On the Insecurity of Parallel Repetition for Leakage Resilience

Allison Lewko *
University of Texas at Austin
alewko@cs.utexas.edu

Brent Waters †
University of Texas at Austin
bwaters@cs.utexas.edu

Abstract

A fundamental question in leakage-resilient cryptography is: can leakage resilience always be amplified by parallel repetition? It is natural to expect that if we have a leakage-resilient primitive tolerating ℓ bits of leakage, we can take n copies of it to form a system tolerating $n\ell$ bits of leakage. In this paper, we show that this is not always true. We construct a public key encryption system which is secure when at most ℓ bits are leaked, but if we take n copies of the system and encrypt a share of the message under each using an n-out-of-n secret-sharing scheme, leaking $n\ell$ bits renders the system insecure. Our results hold either in composite order bilinear groups under a variant of the subgroup decision assumption *or* in prime order bilinear groups under the decisional linear assumption. We note that the n copies of our public key systems share a common reference parameter.

1 Introduction

Traditional security definitions for cryptographic schemes address an adversary who only has black box access to the scheme, assuming that the secret key and other internal state remains completely hidden. In practice, the adversary might gain partial knowledge of the secret key or other internal state through various side-channel and memory attacks [31, 8, 5, 7, 35, 6, 32, 40, 24, 27]. Such attacks might leverage physical phenomena like computation time, power use, etc. to deduce partial information about the secret key or state. The cold-boot attack of [27] also demonstrates that an adversary can learn noisy information about the memory contents of a machine after the machine is powered down.

Devising specific countermeasures for each known kind of attack is an unsatisfying approach, since it may require frequent updates to cryptographic systems and always leaves them potentially vulnerable to new attacks which have not yet been anticipated. A relatively new alternative approach is to develop new cryptographic security definitions that model a wide class of attacks by allowing the adversary to specify a leakage function f and learn the output of f applied to the secret key or other portions of the internal state. Clearly, there must be limits placed on the leakage function, or the adversary could learn the entire secret key and the system would be insecure. Typically, we assume that the leakage function must be efficiently computable and that the size of

*Supported by a National Defense Science and Engineering Graduate Fellowship.

†Supported by NSF CNS-0716199, CNS-0915361, and CNS-0952692, Air Force Office of Scientific Research (AFO SR) under the MURI award for "Collaborative policies and assured information sharing" (Project PRESIDIO), Department of Homeland Security Grant 2006-CS-001-000001-02 (subaward 641), and the Alfred P. Sloan Foundation.

its output is bounded by ℓ bits, where ℓ is a function of the security parameter and is less than the bit-length of the secret key.

This approach has yielded leakage-resilient constructions of many cryptographic primitives, including stream ciphers, signatures, symmetric key encryption, and public key encryption [34, 30, 21, 39, 17, 2, 3, 22, 19, 16, 12, 20, 4]. Given a construction that can tolerate ℓ bits of leakage, it is natural to ask: what if we expect even greater leakage? Recently, Alwen, Dodis, and Wichs [4] and Alwen, Dodis, Naor, Segev, Walfish, and Wichs [3] successfully employed parallel repetition to amplify leakage resilience for particular schemes and raised the fundamental question of whether leakage resilience can *always* be amplified by parallel repetition. More concretely, suppose we are given a public key encryption scheme which remains secure when ℓ bits are leaked (i.e. against an adversary who obtains ℓ bits of information about the secret key before seeing a challenge ciphertext). We can take n independent copies of the system corresponding to n public key, private key pairs. To encrypt, we now split the message into n shares, and encrypt the i^{th} share under the i^{th} public key. One may expect that this new system will remain secure up to $n\ell$ bits of leakage. Alwen et. al. [3] successfully apply this technique for specific schemes. As explained by [4, 3], it would seem quite difficult to prove this works in general, since a general reduction would need to simulate $n\ell$ bits of leakage for the parallel scheme using only ℓ bits of leakage from the original scheme.

We note that parallel repetition does hold generically if we weaken the definition of leakage resilience by restricting the leakage to a be subset of the bits representing the secret key, instead of allowing more complicated functions. This model was previously considered in [11, 18, 29]. In this setting, parallel repetition can be proven via the pigeonhole principle, since if $\leq n\ell$ bits are leaked from n keys, then there is some key for which at most ℓ bits are leaked, and security can then be proven via a reduction. (In fact, if $< n(\ell + 1)$ bits are leaked from n keys, then there is some key for which $\leq \ell$ bits are leaked.)

Though posed in the context of public key encryption, parallel repetition naturally extends to other primitives, and would be a powerful general tool for amplifying leakage resilience while preserving reasonable levels of efficiency. We note that a more basic approach to improving resilience might be to artificially increase the security parameter, λ. The success of this approach will depend on how ℓ grows as a function of λ, and it also leads to an unacceptable loss in efficiency, since many common operations require time $O(\lambda^3)$ to compute.

Our Contribution We show that there exist public key encryption schemes which are ℓ-leakage-resilient, but for which parallel repetition fails to yield an $n\ell$-leakage-resilient system for any $n > 1$. In fact, the parameters of our schemes can be chosen to rule out $\Omega(n\ell)$-leakage-resilience of the parallel schemes. Our results hold *either* under a variant of the subgroup decision assumption in composite order bilinear groups *or* under the decisional linear assumption in prime order bilinear groups. In both cases, our n parallel copies of the system share common setup parameters (i.e. are instantiated over the same group). Assuming a common group is natural in many settings, e.g. when using curves recommended by NIST [37].

Often, leakage resilience is established by employing mostly information-theoretic techniques, e.g. leveraging the fact that a function f with bounded output length cannot leak enough useful information about a key with sufficient min-entropy *even* if f is computationally unbounded. This approach is employed by [30, 34, 4, 3], for example. In the arguments of [3, 34], a computational assumption is used to argue that a valid ciphertext can be replaced by an invalid ciphertext.

However, since the adversary does not receive the ciphertext until after the leakage, it is not clear that even a computationally unbounded leakage function would allow the adversary to distinguish the two cases.[1]For these kinds of arguments, it seems plausible that if ℓ bits of leakage is not enough to compromise the security of one key, then 2ℓ bits of leakage should not be enough to compromise the security of 2 keys. (We consider the case $n = 2$ here for concreteness and will later generalize.) However, security against computationally unbounded functions f is not strictly necessary. It is possible instead to have keys with less than ℓ bits of entropy, but where it is computationally hard to compress all of the information needed for decryption into only ℓ bits.

The main idea of our approach is to design a system where it is computationally hard to represent the needed information about a single key in ℓ bits, but where two keys can be efficiently compressed into 2ℓ bits. As a first attempt at creating keys with less than ℓ bits of entropy which are computationally hard to compress, one might try using pseudorandom generators. However, it is not clear how one might find suitable structure to allow compression of two keys using this approach. Instead, we use the structure of bilinear groups. We describe our approach in terms of composite order groups for ease of exposition. We suppose we have a bilinear group G of order $N = p_1 p_2 q$, which is a product of 3 distinct primes. This group has subgroups G_{p_1}, G_{p_2}, and G_q of orders p_1, p_2, and q respectively, and whenever elements of these different subgroups are paired together under the bilinear map, the result is the identity. In this sense, the subgroups are orthogonal to each other. In our system, keys and ciphertexts will take on one of two types: type 1 keys and ciphertexts will involve only elements of G_{p_1}, while type 2 keys and ciphertexts will involve only elements of G_{p_2}. Ciphertext elements are paired with key elements in order to decrypt. A key of type 1 and a key of type 2 can be efficiently compressed into a single key by multiplying them together in the group. This new key will now decrypt ciphertexts for *both* of the private keys, since the multiplied G_{p_2} elements will not affect the result of the pairing with the type 1 ciphertext, and the multiplied G_{p_1} elements will not affect the result of the pairing with the type 2 ciphertext. Assuming for simplicity that group elements are represented by approximately $\log(N)$ bits, we can set $\ell = \frac{1}{2}\log(N)$ so that 2ℓ bits is enough to leak a group element, but ℓ bits is not.

We now have a system that is attackable when parallelized, but it is not clear that it is leakage-resilient in the first place. To prove that a single key cannot be compromised by the leakage of ℓ bits, we cannot simply make an information-theoretic argument, since either $\log(p_1)$ or $\log(p_2)$ will be less than ℓ (hence there is min-entropy $< \ell$ in at least one type of secret key). To overcome this difficulty, we introduce an expansion technique which leverages the computational bound on the leakage function. More specifically, we use the G_q subgroup as what we call an "expansion space" to argue that the secret keys have sufficiently high pseudo-entropy (i.e. their distribution is computationally indistinguishable from a distribution with high min-entropy). Relying on a close variant of the subgroup decision assumption, we expand the keys into the G_q space, and argue that an attacker cannot distinguish between elements of G_{p_1} and $G_{p_1 q}$, where $G_{p_1 q}$ denotes the subgroup of order $p_1 q$ in G (and similarly cannot distinguish between elements of G_{p_2} and $G_{p_2 q}$). We note that the expansion space G_q is shared by both key types. In this computational step of the proof, it is crucial that the leakage function f is computationally bounded (since a computationally unbounded function could distinguish the subgroups). We next expand the ciphertexts into the G_q subgroup as well, and we are then able to finish our proof with an information-theoretic argument.

[1]In fact, we conjecture that their schemes could be proven secure against a computationally unbounded leakage function under the stronger assumption that the computational problem remains hard against an adversary who is allowed unlimited preprocessing, given only the public parameters.

Prime Order Groups We also provide a system following this framework in prime order bilinear groups, under the decisional linear assumption. Again, we expand keys into an expansion space to obtain sufficient entropy. As in our composite order system, we accomplish this expansion through a computational step, this time relying on the decisional linear assumption.

Extension to Signatures and Other Primitives While we state and prove our formal result in the context of public key encryption, our methodology is broader and can be applied to show negative results for parallel repetition in other contexts. A natural application is to parallel repetition for signature systems, where one realizes parallel repetition by signing the same message under n different signing keys.

We sketch how our technique can be extended to provide a counterexample to parallel repetition for signature schemes, using the framework developed by Katz and Vaikuntanathan [30]. Katz and Vaikuntanathan construct signature schemes such that each public key corresponds to exponentially many secret keys, and given a public key, secret key pair, it is hard to compute a different secret key corresponding to the same public key. They obtain this property by using a universal one-way hash function H with domain $\{0,1\}^n$ and range $\{0,1\}^{n^\epsilon}$. The secret key of the signature scheme is an n-bit string x, and the public key is $(y = H(x), pk, r)$, where pk is a public key for a CPA-secure public key encryption scheme, and r is a common reference string for an unbounded simulation-sound NIZK proof system [15, 41]. To sign a message m, the signer computes $C = Enc(pk, m||x)$, and also supplies a proof π that C is an encryption of $m||x'$ for some x' such that $H(x') = y$. Note that for most choices of x, $H(x)$ will have many pre-images and that finding two secret keys for the same public key here corresponds to finding a collision for H. To show leakage resilience, Katz and Vaikuntanathan prove that the secret key used by the signer has high min-entropy in the adversary's view, even after the adversary has observed signatures on its chosen messages. They then show how to use a forgery to compute a new secret key for the public key. When the leakage is bounded below the min-entropy of the secret key, this new secret key will not match the original with a non-negligible probability, which contradicts the hardness of finding a different secret key.

We define a similar signature scheme as follows. We use a bilinear group G of order $N = p_1p_2q$. Our keys will be either of type 1 or type 2. A type t secret key is a random element $u \in G_{p_t}$, and the corresponding public key is comprised of $A = e(u, g_{p_t})$, pk, and r, where pk is a public key for a CPA-secure public key encryption scheme, and r is a common reference string for an unbounded simulation-sound NIZK proof system. To sign a message, the signer computes $C = Enc(pk, m||u)$, and also supplies a proof π that C is an encryption of $m||u'$ for some u' such that $e(u', g_{p_t}) = A$. We note that there are many such $u' \in G_{p_tq}$, since the part of u' in the subgroup G_q is unconstrained. Also, finding such a $u' \neq u$ violates the subgroup decision assumption, since it would allow us to produce an element u/u' of G_q. To prove leakage resilience of this scheme, we first apply our expansion technique to expand all private keys into G_q. We then employ the proof of [30]. To attack a parallel version of this system with two copies constructed from a common bilinear group G, we assume the first secret key, u_1, is type 1 and the second secret key, u_2, is type 2. We let the leakage function be $f(u_1, u_2) = u_1 \cdot u_2$. This value now serves as a secret key for either copy of the system, since G_{p_1} and G_{p_2} are orthogonal.

An alternate counterexample for signatures (due to Wichs [43]) is also discussed in Section 9.

