# Tight Time-Space Tradeoffs for the Decisional Diffie-Hellman Problem 

Akshima ${ }^{1}$, Tyler Besselman ${ }^{1}$, Siyao Guo* ${ }^{*}$, Zhiye Xie ${ }^{1}$, and Yuping Ye ${ }^{2,1}$<br>${ }^{1}$ NYU Shanghai ${ }^{\dagger}$<br>${ }^{1}\{$ akshima, tyler.william.b, siyao.guo, zhiye.xie, yy3813\}@nyu.edu<br>${ }^{2}$ East China Normal University


#### Abstract

In the (preprocessing) Decisional Diffie-Hellman (DDH) problem, we are given a cyclic group $G$ with a generator $g$ and a prime order $N$, and we want to prepare some advice of size $S$, such that we can efficiently distinguish $\left(g^{x}, g^{y}, g^{x y}\right)$ from $\left(g^{x}, g^{y}, g^{z}\right)$ in time $T$ for uniformly and independently chosen $x, y, z$ from $\mathbb{Z}_{N}$. This is a central cryptographic problem whose computational hardness underpins many widely deployed schemes, such as the Diffie-Hellman key exchange protocol.

We prove that any generic preprocessing DDH algorithm (operating in any cyclic group) achieves advantage at most $O\left(S T^{2} / N\right)$. This bound matches the best known attack up to poly-log factors, and confirms that DDH is as secure as the (seemingly harder) discrete logarithm problem against preprocessing attacks. Our result resolves an open question by Corrigan-Gibbs and Kogan (EUROCRYPT 2018), who proved optimal bounds for many variants of discrete logarithm problems except DDH (with an $\tilde{O}\left(\sqrt{S T^{2} / N}\right)$ bound $)$.

We obtain our results by adopting and refining the approach by Gravin, Guo, Kwok, Lu (SODA 2021) and by Yun (EUROCRYPT 2015). Along the way, we significantly simplified and extended the above techniques which may be of independent interest. The highlights of our techniques are as follows:


- We obtain a simpler reduction from decisional problems against $S$-bit advice to their $S$-wise XOR lemmas against zero-advice, recovering the reduction by Gravin, Guo, Kwok and Lu (SODA 2021).
- We show how to reduce generic hardness of decisional problems to their variants in the simpler hyperplane query model proposed by Yun (EUROCRYPT 2015). This is the first work analyzing a decisional problem in Yun's model, answering an open problem proposed by Auerbach, Hoffman, and Pascual-Perez (TCC 2023).
- We prove an $S$-wise XOR lemma of DDH in Yun's model. As a corollary, we obtain the generic hardness of the $S$-XOR DDH problem.

[^0]
## 1 Introduction and Overview

The Decisional Diffie-Hellman (DDH) problem is a classic problem in cryptography. The assumption that it is a hard problem forms the basis for the security claims of the Diffie-Hellman key exchange protocol, which was introduced in a 1976 paper by Diffie and Hellman DH76, and is still used to this day to secure a large chunk of communication that takes place both over the public internet and between private parties. As such, strong theoretical bounds on algorithms attacking this problem are quintessential for guaranteeing the security of its users.

The DDH problem may be roughly summarized as follows: given a cyclic group $G$ with a generator $g$ and (usually prime) order $N$, sample three random integers $x, y, z$ from $\mathbb{Z}_{N}$. Then, an attacker is tasked with determining, given the values $\left(g^{x}, g^{y}\right)$, whether a third, unknown element $h \in G$ is equal to the value $g^{x y}$ or $g^{z}$, chosen at random in each instance.

When this problem was introduced, there was only an assumption that it was difficult to solve. It is clear that an algorithm able to solve for the discrete logarithm of a given group element would be able to differentiate by computing $g^{x y}$ and comparing it to the value $h$ after finding the discrete logarithm value $x=\log _{G}\left(g^{x}\right)$ of the second input element. But is there a more efficient way? It was later shown by Shoup in Sho97] that no algorithm solving the Diffie-Hellman problem for a generic group $G$ can be more efficient than an algorithm solving for the discrete logarithm of a random element $h=g^{x} \in G$. Thus, in this simple case, solving the Diffie-Hellman problem is as difficult as solving the discrete logarithm problem (without relying on special techniques exploiting the structure of a specific group or class of groups $G$ ).

In the modern age of computation, however, a security bound for this simple setting is not enough. It is not unreasonable for strong adversaries, with lots of computation power, to invest their resources in a one-time large pre-computation, which can significantly reduce the time required to solve an instance of a problem later. In some cases, there can be time limits to solve an instance of a problem, and using the pre-computation may allow the adversary to solve the problem within the given time frame. For instance, while it may be hard to invert a cryptographic hash function, an attacker with enough computation power and time may pre-compute rainbow tables in advance so as to more efficiently invert a single hash value given at a later point in time. In this case, we assume the time for pre-computations is effectively unlimited, and we concern ourselves with the size of the pre-computed input required to find inversions efficiently. Specifically, it becomes interesting to analyze the trade-off between the space used in pre-computation and the amount of time taken when presented with an actual instance.

While auxiliary input was originally defined for the random oracle model in Unr07, DGK17, it was later extended to the Generic Group Model (GGM) by Corrigan-Gibbs and Kogan in CK18 and by Coretti et al. in CDG18. An adversary $\mathcal{A}$ in the auxiliary input(AI) idealized oracle (random oracle/random permutation/ generic group) model can be thought of as a two-stage algorithm $\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$ :

- In the first stage, which we will refer to as the "offline" stage or the "pre-computation" stage, $\mathcal{A}_{0}$ gets unbounded access to the oracle but has to output bounded advice $\gamma$ about the oracle.
- In the second stage, which we will refer to as the "online" stage, $\mathcal{A}_{1}$ gets the advice $\gamma$ from $\mathcal{A}_{0}$ and a challenge about the oracle as input. $\mathcal{A}_{1}$ makes a bounded number of queries to the oracle.

In the Generic Group Model, a group of order $N$ is described in terms of a random injective labeling function $\sigma$ mapping $\mathbb{Z}_{N}$ to a set of labels $\mathcal{L}$ such that $\sigma(1), \ldots, \sigma(N-1)$ correspond to the elements of the group with discrete logarithms $1, \ldots, N-1$ respectively. Group operations are performed by querying a generic group oracle $\mathcal{O}_{\sigma}(\cdot, \cdot)$ corresponding to the labeling function $\sigma$ such that for any $\mathrm{X}, \mathrm{Y} \in \mathcal{L}, \mathcal{O}_{\sigma}(\mathrm{X}, \mathrm{Y})$ outputs

- $\sigma(x+y)$ if there exists $x, y \in \mathbb{Z}_{N}$ such that $\sigma(x)=\mathrm{X}$ and $\sigma(y)=\mathrm{Y}$
- $\perp$ otherwise.

For any adversary $\mathcal{A}=\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$ in AI-GGM, $\mathcal{A}_{0}$ gets input $\sigma(1)$ and unbounded access to the oracle $\mathcal{O}_{\sigma}$. In effect, then, $\mathcal{A}_{0}$ gets access to the entire function table of $\sigma$. However, $\mathcal{A}_{0}$ can output only bounded length
advice, say some $\gamma$ where $|\gamma|=S$ bits. Then $\mathcal{A}_{1}$ takes $\gamma$ and a challenge as input, makes a bounded number of queries, say $T$, to $\mathcal{O}_{\sigma}$ and computes an output.

Note that the challenge and the output of $\mathcal{A}_{1}$ depend on the problem $\mathcal{A}$ is solving. For an algorithm $\mathcal{A}$ solving the discrete logarithm problem, $\mathcal{A}_{1}$ gets $(\sigma(1), \sigma(x))$ as challenge for a random $x \in \mathbb{Z}_{N}$. To "win", $\mathcal{A}_{1}$ needs to output $x$. For an algorithm $\mathcal{A}$ solving the DDH problem, $\mathcal{A}_{1}$ gets a tuple $(\sigma(1), \sigma(x), \sigma(y), \sigma(u))$ as challenge where for random and independent $x, y, z \in \mathbb{Z}_{N}$ and random $b \in\{0,1\}, u:=(1-b) x y+b z$, and $\mathcal{A}_{1}$ outputs a guess, say $\hat{b}$, for $b$ and "wins" if $\hat{b}=b$.

### 1.1 Our Results

In this work, we analyze the Decisional Diffie-Hellman ( DDH ) problem in the AI-GGM setting. We try to understand the following aspects of the DDH problem: for any algorithm with bounded space given access to the cyclic group $G$, what is the lower bound on the number of group operations required to solve a given instance of the problem? Is the DDH problem in this setting still asymptotically as difficult as the discrete logarithm problem? We successfully obtained answers to both questions, positively affirming the latter.

Tight time-space lower bounds for the DDH Problem. The main result of this work is the tight lower bound $S T^{2}=\Omega(\varepsilon N)$ for any generic algorithm with $S$-bit advice/auxiliary input, making $T$ oracle queries on a group of order $N$ and solving the DDH problem in AI-GGM with probability at least $1 / 2+\varepsilon$. Formally,

Theorem 1. Let $N$ be a prime and $S, T \geq 1$ be integers. For any generic algorithm with preprocessing $\mathcal{A}=\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$ over $\mathcal{L}$ for $\mathbb{Z}_{N}$ such that $\mathcal{A}_{0}$ outputs an $S$-bit advice string, and $\mathcal{A}_{1}$ makes at most $T$ oracle queries, it holds that

$$
\mathbb{P}\left[\mathcal{A}_{1}^{\mathcal{O}_{\sigma}}\left(\mathcal{A}_{0}^{\mathcal{O}_{\sigma}}(\sigma(1)), \sigma(1), \sigma(x), \sigma(y), \sigma(u)\right)=b\right] \leq \frac{1}{2}+2^{11} \frac{S T^{2}}{N}
$$

where $x, y, z \leftarrow \mathbb{Z}_{N}, b \leftarrow\{0,1\}, u=b z+(1-b) x y$, and $\sigma: \mathbb{Z}_{N} \rightarrow \mathcal{L}$ is a random injective function.
There is a corresponding algorithm for DDH with advantage $\tilde{\Omega}\left(S T^{2} / N\right)$ based on known algorithms for the discrete logarithm problem with an asymptotically equivalent advantage in the preprocessing setting (see Appendix A). Thus, the above security bound is tight for all values $S, T, N$ up to poly-log factors.

Prior to this work, the best-known security bound in this setting was $\widetilde{O}\left(\sqrt{S T^{2} / N}\right)$ in CK19, leaving a quadratic gap between security upper and lower bounds. Our result closes this gap and confirms that DDH is as secure against preprocessing attacks as the (seemingly more difficult) discrete logarithm problem.

Generic hardness of the $k$-XOR DDH Problem. To prove this theorem, we consider a fundamental variant of DDH , called the $k-X O R$ DDH problem. Let's define the DDH bit of a given DDH instance as a bit indicating whether the third coordinate is $x y$ or $z$. The $k$-XOR DDH problem can be described as follows: given $k$ independently sampled DDH instances, compute the XORed DDH bits of the $k$ DDH instances.

We prove the lower bound $k T^{2}=\Omega\left(\varepsilon^{1 / k} N\right)$ for any generic algorithm (without preprocessing) with $k T$ oracle queries on a group of order $N$ that solves the $k$-XOR DDH problem in GGM with probability at least $1 / 2+\varepsilon$.

