


Everybody’s a Target: Scalability in Public-Key Encryption

Benedikt Auerbach¹ , Federico Giacon², and Eike Kiltz³ 

¹ IST Austria, Klosterneuburg, Austria
`benedikt.auerbach@ist.ac.at`

² Gnosis Service GmbH, Berlin, Germany
`federico.giacon@rub.de`

³ Ruhr-Universität Bochum, Bochum, Germany
`eike.kiltz@rub.de`

Abstract. For $1 \leq m \leq n$, we consider a natural m -out-of- n multi-instance scenario for a public-key encryption (PKE) scheme. An adversary, given n independent instances of PKE, wins if he breaks at least m out of the n instances. In this work, we are interested in the *scaling factor* of PKE schemes, SF, which measures how well the difficulty of breaking m out of the n instances scales in m . That is, a scaling factor $SF = \ell$ indicates that breaking m out of n instances is at least ℓ times more difficult than breaking one single instance. A PKE scheme with small scaling factor hence provides an ideal target for mass surveillance. In fact, the Logjam attack (CCS 2015) implicitly exploited, among other things, an almost constant scaling factor of ElGamal over finite fields (with shared group parameters).

For Hashed ElGamal over elliptic curves, we use the generic group model to describe how the scaling factor depends on the scheme’s granularity. In low granularity, meaning each public key contains its independent group parameter, the scheme has optimal scaling factor $SF = m$; In medium and high granularity, meaning all public keys share the same group parameter, the scheme still has a reasonable scaling factor $SF = \sqrt{m}$. Our findings underline that instantiating ElGamal over elliptic curves should be preferred to finite fields in a multi-instance scenario.

As our main technical contribution, we derive new generic-group lower bounds of $\Omega(\sqrt{mp})$ on the complexity of both the m -out-of- n Gap Discrete Logarithm and the m -out-of- n Gap Computational Diffie-Hellman problem over groups of prime order p , extending a recent result by Yun (EUROCRYPT 2015). We establish the lower bound by studying the hardness of a related computational problem which we call the search-by-hypersurface problem.

1 Introduction

For integers $1 \leq m \leq n$, consider the following natural m -out-of- n multi-instance attack scenario for a public-key encryption scheme PKE⁴. An attacker is given

⁴ Formally, in this work we consider key-encapsulation mechanisms.

n independent instances (public keys) of PKE and would like to *simultaneously break semantic security at least m out of n instances*. Note that this is a different setting from the standard, well studied, multi-user attack scenario by Bellare et al. [7]. In the (security-wise) best possible scenario, running an m -out-of- n multi-instance attack is m times more difficult compared to a (standard) single-instance attack. However, there is no guarantee that breaking m -out-of- n instances is more difficult than breaking a single instance.

This motivates the following question:

How well does the difficulty of breaking m out of n instances of PKE scale with m ?

In order to give a quantitative answer to this question, we define the scaling factor (relative to a fixed security notion) of PKE as

$$\text{SF}_{\text{PKE}}^{m,n} = \frac{\text{resources necessary to break } m \text{ out of } n \text{ instances}}{\text{resources necessary to break 1 instance}}, \quad (1)$$

where “resources” refers to the running time to break PKE in the studied security notion. Clearly, the larger SF_{PKE} , the better are the security guarantees in the multi-instance setting. The best we can hope for is $\text{SF}_{\text{PKE}}^{m,n} = m$, meaning that breaking m out of n instances amounts to breaking m times a single instance of PKE.

SCALING FACTOR AND MASS SURVEILLANCE. In 2012, James Bamford wrote in Wired:

According to another top official also involved with the program, the NSA made an enormous breakthrough several years ago in its ability to cryptanalyze, or break, unfathomably complex encryption systems employed by not only governments around the world but also many average computer users in the US. The upshot, according to this official: **“Everybody’s a target; everybody with communication is a target.”**

This statement should appear as a surprise to the cryptographic community: Parameters for cryptographic schemes are usually chosen to make even compromising a single user a daunting challenge, meaning multi-instance attacks seem out of scope even for adversaries with nation-state capabilities. Unfortunately, the use of outdated parameters is a widespread occurrence in practice [2,19], either as a consequence of legacy infrastructure or hardware restrictions. In this case, a bad scaling factor would tip the scale from single compromised users to full-scale mass surveillance. Even more so, the hardness of several common number-theoretic problems is known to scale sub-optimally in the number of instances. Examples are factoring [11] and computing discrete logarithms in the finite-field [4,5] and elliptic-curve [18,20,22] setting. This sub-optimal scaling is typically inherited by the corresponding cryptographic schemes. It has been exploited in practice by the famous Logjam attack [2], where the authors break many Diffie-Hellman instances in TLS with nearly the same resources as to break

PKE	Setting	Shared param.	Public key pk_i
HEG[GGen $_{\mathbb{E}(\mathbb{F}_\ell)}$, high]	Elliptic curve	$\mathbb{E}(\mathbb{F}_\ell), p, g$	g^{x_i}
HEG[GGen $_{\mathbb{E}(\mathbb{F}_\ell)}$, med]	Elliptic curve	$\mathbb{E}(\mathbb{F}_\ell), p$	$g_i, g_i^{x_i}$
HEG[GGen $_{\mathbb{E}(\mathbb{F}_\ell)}$, low]	Elliptic curve	–	$\mathbb{E}_i(\mathbb{F}_{\ell_i}), p_i, g_i, g_i^{x_i}$
HEG[GGen $_{\mathbb{F}_\ell^*}$, high]	Finite field	\mathbb{F}_ℓ^*, p, g	g^{x_i}
HEG[GGen $_{\mathbb{F}_\ell^*}$, med]	Finite field	\mathbb{F}_ℓ^*, p	$g_i, g_i^{x_i}$
HEG[GGen $_{\mathbb{F}_\ell^*}$, low]	Finite field	–	$\mathbb{F}_{\ell_i}, p_i, g_i, g_i^{x_i}$

Table 1. Shared public system parameters and individual public keys for schemes HEG[GGen $_{\mathbb{E}(\mathbb{F}_\ell)}$, gran] and HEG[GGen $_{\mathbb{F}_\ell^*}$, gran] at different granularities. Here g generates a subgroup of prime order p of either an elliptic curve $\mathbb{E}(\mathbb{F}_\ell)$ or a finite field \mathbb{F}_ℓ^* and ℓ is a prime.

a single Diffie-Hellman instance. Concretely, the Logjam attack could successfully break multiple 512-bit finite-field instances, and the authors also speculate about the feasibility of breaking 1024-bit instances. With our work we aim to deliver positive results by computing (non-trivial lower bounds on) the scaling factors of concrete encryption schemes that are currently employed in practice, thereby providing bounds on the hardness of performing mass surveillance.

CONSIDERED ENCRYPTION SCHEMES. We are able to provide non-trivial bounds on the scaling factor for Hashed ElGamal (HEG), also known as DHIES [1], in the elliptic curve (HEG[GGen $_{\mathbb{E}(\mathbb{F}_\ell)}$]) and the finite field (HEG[GGen $_{\mathbb{F}_\ell^*}$]) setting, the arguably most widely used discrete-logarithm-type encryption schemes. Here GGen $_{\mathbb{E}(\mathbb{F}_\ell)}$ and GGen $_{\mathbb{F}_\ell^*}$ are group-generating algorithms that generate prime-order subgroups of elliptic curves and finite fields respectively. In both cases, ℓ denotes randomly chosen primes of appropriate size. We consider both schemes instantiated in three different granularity settings (low, medium, and high), leading to six schemes, HEG[GGen $_{\mathbb{E}(\mathbb{F}_\ell)}$, low], HEG[GGen $_{\mathbb{E}(\mathbb{F}_\ell)}$, med], HEG[GGen $_{\mathbb{E}(\mathbb{F}_\ell)}$, high], HEG[GGen $_{\mathbb{F}_\ell^*}$, low], HEG[GGen $_{\mathbb{F}_\ell^*}$, med], and HEG[GGen $_{\mathbb{F}_\ell^*}$, high], which offer different trade-offs between public key sizes and scalability. The term *granularity* specifies which parts of the scheme’s parameters belong to the global system parameters (shared among all n users), and which parts belong to the individual, user-specific public keys. Table 1 depicts the shared public system parameters and individual keys in a multi-instance setting with n parties for HEG at different granularities.

1.1 Our Results

FORMAL DEFINITIONS: MULTI-INSTANCE SECURITY. The notion of n -out-of- n multi-instance security for any $n \geq 1$ was first considered and formally defined by Bellare et al. [8] in the setting of secret-key encryption. As our first contribution, we extend their notion to m -out-of- n multi-instance security for public-key

encryption, for arbitrary $1 \leq m \leq n$. In fact, we give two different notions, modeling (m, n) -CPA (passive) and (m, n) -CCA (active) security.

Our (m, n) -CPA experiment provides the adversary with n independent public keys $pk[1], \dots, pk[n]$. Next, it picks n independent challenge bits $b[1], \dots, b[n]$ and grants the adversary access to oracle $\text{Enc}(\cdot, \cdot, \cdot)$ which, given i, M_0, M_1 , returns an encryption of message $M_{b[i]}$ under $pk[i]$. The adversary outputs a single bit b' together with a list $L \subseteq \{1, \dots, n\}$ of cardinality at least m . The advantage function is defined as

$$\text{Adv}_{\text{PKE}}^{(m,n)\text{-cpa}} = \Pr \left[b' = \bigoplus_{i \in L} b[i] \right] - \frac{1}{2} .$$

That is, the adversary wins if it guesses correctly the XOR of at least m (out of n) challenge bits. (Note that the standard multi-user security notion for PKE [7] is different: Most importantly, multi-user security involves only a single challenge bit, in particular limiting this notion to the case of $m = 1$.) Why using XOR for defining the winning condition? Bellare et al. [8] argue that this is a natural metric because its well-known ‘‘sensitivity’’ means that as long as at least one of the challenge bits looks random to the adversary so does their XOR. They further argue that other possible winning conditions such as using AND⁵ are less natural and lead to inconsistencies. We refer to Bellare et al. [8] for an extensive discussion. In (m, n) -CCA security, the adversary is furthermore provided with a decryption oracle $\text{Dec}(\cdot, \cdot)$ which given i, c returns a decryption of c under $sk[i]$. To expand on the characteristics of the multi-instance setting, we determine the relations between the security notions (m, n) -CPA and (m, n) -CCA for different values of m and n . The natural results we are able to show in this regard (among others, the intuitive idea that a single-instance adversary of advantage ϵ and running time t can be extended to an m -out-of- n adversary of advantage ϵ^m and running time mt ; see Theorem 1) give us further confidence on the significance of the chosen multi-instance security definition, and enable us to present a formally sound definition of the scaling factor.

SCALING FACTOR OF $\text{HEG}[\text{GGen}_{\mathbb{E}(\mathbb{F}_\ell)}, \cdot]$ AND $\text{HEG}[\text{GGen}_{\mathbb{F}_\ell^*}, \cdot]$. In order to give a lower bound on $\text{SF}_{\text{PKE}}^{m,n}$ as defined in Eq. (1), we need to lower bound the numerator (i.e., resources required to break m out of n instances) for all possible adversaries and upper bound the denominator (i.e., resources needed to break one instance) by specifying a concrete adversary. Unfortunately, unless the famous P vs. NP problem is settled, all meaningful lower bounds on the resources will require either an unproven complexity assumption or a restricted model of computation. We rely on the generic group model [28] for $\text{HEG}[\text{GGen}_{\mathbb{E}(\mathbb{F}_\ell)}, \cdot]$ (which is considered to be meaningful for elliptic-curve groups) and on a hypothesis on the running time of variants of the number field sieve for $\text{HEG}[\text{GGen}_{\mathbb{F}_\ell^*}, \cdot]$ based on the fastest known attacks on finite fields.

Our main results regarding the scaling factor $\text{SF}_{\text{HEG}}^{m,n}$ in different granularities relative to (m, n) -CCA security are summarized in Table 2. In both considered

⁵ I.e., by letting the adversary output a vector $b'[1], \dots, b'[n]$ and a set I and defining the advantage function as $\text{Adv}_{\text{PKE}}^{(m,n)\text{-cpa}} = \Pr[\bigwedge_{i \in I} b[i] = b'[i]] - 1/2^m$.

PKE	Setting	Scaling factor
HEG[GGen $_{\mathbb{E}(\mathbb{F}_\ell)}$, {high, med}]	Elliptic curve	$\Theta(\sqrt{m})$
HEG[GGen $_{\mathbb{E}(\mathbb{F}_\ell)}$, low]	Elliptic curve	$\Theta(m)$
HEG[GGen $_{\mathbb{F}_\ell^*}$, {high, med}]	Finite field	$\begin{cases} 1 & \delta \leq 0.67 \\ L_\ell(1/3, \delta - 0.67) & \delta > 0.67 \end{cases}$
HEG[GGen $_{\mathbb{F}_\ell^*}$, low]	Finite field	$\begin{cases} L_\ell(1/3, \delta) & 0 \leq \delta < 0.105 \\ L_\ell(1/3, 0.105) & 0.105 \leq \delta < 0.368 \\ L_\ell(1/3, -0.263 + \delta) & 0.368 \leq \delta \end{cases}$

Table 2. Lower bounds on the scaling factor $\text{SF}_{\text{HEG}}^{m,n}$ relative to (m, n) -CCA security. $L_\ell(1/3, c)$ is defined as $\exp((c + o(1))(\log \ell)^{1/3}(\log \log \ell)^{2/3})$. In the finite field case $m = L_\ell(1/3, \delta)$ for some $\delta \geq 0$.

group instantiations, HEG shows the same asymptotic scaling behavior for high and medium granularity. In both cases however, HEG scales better in the low-granularity case. Concretely, Hashed ElGamal over elliptic curves (modeled as generic groups) scales optimally for low-granularity parameters. For medium and high granularity, on the other hand, the scaling factor is of order $\Theta(\sqrt{m})$, where the constants hidden by the Θ -notation are small.

