



An Optimal Distributed Discrete Log Protocol with Applications to Homomorphic Secret Sharing

Itai Dinur¹(✉), Nathan Keller², and Ohad Klein²

¹ Department of Computer Science, Ben-Gurion University, Beersheba, Israel
dinuri@cs.bgu.ac.il

² Department of Mathematics, Bar-Ilan University, Ramat Gan, Israel

Abstract. The distributed discrete logarithm (DDL) problem was introduced by Boyle et al. at CRYPTO 2016. A protocol solving this problem was the main tool used in the share conversion procedure of their homomorphic secret sharing (HSS) scheme which allows non-interactive evaluation of branching programs among two parties over shares of secret inputs.

Let g be a generator of a multiplicative group \mathbb{G} . Given a random group element g^x and an unknown integer $b \in [-M, M]$ for a small M , two parties A and B (that cannot communicate) successfully solve DDL if $A(g^x) - B(g^{x+b}) = b$. Otherwise, the parties err. In the DDL protocol of Boyle et al., A and B run in time T and have error probability that is roughly linear in M/T . Since it has a significant impact on the HSS scheme's performance, a major open problem raised by Boyle et al. was to reduce the error probability as a function of T .

In this paper we devise a new DDL protocol that substantially reduces the error probability to $O(M \cdot T^{-2})$. Our new protocol improves the asymptotic evaluation time complexity of the HSS scheme by Boyle et al. on branching programs of size S from $O(S^2)$ to $O(S^{3/2})$. We further show that our protocol is optimal up to a constant factor for all relevant cryptographic group families, unless one can solve the discrete logarithm problem in a *short* interval of length R in time $o(\sqrt{R})$.

Our DDL protocol is based on a new type of random walk that is composed of several iterations in which the expected step length gradually increases. We believe that this random walk is of independent interest and will find additional applications.

Keywords: Homomorphic secret sharing · Share conversion
Fully homomorphic encryption · Discrete logarithm
Discrete logarithm in a short interval · Random walk

1 Introduction

Homomorphic Secret Sharing. Homomorphic secret sharing (HSS) is a practical alternative approach to fully homomorphic encryption (FHE) [12, 17] that

provides some of its functionalities. It was introduced by Boyle et al. [5] at CRYPTO 2016 and further studied and extended in [4, 6, 7, 10]. The main advantage of HSS over traditional secure multiparty computation protocols [1, 8, 20] is that, similarly to FHE, its communication complexity is smaller than the circuit size of the computed function.

HSS allows homomorphic evaluation to be distributed among two parties who do not interact with each other. A (2-party) HSS scheme randomly splits an input w into a pair of shares (w_0, w_1) such that: (1) each share w_i computationally hides w , and (2) there exists a polynomial-time local evaluation algorithm Eval such that for any program P from a given class (e.g., a boolean circuit or a branching program), the output $P(w)$ can be efficiently reconstructed from $\text{Eval}(w_0, P)$ and $\text{Eval}(w_1, P)$.

The main result of [5] is an HSS scheme for branching programs under the Decisional Diffie-Hellman (DDH) assumption that satisfies $P(w) = \text{Eval}(w_0, P) + \text{Eval}(w_1, P)$. It was later optimized in [4, 6], where the security of the optimized variants relies on other discrete log style assumptions.

Let \mathbb{G} be a multiplicative cyclic group of prime order N in which the discrete log problem is (presumably) hard and let g be a generator of this group. The scheme of [5] allows the parties to locally multiply an encrypted (small) input $w \in \mathbb{Z}$ with an additively secret-shared (small) value $y \in \mathbb{Z}$, such that the result $z = wy$ is shared between the parties. The problem is that at this stage g^z is multiplicatively shared by the parties, so they cannot multiply z with a new encrypted input w' . Perhaps the most innovative idea of [5] allows the parties to convert multiplicative shares of g^z into additive shares of z without any interaction via a share conversion procedure. Once the parties have an additive sharing of z , they can proceed to add it to other additive shares. These operations allow to evaluate restricted multiplication straight-line (RMS) programs which can emulate any branching program of size S using $O(S)$ instructions.

The share conversion procedure of [5] is not perfect in the sense that the parties may err. More specifically, the parties fail to compute correct additive shares of z with some error probability δ that depends on the running time T of the parties and on a small integer M that bounds the intermediate computation values. As share conversion is performed numerous times during the execution of Eval , its total error probability accumulates and becomes roughly $\delta \cdot S$, where S is the number of multiplications performed by the RMS program P . Thus, for the total error probability to be constant one has to set the running time T of the parties in the share conversion procedure such that $\delta \approx 1/S$. Consequently, the running time to error tradeoff has a significant impact on the performance of the HSS scheme.

Since the main motivation behind HSS is to provide a practical alternative to FHE, one of the main open problems posed in [5] was to improve the running time to error tradeoff of the share conversion procedure. Progress on this open problem was made in the followup works [4, 6] which significantly improved the practicality of the HSS scheme. Despite this progress, the asymptotic running time to error tradeoff of the share conversion procedure was not substantially improved and the running time T in all the schemes grows (roughly) linearly

with the inverse error probability $1/\delta$ (or as M/δ in general). Thus, to obtain $\delta \approx 1/S$, one has to set $T \approx S$, and since the total number of multiplications in P is S , the total running time becomes $O(S^2)$.

The Distributed Discrete Log Problem. The focus of this paper is on the “distributed discrete log” (DDL) problem which the parties collectively solve in the share conversion procedure. We now describe the DDL problem and abstract away the HSS details for simplicity. The DDL problem involves two parties A and B . The input of A consists of a group element g^x , where x is chosen uniformly at random from \mathbb{Z}_N . The input of B consists of g^{x+b} , where $b \in [-M, M]$ is an unknown uniformly chosen integer in the interval (for a small fixed integer parameter M). The algorithms A, B are restricted by a parameter T which bounds the number of group operations they are allowed to compute.¹ After executing its algorithm, each party outputs an integer. The parties successfully solve the DDL instance if $A(g^x) - B(g^{x+b}) = b$. We stress that A and B are not allowed to communicate.²

If g^z is multiplicatively shared by A (party 0) and B (party 1), then $g^{z_0} \cdot g^{z_1} = g^z$. In the share conversion procedure party A runs $A(g^{-z_0})$ while party B runs $B(g^{z_1})$. Assuming they correctly solve DDL for $|z| \leq M$, we have $A(g^{-z_0}) - B(g^{z_1}) = z_1 + z_0 = z$, namely, $A(g^{-z_0})$ and $-B(g^{z_1})$ are additive shares of z as required.

It is convenient to view the DDL problem as a synchronization problem: A and B try to agree or *synchronize* on a group element with a known offset in the exponent from their input. If they manage to do so, the parties solve DDL by outputting this offset. For example, if both parties synchronize on g^y , then $A(g^x) = y - x$ while $B(g^{x+b}) = y - (x + b)$, so $A(g^x) - B(g^{x+b}) = b$ as required. In particular, if both A and B can solve the discrete logarithm problem for their input (i.e., compute x and $x + b$, respectively), then they can synchronize on the generator g by outputting $A(g^x) = 1 - x$ and $B(g^{x+b}) = 1 - (x + b)$. Of course, this would violate the security of the HSS scheme, implying that the discrete logarithm problem in \mathbb{G} should be hard and A, B have to find other means to succeed.

Our Goals. The goal of this paper is to devise algorithms for A and B (i.e., a DDL protocol) that maximize their success probability (taken over the randomness of x, b), or equivalently, minimize their error probability δ given T . Our

¹ In all algorithms presented in this paper, the bulk of computation involves performing group operations, hence this is a reasonable complexity measure. Alternatively, the parameter T may bound the complexity of A, B in some reasonable computational model.

² We note that in the applications of [4–6], the distribution of $b \in [-M, M]$ is arbitrary. However (as we show in Lemma 13), our choice to define and analyze DDL for the uniform distribution of $b \in [-M, M]$ is technically justified since the uniform distribution is the hardest for DDL: algorithms for A, B that solve DDL with an error probability δ for the uniform distribution, also solve DDL with an error probability $O(\delta)$ for any distribution of $b \in [-M, M]$.

point of reference is the DDL protocol of [6] (which is a refined version of the original DDL protocol [5]) that achieves a linear tradeoff between the parameter T and error probability δ . More precisely, given that A, B are allowed T group operations, the DDL error probability is roughly M/T . In fact, there are several closely related protocols devised in [4–6] which give similar linear tradeoffs between the parameter T and error probability δ .

Yet another goal of this paper is to better understand the limitations of DDL protocols. More specifically, we aim to prove lower bounds on the error probability of DDL protocols by reducing a well-studied computational problem on groups to DDL. In particular, we are interested in the *discrete log in an interval* (DLI) problem, where the input consists of a group element in a known interval of length R and the goal is to compute its discrete log.

DLI has been the subject of intensive study in cryptanalysis and the best known algorithms for it are adaptations of the classical baby-step giant-step algorithm and the memory-efficient variant of Pollard [15] (see [11, 16] for additional extensions). These algorithms are based on collision finding and have complexity of about \sqrt{R} . They are the best known in concrete prime-order group families (in which discrete log is hard) up to large values of the interval R . In particular, for elliptic curve groups, the best known DLI algorithm has complexity of about \sqrt{R} where R is as large as the size of the group N (which gives the standard discrete logarithm problem). For some other groups (such as prime order subgroups of \mathbb{Z}_p^*), the best known complexity is about \sqrt{R} , where R can be up to subexponential in $\log N$ (as discrete log can be solved in subexponential complexity in these groups [13, 14]). We note that besides its relevance in cryptanalysis, DLI is solved as part of the decryption process of some cryptosystems (notably in the cryptosystem by Boneh, Goh and Nissim [3]).

An alternative approach to establishing error probability lower bounds for DDL is to use the generic group model (GGM), introduced by Shoup [18]. In GGM, an algorithm is not allowed direct access to the bit representation of the group elements, but can only obtain randomized encodings of the elements, available via oracle queries. The generic group model is a standard model for proving computational lower bounds on certain (presumably hard) problems on groups and thus establishing confidence in their hardness. Although the bounds obtained in GGM are relevant to a restricted class of algorithms, it is essentially the only model in which meaningful lower bounds are known for some computational problems on groups (such as discrete log). Moreover, for several problems (such as discrete log computation in some elliptic curve groups), generic algorithms are essentially the best algorithms known. The downside of this alternative proof approach is that it does not directly relate DDL to any hard problem in a group family, but rather establishes a lower bound proof in an abstract model.

Our Contribution. The main result of this work is closing the gap for DDL in many concrete group families by presenting upper and lower bounds that are tight (within a constant factor) based on the hardness of DLI in these families.

We first develop an improved DDL protocol that is applicable in any group \mathbb{G} and achieves a quadratic tradeoff between the parameter T and the error probability, namely $\delta = O(M/T^2)$. This is a substantial improvement over the linear tradeoff $\delta = O(M/T)$ obtained in [4–6]. Therefore, when executing Eval on an RMS program P with multiplicative complexity S , one can set $T = O(S^{1/2})$ to obtain $\delta = O(1/S)$ and the total running time is reduced from $O(S^2)$ in [4–6] to $O(S^{3/2})$. This result directly improves upon the computational complexity of some of the HSS applications given in [4–6]. For example, in private information retrieval [9] (PIR), a client privately searches a database distributed among several servers for the existence of a document satisfying a predicate P . The 1-round 2-server PIR scheme of [5] supports general searches expressed as branching programs of size S applied to each document. The computational complexity per document in the scheme of Boyle et al. is $O(S^2)$ and our result reduces this complexity to $O(S^{3/2})$.

