

Amortized Complexity of Zero-Knowledge Proofs Revisited: Achieving Linear Soundness Slack

Ronald Cramer¹, Ivan Damgård^{2*}, Chaoping Xing^{3**}, and Chen Yuan³

¹ CWI, Amsterdam & Mathematical Institute, Leiden University

² Department of Computer Science, Aarhus University

³ School of Physical & Mathematical Sciences, Nanyang Technological University

Abstract. We propose a new zero-knowledge protocol for proving knowledge of short preimages under additively homomorphic functions that map integer vectors to an Abelian group. The protocol achieves amortized efficiency in that it only needs to send $O(n)$ function values to prove knowledge of n preimages. Furthermore we significantly improve previous bounds on how short a secret we can extract from a dishonest prover, namely our bound is a factor $O(k)$ larger than the size of secret used by the honest prover, where k is the statistical security parameter. In the best previous result, the factor was $O(k^{\log^k n})$.

Our protocol can be applied to give proofs of knowledge for plaintexts in (Ring-)LWE-based cryptosystems, knowledge of preimages of homomorphic hash functions as well as knowledge of committed values in some integer commitment schemes.

1 Introduction

Proofs of Knowledge In a zero-knowledge protocol, a prover demonstrates that some claim is true (and in some cases that he knows a proof) while giving the verifier no other knowledge beyond the fact that the claim is true. Zero-knowledge protocols are essential tools in cryptographic protocol design. For instance, one needs zero-knowledge proofs of knowledge in multiparty computation to have a player demonstrate that he knows the input he is providing.

In this work, we will consider the problem of proving knowledge of a preimage under a one-way functions $f : \mathbb{Z}^r \mapsto G$ where G is an Abelian group (written additively in the following), and where furthermore the function is additively homomorphic, i.e., $f(\mathbf{a}) + f(\mathbf{b}) = f(\mathbf{a} + \mathbf{b})$. We will call such functions *ivOWF*'s (for homomorphic One-Way Functions over Integer Vectors). This problem was

* Supported by The Danish National Research Foundation and The National Science Foundation of China (under the grant 61061130540) for the Sino-Danish Center for the Theory of Interactive Computation, within which part of this work was performed; and by the Advanced ERC grant MPCPRO.

** Supported by the Singapore Ministry of Education under Tier 1 grants RG20/13 and RG25/16

considered in several earlier works, in particular recently in [BDLN16], from where we have borrowed most of the notation and basic definitions we use in the following.

ivOWF turns out to be a very general notion. Examples of ivOWFs include:

- The encryption function of several (Ring-)LWE-based cryptosystems (such as the one introduced in [BGV12] and used in the so-called SPDZ protocol [DPSZ12]).
- The encryption function of any semi-homomorphic cryptosystem as defined in [BDOZ11].
- The commitment function in commitment schemes for committing to integer values (see, e.g., [DF02]).
- Hash functions based on lattice problems such as [GGH96,LMPR08], where it is hard to find a short preimage.

We will look at the scenario where a prover \mathcal{P} and a verifier \mathcal{V} are given $y \in G$ and \mathcal{P} holds a short preimage \mathbf{x} of y , i.e., such that $\|\mathbf{x}\| \leq \beta$ for some β . \mathcal{P} wants to prove in zero-knowledge that he knows such an \mathbf{x} . When f is an encryption function and y is a ciphertext, this can be used to demonstrate that the ciphertext decrypts and \mathcal{P} knows the plaintext. When f is a commitment function this can be used to show that one has committed to a number in a certain interval.

A well-known, simple but inefficient solution is the following protocol π :

- (1) \mathcal{P} chooses \mathbf{r} at random such that $\|\mathbf{r}\| \leq \tau \cdot \beta$ for some sufficiently large τ , the choice of which we return to below.
- (2) \mathcal{P} then sends $a = f(\mathbf{r})$ to \mathcal{V} .
- (3) \mathcal{V} sends a random challenge bit b .
- (4) \mathcal{P} responds with $\mathbf{z} = \mathbf{r} + b \cdot \mathbf{x}$.
- (5) \mathcal{V} checks that $f(\mathbf{z}) = a + b \cdot y$ and that $\|\mathbf{z}\| \leq \tau \cdot \beta$.