Let $L_\ell(1/3, c) := \exp((c + o(1))(\log \ell)^{1/3}(\log \log \ell)^{2/3})$. For HEG in the finite field setting with respect to high and medium granularity, we see that the scaling factor is roughly 1 for up to $m = L_\ell(1/3, 0.67)$ instances, the point starting from which the cumulative cost of breaking m individual instances outweighs the cost of the precomputation. Beyond, the KEM scales linearly with slope $L_\ell(1/3, -0.67)$. Note that $L_\ell(1/3, 0.67)$ is large for typical values of ℓ . Concretely, for 512 bit primes we get that $L_\ell(1/3, 0.67) \approx 2^{22}$ meaning that the effort of breaking 2^{22} instances roughly equals the effort to break a single instance. While the concrete number is obtained ignoring the $o(1)$ terms in L_ℓ , it still matches empirical results [2, Table 2]. For low granularity and for up to $L_\ell(1/3, 0.105)$ instances, HEG[GGen $_{\mathbb{F}_\ell^*}$, low] scales optimally. For $L_\ell(1/3, 0.105) \leq m \leq L_\ell(1/3, 0.368)$, the scaling factor is roughly constant, and for larger numbers of instances, it scales linearly with slope $L_\ell(1/3, -0.263)$ which is far larger than the slope in the case of medium or high granularity.

Summing up, Hashed ElGamal instantiated with elliptic curve groups shows a better scaling behavior than the corresponding instantiation in the finite-field setting. Further, in both cases switching from the high granularity setting to the medium granularity setting does not improve the scaling behavior, while the use of individual groups, i.e., low-granularity parameters does. To illustrate our findings we provide example values of the scaling factor for different numbers of instances m and prime sizes ℓ in Table 3.

While our results imply that the use of low-granularity parameters is preferable with respect to security scaling, we stress that generating cryptographically secure groups is a hard and error prone process. Delegating this task to the individual user as part of the key generation might actually have a negative im-

m	ℓ	Elliptic curve		Finite field	
		high, med	low	high, med	low
2^{20}	512	2^{10}	2^{20}	1.21	$2^{11.26}$
	1024	2^{10}	2^{20}	1.00	$2^{8.26}$
	2048	2^{10}	2^{20}	1.00	$2^{6.64}$
2^{30}	512	2^{15}	2^{30}	$2^{7.73}$	$2^{21.26}$
	1024	2^{15}	2^{30}	1.85	$2^{18.13}$
	2048	2^{15}	2^{30}	1.00	$2^{14.02}$

Table 3. Example values of scaling factor $SF_{\text{HEG}[\text{GGen}, \text{gran}]}^{(m,m)\text{-cca}}$ for different values of m and ℓ , $\text{GGen} \in \{\text{GGen}_{\mathbb{E}(\mathbb{F}_\ell)}, \text{GGen}_{\mathbb{F}_\ell^*}\}$, and $\text{gran} \in \{\text{high}, \text{med}, \text{low}\}$.

fact on the scheme’s security in practice. Further, the use of individual groups negatively impacts the efficiency of the scheme, as key generation requires the sampling of secure groups, and key-sizes increase.

DERIVATION OF THE SCALING FACTORS. As we will explain below in more detail, the bounds from Table 2 are obtained in two steps. In a **first step**, we consider an m -out-of- n multi-instance version of the Gap Computational Diffie-Hellman problem, (m, n) -GapCDH[GGen, gran], where the term “gap” refers to the presence of a Decisional Diffie-Hellman (DDH) oracle. The following theorem holds for all $\text{GGen} \in \{\text{GGen}_{\mathbb{E}(\mathbb{F}_\ell)}, \text{GGen}_{\mathbb{F}_\ell^*}\}$ and $\text{gran} \in \{\text{high}, \text{med}, \text{low}\}$.

Theorem *The (m, n) -CCA security of $\text{HEG}[\text{GGen}, \text{gran}]$ is tightly implied by the hardness of (m, n) -GapCDH[GGen, gran].*

The theorem (described formally in Section 4) is a somewhat straightforward generalization of the single-instance case [1]. We stress that tightness in our previous theorem is an essential ingredient to obtain overall tight bounds on the scaling factor.

In a **second step**, we provide bounds on the (m, n) -GapCDH[GGen, gran] problem. In the finite field case, we rely on the following hypothesis:

Hypothesis 1 *The fastest algorithms to break (m, n) -GapCDH[GGen $_{\mathbb{F}_\ell^*}$, gran] are variants of the number field sieve [4, 5] which require running time*

$$T = \begin{cases} L_\ell(1/3, 1.902) + m \cdot L_\ell(1/3, 1.232) & \text{gran} \in \{\text{high}, \text{med}\} \\ \min\{m \cdot L_\ell(1/3, 1.902), L_\ell(1/3, 2.007) + m \cdot L_\ell(1/3, 1.639)\} & \text{gran} = \text{low} \end{cases} .$$

The lower bounds on $SF^{m,n}$ for $\text{HEG}[\text{GGen}_{\mathbb{F}_\ell^*}, \text{gran}]$ are obtained by combining the previous theorem and Hypothesis 1. The running times specified in the hypothesis stem from the multi-field NFS [5] (high/medium granularity) and the DLOG factory [4] (low granularity). Both variants first require an instance-independent precomputation. Then instances can be solved with a constant computational effort. The values $\delta = 0.67$ and $\delta = 0.368$ of Table 2 correspond to the number of instances starting from which the cumulative cost of breaking the instances outweighs the cost of the precomputation.

m -out-of- n problem	Given	Break m out of	Gap?	Complexity	Ref.
DL[GGen, high]	$\mathbb{G}, p, g, g^{x^1}, \dots, g^{x^n}$	x_1, \dots, x_n	–	$\Theta(\sqrt{mp})$	[30,29,22]
DL[GGen, med]	$\mathbb{G}, p, g_1, g_1^{x^1}, \dots, g_n, g_n^{x^n}$	x_1, \dots, x_n	–	$\Theta(\sqrt{mp})$	full version [3]
DL[GGen, low]	$\mathbb{G}_1, p_1, g_1, g_1^{x^1}, \dots, \mathbb{G}_n, p_n, g_n, g_n^{x^n}$	x_1, \dots, x_n	–	$\Theta(m\sqrt{p})$	full version [3]
GapDL[GGen, high]	$\mathbb{G}, p, g, g^{x^1}, \dots, g^{x^n}$	x_1, \dots, x_n	✓	$\Theta(\sqrt{mp})$	§5.2
GapDL[GGen, med]	$\mathbb{G}, p, g_1, g_1^{x^1}, \dots, g_n, g_n^{x^n}$	x_1, \dots, x_n	✓	$\Theta(\sqrt{mp})$	full version [3]
GapDL[GGen, low]	$\mathbb{G}_1, p_1, g_1, g_1^{x^1}, \dots, \mathbb{G}_n, p_n, g_n, g_n^{x^n}$	x_1, \dots, x_n	✓	$\Theta(m\sqrt{p})$	full version [3]
GapCDH[GGen, high]	$\mathbb{G}, p, g, g^{x^1}, g^{y^1}, \dots, g^{x^n}, g^{y^n}$	$g^{x^1 y^1}, \dots, g^{x^n y^n}$	✓	$\Theta(\sqrt{mp})$	§6.1
GapCDH[GGen, med]	$\mathbb{G}, p, g_1, g_1^{x^1}, g_1^{y^1}, \dots, g_n, g_n^{x^n}, g_n^{y^n}$	$g_1^{x^1 y^1}, \dots, g_n^{x^n y^n}$	✓	$\Theta(\sqrt{mp})$	§6.2
GapCDH[GGen, low]	$\mathbb{G}_1, p_1, g_1, g_1^{x^1}, g_1^{y^1}, \dots, \mathbb{G}_n, p_n, g_n, g_n^{x^n}, g_n^{y^n}$	$g_1^{x^1 y^1}, \dots, g_n^{x^n y^n}$	✓	$\Theta(m\sqrt{p})$	§6.3

Table 4. Definition and generic-group complexity of problems (m, n) -DL[GGen, gran], (m, n) -GapDL[GGen, gran], and (m, n) -GapCDH[GGen, gran], where gran belongs to {high, med, low}. \mathbb{G} and \mathbb{G}_i are generic groups of prime order p and p_i , with generators g and g_i , respectively. The third column defines the problem’s winning condition. The Gap column indicates the presence of a DDH oracle.

In the elliptic-curve case, we make the hypothesis that the fastest adversary attacking the system is a generic-group adversary. Concretely, we prove the following generic-group lower bounds for (m, n) -GapCDH[GGen_{gg}, gran] in different granularities, where GGen_{gg} generates a generic group [28] of prime order p , and the granularity gran determines how much information about the used group is shared amongst the challenge instances (see Table 4).

Theorem *The best generic algorithm to break (m, n) -GapCDH[GGen_{gg}, gran] requires running time*

$$T = \begin{cases} \Theta(\sqrt{mp}) & \text{gran} \in \{\text{high}, \text{med}\} \\ \Theta(m\sqrt{p}) & \text{gran} = \text{low} \end{cases},$$

and the constants hidden by the Θ notation are small (between 0.1 and 6.6).

The lower bounds on SF ^{m, n} for HEG[GGen _{$\mathbb{E}(\mathbb{F}_\ell)$} , gran] are obtained by combining our previous theorems and assuming that elliptic-curve groups behave like generic groups.

1.2 Generic Bounds on Multi-Instance GapCDH: Technical Details

We consider multi-instance variants of three different problems: the discrete logarithm problem ((m, n) -DL[GGen_{gg}, gran]), the gap discrete logarithm problem ((m, n) -GapDL[GGen_{gg}, gran]), and the gap computational Diffie-Hellman problem ((m, n) -GapCDH[GGen_{gg}, gran]) in different granularities, see Table 4.

We now discuss the complexity column of Table 4. It is well known that the running time of solving (m, n) -DL[GGen_{gg}, high] is $\Theta(\sqrt{mp})$, the lower bound being in the generic group model [30,29], the matching upper bound stemming from a concrete generic algorithm [22]. It is not hard to see that the bounds on (m, n) -DL[GGen_{gg}, med] are basically the same because the generators g_i can

be viewed as “high-granularity instances” g^{x_j} . Concerning low granularity, it is noteworthy to mention the bound for the case $m = n$ by Garay et al. [17]. Using different techniques, we are able to improve their bound from \sqrt{mp} to $m\sqrt{p}$. In addition, our bound also holds in the case $m < n$ and in the gap setting.

Our **first main technical result** (Corollary 1) is a non-trivial extension of Yun’s generic lower bound [30] to the gap setting, i.e., a new lower bound of $\Omega(\sqrt{mp})$ on solving (m, m) -GapDL[GGen_{gg}, high]. Based on this result, we also deduce bounds in the case of medium and low granularity.

Our **second main technical result** (Theorem 4) states that, in high granularity, the (m, m) -GapDL and the (m, n) -GapCDH problems are essentially equally hard in the algebraic group model [16], hence implying the required bounds in the generic group model. The results in medium and low granularity follow as in the discrete logarithm setting.

MAIN TECHNICAL RESULT 1: LOWER BOUND ON (m, m) -GapDL[GGen_{gg}, high]. We define a new “hard” problem called the *polycheck discrete logarithm problem*: The security game is the same as that of standard multi-instance DL, but the adversary has additional access to an oracle Eval that behaves as follows: Given as input to Eval a polynomial $f \in \mathbb{Z}_p[X_1, \dots, X_k]$ and group elements g^{x_1}, \dots, g^{x_k} , it returns 1 if and only if $g^{f(x_1, \dots, x_k)} = 1$. This problem is easier than GapDL: In fact, we can simulate the gap oracle DDH(g^x, g^y, g^z) by querying Eval($f := X_1 X_2 - X_3, g^x, g^y, g^z$). In the generic group model, we can bound the advantage of an adversary against the m -out-of- m polycheck discrete logarithm problem that queries polynomial of degree at most d ((m, m) - d -PolyDL[GGen_{gg}, high]) as

$$\text{Adv}^{(m, m)\text{-}d\text{-polydl}} \lesssim \left(\frac{dq^2 + dq_{\text{Eval}}}{mp} \right)^m,$$

where q bounds the queries to the group-operation oracle, q_{Eval} to Eval, and p is the order of the generic group. The bound for high-granularity GapDL follows by setting $d = 2$.

The result is proven by extending the arguments by Yun [30] for the standard multi-instance DL problem. In line with Yun’s approach, we define the *search-by-hypersurface* problem in dimension m (m -SHS _{d} [p]), which requires to find a uniformly sampled point $\mathbf{a} \in \mathbb{Z}_p^m$ while being able to check whether \mathbf{a} is a zero of adaptively chosen polynomials in $\mathbb{Z}_p[X_1, \dots, X_m]$ of degree at most d . Notably, Yun’s *search-by-hyperplane-queries* problem in dimension m is equivalent to m -SHS₁. We stress that the more general case of $d \geq 1$ requires significantly different arguments from commutative algebra/algebraic geometry, compared to the linear algebra argument used for the DL bound.

We show that any generic adversary against (m, m) - d -PolyDL[GGen_{gg}, high] can be transformed into an adversary against m -SHS _{d} , and then proceed to bound the advantage of an adversary against m -SHS _{d} . The key step is observing that an adversary can make at most m useful hypersurface queries, that is, queries that return 1 (hence, identify a hypersurface on which the point \mathbf{a} lies) and whose output is not easy to determine based on previous queries. The key

difference between our result and Yun’s lies in how useful queries are processed and counted. Since Yun considers only polynomials of degree 1, a hypersurface defined by a polynomial of degree 1 is a hyperplane of the affine space \mathbb{Z}_p^m . Each useful query identifies another hyperplane on which the sought point lies. When intersecting another hyperplane with the intersection of the hyperplanes previously found, the dimension of the intersection as an affine subspace is brought down by one. The dimension of the full affine space being m , at most m such queries can be made before identifying a single point (dimension 0). However, generalizing to hypersurfaces generated by polynomials of degree ≥ 2 requires to carry over more sophisticated arguments from commutative algebra. Firstly, intersecting m hypersurfaces does not, in general, identify a single point. Secondly, intersection of two hypersurfaces might give rise to the union of two or more irreducible components. Intersecting further with a hypersurface containing just one of those irreducible components would qualify as a useful query, however would not bring down the dimension of the intersection by one. This impasse is overcome by guessing the correct component at each step. Fortunately, Bézout’s theorem and a discerning choice of the guessing probabilities at each useful query makes the argument go through with just an additional loss of d^m , which is absorbed by the exponential bound in the dimension.