On the practical side, we fully verified our protocol by extensive experiments. We hope that it will render HSS practical for new applications.

Our DDL protocol uses a new type of (pseudo) random walk composed of several iterations. Each one of these iterations resembles Pollard’s “kangaroo” random walk algorithm for solving DLI using limited memory [15]. However DDL is different from DLI as the parties cannot communicate and seek to minimize their error probability (rather than make it constant). This leads to a more complex iterative algorithm, where the parties carefully distribute their time complexity T among several random walks iterations. These iterations use increasingly longer step lengths that gradually reduce the error probability towards $O(M/T^2)$.

The new random walk maximizes the probability that parties with close inputs agree (or synchronize) on a common output without communicating. We believe that this random walk is of independent interest and will find additional applications beyond homomorphic secret sharing schemes and cryptography in general.

After presenting our DDL protocol, we focus on lower bounds and show that any DDL protocol for a family of groups must have error probability of $\delta = \Omega(M/T^2)$, unless DLI (with interval of length R) can be solved in time $T' \approx T = o(\sqrt{R})$ in this family. This is currently not achievable for small (polynomial) T in standard cryptographic groups (for which the group-based HSS scheme is deemed to be secure).

Finally, we analyze DDL protocols in the generic group model. In this model, our DDL protocol is adaptive, as the oracle queries of A and B depend on the answers to their previous queries. This stands in contrast to the protocols of [4–6] in GGM, whose oracle queries are fixed in advance (or selected with high probability from a pre-fixed set of size $O(T)$). It is therefore natural to ask whether adaptivity is necessary to obtain optimal DDL protocols in GGM. Interestingly, we prove that the answer is positive. In fact, we show that the linear tradeoff obtained in [4–6] is essentially the best possible for non-adaptive DDL protocols in GGM.

Paper Organization. The rest of the paper is organized as follows. We describe preliminaries in Sect. 2 and present an overview of our new protocol and related work in Sect. 3. Our new DDL protocol is analyzed in Sect. 4. We prove lower bounds on the DDL error probability in concrete group families in Sect. 5 and finally prove lower bounds on non-adaptive algorithms in GGM in Sect. 6.

2 Preliminaries

In this section we describe the preliminaries required for this work. First we introduce notation that we use throughout the paper and then we present and analyze the DDL algorithm of [5], which will serve as a basis for our algorithms.

2.1 Notation for the Distributed Discrete Log Problem

Recall that the parties A and B successfully solve the DDL instance if $A(g^x) - B(g^{x+b}) = b$. To simplify our notation, we typically do not explicitly write the parameters \mathbb{G}, g, N, M, T in the description of A, B , although some of them will appear in the analysis. We are interested in the success (or error) probability of A and B , taken over the randomness of x, b (and possibly over the randomness of A, B). We denote by $\mathbf{err}(A, B, x, b, T)$ the error event $A(g^x) - B(g^{x+b}) \neq b$, and by $\Pr_{\mathbf{err}}(A, B, [M_1, M_2], T)$ its probability $\Pr_{x,b}[\mathbf{err}(A, B, x, b, T)]$, where $x \in \mathbb{Z}_N$ and $b \in [M_1, M_2]$ are uniform (typically, we are interested in $M_2 = -M_1 = M$). We also denote by $\mathbf{suc}(A, B, x, b, T)$ the complementary success event $A(g^x) - B(g^{x+b}) = b$.

When both parties perform the same algorithm A , we shorten the notation into $\mathbf{err}(A, x, b, T)$, $\Pr_{\mathbf{err}}(A, [M_1, M_2], T)$, and $\mathbf{suc}(A, x, b, T)$, respectively. If the parameters A, B, x, b, T are apparent from the context, we sometimes use \mathbf{err} and \mathbf{suc} instead of $\mathbf{err}(A, B, x, b, T)$ and $\mathbf{suc}(A, B, x, b, T)$, respectively. As mentioned above, A and B can be randomized algorithms and in this case the success (and error) probabilities are taken over their randomness as well. However, to simplify our notation we will typically not refer to this randomness explicitly.

We note that the DDL problem considered in [4–6] is slightly different, as A, B are allowed to perform up to T group operations in expectation. In this alternative definition, one can construct DDL protocols that are more efficient than ours by a small constant factor, while our lower bounds remain the same (again, up to a constant factor).

In the description and analysis of the DDL algorithms, we make frequent use of group elements of the form g^{x+j} . For sake of simplicity, we denote $g_j := g^{x+j}$. In addition, we usually assume $b \geq 0$, as otherwise we can simply exchange the names of the parties A and B when they use the same algorithm. Finally, we refer to a group operation whose output is h as a query to h .

2.2 The Basic DDL Algorithm

Let $\phi : \mathbb{G} \rightarrow [0, N - 1]$ be a pseudo-random function (PRF) that maps group elements to integers. Our protocols evaluate ϕ on $O(T)$ group elements for $T \ll N^{1/2}$. We assume throughout the analysis that ϕ behaves as a truly random permutation on the evaluated group elements, and in particular, we do not encounter collisions in ϕ (i.e., for arbitrary $h \neq h'$, $\phi(h) \neq \phi(h')$). Our probabilistic calculations are taken over the choice of ϕ , even though we do not indicate this explicitly for simplicity.³

We describe the min-based DDL algorithm of [5] in Algorithm 1 and refer to it as the basic DDL algorithm. The algorithm is executed by both A and B . When applied to $g_0 = g^x$, the algorithm scans the T values g_0, g_1, \dots, g_{T-1} and chooses the index i_{min} for which $\phi(g_i)$ is minimal. The output of the algorithm is $\text{Basic}_T(g^x) = (i_{min}, g_{min})$. Note that the algorithm depends also on \mathbb{G}, g ; however, we do not mention them explicitly in the notation. Furthermore, the output g_{min} will only be relevant later, when we use this algorithm as a sub-procedure. For the sake of analysis, we slightly abuse notation below and refer to i_{min} as the (only) output of $\text{Basic}_T(g^x)$.

The motivation behind the algorithm is apparent: if party A applies $\text{Basic}_T(g^x)$ and party B applies $\text{Basic}_T(g^{x+b})$, where $b \ll T$, then the lists of values scanned by the two algorithms (i.e., g_0, g_1, \dots, g_{T-1} and $g_b, g_{b+1}, \dots, g_{b+T-1}$) contain many common values, and thus, with a high probability the minimum is one of the common values, resulting in success of the algorithm.

2.3 Analysis of the Basic DDL Algorithm

Error Probability. The following lemma calculates the error probability of the basic DDL algorithm, as a function of $|b|$ and T .

Lemma 1. *The error probability of the basic DDL algorithm is*

$$\Pr_x[\mathbf{err}(\text{Basic}_T, x, b, T)] = \Pr[(\text{Basic}_T(g^x) - \text{Basic}_T(g^{x+b}) \neq b)] = 2|b|/(|b| + T).$$

Proof. We assume $b \geq 0$, as otherwise we exchange the names of A and B . Since both A and B use Algorithm 1, then A computes the function ϕ on g_0, g_1, \dots, g_{T-1} , while B computes this function on $g_b, g_{b+1}, \dots, g_{b+T-1}$. If the minimum value of ϕ for each party is obtained on an element $g_{min} = g^x \cdot g^{i_{min}}$ which is queried by both, then we have $\text{Basic}_T(g^x) = i_{min}$ and $\text{Basic}_T(g^{x+b}) = i_{min} - b$, implying that $\text{Basic}_T(g^x) - \text{Basic}_T(g^{x+b}) = b$ and the parties are successful. Similarly, they fail when the minimal value of ϕ on the elements

³ The function ϕ (and additional pseudo-random functions defined in this paper) can be implemented by a keyed MAC, where the key is pre-distributed to A and B . Thus, our probabilistic calculations should be formally taken over the choice of the key. They should include an error term that accounts for the distinguishing advantage of an efficient adversary (A or B in our case) that attempts to distinguish the PRF from a truly random permutation. However, for an appropriately chosen PRF, the distinguishing advantage is negligible and we ignore it for the sake of simplicity.

Algorithm 1. $\text{Basic}_T(g^x)$

```

1 begin
2    $h' \leftarrow g^x, i \leftarrow 0, \min \leftarrow \infty;$ 
3   while  $i < T$  do
4      $y \leftarrow \phi(h');$ 
5     if  $y < \min$  then
6        $g_{\min} \leftarrow h';$ 
7        $i_{\min} \leftarrow i, \min \leftarrow y;$ 
8     end
9      $h' \leftarrow h' \cdot g;$ 
10     $i \leftarrow i + 1;$ 
11  end
12  Output  $(i_{\min}, g_{\min});$ 
13 end

```

$g_0, g_1, \dots, g_{b+T-1}$ is obtained on an element computed only by one party, namely on one of the $2b$ elements $g^x \cdot g^i$ for $0 \leq i < b$ or $T \leq i < b+T$. Assuming that the output of ϕ on each element is uniform and the outputs are distinct, this occurs with probability $2b/(b+T)$. Hence $\Pr_x[\mathbf{err}(\text{Basic}_T, x, b, T)] = 2|b|/(|b|+T)$, as asserted. \blacksquare

The output difference in case of failure. An important quantity that plays a role in our improved protocol is the output difference of the parties in case they fail to synchronize on the same element g_{\min} (i.e., their output difference is not b). The following lemma calculates the expectation of this difference, as function of $|b|$ and T .

Lemma 2.

$$\mathbb{E} [|\text{Basic}_T(g^x) - \text{Basic}_T(g^{x+b}) - b| \mid \mathbf{err}] = (|b| + T)/2.$$

Proof. We assume $b \geq 0$. As written above, A computes the function ϕ on g_0, \dots, g_{T-1} , and B computes this function on g_b, \dots, g_{b+T-1} . The ordering of the values $\phi(g_0), \dots, \phi(g_{b+T-1})$ is uniform, and so the permutation π satisfying $\phi(g_{\pi(0)}) < \phi(g_{\pi(1)}) < \dots < \phi(g_{\pi(b+T-1)})$ is uniformly random in the permutation group of $\{0, 1, \dots, b+T-1\}$. The event \mathbf{err} is equivalent to the event $\pi(0) \notin [b, T-1]$. Without loss of generality let us restrict ourselves to the event $\pi(0) < b$, i.e. A encounters the minimal group element, and B does not; the other possibility $\pi(0) \geq T$ is symmetric with respect to reflection. Clearly, $\pi(0)$, which equals $\text{Basic}_T(g^x)$, is uniformly random in $[0, b-1]$. Moreover, $\text{Basic}_T(g^{x+b}) + b$ is $\pi(\min\{i \mid \pi(i) > b\})$ which uniformly distributes in $[b, b+T-1]$. Hence the expected final distance between the parties is $(2b+T+1)/2 - (b+1)/2 = (b+T)/2$. \blacksquare

3 Overview of Our New Protocol and Related Work

3.1 The New DDL Protocol

For the sake of simplicity, we assume in this overview that $M = 1$, hence $|b| \leq 1$. The starting point of our new DDL protocol is Algorithm 1. It makes T queries (i.e., group operations) and fails with probability of roughly $2/T$ according to Lemma 1. Let us assume that we run this algorithm with only $T/2$ queries, which increases the error probability by a factor of 2 to about $4/T$. On the other hand, we still have a budget of $T/2$ queries and we can exploit them to reduce the error probability. Interestingly, simply proceeding to calculate more consecutive group elements is not an optimal way to exploit the remaining budget.

