.* $\mathsf{EG}^\ell_{twist\text{-}rs}$ *as described above is an eeg family with embedding space* $\mathsf{EG}^\ell_{twist\text{-}rs}.\mathsf{ES}(k, \pi) = [2^k]$ *and inversion error* $\mathsf{EG}^\ell_{twist\text{-}rs}.\mathsf{ie}(k) \leq 2^{-\ell(k)}$. *Furthermore* $\mathsf{EG}^\ell_{twist\text{-}rs}$ *has pseudorandom embeddings. More precisely, for every (potentially unbounded) adversary* $\mathcal{A}$ *we have*

$$\mathbf{Adv}^{\text{epr-psa}}_{\mathsf{EG}^\ell_{twist\text{-}rs}, \mathcal{A}}(k) \leq 2^{-\ell(k)}.$$

The proof of the theorem can be found in the full version of this paper [4]. As discussed above, we obtain that—assuming that TGen invoked on randomly sampled prime p returns a secure curve-twist pair—the sCDH-PSA-problem with respect to eeg family $\mathsf{EG}^\ell_{\text{twist-rs}}$ is hard.

Lemma 12. *Let* $\ell : \mathbb{N} \to \mathbb{N}$ *be a polynomial and* $\mathsf{EG}^\ell_{twist\text{-}rs}$ *the eeg family with underlying curve-twist generator* TGen *as described above. If the sCDH assumption for uniform primes holds with respect to* TGen, *then* sCDH-PSA *is hard with respect to* $\mathsf{EG}^\ell_{twist\text{-}rs}$. *Concretely, for every adversary* $\mathcal{A}$ *against game* $\mathbf{G}^{\text{scdh-psa}}_{\mathsf{EG}^\ell_{twist\text{-}rs}, \mathcal{A}}(\cdot)$ *making at most* Q *queries to its* DDH-*oracle there exist adversaries* $\mathcal{B}_0$, $\mathcal{B}_1$ *against* $\mathbf{G}^{\text{twist}_0\text{-up-scdh}}_{\mathsf{TGen}, \mathcal{B}_0}(\cdot)$ *or* $\mathbf{G}^{\text{twist}_1\text{-up-scdh}}_{\mathsf{TGen}, \mathcal{B}_1}(\cdot)$ *respectively, making at most* Q *queries to their* DDH-*oracles and running in the same time as* $\mathcal{A}$, *which satisfy*

$$\mathbf{Adv}^{\text{scdh-psa}}_{\mathsf{EG}^\ell_{twist\text{-}rs}, \mathcal{A}}(k) \leq 3 \left(\mathbf{Adv}^{\text{twist}_0\text{-up-scdh}}_{\mathsf{TGen}, \mathcal{B}_0}(k) + \mathbf{Adv}^{\text{twist}_1\text{-up-scdh}}_{\mathsf{TGen}, \mathcal{B}_1}(k) \right) + 2^{-\ell(k)}$$

for all $k \in \mathbb{N}_{\geq 6}$.

$\mathsf{EG}_{\text{twist-re}}.\mathsf{G}(1^k, \pi)$
_____
$p \leftarrow_{\$} \mathcal{P}_k$
$G' \leftarrow_{\$} \mathsf{EG}_{\text{twist}}.\mathsf{G}(1^k, p); G \leftarrow (G', p)$
Return G

$\mathsf{EG}_{\text{twist-re}}.\mathsf{S}(1^k, \pi, G)$
_____
$(G', p) \leftarrow G$
$z \leftarrow_{\$} [2^{k+1}]$
If $(z \leq 2p)$
$\quad y \leftarrow \psi_G(z)$
$\quad$ If $(\mathsf{EG}_{\text{twist}}.\mathsf{E}(1^k, p, G', g^y) < 2^{k+1} - (2p + 1))$
$\quad\quad$ return y
$\quad$ Else $z \leftarrow \mathsf{EG}_{\text{twist}}.\mathsf{E}(1^k, p, G', g^y)$
Else $z \leftarrow z - (2p + 1)$
$y \leftarrow \psi_G(z)$
Return y

$\mathsf{EG}_{\text{twist-re}}.\mathsf{P}(1^k)$
_____
Return ε

$\mathsf{EG}_{\text{twist-re}}.\mathsf{E}(1^k, \pi, G, h)$
_____
$(G', p) \leftarrow G; b \leftarrow_{\$} \{0, 1\}$
$z \leftarrow \mathsf{EG}_{\text{twist}}.\mathsf{E}(1^k, p, G', h)$
If $z < 2^{k+1} - (2p + 1)$
$\quad z \leftarrow z + b(2p + 1)$
Return z