MAIN TECHNICAL RESULT 2: (m, m) -GapDL[GGen, high] HARDNESS IMPLIES (m, n) -GapCDH[GGen, high]. The algebraic group model [16] is a technique used to extend existing bounds in the generic group model to different problems by means of generic reductions. Our second technical result (Theorem 4) presents a generic reduction between the problems (m, n) -GapCDH[GGen, high] and (m, m) -GapDL[GGen, high] with a tightness loss of 2^m in the algebraic group model. Combining this with the generic-group lower bound we prove as our first main technical result, we obtain, in the generic group model:

$$\text{Adv}_{\text{high}}^{(m, n)\text{-gcdh}} \stackrel{\text{Th. 4}}{\leq} 2^m \cdot \text{Adv}_{\text{high}}^{(m, m)\text{-gdl}} \stackrel{\text{Cor. 1}}{\lesssim} 2^m \left(\frac{q^2 + q_{\text{DDH}}}{mp} \right)^m \approx \left(\frac{2q^2}{mp} \right)^m,$$

where q bounds the queries to the group-operation oracle, q_{DDH} to the gap oracle, and p is the order of the generic group. Note that the reduction’s exponential loss of 2^m gets swallowed by the (m, m) -GapDL[GGen_{gg}, high] bound. More importantly, by the above bound one requires $q \geq \Omega(\sqrt{mp})$ generic-group operations to break (m, n) -GapCDH[GGen_{gg}, high] with overwhelming advantage.

A natural approach to tackle the proof of Theorem 4 would be to adapt the single-instance proof presented by Fuchsbauer et al. [16] to the multi-instance setting. Following this strategy in a reduction, however, one would need to argue about the size of the solution set of a multivariate system of quadratic equations. In this work we employ significantly different proof techniques.

The path we pursue maintains, instead, the linear character of the system. The reduction distributes the i -th DL challenges in either the X or Y components of the i -th challenges to the CDH adversary. The intuition at the core of the proof is that an adversary finding the CDH solution for any one instance must provide the DL of at least one of the two corresponding challenge components

(even if possibly depending on the remaining, unrecovered DLs). If the reduction manages to embed the m DL challenges at the right spot, then it is able to recover all logarithms. The reduction loss of 2^m is consequence of this guess. Moreover, expanding the m DL challenges into n CDH challenges adds a further layer of complexity.

1.3 Related Work and Future Directions

RELATED WORK. Multi-instance security in the sense of breaking m out of m instances was first formally considered in the setting of symmetric encryption by Bellare et al. [8]. We point out that the term is sometimes also used to describe multi-user, multi-challenge generalizations of single-instance security notions [21].

The (single-instance) GapCDH problem was introduced by Okamoto and Pointcheval [25]. Boneh et al. [12] and Rupp et al. [26] provide frameworks in the generic-group model that can be used to derive generic-group lower bounds on the hardness of many single-instance problems, gapCDH amongst others. The generic hardness of (m, m) -DL in the high-granularity setting was first analyzed by Yun [30], the result later generalized to (m, n) -DL by Ying and Kunihiro [29]. Kuhn and Struik [22], and Fouque et al. [15] give generic algorithms matching the lower bounds. The first bound for (m, m) -DL in the low granularity setting was derived by Garay et al. [17]. The algebraic-group model was introduced by Fuchsbauer et al. [16]. Mizuide et al. [24] provide a framework that can be used to reduce single-instance CDH-type problems to the discrete-logarithm problem in the algebraic-group model.

Bartusek et al. [6] and Sadeghi et al. [27] discuss differences between DL-type assumptions depending on whether the used group and group generator are fixed or sampled at random. We stress that in this work groups and group generators, while potentially shared amongst different users, are sampled at the beginning of the game and hence part of its probability space.

FUTURE DIRECTIONS. Corrigan-Gibbs and Kogan [14] consider the multi-instance discrete logarithm problem in a setting where the adversary is allowed to first perform unbounded preprocessing over the group to produce an advice string of bounded size, which in a second stage is used to solve multiple discrete logarithm instances. The resulting lower bounds in the generic group model were also derived by Coretti et al. [13] using a different technique. It would be interesting to compute scaling factors of the considered schemes taking preprocessing into account. Another possible direction is to derive lower bounds on the scaling factor for practical encryption schemes in the RSA setting (e.g., RSA-OAEP [9]) and in the post-quantum setting (e.g., based on lattices and codes).

2 Preliminaries

2.1 Notation

VECTOR NOTATION. We denote vectors with boldface fonts, for example \mathbf{v} . The number of elements of a vector is represented by $|\mathbf{v}|$. Element indexing starts from 1, and the entry at position i is accessed through square brackets: $\mathbf{v}[i]$. To initialize all entries of a vector to some element a we write $\mathbf{v}[\cdot] \leftarrow a$. We may initialize multiple vectors simultaneously, and moreover initialize them through running some (possibly randomized) routine. As an example, we could initialize a vector of public and of secret keys as $(\mathbf{pk}, \mathbf{sk})[\cdot] \leftarrow_{\S} \text{Gen}$ to indicate that for every index i we run Gen with fresh randomness and, denoting the output with (pk, sk) , set $\mathbf{pk}[i] \leftarrow pk$ and $\mathbf{sk}[i] \leftarrow sk$. Given any set of indices I , we denote with $\mathbf{v}[I]$ the vector that contains only the entries indexed with elements in I . For example, if $\mathbf{v} = (a, b, c)$ then $\mathbf{v}[\{1, 3\}] = (a, c)$. We slightly abuse this notation, writing $\mathbf{v}[I] \leftarrow \mathbf{w}$ when replacing each entry of \mathbf{v} whose indices belong to I by the elements of \mathbf{w} in their order. For example, if $\mathbf{v} = (a, b, c)$ and we execute $\mathbf{v}[\{1, 3\}] \leftarrow (d, e)$ then $\mathbf{v} = (d, b, e)$.

GROUP NOTATION. In this paper we consider groups \mathbb{G} of prime order p , generated by g . We call $\mathcal{G} = (\mathbb{G}, p, g)$ a group representation. A group-generating algorithm GGen is a randomized algorithm that outputs a group representation \mathcal{G} . We assume that all groups output by GGen are of the same bit length.

In this work we consider two instantiations $\text{GGen}_{\mathbb{E}(\mathbb{F}_\ell)}$ and $\text{GGen}_{\mathbb{F}_\ell^*}$ of group-generating algorithms. In both cases ℓ denotes a randomly sampled prime of appropriate size. Group descriptions \mathcal{G} output by $\text{GGen}_{\mathbb{E}(\mathbb{F}_\ell)}$ are prime-order p subgroups of elliptic curves defined over the field \mathbb{F}_ℓ . Group descriptions output by the second considered group-generating algorithm $\text{GGen}_{\mathbb{F}_\ell^*}$ are subgroups of the multiplicative group \mathbb{F}_ℓ^* of sufficiently large prime order.

Except for the group generators, all group elements will be denoted with uppercase letters, e.g., X . We use vectors and matrices of elements in \mathbb{Z}_p to compute with group elements: If Y is a group element and \mathbf{x} is a vector of elements in \mathbb{Z}_p , we write $Y^{\mathbf{x}}$ to denote the group element vector $(Y^{\mathbf{x}[1]}, Y^{\mathbf{x}[2]}, \dots)$. Similarly, given some matrix $M = (m_{ij})_{i,j \in [1..n] \times [1..k]}$ and a vector of group elements \mathbf{Y} of size k , we define \mathbf{Y}^M to be the n -size vector $(\mathbf{Y}[1]^{m_{11}} \dots \mathbf{Y}[k]^{m_{1k}}, \dots, \mathbf{Y}[1]^{m_{n1}} \dots \mathbf{Y}[k]^{m_{nk}})$. Note that if $\mathbf{Y} = g^{\mathbf{y}}$ then $\mathbf{Y}^M = g^{M\mathbf{y}}$.

SECURITY GAMES. We define security notions via *code-based games* [10]. A game G consists of a main procedure and zero or more oracles that can be accessed from within the game. The game is defined with respect to an adversary \mathcal{A} , which is invoked within the main procedure. The adversary may have access to some of the oracles of the game: The ability to access oracle \mathcal{O} is represented by invoking the adversary as $\mathcal{A}^{\mathcal{O}}$. When the game stops, it outputs either a success (1) or a failure (0) symbol. With $\Pr[G(\mathcal{A})]$ we denote the probability that adversary \mathcal{A} wins, i.e., that game G , executed with respect to \mathcal{A} , stops with output 1.

2.2 Generic/Algebraic Group Model

GENERIC GROUP MODEL. Intuitively, the Generic Group Model (GGM) is an abstraction to study the behavior of adversaries that do not exploit any specific structure of the group at play, but rather treat the group in a black-box fashion. This is usually modeled by representing group elements exclusively through “opaque” handles, which hide the structure of the group. These handles are used as input to a model-bound oracle, the group-operation oracle, which is the only interface to the group available to the adversary. An algorithm with such restrictions is referred to as a *generic algorithm*. The running time of generic adversaries is normally measured in number of calls to the group-operation oracle. For further details on the GGM we refer to the literature [28,23]. To derive bounds on the hardness of solving certain computational problems with respect to $\text{GGen}_{\mathbb{E}(\mathbb{F}_\ell)}$ we model the output elliptic curves as generic groups. For clarity, in this case we denote the group-generating algorithm by GGen_{gg} .

ALGEBRAIC GROUP MODEL. For every group element Z it returns, an *algebraic algorithm* \mathcal{A} must present a description of this element in terms of the elements it has previously seen. That is, if n is the order of the group and X_1, \dots, X_k are the elements that \mathcal{A} received so far from the game, then \mathcal{A} must return some elements $a_1, \dots, a_k \in \mathbb{Z}_n$ such that $Z = X_1^{a_1} \dots X_k^{a_k}$. We use the algebraic group model to analyze generic reductions:

Note that a generic reduction executed with respect to a generic adversary is itself a generic algorithm. Without loss of generality we may assume that generic adversaries are algebraic, which allows the reduction to exploit the useful algebraic representation of the input group elements. As demonstrated by Fuchsbauer et al. [16], this idea gives a handy technique for carrying over generic lower bounds through generic reductions, as seen in the following lemma.

Lemma 1 ([16, Lemma 1]). *Let α, Δ be constants and let \mathcal{R} be a generic reduction \mathcal{R} from game G_1 to G_0 . Assume that for every generic adversary \mathcal{A} that succeeds with probability ε and makes at most q group-operation queries, reduction \mathcal{R} executed with respect to \mathcal{A} makes at most $q + \Delta$ group-operation queries and succeeds with probability of at least $\alpha\varepsilon$. If there exists a function f such that $\Pr[G_1(\mathcal{B})] \leq f(q)$ for every generic adversary \mathcal{B} making at most q group-operation queries, then for every generic adversary \mathcal{A} making at most q group-operation queries we obtain $\Pr[G_0(\mathcal{A})] \leq \alpha^{-1}f(q + \Delta)$.*

2.3 Key-Encapsulation Mechanisms

A *key-encapsulation mechanism* (KEM) KEM specifies the following. Parameter generation algorithm Par generates public parameters par to be utilized by all users. Key-generation algorithm Gen gets the parameters as input and outputs a pair (pk, sk) consisting of a public and a secret key. Encapsulation algorithm Enc on input of the parameters and a public key outputs a pair (K, c) consisting of an encapsulated key K belonging to the encapsulated key space $\text{KS}(par)$

and a ciphertext c belonging to the ciphertext space $\text{CS}(par)$. Deterministic decapsulation algorithm Dec receives the parameters, a secret key sk and a ciphertext c as input and returns either the symbol \perp indicating failure or an encapsulated key K . For *correctness* we require that for all par output of Par and for every (pk, sk) output of $\text{Gen}(par)$ we obtain $K \leftarrow \text{Dec}(par, sk, c)$ for $(K, c) \leftarrow_s \text{Enc}(par, pk)$.

3 Multi-Instance Security

In this section we investigate the m -out-of- n multi-instance security of key-encapsulation mechanisms. After giving security definitions in Section 3.1, in Section 3.2 we consider the relation between security notions for varying m and n . In Section 3.3 we define the scaling factor, which measures how well the security of KEMs scales with the number of users. Finally, in Section 3.4 we give security definitions for Diffie-Hellman type problems in the multi-instance setting, which will be used in the security analysis of the Hashed-ElGamal KEM in the next section.

3.1 Key Encapsulation in the Multi-Instance Setting

Below we give security definitions for key-encapsulation mechanisms in the multi-instance setting. Our definitions are in the xor metric introduced by Bellare et al. [8] for symmetric encryption schemes. We target m -out-of- n multi-instance indistinguishability of encapsulated keys from random against chosen-plaintext attacks ((m, n) -CPA) or chosen-ciphertext attacks ((m, n) -CCA).

In its most general form, the xor metric models the inability of an adversary to break m out of n instances of a decisional problem. The adversary receives as input n challenges, generated independently of each other with respect to n independent challenge bits \mathbf{b} . The adversary’s task is to output a subset $L \subseteq [1..n]$ of size at least m (representing the “broken instances”) together with a guess for $\bigoplus_{i \in L} \mathbf{b}[i]$; the intuition being that as long as at least one of the challenge bits contained in L is hidden to the adversary, so is $\bigoplus_{i \in L} \mathbf{b}[i]$, reducing the adversary to guessing the final output.