Theorem 2. Let $N$ be a prime and $k, T \geq 1$ be integers. For any generic algorithm $\mathcal{B}$ over $\mathcal{L}$ for $\mathbb{Z}_{N}$ such that $\mathcal{B}$ makes at most $k T$ oracle queries, it holds that

$$
\mathbb{P}\left[\mathcal{B}^{\mathcal{O}_{\sigma}}(\sigma(1), \sigma(\mathbf{w}))=\bigoplus_{i=1}^{k} b_{i}\right] \leq \frac{1}{2}+\frac{1}{2}\left(2^{11} \frac{k T^{2}}{N}\right)^{k}
$$

where $\sigma(\mathbf{w}):=\left\{\sigma\left(x_{i}\right), \sigma\left(y_{i}\right), \sigma\left(u_{i}\right)\right\}_{i=1}^{k}$ for $u_{i}=b_{i} z_{i}+\left(1-b_{i}\right) x_{i} y_{i},\left\{x_{i}, y_{i}, z_{i}\right\}_{i=1}^{k} \leftarrow \mathbb{Z}_{N}^{3 k},\left\{b_{1}, \ldots, b_{k}\right\} \leftarrow$ $\{0,1\}^{k}$, and $\sigma: \mathbb{Z}_{N} \rightarrow \mathcal{L}$ a uniformly random injective function.

Notice that when $k=1$, the security bound from the theorem above is exactly the lower bound for the standard DDH problem, which was proven tight in [Sho97]. For general $k$, our bound confirms that XOR amplifies the hardness of DDH .

Yun Yun15 considered the generic hardness of solving multiple discrete logarithm problems simultaneously and showed a similar lower bound $k T^{2}=\Omega\left(\varepsilon^{1 / k} N\right)$ for solving $k$-discrete logarithm problems using $k T$ queries with success probability at least $\varepsilon$. Our result suggests that the (seemingly easier) $k$-XOR DDH admits similar hardness to the $k$-discrete logarithm problem.

### 1.2 Our Proofs

To obtain this result, we adopt and further improve the techniques from GGKL21 and Yun15. Our simplification and refinement of these techniques not only allow us to obtain the desired time-space lower bounds for the DDH problem but also characterize the potential of these techniques, which could be of independent interest. The proof of our main results will be split into three stages, each corresponding to one of the three lemmas in Section 3

From preprocessing DDH to $k$-XOR DDH. Once a single bit of advice about $\sigma$ is allowed, standard analyzing techniques for the generic group model (such as the commonly used lazy sampling technique) may become inapplicable. One natural idea is reducing preprocessing problems to their variants without preprocessing so we can apply standard techniques.

Thus, in Lemma 1 (see Section 3.1 for details), we reduce the DDH problem with $S$-bit advice to the $S$-XOR DDH problem without advice. In particular, we show that, for any even $S$, a preprocessing adversary $\mathcal{A}$ for DDH with advantage $\epsilon$, by simply guessing the advice uniformly and randomly, and then running its online algorithm $S$ times, we obtain an adversary $\mathcal{B}$ for the $S$-XOR DDH problem with advantage at least $(\epsilon / 2)^{S}$. The restriction on $S$ being even is minor because we can reduce an odd $S$ to an even $S$ before using Lemma 1. Intuitively, if $\sigma$ is fixed, guessing the $S$-bit advice correctly happens with probability $1 / 2^{S}$, and conditioning on that, the advantage of solving $S$-XOR DDH is $\epsilon^{S}$ if $\mathcal{A}$ has advantage $\epsilon$ (for this fixed $\sigma$ ). Thus, the advantage of $\mathcal{B}$ is at least $\epsilon^{S} / 2^{S}$. To obtain the formal argument, we use the fact that $S$ is an even number to argue that the advantage is at least zero when conditioning on any wrong guess, and we apply Jensen's inequality to take care of the distributional case of $\sigma$.

Reducing decisional problems with $S$-bit advice to their $S$-wise XOR variants without advice was first proposed by Gravin et al. GGKL21 in the context of proving tight time-space tradeoffs for the hardcore bit problem. The reduction in GGKL21] requires proving a concentration bound for distributions satisfying a related $k$-wise XOR condition. Our reduction, which provides an algorithm for the $k$-XOR variant by guessing the advice based on the preprocessing algorithm, is considerably simpler and more constructive. The idea of guessing the advice is inspired by the reduction for reducing search problems with $S$-bit advice to their $S$-wise multi-instance variants without advice in CGLQ20 (and its refinement AGL22). Our reduction can be viewed as an analog to their reductions for decisional problems.

From Generic Group Model to Hyperplane Query Model. By Lemma 1 we can shift our attention to the $k$-XOR DDH problem in the generic group model. The encoding $\sigma$ and the oracle $\mathcal{O}_{\sigma}$ are slightly complex to work with for problems involving multiple instances. Inspired by Yun Yun15]'s approach for the multiple discrete logarithm problem, we aim to remove the encoding $\sigma$ and work with a simpler oracle.

Thus, in Lemma 2 (see Section 3.2 for details), we reduce the $k$-XOR DDH problem in the generic group model to a corresponding $k$-XOR DDH problem in the Hyperplane Query Model (HQM) introduced by Yun Yun15. In particular, we show that any algorithm with input $(\sigma(1), \sigma(\mathbf{w}))$, where $\mathbf{w}:=\left\{x_{i}, y_{i}, u_{i}\right\}_{i=1}^{k}$ and with oracle access to $\mathcal{O}_{\sigma}$, can be simulated by an algorithm with only a few hyperplane queries to $H_{\mathbf{w}}$, and incurs a quadratic blow-up in the query complexity. The hyperplane query simply submits a linear equation of the form $L(\cdot)=c$, and the hyperplane query oracle $\mathcal{H}_{\mathbf{w}}$ outputs 1 if $L(\mathbf{w})=c$, otherwise it outputs 0 .

Our proof is largely based on Yun's proof, which utilizes the lazy sampling property of $\sigma$ - i.e., new input-output pairs can be generated on the fly with limited information from previously generated pairs. To record the generated input-output pairs, the simulation maintains a list $\left\{\left(L_{i}, s_{i}\right)\right\}_{i \geq 0}$, where $L_{i}$ is a multi-linear polynomial in variables $\left\{X_{1}, Y_{1}, U_{1}, \ldots, X_{k}, Y_{k}, U_{k}\right\}$ such that $\sigma\left(L_{i}(\mathbf{w})\right)=s_{i}$. To simulate $\sigma(\cdot)$ and $\sigma^{-1}(\cdot)$ (and thus $\mathcal{O}_{\sigma}$ ), we demonstrate how to search an existing pair and how to sample a new pair in the maintained list using a few hyperplane queries to $\mathbf{w}$.

One of the main differences between our proof and Yun's proof is that Yun's reduction assumes that the adversary in the generic group model never queries the oracle on labels that have not appeared previously. Our proof makes no such assumptions and shows how to sample new preimages using only a few hyperplane queries in total (see the subroutine FindPreimage in Section 3.2).
$k$-XOR DDH in the Hyperplane Query Model. In Lemma 3, we show that any algorithm with $q$ queries for the $k$-XOR DDH problem has advantage $\approx(q / N K)^{k}$. Plugging in $q \approx(k T)^{2}$, we obtain that the generic hardness of $k$-XOR DDH against $k T$-query adversaries is $\left(k T^{2} / N\right)^{k}$, and the generic hardness of preprocessing DDH against $S$-bit advice and $T$-query algorithm is $\approx S T^{2} / N$. See Section 3.4 for details.

Our proof proceeds by induction on $(k, q)$. In particular, in Claim 2 and Claim 3 , we bound the best advantage of $q$-query algorithm for $k$-XOR DDH problem by related best advantages achieved by ( $q-1$ )-query algorithms for $k$-XOR problem or $q$-query algorithm for $(k-1)$-XOR DDH problem. Then by carefully choosing an inductive hypothesis $\binom{q+k}{k}(4 / N)^{k}$ and an inductive step, we obtain the desired bound.

Given the best $q$-query algorithm for the $k$-XOR DDH problem, it is helpful to view it as a decision tree of depth $q$ such that each internal node is a hyperplane query, and its subtrees represent the sub-algorithms executed according to the answer from the hyperplane query oracle. Let $H$ denote the first hyperplane query. We will bound the advantage of the algorithm based on whether $H$ is satisfied or not.

Intuitively, if $H$ is not satisfied, this query should be somewhat useless because it provides very little information about the variables. So the contribution of this case should be roughly the same as that of its $(q-1)$ sub-algorithm for the $k$-XOR DDH problem. We capture this intuition in Claim 2 . In particular, we show that, at the cost of an additional item of the best advantage assuming that $H$ holds (which we will handle later anyway), we can bound the advantage when $H$ is not satisfied by the advantage of a ( $q-1$ )-query algorithm for the $k$-XOR DDH problem.

If $H$ is satisfied, the intuition is that this happens with a probability of only $O(1 / N)$. Conditioning on $H$ happens, the sub-algorithm executed can be converted into either a $(q-1)$ algorithm for the $(k-1)$-XOR problem or a $q$ algorithm for the $(k-1)$-XOR problem where the first equation is satisfied. See Claim 3 for the details.

### 1.3 Prior and Other Related Works

Several works BL13, LCH11, Mih10] have unveiled the ability of generic algorithms with preprocessing to tackle the discrete logarithm problem. These works have shown that generic algorithms for every group of order $N$ use $N^{1 / 3}$ bits of group-specific advice and approximately $N^{1 / 3}$ online time to solve the discrete logarithm problem. This beats Shoup's bound without preprocessing of $\Omega(\sqrt{N})$ group operations for a group of order $N$ in Sho97. Inspired by these preprocessing attacks, Corrigan-Gibbs and Kogan CK18 were the first to define and analyze generic algorithms with precomputation for the discrete logarithm, DDH, Computational Diffie Hellman and multiple discrete logarithm problems. Corrigan-Gibbs and Kogan showed the optimal generic hardness $S T^{2}=\widetilde{\Omega}(\varepsilon N)$ for the discrete logarithm problem in the auxiliary input model, where the advice is $S$-bits, the algorithm makes $T$ queries in the online stage, and the success probability is $\varepsilon$ for a group of (prime) size $N$. Later, Bartusek, Ma, and Zhandry [BMZ19] found that the fixed-generator and the random-generator variants of the discrete logarithm problem are not equally hard and proved the security bound of the random-generator variant of the discrete logarithm problem is $S T^{2}=\widetilde{\Theta}(\sqrt{\varepsilon} N)$ using the presampling technique from CDG18.

Corrigan-Gibbs and Kogan also studied the multiple discrete logarithm problem (where the algorithm is tasked with solving multiple instances of the discrete logarithm problem simultaneously) for preprocessing
generic algorithms (with $S$-bit advice and $T$ online queries) and proved a security bound for solving $M$ instances of the discrete logarithm with a probability of at least $\varepsilon$ of $S T^{2}=\Omega\left(\varepsilon^{1 / M} N M\right)$. They were the first to study the multiple discrete logarithm problem with auxiliary input. Several works [FJM14, KS01, Yun15] prior to CK18 studied both upper and lower bounds for the problem without preprocessing. One work of particular importance is Yun15. To show the generic hardness of solving the multiple discrete logarithm problem, Yun introduced a related computational model called the search-by-hyperplane-queries model and reduced the security of the multiple discrete logarithm problems in this model.

For the DDH problem in the auxiliary input model, Corrigan-Gibbs and Kogan showed the bound $S T^{2}=\widetilde{\Omega}\left(\varepsilon^{2} N\right)$ for any generic algorithm solving the problem with a probability of at least $1 / 2+\varepsilon$. There is a gap between this Corrigan-Gibbs and Kogan bound and the best-known preprocessing attack for the DDH problem. At the same time, they propose a new attack for square DDH (distinguishing $g, g^{x}, g^{x^{2}}$ from $g, g^{x}, g^{y}$ for random $x, y \sim \mathbb{Z}_{N}$ ) showing that, unlike in the non-preprocessing setting, square DDH is less secure than the discrete logarithm problem. Their result motivates the question of whether DDH is as hard as the discrete logarithm problem in this setting. Corrigan-Gibbs and Kogan's results are based on the incompressibility argument and a non-trivial batching technique. In the same year, Coretti et al. in CDG18, using their refined presampling techniques (first developed by Unruh Unr07 for the Random Oracle model) proved the same bound. In the same work, Coretti et al. showed hardness bounds for several other security applications in the Auxiliary Input Generic Group model, including the discrete logarithm problem and the DDH problem, using their presampling technique.