After the first $T/2$ queries, we say that A (or B) is placed at group element g^y if $\phi(g^y)$ is the minimal value in its computed set of size $T/2$. Assume that A and B fail to synchronize on the same group element after the first $T/2$ queries (which occurs with probability of roughly $4/T$). Then, by Lemma 2, A and B are placed at elements which are at distance of about $T/4$, i.e., if A is placed at g^y and B is placed at g^z , then $|y - z| \approx T/4$. Our main idea is to use a somewhat different procedure in order to try to synchronize A and B in case they fail to do so after the first $T/2$ queries, while keeping A and B synchronized if they already are.

The next procedure employed by both A and B is a (pseudo) random walk starting from their initial position, whose step length is uniformly distributed in $[1, L - 1]$, where $L \approx \sqrt{T}$. The step length at group element g^y is determined by $\psi_{L-1}(g^y)$, where ψ_{L-1} is a pseudo-random function independent of ϕ that outputs a uniform integer in $[1, L - 1]$.⁴ Assume that after the first $T/2$ queries, B is placed at distance of about $T/4$ in front of A . Then A will pass B 's initial position after about $\sqrt{T}/2$ steps and simple probabilistic analysis shows that A will land on one of B 's steps after an additional expected number of about $\sqrt{T}/2$ steps. From this point, the walks coincide for the remaining $T/2 - \sqrt{T}$ steps. Similarly to Algorithm 1, each party outputs the offset of the minimal $\phi(g^y)$ value visited during its walk. Since both A and B use the same deterministic algorithm, they remain synchronized if they already are at the beginning of the walks. On the other hand, if they are not initially synchronized, their walks are expected to coincide on $T/2 - \sqrt{T}$ elements, and hence the probability that they remain unsynchronized is roughly $\sqrt{T}/(T/2) = 2 \cdot T^{-1/2}$. Thus, the error probability at this stage is about $4 \cdot T^{-1} \cdot 2 \cdot T^{-1/2} = 8 \cdot T^{-3/2}$, which already significantly improves upon the $2 \cdot T^{-1}$ error probability of Algorithm 1 for large T .

However, we can still do better. For the sake of simplicity, let us completely ignore constant factors in rest of this rough analysis. Note that we may reserve an additional number of $O(T)$ queries to be used in another random walk by shortening the first two random walks, without affecting the failure probability

⁴ Our analysis assumes that ψ_{L-1} is a truly random function and our probabilistic calculations are taken over the choice of ψ_{L-1} .

significantly. Hence, assume that the parties fail to synchronize after the random walk (which occurs with probability of about $T^{-3/2}$) and that we still have enough available queries for another random walk with $O(T)$ steps. Since each party covers a distance of about $T^{3/2}$ during its walk, then the expected distance between the parties in case of failure is roughly $T^{3/2}$. We can now perform another random walk with expected step length of $T^{3/4}$ (hence the walks are expected to coincide after about $T^{3/4}$ steps), reducing the error probability to about $T^{-3/2} \cdot (T^{3/4} \cdot T^{-1}) = T^{-7/4}$. This further increases the expected distance between A and B in case of failure to approximately $T^{7/4}$. We continue executing random walk iterations with a carefully chosen step length (distributing a budget of $O(T)$ queries among them). After i random walk iterations, the error probability is reduced to about $T^{-2+2^{-i}}$ (and the expected distance between the parties is roughly $T^{2-2^{-i}}$). Choosing $i \approx \log \log T$ gives an optimal error probability of about $T^{-2+1/\log T} = O(T^{-2})$.

Our new DDL protocol is presented in Algorithms 2 and 3. Algorithm 2 describes a single iteration of the random walk, parameterized by (L, T) which determine the maximal step length and the number of steps, respectively.⁵ Algorithm 3 describes the full protocol which is composed of application of the basic DDL algorithm (using $t_0 < T$ queries, reserving queries for the subsequent random walks), and then I additional random walks, where the i 'th random walk is parameterized by (L_i, t_i) which determine its maximal step length and number of steps. Between each two iterations in Step 6, both parties are moved forward by a large (deterministic) number of steps, in order to guarantee independence between the iterations (the computation time used to perform these calculations is negligible compared to T). We are free to choose the parameters $I, \{L_i, t_i\}$, as long as $\sum_{i=0}^I t_i = T$ is satisfied.

The very rough analysis presented above assumes that we have about T queries in each of the $\log \log T$ iterations, whereas we are only allowed T queries overall. Moreover, it does not accurately calculate the error probability and the distance between the parties in case of failure in each iteration. Taking all of these into account in an accurate analysis results in an error probability of $\Omega(\log T \cdot T^{-2})$. Surprisingly, we can still achieve an error probability of $O(T^{-2})$. This is done by a fine tuning of the parameters which distribute the number of queries among the iterations and select the step length of each random walk. In particular, it is not optimal to independently optimize the step length of each iteration and one has to analyze the subtle dependencies between the iterations in order to achieve an error probability of $O(T^{-2})$.

As the fine tuning of the parameters is rather involved, in addition to the theoretical analysis we verified the failure probability by extensive experiments.

⁵ We assume that the algorithm uses a table containing the pre-computed values g, g^2, \dots, g^{L-1} . Otherwise, it has to compute $g^{z_{i+1}}$ on-the-fly in Step 10, which results in a multiplicative penalty of $O(\log(T))$ on the number of group operations. Of course, it is also possible to obtain a time-memory tradeoff here.

Algorithm 2. $\text{RandW}_{L,T}(h)$

```

1 begin
2    $h' \leftarrow h, i \leftarrow 0, \min \leftarrow \infty, d_0 \leftarrow 0;$ 
3   while  $i < T$  do
4      $y \leftarrow \phi(h');$ 
5     if  $y < \min$  then
6        $h_{\min} \leftarrow h';$ 
7        $d_{\min} \leftarrow d_i, \min \leftarrow y;$ 
8     end
9      $z_{i+1} \leftarrow \psi_{L-1}(h');$ 
10     $h' \leftarrow h' \cdot g^{z_{i+1}};$ 
11     $d_{i+1} \leftarrow d_i + z_{i+1};$ 
12     $i \leftarrow i + 1;$ 
13  end
14  Output  $(d_{\min}, h_{\min});$ 
15 end

```

Algorithm 3. $\text{IteratedRandW}_{I,t_0,\{(L_i,t_i)_{i=1}^I\}}(h)$

```

1 begin
2    $(c_0, h_0) \leftarrow \text{Basic}_{t_0}(h);$ 
3    $p_0 \leftarrow c_0;$ 
4    $i \leftarrow 1;$ 
5   while  $i \leq I$  do
6      $h'_{i-1} \leftarrow h_{i-1} \cdot g^{\sum_{j<i} t_j L_j};$ 
7      $(c_i, h_i) \leftarrow \text{RandW}_{L_i,t_i}(h'_{i-1});$ 
8      $p_i \leftarrow p_{i-1} + c_i;$ 
9      $i \leftarrow i + 1;$ 
10  end
11  Output  $p_I;$ 
12 end

```

3.2 Related Work

The most closely related work to our DDL algorithm is Pollard’s “kangaroo” method for solving the discrete logarithm problem in an interval using limited memory (see [15] and [11, 16] for further analysis and extensions). The kangaroo method launches two random walks (kangaroos), one from the input $h = g^x$ (where x the unknown discrete log) and one from g^y , where y is a known value in an interval of a fixed size R around x . The algorithm is optimized such that the walks meet at a “distinguished point”, which reveals x . The kangaroo method thus resembles a single random walk iteration of our DDL algorithm.

On the other hand, there are fundamental differences between the standard DLI and DDL. These differences result in the iterative structure of our algorithm that differs from Pollard’s method. First, in contrast to the DLI problem, in DDL

A and B cannot communicate and never know if they succeed to synchronize. Hence, the parties cannot abort the computation at any time. Second, the goal in DDL is to minimize the error probability, whereas achieving a constant error probability (as in standard DLI) is unsatisfactory. To demonstrate the effect of these differences, observe that solving the discrete log problem in an interval of size 3 can be trivially done with probability 1 using 3 group operations. On the other hand, our algorithm for solving DDL for $M = 1$ is much more complicated and achieves an error probability of about T^{-2} using T group operations (which is essentially optimal for many concrete group families).

Yet another difference between DLI and DDL is that in DLI the boundaries of the interval of the input h are known, whereas in DDL the input of each party is completely uniform. The knowledge of the interval boundaries in DLI allows to shift it to the origin (using the self-reducibility property of discrete log) and efficiently use preprocessing (with a limited amount of storage) to speed up the online computation [2]. On the other hand, it is not clear how to efficiently exploit preprocessing in DDL.

4 The New Distributed Discrete Log Protocol

In this section we study our new DDL protocol in more detail. In Sect. 4.1 we focus on a single iteration of our DDL protocol (i.e., a single random walk iteration) and analyze its failure probability and the expected distance between its outputs in case of a failure. In Sect. 4.2 we briefly analyze the complete protocol. Some parts of the analysis are quite involved and presented in the extended version of this paper. This includes the proofs of Lemmas 5, 8, 9 and Theorem 1.

The experimental verification of the protocol is presented in Sect. 4.3. We also describe some practical considerations regarding the protocol in the extended version of this paper.

4.1 A Single Iteration of Our DDL Protocol – The Random Walk DDL Algorithm

Recall that in Algorithm 2 applied with parameters (L, T) , both parties perform a random walk of T steps of the form $g^y \rightarrow g^{y+a_i}$, where the length a_i of each step is determined by a (pseudo) random function $\psi_{L-1} : \mathbb{G} \rightarrow \{1, 2, \dots, L-1\}$ which guarantees that the step length is uniformly distributed in the range $[1, L-1]$. Each party then chooses among the elements of \mathbb{G} visited by its walk, the element h_{min} for which $\phi(h_{min})$ is minimal (as in the basic DDL algorithm).⁶

Once the parties synchronize in a given iteration, they remain synchronized in the subsequent ones as each iteration is deterministic. Thus, an application

⁶ We assume in our analysis that during the application of the whole protocol by a single party, each function ϕ, ψ_{L-1} is not evaluated twice on the same input. These constraints are satisfied since $|\mathbb{G}| = N$ is much larger than T (e.g., $|\mathbb{G}| > cT^2$ for a sufficiently large constant c).

of Algorithm 2 is “relevant” only if the two parties failed to synchronize in the previous iterations. In this case, the initial distance b between the parties is the difference between their outputs in the previous iteration. We shall compute the probability of failure as a function of b (which allows us to treat b as a constant throughout the analysis), and then substitute the expectation of b – computed in the analysis of the previous step – into the computation. In particular, as the expected distance between the outputs in case of failure of the basic DDL algorithm was computed in Lemma 2, we will be able to substitute it as $\mathbb{E}[|b|]$ into the computation of the failure probability of the second iteration of our DDL protocol. At the same time, in addition to the failure probability we shall compute the expected distance between the outputs in case of failure of Algorithm 2 in order to be able to link the examined iteration to the subsequent one.