Formally, let KEM be a KEM and let $m, n \in \mathbb{N}$ such that $1 \leq m \leq n$. Consider games $G_{\text{KEM}}^{(m,n)\text{-cpa}}(\mathcal{A})$ and $G_{\text{KEM}}^{(m,n)\text{-cca}}(\mathcal{A})$ of Fig. 1 associated with KEM, m, n , and an adversary \mathcal{A} . In both games, \mathbf{b} is a vector of n challenge bits, which corresponds to vectors \mathbf{pk}, \mathbf{sk} of public and secret keys, which are set up using a single set of global parameters par . The adversary has access to a challenge oracle Enc , which on input of index $i \in [1..n]$ returns a pair consisting of an encapsulated key and a ciphertext generated with $\text{Enc}(par, \mathbf{pk}[i])$ if the challenge bit $\mathbf{b}[i]$ equals 1, or, if $\mathbf{b}[i]$ equals 0, a ciphertext and a randomly sampled element of $\text{KS}(par)$. At the end of the game, adversary \mathcal{A} outputs a list of indices $L \subseteq [1..n]$ and a bit b' . \mathcal{A} wins if L contains at least m elements and if $b' = \bigoplus_{i \in L} \mathbf{b}[i]$. In game $G_{\text{KEM}}^{(m,n)\text{-cca}}(\mathcal{A})$ the adversary additionally has access to a decapsulation oracle Dec , which on input of index $i \in [1..n]$ and ciphertext c returns the

Games $G_{\text{KEM}}^{(m,n)\text{-cpa}}(\mathcal{A}), G_{\text{KEM}}^{(m,n)\text{-cca}}(\mathcal{A})$ 00 $\mathcal{C}^*[\cdot] \leftarrow \emptyset$ 01 $\mathbf{b} \leftarrow_{\mathcal{S}} \{0, 1\}^n$ 02 $par \leftarrow_{\mathcal{S}} \text{Par}$ 03 for $i \in [1..n]$: 04 $(\mathbf{pk}[i], \mathbf{sk}[i]) \leftarrow_{\mathcal{S}} \text{Gen}(par)$ 05 $(L, b') \leftarrow_{\mathcal{S}} \mathcal{A}^{\text{Enc}}(par, \mathbf{pk}) \quad \parallel (m, n)\text{-CPA}$ 06 $(L, b') \leftarrow_{\mathcal{S}} \mathcal{A}^{\text{Enc, Dec}}(par, \mathbf{pk}) \quad \parallel (m, n)\text{-CCA}$ 07 if $ L < m$: return 0 08 if $\bigoplus_{i \in L} \mathbf{b}[i] = b'$: return 1 09 else: return 0	Oracle $\text{Enc}(i)$ 10 $(K_1^*, c^*) \leftarrow_{\mathcal{S}} \text{Enc}(par, \mathbf{pk}[i])$ 11 $K_0^* \leftarrow_{\mathcal{S}} \text{KS}(par)$ 12 $\mathcal{C}^*[i] \leftarrow \mathcal{C}^*[i] \cup \{c^*\}$ 13 return $(K_{\mathbf{b}[i]}^*, c^*)$ Oracle $\text{Dec}(i, c)$ 14 if $c \in \mathcal{C}^*[i]$: return \perp 15 $K \leftarrow \text{Dec}(par, \mathbf{sk}[i], c)$ 16 return K
--	---

Fig. 1. Games $G_{\text{KEM}}^{(m,n)\text{-cpa}}$ and $G_{\text{KEM}}^{(m,n)\text{-cca}}$ modeling m -out-of- n multi-instance indistinguishability of encapsulated keys from random. We assume that $L \subseteq [1..n]$.

decapsulation of c under parameters par and secret key $\mathbf{sk}[i]$ (unless c was output as response to a challenge query $\text{Enc}(i)$ for index i).

We define \mathcal{A} 's advantage in game $G_{\text{KEM}}^{(m,n)\text{-cpa}}$ and $G_{\text{KEM}}^{(m,n)\text{-cca}}$ respectively as

$$\begin{aligned} \text{Adv}_{\text{KEM}}^{(m,n)\text{-cpa}}(\mathcal{A}) &= 2 \Pr[G_{\text{KEM}}^{(m,n)\text{-cpa}}(\mathcal{A})] - 1, \\ \text{Adv}_{\text{KEM}}^{(m,n)\text{-cca}}(\mathcal{A}) &= 2 \Pr[G_{\text{KEM}}^{(m,n)\text{-cca}}(\mathcal{A})] - 1. \end{aligned}$$

The definition we have just presented lends itself naturally to a comparison with the standard multi-user security notion of Bellare et al. [7]. We describe the relationship between multi-user security and $(1, n)$ -CCA in detail in the full version of the paper [3].

3.2 Advantage Relations for Different m and n

The relations between (m', n') -CPA and (m, n) -CPA security are summarized in Fig. 2. They are stated more formally in the following theorem. Its proof is in the full version of the paper [3]

Theorem 1. *Let m, n, m', n' be positive integers such that $m \leq n, m' \leq n'$, and let KEM be any KEM scheme. Then for every adversary \mathcal{A} against game $G_{\text{KEM}}^{(m,n)\text{-cpa}}$ there exists an adversary \mathcal{B} against game $G_{\text{KEM}}^{(m',n')\text{-cpa}}$ such that:*

1. *If $m' \leq m$ and $m'n \leq mn'$ then \mathcal{B} has roughly the same running time of \mathcal{A} and*

$$\text{Adv}_{\text{KEM}}^{(m',n')\text{-cpa}}(\mathcal{B}) \geq \frac{1}{2} \text{Adv}_{\text{KEM}}^{(m,n)\text{-cpa}}(\mathcal{A}).$$

2. *Additionally, if $n' - m' \geq n - m$ then the reduction does not lose the factor $1/2$. If $m' \leq m$ and $m'n > mn'$ then \mathcal{B} has roughly the same running time of \mathcal{A} and*

$$\text{Adv}_{\text{KEM}}^{(m',n')\text{-cpa}}(\mathcal{B}) \geq \frac{1}{2} \binom{n'}{m'} \left(\frac{\lceil nm'/m \rceil}{m'} \right)^{-1} \text{Adv}_{\text{KEM}}^{(m,n)\text{-cpa}}(\mathcal{A}).$$




Fig. 2. Relations between (m', n') -CPA and (m, n) -CPA security. Given \mathcal{A} against (m, n) -CPA with advantage ϵ , one can build \mathcal{B} against (m', n') -CPA with advantage as shown in figure, depending on its position on the plane. The constants in the figure are $k = \lceil m'/m \rceil$ and $\ell = \frac{1}{2} \binom{n'}{m'} \left(\frac{\lceil nm'/m \rceil}{m'} \right)^{-1}$. The same result holds for CCA.

3. If $m' > m$ and $m'n \leq mn'$ then \mathcal{B} has roughly $k = \lceil m'/m \rceil$ times the running time of \mathcal{A} and

$$\text{Adv}_{\text{KEM}}^{(m', n')\text{-cpa}}(\mathcal{B}) \geq \frac{1}{2} \left(\text{Adv}_{\text{KEM}}^{(m, n)\text{-cpa}}(\mathcal{A}) \right)^k .$$

Additionally, if m divides m' then the reduction does not lose the factor $1/2$.

4. If $m' > m$ and $m'n > mn'$ then \mathcal{B} has roughly $k = \lceil m'/m \rceil$ times the running time of \mathcal{A} and

$$\text{Adv}_{\text{KEM}}^{(m', n')\text{-cpa}}(\mathcal{B}) \geq \frac{1}{2} \binom{n'}{m'} \left(\frac{\lceil nm'/m \rceil}{m'} \right)^{-1} \left(\text{Adv}_{\text{KEM}}^{(m, n)\text{-cpa}}(\mathcal{A}) \right)^k .$$

An analogous statement holds between (m, n) -CCA and (m', n') -CCA. If \mathcal{A} queries its decryption oracle q times, then adversary \mathcal{B} queries its decryption oracle at most q , q , kq , and kq times respectively.

3.3 Scaling Factor

We now define the scaling factor of key-encapsulation mechanisms. To be able to give an intuitive and accessible definition we treat the running time and advantages of adversaries as if they were elements of \mathbb{R} and $[0, 1]$ respectively. A formal definition that takes the asymptotic nature of running time and advantage into account as well as rigorous proofs for the bounds on the scaling factor derived in this section can be found in the full version of the paper [3]. We start with a definition for adversaries succeeding with advantage 1 and afterwards give a generalized version for arbitrary advantages.

We fix a computational model that associates each adversary \mathcal{A} with its running time. Let $\text{MinTime}_{\text{KEM}}^{(m,n)\text{-cpa}}$ be the minimal time T for which there exists an adversary \mathcal{A} that runs in at most time T and achieves advantage $\text{Adv}_{\text{KEM}}^{(m,n)\text{-cpa}}(\mathcal{A}) = 1$.

We define the scaling factor of KEM relative to (m, n) -CPA security as

$$\text{SF}_{\text{KEM}}^{(m,n)\text{-cpa}} := \frac{\text{MinTime}_{\text{KEM}}^{(m,n)\text{-cpa}}}{\text{MinTime}_{\text{KEM}}^{(1,1)\text{-cpa}}} .$$

The scaling factor of KEM relative to (m, n) -CCA security, $\text{SF}_{\text{KEM}}^{(m,n)\text{-cca}}$, is defined in the same way relative to advantage $\text{Adv}_{\text{KEM}}^{(m,n)\text{-cca}}(\mathcal{A})$. By the results of Section 3.2 we can give the following bounds on the scaling factor (which also hold in the CCA setting):

$$\text{SF}_{\text{KEM}}^{(m,n)\text{-cpa}} \leq \text{SF}_{\text{KEM}}^{(m,m)\text{-cpa}} \leq m$$

The lower bound follows since any adversary against (m, m) -CPA is also an adversary against (m, n) -CPA with the same advantage (Theorem 1, item 1). The upper bound follows from Theorem 1, item 3. Surprisingly, the scaling factor can be smaller than 1: Being able to choose which users to attack can make the task of breaking multiple instances easier than breaking a single one. An artificial example of a KEM with scaling factor of m/n is sketched in the full version of the paper [3]. This is, however, a phenomenon limited to the case $m \neq n$: For $n = m$, we know that $\text{SF}_{\text{KEM}}^{(n,m)\text{-cpa}} \geq 1$ by Theorem 1, item 1. Importantly, specific KEMs such as HEG or Cramer-Shoup are known to be “random self-reducible”, which implies $\text{MinTime}_{\text{KEM}}^{(1,n)\text{-cpa}} = \text{MinTime}_{\text{KEM}}^{(1,1)\text{-cpa}}$, and hence by Theorem 1, item 1:

$$1 \leq \text{SF}_{\text{KEM}}^{(m,n)\text{-cpa}} \leq m .$$

The definition given above exclusively considers adversaries that achieve advantage 1. This definition generalizes naturally to encompass adversaries with arbitrary advantage as follows. Let $\text{MinTime}_{\text{KEM}}^{(m,n)\text{-cpa}}(\varepsilon)$, associated with $0 \leq \varepsilon \leq 1$, denote the running time of the fastest adversary achieving advantage at least ε in game (m, n) -CPA. Intuitively, an optimally scaling scheme requires m independent execution of a $(1, 1)$ -CPA adversary in order to break m instances of the scheme. Hence, the advantage-dependent scaling factor for advantage ε is defined as

$$\text{SF}_{\text{KEM}}^{(m,n)\text{-cpa}}(\varepsilon) := \text{MinTime}_{\text{KEM}}^{(m,n)\text{-cpa}}(\varepsilon^m) / \text{MinTime}_{\text{KEM}}^{(1,1)\text{-cpa}}(\varepsilon) .$$

Again, we can use Theorem 1 to show that, for every $0 \leq \varepsilon \leq 1$,

$$\text{SF}_{\text{KEM}}^{(m,n)\text{-cpa}}(\varepsilon) \leq \text{SF}_{\text{KEM}}^{(m,m)\text{-cpa}}(\varepsilon) \leq m .$$

3.4 Multi-Instance Diffie-Hellman-Type Problems

GAP DISCRETE LOGARITHM PROBLEM. The m -out-of- n multi-instance gap discrete logarithm problem ((m, n) -GapDL) requires to find the discrete logarithms

of at least m out of n input group elements given access to a decisional Diffie-Hellman oracle. We consider three variants of the problem, which differ in their granularity. For high granularity all discrete logarithm challenges are sampled with respect to a fixed group and group generator, while for medium granularity the challenges are elements of a fixed group but defined with respect to different group generators. Finally, in the case of low granularity a fresh group and generator is used for each challenge.

Formally, let $m, n \in \mathbb{N}$ such that $1 \leq m \leq n$ and consider game $G_{\text{GGen, gran}}^{(m,n)\text{-gdl}}(\mathcal{A})$ of Fig. 3 associated with adversary \mathcal{A} , group-generating algorithm GGen , and granularity $\text{gran} \in \{\text{high, med, low}\}$. In the game, a vector \mathcal{G} of n group descriptions is set up according to the desired level of granularity using parameter generation algorithm $\text{PGen}[\text{gran}]$. Each entry of \mathcal{G} is of the form (\mathbb{G}, p, g) with \mathbb{G} being a group of prime order p generated by g . After the setup of \mathcal{G} the three variants of the game proceed in the same way. A vector \mathbf{x} of length n is sampled, where $\mathbf{x}[i]$ is uniformly distributed in $\mathbb{Z}_{p[i]}$. The corresponding challenge vector contains the group elements $\mathbf{X}[i] = \mathbf{g}[i]^{\mathbf{x}[i]}$. At the end of the game, adversary \mathcal{A} outputs a list of indices $L \subseteq [1..n]$ and a vector \mathbf{x}' of length n , where the i -th entry is in $\mathbb{Z}_{p[i]}$. The adversary wins if L contains at least m elements and if the vector \mathbf{x}' coincides with \mathbf{x} for all indices in L . Additionally, the adversary has access to an oracle DDH, which, on input of index $i \in [1..n]$ and three group elements $\hat{X}, \hat{Y}, \hat{Z}$, behaves as follows. The game computes the discrete logarithms \hat{x}, \hat{y} of input \hat{X}, \hat{Y} with respect to generator $\mathbf{g}[i]$, and then returns 1 if and only if $\mathbf{g}[i]^{\hat{x}\hat{y}} = \hat{Z}$.

We define \mathcal{A} 's advantage in game $G_{\text{GGen, gran}}^{(m,n)\text{-gdl}}(\mathcal{A})$ as

$$\text{Adv}_{\text{GGen, gran}}^{(m,n)\text{-gdl}}(\mathcal{A}) = \Pr[G_{\text{GGen, gran}}^{(m,n)\text{-gdl}}(\mathcal{A})] .$$

The m -out-of- n multi-instance discrete logarithm ((m, n) -DL) problem is defined as (m, n) -GapDL with the restriction that \mathcal{A} cannot query DDH.

GAP COMPUTATIONAL DIFFIE-HELLMAN PROBLEM. The m -out-of- n multi-instance gap computational Diffie-Hellman problem ((m, n) -GapCDH) requires, on input of vectors $\mathbf{g}^{\mathbf{x}}$ and $\mathbf{g}^{\mathbf{y}}$, to compute at least m elements of the form $\mathbf{g}^{\mathbf{x}[i]\mathbf{y}[i]}$ for distinct $i \in [1..n]$. As in the corresponding DL game, the adversary has access to an oracle DDH which computes whether three given group elements are a Diffie-Hellman triple. As in the definition of (m, n) -GapDL, we consider three variants of the problem, which differ in their granularity.