If τ is sufficiently large, the distribution of \mathbf{z} will be statistically independent of \mathbf{x} , and the protocol will be honest verifier statistical zero-knowledge⁴. On the other hand, we can extract a preimage of y from a cheating prover who can produce correct answers $\mathbf{z}_0, \mathbf{z}_1$ to $b = 0, b = 1$, namely $f(\mathbf{z}_1 - \mathbf{z}_0) = y$. Clearly, we have $\|\mathbf{z}_1 - \mathbf{z}_0\| \leq 2 \cdot \tau \cdot \beta$. We will refer to the factor 2τ as the *soundness slack* of the protocol, because it measures the discrepancy between the interval used by the honest prover and what we can force a dishonest prover to do. The value of the soundness slack is important: if f is, e.g., an encryption function, then a large soundness slack will force us to use larger parameters for the underlying cryptosystem to ensure that the ciphertext decrypts even if the input is in the larger interval, and this will cost us in efficiency.

The naive protocol above requires an exponentially large slack to get zero-knowledge, but using Lyubachevsky’s rejection sampling technique, the soundness slack can be made polynomial or even constant (at least in the random oracle

⁴ We will only be interested in honest verifier zero-knowledge here. In applications one would get security for malicious verifiers by generating the challenge in a trusted way, e.g., using a maliciously secure coin-flip protocol.

model, at the cost that even the honest prover may sometimes fail to execute the protocol).

The obvious problem with the naive solution is that one needs to repeat the protocol k times where k is the statistical security parameter, to get soundness error probability 2^{-k} . This means that one needs to generate $\Omega(k)$ auxiliary f -values. We will refer to this as the *overhead* of the protocol and use it as a measure of efficiency.

One wants, of course as small overhead and soundness slack as possible, but as long as we only want to give a proof for a single f -value, we do not know how to reduce the overhead dramatically in general. But if instead we want to give a proof for k or more f -values, then we know how to reduce the *amortised* overhead: Cramer and Damgård ([CD09], see also full version in [CDK14]) show how to get amortised overhead $O(1)$, but unfortunately the soundness slack is $2^{\Omega(k)}$, even if rejection sampling is used⁵. In [DKL⁺13] two protocols were suggested, where one is only covertly secure. The other one can achieve polynomial soundness slack with overhead $\Omega(k)$ and works only in the random oracle model⁶. This was improved in [BDLN16]: a protocol was obtained (without random oracles) that has $O(1)$ overhead and quasi polynomial soundness slack (proportional to $n \cdot (2k + 1)^{\log(k)/2}$).

1.1 Contributions & Techniques

In this paper, we improve significantly the result from [BDLN16] and [DKL⁺13]: we obtain $O(1)$ overhead and soundness slack $O(k)$. All results hold in the standard model (no random oracles are needed). As with any other protocol with amortised efficiency, one needs to amortise over at least some number of instances before the amortisation “kicks in”, i.e., n needs to be large enough in order to achieve the amortized efficiency. Our most basic construction needs n to be $\Theta(k^2)$, and we later improve this to $\Theta(k^{3/2})$, still with the same overhead and soundness slack.

Our protocol uses a high-level strategy similar to [BDLN16]:

- (1) Do a cut-and-choose style protocol for the inputs y_1, \dots, y_n . This is a relatively simple but imperfect proof of knowledge: It only guarantees that the prover knows almost all preimages.
- (2) Let the verifier assign each y_i to one of several *buckets*.

⁵ In [CD09], the main result was first shown for functions dealing with *finite* rings and groups, and then generalised to the integers. The result is optimal for the finite case, while the integer case leaves room for improvement.

⁶ The protocol in [DKL⁺13] is actually stated as a proof of plaintext knowledge for random ciphertexts, but generalizes to a protocol for ivOWFs. It actually offers a tradeoff between soundness slack s and overhead in the sense that the overhead is $M \cdot \log(k)$, where M has to be chosen such that the error probability $(1/s)^M$ is negligible. Thus to get exponentially small error probability in k as we do here, one can choose s to be $\text{poly}(k)$ and hence M will be $\Omega(k/\log k)$.

- (3) For each bucket, add all elements that landed in the bucket and do an imperfect proof of knowledge as in the first step, but now with all the bucket sums as input.