Additional Notation. In our analysis we use some auxiliary notation. Without loss of generality, we assume that $b \geq 0$ (namely, B is located at distance b in front of A). We let S_A be the number of steps of A until its walk lands on an element visited by B (i.e., the number of queries made by A strictly before the first element of A that is included in B ’s path). If this never occurs, we let $S_A = T$. Similarly, we define S_B as the number of steps of B until its walk lands on an element visited by A . Clearly, the walks of A and B coincide for $T - \max(S_A, S_B)$ steps.

We define U_A as the number of steps A performs until it is within reach of a single step from the starting point of B . Namely, $U_A = \min\{i \mid d_i^A > b - L\}$, where d_i^A is the variable d_i in Algorithm 2 applied by A . In addition, we let V_A, V_B denote the numbers of steps performed by A and B , respectively, starting from the point where A is within reach of a single step from the starting point of B , until the walks collide or one of them ends. Furthermore, we denote $V_m = \max\{V_A, V_B\}$. Notice that $S_A = U_A + V_A$ and $S_B = V_B$, and hence

$$\max(S_A, S_B) \leq U_A + V_m. \tag{1}$$

Below, we evaluate the expectations of the random variables U_A, V_m in order to bound the error probability of synchronization based on Algorithm 2.

Finally, while $\text{RandW}_{L,T}(h)$ has two outputs, we slightly abuse notation and refer to d_{\min} as the (only) output of $\text{RandW}_{L,T}(h)$ since only this output is relevant for this analysis.

The Failure Probability of Algorithm 2. First, we bound the expected number of steps performed by A until it reaches the starting point of B .

Lemma 3. $\mathbb{E}[U_A] < 2b/L$.

Proof. By the definition of U_A , we have $d_{U_A}^A < b$. Consider the martingale $d'_i = d_i^A - iL/2$ (which is indeed a martingale, as $\psi_{L-1}(h')$ computed in the algorithm are independent and have expectation $L/2$). The classical *Doob’s martingale theorem* yields

$$0 = d'_0 = \mathbb{E}[d'_{U_A}] = \mathbb{E}[d_{U_A}^A] - L \mathbb{E}[U_A]/2.$$

As $d_{U_A} < b$, we deduce $\mathbb{E}[U_A] < 2b/L$. ■

Our next lemma bounds the expectation of $S_A + S_B$, that is, the total number of steps performed by the two walks together before they meet.

Lemma 4. *Suppose the initial distance between the parties, b , satisfies $0 < b < L$. Then $\mathbb{E}[S_A + S_B] \leq L - 1$.*

Proof. For ease of computation, we do not trim S_A and S_B with T . Of course, this can only make the upper bound larger. One easily sees that $\mathbb{E}[S_A + S_B]$ is finite and depends only on b (and the parameter L). Write E_b for this expectation. We have $E_0 = 0$, and by dividing into cases according to the result of a single step of A , we obtain

$$E_b = 1 + \frac{1}{L-1} (E_{b-1} + E_{b-2} + \dots + E_1 + E_0 + E_1 + \dots + E_{L-1-b}).$$

A valid solution for this system of linear equations is $E_b = L - 1$ for all $0 < b < L$. This is actually the only solution, since the matrix corresponding to this system is strictly diagonally dominant, and thus is invertible by the *Levy-Desplanques theorem*. Therefore, $\mathbb{E}[S_A + S_B] = L - 1$, independently of b . ■

The next lemma bounds the maximum between the numbers of steps performed by A and B between the time A “almost” reached the starting point of B and the meeting of the walks.

Lemma 5. $\mathbb{E}[V_m] \leq (L - 1)/2 + \sqrt{8(L - 1)}$.

The proof of the lemma is a lengthy technical argument, which mainly uses Lemma 4 and Doob’s martingale theorem.

Now we are ready to estimate the failure probability of Algorithm 2.

Lemma 6. *Let $R = 2b/L + L/2 + \sqrt{8L}$ for $0 < b < L$. The error probability of the random walk DDL algorithm satisfies*

$$\Pr_x[\mathbf{err}(\text{RandW}, x, b, T)] = \Pr[(\text{RandW}_{L,T}(g^x) - \text{RandW}_{L,T}(g^{x+b}) \neq b)] \leq 2R/(T + R).$$

Proof. The walks of A and B coincide for $T - \max(S_A, S_B)$ steps. Notice that we have,

$$\mathbb{E}[\max(S_A, S_B)] \leq \mathbb{E}[U_A + V_m] \leq L/2 + 2b/L + \sqrt{8L} = R, \tag{2}$$

where the first inequality uses (1) and the second inequality uses Lemmas 3 and 5. Similarly to the basic DDL algorithm (Lemma 1), the error probability (assuming that the output of ϕ on each element is uniformly random) is

$$\begin{aligned} \Pr_x[\mathbf{err}(\text{RandW}, x, b, T)] &= \mathbb{E}[2 \max(S_A, S_B)/(T + \max(S_A, S_B))] \\ &\leq \frac{\mathbb{E}[2 \max(S_A, S_B)]}{T + \mathbb{E}[\max(S_A, S_B)]} \leq \frac{2R}{T + R}, \end{aligned}$$

where the first inequality is Jensen’s inequality applied to the increasing concave function $x \mapsto 2x/(T + x)$ in the domain $x > 0$, and the second inequality uses the monotonicity of the function $x \mapsto 2x/(T + x)$ and Eq. (2). ■

The Output Difference in Case of Failure. Similarly to Lemma 2 which bounded the expected difference of outputs in case of failure for the basic DDL algorithm, we bound the analogous quantity for Algorithm 2. In order to achieve this result, we need a “conditional” version of the classical *Azuma martingale inequality*.

Lemma 7 (Azuma’s inequality). *Let X_0, X_1, \dots, X_n be a martingale with $|X_i - X_{i-1}| \leq V$. Then for any $t \geq 0$,*

$$\Pr [|X_n - X_0| \geq V \cdot t\sqrt{n}] \leq 2 \exp(-t^2/2).$$

Lemma 8. *Let X_1, \dots, X_n be independent random variables with $|X_i - \mathbb{E}[X_i]| \leq V$ and let \mathcal{E} be an event. Then*

$$\mathbb{E} \left[\max_{k=1}^n \left| \sum_{i \leq k} (X_i - \mathbb{E}[X_i]) \right| \middle| \mathcal{E} \right] \leq V \sqrt{8n \log(2/\Pr[\mathcal{E}])}.$$

To understand the intuition behind the lemma, consider the sum of independent $\{1, -1\}$ random variables $\{X_i\}_{i=1}^n$. By Chernoff’s inequality, $\Pr[|\sum X_i| > t\sqrt{n}] < e^{-t^2/2}$. However, if we condition on the event $\mathcal{E} = \{\sum X_i = n\}$, then we have $\Pr[|\sum X_i| > t\sqrt{n} | \mathcal{E}] = 1$ for all $t < \sqrt{n}$. On the other hand, the probability of the event \mathcal{E} is extremely small. We claim that if one is allowed to condition only on events with not-so-small probability, then all sums $|\sum_{i=1}^k X_i|$ are not much larger than the Chernoff bound, which applies without the conditioning. Our lemma is the martingale version of this intuition.

Finally, we show that either the error probability is “very small” (and so there is no need to continue the random walk iterations), or we can bound the distance between the outputs in case of failure.

Lemma 9. *If $\Pr_{\text{err}}(\text{RandW}_{L,T}, [1, 1], T) \geq \epsilon$, then for $0 < b < L$ and $h_1 \in \mathbb{G}$*

$$\begin{aligned} \mathbb{E} [|\text{RandW}_{L,T}(h_1) - \text{RandW}_{L,T}(h_1 \cdot g^b) - b| \mid \text{err}] \\ \leq b + TL/4 + L\sqrt{32T \log(2/\epsilon)}. \end{aligned} \tag{3}$$

The proof of the lemma is a somewhat lengthy technical argument which mainly uses Lemma 8.

4.2 The Iterated Random Walk DDL Algorithm

As described in Sect. 3, our full DDL protocol (i.e., Algorithm 3) runs iteratively several stages of Algorithm 2. It depends on a set of parameters: I , which is the number of iterations in the algorithm (on top of the basic DDL algorithm), $(t_i)_{i=0}^I$, which represent the number of queries in each of the $I + 1$ iterations, and $(L_i)_{i=1}^I$, which determine the (maximal) sizes of steps performed in each random walk iteration.

Given a set of parameters, Lemmas 6 and 9 allow us to compute the failure probability of IteratedRandW under that set of parameters. A “naive” choice of

parameters leads to a failure probability of $O(T^{-2} \log T)$. However, we show in the following theorem that the parameters can be chosen in such a way that the failure probability becomes $O(T^{-2})$.

Theorem 1. *There exists a parameter set PS for which the error probability of the iterated random walk DDL algorithm is*

$$\begin{aligned} & \Pr_{\text{err}}(\text{IteratedRandW}_{PS}, [1, 1], T) \\ &= \Pr[\text{IteratedRandW}_{PS}(g^x) - \text{IteratedRandW}_{PS}(g^{x+1}) \neq 1] \\ &\leq 2^{10.2+o(1)}/T^2. \end{aligned}$$

We give concrete choices of parameters and experimentally obtained values of the error probability for various values of T (that take into consideration the low-order terms) in the extended version of this paper.

A simple *distance extension* argument (see Sect. 5.5, Lemma 14) allows us to obtain a similar result for larger distances between the starting points.

Corollary 1. *Consider Algorithm 3 with the parameter set PS chosen in Theorem 1. Then for any distribution of the initial distance b that has expectation $\mathbb{E}[|b|]$, the error probability of Algorithm 3 is at most $O(\mathbb{E}[|b|]/T^2)$. In particular, $\Pr_{\text{err}}(\text{IteratedRandW}_{PS}, [-M, M], T) = O(M/T^2)$.*

We note that when $\mathbb{E}[|b|] \gg 1$, it is more efficient to start the sequence of iterations directly with a random walk of expected step length of roughly $\sqrt{\mathbb{E}[|b|]}$ (instead of starting it with Algorithm 1). This reduces the error probability by a constant factor.

Dependency of Iterations and Parameter Selection. In Sect. 3, we noted that in order to minimize the error probability of the protocol it is necessary to consider the dependencies between the different iterations, rather than optimizing their error probability independently. We demonstrate this here by analyzing only two iterations of Algorithm 2. We first simplify the formulas in Lemmas 6 and 9 by ignoring low-order terms.

Lemma 6 asserts that the error probability of each iteration is $2R/(T + R)$, where $R \approx 2b/L + L/2$ (ignoring low-order terms). We can lower bound this probability by $2R/T \approx (4b + L^2)/2TL$. Lemma 9 asserts that the expected distance between the parties in case of error is at most $b + TL/4$ (ignoring low-order terms). Since L is roughly \sqrt{b} and $b < T^2$ at any iteration (as noted in Sect. 3), then we estimate the expected distance as $TL/4$.