Related Work Alwen et. al. [3] provide a counterexample to the security of parallel repetition in a more restricted setting where the public key setup is done by a single trusted authority holding

a master secret key who additionally employs an n-out-of-n secret sharing scheme. In this system, leakage resilience cannot be amplified beyond the size of the master secret key. Such a setup occurs, for example, when an IBE scheme is employed. In contrast, our counterexample requires only that the n copies of the PKE scheme share the same underlying group.

Once our keys occupy the expansion space, the rest of our proof strategy is very similar to the machinery of hash proof systems (HPS), a primitive introduced by Cramer and Shoup [14] and used by Naor and Segev [34] to obtain leakage resilient PKE schemes.

More generally, various forms of leakage resilience have been studied in many previous works [42, 38, 30, 4, 12, 16, 20, 18, 29, 19, 34, 2, 3, 11, 17, 21, 28, 33, 39, 22]. Several models of leakage resilience have been proposed, differing primarily in the restrictions placed on the leakage functions and the internal state they are applied to. We discuss the key features and distinctions of these approaches below, organizing references according to their models.

Exposure-resilient cryptography [11, 18, 29] considered an adversary who could only learn a limited subset of the secret key bits, while [28] considered an adversary who could only learn the values on certain wires of the circuit implementing a computation. For models allowing arbitrary efficiently computable leakage functions f, one can choose to bound the amount of leakage totally (bounding the total leakage over the lifetime of the system) or locally (bounding the amount of leakage per usage, e.g. per signature generated by a leakage-resilient signature scheme). A local bound is only reasonable if the internal state is continually updated, and the amount of leakage between updates is bounded. (If the secret key is unchanging, and one can leak an arbitrary ℓ bits of it many times, then an attacker will eventually learn the entire secret key.) A total bound is employed e.g. by [2, 30, 34, 4, 3], while a local bound is employed e.g. by [22, 21, 39].

There is also a distinction between models which allow the leakage function to depend only on the secret key and models where the leakage function can depend on additional internal state. For schemes where the secret key is the only internal state, the secret key is a natural choice for the input to the leakage functions. For signatures, for example, the signer may maintain additional state. Micali and Reyzin [33] introduce the assumption that "only computation leaks information". Under this assumption, one may define the input to the leakage function to be the portion of the internal state which is accessed on that particular invocation. This approach is employed by [22] for stateful signatures with a local leakage bound, for example.

A general approach to tolerating leakage that is less than the length of the secret key is to guarantee that the secret key will have sufficient min-entropy conditioned on the leakage. The works [30, 2, 4, 3, 21, 39], for example, fall into this framework. Another possibility is considered by [17], who present schemes that can tolerate leakage of arbitrary length if the secret key remains sufficiently difficult to compute from the leakage (in this case, it is possible the secret key is information-theoretically determined by the leakage). One difficulty with this approach is that it may be hard to decide if a particular collection of possible leakage functions satisfy this criterion.

2 Organization

In Section 3, we give the necessary background. In Section 4, we give our PKE system in composite order bilinear groups. In Section 5, we prove it is leakage-resilient up to ℓ bits of leakage. In Section 6, we present an attack on the parallel version of our system with $n\ell$ bits of leakage. In Section 7, we give our PKE system in prime order bilinear groups. In Section 8, we discuss instantiations of our system in specific groups. In Section 9, we discuss variations on our system

and attack and an alternative counterexample for signatures. In Section 10, we discuss possible extensions of our work.

3 Background

3.1 Leakage-resilient Public Key Encryption

We define IND-CPA leakage-resilient public key encryption schemes in terms of the following game between a challenger and an attacker. We let λ denote the security parameter, ℓ denote the leakage parameter, and (KeyGen, Encrypt, Decrypt) denote the algorithms of the PKE scheme. (Typically, ℓ is a function of λ.)

Key Generation The challenger computes $(\mathrm{PK}, \mathrm{SK}) \leftarrow \mathrm{KeyGen}(\lambda, \ell)$ and gives PK to the attacker.

Leakage The attacker chooses a function $f : \{0,1\}^* \rightarrow \{0,1\}^\ell$ that can be computed in polynomial time and receives $f(\mathrm{SK})$ from the challenger.

Challenge The attacker chooses two messages, M_0 and M_1, and gives these to the challenger. The challenger chooses a uniformly random bit $\beta \in \{0,1\}$, and gives the attacker $\mathrm{CT} \rightarrow \mathrm{Encrypt}(M_\beta, PK)$.

Guess The attacker outputs a bit $\beta' \in \{0,1\}$.

The attacker succeeds if $\beta = \beta'$. We define the advantage of an attacker $\mathcal{A}$ in this game to be $Adv_{\mathcal{A}}(\lambda, \ell) := |Pr[\beta = \beta'] - \frac{1}{2}|$.

Definition 3.1. *A public key encryption system (KeyGen, Encrypt, Decrypt) is ℓ-leakage-resilient if all polynomial time attackers $\mathcal{A}$ have a negligible advantage in the above game.*

3.2 Parallel Repetition

We now formally state the parallel repetition question introduced by [4, 3]. We let (KeyGen, Encrypt, Decrypt) denote the algorithms of a PKE scheme. For each positive integer n, we define a new scheme, (KeyGen_n, $\mathrm{Encrypt}_n$, $\mathrm{Decrypt}_n$) as follows. KeyGen_n calls KeyGen n times to produce n pairs of public and secret keys: $(\mathrm{PK}_1, \mathrm{SK}_1), \ldots, (\mathrm{PK}_n, \mathrm{SK}_n)$. The public key is $\overline{\mathrm{PK}} = (\mathrm{PK}_1, \ldots, \mathrm{PK}_n)$ and the secret key is $\overline{SK} = (\mathrm{SK}_1, \ldots, \mathrm{SK}_n)$. $\mathrm{Encrypt}_n(M, \overline{PK})$ first splits the message M into n shares $M_1, \ldots, M_n$, using an n-out-of-n secret-sharing scheme. It then produces the ciphertext as: $\overline{\mathrm{CT}} = (\mathrm{Encrypt}(M_1, PK_1), \ldots, \mathrm{Encrypt}(M_n, PK_n))$. $\mathrm{Decrypt}_n(\overline{\mathrm{CT}}, \overline{\mathrm{SK}})$ calls $\mathrm{Decrypt}(\mathrm{Encrypt}(M_i, PK_i), SK_i)$ for each i to produce M_i, and then reconstructs the secret M from its shares.

Question 3.2. *If (KeyGen, Encrypt, Decrypt) is ℓ-leakage-resilient, then is (KeyGen_n, $\mathrm{Encrypt}_n$, $\mathrm{Decrypt}_n$) necessarily $n\ell$-leakage-resilient for each positive integer n?*

Below, we answer this question in the negative, even for $n\ell$ replaced by $\Omega(n\ell)$.

3.3 Bilinear Groups

We define bilinear groups by using a group generator $\mathcal{G}$, an algorithm which takes a security parameter λ as input and outputs a description of a bilinear group G. For prime order bilinear groups, $\mathcal{G}$ outputs (p, G, G_T, e) where p is prime, G and G_T are cyclic groups of order p, and $e : G^2 \to G_T$ is a map such that:

1. (Bilinear) $\forall g, h \in G, a, b \in \mathbb{Z}_p, e(g^a, h^b) = e(g, h)^{ab}$
2. (Non-degenerate) $\exists g \in G$ such that $e(g, g)$ has order p in G_T.

We will also use composite order bilinear groups (first introduced in [10]), where $\mathcal{G}$ outputs $(N = p_1 p_2 q, G, G_T, e)$ such that p_1, p_2, q are distinct primes, G and G_T are cyclic groups of order N, and $e : G^2 \to G_T$ is a bilinear map.

We further require that the group operations in G and G_T as well as the bilinear map e are computable in polynomial time with respect to λ. Also, we assume the group descriptions of G and G_T include generators of the respective cyclic groups. For composite order groups, we let G_{p_1}, G_{p_2}, and G_q denote the subgroups of order p_1, p_2 and q in G respectively. We note that when two elements coming from different prime order subgroups are paired together under e, the result is the identity element in G_T. In this sense, the subgroups G_{p_1}, G_{p_2}, and G_q are orthogonal to each other.

3.4 Complexity Assumptions

We will first present a version of our system in composite order bilinear groups and prove its security under the following assumption, which is a variant of the subgroup decision problem from [10].

In the assumption below, we let $G_{p_1 q}$ denote the subgroup of order $p_1 q$ in G.

Subgroup Decision Assumption Given a group generator $\mathcal{G}$ for composite order bilinear groups, we define the following distribution:

$$\mathbb{G} = (N = p_1 p_2 q, G, G_T, e) \xleftarrow{R} \mathcal{G},$$

$$g_{p_1}, Y_1 \xleftarrow{R} G_{p_1},\ g_{p_2} \xleftarrow{R} G_{p_2}, Y_q \xleftarrow{R} G_q$$

$$P = (\mathbb{G}, g_{p_1}, g_{p_2}, Y_1 Y_q),$$

$$T_1 \xleftarrow{R} G_{p_1 q},\ T_2 \xleftarrow{R} G_{p_1}.$$

We define the advantage of an algorithm $\mathcal{A}$ in breaking the subgroup decision assumption to be:

$$AdvSD_{\mathcal{G},\mathcal{A}}(\lambda) := \big| Pr[\mathcal{A}(P, T_1) = 1] - Pr[\mathcal{A}(P, T_2) = 1] \big|.$$

Definition 3.3. *We say that $\mathcal{G}$ satisfies the subgroup decision assumption if $AdvSG_{\mathcal{G},\mathcal{A}}(\lambda)$ is a negligible function of λ for any polynomial time algorithm $\mathcal{A}$.*

We also provide a translation of our system into prime order groups, where security is proven from the decisional linear assumption.

Decisional Linear Assumption Given a group generator $\mathcal{G}$ for prime order bilinear groups, we define the following distribution:

$$\mathbb{G} = (p, G, G_T, e) \xleftarrow{R} \mathcal{G},$$

$$g_0, g_1, g_2 \xleftarrow{R} G,\ r_1, r_2 \xleftarrow{R} \mathbb{Z}_p,$$
$$D = (\mathbb{G}, g_0, g_1, g_2, g_1^{r_1}, g_2^{r_2}),$$
$$T_1 = g_0^{r_1+r_2},\ T_2 \xleftarrow{R} G.$$

We define the advantage of an algorithm $\mathcal{A}$ in breaking the decisional linear assumption to be:

$$AdvDLin_{\mathcal{G},\mathcal{A}}(\lambda) := \big|Pr[\mathcal{A}(D, T_1) = 1] - Pr[\mathcal{A}(D, T_2) = 1]\big|.$$

Definition 3.4. *We say that $\mathcal{G}$ satisfies the decisional linear assumption if $AdvDLin_{\mathcal{G},\mathcal{A}}(\lambda)$ is a negligible function of λ for any polynomial time algorithm $\mathcal{A}$.*

3.5 Min-Entropy and Extractors

For a random variable (or distribution) X taking values in a finite set, we define the *min-entropy* of X to be:

$$H_\infty(X) := -\log(\max_x Pr[X = x]).$$

We define the min-entropy of X conditioned on an event E to be:

$$H_\infty(X|E) := -\log(\max_x Pr[X = x|E]).$$

We define the statistical distance between two distributions X_1 and X_2 taking values in the same finite set to be:

$$\frac{1}{2}\sum_x |Pr[X_1 = x] - Pr[X_2 = x]|.$$

If the statistical distance between X_1 and X_2 is at most ϵ, we write $X_1 \approx_\epsilon X_2$.

Definition 3.5. *A function $E : \{0,1\}^S \times \{0,1\}^d \to \{0,1\}^m$ is called a strong (k, ϵ)-*extractor *if for any distribution X on $\{0,1\}^S$ with min-entropy $\geq k$, the distribution $(E(X, D), D)$ (where D is uniformly distributed over $\{0,1\}^d$) is within statistical distance ϵ from the uniform distribution of $\{0,1\}^m \times \{0,1\}^d$.*

The existence of extractors sufficient for our purposes is established in the following lemma from [36]:

Lemma 3.6. *[36] For any parameters $\delta = \delta(S)$ and $\epsilon = \epsilon(S)$ with $\frac{1}{S} \leq \delta \leq \frac{1}{2}$ and $2^{-\delta S} \leq \epsilon \leq \frac{1}{S}$, there exists a polynomial-time computable strong $(\delta S, \epsilon)$-extractor*

$$E : \{0,1\}^S \times \{0,1\}^d \to \{0,1\}^m,$$

where $d = O(\log \epsilon^{-1} \log^2 S \log \delta^{-1}/\delta)$, and $m = \Omega(\delta^2 S / \log \delta^{-1})$.

We will also need the following lemma about min-entropy [30]. (This lemma is standard, but we include its short proof for completeness.)