Formally, for $m, n \in \mathbb{N}$ s.t. $1 \leq m \leq n$ consider game $G_{\text{GGen, gran}}^{(m,n)\text{-gcdh}}(\mathcal{A})$ of Fig. 4 associated with adversary \mathcal{A} , group-generating algorithm GGen , and granularity $\text{gran} \in \{\text{high, med, low}\}$. In the game, a vector \mathcal{G} of n group descriptions is set up according to parameter generation algorithm $\text{PGen}[\text{gran}]$. After the setup of \mathcal{G} the three variants of the game proceed in the same way. Two vectors \mathbf{x}, \mathbf{y} of length n are sampled, where $\mathbf{x}[i], \mathbf{y}[i]$ are uniformly distributed in $\mathbb{Z}_{p[i]}$. The corresponding challenge vectors contain the group elements $\mathbf{X}[i] = \mathbf{g}[i]^{\mathbf{x}[i]}$ and $\mathbf{Y}[i] = \mathbf{g}[i]^{\mathbf{y}[i]}$. Additionally, the adversary has access to an oracle DDH, which behaves as described for $G_{\text{GGen, gran}}^{(m,n)\text{-gdl}}(\mathcal{A})$. At the end of the game, adversary \mathcal{A}

Games $G_{\text{GGen,gran}}^{(m,n)\text{-gdl}}(\mathcal{A})$		Oracle $\text{DDH}(i, \hat{X}, \hat{Y}, \hat{Z})$	
00 $\mathcal{G} \leftarrow_{\S} \text{PGen}[\text{gran}]$		06 parse \hat{X}, \hat{Y} as $g[i]^{\hat{x}}, g[i]^{\hat{y}}$	
01 $\mathbf{x}[\cdot] \leftarrow_{\S} \mathbb{Z}_p[\cdot]; \mathbf{X}[\cdot] \leftarrow g[\cdot]^{\mathbf{x}[\cdot]}$		07 if $g[i]^{\hat{x}\hat{y}} = \hat{Z}$:	
02 $(L, \mathbf{x}') \leftarrow_{\S} \mathcal{A}^{\text{DDH}}(\mathcal{G}, \mathbf{X})$		08 return 1	
03 if $ L < m$: return 0		09 else: return 0	
04 if $\mathbf{x}'[L] = \mathbf{x}[L]$: return 1			
05 else: return 0			
Procedure $\text{PGen}[\text{high}]$	Procedure $\text{PGen}[\text{med}]$	Procedure $\text{PGen}[\text{low}]$	
10 $\mathcal{G} = (\mathbb{G}, p, g) \leftarrow_{\S} \text{GGen}$	13 $(\mathbb{G}, p, g) \leftarrow_{\S} \text{GGen}$	17 $\mathcal{G}[\cdot] \leftarrow_{\S} \text{GGen}$	
11 $\mathcal{G}[\cdot] \leftarrow \mathcal{G}$	14 $g \leftarrow_{\S} (\mathbb{G} \setminus \{1\})^n$	18 return \mathcal{G}	
12 return \mathcal{G}	15 $\mathcal{G}[\cdot] \leftarrow (\mathbb{G}, p, g[\cdot])$		
	16 return \mathcal{G}		

Fig. 3. Security game $G_{\text{GGen,gran}}^{(m,n)\text{-gdl}}(\mathcal{A})$ for $\text{gran} \in \{\text{high}, \text{med}, \text{low}\}$ modeling the m -out-of- n multi-instance gap discrete logarithm problem.

Game $G_{\text{GGen,gran}}^{(m,n)\text{-gcdh}}(\mathcal{A})$		Oracle $\text{DDH}(i, \hat{X}, \hat{Y}, \hat{Z})$	
00 $\mathcal{G} \leftarrow_{\S} \text{PGen}[\text{gran}]$		08 parse \hat{X}, \hat{Y} as $g[i]^{\hat{x}}, g[i]^{\hat{y}}$	
01 $\mathbf{x}[\cdot] \leftarrow_{\S} \mathbb{Z}_p[\cdot]; \mathbf{X}[\cdot] \leftarrow g[\cdot]^{\mathbf{x}[\cdot]}$		09 if $g[i]^{\hat{x}\hat{y}} = \hat{Z}$:	
02 $\mathbf{y}[\cdot] \leftarrow_{\S} \mathbb{Z}_p[\cdot]; \mathbf{Y}[\cdot] \leftarrow g[\cdot]^{\mathbf{y}[\cdot]}$		10 return 1	
03 $\mathbf{Z}[\cdot] \leftarrow g[\cdot]^{\mathbf{x}[\cdot]\mathbf{y}[\cdot]}$		11 else: return 0	
04 $(L, \mathbf{Z}') \leftarrow_{\S} \mathcal{A}^{\text{DDH}}(\mathcal{G}, \mathbf{X}, \mathbf{Y})$			
05 if $ L < m$: return 0			
06 if $\mathbf{Z}[L] = \mathbf{Z}'[L]$: return 1			
07 else: return 0			

Fig. 4. Security game $G_{\text{GGen,gran}}^{(m,n)\text{-gcdh}}(\mathcal{A})$ for $\text{gran} \in \{\text{high}, \text{med}, \text{low}\}$ modeling the m -out-of- n multi-instance gap computational Diffie-Hellman problem. PGen is defined in Fig. 3.

outputs a list of indices $L \subseteq [1..n]$ and a vector \mathbf{Z}' of length n , where the i -th entry is an element of the group represented by $\mathcal{G}[i]$. The adversary wins if L contains at least m elements and if the vector \mathbf{Z}' coincides with \mathbf{Z} for all indices in L . We define \mathcal{A} 's advantage in game $G_{\text{GGen,gran}}^{(m,n)\text{-gcdh}}(\mathcal{A})$ as

$$\text{Adv}_{\text{GGen,gran}}^{(m,n)\text{-gcdh}}(\mathcal{A}) = \Pr[G_{\text{GGen,gran}}^{(m,n)\text{-gcdh}}(\mathcal{A})] .$$

The m -out-of- n multi-instance computational Diffie-Hellman ((m, n) -CDH) problem is defined as (m, n) -GapCDH with the restriction that \mathcal{A} cannot query oracle DDH.

4 Hashed ElGamal in the Multi-Instance Setting

We investigate the multi-instance security of the well-known Hashed-ElGamal key-encapsulation mechanism [1]. We consider three variants, $\text{HEG}[\text{GGen, high}]$,

HEG[GGen, med], and HEG[GGen, low], corresponding to high, medium, and low granularity respectively. After giving formal definitions of these variants in Section 4.1, in Section 4.2 we prove the main result of this section: The multi-instance security of each variant of the KEM in the random oracle model is tightly implied by the hardness of (m, n) -GapCDH[GGen, gran] for the corresponding granularity. Finally, in Section 4.3 we compute lower bounds on the scaling factor of HEG[GGen, gran] for $\text{GGen} \in \{\text{GGen}_{\mathbb{F}_\ell^*}, \text{GGen}_{\mathbb{E}(\mathbb{F}_\ell)}\}$ and $\text{gran} \in \{\text{high}, \text{med}, \text{low}\}$.

4.1 Hashed-ElGamal Key Encapsulation

We consider three variants of the Hashed-ElGamal KEM, defined relative to a hash function H and differing in the way parameters and key pairs are generated. For high granularity the parameters specify a group description $\mathcal{G} = (\mathbb{G}, p, g)$ with a fixed generator g . Key pairs (pk, sk) are of the form $pk = X = g^x$ and $sk = x$, where x is randomly sampled in \mathbb{Z}_p . For medium granularity the parameters consist of a group \mathbb{G} of order p , but no fixed generator. In this case $pk = (g, g^x)$ and $sk = (g, x)$, where g is a randomly chosen generator of the group \mathbb{G} . Finally, for low granularity empty parameters are used. Correspondingly, in this case public keys are of the form $pk = (\mathcal{G}, g^x)$ and secret keys of the form $sk = (\mathcal{G}, x)$, where $\mathcal{G} = (\mathbb{G}, p, g)$ is a freshly sampled group description.

Note that in all three cases the parameters par and a key pair (pk, sk) generated with respect to par determine a group description (\mathbb{G}, p, g) as well as x and X . In all three variants encapsulated keys are of the form $H(pk, g^y, X^y)$ with corresponding ciphertext g^y , where the y is sampled at random in \mathbb{Z}_p . The decapsulation of a ciphertext c is given by $H(pk, c, c^x)$. A formal description of the algorithms describing the Hashed-ElGamal key-encapsulation mechanism for each of the three considered variants can be found in Fig. 5.

4.2 Multi-Instance Security of Hashed ElGamal

The following theorem shows that (m, n) -GapCDH tightly reduces to the security against chosen-ciphertext attacks of HEG in the multi-instance setting for the corresponding granularity⁶. Its proof is a generalization of the single-instance version [1] and can be found in the full version of the paper [3].

Theorem 2. *Let $m, n \in \mathbb{N}$ with $1 \leq m \leq n$, let $\text{gran} \in \{\text{high}, \text{med}, \text{low}\}$, let GGen be a group-generating algorithm, and let $\text{HEG}[\text{GGen}, \text{gran}]$ be the Hashed-ElGamal KEM of Fig. 5 relative to hash function H . If H is modeled as a random oracle and if the (m, n) -GapCDH[GGen, gran] problem is hard, then $\text{HEG}[\text{GGen}, \text{gran}]$ is (m, n) -CCA secure. Formally, for every adversary \mathcal{A} against game $\mathsf{G}_{\text{HEG}[\text{GGen}, \text{gran}]}^{(m, n)\text{-cca}}$ making at most q queries to random oracle RO there exists*

⁶ The same result holds under the multi-instance version of the strong Diffie-Hellman assumption [1], a falsifiable assumption that is implied by (m, n) -GapCDH.

gran = high	gran = med	gran = low
Alg. Par[high]	Alg. Par[med]	Alg. Par[low]
00 $\mathcal{G} = (\mathbb{G}, p, g) \leftarrow_{\mathcal{S}} \text{GGen}$	06 $\mathcal{G} = (\mathbb{G}, p, g) \leftarrow_{\mathcal{S}} \text{GGen}$	13 $par \leftarrow \perp$
01 $par \leftarrow \mathcal{G}$	07 $par \leftarrow (\mathbb{G}, p)$	14 return par
02 return par	08 return par	
Alg. Gen[high](par)	Alg. Gen[med](par)	Alg. Gen[low](par)
03 $x \leftarrow_{\mathcal{S}} \mathbb{Z}_p; X \leftarrow g^x$	09 $g \leftarrow_{\mathcal{S}} \mathbb{G} \setminus \{1\}$	15 $\mathcal{G} = (\mathbb{G}, p, g) \leftarrow_{\mathcal{S}} \text{GGen}$
04 $pk \leftarrow X; sk \leftarrow x$	10 $x \leftarrow_{\mathcal{S}} \mathbb{Z}_p; X \leftarrow g^x$	16 $x \leftarrow_{\mathcal{S}} \mathbb{Z}_p; X \leftarrow g^x$
05 return (pk, sk)	11 $pk \leftarrow (g, X); sk \leftarrow (g, x)$	17 $pk \leftarrow (\mathcal{G}, X); sk \leftarrow (\mathcal{G}, x)$
	12 return (pk, sk)	18 return (pk, sk)
	Alg. Enc(par, pk)	Alg. Dec(par, sk, c)
	19 $y \leftarrow_{\mathcal{S}} \mathbb{Z}_p$	23 $K \leftarrow H(pk, c, c^x)$
	20 $c \leftarrow g^y$	24 return K
	21 $K \leftarrow H(pk, c, X^y)$	
	22 return (K, c)	

Fig. 5. Variants of Hashed-ElGamal KEM $\text{HEG}[\text{GGen}, \text{high}]$, $\text{HEG}[\text{GGen}, \text{med}]$, and $\text{HEG}[\text{GGen}, \text{low}]$ relative to hash function H and group-generating algorithm GGen . The KEMs share the same encapsulation and decapsulation algorithms. Note that both (par, pk) or (par, sk) determine group description (\mathbb{G}, p, g) and key pk .

an adversary \mathcal{B} against game $\text{G}_{\text{GGen}, \text{gran}}^{(m,n)\text{-gcdh}}$ that makes at most q queries to DDH and runs in essentially the same time as \mathcal{A} and satisfies

$$\text{Adv}_{\text{GGen}, \text{gran}}^{(m,n)\text{-gcdh}}(\mathcal{B}) \geq \text{Adv}_{\text{HEG}[\text{GGen}, \text{gran}]}^{(m,n)\text{-cca}}(\mathcal{A}) .$$

4.3 Scaling Factor of Hashed ElGamal for Different Parameters

Below we compute the scaling factor of Hashed-ElGamal key encapsulation for different parameter choices. Recall that the scaling factor is given by

$$\text{SF}_{\text{HEG}[\text{GGen}, \text{gran}]}^{(m,n)\text{-cca}} = \text{MinTime}_{\text{HEG}[\text{GGen}, \text{gran}]}^{(m,n)\text{-cca}} / \text{MinTime}_{\text{HEG}[\text{GGen}, \text{gran}]}^{(1,1)\text{-cca}} .$$

Note that the multi-instance security of HEG can be broken by computing m public keys, which corresponds to computing m DL instances. On the other hand, from Theorem 2 we know that the (m, n) -CCA-security of HEG is tightly implied by (m, n) -GapCDH. Thus,

$$\text{MinTime}_{\text{GGen}, \text{gran}}^{(m,n)\text{-gcdh}} \leq \text{MinTime}_{\text{HEG}[\text{GGen}, \text{gran}]}^{(m,n)\text{-cca}} \leq \text{MinTime}_{\text{GGen}, \text{gran}}^{(m,n)\text{-dl}} .$$

Hence, we can bound the scaling factor of Hashed ElGamal as

$$\text{SF}_{\text{HEG}[\text{GGen}, \text{gran}]}^{(m,n)\text{-cca}} \geq \text{MinTime}_{\text{GGen}, \text{gran}}^{(m,n)\text{-gcdh}} / \text{MinTime}_{\text{GGen}, \text{gran}}^{(1,1)\text{-dl}} .$$

Below we consider two instantiations of group-generating algorithms: $\text{GGen}_{\mathbb{F}_\ell^*}$ and $\text{GGen}_{\mathbb{E}(\mathbb{F}_\ell)}$. Due to either Hypothesis 1 from the introduction or the results of

Sections 5 and 6 respectively, for both instantiations solving (m, n) -GapCDH is as hard as (m, n) -GapDL. Thus, the lower bounds on the scaling factor derived below are sharp.