The reason why one might hope this would work is as follows: as mentioned, the first step will ensure that we can extract *almost* all of the required n preimages, in fact we can extract all but k preimages (we assume throughout that $n \gg k$). In the second step, since we only have k elements left that were “bad” in the sense that we could not yet extract a preimage, then if we have many more than k buckets and distribute them in buckets according to a carefully designed strategy, we may hope that with overwhelming probability, all the bad elements will be alone in one of those buckets for which we can extract a preimage of the bucket sum. This seems plausible because we can extract almost all such preimages. If indeed this happens, we can extract all remaining preimages by linearity of f : each bad element can be written as a sum of elements for which the extractor already knows a preimage.

Furthermore, the overall cost of doing the protocol would be $O(n)$, and the soundness slack will be limited by the maximal number of items in a bucket. In fact, if each bucket contains $O(k)$ elements, then the soundness slack is $O(k)$ as well. Our main technical contribution is a construction of a strategy for assignment to buckets with properties as we just outlined. We explain more about the intuition below.

In comparison, the protocol from [BDLN16] also plays a “balls and buckets” game. The difference is that they use only $O(k)$ buckets, but repeat the game $\Omega(\log k)$ times. This means that their extraction takes place in $\Omega(\log k)$ stages, which leads to the larger soundness slack. Also, they use a randomised strategy for assignment to buckets. While this makes the protocol and analysis somewhat more complicated, the randomization seems critical to make the proof go through: it makes essential use of the fact that the adversary does not know how elements are distributed in buckets until after the “bad” elements from Step 1 have been fixed. It is therefore somewhat surprising that the problem can be solved with a deterministic strategy, as we do here.

We also show a probabilistic strategy which is inferior to our deterministic one in that it requires k^3 input instances to work. On the other hand, it differs from the deterministic strategy by being more flexible: if the number of instances is less than k^3 , then the protocol will not remove all bad elements, but it will reduce the number of bad elements significantly. We can therefore combine the deterministic and probabilistic methods to get a protocol that works already for $k^{3/2}$ input instances, still with the same overhead and soundness slack.

Our protocol is honest verifier zero-knowledge and is sound in the sense of a standard proof of knowledge, i.e., we extract the prover’s witness by rewinding. Nevertheless, the protocol can be readily used as a tool in a bigger protocol that is intended to be UC secure against malicious adversaries. Such a construction is already known from [DPSZ12].

We now explain how we arrive at our construction of the strategy for assigning elements to buckets: We define the buckets via a bipartite graph. Consider a

finite, undirected, bipartite graph $G = (L, R, E)$ without multi-edges, where L denotes the set of vertices “on the left,” R those “on the right” and E the set of edges. Write $n = |L|$ and $m = |R|$. Each vertex $w \in R$ on the right gives a “bucket of vertices” $N(\{w\}) \subset L$ on the left, where $N(\{w\})$ denotes the neighborhood of w .

We say that the bipartite graph G has the (f_1, f_2) -strong unique neighbour property if the following holds. For each set $N_1 \subset L$ with $|N_1| = f_1$, for each set $N_2 \subset R$ with $|N_2| = f_2$, and for each $i \in N_1$, there is $w \in R \setminus N_2$ such that $N_1 \cap N(\{w\}) = \{i\}$. Note that this property is anti-monotonous in the sense that if it holds for parameters (f_1, f_2) it also holds for parameters (f'_1, f'_2) with $f'_1 \leq f_1$ and $f'_2 \leq f_2$.

With f_1 corresponding to the failures in step 1 and f_2 corresponding to those in step 3, it should be clear that this property on (an infinite family of bipartite graphs) G , together with the conditions that $n = \text{poly}(k)$, $m = O(n)$, $f_1 = O(k)$, $f_2 = O(k)$ and the condition that the right-degrees in G are all in $O(k)$, is sufficient to pull off our claimed result. Of course, in addition, this requires *efficient* construction of G . We propose two approaches satisfying each of these requirements. The first one, based on a construction from universal hash functions, achieves $n = O(k^2)$. A second approach, based on certain excellent (nonconstant-degree) expander graphs achieves $n = O(k^3)$, but also achieves a weaker (but still useful) “neighbour property” even if n is much smaller than k^3 .