Let us assume that after iteration $i - 1$ the parties are at distance b_i (which is a random variable). Assume we are allowed T' queries in the next two iterations, where in iteration i we use step length parameter L_i and query parameter t_i . After this iteration, the expected distance between the parties is b_{i+1} , which we estimated above as $t_i L_i/4$. Similarly, in iteration $i + 1$ we use parameters L_{i+1} and t_{i+1} (under the restriction that $t_i + t_{i+1} = T'$). Thus, the estimated

error probability of the two iterations (assuming that the previous ones failed to synchronize) is

$$\frac{4b_i + L_i^2}{2t_i L_i} \cdot \frac{4b_{i+1} + L_{i+1}^2}{2t_{i+1} L_{i+1}}.$$

To simplify this expression, for parameters α, β write $L_i = \alpha\sqrt{b_i}$ and $L_{i+1} = \beta\sqrt{b_{i+1}} \approx \beta/2 \cdot \sqrt{t_i L_i}$. Then, the error probability expression simplifies to

$$\frac{b_i^{3/4}}{8} \cdot \frac{(4 + \alpha^2)(4 + \beta^2)}{\alpha^{1/2}\beta \cdot t_i^{1/2}t_{i+1}}.$$

It remains to minimize this expression by appropriately selecting $\alpha, \beta, t_i, t_{i+1}$ (such that $t_i + t_{i+1} = T'$). Taking partial derivatives, it is clear that the global minimum is obtained by selecting different parameters for the two iterations, namely, the values of α, t_i for iteration i are different from β, t_{i+1} for iteration $i + 1$. In other words, in order to minimize the error probability it is necessary to solve the global optimization problem rather than optimize the error probability of each iteration independently.

4.3 Experimental Verification

In order to evaluate Algorithm 3 in practice, we programmed a simulator which simulates $\text{IteratedRandW}_{PS}(g^x)$ and $\text{IteratedRandW}_{PS}(g^{x+1})$, and empirically approximates $\text{Pr}[\text{err}]$ as the percentage of pairs of simulations which disagree. The parameters for these simulations were chosen using a numerical optimizer based on the analysis above. The results are given in Table 1.⁷ Since performing full simulations is too expensive for large values of T , we had to use three optimizations that do not affect the simulator’s reliability. These are detailed in the extended version of this paper.

5 Error Probability Lower Bounds in Concrete Group Families

5.1 Overview of the Lower Bound Proof

We outline the main ideas of the lower bound proof in concrete family groups. For the sake of simplicity, we only consider DDL with $M = 1$ in this overview (whereas the proof considers general M).

⁷ We note that according to the proof of Theorem 1 (given in the extended version of this paper), I can be taken to be any value between $\log \log(T) + \omega(1)$ and $o(\sqrt{\log(T)})$, without a significant effect on the provable performance of the algorithm. On the other hand, according to Table 1, I seems to grow more sharply. However, the restriction of $I = o(\sqrt{\log(T)})$ is merely an artifact of the proof and the optimal value of I could be asymptotically larger. Furthermore, the sharp increase of the values of I in Table 1 could be attributed to low-order terms that have a more noticeable effect for small T values.

Table 1. Experimental results

T	I	$\sim \log_2(t_k)$ $\sim \log_2(L_k)$	$T^2 \cdot \Pr[\text{err}]$ $\sigma(SD)$	$T^2 \cdot \Pr[\text{err}]$ $\sigma(SD)$
2^{13}	5	6.0, 8.6, 9.6, 10.4, 11.1, 11.7	334	336.6
		1.6, 3.6, 5.6, 7.5, 9.4	2	0.3
2^{16}	6	7.1, 10.3, 11.6, 12.5, 13.3, 14.1, 14.7	390	382.5
		1.6, 3.9, 6.2, 8.3, 10.3, 12.2	10	1
2^{19}	7	7.0, 10.3, 12.6, 13.7, 14.7, 15.5, 16.3, 17.1, 17.7	391	394
		1.6, 2.8, 5.2, 7.4, 9.6, 11.5, 13.4, 15.2	25	2
2^{22}	8	8.2, 12.6, 15.1, 16.4, 17.5, 18.5, 19.3, 20.1, 20.7	—	420
		1.6, 3.7, 6.7, 9.4, 11.8, 14.1, 16.2, 18.1	—	4
2^{25}	9	8.4, 13.0, 16.5, 18.0, 19.3, 20.5, 21.5, 22.3, 23.1, 23.8	—	427
		1.6, 3.2, 6.4, 9.4, 12.2, 14.7, 17.0, 19.1, 21.1	—	10

The fourth column gives the result of simulations without the third optimization (detailed in the extended version of this paper), and the last column uses that optimization.

We first prove in Lemma 10 that in a DDL protocol, using different algorithms for A, B cannot give a significant advantage in the error probability. As a result, we can assume that both A and B use A 's algorithm, which simplifies the analysis.

Let us assume that we can solve DDL with error probability $\delta \ll 1$ in time T for $M = 1$. Our main reduction shows how to use A 's algorithm to solve DLI in an interval of length about $R \approx 1/\delta$ in time less than $4T$ with probability $1/2$. If we assume that in a specific family of groups, DLI in an interval of length $c \cdot T^2$ (for a sufficiently large constant c) cannot be solved in complexity lower than $4T$ with probability $1/2$,⁸ we must have $R \approx 1/\delta < c \cdot T^2$ or $\delta = \Omega(T^{-2})$, which gives our main lower bound for the case of $M = 1$. It is important to stress that A is a DDL algorithm for $M = 1$ that is not explicitly given the DLI interval length parameter R . Yet the reduction below will apply A 's algorithm to solve DLI with parameter R in a black-box manner.

Recall that a DLI algorithm obtains as input a group element $h = g^x$, where x is in a known interval of length $R \approx 1/\delta$. By the self-reducibility of discrete log, we can assume that h is a uniform group element (i.e., we can multiply the input by a randomly chosen group element). Our reduction picks a point g^z in the interval (for a known z) and runs A on inputs g^x and g^z , where $|x - z| \leq R$. We hope that $A(g^x) - A(g^z) = z - x$ and thus we return $z - (A(g^x) - A(g^z)) = x$.

Clearly, the DLI algorithm runs in time less than $4T$ and it remains to upper bound its error probability by $1/2$. In other words, we need to upper bound

⁸ We consider only uniform algorithms that can be applied to families of groups (such as elliptic curve groups) and not non-uniform algorithms that are specialized to a specific group \mathbb{G} . Indeed, in the non-uniform model, there exist algorithms that solve DLI in an interval of length R in time $o(\sqrt{R})$ for any specific group (see, e.g., [2]).

the probability of $A(g^x) - A(g^z) \neq z - x$ by $1/2$. We know that the DDL error probability is δ for $M = 1$, namely, if $|x - z| \leq 1$, then the required probability is⁹ δ . Next, assume that $z = x + 2$. Then, if $A(g^x) - A(g^{x+2}) \neq 2$ this implies that either $A(g^x) - A(g^{x+1}) \neq 1$ or $A(g^{x+1}) - A(g^{x+2}) \neq 1$ (or both). Since the probability of each of these two events is δ , we can use a union bound to upper bound the probability that $A(g^x) - A(g^{x+2}) \neq 2$ by 2δ . Using a similar argument (which we refer to as *distance extension*, formalized in Lemma 14), we can upper bound the probability of the event $A(g^x) - A(g^z) \neq z - x$ for $|x - z| \leq R$ by $O(R \cdot \delta)$ and for $R = O(1/\delta)$, this gives error probability $1/2$, as required. Note that the same algorithm A is used for any distance $|x - z| \leq R$ (which is unknown in advance) and conditioning on this distance is only done for the sake of analysis.

5.2 The Single Algorithm Distributed Discrete Log Problem

We now define the *single algorithm* DDL problem, which is the same problem as general DDL with the restriction that the algorithms of the parties are the same (i.e., both parties use A 's algorithm). Denote by $\mathbf{err}(A, x, b, T)$ the event $A(g^x) - A(g^{x+b}) \neq b$ and by $\Pr_{\mathbf{err}}(A, [M_1, M_2], T)$ its probability (over $x \in \mathbb{Z}_N, b \in [M_1, M_2]$). Obviously, the optimal 2-party DDL error probability is a lower bound on the optimal single algorithm DDL error probability. In this section, we prove that the bound in the other direction holds as well up to a constant factor in case $M_2 = -M_1 = M$.¹⁰

Lemma 10. $\Pr_{\mathbf{err}}(A, B, [-M, M], T) \geq 1/8 \cdot \Pr_{\mathbf{err}}(A, [-M, M], T)$.

Proof. Note that if $A(g^{x+b_1}) - A(g^{x+b_2}) \neq b_2 - b_1$, then $A(g^{x+b_1}) - B(g^x) \neq -b_1$ or $A(g^{x+b_2}) - B(g^x) \neq -b_2$ (or both). Therefore, for uniform $b_1, b_2 \in [-M, M]$,

$$\begin{aligned} \Pr_{x, b_1, b_2} [\mathbf{err}(A, x + b_1, b_2 - b_1, T)] &= \Pr_{x, b_1, b_2} [A(g^{x+b_1}) - A(g^{x+b_2}) \neq b_2 - b_1] \\ &\leq \Pr_{x, b_1, b_2} [(A(g^{x+b_1}) - B(g^x) \neq -b_1) \cup (A(g^{x+b_2}) - B(g^x) \neq -b_2)] \\ &\leq \Pr_{x, b_1} [A(g^{x+b_1}) - B(g^x) \neq -b_1] + \Pr_{x, b_2} [A(g^{x+b_2}) - B(g^x) \neq -b_2] \\ &= 2 \cdot \Pr_{\mathbf{err}}(A, B, [-M, M], T) \end{aligned}$$

It remains to relate $\Pr_{\mathbf{err}}(A, [-M, M], T)$ to $\Pr_{x, b_1, b_2} [\mathbf{err}(A, x + b_1, b_2 - b_1, T)]$. Denote the event $|b_2 - b_1| \leq M$ by \mathcal{E} and note that $\Pr_{b_1, b_2} [\mathcal{E}] \geq 1/2$. Conditioned on \mathcal{E} , if $b_2 - b_1$ was uniform in $[-M, M]$, then we would have $\Pr_{x, b_1, b_2} [\mathbf{err}(A, x + b_1, b_2 - b_1, T) | \mathcal{E}] = \Pr_{\mathbf{err}}(A, [-M, M], T)$. Although it is not uniform, $b_2 - b_1$ is almost uniform in the sense that for each $i \in [-M, M]$, we have $\Pr_{b_1, b_2} [b_2 - b_1 = i] \approx 1/(2M)$.

⁹ More accurately, it is $O(\delta)$, as δ is the average error probability in the interval $[-1, 1]$.
¹⁰ It is also possible to prove a similar bound in case the interval $[M_1, M_2]$ is not symmetric around the origin.