$\mathsf{EG}_{\text{twist-re}}.\mathsf{I}(1^k, \pi, G, z)$
_____
$(G', p) \leftarrow G$
If $(z \geq 2p + 1)$
$\quad z \leftarrow z - (2p + 1)$
$h \leftarrow \mathsf{EG}_{\text{twist}}.\mathsf{I}(1^k, p, G', z)$
Return h

Fig. 13. Definition of eeg family $\mathsf{EG}_{\text{twist-re}}$ with embedding space $\mathsf{EG}_{\text{twist-re}}.\mathsf{ES}(k, \pi) :=$ $[2^{k+1}]$. ψ_G denotes the bijection $[2p + 1] \rightarrow [\mathsf{EG}_{\text{twist}}.\mathsf{S}(1^k, p, G')]$ defined in Sect. 5.4.

The proof of the lemma can be found in the full version of this paper [4].

5.4 A Parameter-Free Family Using Range Expansion

In this section we modify the algorithms of $\mathsf{EG}_{\text{twist}}$ to obtain an embeddable group family $\mathsf{EG}_{\text{twist-re}}$ with embedding space $\mathsf{EG}_{\text{twist-re}}.\mathsf{ES}(k, \pi) = [2^{k+1}]$. The eeg family has inversion error $\mathsf{EG}_{\text{twist-re}}.\mathsf{ie}(k) = 0$ and achieves uniformly distributed embeddings. The construction is building on a technique introduced by Hayashi et al. [33], where it is used to expand the range of one way permutations. As in Sect. 5.3, the hardness sCDH-PSA with respect to $\mathsf{EG}_{\text{twist-re}}$ is based on the hardness of the sCDH problem for uniform primes with respect to TGen. The sampling algorithm—in contrast to the construction based on rejection sampling—needs access to only one uniformly random sampled integer, performs at most one exponentiation in the group and uses at most one evaluation of $\mathsf{EG}_{\text{twist}}.\mathsf{E}$ to output y with the correct distribution. Furthermore, exponents sampled by $\mathsf{EG}_{\text{twist-re}}.\mathsf{S}$ are distributed such that the eeg family achieves $\mathsf{EG}_{\text{twist-re}}.\mathsf{ie}(k) = 0$ and for every (potentially unbounded) adversary $\mathcal{A}$ we additionally have $\mathbf{Adv}_{\mathsf{EG}_{\text{twist-re}}, \mathcal{A}}^{\text{epr-psa}}(k) = 0$.

The description of $\mathsf{EG}_{\text{twist-re}}$ may be found in Fig. 13. We now discuss the construction in greater detail. Let $(G', p) = G \in [\mathsf{EG}_{\text{twist-re}}.\mathsf{G}(k, \pi)]$, where $G' = (\langle \mathbb{G} \rangle, n, g)$. The idea of the construction is to partition $[\mathsf{EG}_{\text{twist}}.\mathsf{S}(1^k, p, G')]$ into two sets M_1, M_2 with $M_1 \cup M_2 = [\mathsf{EG}_{\text{twist}}.\mathsf{S}(1^k, p, G')]$, $\{\mathsf{EG}_{\text{twist}}.\mathsf{E}(1^k, p, G', g^y) \mid y \in M_1\} = \{2^{k+1} - (2p + 1), \cdots, 2p\}$ and $\{\mathsf{EG}_{\text{twist}}.\mathsf{E}(1^k, p, G', g^y) \mid y \in M_2\} = \{0, \cdots, 2^{k+1} - (2p + 2)\}$. The sampling algorithm $\mathsf{EG}_{\text{twist-re}}.\mathsf{S}$ is constructed such that for y sampled by $\mathsf{EG}_{\text{twist-re}}.\mathsf{S}(1^k, \pi, G)$, the probability $\Pr[y = y']$ equals 2^{-k} for all $y' \in M_2$ and $2^{-(k+1)}$ for all $y' \in M_1$. Embedding algorithm $\mathsf{EG}_{\text{twist-re}}.\mathsf{E}$

on input $(1^k, \pi, G, h)$ first computes $c \leftarrow \mathsf{EG}_{\text{twist}}.\mathsf{E}(1^k, p, G', h)$. If $c \in \{2^{k+1} - (2p+1), \cdots, 2p\}$ its output remains unchanged. Otherwise it is shifted to $\{2p + 1, \cdots, 2^{k+1} - 1\}$ with probability $1/2$. In this way we achieve embeddings, which are uniformly distributed on $\mathsf{EG}_{\text{twist-re}}.\mathsf{ES}(k, \pi) = [2^{k+1}]$.