Gravin et al. GGKL21, inspired by the elegant and short application of union-bounds in Impagliazzo's proof in Imp11, reduced the security for algorithms with $S$-bit advice to showing a roughly $2^{-S}$ concentration bound on the advantage of algorithms with no advice. Thus, they derived concentration bounds for the $k$-wise XOR version of the problem. For a decision problem $G$, if no algorithm with zero advice and $k T$ queries can solve the $k$-wise XOR version of problem $G$ with a probability of more than $1 / 2+\varepsilon^{k}$, then for any $k \leq S+\log N$, no algorithm with $S$-bit advice making $T$ online queries can solve $G$ with a probability of more than $\varepsilon^{\prime}$ where $\varepsilon^{\prime}=2 \varepsilon+O(\sqrt{S / N})$.

We note several previous works Hel80, FN91, DGK17, CDG18, CDGS18, CK19, CHM20, ACDW20, CGLQ20, Liu23, AGL22, GK22, FGK22, FGK23, GGPS23, ADGL23, GLLZ21 that have explored the time-space lower bounds for various cryptographic primitives and applications in idealized models.

## 2 Preliminaries

In this paper, we write $\mathbb{Z}_{N}$ for $N \geq 1$ to denote the ring of integers modulo $N$. Let $[n]=\{1, \ldots, n\}$ for any integer $n \geq 1$. We use the notation $\widetilde{O}, \widetilde{\Omega}$, and $\widetilde{\Theta}$ to denote asymptotic bounds up to a poly-log factor. An algorithm $\mathcal{A}$ taking an input the value $x$ is denoted $\mathcal{A}(x)$. Given a set $S$, let $x \leftarrow S$ denote that $x$ is sampled randomly from the uniform distribution over $S$.

Next, we formally define the idealized models used to prove our results.
Definition (Generic Group Model (GGM) [Sho97, CK18]). For any cyclic group $G$ of order $N$ with generator $g$, let $\sigma: \mathbb{Z}_{N} \rightarrow \mathcal{L}$ be a random injective"labeling" function, where $\mathcal{L}$ is an arbitrary set of "labels" for encoding the set of group elements. We identify group elements with their images $\left\{g^{i}\right\}_{i=0}^{N-1} \cong\{\sigma(i)\}_{i=0}^{N-1}$. As such, we will say that an element $i \in \mathbb{Z}_{N}$ is the discrete logarithm of the label $\sigma(i)$. The generic group oracle $\mathcal{O}_{\sigma}: \mathcal{L} \times \mathcal{L} \rightarrow \mathcal{L}$ is a function that takes two elements $x, y \in \mathcal{L}$ as input and response as follows:

- If there exists $i, j \in \mathbb{Z}_{N}$ such that $\sigma(i)=x$ and $\sigma(j)=y$ respectively, then the oracle responds with $\sigma(i+j)$, where addition is modulo the group order $N$.
- If there exists no such $i$ or $j$, then the oracle responds with $\perp$.

Definition (Generic Algorithm). A generic algorithm over a label set $\mathcal{L}$ for $\mathbb{Z}_{N}$ is a probabilistic algorithm $\mathcal{A}$ which has access to the generic group oracle $\mathcal{O}_{\sigma}$ and takes as input a list of labels $\left\{\sigma\left(x_{i}\right)\right\}_{i=1}^{\ell}$. Furthermore, for such an algorithm $\mathcal{A}$, we define the query complexity of $\mathcal{A}$ to be the number of calls the algorithm makes to the oracle $\mathcal{O}_{\sigma}$.

Definition (Generic Algorithm with Preprocessing). A generic algorithm with preprocessing over a label set $\mathcal{L}$ for $\mathbb{Z}_{N}$ consists of a pair of generic group algorithms $\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$ :

- Offline stage: The algorithm $\mathcal{A}_{0}$ takes $\sigma(1)$ as input, can make an unlimited number of queries to the oracle $\mathcal{O}_{\sigma}$, and computes an advice string $s_{\sigma} \in\{0,1\}^{*}$.
- Online stage: The algorithm $\mathcal{A}_{1}$ takes an advice string $s_{\sigma}$ output in the offline stage and a list of labels $\left\{\sigma\left(x_{i}\right)\right\}_{i=1}^{\ell}$ as input, makes some number of queries to oracle $\mathcal{O}_{\sigma}$, and outputs some result. The query complexity of the online algorithm $\mathcal{A}_{1}$ is called the online query complexity.

In particular, we focus on generic algorithms with preprocessing where the advice string $s_{\sigma}$ is limited to $S$-bits, i.e. $s_{\sigma} \in\{0,1\}^{S}$ and with online query complexity $T$. Here $S, T$ may be functions of the group order $N$.

Definition (Hyperplane Query Model (HQM) Yun15). Let $N$ be a prime number and $\mathbb{Z}_{N}^{\ell}$ be the $\ell$ dimensional affine space over finite field $\mathbb{Z}_{N}$. Let $X_{1}, \ldots, X_{\ell}$ be the canonical coordinates of $\mathbb{Z}_{N}^{\ell}$. Let $\Lambda$ be the set of all affine hyperplanes in $\mathbb{Z}_{N}^{\ell}$. Then an affine hyperplane $H \in \Lambda$ can be described as the set of points in $\mathbb{Z}_{N}^{\ell}$ satisfying linear equation $\sum_{i=1}^{\ell} \alpha_{i} X_{i}=\alpha_{0}$ for some $\left\{\alpha_{i}\right\}_{i=0}^{\ell} \subseteq \mathbb{Z}_{N}$. Let $\mathbf{w} \in \mathbb{Z}_{N}^{\ell}$ be a point in the affine space. A hyperplane query oracle $\mathcal{H}_{\mathbf{w}}: \Lambda \rightarrow\{0,1\}$ is a function that takes as input a hyperplane $H$ and response as follows:

$$
\mathcal{H}_{\mathbf{w}}(H)= \begin{cases}1, & \text { if } \mathbf{w} \in H \\ 0, & \text { if } \mathbf{w} \notin H\end{cases}
$$

Definition (Hyperplane Query Algorithm). A hyperplane query algorithm for $\mathbb{Z}_{N}^{\ell}$ is a probabilistic algorithm $\mathcal{B}$ that has access to the oracle $\mathcal{H}_{\mathbf{w}}$. The hyperplane query complexity of such an algorithm is the number of calls made by the algorithm to the hyperplane query oracle $\mathcal{H}_{\mathbf{w}}$.

## 3 Tight Time-Space Tradeoffs for DDH

### 3.1 Reducing Preprocessing DDH to $k$-XOR DDH Problem

In this section, we show that to prove a security bound on the preprocessing DDH problem, it suffices to prove a security bound on the $k$-XOR DDH problem.

Lemma 1. Let $N$ be a prime and $S \geq 2$ be an even integer. For any generic algorithm with preprocessing $\mathcal{A}=\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$ over $\mathcal{L}$ for $\mathbb{Z}_{N}$ such that $\mathcal{A}_{0}$ outputs an $S$-bit advice string, $\mathcal{A}_{1}$ makes at most $T$ oracle queries, and

$$
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{A}}:=2 \mathbb{P}\left[\mathcal{A}_{1}^{\mathcal{O}_{\sigma}}\left(\mathcal{A}_{0}^{\mathcal{O}_{\sigma}}(\sigma(1)), \sigma(1), \sigma(x), \sigma(y), \sigma(u)\right)=b\right]-1
$$

where $x, y, z \leftarrow \mathbb{Z}_{N}, b \leftarrow\{0,1\}, u=b z+(1-b) x y$, and $\sigma: \mathbb{Z}_{N} \rightarrow \mathcal{L}$ is a random injective function, there exists a generic algorithm $\mathcal{B}$ which makes at most $S T$ queries, such that for

$$
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{B}}:=2 \mathbb{P}\left[\mathcal{B}^{\mathcal{O}_{\sigma}}\left(\sigma(1),\left\{\sigma\left(x_{i}\right), \sigma\left(y_{i}\right), \sigma\left(u_{i}\right)\right\}_{i=1}^{S}\right)=\bigoplus_{i=1}^{S} b_{i}\right]-1
$$

where $\left\{x_{i}, y_{i}, z_{i}\right\}_{i=1}^{S} \leftarrow \mathbb{Z}_{N}^{3 S},\left\{b_{1}, \ldots, b_{S}\right\} \leftarrow\{0,1\}^{S}, u_{i}=b_{i} z_{i}+\left(1-b_{i}\right) x_{i} y_{i}$ and $\sigma: \mathbb{Z}_{N} \rightarrow \mathcal{L}$ is a random injective function, it holds that

$$
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{B}} \geq \frac{\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{A}}^{S}}{2^{S}}
$$

Proof. Without loss of generality, we assume that $\mathcal{A}$ is deterministic otherwise we fix the best randomness of $\mathcal{A}$. Consider the following algorithm $\mathcal{B}$ based on $\mathcal{A}$ : sample a random advice string $s \leftarrow\{0,1\}^{S}$, then for all $i \in[S]$, compute

$$
\hat{b}_{i}^{(\sigma, s)}:=\mathcal{A}_{1}^{\mathcal{O}_{\sigma}}\left(s, \sigma(1), \sigma\left(x_{i}\right), \sigma\left(y_{i}\right), \sigma\left(u_{i}\right)\right)
$$

and output a single bit $\hat{b}:=\bigoplus_{i=1}^{S} \hat{b}_{i}^{(\sigma, s)}$. It is clear that $\mathcal{B}$ makes at most $S T$ queries. Now, we show that such an algorithm has the desired advantage. Fix a labeling function $\sigma$ and some advice string $a \in\{0,1\}^{S}$. The advantage of $\mathcal{A}$ and $\mathcal{B}$ on $a$ are

$$
\begin{gathered}
\varepsilon_{\mathcal{A}}(\sigma, a):=2 \mathbb{P}\left[\hat{b}^{(\sigma, a)}=b\right]-1=\mathbb{E}\left[(-1)^{\hat{b}^{(\sigma, a)} \oplus b}\right], \\
\varepsilon_{\mathcal{B}}(\sigma, a):=2 \mathbb{P}\left[\bigoplus_{i=1}^{S} \hat{b}_{i}^{(\sigma, a)}=\bigoplus_{i=1}^{S} b_{i}\right]-1=\mathbb{E}\left[(-1)^{\oplus_{i=1}^{S} b_{i} \oplus \hat{b}_{i}^{(\sigma, a)}}\right] .
\end{gathered}
$$

Because $\sigma, a$ are fixed and $\left\{x_{i}, y_{i}, u_{i}\right\}$ are independently sampled, we can derive that

$$
\varepsilon_{\mathcal{B}}(\sigma, a)=\mathbb{E}\left[\Pi_{i=1}^{S}(-1)^{\hat{b}_{i}^{(\sigma, a)} \oplus b_{i}}\right]=\Pi_{i=1}^{S} \mathbb{E}\left[(-1)^{\hat{b}_{i}^{(\sigma, a)} \oplus b_{i}}\right]=\varepsilon_{\mathcal{A}}(\sigma, a)^{S} .
$$

For a fixed $\sigma$, let $s_{\sigma}:=\mathcal{A}_{0}^{\mathcal{O}_{\sigma}}(\sigma(1))$. Note that

$$
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{B}}=\mathbb{E}_{\sigma, a}\left[\varepsilon_{\mathcal{B}}(\sigma, a)\right]=\mathbb{E}_{\sigma, a}\left[\varepsilon_{\mathcal{A}}(\sigma, a)^{S}\right] \geq \frac{1}{2^{S}} \mathbb{E}_{\sigma}\left[\varepsilon_{\mathcal{A}}\left(\sigma, s_{\sigma}\right)^{S}\right] \geq \frac{1}{2^{S}} \mathbb{E}_{\sigma}\left[\varepsilon_{\mathcal{A}}\left(\sigma, s_{\sigma}\right)\right]^{S}=\frac{1}{2^{S}} \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{A}}^{S}
$$

where the first inequality follows from $a$ being uniformly sampled (so $\operatorname{Pr}\left[a=s_{\sigma}\right]=1 / 2^{S}$ ) and $S$ being even (so $\varepsilon_{\mathcal{A}}(\sigma, a)^{S} \geq 0, \forall a \in\{0,1\}^{S}$ ), and the second inequality follows from taking the expectation over $\sigma$ and using Jensen's inequality.