$i] \geq \Pr_{b_1, b_2} [b_2 - b_1 = M] \geq (M + 1)/(4M^2)$ and $\Pr_{b_1, b_2} [b_2 - b_1 = i] \leq \Pr_{b_1, b_2} [b_2 - b_1 = 0] \leq (2M + 1)/(4M^2)$. As the minimal and maximal probabilities assigned to $|b_2 - b_1|$ in $[-M, M]$ are within a factor of 2,

$$\Pr_{x, b_1, b_2} [\mathbf{err}(A, x + b_1, b_2 - b_1, T) | \mathcal{E}] \geq 1/2 \cdot \Pr_{\mathbf{err}}(A, [-M, M], T)$$

and

$$\begin{aligned} & \Pr_{x, b_1, b_2} [\mathbf{err}(A, x + b_1, b_2 - b_1, T)] \\ & \geq \Pr_{x, b_1, b_2} [\mathbf{err}(A, x + b_1, b_2 - b_1, T) | \mathcal{E}] \cdot \Pr_{b_1, b_2} [\mathcal{E}] \geq 1/4 \cdot \Pr_{\mathbf{err}}(A, [-M, M], T). \end{aligned}$$

Finally,

$$\begin{aligned} \Pr_{\mathbf{err}}(A, B, [-M, M], T) & \geq 1/2 \cdot \Pr_{x, b_1, b_2} [\mathbf{err}(A, x + b_1, b_2 - b_1, T)] \\ & \geq 1/8 \cdot \Pr_{\mathbf{err}}(A, [-M, M], T), \end{aligned}$$

concluding the proof. ■

The consequence of the lemma is that for the sake of proving lower bounds on the error probability, we can restrict our attention to A 's algorithm by analyzing $\Pr_{\mathbf{err}}(A, [-M, M], T)$. The lemma immediately gives us the same lower bound on $\Pr_{\mathbf{err}}(A, B, [-M, M], T)$, up to a constant factor.

Furthermore, note that by symmetry we have $\Pr_{\mathbf{err}}(A, B, [-M, M], T) \geq 1/8 \cdot \Pr_{\mathbf{err}}(B, [-M, M], T)$, hence the general DDL error probability is lower bounded by the maximal error probability of the (single) algorithms of the two parties (up to constant factors). Therefore, running different algorithms for the two parties cannot give a much better result than simply having both players run the best algorithm in the single algorithm setting.

5.3 Limitation on Randomness

The effect of the internal randomness of a DDL algorithm A on its outcome is quantified by $\Pr_{\mathbf{err}}(A, [0, 0], T)$. This quantity measures the probability that two different executions of A on the same input differ, where the probability is taken over A 's input g^x and its internal randomness. We prove that A 's internal randomness cannot significantly influence its outcome.

Lemma 11. *Assume $\Pr_{\mathbf{err}}(A, [-M, M], T) = \delta$. Then $\Pr_{\mathbf{err}}(A, [0, 0], T) \leq 2\delta$.*

Proof. To be more explicit, we denote by $A(r, g^x)$ the execution of A with a randomness string r . Assume we fix the output of $A(r, g^{x+b})$ for some $b \in [-M, M]$ and randomness string r . Then, if $A(r_1, g^x) \neq A(r_2, g^x)$ for r_1, r_2 , either

$A(r_1, g^x) - A(r, g^{x+b}) \neq b$ or $A(r_2, g^x) - A(r, g^{x+b}) \neq b$ (or both). Hence,

$$\begin{aligned} & \Pr_{\mathbf{err}}(A, [0, 0], T) \\ &= \Pr_{r_1, r_2, x} [A(r_1, g^x) \neq A(r_2, g^x)] \\ &\leq \Pr_{r_1, r_2, r, x, b} [(A(r_1, g^x) - A(r, g^{x+b}) \neq b) \cup (A(r_2, g^x) - A(r, g^{x+b}) \neq b)] \\ &\leq \Pr_{r_1, r, x, b} [A(r_1, g^x) - A(r, g^{x+b}) \neq b] + \Pr_{r_2, r, x, b} [A(r_2, g^x) - A(r, g^{x+b}) \neq b] \\ &= 2\delta. \end{aligned}$$

■

5.4 Symmetry

We prove the following symmetric property.

Lemma 12. For $M_2 \geq M_1$, $\Pr_{\mathbf{err}}(A, [M_1, M_2], T) = \Pr_{\mathbf{err}}(A, [-M_2, -M_1], T)$.

Proof. It is sufficient to prove that for any positive integer b , $\Pr_{\mathbf{err}}(A, [b, b], T) = \Pr_{\mathbf{err}}(A, [-b, -b], T)$. This indeed holds, since $\mathbf{err}(A, x, b, T)$ and $\mathbf{err}(A, x + b, -b, T)$ are identical events, $\Pr_{\mathbf{err}}(A, [b, b], T) = \Pr_x[\mathbf{err}(A, x, b, T)] = \Pr_x[\mathbf{err}(A, x + b, -b, T)] = \Pr_{\mathbf{err}}(A, [-b, -b], T)$. ■

5.5 Distance Extension

In this section, we show that the distance parameter M of any DDL algorithm A can be extended at the expense of a linear loss in the error probability.

First, we prove the following lemma which reduces the error probability of $\Pr_{\mathbf{err}}(A, [-M, M], T)$ to each one of the indices in the interval. As mentioned in the Introduction (see Footnote 2), it proves that a DDL algorithm with error probability δ for uniform $b \in [-M, M]$ also solves DDL with an error probability $O(\delta)$ for any distribution on $b \in [-M, M]$.

Lemma 13. Assume that $\Pr_{\mathbf{err}}(A, [-M, M], T) = \delta$. Then, for every $b \in [-M, M]$, $\Pr_{\mathbf{err}}(A, [b, b], T) \leq 4\delta$.

Proof. We first assume that $b \in [1, M]$ and let $i \in [-M, M - b]$. Clearly, if g^x is a uniform group element, then so is g^{x+i} . Therefore, $\Pr_{\mathbf{err}}(A, [b, b], T) = \Pr_x[\mathbf{err}(A, x + i, b, T)]$. Furthermore, if the event $\mathbf{err}(A, x + i, b, T)$ occurs, then $A(g^{x+i}) - A(g^{x+i+b}) \neq b$ implying that at least one of the events $A(g^x) - A(g^{x+i+b}) \neq i + b$ and $A(g^x) - A(g^{x+i}) \neq i$ must occur. Consequently, $\Pr_{\mathbf{err}}(A, [b, b], T)$ (the error probability associated with index b) is upper bounded

by $\Pr_{\text{err}}(A, [i+b, i+b], T) + \Pr_{\text{err}}(A, [i, i], T)$ (the sum of error probabilities associated with indices i and $i + b$). Formally,

$$\begin{aligned} \Pr_{\text{err}}(A, [b, b], T) &= \Pr_x[\text{err}(A, x + i, b, T)] \\ &\leq \Pr_x[\text{err}(A, x, i + b, T) \cup \text{err}(A, x, i, T)] \\ &\leq \Pr_x[\text{err}(A, x, i + b, T)] + \Pr_x[\text{err}(A, x, i, T)] \\ &= \Pr_{\text{err}}(A, [i + b, i + b], T) + \Pr_{\text{err}}(A, [i, i], T). \end{aligned}$$

We map the indices in $[-M, M]$ into disjoint pairs of the form $i, i + b$ (this implies that $i \in [-M, M - b]$). We can obtain at least $\lfloor (2M - b + 1)/2 \rfloor$ such pairs, which is at least $(M + 1)/2$ for $b \in [1, M - 1]$. On the other hand, for $b = M$, the number of pairs is $M \geq (M + 1)/2$. We apply the above inequality to each of the pairs:

$$\begin{aligned} \delta &= \Pr_{\text{err}}(A, [-M, M], T) \\ &= 1/(2M + 1) \cdot \sum_{i=-M}^M \Pr_{\text{err}}(A, [i, i], T) \\ &\geq 1/(2M + 1) \cdot (M + 1)/2 \cdot \Pr_{\text{err}}(A, [b, b], T) \\ &\geq \Pr_{\text{err}}(A, [b, b], T)/4. \end{aligned}$$

Thus, $\Pr_{\text{err}}(A, [b, b], T) \leq 4\delta$ for $b \in [1, M]$.

The proof for $b \in [-M, -1]$ follows by symmetry (Lemma 12). Finally, for $b = 0$, we have $\Pr_{\text{err}}(A, [0, 0], T) \leq 2\delta$ by Lemma 11. ■

Lemma 14. *Let $\Pr_{\text{err}}(A, [-M, M], T) = \delta$. Then for any $\beta > 1$,*

$$\Pr_{\text{err}}(A, [-\beta M, \beta M], T) \leq 8\beta \cdot \delta.$$

For the sake of simplicity we assume that βM is an integer (otherwise, we only consider integer values in $[-\beta M, \beta M]$).

Proof. First, we analyze $\Pr_{\text{err}}(A, [1, \beta M], T)$. Let $b \in [1, \beta M]$ and divide it by M , writing $b = b_1 \cdot M + b_2$, for integers $b_1 \leq \beta$ and $b_2 \in [0, M)$. We examine the following $b_1 + 1$ success events:

$$\begin{aligned} \mathcal{E}_1 &: \text{succ}(A, x, M, T) \\ \mathcal{E}_2 &: \text{succ}(A, x + M, M, T) \\ &\dots \\ \mathcal{E}_{b_1} &: \text{succ}(A, x + (b_1 - 1)M, M, T) \\ \mathcal{E}_{b_1+1} &: \text{succ}(A, x + b_1 M, b_2, T) \end{aligned}$$

Observe that if all $b_1 + 1$ events hold (i.e. $\bigcap_{i=1}^{b_1+1} \mathcal{E}_i$), then $A(g^x) - A(g^{x+b_1 M + b_2}) = A(x) - A(g^{x+b}) = b$, i.e., $\text{succ}(A, x, b, T)$ holds.

By Lemma 13, $\Pr_x[\bar{\mathcal{E}}_i] \leq 4\delta$ holds for each $i \in 1, 2, \dots, b_1$, while $\Pr_{x,b}[\bar{\mathcal{E}}_{b_1+1}] \leq 4\delta$ holds as well by the same lemma. Hence,

$$\begin{aligned} 1 - \Pr_{\text{err}}(A, [1, \beta M], T) &\geq \Pr_{x,b}[\bigcap_{i=1}^{b_1+1} \mathcal{E}_i] = 1 - \Pr_{x,b}[\bigcup_{i=1}^{b_1+1} \bar{\mathcal{E}}_i] \\ &\geq 1 - \sum_{i=1}^{b_1+1} \Pr_{x,b}[\bar{\mathcal{E}}_i] \geq 1 - (b_1 + 1)4\delta \geq 1 - (\beta + 1)4\delta \geq 1 - 8\beta \cdot \delta. \end{aligned}$$

Therefore, $\Pr_{\text{err}}(A, [1, \beta M], T) \leq 8\beta \cdot \delta$. By symmetry (Lemma 12), we have $\Pr_{\text{err}}(A, [-\beta M, -1], T) \leq 8\beta \cdot \delta$ as well. Since $\Pr_{\text{err}}(A, [0, 0], T) \leq 2\delta$ by Lemma 11, we conclude that $\Pr_{\text{err}}(A, [-\beta M, \beta M], T) \leq 8\beta \cdot \delta$, as claimed. ■

Remark 1. An open question of [5] asked whether the DDL error probability can be eliminated completely. If we apply the above lemma with no error (i.e., $\Pr_{\text{err}}(A, [-M, M], T) = \delta = 0$), we obtain $\Pr_{\text{err}}(A, [-\beta M, \beta M], T) = 0$, implying that the two parties (running A 's algorithm) never err for any distance. This allows the parties to collectively solve the discrete log problem in \mathbb{G} with probability 1 (a similar reduction will be formally presented in Algorithm 4), thus violating the security assumption of the underlying HSS scheme. Namely, the DDL error probability cannot be eliminated (in fact it is easy to show that it must be superpolynomial in $1/N$), answering negatively the open question of Boyle et al.