Lemma 3.7. *Let X be a random variable with min-entropy h and let f be an arbitrary function with range $\{0,1\}^\ell$. For any $\tau \in [0, h-\ell]$, we define the set*

$$V_\tau := \{v \in \{0,1\}^\ell \mid H_\infty(X \mid v = f(X)) \leq h - \ell - \tau\}.$$

Then:

$$Pr[f(X) \in V_\tau] \leq 2^{-\tau}.$$

Proof. We note that for any x in the range of X and $v = f(x)$, we have:

$$Pr[X = x] = Pr[X = x \mid v = f(X)]Pr[v = f(X)].$$

(This follows from the definition of conditional probability and the fact that f is deterministic.) If $v \in V_\tau$, this yields:

$$Pr[v = f(X)] \leq 2^{-h}2^{h-\ell-\tau} = 2^{-\ell-\tau}.$$

Since $f(X)$ can only take on 2^ℓ values, we may conclude that $Pr[f(X) \in V_\tau] \leq 2^{-\tau}$. □

4 Construction for Composite Order Groups

Our PKE system for composite order groups consists of four algorithms, (Setup, KeyGen, Encrypt, Decrypt). We assume the messages to be encrypted are elements of $\{0,1\}^m$. We build our system in a bilinear group whose order is a product of three primes, $N = p_1p_2q$, and prove its security for leakage parameter ℓ (sufficiently smaller than $\log(q)$) based on the subgroup decision assumption. We will state the precise conditions on ℓ that are needed in the following section. The main idea is that secret keys in the system will have *less than* ℓ bits of entropy, but since the leakage function is constrained to be computationally efficient, it cannot be used to break the system with only ℓ bits of leakage. More specifically, the secret key and ciphertext will be in subgroup G_{p_1} (or G_{p_2}) in the real system, but the subgroup decision assumption allows us to expand them to be in G_{p_1q} (or G_{p_2q}), and now we have min-entropy greater than ℓ. However, if we make two copies of our system with a common setup phase, leakage of 2ℓ bits will allow us to leak a complete element of the group G, and we can obtain a "compressed" secret key that can decrypt for both copies of the system at once. The details of this will be given Section 6.

Setup(λ) $\rightarrow$ PP The setup algorithm takes in the security parameter λ, and chooses distinct, sufficiently large primes p_1, p_2, q (it will choose q to be much larger p_1, p_2 - we will discuss this more precisely later). It chooses a bilinear group G of order $N = p_1p_2q$, along with generators g_{p_1}, g_{p_2} of the subgroups G_{p_1} and G_{p_2} respectively. (We let G_{p_1} denote the subgroup of G of order p_1, and G_{p_2} denote the subgroup of G of order p_2.) We let $S(G)$ denote the number of bits used to represent an element of G, and $S(G_T)$ denote the number of bits used to represent an element of G_T. It also chooses a uniformly random seed $D \in \{0,1\}^d$ for a (k,ϵ)-extractor $E : \{0,1\}^{S(G_T)} \times \{0,1\}^d \rightarrow \{0,1\}^m$, where m is the bit length of the messages. We assume the parameters k, ϵ, m are chosen so that ϵ is negligible with respect to the security parameter λ, and k is significantly smaller than $\log(q)$. It outputs the public parameters:

$$\text{PP} := \{N, G, g_{p_1}, g_{p_2}, E, D\}.$$

KeyGen(PP) $\rightarrow$ (PK, SK) The key generation algorithm produces two types of keys: type 1 and type 2. It chooses a type $t \in \{1, 2\}$ randomly, and then chooses a random $u \in G_{p_t}$. It sets $A = e(u, g_{p_t})$. The public key is PK $= (t, A)$, and the secret key is SK $= u$.

Encrypt(PP, PK, M) $\rightarrow$ CT The encryption algorithm takes in the public parameters PP, a public key PK, and a message $M \in \{0, 1\}^m$. It chooses a random $s \in \mathbb{Z}_N$ and computes $C_1 = g_{p_t}^s$ (where t is the type of the public key PK) and $C_2 = E(A^s, D) \oplus M$. It outputs the ciphertext CT $= (C_1, C_2)$.

Decrypt(CT, SK) $\rightarrow M$ The decryption algorithm computes:

$$E(e(C_1, \text{SK}), D) \oplus C_2 = M.$$

5 Leakage Resilience of our PKE system in Composite Order Groups

We will prove our system is ℓ-leakage-resilient for $\ell \leq \log(q) - k - \tau$ (where λ is our security parameter, k is the min-entropy required by our extractor E, and τ is a parameter chosen so that $2^{-\tau}$ is negligible in λ). We note that an attacker $\mathcal{A}$ who has a non-negligible advantage against the scheme must have a non-negligible advantage against type 1 or type 2 keys. Since we treat p_1 and p_2 symmetrically, we can assume without loss of generality that such an adversary has a non-negligible advantage against type 1 keys. Therefore, it suffices to prove security considering only type 1 keys (the proof for type 2 keys is exactly the same, with the roles of p_1 and p_2 interchanged). Our proof will proceed by a hybrid argument over the following sequence of games:

Game$_0$: The real security game. Here the private key SK is a random group element in $\mathbb{G}_{p_1}$.

Game$_1$: The private key SK is chosen as a random element in $\mathbb{G}_{p_1 q}$.

Game$_2$: The challenge CT is generated by choosing a random $C_1 \in \mathbb{G}_{p_1}$ and setting $C_2 = E(e(\text{SK}, C_1), D) \oplus M_\beta$. (SK is generated as in Game 1.)

Game$_3$: The challenge CT is generated by choosing a random $C_1 \in \mathbb{G}_{p_1 q}$ and setting $C_2 = E(e(\text{SK}, C_1), D) \oplus M_\beta$. (SK is generated as in Game 1.)

Game$_4$: The challenge CT is generated by choosing a random $C_1 \in \mathbb{G}_{p_1 q}$ and setting C_2 as a uniformly random string.

To transition from Game$_0$ to Game$_1$, we employ our expansion technique: private keys are expanded into the G_q space, and an adversary cannot notice this without violating the subgroup decision assumption. The transition to Game$_2$ is made easily, since it is identically distributed to Game$_1$. To transition from Game$_2$ to Game$_3$, we again employ our expansion technique, this time expanding the ciphertext into the G_q space. We note that $e(\text{SK}, C_1)$ now involves a term from the G_q subgroup, and this term provides sufficient min-entropy to transition to Game$_4$ via an information-theoretic argument. The attacker has advantage 0 in Game$_4$, since the ciphertext is independent of the bit β. We prove these games are indistinguishable in the following lemmas.

Lemma 5.1. *Suppose there exists a polynomial time algorithm $\mathcal{A}$ such that $Game_0Adv_{\mathcal{A}} - Game_1Adv_{\mathcal{A}} = \delta$. Then we can build a polynomial time algorithm $\mathcal{B}$ with advantage δ in breaking the subgroup decision assumption.*

Proof. $\mathcal{B}$ receives $N, G, g_{p_1}, g_{p_2}, Y_1Y_q, T$. It chooses a uniformly random seed $D \in \{0,1\}^d$ and a (k,ϵ)-extractor $E : \{0,1\}^{S(G_T)} \times \{0,1\}^d \to \{0,1\}^m$. It sets the public parameters as

$$\text{PP} := \{N, G, g_{p_1}, g_{p_2}, E, D\}$$

and gives these to $\mathcal{A}$. Next, it sets $u = T$, i.e. SK $= T$, and PK $= (1, A = e(u, g_{p_1}))$. It gives PK to $\mathcal{A}$. When $\mathcal{A}$ chooses the leakage function, $\mathcal{B}$ computes $f(T)$ and sends this to $\mathcal{A}$. $\mathcal{A}$ then sends $\mathcal{B}$ two messages, M_0 and M_1. $\mathcal{B}$ chooses a random bit β and a random $s \in \mathbb{Z}_N$. It sets $C_1 = g_{p_1}^s$ and $C_2 = E(A^s, D) \oplus M_\beta$. It gives $\mathcal{A}$ the ciphertext CT $= (C_1, C_2)$.

If $T \in G_{p_1}$, then $\mathcal{B}$ has properly simulated Game_0. If $T \in G_{p_1q}$, then $\mathcal{B}$ has properly simulated Game_1. Thus, $\mathcal{B}$ can use the output of $\mathcal{A}$ to achieve advantage δ in breaking the subgroup decision assumption. □

Lemma 5.2. *For any algorithm $\mathcal{A}$, $Game_1Adv_{\mathcal{A}} = Game_2Adv_{\mathcal{A}}$.*

Proof. This simply follows from the fact that Game_1 and Game_2 are identically distributed (note that the additional component of G_q now included in SK will not effect the value of $e(\text{SK}, C_1)$, since $C_1 \in G_{p_1}$, and G_{p_1} and G_q are orthogonal under the pairing e). □

Lemma 5.3. *Suppose there exists a polynomial time algorithm $\mathcal{A}$ such that $Game_2Adv_{\mathcal{A}} - Game_3Adv_{\mathcal{A}} = \delta$. Then we can build a polynomial time algorithm $\mathcal{B}$ with advantage δ in breaking the subgroup decision assumption.*

Proof. $\mathcal{B}$ receives $N, G, g_{p_1}, g_{p_2}, Y_1Y_q, T$. It chooses a uniformly random seed $D \in \{0,1\}^d$ and a (k,ϵ)-extractor $E : \{0,1\}^{S(G_T)} \times \{0,1\}^d \to \{0,1\}^m$. It sets the public parameters as

$$\text{PP} := \{N, G, g_{p_1}, g_{p_2}, E, D\}$$

and gives these to $\mathcal{A}$. It sets $u = Y_1Y_q$, i.e. SK $= Y_1Y_q$, and PK $= (1, e(u, g_{p_1}))$, and gives PK to $\mathcal{A}$. When $\mathcal{A}$ chooses the leakage function, $\mathcal{B}$ computes $f(u)$ and sends this to $\mathcal{A}$. $\mathcal{A}$ then sends $\mathcal{B}$ two messages, M_0 and M_1. $\mathcal{B}$ chooses a random bit β and a random $s' \in \mathbb{Z}_N$. It sets $C_1 = T^{s'}$ and $C_2 = E(e(\text{SK}, C_1), D) \oplus M_\beta$. It gives CT $= (C_1, C_2)$ to $\mathcal{A}$.

If $T \in G_{p_1}$, then $\mathcal{B}$ has properly simulated Game_2. If $T \in G_{p_1q}$, then $\mathcal{B}$ has properly simulated Game_3. Thus, $\mathcal{B}$ can use the output of $\mathcal{A}$ to achieve advantage δ in breaking the subgroup decision assumption. □

Lemma 5.4. *For any polynomial time algorithm $\mathcal{A}$, $Game_3Adv_{\mathcal{A}} - Game_4Adv_{\mathcal{A}}$ is negligible.*

Proof. We prove that with all but negligible probability, the distributions of Game_3 and Game_4 are negligibly close in $\mathcal{A}$'s view. In Game_3, the value u can be written as $u = u_1u_q$, where $u_1 \in G_{p_1}$ and $u_q \in G_q$. By the orthogonality of G_{p_1} and G_q, the public key element $e(u, g_{p_1})$ only contains information about u_1, and reveals no information about u_q, which is distributed as a random element of G_q. Thus, even conditioned on PK and C_1, the value $e(\text{SK}, C_1)$ has min-entropy $\log(q)$ from the attacker's perspective.

Since the attacker also receives $f(u)$ for the leakage function f with range $\{0,1\}^\ell$, we invoke Lemma 3.7 to assert that with probability $\geq 1 - 2^{-\tau}$, $e(\mathrm{SK}, C_1)$ will still have min-entropy $\geq k$ as long as $\log(q) - \ell - \tau \geq k$, i.e. $\ell \leq \log(q) - k - \tau$. Therefore, when $\ell \leq \log(q) - k - \tau$, $E(e(\mathrm{SK}, C_1), D)$ has statistical distance at most ϵ from a uniformly random string of length m with all but negligible probability. Since ϵ and $2^{-\tau}$ are chosen to be negligible in the security parameter λ, we have that $Game_3 Adv_{\mathcal{A}} - Game_4 Adv_{\mathcal{A}}$ is negligible as well. (We note here that it is crucial for the adversary to choose the leakage function *before* seeing the ciphertext. Otherwise, the leakage function f could change when we replace $E(e(\mathrm{SK}, C_1), D)$ with a random string of length m, and we would not be able to argue the indistinguishability of the games.) □

This completes our proof of the following theorem:

Theorem 5.5. *For $\ell \leq \log(q) - k - \tau$, our PKE system is ℓ-leakage-resilient.*

6 Attack on the Parallel System for $n > 1$ with Leakage $n\ell$

We first describe our attack for $n = 2$ when 2ℓ bits of leakage is sufficient to return a whole group element. We then generalize the attack to higher values of n. We recall that $S(G)$ denotes the number of bits representing an element of G, and we will discuss the value of $S(G)$ for particular groups in Section 8.