### 3.2 Reducing $k$-XOR DDH from GGM to HQM

In this section, we prove a reduction of the $k$-XOR DDH problem from GGM to HQM which increases in query complexity while maintaining the advantage. Here, we effectively create a hyperplane simulation of a $k$-XOR algorithm in GGM which exploits the lazy sampling property of the labeling function to keep the number of hyperplane queries low.

Lemma 2. Let $N$ be a prime, and $k \geq 1$ be an integer. For any generic algorithm $\mathcal{B}$ over $\mathcal{L}$ for $\mathbb{Z}_{N}$ such that $\mathcal{B}$ makes at most $Q$ oracle queries, and

$$
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{B}}:=2 \mathbb{P}\left[\mathcal{B}^{\mathcal{O}_{\sigma}}(\sigma(1), \sigma(\mathbf{w}))=\bigoplus_{i=1}^{k} b_{i}\right]-1
$$

where $\sigma(\mathbf{w}):=\left\{\sigma\left(x_{i}\right), \sigma\left(y_{i}\right), \sigma\left(u_{i}\right)\right\}_{i=1}^{k}, u_{i}=b_{i} z_{i}+\left(1-b_{i}\right) x_{i} y_{i},\left\{x_{i}, y_{i}, z_{i}\right\}_{i=1}^{k} \leftarrow \mathbb{Z}_{N}^{3 k},\left\{b_{1}, \ldots, b_{k}\right\} \leftarrow\{0,1\}^{k}$, and $\sigma: \mathbb{Z}_{N} \rightarrow \mathcal{L}$ is a random injective function, there exists a hyperplane query algorithm $\mathcal{C}$ which makes at most $2(3 k+2 Q+1)(3 k+3 Q+1)$ queries, and

$$
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{C}}:=2 \mathbb{P}\left[\mathcal{C}^{\mathcal{H}_{\mathrm{w}}}=\bigoplus_{i=1}^{k} b_{i}\right]-1=\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{B}}
$$

where $\mathbf{w}:=\left(x_{1}, y_{1}, u_{1}, \ldots, x_{k}, y_{k}, u_{k}\right), u_{i}=b_{i} z_{i}+\left(1-b_{i}\right) x_{i} y_{i},\left\{x_{i}, y_{i}, z_{i}\right\}_{i=1}^{k} \leftarrow \mathbb{Z}_{N}^{3 k}$, and $\left\{b_{1}, \ldots, b_{k}\right\} \leftarrow$ $\{0,1\}^{k}$.

Proof. We show how to simulate $\mathcal{B}^{\mathcal{O}_{\sigma}}(\sigma(1), \sigma(\mathbf{w}))$ given only oracle access to $\mathcal{H}_{\mathbf{w}}$. We rely on the fact that a random injective function $\sigma$ can be sampled lazily (i.e. new input-output pairs can be generated on the fly) with limited information from previously generated pairs. Specifically, the simulation algorithm $\mathcal{C}$ will maintain a global list $\operatorname{Cur}_{\sigma}:=\left\{\left(L_{i}, s_{i}\right)\right\}_{i \geq 0}$ to help record generated input-output pairs of $\sigma$ during a lazy sampling process, where $L_{i}$ is a multi-linear polynomial in variables $\left\{X_{i}, Y_{i}, U_{i}\right\}_{i=1}^{k}$ representing an implicit preimage of $s_{i}$ so that $s_{i}=\sigma\left(L_{i}(\mathbf{w})\right)$. Additionally, $\mathcal{C}$ will maintain a global set $I \subseteq \mathbb{Z}_{N}$ for tracking explicitly generated preimages, and a global set $J$ for the image set of $\sigma$.

The algorithm $\mathcal{C}$ is initialized by randomly choosing a set $J \subseteq \mathcal{L}$ of size $N$ as the image set of $\sigma$ and setting $\mathbf{C u r}_{\sigma}=\varnothing, I=\varnothing$. It then performs the following simulation with the help of two subroutines FindImage and FindPreimage defined later:

1. Compute challenges $\hat{\sigma}_{0}, \ldots, \hat{\sigma}_{3 k}$ as $\operatorname{FindImage}(H)$ for every polynomial $H \in\left\{1, X_{1}, Y_{1}, U_{1}, \ldots X_{k}, Y_{k}, U_{k}\right\}$.
2. Run $\mathcal{B}$ on the input $\hat{\sigma}_{0}, \ldots, \hat{\sigma}_{3 k}$ with the following oracle $\hat{\mathcal{O}}_{\sigma}\left(t_{1}, t_{2}\right)$ : If $t_{1} \notin J$ or $t_{2} \notin J$, return $\perp$, otherwise, return

## FindImage $\left(\right.$ FindPreimage $\left(t_{1}\right)+$ FindPreimage $\left.\left(t_{2}\right)\right)$.

Intuitively, the two subroutines simulate oracle accesses to a (lazily sampled) $\sigma$ and its inverse function $\sigma^{-1}$. The two subroutines first determine whether the queried input or output has appeared in $\mathbf{C u r}_{\sigma}$, if yes, then return the assigned (implicit) preimage or image accordingly, and otherwise generate a new pair on the fly. The main observation is that searching existence (in the form of a polynomial) and lazily generating new images or preimages can be done with only a few hyperplane queries. We use $\operatorname{Pre}\left(\mathbf{C u r} \boldsymbol{C}_{\sigma}\right)$ (respectively $\left.\operatorname{Im}\left(\mathbf{C u r}_{\sigma}\right)\right)$ to denote the set of $L_{i}$ 's (respectively $s_{i}$ 's) appearing in $\mathbf{C u r}_{\sigma}$. Formally,

- FindImage $(H)$ :

1. If $\mathcal{H}_{\mathbf{w}}(H-L)=1$ for some $(L, s) \in \mathbf{C u r}_{\sigma}$, return $s$.
2. Otherwise, sample $s_{\text {new }} \leftarrow J \backslash \operatorname{Im}\left(\mathbf{C u r}_{\sigma}\right)$, add $\left(H, s_{\text {new }}\right)$ to $\mathbf{C u r}_{\sigma}$, and return $s_{\text {new }}$.

- FindPreimage $(t)$ :

1. If $t=s$ for some $(L, s) \in \mathbf{C u r}_{\sigma}$, return $L$.
2. Otherwise, sample $\beta \leftarrow \mathbb{Z}_{N} \backslash I$ and let $I=I \cup\{\beta\}$ until return:

- If $\mathcal{H}_{\mathbf{w}}(L-\beta)=1$ for some $(L, s) \in \mathbf{C u r}_{\sigma}$, continue.
- Otherwise, let $L_{\text {new }}=\beta$, add $\left(L_{\text {new }}, t\right)$ to $\mathbf{C u r}_{\sigma}$, return $L_{\text {new }}$.

We claim that such an algorithm $\mathcal{C}$ fulfills the criteria in the theorem.

Query Complexity: The algorithm $\mathcal{C}$ makes at most $3 k+3 Q+1$ calls to FindImage and FindPreimage, each call contributing at most 1 new pair to $\mathbf{C u r}_{\sigma}$. Therefore the size of $\mathbf{C u r}_{\sigma}$ and the set of explicitly generated preimages $I$ are at any time upper bounded by $3 k+3 Q+1$. The calls of FindImage search the entire list $\mathbf{C u r}_{\sigma}$ for $(3 k+Q+1)$ times, and the calls of FindPreimage search $\mathbf{C u r}_{\sigma}$ for at most $|I|$ times at the end of the simulation. Because searching in $\mathbf{C u r}_{\sigma}$ takes at most $(3 k+3 Q+1)$ hyperplane queries, the simulation makes at most $(|I|+3 k+Q+1)(3 k+3 Q+1) \leq 2(3 k+2 Q+1)(3 k+3 Q+1)$ hyperplane queries.

Correctness: We will show that, for any fixed $\mathbf{w}$ (and $b_{1}, \ldots, b_{k} \in\{0,1\}$ ), $\mathcal{C}^{\mathcal{H}_{\mathbf{w}}}$ distributes the same as $\mathcal{B}^{\mathcal{O}_{\sigma}}(\sigma(1), \sigma(\mathbf{w}))$ over the randomness of $\sigma: \mathbb{Z}_{N} \rightarrow \mathcal{L}$ and the internal randomness of $\mathcal{B}$ and $\mathcal{C}$. Notice that $\mathcal{B}^{\mathcal{O}_{\sigma}}(\sigma(1), \sigma(\mathbf{w}))$ can be computed by the following process $\mathcal{D}$ : first randomly choose a set $J \subseteq \mathcal{L}$ of size $N$ as the simulated image set of $\sigma$ and set the list of generated input-output pairs $\mathbf{L i s t}_{\sigma}=\varnothing$, then perform the following,

1. Compute challenges $\sigma_{0}, \ldots, \sigma_{3 k}$ as $\operatorname{Lazy}_{\sigma}(\alpha)$ for every $\alpha \in\left\{1, x_{1}, y_{1}, u_{1}, \ldots x_{k}, y_{k}, u_{k}\right\}$.
2. Run $\mathcal{B}$ on the input $\sigma_{0}, \ldots, \sigma_{3 k}$ with access to $\operatorname{Lazy} \mathcal{O}_{\sigma}\left(t_{1}, t_{2}\right)$ : If $t_{1} \notin J$ or $t_{2} \notin J$, return $\perp$; otherwise, return

$$
\mathbf{L a z y}_{\sigma}\left(\mathbf{L a z y}_{\sigma^{-1}}\left(t_{1}\right)+\mathbf{L a z y}_{\sigma^{-1}}\left(t_{2}\right)\right)
$$

where $\operatorname{Lazy}_{\sigma}(\alpha)$ returns $s$ (respectively $\operatorname{Lazy}_{\sigma^{-1}}(s)$ returns $\alpha$ ) if $(\alpha, s)$ appears in $\mathbf{L i s t}_{\sigma}$, and otherwise samples a new $s$ from $J \backslash \operatorname{lm}\left(\right.$ List $\left._{\sigma}\right)$ (respectively a new $\alpha$ from $\mathbb{Z}_{N} \backslash \operatorname{Pre}\left(\operatorname{List}_{\sigma}\right)$ ), then adds the new pair $(\alpha, s)$ to $\mathbf{L i s t}_{\sigma}$.

It suffices to prove the following claim:
Claim 1. For any fixed $\mathbf{w}$ and $J$, the following two distributions $\left(\hat{\sigma}_{0}, \ldots, \hat{\sigma}_{3 k}, \hat{\sigma}_{3 k+1}, \ldots, \hat{\sigma}_{3 k+Q}, \mathbf{C u r}_{\sigma}^{3 k+Q}(\mathbf{w})\right)$ and $\left(\sigma_{0}, \ldots, \sigma_{3 k}, \sigma_{3 k+1}, \ldots, \sigma_{3 k+Q}\right.$, List $\left._{\sigma}^{3 k+Q}\right)$ are distributed identically, where $\sigma_{3 k+i}$ (respectively $\hat{\sigma}_{3 k+i}$ ) is the $i$-th response from $\mathbf{L a z y} \mathcal{O}_{\sigma}$ (respectively $\hat{\mathcal{O}}_{\sigma}$ ), and $\mathbf{C u r}_{\sigma}^{i}(\mathbf{w}):=\left\{\left(L_{j}(\mathbf{w}), \hat{s}_{j}\right)\right\}_{j \geq 0}$ (respectively List ${ }_{\sigma}^{i}:=$ $\left\{\left(\alpha_{j}, s_{j}\right)\right\}_{j \geq 0}$ ) denote the maintained list after $\hat{\sigma}_{i}$ (respectively $\sigma_{i}$ ) has been generated and all polynomials (if applicable) have been evaluated at input $\mathbf{w}$.