Our construction relies on the existence of a bijection $\psi_G : [2p + 1] \to [\mathsf{EG}_{\text{twist}}.\mathsf{S}(1^k, p, G')]$ for all $(G', p) = G \in [\mathsf{EG}_{\text{twist-re}}.\mathsf{G}(1^k, \pi)]$. We use the bijection, which was implicitly given in the definition of $\mathsf{EG}_{\text{twist}}.\mathsf{S}$. That is, for $z \in [2p + 1]$ we define

$$\psi_G(z) := \begin{cases} \varphi_{\text{crt}}(z, 0) & \text{if } z < n_0 \\ \varphi_{\text{crt}}(0, z - n_0 - 1) & \text{else,} \end{cases}$$

where φ_{crt} denotes the canonical isomorphism $\mathbb{Z}_{n_0} \times \mathbb{Z}_{n_1} \to \mathbb{Z}_n$.

Theorem 13. $\mathsf{EG}_{twist\text{-}re}$ *as specified in Fig. 13 is an embeddable group family with embedding space* $\mathsf{EG}_{twist\text{-}re}.\mathsf{ES}(k, \pi) = [2^{k+1}]$ *and inverson error* $\mathsf{EG}_{twist\text{-}re}.\mathsf{ie}(k) = 0$. *Furthermore* $\mathsf{EG}_{twist\text{-}re}$ *has pseudorandom embeddings. More precisely, for every (potentially unbounded) adversary $\mathcal{A}$ we have*

$$\mathbf{Adv}^{\text{epr-psa}}_{\mathsf{EG}_{twist\text{-}re}, \mathcal{A}}(k) = 0.$$

The proof of the theorem can be found in the full version of this paper [4]. As in the case of $\mathsf{EG}^{\ell}_{\text{twist-rs}}$, we obtain that—assuming that TGen invoked on randomly sampled prime p returns a secure curve-twist pair—sCDH-PSA with respect to eeg family $\mathsf{EG}_{\text{twist-re}}$ is hard.

Lemma 14. *Let* $\mathsf{EG}_{twist\text{-}re}$ *be the eeg family defined above with underlying curve-twist generator* TGen. *If the sCDH assumption holds with respect to* TGen, *then* sCDH-PSA *is hard with respect to* $\mathsf{EG}_{twist\text{-}re}$. *Concretely, for every adversary $\mathcal{A}$ against* $\mathbf{G}^{\text{scdh-psa}}_{\mathsf{EG}_{twist\text{-}re}, \mathcal{A}}(\cdot)$ *making at most Q queries to its DDH-oracle there exist adversaries $\mathcal{B}_0$, $\mathcal{B}_1$ against* $\mathbf{G}^{\text{twist}_0\text{-up-scdh}}_{\mathsf{TGen}, \mathcal{B}_0}(\cdot)$ *or* $\mathbf{G}^{\text{twist}_1\text{-up-scdh}}_{\mathsf{TGen}, \mathcal{B}_1}(\cdot)$ *respectively running in the same time as $\mathcal{A}$ and making at most Q queries to their DDH-oracles, which satisfy*

$$\mathbf{Adv}^{\text{scdh-psa}}_{\mathsf{EG}_{twist\text{-}re}, \mathcal{A}}(k) \leq 2 \left(\mathbf{Adv}^{\text{twist}_0\text{-up-scdh}}_{\mathsf{TGen}, \mathcal{B}_0}(k) + \mathbf{Adv}^{\text{twist}_1\text{-up-scdh}}_{\mathsf{TGen}, \mathcal{B}_1}(k) \right).$$

The proof of the lemma can be found in the full version of this paper [4].