6.1 Attack for $n = 2$ when $2\ell \geq S(G)$

We define the system (Setup, KeyGen_2, $\mathrm{Encrypt}_2$, $\mathrm{Decrypt}_2$) as before, except we additionally assume that the two copies of the system are generated by a common setup phase, so the public parameters are shared. In other words, Setup is called only once, and outputs a single set of public parameters PP. KeyGen_2 then calls KeyGen twice on the *same* public parameters PP to generate two keys, $(\mathrm{PK}_1, \mathrm{SK}_1)$ and $(\mathrm{PK}_2, \mathrm{SK}_2)$. $\mathrm{Encrypt}_2$ splits the message M into two shares M_1 and M_2. It encrypts the first share by calling $\mathrm{Encrypt}(\mathrm{PK}_1, M_1)$ and encrypts the second share by calling $\mathrm{Encrypt}(\mathrm{PK}_2, M_2)$. $\mathrm{Decrypt}_2$ calls Decrypt on the first encrypted share with SK_1 to obtain the share M_1, and calls Decrypt on the second encrypted share with SK_2 to obtain M_2. It then reconstructs M from its shares.

The attacker receives the public parameters, PP, and two public keys, PK_1 and PK_2. If the two public keys are of the same type, the attacker aborts (this occurs with probability $\frac{1}{2}$) and guesses β' randomly. Otherwise, we assume PK_1 is of type 1 and PK_2 is of type 2. The attacker chooses the leakage function

$$f(\mathrm{SK}_1, \mathrm{SK}_2) = \mathrm{SK}_1 \cdot \mathrm{SK}_2,$$

which is permitted as long as $2\ell \geq S(G)$ (recall that $S(G)$ denotes the number of bits used to represent an element of the group G). Assuming this holds, the attacker receives the value $\mathrm{SK}_1 \cdot \mathrm{SK}_2$, and can use this to decrypt *both* ciphertexts. For example, to decrypt the ciphertext (C_1, C_2) for PK_1, the attacker computes:

$$E(e(C_1, \mathrm{SK}_1 \cdot \mathrm{SK}_2), D) \oplus C_2 = E(e(C_1, \mathrm{SK}_1), D) \oplus C_2,$$

since $C_1 \in G_{p_1}$ and $\mathrm{SK}_2 \in G_{p_2}$. This yields the first message share, and the second message share is obtained similarly, since the first ciphertext value for PK_2 will be orthogonal to SK_1. Hence, the

attacker can reconstruct the encrypted message and succeed with probability 1 when the keys are of different types. This gives the attacker an overall advantage of $\frac{1}{4}$.

When $S(G) \leq 2(\log(q) - k - \tau)$, we can choose a single value of ℓ such that $\ell \leq \log(q) - k - \tau$ and $2\ell \geq S(G)$. This will yield a PKE system (Setup, KeyGen, Encrypt, Decrypt) that is ℓ-leakage-resilient, but (Setup, KeyGen_2, Encrypt_2, Decrypt_2) is *not* 2ℓ-leakage-resilient. We note that we will choose q to be significantly larger than p_1 and p_2 in order to satisfy $S(G) \leq 2(\log(q) - k - \tau)$ for groups G of order $N = p_1 p_2 q$.

6.2 General Values of n

We now consider attacking the more general case, where we have n public and secret key pairs, with a common setup phase. On average, the attacker can expect that close to half of the n keys will be of type 1 and the others will be of type 2. More specifically, a Chernoff bound implies that for any positive constant c, there will be at least $\frac{n}{2} - c\sqrt{n}$ keys of each type with probability $\geq 1 - 2e^{-2c^2}$.

When this distribution of keys occurs, the attacker can organize the keys into at least $\frac{n}{2} - c\sqrt{n}$ pairs of keys and at most $2c\sqrt{n}$ individual keys, where each pair of keys contains one key of type 1 and one key of type 2. The attacker then defines the leakage function so that it reveals the product of each pair of secret keys and all of the remaining unpaired secret keys. This requires $(\frac{n}{2} - c\sqrt{n})S(G) + 2c\sqrt{n}S(G) = S(G)(\frac{n}{2} + c\sqrt{n})$ bits of leakage. We note that the attacker can now decrypt messages encrypted under *any* of the n public keys.

As long as

$$n\ell \geq S(G)\left(\frac{n}{2} + c\sqrt{n}\right) \Leftrightarrow \ell \geq S(G)\left(\frac{1}{2} + \frac{c}{\sqrt{n}}\right)$$

holds, this is a valid attack with $n\ell$ bits of leakage that succeeds with probability $\geq 1 - 2e^{-2c^2}$.

We recall that $\ell \leq \log(q) - k - \tau$ is required for our proof that a single copy of our system is ℓ-leakage-resilient. For this condition and the attack condition $\ell \geq S(G)(\frac{1}{2} + \frac{c}{\sqrt{n}})$ to hold simultaneously for some value of ℓ, it suffices to have a group G such that:

$$S(G) \leq \frac{2}{1 + \frac{2c}{\sqrt{n}}}(\log(q) - k - \tau).$$

Improving the Compression Factor We have thus far described a system with two key types and a corresponding attack where two keys of different types can essentially be "compressed" into one key of the same size. This will work well for leakage $n\ell$ when $S(G) < 2\log(N)$ and q is chosen to be sufficiently large with respect to p_1 and p_2. To adapt our attack to work with leakage $\gamma\ell n$ for any fixed constant $\gamma > 0$ or for groups where $S(G)$ may be larger than $2\log(N)$, we can improve the compression factor by allowing $T > 2$ key types, where T keys of distinct types can be compressed into one key (assuming $n \geq T$). We discuss this more in Section 9.

7 Realizing our Construction in Prime Order Groups

We now describe a realization of our system in prime order groups, and prove it is ℓ-leakage-resilient from the decisional linear assumption. The main idea of our construction is unchanged: we begin with keys that have less than ℓ bits of entropy, and we use the decisional linear assumption to expand them into a larger space, where they will have entropy much greater than ℓ.

We start with a bilinear group G of prime order p, generated by g. Previously, we made use of the orthogonal subgroups G_{p_1}, G_{p_2}, and G_q in our bilinear group of order N. To make suitable analogs of these subgroups using the prime order group G, we will construct "orthogonal" subgroups of G^j for some (relatively small) positive integer j. For a vector $\vec{x} = (x_1, \ldots, x_j) \in \mathbb{Z}_p^j$, we write $g^{\vec{x}}$ to denote the j-tuple of elements $(g^{x_1}, \ldots, g^{x_j})$ in G^j. We let

$$< g^{\vec{x}}, g^{\vec{y}} > := \{(g^{ax_1+by_1}, g^{ax_2+by_2}, \ldots, g^{ax_j+by_j}) \mid a, b \in \mathbb{Z}_j\}$$

denote the subgroup of G^j generated by $g^{\vec{x}}$ and $g^{\vec{y}}$. When we write an expression of the form $g^{\vec{x}}g^{\vec{y}}$, we mean componentwise multiplication, i.e. $g^{\vec{x}}g^{\vec{y}} = g^{\vec{x}+\vec{y}}$.

We define the map $e_j : G^j \times G^j \to G_T$ by:

$$e_j(g^{\vec{x}}, g^{\vec{y}}) = \prod_{i=1}^{j} e(g^{x_i}, g^{y_i}) = e(g, g)^{\vec{x}\cdot\vec{y}},$$

where e is the bilinear map from $G \times G$ into G_T. We note that $e_j(g^{\vec{x}}, g^{\vec{y}})$ is the identity element in G_T when $\vec{x}$ and $\vec{y}$ are orthogonal as vectors over $\mathbb{Z}_p$. The orthogonal subgroups G_{p_1} and G_{p_2} in composite order groups can now be replaced with subgroups $< g^{\vec{x}}, g^{\vec{y}} >$ and $< g^{\vec{v}}, g^{\vec{z}} >$ in G^j where $\vec{v}, \vec{z}$ are each orthogonal to both $\vec{x}$ and $\vec{y}$. G_q will be replaced by the subgroup comprised of elements $g^{\vec{w}} \in G^j$ where $\vec{w}$ is orthogonal to all of $\vec{x}, \vec{y}, \vec{v}, \vec{z}$. The technique of using orthogonal vectors over $\mathbb{Z}_p$ to simulate orthogonal subgroups in a prime order group was also employed by Freeman [23], though his results do not encompass our construction.

7.1 Construction

Our system in a prime order group G can now be described as follows:

Setup$(\lambda) \to$ PP The setup algorithm takes in the security parameter, and chooses a sufficiently large prime p. It chooses a bilinear group G of order p, along with a generator g and a suitable integer $j > 5$. We let $S(G)$ denote the number of bits used to represent an element of G, and $S(G_T)$ denote the number of bits used to represent an element of G_T. It also chooses a uniformly random seed $D \in \{0,1\}^d$ for a (k, ϵ)-extractor $E : \{0,1\}^{S(G_T)} \times \{0,1\}^d \to \{0,1\}^m$, where m is the bit length of the messages. We assume the parameters k, ϵ, m are chosen so that ϵ is negligible with respect to the security parameter λ, and k is $\leq \frac{1}{4} \log p$. It next chooses uniformly random vectors $\vec{x}, \vec{y} \in \mathbb{Z}_p^j$. Vectors $\vec{v}, \vec{z}$ are then chosen uniformly at random from the space of vectors which are orthogonal to both $\vec{x}$ and $\vec{y}$.

It outputs the public parameters:

$$\text{PP} := \{G, p, g, j, g^{\vec{x}}, g^{\vec{y}}, g^{\vec{v}}, g^{\vec{z}}, E, D\}.$$

KeyGen$(\text{PP}) \to (\text{PK}, \text{SK})$ The key generation algorithm produces two types of keys: type 1 and type 2. It chooses a type $t \in \{1, 2\}$ randomly, and then chooses random values $u_1, u_2 \in \mathbb{Z}_p$. If the type $t = 1$, it sets

$$A_1 = e_j(g^{u_1\vec{x}+u_2\vec{y}}, g^{\vec{x}}),\ A_2 = e_j(g^{u_1\vec{x}+u_2\vec{y}}, g^{\vec{y}}).$$

The public key is $\text{PK} = (1, A_1, A_2)$, and the secret key is $\text{SK} = g^{(u_1\vec{x}+u_2\vec{y})}$. If the type $t = 2$, it sets

$$A_1 = e_j(g^{u_1\vec{v}+u_2\vec{z}}, g^{\vec{v}}),\ A_2 = e_j(g^{u_1\vec{v}+u_2\vec{z}}, g^{\vec{z}}).$$

The public key is $\text{PK} = (2, A_1, A_2)$, and the secret key is $\text{SK} = g^{(u_1\vec{v}+u_2\vec{z})}$.

Encrypt$(\mathrm{PP}, \mathrm{PK}, M) \to \mathrm{CT}$ The encryption algorithm takes in the public parameters PP, a public key PK, and a message $M \in \{0,1\}^m$. It chooses random values $s_1, s_2 \in \mathbb{Z}_p$. If the type $t = 1$, it computes

$$C_1 = g^{s_1\vec{x}+s_2\vec{y}},\ C_2 = E(A_1^{s_1} \cdot A_2^{s_2}, D) \oplus M.$$

If the type $t = 2$, it computes

$$C_1 = g^{s_1\vec{v}+s_2\vec{z}},\ C_2 = E(A_1^{s_1} \cdot A_2^{s_2}, D) \oplus M.$$

Decrypt$(\mathrm{CT}, \mathrm{SK}) \to M$ The decryption algorithm computes:

$$E(e_j(C_1, \mathrm{SK}), D) \oplus C_2 = M.$$

This system is leakage-resilient up to $\ell = (j-5)\log p$ bits of leakage:

Theorem 7.1. *For $\ell \leq (j-5)\log(p)$, our PKE system in prime order groups is ℓ-leakage-resilient under the decisional linear assumption.*

The proof of this theorem is contained in Appendix A.1, and employs the same strategy as our proof for the composite order case. The attack on the parallel system for $n\ell$ bits of leakage is described in Appendix A.2, and is analogous to the attack for composite order groups. For $n = 2$ (for example), the attack is applicable whenever $2\ell \geq jS(G)$.