Proof. We will prove by induction. $\left(\hat{\sigma}_{0}, \operatorname{Cur}_{\sigma}^{0}(\mathbf{w})\right)=(\hat{s},(1, \hat{s}))$, and $\left(\sigma_{0}, \mathbf{L i s t}_{\sigma}^{0}\right)=(s,(1, s))$ where both $s$ and $s^{\prime}$ are uniformly distributed over $J$. Therefore the above claim holds for $i=0$. For $i \geq 1$, assume that the claim holds for $0 \leq j<i$. Now conditioning on, $\sigma_{j}=\hat{\sigma}_{j}$ for any $j<i$ and $\mathbf{C u r}_{\sigma}^{i-1}(\mathbf{w})=\mathbf{L i s t}{ }_{\sigma}^{i-1}$, we prove that $\left(\hat{\sigma}_{i}, \mathbf{C u r}_{\sigma}^{i}(\mathbf{w})\right)$ distributes the same as $\left(\sigma_{i}\right.$, List $\left._{\sigma}^{i}\right)$.

If $i \leq 3 k$, then $\sigma_{i}=\operatorname{Lazy}_{\sigma}(\alpha)$ and $\hat{\sigma}_{i}=\operatorname{FindImage}(H)$ for some $\alpha \in\left\{x_{1}, y_{1}, u_{1}, \ldots, x_{k}, y_{k}, u_{k}\right\}$ and $H \in\left\{X_{1}, Y_{1}, U_{1} \ldots, X_{K}, Y_{K}, U_{K}\right\}$ such that $H(\mathbf{w})=\alpha$. Given that $\mathbf{C u r}_{\sigma}^{i-1}(\mathbf{w})=$ List $_{\sigma}^{i-1},(\alpha, s)$ appears in List ${ }_{\sigma}^{i-1}$ if and only if there exists an $L$ such that $(L, s)$ appears in $\mathbf{C u r}_{\sigma}^{i-1}$ and $H(\mathbf{w})=L(\mathbf{w})=\alpha$ (i.e., $\left.\mathcal{H}_{\mathbf{w}}(H-L)=1\right)$. If $(\alpha, s)$ appears in $\operatorname{List}_{\sigma}^{i-1}$, then $\sigma=\hat{\sigma}_{i}=s$ and $\mathbf{C u r}_{\sigma}^{i}(\mathbf{w})=\mathbf{C u r}_{\sigma}^{i-1}(\mathbf{w})=\mathbf{L i s t}_{\sigma}^{i-1}=\mathbf{L i s t}_{\sigma}^{i}$. If $\alpha$ is new, then $\left(\sigma_{i}, \operatorname{List}_{\sigma}^{i}\right)=\left(s, \operatorname{List}_{\sigma}^{i-1} \cup\{(\alpha, s)\}\right)$ and $\left(\hat{\sigma}_{i}, \mathbf{C u r}^{i}(\mathbf{w})\right)=\left(\hat{s}, \mathbf{C u r}^{i-1}(\mathbf{w}) \cup\{(H(\mathbf{w}), \hat{s})\}\right)$ are distributed identically because both $s, \hat{s}$ are uniformly distributed over $J \backslash \operatorname{lm}\left(\operatorname{List}_{\sigma}^{i-1}\right)$ and $H(\mathbf{w})=\alpha$.

If $3 k<i \leq 3 k+Q$, then $\sigma_{i}=\mathbf{L a z y} \mathcal{O}_{\sigma}\left(t_{1}, t_{2}\right)$ and $\hat{\sigma}_{i}=\hat{\mathcal{O}}_{\sigma}\left(t_{1}, t_{2}\right)$ for some $t_{1}$ and $t_{2}$ computed based on $\sigma_{0}, \ldots, \sigma_{i-1}$ and the internal randomness of $\mathcal{B}$ (here we implicitly use that $\hat{\sigma}_{j}=\sigma_{j}$ for $j<i$ so the $i$-th query admits the same conditional distribution in both simulations). Notice that $\mathbf{C u r}_{\sigma}^{i-1}(\mathbf{w})=\mathbf{L i s t} \boldsymbol{t}_{\sigma}^{i-1}$. If $t_{1}$ or $t_{2}$ is not in $J$, then $\sigma=\hat{\sigma}=\perp$ and no updates are made to either list. If $t_{1} \in J \backslash \operatorname{lm}\left(\operatorname{List}_{\sigma}^{i-1}\right)$, then $\left(\alpha_{1}, t_{1}\right)$ will be added to List $\sigma_{\sigma}^{i-1}$ by calling $\operatorname{Lazy}_{\sigma^{-1}}\left(t_{1}\right)$, and $\left(\alpha_{1}^{\prime}, t_{1}\right)$ will be added into $\mathbf{C u r}_{\sigma}^{i-1}$ by calling FindPreimage $\left(t_{1}\right)$ where both $\alpha_{1}, \alpha_{1}^{\prime}$ are uniformly distributed over $\mathbb{Z}_{N} \backslash \operatorname{Pre}\left(\operatorname{List}{ }_{\sigma}^{i-1}\right)$. Similarly, if $t_{2} \in J \backslash \operatorname{Im}\left(\right.$ List $\left._{\sigma}^{i-1}\right)$, we can update both lists in a similar way. Moreover, importantly, the two updated lists, denoted as $\mathbf{L i s t} \mathbf{t}_{\sigma}^{*}$ and $\mathbf{C u r}_{\sigma}^{*}(\mathbf{w})$ (with $t_{1}, t_{2}$ in the list now), are still distributed identically. Conditioning on that $\mathbf{L i s t}_{\sigma}^{*}=\mathbf{C u r}_{\sigma}^{*}(\mathbf{w})$ and $t_{1}, t_{2}$ appear in List ${ }_{\sigma}^{*}$, we have that $L_{1}(\mathbf{w})=\alpha_{1}$ and $L_{2}(\mathbf{w})=\alpha_{2}$ for some $\left(L_{1}, t_{1}\right)$, $\left(L_{2}, t_{2}\right) \in \mathbf{C u r}_{\sigma}^{*}$ and some $\left(\alpha_{1}, t_{1}\right),\left(\alpha_{2}, t_{2}\right) \in \mathbf{L i s t}{ }_{\sigma}^{*}$. If for some $s,\left(\alpha_{1}+\alpha_{2}, s\right) \in \mathbf{L i s t}_{\sigma}^{*}$, then $\sigma=\hat{\sigma}=s$ and no updates are made to either list. Otherwise, $\left(\alpha_{1}+\alpha_{2}, s\right)$ will be added to $\mathbf{L i s t}_{\sigma}^{*}$ and $\left(L_{1}+L_{2}, s^{\prime}\right)$ will be added to $\mathbf{C u r}_{\sigma}^{*}$ where both $s, s^{\prime}$ are distributed uniformly over $J \backslash \operatorname{Im}\left(\mathbf{L i s t}_{\sigma}^{*}\right)$ and $L_{1}(\mathbf{w})+L_{2}(\mathbf{w})=\alpha_{1}+\alpha_{2}$. For both cases, $\left(\hat{\sigma}_{i}, \operatorname{Cur}_{\sigma}^{i}(\mathbf{w})\right)$ is distributed the same as $\left(\sigma_{i}, \mathbf{L i s t}_{\sigma}^{i}\right)$.

### 3.3 Hardness of $\boldsymbol{k}$-XOR DDH Problem in the HQM

In this section, we prove the generic hardness of $k$-XOR DDH problem in the hyperplane query model. The following lemma will be proven by induction over the number of DDH instances and queries needed by the optimal algorithm for any given case.

Lemma 3. Let $N$ be a prime, and $k \geq 1$ be an integer. For any hyperplane query algorithm $\mathcal{C}$ which makes at most $q$ queries, it holds that

$$
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{C}}:=2 \mathbb{P}_{\mathbf{w}}\left[\mathcal{C}^{\mathcal{H}_{\mathbf{w}}}=\bigoplus_{i=1}^{k} b_{i}\right]-1 \leq\left(\frac{12 q}{N k}+\frac{12}{N}\right)^{k}
$$

where $\mathbf{w}=\left(x_{1}, y_{1}, u_{1}, \ldots, x_{k}, y_{k}, u_{k}\right), u_{i}=b_{i} z_{i}+\left(1-b_{i}\right) x_{i} y_{i},\left\{x_{i}, y_{i}, z_{i}\right\}_{i=1}^{k} \leftarrow \mathbb{Z}_{N}^{3 k}$, and $\left\{b_{1}, \ldots, b_{k}\right\} \leftarrow\{0,1\}^{k}$.

For any integer $k \geq 1$ and $q \geq 0$, we let $\mathcal{C}_{k, q}$ denote the hyperplane query algorithm with $q$ queries for the $k$-XOR DDH problem in the HQM which maximizes the following quantity

$$
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, q):=\mathbb{P}_{\mathbf{w}}\left[\mathcal{C}_{k, q}^{\mathcal{H}_{\mathbf{w}}}=\bigoplus_{i=1}^{k} b_{i}\right]-\mathbb{P}_{\mathbf{w}}\left[\mathcal{C}_{k, q}^{\mathcal{H}_{\mathbf{w}}} \neq \bigoplus_{i=1}^{k} b_{i}\right]
$$

Similarly for any $k, q \geq 1$, we consider the optimal algorithm $\mathcal{D}_{k, q}$, making $q$ queries where its first hyperplane query is denoted by $H$, which maximizes the following quantity

$$
\Delta(k, q):=\mathbb{P}_{\mathbf{w}}\left[\mathcal{D}_{k, q}^{\mathcal{H}_{\mathbf{w}}}=\bigoplus_{i=1}^{k} b_{i} \cap E_{H}\right]-\mathbb{P}_{\mathbf{w}}\left[\mathcal{D}_{k, q}^{\mathcal{H}_{\mathbf{w}}} \neq \bigoplus_{i=1}^{k} b_{i} \cap E_{H}\right]
$$

where $E_{H}:=\left\{\mathcal{H}_{\mathbf{w}}(H)=1\right\}$. This is to say that $\mathcal{D}_{k, q}$ is the optimal $q$ query algorithm in the case that the first query output is 1 .