Notation

Throughout this work we will format vectors such as \mathbf{b} in lower-case bold face letters, whereas matrices such as \mathbf{B} will be in upper case. We refer to the i th position of vector \mathbf{b} as $\mathbf{b}[i]$, let $[r] := \{1, \dots, r\}$ and define for $\mathbf{b} \in \mathbb{Z}^r$ that $\|\mathbf{b}\| = \max_{i \in [r]} \{\mathbf{b}[i]\}$. To sample a variable g uniformly at random from a set G we use $g \xleftarrow{\$} G$. Throughout this work we will let λ be a computational and k be a statistical security parameter. Moreover, we use the standard definition for polynomial and negligible functions and denote those as $\text{poly}(\cdot)$, $\text{negl}(\cdot)$.

2 Homomorphic OWFs and Zero-Knowledge Proofs

We first define a primitive called *homomorphic one-way functions over integer vectors*. It is an extension of the standard definition of a OWF found in [KL14]. Let $\lambda \in \mathbb{N}$ be the security parameter, we consider a probabilistic polynomial time algorithm Gen which on input 1^λ outputs: an Abelian group G , natural numbers β, r , and a function $f : \mathbb{Z}^r \rightarrow G$. Let \mathcal{A} be any algorithm. Consider the following game:

$\text{Invert}_{\mathcal{A}, Gen}(\lambda)$:

- (1) Run $Gen(1^\lambda)$ to get G, β, r and f .
- (2) Choose $\mathbf{x} \in \mathbb{Z}^r$, $\|\mathbf{x}\| \leq \beta$ and compute $y = f(\mathbf{x})$.
- (3) On input $(1^\lambda, y, G, \beta, r, f)$ the algorithm \mathcal{A} computes an \mathbf{x}' .

(4) Output 1 iff $f(\mathbf{x}') = y, \|\mathbf{x}'\| \leq \beta$, and 0 otherwise.

Definition 1 (Homomorphic OWF over Integer Vectors (ivOWF)). *The algorithm Gen producing functions of form $f : \mathbb{Z}^r \rightarrow G$ is called a homomorphic one-way function generator over the integers if the following conditions hold:*

- (1) *There exists a polynomial-time algorithm $eval_f$ such that $eval_f(\mathbf{x}) = f(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{Z}^r$.*
- (2) *For all $\mathbf{x}, \mathbf{x}' \in \mathbb{Z}^r$ it holds that $f(\mathbf{x}) + f(\mathbf{x}') = f(\mathbf{x} + \mathbf{x}')$.*
- (3) *For every probabilistic polynomial-time algorithm \mathcal{A} there exists a negligible function $negl(\lambda)$ such that*

$$\Pr[\text{Invert}_{\mathcal{A}, Gen}(\lambda) = 1] \leq \text{negl}(\lambda)$$

In the following, we will abuse terminology slightly by referring to a fixed function $f : \mathbb{Z}^r \rightarrow G$ as an ivOWF. As mentioned in the introduction, this abstraction captures, among other primitives, lattice-based encryption schemes such as [BGV12, GSW13, BV14] where the one-way property is implied by IND-CPA and β is as large as the plaintext space. Moreover it also captures hash functions such as [GGH96, LMPR08], where it is hard to find a preimage for all *sufficiently short* vectors that have norm smaller than β .

2.1 Proving Knowledge of Preimage

We consider two parties, the prover \mathcal{P} and the verifier \mathcal{V} . \mathcal{P} holds values $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{Z}^r$, both parties have values $y_1, \dots, y_n \in G$ and \mathcal{P} wants to prove to \mathcal{V} that $y_i = f(\mathbf{x}_i)$ and that \mathbf{x}_i is *short*, while giving away no extra knowledge on the \mathbf{x}_i . More formally, the relation that we want to give a zero-knowledge proof of knowledge for is

$$R_{\text{KSP}} = \left\{ (G, \beta, v, w) \left| \begin{array}{l} v = (y_1, \dots, y_n) \wedge w = (\mathbf{x}_1, \dots, \mathbf{x}_n) \wedge \\ [y_i = f(\mathbf{x}_i) \wedge \|\mathbf{x}_i\| \leq \beta]_{i \in [n]} \end{array} \right. \right\}$$

However, like all other protocols for this type of relation, we will have to live with a *soundness slack* τ as explained in the introduction. What this means more precisely is that there must exist a knowledge extractor with properties exactly as in the standard definition of knowledge soundness, but the extracted values only have to satisfy $[y_i = f(\mathbf{x}_i) \wedge \|\mathbf{x}_i\| \leq \tau \cdot \beta]_{i \in [n]}$.