HASHED ELGAMAL IN THE FINITE-FIELD SETTING. Assuming the correctness of Hypothesis 1, we conclude that $\text{MinTime}_{\mathbb{F}_\ell^*, \text{gran}}^{(m, n)\text{-gcdh}} = \text{MinTime}_{\mathbb{F}_\ell^*, \text{gran}}^{(m, n)\text{-dl}}$ is given by

$$L_\ell(1/3, 1.902) + m \cdot L_\ell(1/3, 1.232) \quad \text{for gran} \in \{\text{high}, \text{med}\}, \text{ and} \\ \min\{m \cdot L_\ell(1/3, 1.902), L_\ell(1/3, 2.007) + m \cdot L_\ell(1/3, 1.639)\} \quad \text{for gran} = \text{low}.$$

We obtain the scaling factor by dividing by $\text{MinTime}_{\mathbb{F}_\ell^*, \text{gran}}^{(1, 1)\text{-dl}} = L_\ell(1/3, 1.902)$. Defining δ via $m = L_\ell(1/3, \delta)$ we can rewrite $m \cdot L_\ell(1/3, 1.232)$ as $L_\ell(1/3, \delta + 1.232)$. For $\delta \leq 0.67$ we get $L_\ell(1/3, 1.902) \geq L_\ell(1/3, \delta + 1.232)$. Hence for these values of δ the scaling factor for medium and high granularity is roughly 1. For larger m , on the other hand, it is of order $L_\ell(1/3, \delta - 0.67)$.

Summing up for $\text{gran} \in \{\text{med}, \text{high}\}$ we obtain

$$\text{SF}_{\text{HEG}[\text{GGen}_{\mathbb{F}_\ell^*, \text{gran}}]}^{(m, n)\text{-cca}} = \begin{cases} 1 & \delta \leq 0.67 \\ L_\ell(1/3, \delta - 0.67) & \delta > 0.67 \end{cases}.$$

Further, we get $L_\ell(1/3, \delta + 1.902) \leq L_\ell(1/3, 2.007)$ for $\delta \leq 0.105$. Hence in this case for low granularity the scaling factor is given by $m = L_\ell(1/3, \delta)$. Moreover, we obtain $L_\ell(1/3, \delta + 1.639) = L_\ell(1/3, 2.007)$ for $\delta = 0.368$ implying that for $0.105 \leq \delta \leq 0.368$ the scaling factor is of order $L_\ell(1/3, 2.007 - 1.902)$ and of order $L_\ell(1/3, \delta + 1.639 - 1.902)$ for larger values of δ . Summing up:

$$\text{SF}_{\text{HEG}[\text{GGen}_{\mathbb{F}_\ell^*, \text{low}}]}^{(m, n)\text{-cca}} = \begin{cases} L_\ell(1/3, \delta) & 0 \leq \delta < 0.105 \\ L_\ell(1/3, 0.105) & 0.105 \leq \delta < 0.368 \\ L_\ell(1/3, -0.263 + \delta) & 0.368 \leq \delta \end{cases}.$$

Formally, the asymptotic behavior of the scaling factor computed above is linear⁷ in m and hence, at first glance, seems optimal. However, as discussed in the introduction, the numbers of $L_\ell(1/3, 0.67)$ or $L_\ell(1/3, 0.368)$ instances starting from which the cumulative cost of breaking the instances outweighs the cost of the precomputation are typically large.

HASHED ELGAMAL IN THE ELLIPTIC-CURVE SETTING. Recall that $\text{GGen}_{\mathbb{F}_\ell}$ generates elliptic curves of size $p \approx \ell$ defined over the field \mathbb{F}_ℓ for randomly chosen ℓ . If we model elliptic curves as generic groups we can derive the scaling factor as follows. Ignoring constants, a single DL instance can be solved in time $O(\sqrt{p})$. The lower bounds derived in Section 6 (Corollaries 2 and 3 and Theorem 5) imply the following: A generic algorithm solving (m, n) -GapCDH for high and medium

⁷ For fixed ℓ and very large values of m and n generic attacks start to outperform the NFS and the scaling factor actually becomes $\Theta(\sqrt{m})$.

granularity performs at least $\Omega(\sqrt{mp})$ group operations; the low-granularity case requires at least $\Omega(m\sqrt{p})$ group operations. (In the low-granularity case we formally consider n groups of differing group orders p_1, \dots, p_n , where all p_i are roughly of size p .) Summing up, we obtain

$$\text{SF}_{\text{HEG}[\text{GGen}_{\mathbb{F}_\ell}, \text{gran}]}^{(m,n)\text{-cca}} = \begin{cases} \Theta(\sqrt{mp}/\sqrt{p}) = \Theta(\sqrt{m}) & \text{gran} \in \{\text{high}, \text{med}\} \\ \Theta(m\sqrt{p}/\sqrt{p}) = \Theta(m) & \text{gran} = \text{low} \end{cases}.$$

(The constants hidden within the Θ notation can be made explicit from our results, and are between 0.1 and 6.6.) In the full version of the paper [3] we additionally illustrate how the scaling factors computed above could be taken into account when choosing parameters for HEG.

5 Generic Hardness of the Multi-Instance Gap Discrete Logarithm Problem

In this section we define a new hard problem, namely the polycheck discrete logarithm problem (PolyDL), in the multi-instance setting. Then, we proceed to show a concrete bound on its security in the generic group model (Theorem 3). Most notably, from this bound we present a concrete bound on the security of GapDL. To prove the bound we define an additional problem, the *search-by-hypersurface* problem (SHS). In Section 5.1 we define the PolyDL and SHS problems. In Section 5.2 we derive the bound on the security of GapDL in the high granularity setting, and further argue that it is optimal. Bounds for the cases of medium and low granularity can be found in the full version of the paper [3].

5.1 Polycheck Discrete Logarithm and Search-by-Hypersurface Problem

POLYCHECK DISCRETE LOGARITHM PROBLEM. The m -out-of- n multi-instance polycheck discrete logarithm problem ((m, n) - d -PolyDL) for polynomials of degree at most d requires to find the discrete logarithms of at least m out of n input group elements given access to a decisional oracle Eval which behaves as follows. Eval takes as input a polynomial $f \in \mathbb{Z}_p[X_1, \dots, X_k]$ of degree at most d and a list of group elements $(g^{\hat{x}_1}, \dots, g^{\hat{x}_k})$, where k is an arbitrary integer, and returns 1 if and only if $g^{f(\hat{x}_1, \dots, \hat{x}_k)} = 1$. As usual, we consider three variants of the problem, which differ in their granularity.

Formally, let $m, n, d \in \mathbb{N}$ such that $1 \leq m \leq n$, $d \geq 1$, and consider game $\text{G}_{\text{GGen}, \text{gran}}^{(m,n)\text{-}d\text{-polydl}}(\mathcal{A})$ of Fig. 6 associated with adversary \mathcal{A} and granularity $\text{gran} \in \{\text{high}, \text{med}, \text{low}\}$. In the game, a vector \mathcal{G} of n group descriptions is set up according to the desired level of granularity using $\text{PGen}[\text{gran}]$. After the setup of \mathcal{G} the three variants of the game proceed in the same way. A vector \mathbf{x} of length n is sampled, where $\mathbf{x}[i]$ is uniformly distributed in $\mathbb{Z}_{p[i]}$. The

Game $G_{\text{GGen, gran}}^{(m,n)-d\text{-polydl}}(\mathcal{A})$	Oracle $\text{Eval}(i, f, \hat{\mathbf{X}})$
00 $\mathcal{G} \leftarrow_{\$} \text{PGen}[\text{gran}]$	06 if $\deg f > d$: return 0
01 $\mathbf{x}[\cdot] \leftarrow_{\$} \mathbb{Z}_p[\cdot]; \mathbf{X}[\cdot] \leftarrow \mathbf{g}[\cdot]^{\mathbf{x}[\cdot]}$	07 parse $\hat{\mathbf{X}}$ as $\mathbf{g}[i]^{\hat{\mathbf{x}}}$
02 $(L, \mathbf{x}') \leftarrow_{\$} \mathcal{A}^{\text{Eval}}(\mathcal{G}, \mathbf{X})$	08 if $\mathbf{g}[i]^{f(\hat{\mathbf{x}})} = 1$:
03 if $ L < m$: return 0	09 return 1
04 if $\mathbf{x}'[L] = \mathbf{x}[L]$: return 1	10 else: return 0
05 else: return 0	

Fig. 6. Security game $G_{\text{GGen, gran}}^{(m,n)-d\text{-polydl}}(\mathcal{A})$ relative to GGen, gran , modeling the m -out-of- n multi-instance polycheck discrete logarithm problem for polynomials of degree at most d . We assume that polynomial f input to Eval has $|\hat{\mathbf{X}}|$ indeterminates. PGen is defined in Fig. 3.

corresponding challenge vector contains the group elements $\mathbf{X}[i] = \mathbf{g}[i]^{\mathbf{x}[i]}$. At the end of the game, adversary \mathcal{A} outputs a list of indices $L \subseteq [1..n]$ and a vector \mathbf{x}' of length n , where the i -th entry is in $\mathbb{Z}_p[i]$. The adversary wins if L contains at least m elements and if the vector \mathbf{x}' coincides with \mathbf{x} for all indices in L . Additionally, the adversary has access to an evaluation oracle Eval , which on input of an index $i \in [1..n]$, a polynomial $f \in \mathbb{Z}_p[X_1, \dots, X_k]$, and a list of group elements $\hat{\mathbf{X}} = (\hat{\mathbf{X}}[1], \dots, \hat{\mathbf{X}}[k])$, where k is an arbitrary integer which might be different on different calls, behaves as follows. If $\deg f > d$, then Eval returns 0. Otherwise, the game computes the discrete logarithms $\hat{\mathbf{x}}$ of the input elements $\hat{\mathbf{X}}$ with respect to generator $\mathbf{g}[i]$, and then returns 1 if and only if $\mathbf{g}[i]^{f(\hat{\mathbf{x}}[1], \dots, \hat{\mathbf{x}}[k])} = 1$.

We define the advantage of \mathcal{A} in game $G_{\text{GGen, gran}}^{(m,n)-d\text{-polydl}}(\mathcal{A})$ as

$$\text{Adv}_{\text{GGen, gran}}^{(m,n)-d\text{-polydl}}(\mathcal{A}) = \Pr[G_{\text{GGen, gran}}^{(m,n)-d\text{-polydl}}(\mathcal{A})] .$$

The next definition extends the search-by-hyperplane-query problem (SHQ) by Yun [30].

SEARCH-BY-HYPERSURFACE PROBLEM. The search-by-hypersurface problem in dimension n for polynomials of degree at most d ($n\text{-SHS}_d$) requires to find a randomly sampled point \mathbf{a} of the space by adaptively checking whether point \mathbf{a} is contained in the queried hypersurface (i.e., the set of zeroes of a polynomial).

Formally, let $n, d, p \in \mathbb{N}$ such that p is prime and $d, n \geq 1$, and consider game $G_p^{n\text{-shs}_d}(\mathcal{A})$ of Fig. 7 associated with adversary \mathcal{A} . In the game, a vector \mathbf{a} of length n is sampled, where $\mathbf{a}[i]$ is uniformly distributed in \mathbb{Z}_p . At the end of the game, adversary \mathcal{A} outputs a vector $\mathbf{a}' \in \mathbb{Z}_p^n$. The adversary wins if $\mathbf{a}' = \mathbf{a}$. Additionally, the adversary has access to an evaluation oracle Eval , which on input of a polynomial $f \in \mathbb{Z}_p[X_1, \dots, X_n]$ behaves as follows. If $\deg f > d$, then Eval returns 0. Otherwise, the oracle returns 1 if and only if $f(\mathbf{a}) = 0$.

We define the advantage of \mathcal{A} in game $G_p^{n\text{-shs}_d}(\mathcal{A})$ as

$$\text{Adv}_p^{n\text{-shs}_d}(\mathcal{A}) = \Pr[G_p^{n\text{-shs}_d}(\mathcal{A})] .$$

Game $G_p^{n\text{-shsd}}(\mathcal{A})$	Oracle Eval(f)
00 $\mathbf{a} \leftarrow_{\S} \mathbb{Z}_p^n$	04 if $\deg(f) > d$: return 0
01 $\mathbf{a}' \leftarrow_{\S} \mathcal{A}^{\text{Eval}}(p)$	05 if $f(\mathbf{a}) = 0$: return 1
02 if $\mathbf{a}' = \mathbf{a}$: return 1	06 else: return 0
03 else: return 0	

Fig. 7. Security game $G_p^{n\text{-shsd}}(\mathcal{A})$ with respect to integer d and prime p modeling the search-by-hypersurface problem on dimension n for polynomials of degree at most d . All inputs f to oracle Eval are elements of the polynomial ring $\mathbb{Z}_p[X_1, \dots, X_n]$.

5.2 Generic Hardness of High-Granularity (m, n) - d -PolyDL

Below, we state the main result of this section, an explicit upper bound on the security of high-granularity (n, n) - d -PolyDL in the generic group model.

Note that this bound is of particular interest in the context of generic bilinear (or even multilinear) maps. In fact, a d -linear map yields a natural way to compute any answer of oracle Eval for polynomials of degree at most d in the base group.

Theorem 3. *Let n, d be positive integers and p a prime number. Let GGen_{gg} be a group-generating algorithm that generates generic groups of exactly size p . Then for every generic adversary \mathcal{A} against (n, n) - d -PolyDL[$\text{GGen}_{\text{gg}}, \text{high}$] that makes at most q queries to the group-operation oracle and q_{Eval} queries to oracle Eval:*

$$\text{Adv}_{\text{GGen}_{\text{gg}}, \text{high}}^{(n, n)\text{-}d\text{-polydl}}(\mathcal{A}) \leq \left(\frac{d}{p}\right)^n + \frac{1}{2} \left(\frac{ed(q+n+1)^2 + 2edq_{\text{Eval}}}{2np} \right)^n.$$

This extends [30, Corollary 2] from standard DL to the polycheck case. Most importantly, it allows us to prove the following corollary.