Our proof of the lemma will follow by induction on $(k, q)$. We first state our claims (see claim 2 and 3) for the bounds on the $\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}$ and $\Delta$ functions, which we will be using in the inductive step, present our inductive proof of the lemma, and then prove the claims. We consider our proof for claim 3 to be one of the main technical contributions of this work, which intuitively states that the advantage of any algorithm for $k$-XOR DDH problem, making a single successful hyperplane query (oracle output 1), is bounded by the product of the probability of finding the one successful hyperplane query and the advantage of some algorithm for $(k-1)$-XOR DDH problem.
Claim 2. For any $k \geq 1$, and $q \geq 1, \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, q) \leq \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, q-1)+2 \Delta(k, q)$.
Claim 3. For any $k \geq 2$ and $q \geq 1, \Delta(k, q) \leq 3 \boldsymbol{A d v}(k-1, q-1) / 2 N+\Delta(k-1, q) / 2 N$.
Proof of Lemma 3. Given Claim 2 and 3, we prove the following by induction on $(k, q)$.

$$
\begin{equation*}
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, q) \leq\left(\frac{4}{N}\right)^{k}\binom{q+k}{k} \tag{1}
\end{equation*}
$$

then Lemma 3 follows from the fact that $\binom{q+k}{k} \leq(3(q+k) / k)^{k}$.
The base cases $(k, 0)$ and $(1, q)$. For any $k \geq 1$, (1) holds trivially for $(k, 0)$ by the fact that $\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, 0)=0$. For any $q \geq 1$, by representing the first hyperplane query $H \in \mathbb{Z}_{N}\left[x_{1}, y_{1}, u_{1}\right]$ as a polynomial with coefficients $\alpha, \beta, \gamma, \lambda \in \mathbb{Z}_{N}$, we have that

$$
\begin{equation*}
\Delta(1, q) \leq \mathbb{P}_{\mathbf{w}}\left[\mathcal{D}_{1, q}^{\mathcal{H}_{\mathbf{w}}}=b \cap \mathcal{H}_{\mathbf{w}}(H)=1\right] \leq \mathbb{P}_{\mathbf{w}}\left[\mathcal{H}_{\mathbf{w}}(H)=1\right]=\mathbb{P}_{\mathbf{w}}\left[\alpha x_{1}+\beta y_{1}+\gamma u_{1}=\lambda\right] \leq \frac{3}{2 N} \tag{2}
\end{equation*}
$$

where the last inequality is because by the Schwartz-Zippel lemma Sch80, Zip79, the hyperplane query $H$ holds with probability at most $1 / N$ conditioning on $u_{1}=z_{1}$ and holds with probability at most $2 / N$ conditioning on $u_{1}=x_{1} y_{1}$ (as $x_{1}, y_{1}, z_{1} \leftarrow \mathbb{Z}_{N}$ and $H$ is not a constant polynomial by the optimality of $\mathcal{D}_{1, q}$ ). Then, by applying Claim 2 with $k=1$ recursively together with $\sqrt[2]{2}$, we obtain that

$$
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(1, q) \leq \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(1,0)+2 \sum_{i=0}^{q-1} \Delta(1, q-i)=2 \sum_{i=0}^{q-1} \Delta(1, q-i) \leq \frac{3 q}{N} \leq \frac{4}{N}\binom{q+1}{1}
$$

The inductive step. For any fixed $k \geq 2$ and $q \geq 1$, assume that holds for all tuples $(\ell, r)$ such that $\ell \in\{1, \ldots, k-1\}, r \in\{0, \ldots, q-1\}$. Then, recursively apply Claim $3(k-1)$ times to the final term in Claim 2, yielding

$$
\begin{aligned}
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, q) & \leq \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, q-1)+3\left(\sum_{i=1}^{k-1} \frac{\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k-i, q-1)}{2^{i-1} N^{i}}\right)+\frac{\Delta(1, q)}{2^{k-2} N^{k-1}} \\
& \leq \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, q-1)+3\left(\sum_{i=1}^{k-1} \frac{\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k-i, q-1)}{N^{i}}\right)+\frac{3}{N^{k}}
\end{aligned}
$$

where the last inequality uses $(2)$. Then, applying the inductive assumption,

$$
\begin{aligned}
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, q) & \leq\left(\frac{4}{N}\right)^{k}\binom{q+k-1}{k}+3\left(\sum_{i=1}^{k-1} \frac{1}{N^{i}} \cdot \frac{4^{k-i}}{N^{k-i}}\binom{q+(k-i)-1}{k-i}\right)+\frac{3}{N^{k}} \\
& \leq\left(\frac{4}{N}\right)^{k}\left(\binom{q+k-1}{k}+\binom{q+k-1}{k-1} \sum_{i=1}^{k} \frac{3}{4^{i}}\right) \\
& \leq\left(\frac{4}{N}\right)^{k}\left(\binom{q+k-1}{k}+\binom{q+k-1}{k-1}\right) \\
& =\left(\frac{4}{N}\right)^{k}\binom{q+k}{k}
\end{aligned}
$$

where the second inequality uses the fact that $\binom{q+i-1}{i-1} \leq\binom{ q+k-1}{k-1}$ for $1 \leq i \leq k-1$, the third inequality uses the fact that $\sum_{i=1}^{k}\left(1 / 4^{i}\right) \leq 1 / 3$, and the last equality uses the fact that $\binom{q+k}{k}=\binom{q+k-1}{k-1}+\binom{q+k-1}{k}$.

Proof of $\operatorname{Claim}$ 2. For any $k, q \geq 1$, let $\mathcal{C}_{k, q}$ be as in the definition of $\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, q)$ above, and let $H$ denote the first hyperplane query made by $\mathcal{C}_{k, q}$. We denote the XORed bits of $k \operatorname{DDH}$ instances by $b$, i.e., $b=\bigoplus_{i=1}^{k} b_{i}$.

It is useful to consider a hyperplane algorithm as a depth- $q$ decision tree, with each branch corresponding to the output of a single oracle query. Once the branch is selected at the root of a depth- $q$ tree (depending on the response to the first query), the remaining depth- $(q-1)$ subtree corresponds to a sub-algorithm making $q-1$ queries. When analyzing $\mathcal{C}_{k, q}$, we define $\mathcal{C}_{k, q-1}$ to be the sub-algorithm that gets executed when $\mathcal{H}_{\mathbf{w}}(H)=0$ and $\overline{\mathcal{C}}_{k, q-1}$ to be the sub-algorithm that gets executed when $\mathcal{H}_{\mathbf{w}}(H)=1$. In other words, the algorithms $\underline{\mathcal{C}}_{k, q-1}$ and $\overline{\mathcal{C}}_{k, q-1}$ can be thought of as the "left" and "right" subtrees of $\mathcal{C}_{k, q}$ corresponding to the output of the first query $H$ being 0 or 1 respectively. Note that by the definition of the subtrees,

$$
\mathbb{P}_{\mathbf{w}}\left[\mathcal{C}_{k, q}^{\mathcal{H}_{\mathbf{w}}}=b \cap E_{H}\right]=\mathbb{P}_{\mathbf{w}}\left[\overline{\mathcal{C}}_{k, q-1}^{\mathcal{H}_{\mathbf{w}}}=b \cap E_{H}\right] \text { and } \mathbb{P}_{\mathbf{w}}\left[\mathcal{C}_{k, q}^{\mathcal{H}_{\mathbf{w}}}=b \cap \bar{E}_{H}\right]=\mathbb{P}_{\mathbf{w}}\left[\mathcal{\mathcal { C }}_{k, q-1}^{\mathcal{H}_{\mathbf{w}}}=b \cap \bar{E}_{H}\right]
$$

Thus, the law of total probability gives the relationship

$$
\begin{aligned}
\mathbb{P}_{\mathbf{w}}\left[\mathcal{C}_{k, q}^{\mathcal{H}_{\mathbf{w}}}=b\right] & =\mathbb{P}_{\mathbf{w}}\left[\mathcal{C}_{k, q}^{\mathcal{H}_{\mathbf{w}}}=b \cap E_{H}\right]+\mathbb{P}_{\mathbf{w}}\left[\mathcal{C}_{k, q}^{\mathcal{H}_{\mathbf{w}}}=b \cap \bar{E}_{H}\right] \\
& =\mathbb{P}_{\mathbf{w}}\left[\overline{\mathcal{C}}_{k, q-1}^{\mathcal{H}_{\mathbf{w}}}=b \cap E_{H}\right]+\mathbb{P}_{\mathbf{w}}\left[\underline{\mathcal{C}}_{k, q-1}^{\mathcal{H}_{\mathbf{w}}}=b \cap \bar{E}_{H}\right] \\
& =\mathbb{P}_{\mathbf{w}}\left[\overline{\mathcal{C}}_{k, q-1} \mathcal{H}_{\mathbf{w}}=b \cap E_{H}\right]+\mathbb{P}_{\mathbf{w}}\left[\underline{\mathcal{C}}_{k, q-1}^{\mathcal{H}_{\mathbf{w}}}=b\right]-\mathbb{P}_{\mathbf{w}}\left[\mathcal{\mathcal { C }}_{k, q-1}^{\mathcal{H}_{\mathbf{w}}}=b \cap E_{H}\right] .
\end{aligned}
$$

Similarly,

$$
\mathbb{P}_{\mathbf{w}}\left[\mathcal{C}_{k, q}^{\mathcal{H}_{\mathbf{w}}} \neq b\right]=\mathbb{P}_{\mathbf{w}}\left[\overline{\mathcal{C}}_{k, q-1}^{\mathcal{H}_{\mathbf{w}}} \neq b \cap E_{H}\right]+\mathbb{P}_{\mathbf{w}}\left[\underline{\mathcal{C}}_{k, q-1}^{\mathcal{H}_{\mathbf{w}}} \neq b\right]-\mathbb{P}_{\mathbf{w}}\left[\underline{\mathcal{C}}_{k, q-1}^{\mathcal{H}_{\mathbf{w}}} \neq b \cap E_{H}\right]
$$

By the definition of $\Delta(k, q)$, we have that for any depth- $(q-1)$ subtree algorithm $C \in\left\{\underline{\mathcal{C}}_{k, q-1}, \overline{\mathcal{C}}_{k, q-1}\right\}$,

$$
\left|\mathbb{P}_{\mathbf{w}}\left[\mathcal{C}^{\mathcal{H}_{\mathbf{w}}}=b \cap E_{H}\right]-\mathbb{P}_{\mathbf{w}}\left[\mathcal{C}^{\mathcal{H}_{\mathbf{w}}} \neq b \cap E_{H}\right]\right| \leq \Delta(k, q)
$$

By the definition of $\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, q-1)$, we have that,

$$
\left|\mathbb{P}_{\mathbf{w}}\left[\underline{\mathcal{C}}_{k, q-1}^{\mathcal{H}_{\mathbf{w}}}=b\right]-\mathbb{P}_{\mathbf{w}}\left[\underline{\mathcal{C}}_{k, q-1}^{\mathcal{H}_{\mathbf{w}}} \neq b\right]\right| \leq \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, q-1)
$$

Therefore, by the definition of $\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, q)$ and the triangle inequality,

$$
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, q):=\left|\mathbb{P}_{\mathbf{w}}\left[\mathcal{C}_{k, q}^{\mathcal{H}_{\mathbf{w}}}=b\right]-\mathbb{P}_{\mathbf{w}}\left[\mathcal{C}_{k, q}^{\mathcal{H}_{\mathbf{w}}} \neq b\right]\right| \leq 2 \Delta(k, q)+\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k, q-1)
$$

Proof of Claim 3. Let $\mathcal{D}$ denote the sub-algorithm of that gets executed when $\mathcal{H}_{\mathbf{w}}(H)=1$ by $\mathcal{D}_{k, q}$, and $H$ denote the first hyperplane query made by $\mathcal{D}_{k, q}$, described by a linear equation $L\left(x_{1}, \ldots, u_{k}\right)=c$ for some constant $c$. Notice that $\mathcal{D}$ makes at most $q-1$ queries. Let $b:=\bigoplus_{i=1}^{k} b_{i}$, we further denote for any events $E_{1}, E_{2}$,

$$
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{1} \mid E_{2}\right)=\mathbb{P}_{\mathbf{w}}\left[\mathcal{D}^{\mathcal{H}_{\mathbf{w}}}=b \cap E_{1} \mid E_{2}\right]-\mathbb{P}_{\mathbf{w}}\left[\mathcal{D}^{\mathcal{H}_{\mathbf{w}}} \neq b \cap E_{1} \mid E_{2}\right]
$$