7.2 Correctness

Correctness of the algorithm is verified by observing (e.g. for type 1 keys):

$$\begin{aligned} e_j(C_1, \mathrm{SK}) &= e_j(g^{s_1\vec{x}+s_2\vec{y}}, g^{u_1\vec{x}+u_2\vec{y}}) \\ &= e(g,g)^{(s_1\vec{x}+s_2\vec{y})\cdot(u_1\vec{x}+u_2\vec{y})} \\ &= e(g,g)^{s_1((u_1\vec{x}+u_2\vec{y})\cdot\vec{x})+s_2((u_1\vec{x}+u_2\vec{y})\cdot\vec{y})} \\ &= A_1^{s_1}A_2^{s_2}. \end{aligned}$$

8 Achieving Suitable Parameters in Practical Groups

To find suitable bilinear groups in which to instantiate our prime order construction, we use supersingular elliptic curves as in Boneh and Franklin [9]. To construct these curves, we choose primes p and q such that $q = 6p - 1$. (This implies $q \equiv 2 \mod 3$.) We then consider the elliptic curve defined by the equation $y^2 = x^3 + 1$ over $\mathbb{F}_q$. As noted in [9], this curve has the following properties:

- It has $q + 1$ points over $\mathbb{F}_q$.
- The points of order $p = (q+1)/6$ form a group of order p, which we denote by G.
- Since $q \equiv 2 \mod 3$, every element of $\mathbb{Z}_q$ has precisely one cube root. Thus, for each $y_0 \in \mathbb{F}_q$, there is exactly one point (x_0, y_0) on the curve, where $x_0 = (y_0^2 - 1)^{\frac{1}{3}}$.
- We let G_T be the subgroup of $\mathbb{F}_{q^2}^*$ containing all elements of order $p = (q+1)/6$. Then the modified Weil pairing gives an efficiently computable bilinear map $e : G \times G \to G_T$.

Since cube roots are efficiently computable modulo q, we can represent a group element (x_0, y_0) of G by simply storing y_0 (since x_0 can be efficiently computed from y_0). This means that $S(G) = \log q = \log(6p-1)$, which is $\log(p)$ plus a small constant.

Choosing $j > 10$, we then have that $S(G) \leq \frac{2(j-5)}{j} \log p$ holds. This allows us to set $\ell = (j-5)\log p$ and obtain a system which is ℓ-leakage-resilient, but which is not 2ℓ-leakage-resilient when we take two parallel copies of the system. We can choose the parameters of our extractor E as $k = \frac{1}{4}\log p$ and $\epsilon = 2^{-k}$. This ensures that ϵ is negligible in λ. We can then set d, m appropriately by relying on Lemma 3.6.

To instantiate our composite order group construction, we can use the similar algorithm given in [10] for constructing a bilinear group G of a particular order N:

1. Find the smallest positive integer $a \in \mathbb{Z}$ such that $p = aN - 1$ is prime and $p \equiv 2 \mod 3$.
2. Consider the group of points on the elliptic curve $y^2 = x^3 + 1$ over $\mathbb{F}_p$. This curve has $p + 1 = aN$ points over $\mathbb{F}_p$. Thus, there is a subgroup of order N, which we will designate as G.
3. We let G_T be the subgroup of $\mathbb{F}^*_{p^2}$ of order N. The modified Weil pairing on the curve gives an efficiently computable bilinear map $e : G \times G \to G_T$.

We consider our attack for $N = p_1 p_2 q$. Our attack is applicable as long as $2\ell \geq S(G)$ (i.e. $\ell \geq \frac{S(G)}{2}$), where $\ell \leq \log q - k - \tau$, for some k and τ such that $2^{-k}, 2^{-\tau}$ are negligible in the security parameter. Now, when the a in the algorithm above happens to be quite small, $\log p$ will be approximately $\log(N) = \log p_1 + \log p_2 + \log q$. To represent a point (x_0, y_0) on the curve, we can use a single element in $\mathbb{F}_p$ (we can compute x_0 from y_0 in $\mathbb{F}_p$), which will give us $S(G)$ approximately $\log(N)$. If we choose the size of q to be sufficiently large with respect to the size of p_1, p_2 (and a is sufficiently small), we can obtain $S(G) \leq 2(\log q - k - \tau)$. This gives us a system which is secure for $\ell = \log q - k - \tau$ bits of leakage, but the parallel version of the system for $n = 2$ is insecure with 2ℓ bits of leakage.

9 Discussion

Improving the Compression Factor We could improve the compression factor for both prime and composite order groups by constructing our systems with T key types, allowing T keys of distinct types to be compressed at once. In composite order groups, this would be done by choosing a group order $N = p_1 p_2 \cdots p_T q$, where q is much larger than each of the p_i's. A key of type i would be in the subgroup G_{p_i}. The product of T keys, one of each type, would yield a single element that could be substituted for *any* of the T input keys in the decryption algorithm.

This system would still be ℓ-leakage-resilient for $\ell \leq \log(q) - k - \tau$, under the analog of the subgroup decision assumption for groups of order $N = p_1 \cdots p_T q$. For the parallel version of the system with $n \geq T$ keys, the attacker can expect roughly $\frac{n}{T}$ keys of each type. More precisely, for any constant $c > 0$, there will be at least $\frac{n}{T} - \frac{c}{T}\sqrt{n}$ keys of each type with probability $\geq 1 - Te^{-\frac{c^2}{2T}}$ (by a Chernoff bound). When this occurred, the attacker could group the n keys into $\geq \frac{n}{T} - \frac{c}{T}\sqrt{n}$ sets of T keys of distinct types, with at most $c\sqrt{n}$ individual keys remaining. With leakage at least $S(G)(\frac{n}{T} + c(1 - \frac{1}{T})\sqrt{n})$, the attacker could learn the products of all the groups of keys and all the remaining individual keys, and hence decrypt all the shares of the ciphertext.

For any fixed constant $\gamma > 0$, we can then mount our attack on the parallel system with leakage $\gamma \ell n$ as long as:

$$S(G) \leq \frac{\gamma T}{1 + \frac{c(T-1)}{\sqrt{n}}} \left(\log(q) - k - \tau\right).$$

For prime order groups (of order p), having T key types instead of T would simply require setting $j > 2T + 1$ to create enough space to simulate $T + 1$ orthogonal subgroups in $\mathbb{Z}_p^j$ with our method. Our system is leakage resilient up to ℓ bits of leakage for $\ell = (j - (2T + 1)) \log p$, for each value of T. We note the underlying security assumption (decision linear) is independent of the number of simulated subgroups and the value of j. Our attack then requires $jS(G)\left(\frac{n}{T} + c\left(1 - \frac{1}{T}\right)\sqrt{n}\right)$ bits of leakage, and this will be $\leq \gamma n \ell$ for $\ell = (j - (2T + 1)) \log p$ as long as:

$$S(G) \leq \frac{\gamma(j - (2T + 1))T}{j\left(1 + \frac{c(T-1)}{\sqrt{n}}\right)} \log(p).$$

For any fixed $\gamma > 0$, we can choose T and j large enough to meet this requirement. Hence, our counterexample shows that leakage $\Omega(n\ell)$ is not always achieved by parallel repetition.

An Alternate Counterexample for Signatures One can also obtain a counterexample to parallel repetition for signatures from any multi-signature scheme secure against "rogue-key attacks", as observed by Wichs [43]. One defines an ordinary signature scheme by using the same key generation algorithm, and having a signer simply sign each message under a set of size 1, containing only its own public key. The verification algorithm will accept as valid any signature under a set containing the correct public key which verifies under the multi-signature scheme verification algorithm (note that this accepts signatures which would never be generated by a honest signer). This scheme is secure up to some ℓ bits of leakage (at least logarithmic). The parallel system (with n copies) can then be broken for any leakage exceeding the size of a single signature under the multi-signature scheme (which is independent of n). The attacker simply asks for leakage which is a multi-signature of some message under the set containing all n public keys: this will then verify as a valid signature for each of the n signers individually. This allows the attacker to forge on only one message. In contrast, a counterexample for signatures obtained using our techniques will allow the attacker to forge as many messages as desired.

10 Future Directions

We now discuss a few approaches for avoiding or extending our counterexample.

Removing the Common Reference String The assumption of a common setup for the n copies of a parallel system is natural, since it is typical to create several public keys from one group in practical systems. For example, NIST recommends using certain elliptic curves [37]. However, it would be interesting to construct a counterexample to parallel repetition that does not rely on a common setup.

A Black-box Separation Alwen et. al. [3] suggest the potential direction of showing a black-box separation for parallel repetition. Such a result would rule out the possibility of a general

reduction using an attacker who can break the parallel scheme with $n\ell$ bits of leakage to break the original scheme with ℓ bits of leakage. This would be incomparable to our result, since a black-box separation does not imply the existence of a counterexample, and our result relies on unproven (though commonly used) assumptions for security.

Alternate Assumptions We rely on either a variant of the subgroup decision assumption in composite order bilinear groups *or* the decisional linear assumption in prime order bilinear groups. We also suspect that our results could be adapted to provide a counterexample under lattice-based assumptions, given that many results obtained using bilinear groups have also been instantiated with lattices (e.g. [26, 25, 1, 13]). Obtaining additional counterexamples under a variety of assumptions would provide stronger evidence for the insecurity of parallel repetition as a generic tool.

A Looser Bound One might ask if parallel repetition holds for some sublinear bound on the leakage as a function of n. In other words, can we take n copies of an ℓ-leakage-resilient system and always build an $f(n, \ell)$-leakage-resilient system for some sufficiently growing function $f(n, \ell) = o(n\ell)$?

Alternative Leakage Models One can also ask if parallel repetition holds for other interesting leakage models. For instance, we might strengthen the definition of leakage resilience by allowing the leakage function f to be computationally unbounded. Our counterexample no longer applies in this case. If we instead weaken the definition of leakage resilience by restricting the leakage to being a subset of the bits representing the secret key, we recall that parallel repetition *does* hold.

References

[1] S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (h)ibe in the standard model. In *EUROCRYPT*, pages 553–572, 2010.

[2] A. Akavia, S. Goldwasser, and V. Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In *TCC*, pages 474–495, 2009.

[3] J. Alwen, Y. Dodis, M. Naor, G. Segev, S. Walfish, and D. Wichs. Public-key encryption in the bounded-retrieval model. In *EUROCRYPT*, pages 113–134, 2010.

[4] J. Alwen, Y. Dodis, and D. Wichs. Leakage-resilient public-key cryptography in the bounded-retrieval model. In *CRYPTO*, pages 36–54, 2009.

[5] E. Biham, Y. Carmeli, and A. Shamir. Bug attacks. In *CRYPTO*, pages 221–240, 2008.

[6] E. Biham and A. Shamir. Differential fault analysis of secret key cryptosystems. In *CRYPTO*, pages 513–525, 1997.

[7] D. Boneh and D. Brumley. Remote timing attacks are practical. In *Computer Networks*, volume 48, pages 701–716, 2005.

[8] D. Boneh, R. A. DeMillo, and R. J. Lipton. On the importance of checking cryptographic protocols for faults. In *EUROCRYPT*, pages 37–51, 1997.

[9] D. Boneh and M. Franklin. Identity based encryption from the weil pairing. In *CRYPTO*, pages 213–229, 2001.

[10] D. Boneh, E. Goh, and K. Nissim. Evaluating 2-dnf formulas on ciphertexts. In *TCC*, pages 325–341, 2005.

[11] R. Canetti, Y. Dodis, S. Halevi, E. Kushilevitz, and A. Sahai. Exposure-resilient functions and all-or-nothing transforms. In *EUROCRYPT*, pages 453–469, 2000.

[12] D. Cash, Y. Z. Ding, Y. Dodis, W. Lee, R. J. Lipton, and S. Walfish. Intrusion-resilient key exchange in the bounded retrieval model. In *TCC*, pages 479–498, 2007.

[13] D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai trees, or how to delegate a lattice basis. In *EUROCRYPT*, pages 523–552, 2010.

[14] R. Cramer and V. Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In *EUROCRYPT*, pages 45–64, 2002.

[15] A. De Santis, G. Di Crescenzo, R. Ostrovsky, G. Persiano, and A. Sahai. Robust non-interactive zero knowledge. In *CRYPTO*, pages 566–598, 2001.

[16] D. Di Crescenzo, R. J. Lipton, and S. Walfish. Perfectly secure password protocols in the bounded retrieval model. In *TCC*, pages 225–244, 2006.

[17] Y. Dodis, Y. Kalai, and S. Lovett. On cryptography with auxiliary input. In *STOC*, pages 621–630, 2009.

[18] Y. Dodis, A. Sahai, and A. Smith. On perfect and adaptive security in exposure-resilient cryptography. In *EUROCRYPT*, pages 301–324, 2001.

[19] S. Dziembowski. Intrusion-resilience via the bounded-storage model. In *TCC*, pages 207–224, 2006.

[20] S. Dziembowski and K. Pietrzak. Intrusion-resilient secret sharing. In *FOCS*, pages 227–237, 2007.

[21] S. Dziembowski and K. Pietrzak. Leakage-resilient cryptography. In *FOCS*, pages 293–302, 2008.

[22] S. Faust, E. Kiltz, K. Pietrzak, and G. N. Rothblum. Leakage-resilient signatures. In *TCC*, pages 343–360, 2010.

[23] D. M. Freeman. Converting pairing-based cryptosystems from composite-order groups to prime-order groups. In *EUROCRYPT*, pages 44–61, 2010.

[24] K. Gandolfi, C. Mourtel, and F. Olivier. Electromagnetic analysis: Concrete results. In *CHES*, number Generators, pages 251–261, 2001.