5.6 Reduction from Discrete Log in an Interval to Distributed Discrete Log

Recall that the discrete log problem in an interval (DLI) is parametrized by an interval length R for a cyclic multiplicative group \mathbb{G} of size N with generator g . The input to the problem is a group element $h = g^x$, where¹¹ $x \in [0, R - 1]$ and the goal is to recover x with high probability (which is at least a constant).

The following lemma reduces the DLI problem to DDL.

Lemma 15. *For a family of groups, assume that $\Pr_{\text{err}}(A, [-M, M], T) = \delta$, where $T \geq \log N$ and $\delta < 1/32$. Then discrete log in an interval of length $R = M/(32\delta)$ can be solved in complexity $4T$ with probability $1/2$.*

Proof. Consider Algorithm 4 for solving DLI on input $h = g^x$ for $x \in [0, R - 1]$. For the sake of simplicity we assume that R is even.

The algorithm computes $h \cdot g^y$ and $g^{y+(R/2)}$, which can be carried out by performing $2 \log N$ group operations using the square-and-multiply algorithm. It further invokes A twice in complexity $2T$ and therefore its total complexity is $2T + 2 \log N \leq 4T$ (since $T \geq \log N$).

¹¹ Alternatively, x could be in any fixed interval of length R . The exact interval is not important as one can easily reduce the problem in a given interval to any other interval.

Algorithm 4. DLI(h)

```

1 begin
2    $y \xleftarrow{R} \mathbb{Z}_N$ ;
3    $d_1 \leftarrow A(h \cdot g^y)$ ;
4    $d_2 \leftarrow A(g^{y+(R/2)})$ ;
5   Output  $(R/2) - (d_1 - d_2)$ ;
6 end
    
```

It remains to upper bound the error probability of the algorithm by $1/2$. The algorithm succeeds to return x if $(R/2) - (d_1 - d_2) = x$, namely $d_1 - d_2 = (R/2) - x$ or equivalently $A(g^{y+x}) - A(g^{y+(R/2)}) = (R/2) - x$. Since $y \in \mathbb{Z}_N$ is uniform, then g^{y+x} is a uniform group element. Moreover, since $x \in [0, R - 1]$, then $(R/2) - x \in [-R/2, R/2]$. Therefore, by Lemma 13, the error probability of the algorithm is at most

$$\begin{aligned}
 4 \cdot \Pr_{\text{err}}(A, [-R/2, R/2], T) &\leq 4 \cdot 8 \cdot R/(2M) \cdot \Pr_{\text{err}}(A, [-M, M], T) \\
 &= 16 \cdot R/M \cdot \delta = 1/2,
 \end{aligned}$$

where the first inequality is due to Lemma 14. Note that we use Lemma 13 (and pay a factor of 4 in the error probability), as $x \in [0, R - 1]$ may be selected by an adversary (whereas $\Pr_{\text{err}}(A, [-R/2, R/2], T)$ averages the error probability). ■

Theorem 2 is a simple corollary of Lemma 15.

Theorem 2. *For a specific family of groups, assume there exists a constant c such that for any group in the family of size N , DLI in an interval of length at least $c \cdot T^2$ cannot be solved in complexity $4T$ with probability at least $1/2$ (where $\log N \leq T < \mathcal{B}$ for a bound \mathcal{B}). Moreover, assume that there is a DDL protocol A for this family with time complexity parameter T , maximal distance parameter M and error probability $\Pr_{\text{err}}(A, [-M, M], T) = \delta$ for $\delta < 1/32$. Then $\delta = \Omega(M \cdot T^{-2})$.*

We note that the bound \mathcal{B} depends on N according to the concrete group family. For example, for some subgroups of \mathbb{Z}_p^* , \mathcal{B} is subexponential in $\log N$.

Proof. By Lemma 15, discrete log in an interval of length $R = M/(32\delta)$ can be solved in complexity $4T$ with probability $1/2$. By our assumption, $R = M/(32\delta) < c \cdot T^2$ implying that $\delta = \Omega(M \cdot T^{-2})$ as claimed. ■

6 Error Probability Lower Bounds for Non-Adaptive Algorithms in the Generic Group Model

In this section, we prove lower bounds on DDL algorithms in the generic group model (GGM), focusing on non-adaptive algorithms. We first review the generic

group model (GGM) we consider (which is slightly different than the one proposed by Shoup [18]) and formulate DDL in this model. This formulation is given for the additive group \mathbb{Z}_N , which is isomorphic to the multiplicative group \mathbb{G} of size N . We note that the proofs of most of the statements in this section are given in the extended version of this paper.

6.1 Distributed Discrete Log in the Generic Group Model

Let \mathbb{Z}_N be the additive group of integers, and let S be a set of bit strings of cardinality at least N . An encoding function of \mathbb{Z}_N on S is an injective map $\sigma : \mathbb{Z}_N \rightarrow S$.

A generic algorithm A for \mathbb{Z}_N on S for the discrete logarithm problem is a probabilistic algorithm that takes as input an *encoding list* of size 2, $\sigma(1), \sigma(x)$, namely, the encodings of a generator of \mathbb{Z}_N and a uniform $x \in \mathbb{Z}_N$, where σ is an encoding function of \mathbb{Z}_N on S . Throughout its execution, A continues to maintain the encoding list, and is allowed to extend it using oracle queries. An oracle query in our model specifies two indices $i, j \in \mathbb{Z}_N$. The oracle computes $\sigma(i \cdot x + j)$ and the returned bit string is appended to the encoding list. A succeeds to solve the discrete log problem if $A(\sigma; 1, x) = x$, and its success probability is taken over the uniform choices of $\sigma : \mathbb{Z}_N \rightarrow S$ and $x \in \mathbb{Z}_N$ (and perhaps additional randomness of its own coin tosses). We measure the complexity of A according to the number of oracle queries it makes. The following success probability upper bound was proved in [18].

Theorem 3 ([18]). *If a generic discrete log algorithm A is allowed T oracle queries, then $\Pr_{\sigma, x}[A(\sigma; 1, x) = x] = O(T^2/N)$, assuming that N is prime.*

We note that our GGM formulation is slightly stronger than the one of [18], where the queries of A are limited to linear combinations with coefficients of ± 1 of elements in its encoding list. Since any query (i, j) can be issued in Shoup’s original GGM after at most $O(\log N)$ queries using the double-and-add algorithm, a stronger GGM algorithm can be simulated by a standard one by increasing the query complexity by a multiplicative factor of $\log N$. However, by following its original proof in [18], it is easy to verify that Theorem 3 actually holds with no modification in our stronger GGM. Obviously, any algorithm in the original GGM is also an algorithm in the stronger GGM. Therefore, any lower bounds we obtain in the stronger GGM also apply in the original GGM.

We now describe the basic game of distributed discrete log in GGM. Obviously, all the results of Sect. 5 also hold in the generic group model. In particular, by Lemma 10 it is sufficient to consider single algorithm DDL to obtain general DDL lower bounds.

A party (algorithm) A is given as input $\sigma(1)$ and the encoding of an additional group element $\sigma(x)$ for $x \in \mathbb{Z}_N$, selected uniformly at random. Algorithm A is allowed to make T oracle queries. After obtaining the answers from the oracle, A returns an integer value. Two parties (both running A ’s algorithm) win the

DDL game in GGM if

$$A(\sigma; 1, x) - A(\sigma; 1, x + b) = b,$$

otherwise, they lose the game, or err.

We are interested in proving lower bounds on the DDL error probability as a function of T , namely

$$\Pr_{\sigma, x, b} [A(\sigma; 1, x) - A(\sigma; 1, x + b) \neq b].$$

Analogously to our notation for multiplicative groups, we denote by $\mathbf{err}(A, \sigma, x, b, T)$ the error event $A(\sigma; 1, x) - A(\sigma; 1, x + b) \neq b$, and by $\Pr_{\mathbf{err}}(A, \sigma, [M_1, M_2], T)$ its probability $\Pr_{\sigma, x, b}[\mathbf{err}(A, \sigma, x, b, T)]$, where $b \in [M_1, M_2]$ is a uniform integer. We further denote by $\mathbf{succ}(A, \sigma, x, b, T)$ the complementary success event.

6.2 An Error Probability Lower Bound for Arbitrary Generic Algorithms

The following theorem gives a DDL error probability lower bound in GGM. The theorem is a somewhat weaker statement than Theorem 2 (which has implications in concrete group families).

Theorem 4. *For any generic DDL algorithm A , $\Pr_{\mathbf{err}}(A, \sigma, [-M, M], T) = \Omega(M \cdot T^{-2})$, given that $M = O(T^2)$, $T = o(\sqrt{N})$, and N is prime.*

We omit the proof, as it is similar to the one of Theorem 2. It applies a reduction to discrete log, while using Theorem 3 to obtain the error probability lower bound. Alternatively, one could obtain a hardness result for DLI in GGM (extending Theorem 3 to smaller intervals) and apply Theorem 2 directly.

6.3 An Error Probability Lower Bound for Non-Adaptive Generic Algorithms

In this section, we prove a lower bound on the DDL error probability of non-adaptive generic algorithms, whose oracle queries $\{(i_1, j_1), (i_2, j_2), \dots, (i_T, j_T)\}$ are fixed in advance and do not depend on previous answers.

We will prove the following lower bound:

Theorem 5. *Any non-adaptive DDL algorithm A satisfies $\Pr_{\mathbf{err}} A, \sigma, [-1, 1], T = \Omega(1/T)$, given that $T = o(N^{1/2})$, and N is prime.*

Overview of the Lower Bound Proof on Non-Adaptive Algorithms in the Generic Group Model. Let us first consider the class of algorithms that make T consecutive oracle queries to group elements (such as Algorithm 1 and the ones of [4–6] in general). Consider the executions $A(\sigma; 1, x)$ and $A(\sigma; 1, x + T)$, which query $2T$ disjoint group elements. In GGM, algorithm executions

that query disjoint elements are essentially independent (as each group element is associated with a random string), which implies that the probability that $A(\sigma; 1, x) - A(\sigma; 1, x + T) \neq T$ is at least $1/2$. Recall that we are interested in the probability that $A(\sigma; 1, x) - A(\sigma; 1, x + 1) \neq 1$ and it can be lower bounded by $\Omega(T^{-1})$ using distance extension (Lemma 14). A similar lower bound applies if A only queries group elements in a short interval of length $O(T)$.¹²

Of course, we are interested in proving the $\Omega(T^{-1})$ lower bound for arbitrary non-adaptive algorithms. The main idea that allows us to achieve this is to define a transformation that takes an arbitrary non-adaptive algorithm A' and maps its T queries to a small interval of size $O(T)$, obtaining a new algorithm A (for which the error lower bound $\Omega(T^{-1})$ holds). We require that the query mapping preserves the error probability of A' , thus proving that the error probability lower bound $\Omega(T^{-1})$ above also applies to non-adaptive algorithms in general. In order to preserve the error probability of A' , the mapping will ensure that the joint input distribution of $A'(\sigma; 1, x)$ and $A'(\sigma; 1, x + 1)$ is equal to that of $A(\sigma; 1, x)$ and $A(\sigma; 1, x + 1)$. In the generic group model, this means that the mapping should preserve joint queries, namely, satisfy the condition that query i of $A'(\sigma; 1, x)$ and query j of $A'(\sigma; 1, x + 1)$ evaluate the same group element if and only if query i of $A(\sigma; 1, x)$ and query j of $A(\sigma; 1, x + 1)$ evaluate the same group element.¹³ Based on this observation, it is possible to define an appropriate query mapping and complete the proof, since for non-adaptive algorithms we know in advance (independently of σ) if query i of $A'(\sigma; 1, x)$ and query j of $A'(\sigma; 1, x + 1)$ evaluate the same group element.