References

1. Abdalla, M., et al.: Searchable encryption revisited: consistency properties, relation to anonymous IBE, and extensions. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 205–222. Springer, Heidelberg (2005). https://doi.org/10.1007/11535218_13

2. Abdalla, M., Bellare, M., Rogaway, P.: The oracle Diffie-Hellman assumptions and an analysis of DHIES. In: Naccache, D. (ed.) CT-RSA 2001. LNCS, vol. 2020, pp. 143–158. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-45353-9_12

3. Ateniese, G., Magri, B., Venturi, D.: Subversion-resilient signature schemes. In: Ray, I., Li, N., Kruegel, N. (eds.) ACM CCS 15: 22nd Conference on Computer and Communications Security, pp. 364–375. ACM Press, October 2015

4. Auerbach, B., Bellare, M., Kiltz, E.: Public-key encryption resistant to parameter subversion and its realization from efficiently-embeddable groups. Cryptology ePrint Archive, Report 2018/023 (2018). http://eprint.iacr.org/2018/023

5. Baignères, T., Delerablée, C., Finiasz, M., Goubin, L., Lepoint, T., Rivain, M.: Trap me if you can - million dollar curve. Cryptology ePrint Archive, Report 2015/1249 (2015). http://eprint.iacr.org/2015/1249

6. Bellare, M., Boldyreva, A., Desai, A., Pointcheval, D.: Key-privacy in public-key encryption. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 566–582. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-45682-1_33

7. Bellare, M., Boldyreva, A., Micali, S.: Public-key encryption in a multi-user setting: security proofs and improvements. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 259–274. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-45539-6_18

8. Bellare, M., Desai, A., Pointcheval, D., Rogaway, P.: Relations among notions of security for public-key encryption schemes. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 26–45. Springer, Heidelberg (1998). https://doi.org/10.1007/BFb0055718

9. Bellare, M., Fuchsbauer, G., Scafuro, A.: NIZKs with an untrusted CRS: security in the face of parameter subversion. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016 Part II. LNCS, vol. 10032, pp. 777–804. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53890-6_26

10. Bellare, M., Hoang, V.T.: Resisting randomness subversion: fast deterministic and hedged public-key encryption in the standard model. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015 Part II. LNCS, vol. 9057, pp. 627–656. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46803-6_21

11. Bellare, M., Jaeger, J., Kane, D.: Mass-surveillance without the state: strongly undetectable algorithm-substitution attacks. In: Ray, I., Li, N., Kruegel, C. (eds.) ACM CCS 15: 22nd Conference on Computer and Communications Security, pp. 1431–1440. ACM Press, October 2015

12. Bellare, M., Paterson, K.G., Rogaway, P.: Security of symmetric encryption against mass surveillance. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014 Part I. LNCS, vol. 8616, pp. 1–19. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44371-2_1

13. Bellare, M., Rogaway, P.: Random oracles are practical: a paradigm for designing efficient protocols. In: Ashby, V. (ed.) ACM CCS 93: 1st Conference on Computer and Communications Security, pp. 62–73. ACM Press, November 1993

14. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 409–426. Springer, Heidelberg (2006). https://doi.org/10.1007/11761679_25

15. Bernstein, D.J., Chou, T., Chuengsatiansup, C., Hülsing, A., Lange, T., Niederhagen, R., van Vredendaal, C.: How to manipulate curve standards: a white paper for the black hat. Cryptology ePrint Archive, Report 2014/571 (2014). http://eprint.iacr.org/2014/571

16. Bernstein, D.J., Duif, N., Lange, T., Schwabe, P., Yang, B.-Y.: High-speed high-security signatures. In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 124–142. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-23951-9_9

17. Bernstein, D.J., Hamburg, M., Krasnova, A., Lange, T.: Elligator: elliptic-curve points indistinguishable from uniform random strings. In: Sadeghi, A.-R., Gligor, V.D., Yung, M. (eds.) ACM CCS 13: 20th Conference on Computer and Communications Security, pp. 967–980. ACM Press, November 2013

18. Bernstein, D.J., Lange, T.: SafeCurves: choosing safe curves for elliptic-curve cryptography. https://safecurves.cr.yp.to. Accessed 18 May 2016

19. Bernstein, D.J., Lange, T., Niederhagen, R.: Dual EC: a standardized back door. Cryptology ePrint Archive, Report 2015/767 (2015). http://eprint.iacr.org/2015/767

20. Canetti, R., Pass, R., Shelat, A.: Cryptography from sunspots: how to use an imperfect reference string. In: 48th Annual Symposium on Foundations of Computer Science, pp. 249–259. IEEE Computer Society Press, October 2007