Corollary 1. *Let n be any positive integer and GGen_{gg} be a group-generating algorithm that generates generic groups of at least size p . Then for every generic adversary \mathcal{A} against (n, n) -GapDL[$\text{GGen}_{\text{gg}}, \text{high}$] that makes at most q queries to the group-operation oracle and q_{DDH} queries to the DDH oracle:*

$$\text{Adv}_{\text{GGen}_{\text{gg}}, \text{high}}^{(n, n)\text{-}gdl}(\mathcal{A}) \leq \left(\frac{2}{p}\right)^n + \frac{1}{2} \left(\frac{e(q+n+1)^2 + 2eq_{\text{DDH}}}{np} \right)^n \approx \left(\frac{q^2}{np}\right)^n.$$

Proof (Corollary 1). Note that oracle DDH of game (n, n) -GapDL can be simulated using oracle Eval from game (n, n) -2-PolyDL. In fact, $g^{xy} = g^z$ if and only if $g^{f(x, y, z)} = 1$, with $f(X_1, X_2, X_3) := X_1X_2 - X_3$. Then apply Theorem 3 with $d = 2$. \square

The result is optimal. Concretely, in the full version of the paper [3] we construct an algorithm that solves (n, n) -GapDL[$\text{GGen}_{\text{gg}}, \text{high}$] in q group operations with success probability $(q^2/4np)^n$. Thus, for large p the fastest generic adversary solving (n, n) -GapDL[$\text{GGen}_{\text{gg}}, \text{high}$] with overwhelming success probability requires $\sqrt{np/e} \leq q \leq 2\sqrt{np}$ group operations.

The proof of Theorem 3 follows a structure similar to Yun [30]. First we prove the equivalence of n -SHS $_d[p]$ and (n, n) - d -PolyDL[GGen $_{\text{gg}}$, high], and then we bound the success probability of an adversary against n -SHS $_d[p]$. The equivalence of the two problems corresponds to the lemma below.

Statement and proof closely follow [30, Theorem 1] while additionally handling Eval queries. The proof can be found the full version of the paper [3].

Lemma 2. *Let n, d be positive integers and p a prime number. Let GGen $_{\text{gg}}$ be a group-generating algorithm that generates generic groups of exactly size p . Then for every adversary \mathcal{A} against game (n, n) - d -PolyDL[GGen $_{\text{gg}}$, high] there exists an adversary \mathcal{B} against n -SHS $_d[p]$ such that*

$$\text{Adv}_p^{n\text{-shs}_d}(\mathcal{B}) \geq \text{Adv}_{\text{GGen}_{\text{gg}}, \text{high}}^{(n, n)\text{-}d\text{-polydl}}(\mathcal{A}) .$$

Moreover, if \mathcal{A} makes q group-operation queries and q_{Eval} queries to Eval, then \mathcal{B} makes at most $q_{\text{Eval}} + (n + q)(n + q + 1)/2$ queries to Eval.

We start working on n -SHS $_d[p]$ with the next lemma. Here we express that, up to a loss of d^n , an adversary against n -SHS $_d[p]$ does not need more than n hypersurface queries which return 1 to identify a solution.

Importantly, observe how we limit the resources of an adversary against n -SHS $_d[p]$ exclusively in terms of its queries to Eval. Our adversaries are otherwise unbounded. For this reason, the following reduction does not consider the computational resources needed by the adversary to perform its operations. The proof is in the full version of the paper [3]

Lemma 3. *Let n, d be positive integers and p a prime number. For every adversary \mathcal{A} against n -SHS $_d[p]$ that makes at most q queries to Eval there exists an adversary \mathcal{B} against n -SHS $_d[p]$ that makes at most q queries to Eval such that at most n of them return 1 and*

$$\text{Adv}_p^{n\text{-shs}_d}(\mathcal{B}) \geq d^{-n} \text{Adv}_p^{n\text{-shs}_d}(\mathcal{A}) .$$

PROOF IDEA. Intuition for the proof is simple for the case $n = 1$: All queries of \mathcal{A} to SimEval are forwarded to Eval. The first time Eval(g) returns 1, we know that the secret \mathbf{a} must be a zero of g . Since g has degree at most d , there can be at most d distinct zeroes. The reduction guesses which zero is the correct one (this is the reduction loss) and then simulates the remaining queries of \mathcal{A} to SimEval accordingly. The proof is similar for $n > 1$. We know that, in general, n polynomials in $\mathbb{Z}_p[X_1, \dots, X_n]$ of degree d have at most d^n zeroes in common, one of which the reduction can use to simulate remaining queries to SimEval. However, the n queried polynomials must be in general position: For example, the zeroes of $x_1 + x_2$ are the same as those of $2x_1 + 2x_2$, and querying both polynomials would not help the reduction. To resolve this issue, the reduction keeps a set Z of common zeroes to all polynomials seen so far which, when forwarded to Eval, make the oracle return 1 (i.e., polynomials which vanish on \mathbf{a}).

This set has a rich structure: In fact, the study of zero sets of polynomial is the raison d'être of the field of algebraic geometry. If the polynomial g queried by \mathcal{A} carries no new information (i.e., $g(Z) = \{0\}$) then the simulated oracle returns 1 without forwarding. Otherwise, the polynomial is forwarded. If the answer is 1, then the reduction updates the set Z and then guesses which one of its irreducible components contains \mathbf{a} , which becomes the updated Z . The identification of irreducible components is made possible by the underlying structure of the set Z . Selecting an irreducible component guarantees that, on a following evaluation query, intersecting the now irreducible Z with another hypersurface not containing Z brings down the dimension of Z by 1. Since the dimension of \mathbb{Z}_p^n is n , we can have at most n such queries. With a careful choice of the guessing probability of each irreducible component, Bézout's theorem ensures that the probability of always making the right guess is again d^{-n} . \square

Remark 1. The bound on the advantage against (n, n) - d -PolyDL[GGen_{gg}, high] of Theorem 3 extends to (m, n) - d -PolyDL[GGen_{gg}, high], for $m \leq n$. This is done by a simple tight reduction between problems (m, n) - d -PolyDL[GGen_{gg}, high] and (m, m) - d -PolyDL[GGen_{gg}, high]. The reduction extends the one for standard multi-instance discrete logarithm [29, Section 3] by also simulating oracle Eval: It simply forwards every query to its own oracle.

6 Generic Hardness of the Multi-Instance Gap Computational Diffie-Hellman Problem

In this section we derive lower bounds on the hardness of the m -out-of- n gap computational Diffie-Hellman problem in the generic group model for different granularities. We further argue that all derived bounds are optimal. Section 6.1 covers high, Section 6.2 medium, and Section 6.3 low granularity.

6.1 Generic Hardness of High-Granularity (m, n) -GapCDH

We work in the algebraic group model to show that the generic lower bound on the hardness of high-granularity (m, m) -GapDL carries over to high-granularity (m, n) -GapCDH. Concretely, in Theorem 4 we provide a generic reduction from (m, n) -GapCDH[GGen, high] to (m, m) -GapDL[GGen, high]. Then, an application of Corollary 1 establishes the desired bound on (m, n) -GapCDH.

In this section we work with high-granularity problems, in which the group description $\mathcal{G} = (\mathbb{G}, p, g)$ is shared by all instances. For ease of notation, we treat \mathcal{G} as an implicit parameter of the system until the end of this section.

The generic reduction from (m, n) -GapCDH to (m, m) -GapDL in the high-granularity setting is sketched below. The full proof can be found in the full version of the paper [3]

Theorem 4. *Let GGen be a group-generating algorithm that generates groups of at least size p , and let m, n be two positive integers such that $m \leq n \leq p$. Then*

there exists a generic reduction that constructs from any algebraic adversary \mathcal{A} against game $G_{\text{Gen,high}}^{(m,n)\text{-gcdh}}$ an algebraic adversary \mathcal{B} against $G_{\text{Gen,high}}^{(m,m)\text{-gdl}}$ such that

$$\text{Adv}_{\text{Gen,high}}^{(m,m)\text{-gdl}}(\mathcal{B}) \geq 2^{-m} \text{Adv}_{\text{Gen,high}}^{(m,n)\text{-gcdh}}(\mathcal{A}) .$$

Moreover, \mathcal{B} makes at most $2n(m+2)(\log p + 1)$ group operations in addition to those made by \mathcal{A} , and the same amount of queries to DDH.

Despite the seemingly sizeable reduction loss of 2^m , we argue that the factor is small in the context of the final security bounds. In fact, as seen in Section 5, the advantage in breaking (m, m) -GapDL decreases exponentially with m . This renders the exponential contribution of the factor 2^m irrelevant, as the following concrete bound on the hardness of (m, n) -GapCDH[$\text{GGen}_{\text{gg,high}}$] shows. Its proof can be found in the full version of the paper [3].

Corollary 2. *Let GGen_{gg} be a group-generating algorithm that generates groups of at least size p , and let m, n be two positive integers such that $m \leq n \leq p$. Then for every generic adversary \mathcal{A} against (m, n) -GapCDH[$\text{GGen}_{\text{gg,high}}$] that makes at most q queries to the group-operation oracle and q_{DDH} queries to the gap oracle:*

$$\text{Adv}_{\text{GGen}_{\text{gg,high}}}^{(m,n)\text{-gcdh}}(\mathcal{A}) \leq \left(\frac{2e(q + 12mn \log p)^2 + 4eq_{\text{DDH}}}{mp} \right)^m \approx \left(\frac{q^2}{mp} \right)^m .$$

Similarly to the bound for computing discrete logarithms, this result is optimal. Namely, problem (m, n) -GapCDH[$\text{GGen}_{\text{gg,high}}$] can be solved computing q group operations with success probability $(q^2/4mp)^m$ by using the generic adversary against high-granularity DL provided in the full version [3]. Thus, for large p the fastest generic adversary solving (m, n) -GapCDH[$\text{GGen}_{\text{gg,high}}$] with overwhelming success probability requires $\sqrt{mp/2e} \leq q \leq 2\sqrt{mp}$ group operations.

PROOF IDEA OF THEOREM 4. This proof extends the following simple single-instance reduction \mathcal{B} , in turn built from two reductions \mathcal{B}_0 and $\mathcal{B}_{\{1\}}$. The reductions build upon a CDH adversary \mathcal{A} . Adversary \mathcal{A} receives $X = g^x$ and $Y = g^y$, and is tasked with computing $W = g^{xy}$. In the algebraic group model, \mathcal{A} must return a representation of the output as a combination of its input, i.e., some elements $a, b, c \in \mathbb{Z}_p$ such that $W = X^a Y^b g^c$. Rewriting this expression in the exponents, we obtain that, if \mathcal{A} wins,

$$xy = ax + by + c .$$

Given a DL challenge $Z = g^z$, reduction \mathcal{B}_0 embeds the challenge as $X = Z$ and generates $Y = g^y$ by picking a random y . Then, \mathcal{B}_0 can compute the DL as $z = x = (y - a)^{-1}(by + c)$. However, $y - a$ might not be invertible. In this case, adversary $\mathcal{B}_{\{1\}}$ would be successful: It embeds the challenge as $Y = Z$ and returns a , which is a correct solution if $y - a$ is not invertible. Reduction \mathcal{B} picks

one of the two subsets $I \subseteq \{1\}$ at random and runs \mathcal{B}_I . If the CDH adversary is successful, then \mathcal{B} has at least probability $1/2$ of succeeding.

Case $n = m > 1$ is approached as follows. Again the reduction \mathcal{B} is composed of components \mathcal{B}_I , where $I \subseteq [1..n]$. The DL challenge $\mathbf{Z}[i]$ is distributed as either $\mathbf{X}[i]$ or $\mathbf{Y}[i]$ according to whether $i \in I$, and all remaining values are picked by the reduction. The CDH adversary—if successful—returns square matrices A, B and vector \mathbf{c} such that $\text{diag}(\mathbf{y})\mathbf{x} = A\mathbf{x} + B\mathbf{y} + \mathbf{c}$, where $\text{diag}(\mathbf{y})$ is the diagonal matrix with the elements of \mathbf{y} on the diagonal. Rearranging, we obtain

$$(\text{diag}(\mathbf{y}) - A)\mathbf{x} = B\mathbf{y} + \mathbf{c} .$$

Our goal is to iteratively decrease the dimension of this matrix equation. If $n \notin I$ adversary \mathcal{B}_I expresses $\mathbf{x}[n]$ in terms of $\mathbf{x}[1..n-1]$. On the other hand, if $n \in I$ then it computes $\mathbf{y}[n]$. Whether this computation is correct depends on whether I is the right choice for A, B , and \mathbf{c} . More explicitly, from the last row of the previous matrix equation we get the expression

$$\begin{aligned} \mathbf{x}[n](\mathbf{y}[n] - A_{nn}) = & (A_{n1}, \dots, A_{n(n-1)})\mathbf{x}[1..n-1] + \\ & + (B_{n1}, \dots, B_{n(n-1)})\mathbf{y}[1..n-1] + B_{nn}\mathbf{y}[n] + \mathbf{c}[n] . \end{aligned}$$

If the number $\mathbf{y}[n] - A_{nn}$ is not invertible (case $n \in I$), then adversary \mathcal{B}_I can set $\mathbf{y}[n] = A_{nn}$. In the other case (case $n \notin I$) the adversary can replace the expression for $\mathbf{x}[n]$ into the remaining $n-1$ rows of the matrix. In this case, $\mathbf{y}[n]$ is known, and calling $\mathbf{x}' = (\mathbf{x}[1], \dots, \mathbf{x}[n-1])$, $\mathbf{y}' = (\mathbf{y}[1], \dots, \mathbf{y}[n-1])$, we have recovered again a matrix equation of the form

$$\text{diag}(\mathbf{y}')\mathbf{x}' = A'\mathbf{x}' + B'\mathbf{y}' + \mathbf{c}'$$

of decreased dimension $n-1$. Repeating this argument, we arrive at an equation of dimension 1. At this point all elements of \mathbf{y} are known to \mathcal{B}_I , which is then able to recover the elements of \mathbf{x} .