3 Proofs of Preimage

3.1 Imperfect Proof of Knowledge

The first tool we need for our protocol is a subprotocol which we borrow from [BDLN16], a so-called *imperfect proof of knowledge*. This protocol is proof of

knowledge for the above relation with a certain soundness slack, however, the knowledge extractor is only required to extract almost all preimages. We note that to show knowledge soundness later for our full protocol, Goldreich and Bellare [BG93] have shown that it is sufficient to consider deterministic provers, therefore we only need to consider deterministic provers in the following.

The idea for the protocol is that the prover constructs $T = 3n$ auxiliary values of form $z_i = f(\mathbf{r}_i)$ where \mathbf{r}_i is random and short. The verifier asks the prover to open half the values (chosen at random) and aborts if the preimages received are not correct and short. One can show that this means the prover must know correct preimages of almost all the unopened values. The prover must now reveal, for each y_i in the input, a short preimage of the sum $y_i + z_j$ for some unopened z_j . By the homomorphic property of f this clearly means we can extract from the prover also a short preimage of most of the y_i 's.

The reason one needs to have more than $2n$ auxiliary values is that the protocol makes use of Lyubashevsky's rejection sampling technique [Lyu08,Lyu09], where the prover is allowed to refuse to use some of the auxiliary values. This allows for a small soundness slack while still maintaining the zero-knowledge property. For technical reasons the use of rejection sampling means that the prover should not send the auxiliary values z_i in the clear at first but should commit to them, otherwise we cannot show zero-knowledge.

The following theorem is proved in [BDLN16] (their Theorem 1):

Theorem 1. *Let f be an iwOWF, k be a statistical security parameter, Assume we are given C_{aux} , a perfectly binding/computationally hiding commitment scheme over G , $\tau = 100 \cdot r$ and $T = 3 \cdot n$, $n \geq \max\{10, k\}$. Then there exists a protocol $\mathcal{P}_{\text{IMPERFECTPROOF}}$ with the following properties:*

Efficiency: *The protocol requires communication of at most $T = 3n$ f -images and preimages.*

Completeness: *If \mathcal{P}, \mathcal{V} are honest and run on an instance of R_{KSP} , then the protocol succeeds with probability at least $1 - \text{negl}(k)$.*

Soundness: *For every deterministic prover $\hat{\mathcal{P}}$ that succeeds to run the protocol with probability $p > 2^{-k+1}$ one can extract at least $n - k$ values \mathbf{x}'_i such that $f(\mathbf{x}'_i) = y_i$ and $\|\mathbf{x}'_i\| \leq 2 \cdot \tau \cdot \beta$, in expected time $O(\text{poly}(s) \cdot k^2/p)$ where s is the size of the input to the protocol.*

Zero-Knowledge: *The protocol is computational honest-verifier zero-knowledge.*

In the following we will use $\mathcal{P}_{\text{IMPERFECTPROOF}}(v, w, T, \tau, \beta)$ to denote an invocation of the protocol from this theorem with inputs $v = (y_1, \dots, y_n), w = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ and parameters τ, β .

3.2 The Full Proof of Knowledge

The above imperfect protocol will be used as a building block. After executing it with the (\mathbf{x}_i, y_i) as input, we may assume that a preimage of most of the y_i 's (in fact, all but k) can be extracted from the prover.

The strategy for the last part of the protocol is as follows: each y_i is assigned to one of several *buckets*. Then, for each bucket, we add all elements that landed in the bucket and have the prover demonstrate that he knows a preimage of the sum. The observation (made in [BDLN16]) is that we can now extract a preimage of every bad elements that is alone in a bucket. The question, however, is how we distribute items in buckets to maximize our chance of extracting all the missing preimages, and how many buckets we should use. One solution to this was given in [BDLN16], but it requires repeating the experiment $\log k$ times before all bad elements have been handled with good probability.

Here we propose a new strategy that achieves much better results: we need just one repetition of the game and each bucket will contain only $O(k)$ items which gives us the soundness slack of $O(k)$.

Before we can describe the protocol, we need to define a combinatorial object we use in the protocol, namely a good set system:

Definition 2. A set system \mathcal{S} with parameters n, m is a collection of m index sets B_1, \dots, B_m , where each $B_j \subset [n]$, and $[n] = \{1, \dots, n\}$. Both n and m depend on a security parameter k . The set system is good if the maximal size of a set B_j is $O(k)$, m is $O(n)$ and if for every set $N_1 \subset [n]$ of size k , every set $N_2 \subset [m]$ of size k and every $i \in N_1$, there exists $j \in [m] - N_2$ such that $B_j \cap N_1 = \{i\}$.