[25] C. Gentry, S. Halevi, and V. Vaikuntanathan. A simple bgn-type cryptosystem from lwe. In *EUROCRYPT*, pages 506–522, 2010.

[26] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.

[27] A. Halderman, S. Schoen, N. Heninger, W. Clarkson, W. Paul, J. Calandrino, A. Feldman, J. Applebaum, and E. Felten. Lest we remember: Cold boot attacks on encryption keys. In *USENIX Security Symposium*, pages 45–60, 2008.

[28] Y. Ishai, A. Sahai, and D. Wagner. Private circuits: Securing hardware against probing attacks. In *CRYPTO*, pages 463–481, 2003.

[29] J. Kamp and D. Zuckerman. Deterministic extractors for bit-fixing sources and exposure-resilient cryptography. In *FOCS*, pages 92–101, 2003.

[30] J. Katz and V. Vaikuntanathan. Signature schemes with bounded leakage resilience. In *ASIACRYPT*, pages 703–720, 2009.

[31] P. C. Kocher. Timing attacks on implementations of diffie-hellman, rsa, dss, and other systems. In *CRYPTO*, pages 104–113, 1996.

[32] P. C. Kocher, J. Jaffe, and B. Jun. Differential power analysis. In *CRYPTO*, pages 388–397, 1999.

[33] S. Micali and L. Reyzin. Physically observable cryptography. In *TCC*, pages 278–296, 2004.

[34] M. Naor and G. Segev. Public-key cryptosystems resilient to key leakage. In *CRYPTO*, pages 18–35, 2009.

[35] P. Q. Nguyen and I. Shparlinski. The insecurity of the digital signature algorithm with partially known nonces. volume 15, pages 151–176, 2002.

[36] N. Nisan and D. Zuckerman. Randomness is linear in space. In *J. Comput. Syst. Sci.*, volume 52, pages 43–52, 1996.

[37] National Institute of Standards and Technology. Digital signature standard (dss), June 2009. http://csrc.nist.gov/publications/fips/fips186-3/fips_186-3.pdf.

[38] C. Petit, F.X. Standaert, O. Pereira, T. Malkin, and M. Yung. A block cipher based pseudo random number generator secure against side-channel key recovery. In *ASIACCS*, pages 56–65, 2008.

[39] K. Pietrzak. A leakage-resilient mode of operation. In *EUROCRYPT*, pages 462–482, 2009.

[40] J. Quisquater and D. Samyde. Electromagnetic analysis (ema): Measures and countermeasures for smart cards. In *E-smart*, pages 200–210, 2001.

[41] A. Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *FOCS*, pages 543–553, 1999.

[42] F.X. Standaert, T. Malkin, and M. Yung. A unified framework for the analysis of side-channel key recovery attacks. In *EUROCRYPT*, pages 443–461, 2009.

[43] D. Wichs. personal communication, 2010.

A Properties of Our System in Prime Order Groups

A.1 Leakage Resilience

We prove our system in prime order groups is ℓ-leakage-resilient for $\ell \leq (j-5)\log(p)$ under the decisional linear assumption. We first present the simulated subgroup decision assumption, which is a prime order analog of the subgroup decision assumption. We then show that the simulated subgroup decision assumption follows from the decisional linear assumption. The proof of security for our prime order group system under the simulated subgroup decision assumption then proceeds similarly to the proof for our composite order group system under the subgroup decision assumption.

Simulated Subgroup Decision Assumption Given a group generator $\mathcal{G}$ for prime order bilinear groups, we define the following distribution. Below, $j > 5$ is a fixed integer, and $\mathcal{O}_{\vec{x},\vec{y}} \subseteq \mathbb{Z}_p^j$ denotes the space of vectors orthogonal to both $\vec{x}, \vec{y} \in \mathbb{Z}_p^j$.

$$\mathbb{G} = (p, G, G_T, e) \xleftarrow{R} \mathcal{G},$$

$$\vec{x}, \vec{y} \xleftarrow{R} \mathbb{Z}_p^j,\ \vec{v}, \vec{z} \xleftarrow{R} \mathcal{O}_{\vec{x},\vec{y}},$$

$$g \xleftarrow{R} G,\ a, b \xleftarrow{R} \mathbb{Z}_p,\ \vec{w}, \vec{\gamma} \xleftarrow{R} \mathcal{O}_{\vec{v},\vec{z}},$$

$$P = (\mathbb{G}, j, g, g^{\vec{x}}, g^{\vec{y}}, g^{\vec{v}}, g^{\vec{z}}, g^{\vec{\gamma}}),$$

$$T_1 = g^{a\vec{x}+b\vec{y}},\ T_2 = g^{\vec{w}}.$$

We define the advantage of an algorithm $\mathcal{A}$ in breaking the simulated subgroup decision assumption to be:

$$AdvSSD_{\mathcal{G},\mathcal{A}}(\lambda) := \left|Pr[\mathcal{A}(P, T_1) = 1] - Pr[\mathcal{A}(P, T_2) = 1]\right|.$$

Definition A.1. *We say that $\mathcal{G}$ satisfies the simulated subgroup decision assumption if $AdvSSD_{\mathcal{G},\mathcal{A}}(\lambda)$ is a negligible function of λ for any polynomial time algorithm $\mathcal{A}$.*

Theorem A.2. *Suppose there exists a polynomial time algorithm $\mathcal{A}$ with non-negligible advantage in breaking the simulated subgroup decision assumption. Then we can create a polynomial time algorithm $\mathcal{B}$ with non-negligible advantage in breaking the decisional linear assumption.*

Proof. Our proof will proceed by a hybrid argument with two steps. We let D_1 denote the distribution (P, T_1) as defined in the simulated subgroup decision assumption above, and we let D_3 denote the distribution (P, T_2). We define an intermediate distribution, D_2, and then we will show that D_1 and D_2 are computationally indistinguishable under the decisional linear assumption, and that the statistical distance between D_2 and D_3 is negligible. This shows that the simulated subgroup decision assumption holds under the decisional linear assumption.

We define the distribution D_2 as follows:

$$\mathbb{G} = (p, G, G_T, e) \xleftarrow{R} \mathcal{G},$$

$$\vec{x}_1, \vec{y}_2, \vec{s} \xleftarrow{R} \mathbb{Z}_p^j,\ \vec{v}, \vec{z} \xleftarrow{R} \mathcal{O}_{\vec{x}_1,\vec{y}_2,\vec{s}},$$

$$g \xleftarrow{R} G,\ r_1', r_2', b_1, b_2, b_3 \xleftarrow{R} \mathbb{Z}_p,\ \vec{\gamma} \xleftarrow{R} \mathcal{O}_{\vec{v},\vec{z}},$$

$$\vec{x} := \vec{x}_1 + \frac{r'_1}{b_1}\vec{s},\ \vec{y} := \vec{y}_2 + \frac{r'_2}{b_2}\vec{s},$$
$$P_2 = (\mathbb{G}, j, g, g^{\vec{x}}, g^{\vec{y}}, g^{\vec{v}}, g^{\vec{z}}, g^{\vec{\gamma}}),$$
$$T = g^{b_1\vec{x}+b_2\vec{y}+b_3\vec{s}}.$$
$$D_2 := (P_2, T).$$

Lemma A.3. *Suppose there exists a polynomial time algorithm $\mathcal{A}$ with non-negligible advantage in distinguishing between distributions D_1 and D_2. Then there is a polynomial time algorithm $\mathcal{B}$ with non-negligible advantage in breaking the decisional linear assumption.*

Proof. $\mathcal{B}$ is given $G, p, g_0, g_1, g_2, g_1^{r_1}, g_2^{r_2}, T$. It sets $g = g_0$, and chooses random vectors $\vec{x}_1, \vec{y}_2, \vec{s} \in \mathbb{Z}_p^j$. It chooses $b_1, b_2 \in \mathbb{Z}_p$ randomly. It implicitly sets $\vec{x}$ to be a multiple of $\vec{x_1} + \frac{r_1}{b_1}\vec{s}$ by setting $g^{\vec{x}} = g_1^{\vec{x_1}+\frac{r_1}{b_1}\vec{s}}$ (note it can compute this because it knows $g_1, g_1^{r_1}, \vec{x_1}, \vec{s}$). It implicitly sets $\vec{y}$ to be a multiple of $\vec{y_2} + \frac{r_2}{b_2}\vec{s}$ by setting $g^{\vec{y}} = g_2^{\vec{y_2}+\frac{r_2}{b_2}\vec{s}}$.

$\mathcal{B}$ next chooses random vectors $\vec{v}, \vec{z}$ from the space of vectors orthogonal to $\vec{x_1}, \vec{y_2}, \vec{s}$. We note that $\vec{v}, \vec{z}$ will also be orthogonal to $\vec{x}, \vec{y}$. $\mathcal{B}$ also chooses a random vector $\vec{\gamma}$ which is orthogonal to $\vec{v}$ and $\vec{z}$. It sets

$$T' = g_0^{b_1\vec{x_1}+b_2\vec{y_2}}T^{\vec{s}}.$$

It gives $g, g^{\vec{x}}, g^{\vec{y}}, g^{\vec{v}}, g^{\vec{z}}, g^{\vec{\gamma}}, T'$ to $\mathcal{A}$.

If $T = g_0^{r_1+r_2}$, then T' is a uniformly random element of $< g^{\vec{x}}, g^{\vec{y}} >$. If T is a uniformly random element of G, then $T' = g^{\vec{w}}$, where $\vec{w}$ is a uniformly random element of the 3 dimensional space in $\mathbb{Z}_p^j$ spanned by $\vec{x_1}, \vec{y_2}, \vec{s}$, and we note that is the *same* space spanned by $\vec{x}, \vec{y}, \vec{s}$ (linear independence of these vectors holds with all but negligible probability, so we will implicitly assume this).

It is clear that $\vec{x}, \vec{y}, \vec{v}, \vec{z}, \vec{\gamma}$ are distributed identically to their distribution in D_2. We now show that their distribution also matches D_1: $\vec{x}$ and $\vec{y}$ are distributed as uniformly random vectors, and $\vec{v}$ and $\vec{z}$ are distributed as random vectors from $\mathcal{O}_{\vec{x},\vec{y}}$ ($\vec{\gamma}$ can then be ignored, as the equivalence of its distribution follows automatically). To see this, consider vectors $\vec{x}, \vec{y}, \vec{v}, \vec{z}$ such that $\vec{v}, \vec{z} \in \mathcal{O}_{\vec{x},\vec{y}}$. With all but negligible probability, $\vec{x}, \vec{y}$ are independent, so $|\mathcal{O}_{\vec{v},\vec{z}}| = p^{j-2}$. In this case, the probability of obtaining these particular vectors from D_1 is:

$$\frac{1}{p^j} \cdot \frac{1}{p^j} \cdot \frac{1}{p^{j-2}} \cdot \frac{1}{p^{j-2}} = \frac{1}{p^{4j-4}}.$$

It is clear that $\vec{x}$, $\vec{y}$ still occur with probability $\frac{1}{p^j} \cdot \frac{1}{p^j}$ under D_2, since x_1, y_2 are uniformly random. Now, $\vec{v}$ and $\vec{z}$ can be obtained from D_2 whenever $\vec{s}$ happens to be in $\mathcal{O}_{\vec{v},\vec{z}}$. With all but negligible probability, $\vec{v}$ and $\vec{z}$ are linearly independent, in which case $|\mathcal{O}_{\vec{v},\vec{z}}| = p^{j-2}$. Thus, the probability of $\vec{s}$ being chosen in this set is $\frac{p^{j-2}}{p^j}$. Conditioning on the choice of $\vec{x}, \vec{y}$, and such a $\vec{s}$, the probability of choosing $\vec{v}, \vec{z}$ from $\mathcal{O}_{\vec{x},\vec{y},\vec{s}}$ is $\frac{1}{p^{j-3}} \cdot \frac{1}{p^{j-3}}$ (assuming that $\vec{x}, \vec{y}, \vec{s}$ are linearly independent, which occurs with all but negligible probability). Thus, the total probability of $\vec{x}, \vec{y}, \vec{v}, \vec{z}$ occurring under D_2 (ignoring negligible events) is:

$$\frac{p^{j-2}}{p^j} \cdot \frac{1}{p^{j-3}} \cdot \frac{1}{p^{j-3}} \cdot \frac{1}{p^j} \cdot \frac{1}{p^j} = \frac{1}{p^{4j-4}}.$$

This shows the distributions of $\vec{x}, \vec{y}, \vec{v}, \vec{z}$ under D_1 and D_2 are equivalent.