Thus, by definition $\Delta(k, q)=\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{H} \mid \varnothing\right.$ ). Because $L$ is not a constant polynomial (by the optimality of $\mathcal{D}_{k, q}$ ), there must be an $i \in[k]$ such that the coefficients of $x_{i}, y_{i}, u_{i}$ cannot be all zeros. Without loss of generality, we assume that $i=1$, i.e., the hyperplane query $H$ can be written as

$$
\alpha x_{1}+\beta y_{1}+\gamma u_{1}+L^{\prime}\left(\left\{x_{i}, y_{i}, u_{i}\right\}_{i=2}^{k}\right)=c
$$

where $(\alpha, \beta, \gamma) \neq 0^{3}$ and $L^{\prime}$ is some linear polynomial in $\left\{x_{i}, y_{i}, u_{i}\right\}_{i=2}^{k}$. It suffices to show

$$
\begin{equation*}
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{H} \mid b_{1}=1\right) \leq \frac{1}{N} \cdot \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k-1, q-1) \tag{3}
\end{equation*}
$$

and

$$
\begin{equation*}
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{H} \mid b_{1}=0\right) \leq \frac{2}{N} \cdot \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k-1, q-1)+\frac{1}{N} \cdot \Delta(k-1, q) \tag{4}
\end{equation*}
$$

Then the desired conclusion follows from averaging over $b_{1}$ and applying (3) and (4).
When $b_{1}=1$, then $u_{1}=z_{1}$ where $z_{1}$ is independent of $x_{1}, y_{1}, x_{2}, y_{2}, u_{2}, \ldots, x_{k}, y_{k}, u_{k}$. Without loss of generality, we assume that $\alpha \neq 0$ (the cases of $\beta \neq 0$ or $\gamma \neq 0$ are the same reasoning by conditioning on fixed values of other two variables),

$$
\begin{aligned}
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{H} \mid b_{1}=1\right) & =\mathbb{E}_{\delta_{1}, \delta_{2} \sim \mathbb{Z}_{N}}\left[\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{H} \mid y_{1}=\delta_{1} \cap z_{1}=\delta_{2}, b_{1}=1\right)\right] \\
& =\mathbb{E}_{\delta_{1}, \delta_{2} \sim \mathbb{Z}_{N}}\left[\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(\mid E_{H} \cap y_{1}=\delta_{1} \cap z_{1}=\delta_{2} \cap b_{1}=1\right) \mathbb{P}\left[E_{H} \mid y_{1}=\delta_{1} \cap z_{1}=\delta_{2} \cap b_{1}=1\right]\right] \\
& =\frac{1}{N} \cdot \mathbb{E}_{\delta_{1}, \delta_{2} \sim \mathbb{Z}_{N}}\left[\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(\mid E_{H} \cap y_{1}=\delta_{1} \cap z_{1}=\delta_{2} \cap b_{1}=1\right)\right] \\
& \leq \frac{1}{N} \cdot \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k-1, q-1)
\end{aligned}
$$

where the third line is because that for $b_{1}=1$ and any fixed $y_{1}=\delta_{1}, z_{1}=\delta_{2},\left(\alpha x_{1}+\beta y_{1}+\gamma z_{1}\right)$ is uniformly distributed over $\mathbb{Z}_{N}$ and independent from any $\left\{x_{i}, y_{i}, u_{i}\right\}_{i=2}^{k}$ so $E_{H}$ holds with probability $1 / N$; and the fourth line is because that further conditioning on that $E_{H}$ happens, by substituting $x_{1}, y_{1}, z_{1}$ with either constants or linear polynomials in variables $\left\{x_{i}, y_{i}, u_{i}\right\}_{i=2}^{k}$ in its hyperplane queries, $\mathcal{D}$ becomes an algorithm for solving $(k-1)$-XOR DDH with $(q-1)$ queries, and thus has advantage bounded by $\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k-1, q-1)$. (3) follows.

When $b=0$, then $u_{1}=x_{1} y_{1}$. Without loss of generality, we assume that at least one of $\gamma, \alpha$ is non-zero (if not we work with $\beta$ instead of $\gamma y_{1}+\alpha$ in the following reasoning). Then we can describe $H$ by the following equation

$$
\left(\gamma y_{1}+\alpha\right) x_{1}+\beta y_{1}+L^{\prime}\left(\left\{x_{i}, y_{i}, u_{i}\right\}_{i=2}^{k}\right)=c
$$

By similar argument as before, for any $\delta_{1}$ such that $\gamma \delta_{1}+\alpha \neq 0$, we have

$$
\begin{aligned}
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{H} \mid y_{1}=\delta_{1} \cap b_{1}=0\right) & =\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(\mid E_{H} \cap y_{1}=\delta_{1} \cap b_{1}=0\right) \mathbb{P}\left[E_{H} \mid y_{1}=\delta_{1} \cap b_{1}=0\right] \\
& \leq \frac{1}{N} \cdot \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k-1, q-1)
\end{aligned}
$$

where the inequality is because for $\gamma \delta_{1}+\alpha \neq 0,\left(\gamma \delta_{1}+\alpha\right) x_{1}$ is uniformly and distributed over $\mathbb{Z}_{N}$ and independent from any $\left\{x_{i}, y_{i}, u_{i}\right\}_{i=2}^{k}$ so $E_{H}$ holds with probability $1 / N$; and further conditioning on that $E_{H}$ happens, by replacing $y_{1}, x_{1}, u_{1}$ by either a constant or a linear polynomial only in variables $\left\{x_{i}, y_{i}, u_{i}\right\}_{i=2}^{k}$ (according to $y_{1}=\delta_{1}$ and the equation by $H$ and $u_{1}=\delta_{1} x_{1}$ ) in its hyperplane queries, $\mathcal{D}$ becomes an algorithm for $(k-1)$-XOR DDH problem, and thus has advantage bounded by $\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k-1, q-1)$.

For $\delta^{\prime}$ such that $\gamma \delta^{\prime}+\alpha=0$, let $H^{\prime}$ denote a new hyperplane query $L^{\prime}\left(\left\{x_{i}, y_{i}, u_{i}\right\}_{i=2}^{k}\right)=c-\delta^{\prime} \beta$. Then,

$$
\begin{aligned}
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{H} \mid y_{1}=\delta^{\prime} \cap b_{1}=0\right) & =\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{H^{\prime}} \mid y_{1}=\delta^{\prime} \cap b_{1}=0\right) \\
& =\mathbb{E}_{\delta_{2} \sim \mathbb{Z}_{N}}\left[\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{H^{\prime}} \mid y_{1}=\delta^{\prime} \cap x_{1}=\delta_{2} \cap b_{1}=0\right)\right] \\
& \leq \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k-1, q-1)+\Delta(k-1, q)
\end{aligned}
$$

where the inequality is because by substituting $y_{1}=\delta^{\prime}, x_{1}=\delta_{2}$ and $u_{2}=\delta^{\prime} \cdot \delta_{2}, \mathcal{D}$ becomes a $(k-1)$-XOR DDH hyperplane query algorithm in $\left\{x_{i}, y_{i}, u_{i}\right\}_{i=2}^{k}$. If $E_{H}^{\prime}$ is a non-trivial hyperplane query in $\left\{x_{i}, y_{i}, u_{i}\right\}_{i=2}^{k}$ (i.e., the degree of $L^{\prime}$ is non-zero), then its advantage is upper bounded by $\Delta(k-1, q)$ otherwise is at most $\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k-1, q-1)$. Combining both cases and the fact that $\operatorname{Pr}\left[\gamma y_{1}+\alpha=0\right] \leq 1 / N$ given $(\gamma, \alpha) \neq 0^{2}$.

$$
\begin{aligned}
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{H} \mid b_{1}=0\right) & =\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{H} \cap \gamma y_{1}+\alpha \neq 0 \mid b_{1}=0\right)+\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{H} \cap \gamma y_{1}+\alpha=0 \mid b_{1}=0\right) \\
& \leq \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{H} \mid \gamma y_{1}+\alpha \neq 0 \cap b_{1}=0\right)+\operatorname{Pr}\left[\gamma y_{1}+\alpha=0\right] \cdot \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{D}}\left(E_{H} \mid \gamma y_{1}+\alpha=0 \cap b_{1}=0\right) \\
& \leq \frac{1}{N} \cdot \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k-1, q-1)+\frac{1}{N}(\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}(k-1, q-1)+\Delta(k-1, q))
\end{aligned}
$$

implying (4).

### 3.4 Putting It All Together

We can use the above reductions and the final security bound to prove our two main theorems. At this point, both proofs follow in a relatively straightforward manner.

Theorem 1. Let $N$ be a prime and $S, T \geq 1$ be integers. For any generic algorithm with preprocessing $\mathcal{A}=\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$ over $\mathcal{L}$ for $\mathbb{Z}_{N}$ such that $\mathcal{A}_{0}$ outputs an $S$-bit advice string, and $\mathcal{A}_{1}$ makes at most $T$ oracle queries, it holds that

$$
\mathbb{P}\left[\mathcal{A}_{1}^{\mathcal{O}_{\sigma}}\left(\mathcal{A}_{0}^{\mathcal{O}_{\sigma}}(\sigma(1)), \sigma(1), \sigma(x), \sigma(y), \sigma(u)\right)=b\right] \leq \frac{1}{2}+2^{11} \frac{S T^{2}}{N}
$$

where $x, y, z \leftarrow \mathbb{Z}_{N}, b \leftarrow\{0,1\}, u=b z+(1-b) x y$, and $\sigma: \mathbb{Z}_{N} \rightarrow \mathcal{L}$ is a random injective function.
Proof. Let $\mathcal{A}^{\prime}$ be a generic algorithm with preprocessing such that $\mathcal{A}^{\prime}$ produces a $(2 S)$-bit advice string and makes at most $T$ online queries. The advantage of $\mathcal{A}^{\prime}$ is at least that of $\mathcal{A}$.

Now, let $k=2 S, Q=2 S T, q=2(4 S T+6 S+1)(6 S T+6 S+1)$, apply Lemma 1 and Lemma 2 to construct a $(2 S)$-XOR DDH algorithm $\mathcal{B}$ and a hyperplane query algorithm $\mathcal{C}$, then apply Lemma 3, such that

$$
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{A}^{\prime}}^{2 S} \leq 2^{2 S} \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{B}}=2^{2 S} \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{C}} \leq\left(\frac{24(4 S T+6 S+1)(6 S T+6 S+1)}{S N}+\frac{24}{N}\right)^{2 S}
$$

and thus

$$
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{A}} \leq \boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{A}^{\prime}} \leq \frac{3432(S T)^{2}}{S N}+\frac{24}{N} \leq 2^{12}\left(\frac{S T^{2}}{N}\right)
$$

Theorem 2. Let $N$ be a prime and $k, T \geq 1$ be integers. For any generic algorithm $\mathcal{B}$ over $\mathcal{L}$ for $\mathbb{Z}_{N}$ such that $\mathcal{B}$ makes at most $k T$ oracle queries, it holds that

$$
\mathbb{P}\left[\mathcal{B}^{\mathcal{O}_{\sigma}}(\sigma(1), \sigma(\mathbf{w}))=\bigoplus_{i=1}^{k} b_{i}\right] \leq \frac{1}{2}+\frac{1}{2}\left(2^{11} \frac{k T^{2}}{N}\right)^{k}
$$

where $\sigma(\mathbf{w}):=\left\{\sigma\left(x_{i}\right), \sigma\left(y_{i}\right), \sigma\left(u_{i}\right)\right\}_{i=1}^{k}$ for $u_{i}=b_{i} z_{i}+\left(1-b_{i}\right) x_{i} y_{i},\left\{x_{i}, y_{i}, z_{i}\right\}_{i=1}^{k} \leftarrow \mathbb{Z}_{N}^{3 k},\left\{b_{1}, \ldots, b_{k}\right\} \leftarrow$ $\{0,1\}^{k}$, and $\sigma: \mathbb{Z}_{N} \rightarrow \mathcal{L}$ a uniformly random injective function.