The idea in the definition is that the buckets are defined by the sets $\{B_j\}$. Then, if the set system is good, and if we can extract preimage sums over all bucket except k , then we will be in business.

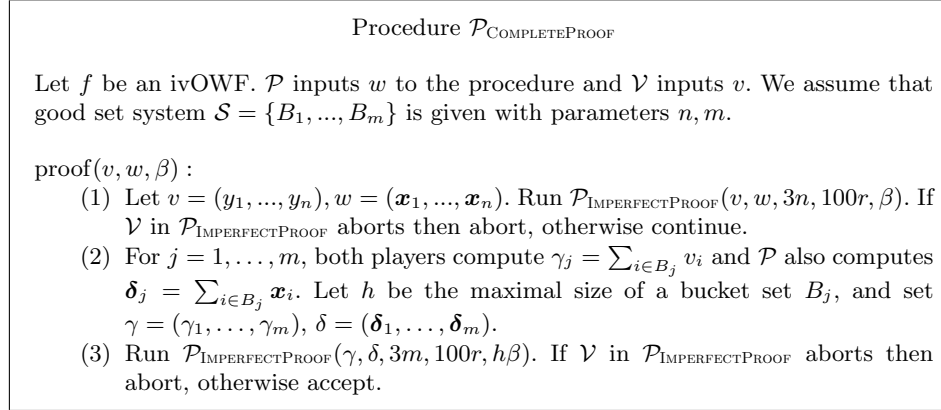


Fig. 1. A protocol to prove the relation R_{KSP}

Theorem 2. Let f be an ivOWF, k be a statistical security parameter, and β be a given upper bound on the size of the finest prover's secrets. If $\mathcal{P}_{\text{COMPLETEPROOF}}$ is executed using a good set system \mathcal{S} , then it is an interactive honest-verifier zero-knowledge proof of the relation R_{KSP} with knowledge error 2^{-k+1} . More specifically, it has the following properties:

Efficiency *The protocol has overhead $O(1)$.*

Correctness: *If \mathcal{P}, \mathcal{V} are honest then the protocol succeeds with probability at least $1 - 2^{-O(k)}$.*

Soundness: *For every deterministic prover $\hat{\mathcal{P}}$ that succeeds to run the protocol with probability $p > 2^{-k+1}$ one can extract n values \mathbf{x}'_i such that $f(\mathbf{x}'_i) = y_i$ and $\|\mathbf{x}'_i\| \leq O(k \cdot r \cdot \beta)$ except with negligible probability, in expected time $\text{poly}(s, k)/p$, where s is the size of the input to the protocol.*

Zero-Knowledge: *The protocol is computational honest-verifier zero-knowledge.*

Proof. Efficiency is immediate from Theorem 1 and the fact that we use a good set system, so that m is $O(n)$. Note also that the verifier can specify the set system for the prover using $O(m \cdot k \cdot \log n)$ bits. This will be dominated by the communication of m preimages if a preimage is larger than $k \log n$ bits, which will be the case for any realistic setting.

Correctness is immediate from correctness of $\mathcal{P}_{\text{IMPERFECTPROOF}}$.

The extractor required for knowledge soundness will simply run the extractor for $\mathcal{P}_{\text{IMPERFECTPROOF}}$ twice, corresponding to the 2 invocations of $\mathcal{P}_{\text{IMPERFECTPROOF}}$. Let N_1 be the set of k preimages we fail to extract in the first invocation, and let N_2 be the set of bucket sums we fail to extract in the second invocation. The properties of a good set system distribution now guarantee that no matter what set N_2 turns out to be, we can find, for each $i \in N_1$, a set B_j where we know a preimage of the sum over the bucket ($j \in [m] - N_2$), and furthermore $B_j \cap N_1 = \{i\}$. Concretely, we know δ_j such that $f(\delta_j) = \sum_{l \in B_j} y_l$ and we know preimages of all summands except for y_i . By the homomorphic property of f we can solve for a preimages of y_i , and the size of the preimage found follows immediately from Theorem 1 and the fact that buckets have size $O(k)$.