Therefore, when $T = g_0^{r_1+r_2}$, $\mathcal{B}$ has properly simulated distribution D_1. When T is random, $\mathcal{B}$ has properly simulated distribution D_2. Thus, $\mathcal{B}$ can use the output of $\mathcal{A}$ to break the decisional linear assumption with non-negligible advantage. □

Lemma A.4. *The statistical distance between distributions D_2 and D_3 is negligible.*

Proof. We consider fixed vectors $\vec{x}, \vec{y}, \vec{v}, \vec{z}, \vec{\gamma}, \vec{w}$ such that $\vec{v}, \vec{z} \in \mathcal{O}_{\vec{x},\vec{y}}$, $\vec{\gamma}, \vec{w} \in \mathcal{O}_{\vec{v},\vec{z}}$. We may assume that $\vec{x}, \vec{y}, \vec{w}$ are linearly independent and that $\vec{v}, \vec{z}$ are linearly independent, since this happens with all but negligible probability under both distributions (recall that $j > 4$). The probability of these vectors occurring under distribution D_3 is:

$$\frac{1}{p^j} \cdot \frac{1}{p^j} \cdot \frac{1}{p^{j-2}} \cdot \frac{1}{p^{j-2}} \cdot \frac{1}{p^{j-2}} \cdot \frac{1}{p^{j-2}} = \frac{1}{p^{6j-8}},$$

since $|\mathcal{O}_{\vec{x},\vec{y}}| = p^{j-2}$ and $|\mathcal{O}_{\vec{v},\vec{z}}| = p^{j-2}$.

For distribution D_2, we consider the probability of $\vec{x}, \vec{y}, \vec{v}, \vec{z}, \vec{\gamma}, \vec{w}$ conditioned on the choice of $\vec{s}$. This probability will be zero except when $\vec{s} \in < \vec{x}, \vec{y}, \vec{w} > \setminus < \vec{x}, \vec{y} >$. The probability that such an $\vec{s}$ is chosen is $\frac{p^3-p^2}{p^j}$. Conditioning on $\vec{x}, \vec{y}$, and such an $\vec{s}$, the probability of $\vec{w}$ is then $\frac{1}{p^3}$. Also, since $\vec{w}$ is orthogonal to $\vec{v}, \vec{z}$, for any such $\vec{s}$ we also have that $\vec{v}, \vec{z}$ are orthogonal to $\vec{s}$. We note that $\vec{x}, \vec{y}$ are still uniformly distributed under D_2. Thus, the probability of $\vec{x}, \vec{y}, \vec{v}, \vec{z}, \vec{\gamma}, \vec{w}$ occurring under distribution D_2 is:

$$\frac{1}{p^j} \cdot \frac{1}{p^j} \cdot \frac{p^3 - p^2}{p^j} \cdot \frac{1}{p^{j-3}} \cdot \frac{1}{p^{j-3}} \cdot \frac{1}{p^{j-2}} \cdot \frac{1}{p^3},$$

since $|\mathcal{O}_{\vec{x},\vec{y},\vec{s}}| = p^{j-3}$ when $\vec{x}, \vec{y}, \vec{s}$ are linearly independent. This equals:

$$\frac{1}{p^{6j-5}} \left(p^3 - p^2\right) = \frac{1}{p^{6j-8}} \left(1 - \frac{1}{p}\right).$$

This shows that the statistical distance between D_2 and D_3 is negligible. □

The theorem then follows. □

In summary, we have shown that if the decisional linear assumption holds, then the simulated subgroup decision assumption holds as well. The proof of security for our system in prime order groups under the simulated subgroup decision assumption now proceeds much like the proof of security for our system in composite order groups under the subgroup decision assumption. We note that an attacker $\mathcal{A}$ who has a non-negligible advantage against the scheme must have a non-negligible advantage against type 1 or type 2 keys. Since we treat these types symmetrically, we can assume without loss of generality that such an adversary has a non-negligible advantage against type 1 keys. Therefore, it suffices to prove security considering only type 1 keys. We define the following sequence of games:

Game$_0$: The real security game. Here the private key SK is a random group element in $< g^{\vec{x}}, g^{\vec{y}} >$.

Game$_1$: The private key SK is chosen as $g^{\vec{w}}$, where $\vec{w}$ is chosen randomly from $\mathcal{O}_{\vec{v},\vec{z}}$. The public key is still computed as $A_1 = e_j(\mathrm{SK}, g^{\vec{x}}), A_2 = e_j(\mathrm{SK}, g^{\vec{y}})$.

Game$_2$: The challenge CT is generated by choosing a random $C_1 \in < g^{\vec{x}}, g^{\vec{y}} >$ and setting $C_2 = E(e_j(\mathrm{SK}, C_1), D) \oplus M_\beta$. ($\mathrm{SK}, \mathrm{PK}$ are generated as in Game$_1$.)

Game$_3$: The challenge CT is generated by choosing $C_1 = g^{\vec{s}}$, where $\vec{s}$ is chosen randomly from $\mathcal{O}_{\vec{v},\vec{z}}$, and setting $C_2 = E(e_j(\mathrm{SK}, C_1), D) \oplus M_\beta$. ($\mathrm{SK}, \mathrm{PK}$ are generated as in Game$_1$.)

Game$_4$: The challenge CT is generated by choosing $C_1 = g^{\vec{s}}$, where $\vec{s}$ is chosen randomly from $\mathcal{O}_{\vec{v},\vec{z}}$, and setting C_2 as a uniformly random string. (SK, PK are generated as in Game$_1$).

In Game$_4$, the ciphertext no longer depends on the message, so the attacker has advantage 0. We will prove these games are indistinguishable in the following lemmas.

Lemma A.5. *Suppose there exists a polynomial time algorithm $\mathcal{A}$ such that $Game_0 Adv_{\mathcal{A}} - Game_1 Adv_{\mathcal{A}} = \delta$. Then we can build a polynomial time algorithm $\mathcal{B}$ with advantage δ in breaking the simulated subgroup decision assumption.*

Proof. $\mathcal{B}$ receives $G, p, j, g, g^{\vec{x}}, g^{\vec{y}}, g^{\vec{v}}, g^{\vec{z}}, g^{\vec{\gamma}}, T$. It chooses a uniformly random seed $D \in \{0,1\}^d$ and a (k, ϵ)-extractor $E : \{0,1\}^{S(G_T)} \times \{0,1\}^d \to \{0,1\}^m$. It sets the public parameters as

$$\mathrm{PP} := \{G, p, g, j, g^{\vec{x}}, g^{\vec{y}}, g^{\vec{v}}, g^{\vec{z}}, E, D\}.$$

and gives these to $\mathcal{A}$.

Next, it sets $\mathrm{SK} = T$, and $\mathrm{PK} = (1, A_1 = e_j(T, g^{\vec{x}}), A_2 = e_j(T, g^{\vec{y}}))$. It gives PK to $\mathcal{A}$. When $\mathcal{A}$ chooses the leakage function, $\mathcal{B}$ computes $f(T)$ and sends this to $\mathcal{A}$. $\mathcal{A}$ then sends $\mathcal{B}$ two messages, M_0 and M_1. $\mathcal{B}$ chooses a random bit β and random values $s_1, s_2 \in \mathbb{Z}_p$. It sets $C_1 = g^{s_1\vec{x} + s_2\vec{y}}$ and $C_2 = E(A_1^{s_1} A_2^{s_2}, D) \oplus M_\beta$. It gives $\mathcal{A}$ the ciphertext $\mathrm{CT} = (C_1, C_2)$.

It $T \in < g^{\vec{x}}, g^{\vec{y}} >$, then $\mathcal{B}$ has properly simulated Game$_0$. If $T = g^{\vec{w}}$ for $\vec{w}$ randomly chosen from $\mathcal{O}_{\vec{v},\vec{z}}$, then $\mathcal{B}$ has properly simulated Game$_1$. Thus, $\mathcal{B}$ can use the output of $\mathcal{A}$ to achieve advantage δ in breaking the simulated subgroup decision assumption. □

Lemma A.6. *For any algorithm $\mathcal{A}$, $Game_1 Adv_{\mathcal{A}} = Game_2 Adv_{\mathcal{A}}$.*

Proof. This simply follows from the fact that Game$_1$ and Game$_2$ are identically distributed. Since A_1 is computed as $e_j(g^{\vec{w}}, g^{\vec{x}})$ and A_2 is computed as $e_j(g^{\vec{w}}, g^{\vec{y}})$, we have that

$$e_j(\mathrm{SK}, C_1) = e_j(g^{\vec{w}}, g^{s_1\vec{x} + s_2\vec{y}}) = e_j(g^{\vec{w}}, g^{\vec{x}})^{s_1} \cdot e_j(g^{\vec{w}}, g^{\vec{y}})^{s_2} = A_1^{s_1} A_2^{s_2}.$$

□

Lemma A.7. *Suppose there exists a polynomial time algorithm $\mathcal{A}$ such that $Game_2 Adv_{\mathcal{A}} - Game_3 Adv_{\mathcal{A}} = \delta$. Then we can build a polynomial time algorithm $\mathcal{B}$ with advantage δ in breaking the simulated subgroup decision assumption.*

Proof. $\mathcal{B}$ receives $G, p, j, g, g^{\vec{x}}, g^{\vec{y}}, g^{\vec{v}}, g^{\vec{z}}, g^{\vec{\gamma}}, T$. It chooses a uniformly random seed $D \in \{0,1\}^d$ and a (k, ϵ)-extractor $E : \{0,1\}^{S(G_T)} \times \{0,1\}^d \to \{0,1\}^m$. It sets the public parameters as

$$\mathrm{PP} := \{G, p, g, j, g^{\vec{x}}, g^{\vec{y}}, g^{\vec{v}}, g^{\vec{z}}, E, D\}.$$

and gives these to $\mathcal{A}$.

$\mathcal{B}$ sets SK $= g^{\vec{\gamma}}$, and PK $= (1, A_1 = e_j(g^{\vec{\gamma}}, g^{\vec{x}}), A_2 = e_j(g^{\vec{\gamma}}, g^{\vec{y}}))$. It gives PK to $\mathcal{A}$. When $\mathcal{A}$ chooses the leakage function, $\mathcal{B}$ computes $f(g^{\vec{\gamma}})$ and sends this to $\mathcal{A}$. $\mathcal{A}$ then sends $\mathcal{B}$ two messages, M_0 and M_1. $\mathcal{B}$ chooses a random bit β, and sets $C_1 = T$ and $C_2 = E(e_j(\mathrm{SK}, C_1), D) \oplus M_\beta$. It gives CT $= (C_1, C_2)$ to $\mathcal{A}$.

If T is a random element of $< g^{\vec{x}}, g^{\vec{y}} >$, then $\mathcal{B}$ has properly simulated Game$_2$. If $T = g^{\vec{w}}$ for $\vec{w}$ chosen randomly from $\mathcal{O}_{\vec{v},\vec{z}}$, then $\mathcal{B}$ has properly simulated Game$_3$. Thus, $\mathcal{B}$ can use the output of $\mathcal{A}$ to achieve advantage δ in breaking the simulated subgroup decision assumption. □

Lemma A.8. *For any polynomial time algorithm $\mathcal{A}$, Game$_3$Adv$_\mathcal{A}$ − Game$_4$Adv$_\mathcal{A}$ is negligible.*

Proof. We prove that with all but negligible probability, the distributions of Game$_3$ and Game$_4$ are negligibly close in $\mathcal{A}$'s view. In both games, the secret key SK is distributed as $g^{\vec{w}}$, where $\vec{w}$ is chosen randomly from the space $\mathcal{O}_{\vec{v},\vec{z}}$. We recall that $j > 4$, and that $\vec{x}, \vec{y}, \vec{v}, \vec{z}$ are linearly independent with high (all but negligible) probability, so we may assume this. We now show that $\mathcal{O}_{\vec{x},\vec{y},\vec{v},\vec{z}} \cap < \vec{x}, \vec{y}, \vec{v}, \vec{z} >= \{0\}$ holds with all but negligible probability. This will allow us to fix a basis of $\mathbb{Z}_p^j$ consisting of $\vec{x}, \vec{y}, \vec{v}, \vec{z}$ and a basis for the $\mathcal{O}_{\vec{x},\vec{y},\vec{v},\vec{z}}$.

We note that $\mathcal{O}_{\vec{x},\vec{y}} \cap < \vec{x}, \vec{y} >\neq \{0\}$ can only happen if the matrix

$$\begin{pmatrix} \vec{x}\cdot\vec{x} & \vec{x}\cdot\vec{y} \\ \vec{x}\cdot\vec{y} & \vec{y}\cdot\vec{y} \end{pmatrix}$$

has determinant equal to 0 in $\mathbb{Z}_p$. This determinant is a nonzero degree 4 multivariate polynomial in the coordinates of $\vec{x}, \vec{y}$, and since $\vec{x}, \vec{y}$ are chosen randomly over $\mathbb{Z}_p^j$, the probability that this determinant is 0 is at most $\frac{4}{p}$ by the Schwartz-Zippel Lemma, which is negligible. The same holds for $\vec{v}, \vec{z}$, since these are also uniformly randomly distributed (note that choosing $\vec{v}, \vec{z}$ uniformly randomly and then choosing $\vec{x}, \vec{y}$ randomly from $\mathcal{O}_{\vec{v},\vec{z}}$ yields the same distribution of $\vec{x}, \vec{y}, \vec{v}, \vec{z}$). Thus, with all but negligible probability, we have that $\mathcal{O}_{\vec{x},\vec{y}} \cap < \vec{x}, \vec{y} >= \{0\}$ and $\mathcal{O}_{\vec{v},\vec{z}} \cap < \vec{v}, \vec{z} >= \{0\}$. Since $\vec{v}, \vec{z}$ are orthogonal to $\vec{x}, \vec{y}$, any vector in $< \vec{x}, \vec{y}, \vec{v}, \vec{z} >$ which is also in $\mathcal{O}_{\vec{x},\vec{y},\vec{v},\vec{z}}$ can be written as a sum of a vector in $< \vec{x}, \vec{y} > \cap \mathcal{O}_{\vec{x},\vec{y}}$ and a vector in $< \vec{v}, \vec{z} > \cap \mathcal{O}_{\vec{v},\vec{z}}$, but such a vector must be 0.