Proof. Let $\mathcal{B}$ be a $k$-XOR DDH algorithm such that $\mathcal{B}$ makes at most $k T$ queries. Let $q=2(2 k T+3 k+$ $1)(3 k T+3 k+1)$, apply Lemma 2 to construct a hyperplane query algorithm $\mathcal{C}$ and apply Lemma 3 , such that

$$
\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{B}}=\boldsymbol{A} \boldsymbol{d} \boldsymbol{v}_{\mathcal{C}} \leq\left(\frac{1008 k^{2} T^{2}}{k N}+\frac{12}{N}\right)^{k} \leq\left(2^{11} \frac{k T^{2}}{N}\right)^{k}
$$

## References

[ACDW20] Akshima, David Cash, Andrew Drucker, and Hoeteck Wee. Time-space tradeoffs and short collisions in merkle-damgård hash functions. In CRYPTO, 2020.
[ADGL23] Akshima, Xiaoqi Duan, Siyao Guo, and Qipeng Liu. On time-space lower bounds for finding short collisions in sponge hash functions. In $T C C, 2023$.
[AGL22] Akshima, Siyao Guo, and Qipeng Liu. Time-space lower bounds for finding collisions in merkledamgård hash functions. In CRYPTO, 2022.
[BL13] Daniel J. Bernstein and Tanja Lange. Non-uniform cracks in the concrete: the power of free precomputation. In ASIACRYPT, 2013.
[BMZ19] James Bartusek, Fermi Ma, and Mark Zhandry. The distinction between fixed and random generators in group-based assumptions. In CRYPTO, 2019.
[CDG18] Sandro Coretti, Yevgeniy Dodis, and Siyao Guo. Non-uniform bounds in the random-permutation, ideal-cipher, and generic-group models. In CRYPTO, 2018.
[CDGS18] Sandro Coretti, Yevgeniy Dodis, Siyao Guo, and John P. Steinberger. Random oracles and non-uniformity. In EUROCRYPT, 2018.
[CGLQ20] Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. Tight quantum time-space tradeoffs for function inversion. In FOCS, 2020.
[CHM20] Dror Chawin, Iftach Haitner, and Noam Mazor. Lower bounds on the time/memory tradeoff of function inversion. In TCC, 2020.
[CK18] Henry Corrigan-Gibbs and Dmitry Kogan. The discrete-logarithm problem with preprocessing. In EUROCRYPT, 2018.
[CK19] Henry Corrigan-Gibbs and Dmitry Kogan. The function-inversion problem: Barriers and opportunities. In $T C C, 2019$.
[DGK17] Yevgeniy Dodis, Siyao Guo, and Jonathan Katz. Fixing cracks in the concrete: Random oracles with auxiliary input, revisited. In EUROCRYPT, 2017.
[DH76] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. IEEE Transactions on Information Theory, 22(6):644-654, 1976.
[FGK22] Cody Freitag, Ashrujit Ghoshal, and Ilan Komargodski. Time-space tradeoffs for sponge hashing: Attacks and limitations for short collisions. In CRYPTO, 2022.
[FGK23] Cody Freitag, Ashrujit Ghoshal, and Ilan Komargodski. Optimal security for keyed hash functions: Avoiding time-space tradeoffs for finding collisions. In EUROCRYPT, 2023.
[FJM14] Pierre-Alain Fouque, Antoine Joux, and Chrysanthi Mavromati. Multi-user collisions: Applications to discrete logarithm, even-mansour and PRINCE. In ASIACRYPT, 2014.
[FN91] Amos Fiat and Moni Naor. Rigorous time/space tradeoffs for inverting functions. In STOC, 1991.
[GGKL21] Nick Gravin, Siyao Guo, Tsz Chiu Kwok, and Pinyan Lu. Concentration bounds for almost $k$-wise independence with applications to non-uniform security. In $S O D A, 2021$.
[GGPS23] Alexander Golovnev, Siyao Guo, Spencer Peters, and Noah Stephens-Davidowitz. Revisiting time-space tradeoffs for function inversion. In CRYPTO, 2023.
[GK22] Ashrujit Ghoshal and Ilan Komargodski. On time-space tradeoffs for bounded-length collisions in merkle-damgård hashing. In CRYPTO, 2022.
[GLLZ21] Siyao Guo, Qian Li, Qipeng Liu, and Jiapeng Zhang. Unifying presampling via concentration bounds. In TCC, 2021.
[Hel80] M. Hellman. A cryptanalytic time-memory trade-off. IEEE Transactions on Information Theory, 26(4):401-406, 1980.
[Imp11] Russell Impagliazzo. Relativized separations of worst-case and average-case complexities for NP. In $C C C, 2011$.
[KS01] Fabian Kuhn and René Struik. Random walks revisited: Extensions of pollard's rho algorithm for computing multiple discrete logarithms. In $S A C, 2001$.
[LCH11] Hyung Tae Lee, Jung Hee Cheon, and Jin Hong. Accelerating id-based encryption based on trapdoor dl using pre-computation. Cryptology ePrint Archive, 2011.
[Liu23] Qipeng Liu. Non-uniformity and quantum advice in the quantum random oracle model. In EUROCRYPT, 2023.
[Mih10] Joseph P Mihalcik. An analysis of algorithms for solving discrete logarithms in fixed groups. PhD thesis, Citeseer, 2010.
[Sch80] J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. Journal of the $A C M, 27(4): 701-717,1980$.
[Sho97] Victor Shoup. Lower bounds for discrete logarithms and related problems. In EUROCRYPT, 1997.
[Unr07] Dominique Unruh. Random oracles and auxiliary input. In CRYPTO, 2007.
[Yun15] Aaram Yun. Generic hardness of the multiple discrete logarithm problem. In EUROCRYPT, 2015.
[Zip79] R. Zippel. Probabilistic algorithms for sparse polynomials. In EUROSAM, 1979.

## A Preprocessing algorithms for DDH

Here, we show that any preprocessing algorithm for the discrete logarithm problem yields a preprocessing algorithm for the DDH problem with roughly the same advantage and complexity.

Proposition 1. Let $N$ be a prime and $S, T \geq 1$ be integers. For any generic algorithm with preprocessing $\mathcal{A}=\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$ over $\mathcal{L}$ for $\mathbb{Z}_{N}$ such that $\mathcal{A}_{0}$ outputs an $S$-bit advice string, $\mathcal{A}_{1}$ makes at most $T$ oracle queries, and

$$
\mathbb{P}\left[\mathcal{A}_{1}^{\mathcal{O}_{\sigma}}\left(\mathcal{A}_{0}^{\mathcal{O}_{\sigma}}(\sigma(1)), \sigma(1), \sigma(x)\right)=x\right]=\varepsilon
$$

where $x \leftarrow \mathbb{Z}_{N}$ and $\sigma: \mathbb{Z}_{N} \rightarrow \mathcal{L}$ is a random injective function, there exists a generic algorithm with preprocessing $\mathcal{B}=\left(\mathcal{B}_{0}, \mathcal{B}_{1}\right)$ over $\mathcal{L}$ for $\mathbb{Z}_{N}$ such that $\mathcal{B}_{0}$ outputs an $S$-bit advice string, $\mathcal{B}_{1}$ makes at most $T+4\lceil\log N\rceil$ oracle queries, and

$$
\mathbb{P}\left[\mathcal{B}_{1}^{\mathcal{O}_{\sigma}}\left(\mathcal{B}_{0}^{\mathcal{O}_{\sigma}}(\sigma(1)), \sigma(1), \sigma(x), \sigma(y), \sigma(u)\right)=b\right]=\frac{1}{2}+\frac{\varepsilon}{2}\left(1-\frac{1}{N}\right)
$$

where $x, y, z \leftarrow \mathbb{Z}_{N}, b \leftarrow\{0,1\}, u=b z+(1-b) x y$, and $\sigma: \mathbb{Z}_{N} \rightarrow \mathcal{L}$ is a random injective function.
Proof. Let $\mathcal{B}_{0}$ be the same as $\mathcal{A}_{0}$ which outputs an $S$-bit advice, and $\mathcal{B}_{1}$ be the following algorithm:

1. Simulate $\mathcal{A}_{1}$ with the output from $\mathcal{B}_{0}$ on $(\sigma(1), \sigma(x))$ and obtain $\hat{x}$.
2. Compute $\sigma(\hat{x})$ given $\sigma(1)$ and $\hat{x}$. If $\sigma(x) \neq \sigma(\hat{x})$, return a random bit $\hat{b} \leftarrow\{0,1\}$.
3. Compute $\sigma(\hat{x} y)$ given $\sigma(y)$ and $\hat{x}$. If $\sigma(u)=\sigma(\hat{x} y)$, return $\hat{b}:=0$ otherwise $\hat{b}:=1$.

Note that for two integers $x, y \in \mathbb{Z}_{N}$, we can compute the value of $\sigma(x y)$ given $x$ and $\sigma(y)$ with at most $2\lceil\log N\rceil$ queries to $\mathcal{O}_{\sigma}{ }^{\dagger}$. Therefore $\mathcal{B}_{1}$ makes at most $T+4\lceil\log N\rceil$ oracle queries.

Observe that, conditioning on $\sigma(x) \neq \sigma(\hat{x}), \hat{b}=b$ happens with probability exactly $1 / 2$. Conditioning on $\sigma(x)=\sigma(\hat{x})$, by the injectivity of $\sigma$, and the independence of $y, z, b$ from $\hat{x}$ and $\sigma$,

$$
\mathbb{P}[\hat{b}=b \mid \hat{x}=x]=\mathbb{P}[b=0]+\mathbb{P}[z \neq x y \cap b=1]=\frac{1}{2}+\frac{1}{2}\left(1-\frac{1}{N}\right)=1-\frac{1}{2 N} .
$$

Therefore, the overall success probability of $\mathcal{B}$ is

$$
\mathbb{P}[\hat{b}=b \mid \hat{x} \neq x] \mathbb{P}[\hat{x} \neq x]+\mathbb{P}[\hat{b}=b \mid \hat{x}=x]=\frac{1}{2}(1-\varepsilon)+\left(1-\frac{1}{2 N}\right) \varepsilon=\frac{1}{2}+\frac{\varepsilon}{2}\left(1-\frac{1}{N}\right)
$$

Combine this proposition with known generic preprocessing algorithms for discrete-log problem BL13, LCH11, Mih10, we obtain the following corollary:

Corollary. Let $N$ be a prime and $S, T \geq 1$ be integers, there exists a generic preprocessing algorithm $\mathcal{A}=\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$ over $\mathcal{L}$ for $\mathbb{Z}_{N}$ for the DDH problem such that $\mathcal{A}_{0}$ outputs an $S$-bit advice string, $\mathcal{A}_{1}$ makes at most $T$ oracle queries with advantage $\widetilde{\Omega}\left(S T^{2} / N\right)$.

[^1]
[^0]:    *Supported by National Natural Science Foundation of China Grant No. 62102260, NYTP Grant No. 20121201 and NYU Shanghai Boost Fund. Parts of the work are done while visiting CIFRA Institute, Bocconi University supported by Fondazione Cariplo (n. 2021-3641).
    ${ }^{\dagger}$ Shanghai Frontiers Science Center of Artificial Intelligence and Deep Learning

[^1]:    ${ }^{1}$ Given $\sigma(y)$ and an integer $x \in \mathbb{Z}_{N}$, we can compute $\sigma(x y)$ using recursion: if $x=1$, return $\sigma(y)$; otherwise first compute $Z:=\sigma(\lfloor x / 2\rfloor y)$ using recursion, then return $\mathcal{O}_{\sigma}(Z, Z)$ if $x$ is even, and return $\mathcal{O}_{\sigma}\left(\sigma(y), \mathcal{O}_{\sigma}(Z, Z)\right)$ if $x$ is odd. Because each recursion reduces computing $\sigma(x y)$ to a single sub-problem of computing $\sigma(\lfloor x / 2\rfloor y)$, and makes at most two oracle queries to $O_{\sigma}$. The number of recursion calls is at most $\lceil\log N\rceil$ and the number of oracle queries is at most $2\lceil\log N\rceil$.