Honest-verifier zero-knowledge follows immediately from Theorem 1. We do the simulation by first invoking the simulator $\mathcal{P}_{\text{IMPERFECTPROOF}}$ with the input parameters for the first step. We then sample according to \mathcal{D} , compute the input parameters for the second invocation and run the simulator for $\mathcal{P}_{\text{IMPERFECTPROOF}}$ again. \square

To make this theorem be useful, we need of course that good set systems exist. This is taken care of in the following theorem which we prove in the next section.

Theorem 3. *Good set systems exist with parameters $n = m \in O(k^2)$ and can be constructed in time polynomial in k .*

This theorem implies that we need to have at least $\Omega(k^2)$ instances to amortise over to get an efficient protocol. Of course, for the applicability of the protocol it is better if one could make do with less. We now sketch how to get the same overhead and soundness slack using only $O(k^{3/2})$ inputs.

This is based on a weaker, but more flexible notion of set system, namely an (k, d, s) -good set system:

Definition 3. *A set system \mathcal{S} with parameters n, m is a collection of m index sets B_1, \dots, B_m with each $B_j \subseteq [n]$. Both parameters n, m depend on a security*

parameter k . We say a set system is (k, d, s) -good for N_1 if m is $O(n)$, the maximal size of a set B_j is d and if $N_1 \subseteq [n]$ of size k satisfies the following: for every set $N_2 \subseteq [m]$ of size k , there exists a subset $T \subseteq N_1$ of size at least $k - s$ such that for every $i \in T$, there exists $j \in [m] - N_2$ satisfying $B_j \cap N_1 = \{i\}$.

As before, the idea is that the system can be used to design a protocol based on a balls-and-buckets game similar to the above, where the B_j 's define the buckets, and N_1, N_2 correspond to the subset of instances we fail to extract via the weak zero-knowledge protocol. The final requirement now says that if the system is good for N_1 , then we can extract witnesses for $k - s$ of the remaining bad items in N_1 using the witnesses we have for the bucket sums.

While it seem like bad news that we will not be able to kill all the bad items in N_1 , the point is that this relaxed requirement enables us to construct such set systems with different parameters, in particular with much smaller n, m compared to k that we can get for a regular set system. In particular we have the following theorem which is proved in the next section.

Theorem 4. *For any constant $0 < c < 1$, there is a probabilistic polynomial time algorithm for constructing set systems where $m = n = O(k^{1+2c})$, such that for any fixed $N_1 \subseteq [n]$ of size k , the resulting system is $(k, k^c, 5k^{1-c})$ -good for N_1 except with probability exponentially small in k ,*

In our protocol, we set $c = 0.25$, so we get that we can construct a set system $\mathcal{S}_1 = \{A_1, \dots, A_m\}$ with $m = n = O(k^{1.5})$, such that for any fixed N_1 , it will be $(k, k^{0.25}, 5k^{0.75})$ -good for N_1 , except with exponentially small probability. Note that this property does not guarantee that the system will be good for every N_1 simultaneously.

On the other hand, this property *is* guaranteed by the good set systems from Theorem 3. It is easy to see that these are *simultaneously* $(r, 2r, 0)$ -good for all N_1 of size k . We are going to set $r = 5k^{0.75}$. So we obtain a $(5k^{0.75}, 10k^{0.75}, 0)$ -good set system $\mathcal{S}_2 = \{B_1, \dots, B_m\}$ with $m = n = O(k^{1.5})$.

Here follows an informal sketch of the protocol we can now construct for an input consisting of $n = O(k^{1.5})$ f -images $\mathbf{y} = (y_1, \dots, y_n)$:

- (1) Both players compute bucket sums $\boldsymbol{\delta} = (\delta_1, \dots, \delta_m)$ of the y_i 's according to the set system \mathcal{S}_2 .
- (2) Run the imperfect zero-knowledge proof for both \mathbf{y} and $\boldsymbol{\delta}$. Note that at this point we cannot hope to extract all witnesses. This would require that only $5k^{0.75}$ witnesses were left unknown by the imperfect proofs. But this is not the case. Therefore we extend the protocol to reduce this number:
- (3) The verifier constructs a set system \mathcal{S}_1 according to Theorem 4 with parameters as defined above. Both players compute bucket sums $\mathbf{u} = (u_1, \dots, u_m)$ of the y_i 's according to the set system \mathcal{S}_1 . Moreover, the players compute bucket sums $\boldsymbol{\omega} = (\omega_1, \dots, \omega_m)$ of the δ_i 's according to the system \mathcal{S}_1 .
- (4) Run the imperfect zero-knowledge proof for \mathbf{u} and $\boldsymbol{\omega}$.