Thus, we may assume that we can express any $\vec{w} \in \mathcal{O}_{\vec{v},\vec{z}}$ as $\vec{w} = a\vec{x} + b\vec{y} + \vec{u}$, where $\vec{u} \in \mathcal{O}_{\vec{x},\vec{y},\vec{v},\vec{z}}$. We note that the public key information-theoretically reveals a and b, but reveals nothing about $\vec{u}$. Conditioned on the public key, there are $p^{j-4} = |\mathcal{O}_{\vec{x},\vec{y},\vec{v},\vec{z}}|$ equally likely possibilities for $\vec{w}$. As a consequence, $\vec{u}$ is a random variable with min-entropy $(j-4)\log p$ in the adversary's view, before the adversary gets the leakage $f(\mathrm{SK})$.

The leakage function f has range $\{0,1\}^\ell$, where $\ell = (j-5)\log p$. Applying Lemma 3.7 (with $\tau = \frac{1}{4}\log p$ for concreteness), we have that with probability $\geq 1 - \frac{1}{p^{1/4}}$, SK conditioned on $f(\mathrm{SK})$ will have min-entropy at least $\frac{3}{4}\log p$. More precisely, the possible values of SK which are consistent with PK and $f(\mathrm{SK})$ form a set of size $\geq p^{\frac{3}{4}}$, and each value is equally likely.

We let W denote the set of possible values of SK in the adversary's view, after the adversary has seen PP, PK, and $f(\mathrm{SK})$. The distribution on this set is uniform, and we have shown that $|W| \geq p^{\frac{3}{4}}$ with all but negligible probability. Now, $C_1 = g^{\vec{s}}$, where $\vec{s}$ is chosen randomly from $\mathcal{O}_{\vec{v},\vec{z}}$. We let $Y_{\vec{s}}$ denote the random variable defined by $\vec{s}\cdot\vec{w}$, where $\vec{w}$ is chosen randomly from W. We will show that for any W of size $\geq p^{\frac{3}{4}}$, with all but negligible probability over the choice of $\vec{s}$, the random variable $Y_{\vec{s}}$ has min-entropy at least $\frac{1}{4}\log p$.

We first note that we can choose $\vec{s}$ randomly from $\mathbb{Z}_p^j$ instead of $\mathcal{O}_{\vec{v},\vec{z}}$ without changing the distribution. This is because we can write any $\vec{s} \in \mathbb{Z}_p^j$ as the sum of an element in the span of $\vec{v}, \vec{z}$

and an element in $\mathcal{O}_{\vec{v},\vec{z}}$, and the part in $< \vec{v}, \vec{z} >$ has no effect on $\vec{s} \cdot \vec{w}$ when $\vec{w} \in \mathcal{O}_{\vec{v},\vec{z}}$. (Here we have relied on our assumption that $\mathcal{O}_{\vec{v},\vec{z}} \cap < \vec{v}, \vec{z} >= \{0\}$.)

Now, every $\vec{w} \in W$ is of the form $\vec{w} = a\vec{x} + b\vec{y} + \vec{u}$, where a, b are fixed (needed for consistency with PK), and $\vec{u} \in \mathcal{O}_{\vec{v},\vec{z}}$ varies. As long as one of a, b is nonzero, we have that no two distinct vectors in W are multiples of each other. Since one of a, b will be nonzero with all but negligible probability, we may assume this is the case. We enumerate the elements of W as $\vec{w}_1, \ldots, \vec{w}_{|W|}$. We define random variables $X_1, \ldots, X_{|W|}$ by $X_i := \vec{s} \cdot \vec{w}_i$, where W is considered fixed and the randomness is over the choice of $\vec{s} \in \mathbb{Z}_p^j$. We note the relevant properties of these random variables:

- $\forall i = 1, \ldots, |W|,\ \forall c \in \mathbb{Z}_p,\ Pr_{\vec{s}}[X_i = c] = \frac{1}{p}$.
- $\forall i \neq h,\ \forall c_1, c_2 \in \mathbb{Z}_p,\ Pr_{\vec{s}}[X_i = c_1 \wedge X_h = c_2] = \frac{1}{p^2}$.

The first property follows from the fact that a linear equation over j variables in $\mathbb{Z}_p$ has p^{j-1} solutions in $\mathbb{Z}_p^j$, and hence the probability that a randomly chosen vector satisfies the linear equation is $\frac{p^{j-1}}{p^j} = \frac{1}{p}$. The second property follows from the fact that $\vec{w}_i$ and $\vec{w}_h$ are not multiples of each other: hence the two linear equations $\vec{s} \cdot \vec{w}_i = c_1$ and $\vec{s} \cdot \vec{w}_h = c_2$ are linearly independent, so the probability of a randomly chosen $\vec{s}$ satisfying them simultaneously is $\frac{1}{p^2}$. This is equivalent to saying that the random variables X_i are pairwise independent.

For each c in $\mathbb{Z}_p$, we define the random variable X_i^c to be 1 when $X_i = c$ and 0 otherwise. Then, $Pr_{\vec{s}}[X_i^c = 1] = \mathbb{E}[X_i^c] = \frac{1}{p}$, for all i and c. We also define the random variable

$$X^c := \sum_{i=1}^{|W|} X_i^c,$$

which counts the number of elements of W whose dot product with $\vec{s}$ is equal to c. By linearity of expectation, we have:

$$\mathbb{E}[X^c] = \sum_{i=1}^{|W|} X_i^c = \frac{|W|}{p}.$$

By pairwise independence of the X_i^c's, we also have:

$$Var[X^c] = \sum_{i=1}^{|W|} Var[X_i^c].$$

For each i, $Var[X_i^c] = \mathbb{E}[(X_i^c)^2)] - (\mathbb{E}[X_i^c])^2 = \frac{1}{p} - \frac{1}{p^2} = \frac{1}{p}\left(1 - \frac{1}{p}\right)$, since $(X_i^c)^2 = X_i^c$. Thus, $Var[X^c] = |W| \cdot \frac{1}{p}\left(1 - \frac{1}{p}\right)$. We note that $Pr[X^c \geq \frac{|W|}{p} \cdot p^{\frac{3}{4}}] = Pr[X^c - \mathbb{E}[X^c] \geq \frac{|W|}{p}\left(p^{\frac{3}{4}} - 1\right)]$. By Chebyshev's inequality, we then obtain:

$$Pr\left[X^c \geq \frac{|W|}{p} \cdot p^{\frac{3}{4}}\right] \leq \frac{p^2 Var[X^c]}{|W|^2 (p^{\frac{3}{4}} - 1)^2} = \frac{p\left(1 - \frac{1}{p}\right)}{|W|(p^{\frac{3}{4}} - 1)}.$$

By recalling our assumption that $|W| \geq p^{\frac{3}{4}}$ and employing the trivial bounds $(p^{\frac{3}{4}} - 1)^2 \geq \frac{1}{2} p^{\frac{3}{2}}$ and $1 - \frac{1}{p} < 1$, we can upper bound this quantity by $2p^{-1-\frac{1}{4}}$.

Since this holds for each value of $c \in \mathbb{Z}_p$, we can apply the union bound to obtain:

$$Pr_{\vec{s}}\left[\max_{c\in\mathbb{Z}_p} X^c \geq \frac{|W|}{p}\cdot p^{\frac{3}{4}}\right] \leq 2p^{-\frac{1}{4}}.$$

Therefore, will all but negligible probability over the choice of $\vec{s}$, we will obtain a random variable $Y_{\vec{s}}$ which takes on each value in $\mathbb{Z}_p$ with probability at most $p^{-\frac{1}{4}}$, and so has min-entropy at least $\frac{1}{4}\log p$ (recall the for a fixed $\vec{s}$ and W, $Y_{\vec{s}}$ is the random variable defined by choosing a vector $\vec{w} \in W$ uniformly at random and computing $\vec{s}\cdot\vec{w}$).

In summary, this means that with all but negligible probability, after $\mathrm{PP}, \mathrm{PK}, f(\mathrm{SK}), C_1$ have been revealed, the input to the extractor E will still have min-entropy at least $\frac{1}{4}\log p$, and so C_2 can be replaced by a uniformly random string, and the resulting statistical distance from the Game_3 distribution will be negligible. □

This completes our proof of Theorem 7.1

A.2 Attack on the Parallel System for $n > 1$ with Leakage $n\ell$

The same attack that applied to our composite order system applies here as well. For $n = 2$, the attacker receives the public parameters, PP, and two public keys, PK_1 and PK_2. If the two public keys are of the same type, the attacker aborts and guesses β' randomly (this occurs with probability $\frac{1}{2}$). Otherwise, we assume PK_1 is of type 1 and PK_2 is of type 2. The attacker chooses the leakage function

$$f(\mathrm{SK}_1, \mathrm{SK}_2) = \mathrm{SK}_1 \cdot \mathrm{SK}_2,$$

which multiples the two secret keys in G^j componentwise. This is permitted as long as $2\ell \geq jS(G)$ (since it will take $jS(G)$ bits to represent an element of G^j). Assuming this holds, the attacker receives the value $\mathrm{SK}_1 \cdot \mathrm{SK}_2$, and can use this to decrypt *both* ciphertexts. For example, to decrypt the ciphertext (C_1, C_2) for PK_1, the attacker computes:

$$E(e_j(C_1, \mathrm{SK}_1 \cdot \mathrm{SK}_2), D) \oplus C_2 = E(e_j(C_1, \mathrm{SK}_1), D) \oplus C_2,$$

since C_1 is orthogonal to SK_2 under e_j. This yields the first message share, and the second message share is obtained similarly, since the first ciphertext value for PK_2 will be orthogonal to SK_1. Hence, the attacker can reconstruct the encrypted message and succeed with probability 1 when the keys are of different types. This gives the attacker an overall advantage of $\frac{1}{4}$.

Whenever $S(G) \leq \frac{2(j-5)}{j}\log p$, we can choose an ℓ such that $\ell \leq (j-5)\log p$ and $2\ell \geq jS(G)$ simultaneously hold. We note that we can choose j freely, so as long as $S(G) \leq (2-\epsilon)\log p$ for some $\epsilon > 0$, we can choose j sufficiently large and obtain $S(G) \leq \frac{2(j-5)}{j}\log p$.

For general values of n, the attacker can expect that close to half of the n keys will be of type 1 and the others will be of type 2. More specifically, a Chernoff bound implies that for any positive constant c, there will be at least $\frac{n}{2} - c\sqrt{n}$ keys of each type with probability $\geq 1 - 2e^{-2c^2}$.

When this distribution of keys occurs, the attacker can organize the keys into at least $\frac{n}{2} - c\sqrt{n}$ pairs of keys and at most $2c\sqrt{n}$ individual keys, where each pair of keys contains one key of type 1 and one key of type 2. The attacker then defines the leakage function so that it reveals the (componentwise) product of each pair of secret keys and all of the remaining unpaired secret keys. This requires $(\frac{n}{2} - c\sqrt{n})j \cdot S(G) + 2c\sqrt{n} \cdot j \cdot S(G) = j \cdot S(G)(\frac{n}{2} + c\sqrt{n})$ bits of leakage. We note that the attacker can now decrypt messages encrypted under *any* of the n public keys.

As long as

$$n\ell \geq j \cdot S(G)\left(\frac{n}{2} + c\sqrt{n}\right) \Leftrightarrow \ell \geq j \cdot S(G)\left(\frac{1}{2} + \frac{c}{\sqrt{n}}\right)$$

holds, this is a valid attack with $n\ell$ bits of leakage that succeeds with probability $\geq 1 - 2e^{-2c^2}$.

For $\ell \leq (j-5)\log p$ and the attack condition $\ell \geq j \cdot S(G)(\frac{1}{2} + \frac{c}{\sqrt{n}})$ to hold simultaneously for some value of ℓ, it suffices to have a group G such that:

$$S(G) \leq \frac{2(j-5)}{j} \cdot \frac{1}{1 + \frac{2c}{\sqrt{n}}} \log p.$$