21. Checkoway, S., Cohney, S., Garman, C., Green, M., Heninger, N., Maskiewicz, J., Rescorla, E., Shacham, H., Weinmann, R.-P.: A systematic analysis of the juniper dual EC incident. In: Proceedings of the 23rd ACM conference on Computer and communications security. ACM (2016)

22. Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 13–25. Springer, Heidelberg (1998). https://doi.org/10.1007/BFb0055717

23. Cramer, R., Shoup, V.: Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. SIAM J. Comput. **33**(1), 167–226 (2003)

24. Degabriele, J.P., Farshim, P., Poettering, B.: A more cautious approach to security against mass surveillance. In: Leander, G. (ed.) FSE 2015. LNCS, vol. 9054, pp. 579–598. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48116-5_28

25. Degabriele, J.P., Paterson, K.G., Schuldt, J.C.N., Woodage, J.: Backdoors in pseudorandom number generators: possibility and impossibility results. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016 Part I. LNCS, vol. 9814, pp. 403–432. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53018-4_15

26. Dodis, Y., Ganesh, C., Golovnev, A., Juels, A., Ristenpart, T.: A formal treatment of backdoored pseudorandom generators. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015 Part I. LNCS, vol. 9056, pp. 101–126. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46800-5_5

27. Flori, J.-P., Plût, J., Reinhard, J.-R., Ekerå, M.: Diversity and transparency for ECC. Cryptology ePrint Archive, Report 2015/659 (2015). http://eprint.iacr.org/

28. Frey, G.: How to disguise an elliptic curve (Weil descent). Talk given at ECC 1998 (1998)

29. Galbraith, S.D., McKee, J.: The probability that the number of points on an elliptic curve over a finite field is prime. J. Lond. Math. Soc. **62**(3), 671–684 (2000)
30. Garg, S., Goyal, V., Jain, A., Sahai, A.: Bringing people of different beliefs together to do UC. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 311–328. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19571-6_19
31. Goldwasser, S., Micali, S.: Probabilistic encryption. J. Comput. Syst. Sci. **28**(2), 270–299 (1984)
32. Groth, J., Ostrovsky, R.: Cryptography in the multi-string model. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 323–341. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74143-5_18
33. Hayashi, R., Okamoto, T., Tanaka, K.: An RSA family of trap-door permutations with a common domain and its applications. In: Bao, F., Deng, R., Zhou, J. (eds.) PKC 2004. LNCS, vol. 2947, pp. 291–304. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24632-9_21
34. Kaliski Jr., B.S.: One-way permutations on elliptic curves. J. Cryptol. **3**(3), 187–199 (1991)
35. Katz, J., Kiayias, A., Zhou, H.-S., Zikas, V.: Distributing the setup in universally composable multi-party computation. In: Halldórsson, M.M., Dolev, S. (eds.) 33rd ACM Symposium Annual on Principles of Distributed Computing, pp. 20–29. Association for Computing Machinery, July 2014
36. Lochter, M., Mekle, J.: RFC 5639: ECC Brainpool Standard Curves & Curve Generation. Internet Engineering Task Force, March 2010
37. Möller, B.: A public-key encryption scheme with pseudo-random ciphertexts. In: Samarati, P., Ryan, P., Gollmann, D., Molva, R. (eds.) ESORICS 2004. LNCS, vol. 3193, pp. 335–351. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30108-0_21
38. NIST: Digital signature standard (DSS) 2013. FIPS PUB 186–4
39. Orman, H.: The OAKLEY key determination protocol (1998)
40. Petit, C., Quisquater, J.-J.: On polynomial systems arising from a Weil descent. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 451–466. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34961-4_28
41. Russell, A., Tang, Q., Yung, M., Zhou, H.-S.: Cliptography: clipping the power of kleptographic attacks. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016 Part II. LNCS, vol. 10032, pp. 34–64. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53890-6_2
42. Russell, A., Tang, Q., Yung, M., Zhou, H.-S.: Generic semantic security against a kleptographic adversary. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 17: 24th Conference on Computer and Communications Security, pp. 907–922. ACM Press, October 2017
43. Young, A., Yung, M.: The dark side of "black-box" cryptography, or: should we trust capstone? In: Koblitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 89–103. Springer, Heidelberg (1996)
44. Young, A., Yung, M.: Kleptography: using cryptography against cryptography. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 62–74. Springer, Heidelberg (1997). https://doi.org/10.1007/3-540-69053-0_6