Additional Notation. We begin by defining additional notation. Given a query (i, j) , denote its evaluation on x as $(i, j)[x] = ix + j$. Thus, its oracle answer is $\sigma(ix + j)$. We denote by $Q(A(\sigma; 1, x))$ the query set of $A(\sigma; 1, x)$, excluding queries (i, j) for which $i = 0$ (which we call *constant queries*). Denote by $QE(A(\sigma; 1, x))$ the set of evaluations of all (non-constant) queries $Q(A(\sigma; 1, x))$.

We further denote by $Q(A)$ the set containing all of the potential (non-constant) queries of A on any input x and encoding σ . Note that for non-adaptive algorithms, $|Q(A)| \leq T$ and any adaptive algorithm A' can be simulated by a non-adaptive algorithm that makes $T' \stackrel{\text{def}}{=} |Q(A')|$ queries.

For the rest of this section, we focus on non-adaptive algorithms. For such algorithms, we can write $QE(A, x)$ (instead of $QE(A(\sigma; 1, x))$), as the query evaluations are independent of σ .

Restricted Queries. We examine pairs of executions $A(\sigma; 1, x)$ and $A(\sigma; 1, x + b)$ for some $b \in [-M, M]$. For such a pair, we define a (non-trivial) collision as the event that two queries issued by these executions (i, j) and (i', j') with $i \neq i'$

¹² Our actual proof is slightly more general than outlined here and uses the notion of query chains.

¹³ Our proof relaxes this strong condition, and requires that it holds unless a low-probability event (called a collision) occurs.

have the same evaluation. The actual evaluations depend on which algorithm issued the queries and there are 4 cases, e.g., $ix + j \bmod N = i'x + j' \bmod N$ if $A(\sigma; 1, x)$ issued both and $ix + j \bmod N = i'(x+b) + j' \bmod N$ if $A(\sigma; 1, x)$ issued (i, j) and $A(\sigma; 1, x+b)$ issued (i', j') , etc. In each of these 4 cases, both algorithms can exploit the collision to jointly solve the discrete logarithm problem using at most $2T$ queries (e.g., in the first case above, $x = (j' - j) \cdot (i - i')^{-1} \bmod N$). According to Theorem 3, the probability of this event is $O(T^2/N) = o(1)$ (by our assumption $T = o(N^{1/2})$), which is negligible. In the following we generally denote collision events by COL .

Most of the analysis below will be conditioned on the event \overline{COL} (whose probability is $1 - o(1)$), but we will omit this explicit conditioning for simplicity, while ignoring a negligible factor in the probability calculation.

Lemma 16. *Assume that $T = o(N^{1/2})$. Then, any non-adaptive algorithm A' with $\Pr_{\text{err}}(A', \sigma, [-M, M], T) = \delta$ can be transformed into a non-adaptive query-restricted algorithm A with $\Pr_{\text{err}}(A, \sigma, [-M, M], T) \leq \delta \cdot (1 + o(1))$ such that A only issues restricted queries of the form (i, j) with $i \in \{0, 1\}$.*

Query Disjoint Indices. We say that a non-adaptive DDL algorithm A has a query disjoint index b if $QE(A, x) \cap QE(A, x + b) = \emptyset$ for any $x \in \mathbb{Z}_N$. We note that A can have many query disjoint indices. We prove the following error probability lower bound on algorithms with a (small) query disjoint index.

Lemma 17. *Any non-adaptive algorithm which is query disjoint on index $b \geq 1$ satisfies $\Pr_{\text{err}}(A, \sigma, [-1, 1], T) = \Omega(1/b)$.*

Query Chains. Given a query $(1, j)$, we refer to the value j as a query offset. For a non-adaptive algorithm A , we define a query chain of length c as a sequence of $c + 1$ query offsets $j, j + 1, j + 2, \dots, j + c$ such that for each $k \in \{0, 1, \dots, c\}$, $(1, j + k) \in Q(A)$, while $(1, j + c + 1) \notin Q(A)$ and $(1, j - 1) \notin Q(A)$ (i.e., the sequence is maximal).

Denote the length of the longest query chain of A by $C(A)$.

Lemma 18. *Any non-adaptive query-restricted algorithm A satisfies*

$$\Pr_{\text{err}}(A, \sigma, [-1, 1], T) \geq \Omega(1/C(A)).$$

Proof (of Theorem 5). The theorem is a simple corollary of Lemmas 16 and 18. Given a non-adaptive algorithm A , transform it into a query-restricted algorithm A' using Lemma 16, with a multiplicative loss of $1 + o(1)$ in error probability. Clearly, $C(A') \leq T$, hence by Lemma 18 we have $\Pr_{\text{err}}(A, \sigma, [-1, 1], T) \geq \Pr_{\text{err}}(A', \sigma, [-1, 1], T) \cdot 1/(1 + o(1)) = \Omega(1/C(A')) = \Omega(1/T)$, concluding the proof. ■

A Generalization of Theorem 5. The theorem above does not completely render non-adaptive algorithms as inefficient since (for example) it does not rule out

the possibility that $\Pr_{\text{err}}(A, \sigma, [-T, T], T) = O(1/T)$ (which is optimal according to Theorem 4). However, the following theorem states that this is impossible and non-adaptive algorithms have a linear query-error tradeoff at best.

Theorem 6. *For all $1 \leq M \leq T$, any non-adaptive generic DDL algorithm A satisfies $\Pr_{\text{err}}(A, \sigma, [-M, M], T) = \Omega(M/T)$ given that $T = o(N^{1/2})$ and N is prime. In particular, for $M = T$, $\Pr_{\text{err}}(A, \sigma, [-T, T], T) = \Omega(1)$.*

Note that Theorem 5 is a special case of the one above, obtained for $M = 1$. The proof of Theorem 6 uses Fourier analysis and is given in the extended version of this paper.

Acknowledgements. The authors would like to thank Elette Boyle, Niv Gilboa, Yuval Ishai and Yehuda Lindell for discussions and helpful suggestions regarding this work.

This research was supported by the European Research Council under the ERC starting grant agreement no. 757731 (LightCrypt) and by the BIU Center for Research in Applied Cryptography and Cyber Security in conjunction with the Israel National Cyber Bureau in the Prime Minister's Office.

The first author was additionally supported by the Israeli Science Foundation through grant No. 573/16.

References

1. Ben-Or, M., Goldwasser, S., Wigderson, A.: Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In: Simon [19], pp. 1–10 (1988)
2. Bernstein, D.J., Lange, T.: Computing small discrete logarithms faster. In: Galbraith, S., Nandi, M. (eds.) INDOCRYPT 2012. LNCS, vol. 7668, pp. 317–338. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34931-7_19
3. Boneh, D., Goh, E.-J., Nissim, K.: Evaluating 2-DNF formulas on ciphertexts. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 325–341. Springer, Heidelberg (2005). https://doi.org/10.1007/978-3-540-30576-7_18
4. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Orrù, M.: Homomorphic secret sharing: optimizations and applications. In: Thuringham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, 30 October–03 November 2017, pp. 2105–2122. ACM (2017)
5. Boyle, E., Gilboa, N., Ishai, Y.: Breaking the circuit size barrier for secure computation under DDH. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part I. LNCS, vol. 9814, pp. 509–539. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53018-4_19
6. Boyle, E., Gilboa, N., Ishai, Y.: Group-based secure computation: optimizing rounds, communication, and computation. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017, Part II. LNCS, vol. 10211, pp. 163–193. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56614-6_6
7. Boyle, E., Gilboa, N., Ishai, Y., Lin, H., Tessaro, S.: Foundations of homomorphic secret sharing. In: Karlin, A.R. (ed.) 9th Innovations in Theoretical Computer Science Conference, ITCS 2018, Cambridge, MA, USA, 11–14 January 2018. LIPIcs, vol. 94, pp. 21:1–21:21. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2018)

8. Chaum, D., Crépeau, C., Damgård, I.: Multipartly unconditionally secure protocols (extended abstract). In: Simon [19], pp. 11–19 (1988)
9. Chor, B., Kushilevitz, E., Goldreich, O., Sudan, M.: Private information retrieval. *J. ACM* **45**(6), 965–981 (1998)
10. Fazio, N., Gennaro, R., Jafarikhah, T., Skeith III, W.E.: Homomorphic secret sharing from Paillier encryption. In: Okamoto, T., Yu, Y., Au, M.H., Li, Y. (eds.) *ProvSec 2017*. LNCS, vol. 10592, pp. 381–399. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-68637-0_23
11. Galbraith, S.D., Pollard, J.M., Ruprai, R.S.: Computing discrete logarithms in an interval. *Math. Comput.* **82**(282), 1181–1195 (2013)
12. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: Mitzenmacher, M. (ed.) *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009*, Bethesda, MD, USA, 31 May–2 June 2009, pp. 169–178. ACM (2009)
13. Gordon, D.M.: Discrete logarithms in $GF(P)$ using the number field sieve. *SIAM J. Discret. Math.* **6**(1), 124–138 (1993)
14. Lenstra, A.K., Lenstra, H.W. (eds.): *The Development of the Number Field Sieve*. LNM, vol. 1554. Springer, Heidelberg (1993). <https://doi.org/10.1007/BFb0091534>
15. Pollard, J.M.: Monte Carlo methods for index computation (mod p). *Math. Comput.* **32**(143), 918–924 (1978)
16. Pollard, J.M.: Kangaroos, monopoly and discrete logarithms. *J. Cryptol.* **13**(4), 437–447 (2000)
17. Rivest, R.L., Adleman, L., Dertouzos, M.L.: On data banks and privacy homomorphisms. In: *Foundations of Secure Computation*, pp. 169–179 (1978)
18. Shoup, V.: Lower bounds for discrete logarithms and related problems. In: Fumy, W. (ed.) *EUROCRYPT 1997*. LNCS, vol. 1233, pp. 256–266. Springer, Heidelberg (1997). https://doi.org/10.1007/3-540-69053-0_18
19. Simon, J. (ed.): *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, Chicago, Illinois, USA, 2–4 May 1988. ACM (1988)
20. Yao, A.C.: Protocols for secure computations (extended abstract). In: *23rd Annual Symposium on Foundations of Computer Science*, Chicago, Illinois, USA, 3–5 November 1982, pp. 160–164. IEEE Computer Society (1982)