Note that there always exists, for every possible A, B , and \mathbf{c} , a set I for which the above procedure is successful, i.e., a set I such that, for every $i \in [1..n]$, the expression $i \in I$ is satisfied exactly if $\mathbf{y}[i] = (A_{(i)})_{ii}$, where $A_{(i)}$ is the i -th update of matrix A . Since adversary \mathcal{B} picks $I \subseteq [1..n]$ at random and runs \mathcal{B}_I , the reduction loses a factor of 2^n .

The case $n \neq m$ adds more complexity to the proof. The reduction first expands the m DL challenges $\hat{\mathbf{Z}}$ to a vector $\mathbf{Z} = \hat{\mathbf{Z}}^V$ (plus some rerandomization) of length n . Here V is a $n \times m$ matrix for which each $m \times m$ submatrix is invertible.⁸ This has two important consequences: Firstly, we can express any element of \mathbf{Z} as a combination of any other fixed m elements of \mathbf{Z} . Secondly, retrieving any m DLs of \mathbf{Z} allows the reduction to compute the DLs of the original $\hat{\mathbf{Z}}$. This has, however, an unintended side effect: We can still obtain an equation of the form $\text{diag}(\mathbf{y})\mathbf{x} = A\mathbf{x} + B\mathbf{y} + \mathbf{c}$, where all terms are of size m (this is the role, in the reduction code, of the function `reduceMatrices`), but

⁸ This expansion technique is originally from the work of Ying and Kunihiro [29].

now A, B, \mathbf{c} depend on the distribution of the challenges to \mathbf{X} and \mathbf{Y} , that is, on the set I . This means that the reduction cannot simply compute the element $\mathbf{y}[i]$ as A_{ii} at each step. It has to answer the question: “Assuming the reduction was not trying to compute $\mathbf{y}[m]$, what would be the value for $\mathbf{y}[m]$ which would make it unable to compute $\mathbf{x}[m]$?” (In the reduction code, the answer is yielded by the function `computeDlog`.)

In the proof, the gap oracle of \mathcal{A} is simply simulated by forwarding all queries to DDH. \square

Remark 2. Note that using Corollary 2 with $q_{\text{DDH}} = 0$ yields a generic lower bound on the hardness of the “standard” multi-instance CDH problem.

Further, oracle DDH plays a modest role in the proof of Theorem 4. One could define a “polycheck CDH” problem in the same fashion as it is done for discrete logarithm in Section 5 (in short, (m, n) - d -PolyCDH). It is then immediate to extend Theorem 4 to show the equivalence of games (m, n) - d -PolyCDH[GGen, high] and (m, n) - d -PolyDL[GGen, high] in the algebraic group model with the same loss of 2^m . Hence, with an additional multiplicative factor of $(d/2)^m$ the advantage of any adversary against game (m, n) - d -PolyCDH[GGen_{gg}, high] can be bounded as in Corollary 2.

6.2 Generic Hardness of Medium-Granularity (m, n) -GapCDH

We present an explicit bound on the concrete security of m -out-of- n gap computational Diffie-Hellman in the generic group model in the medium-granularity setting. The main result of this section is similar to that in Section 6.1. The bound follows from observing that we can simulate the medium-granularity game starting from the high-granularity one. Then, we can apply Corollary 2 after counting the additional group queries by the simulation. For more details, we refer to the full version of the paper [3].

Corollary 3. *Let GGen_{gg} be a group-generating algorithm that generates generic groups of at least size p , and let m, n be two positive integers such that $m \leq n \leq p$. Then for every generic adversary \mathcal{A} against (m, n) -GapCDH[GGen_{gg}, med] that makes at most q queries to the group-operation oracle and q_{DDH} queries to oracle DDH:*

$$\text{Adv}_{\text{GGen}_{\text{gg}}, \text{med}}^{(m, n)\text{-gcdh}}(\mathcal{A}) \leq \left(\frac{2e(q + 6(q_{\text{DDH}} + 5mn) \log p)^2}{mp} \right)^m \approx \left(\frac{q^2}{mp} \right)^m.$$

Similarly to the previous concrete bounds, this result is optimal, namely there exists a generic adversary against (m, n) -GapCDH[GGen_{gg}, med] which needs $2\sqrt{2mp}$ group operations and achieves success probability 1. In fact, we can build an adversary against (m, n) -GapCDH[GGen_{gg}, med] starting from an adversary against $(2m, 2m)$ -DL[GGen_{gg}, high] that requires about the same amount of oracle queries. Summing up, we obtain that for large p the fastest generic adversary achieving overwhelming success probability in game (m, n) -GapCDH[GGen_{gg}, med] requires $\sqrt{mp}/(2e) \leq q \leq 2\sqrt{2mp}$ group operations.

6.3 Generic Hardness of Low-Granularity (m, n) -GapCDH

In this section we present an explicit bound on the concrete security of m -out-of- n gap computational Diffie-Hellman in the generic group model in the low-granularity setting. The bound is stated in the following theorem and is computed directly. The proof can be found in the full version of the paper [3].

Theorem 5. *Let GGen_{gg} be a group-generating algorithm that generates generic groups of at least size p , and let m, n, q, q_{DDH} and $q_i, i \in [1..n]$, be integers such that $1 \leq m \leq n, q = q_1 + \dots + q_n$, and q_i is large ($q_i \geq 60 \log p$ and $4q_i^2 \geq q_{\text{DDH}}$). Then for every generic adversary \mathcal{A} against the low-granularity m -out-of- n multi-instance computational Diffie-Hellman problem that makes at most q_i queries to the i -th group-operation oracle and q_{DDH} queries to the gap oracle:*

$$\text{Adv}_{\text{GGen}_{\text{gg}}, \text{low}}^{(m, n)\text{-gcdh}}(\mathcal{A}) \leq \left(\frac{4eq^2}{m^2p} \right)^m .$$

Since the number of group operations performed by a (m, n) -GapCDH adversary is typically large, we reckon the requirements $q_i \geq 60 \log p$ and $4q_i^2 \geq q_{\text{DDH}}$ are rather mild.

We argue that this result is optimal. In fact, each of the first m instances can be solved in time q/m with success probability $(q/m)^2/4p$ using the algorithm provided in the full version of the paper [3]. Thus, (m, n) -GapCDH[$\text{GGen}_{\text{gg}}, \text{low}$] can be solved in time q by independently running the single-instance adversary on the first m instances which results in a success probability of $(q^2/4m^2p)^m$. Further, for large p the fastest generic adversary achieving overwhelming success probability in game (m, n) -GapCDH[$\text{GGen}_{\text{gg}}, \text{low}$] requires $m\sqrt{p/8e} \leq q \leq 2m\sqrt{p}$ group operations.

Acknowledgments

The authors are grateful to Masayuki Abe, Razvan Barbulescu, Mihir Bellare, Dan Boneh, Nadia Heninger, Tanja Lange, Alexander May, Bertram Poettering, Maximilian Rath, Sven Schäge, Nicola Turchi, and Takashi Yamakawa for their helpful comments. Benedikt Auerbach was supported by the European Research Council, ERC consolidator grant (682815-TOCNeT), and conducted part of this work at Ruhr University Bochum, supported by the ERC Project ERCC (FP7/615074) and the NRW Research Training Group SecHuman. Federico Giaccon conducted part of this work at Ruhr University Bochum, supported by the ERC Project ERCC (FP7/615074). Eike Kiltz was supported by the ERC Project ERCC (FP7/615074), DFG SPP 1736 Big Data, and the DFG Cluster of Excellence 2092 CASA.

References

1. Abdalla, M., Bellare, M., Rogaway, P.: The oracle Diffie-Hellman assumptions and an analysis of DHIES. In: Naccache, D. (ed.) CT-RSA 2001. LNCS, vol. 2020, pp. 143–158. Springer, Heidelberg (Apr 2001)

2. Adrian, D., Bhargavan, K., Durumeric, Z., Gaudry, P., Green, M., Halderman, J.A., Heninger, N., Springall, D., Thomé, E., Valenta, L., VanderSloot, B., Wustrow, E., Zanella-Béguelin, S., Zimmermann, P.: Imperfect forward secrecy: How Diffie-Hellman fails in practice. In: Ray, I., Li, N., Kruegel, C. (eds.) ACM CCS 2015. pp. 5–17. ACM Press (Oct 2015)
3. Auerbach, B., Giacon, F., Kiltz, E.: Everybody’s a target: Scalability in public-key encryption. Cryptology ePrint Archive, Report 2019/364 (2019), <https://eprint.iacr.org/2019/364>
4. Barbulescu, R.: Algorithms for discrete logarithm in finite fields. Ph.D. thesis, University of Lorraine, Nancy, France (2013)
5. Barbulescu, R., Pierrot, C.: The multiple number field sieve for medium- and high-characteristic finite fields. LMS Journal of Computation and Mathematics 17(A), 230–246 (2014)
6. Bartusek, J., Ma, F., Zhandry, M.: The distinction between fixed and random generators in group-based assumptions. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part II. LNCS, vol. 11693, pp. 801–830. Springer, Heidelberg (Aug 2019)
7. Bellare, M., Boldyreva, A., Micali, S.: Public-key encryption in a multi-user setting: Security proofs and improvements. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 259–274. Springer, Heidelberg (May 2000)
8. Bellare, M., Ristenpart, T., Tessaro, S.: Multi-instance security and its application to password-based cryptography. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 312–329. Springer, Heidelberg (Aug 2012)
9. Bellare, M., Rogaway, P.: Optimal asymmetric encryption. In: Santis, A.D. (ed.) EUROCRYPT’94. LNCS, vol. 950, pp. 92–111. Springer, Heidelberg (May 1995)
10. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 409–426. Springer, Heidelberg (May / Jun 2006)
11. Bernstein, D.J., Lange, T.: Batch NFS. In: Joux, A., Youssef, A.M. (eds.) SAC 2014. LNCS, vol. 8781, pp. 38–58. Springer, Heidelberg (Aug 2014)
12. Boneh, D., Boyen, X., Goh, E.J.: Hierarchical identity based encryption with constant size ciphertext. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 440–456. Springer, Heidelberg (May 2005)
13. Coretti, S., Dodis, Y., Guo, S.: Non-uniform bounds in the random-permutation, ideal-cipher, and generic-group models. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part I. LNCS, vol. 10991, pp. 693–721. Springer, Heidelberg (Aug 2018)
14. Corrigan-Gibbs, H., Kogan, D.: The discrete-logarithm problem with preprocessing. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part II. LNCS, vol. 10821, pp. 415–447. Springer, Heidelberg (Apr / May 2018)
15. Fouque, P.A., Joux, A., Mavromati, C.: Multi-user collisions: Applications to discrete logarithm, Even-Mansour and PRINCE. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014, Part I. LNCS, vol. 8873, pp. 420–438. Springer, Heidelberg (Dec 2014)
16. Fuchsbauer, G., Kiltz, E., Loss, J.: The algebraic group model and its applications. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part II. LNCS, vol. 10992, pp. 33–62. Springer, Heidelberg (Aug 2018)
17. Garay, J.A., Johnson, D.S., Kiayias, A., Yung, M.: Resource-based corruptions and the combinatorics of hidden diversity. In: Kleinberg, R.D. (ed.) ITCS 2013. pp. 415–428. ACM (Jan 2013)

18. Guillevic, A., Morain, F.: Discrete logarithms. In: Mrabet, N.E., Joye, M. (eds.) Guide to pairing-based cryptography. CRC Press - Taylor and Francis Group (Dec 2016)
19. Heninger, N., Durumeric, Z., Wustrow, E., Halderman, J.A.: Mining your Ps and Qs: Detection of widespread weak keys in network devices. In: 21st USENIX Security Symposium (2012)
20. Hitchcock, Y., Montague, P., Carter, G., Dawson, E.: The efficiency of solving multiple discrete logarithm problems and the implications for the security of fixed elliptic curves. *International Journal of Information Security* 3(2), 86–98 (Nov 2004)
21. Hofheinz, D., Nguyen, N.K.: On tightly secure primitives in the multi-instance setting. In: Lin, D., Sako, K. (eds.) PKC 2019, Part I. LNCS, vol. 11442, pp. 581–611. Springer, Heidelberg (Apr 2019)
22. Kuhn, F., Struik, R.: Random walks revisited: Extensions of Pollard’s rho algorithm for computing multiple discrete logarithms. In: Vaudenay, S., Youssef, A.M. (eds.) SAC 2001. LNCS, vol. 2259, pp. 212–229. Springer, Heidelberg (Aug 2001)
23. Maurer, U.M.: Abstract models of computation in cryptography (invited paper). In: Smart, N.P. (ed.) 10th IMA International Conference on Cryptography and Coding. LNCS, vol. 3796, pp. 1–12. Springer, Heidelberg (Dec 2005)
24. Mizuide, T., Takayasu, A., Takagi, T.: Tight reductions for Diffie-Hellman variants in the algebraic group model. In: Matsui, M. (ed.) CT-RSA 2019. LNCS, vol. 11405, pp. 169–188. Springer, Heidelberg (Mar 2019)
25. Okamoto, T., Pointcheval, D.: The gap-problems: A new class of problems for the security of cryptographic schemes. In: Kim, K. (ed.) PKC 2001. LNCS, vol. 1992, pp. 104–118. Springer, Heidelberg (Feb 2001)
26. Rupp, A., Leander, G., Bangertner, E., Dent, A.W., Sadeghi, A.R.: Sufficient conditions for intractability over black-box groups: Generic lower bounds for generalized DL and DH problems. In: Pieprzyk, J. (ed.) ASIACRYPT 2008. LNCS, vol. 5350, pp. 489–505. Springer, Heidelberg (Dec 2008)
27. Sadeghi, A.R., Steiner, M.: Assumptions related to discrete logarithms: Why subtleties make a real difference. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 244–261. Springer, Heidelberg (May 2001)
28. Shoup, V.: Lower bounds for discrete logarithms and related problems. In: Fumy, W. (ed.) EUROCRYPT’97. LNCS, vol. 1233, pp. 256–266. Springer, Heidelberg (May 1997)
29. Ying, J.H.M., Kunihiro, N.: Bounds in various generalized settings of the discrete logarithm problem. In: Gollmann, D., Miyaji, A., Kikuchi, H. (eds.) ACNS 17. LNCS, vol. 10355, pp. 498–517. Springer, Heidelberg (Jul 2017)
30. Yun, A.: Generic hardness of the multiple discrete logarithm problem. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part II. LNCS, vol. 9057, pp. 817–836. Springer, Heidelberg (Apr 2015)