We sketch the argument that this is sound as a proof of knowledge: after we run the extractor for the first two imperfect proofs, we know witnesses for all y_i

except for a set N_1 and for all δ_i except for a set N'_1 . Now, we know that except with negligible probability the set system \mathcal{S}_1 will be good for both N_1 and N'_1 (by a union bound). And we can run the knowledge extractor for the last two imperfect proofs so we will get witnesses for all u_i except a set N_2 and for all ω_i except a set N'_2 . All these sets have size k .

Now, by Definition 3, and because we can assume that \mathcal{S}_1 is $(k, k^{0.25}, 5k^{0.75})$ -good for both N_1 and N'_1 , we can use the homomorphic property of f and the known witnesses for \mathbf{y}, \mathbf{u} in the usual way to reduce the set of unknown witnesses for \mathbf{y} (in N_1) to a set M_1 of size $5k^{0.75}$. Like wise, we can reduce the set of unknown witnesses (in N'_1) for δ to a set M_2 of size $5k^{0.75}$.

Finally, we are in a position to use that \mathcal{S}_2 is a $(5k^{0.75}, 10k^{0.75}, 0)$ -good set system, where M_1, M_2 are the set of unknown witnesses. This will allow us to extract all witnesses. Note that the set M_1 is not fixed when \mathcal{S}_2 is constructed but this is fine since \mathcal{S}_2 is simultaneously good for all sets of size $5k^{0.75}$.

We leave it to the reader to verify that this protocol has overhead $O(1)$ and soundness slack $O(k)$.

4 Proof of Theorem 3 and Theorem 4

4.1 Definitions and Conventions

Let $G = (L, R, E)$ be a finite, undirected bipartite graph. For simplicity we also assume G has no multi-edges.⁷ Here, L denotes the set of vertices “on the left,” R the set of vertices “on the right” and E the set of edges. A vertex v is said to be *adjacent* to a vertex w if $(v, w) \in E$. An edge $e \in E$ is *incident* to a vertex v if there is a vertex w such that $e = (v, w)$. Suppose $S \subset L$ and $T \subset R$. The *neighborhood* of S , denoted $N(S)$, consists of all vertices adjacent to some vertex in S . Note that

$$N(S) \subset R$$

since G is bipartite. If $S = \emptyset$ then $N(S) = \emptyset$. The neighborhood $N(T) \subset L$ of $T \subset R$ is defined similarly.

The *unique neighbor set* $U(S) \subset R$ of the set $S \subset L$ consists of all $w \in R$ such that

$$|N(\{w\}) \cap S| = 1,$$

i.e., it consists of all vertices “on the right” whose respective neighborhoods have “a single vertex” intersection with S “on the left.” We make extensive use of the following refinement that “prescribes” that intersection. For $v \in S$, the set $U(S, v)$ consists of all $w \in R$ such that

$$N(\{w\}) \cap S = \{v\}.$$

Note that

$$U(S) \subset N(S),$$

⁷ We do not necessarily require that each of L, R is nonempty. But, of course, if at least one of them is, then also $E = \emptyset$.

and that

$$U(S, v) \subset N(\{v\}).$$

Also note that, if $v, v' \in S$ and if $v \neq v'$, then

$$U(S, v) \cap U(S, v') = \emptyset.$$

The corresponding notions for $T \subset R$ may be defined similarly, but we will not need any of these.

Let $d, d', f_1, f'_1, f'_2, f_2, f, f'$ be nonnegative integers.

We say that the graph G is *d-left-bounded* if, for each $v \in L$, it holds that $|N(\{v\})| \leq d$. In other words, each of “the degrees on the left” is at most d . If there is equality for each vertex, i.e., each of the degrees on the left equals d , we say that the graph G is *d-left-regular*. Similarly for *d'-right-bounded*. The graph G is *(d, d')-bi-bounded* if it is *d-left-bounded* and *d'-right-bounded*. Finally, the graph G is *d-biregular* if it is *d-left-regular* and *d-right-regular*.

Definition 4 (Unique Neighbor Property). *The set S has the unique neighbor property if it holds that $U(S) \neq \emptyset$.*

Definition 5 (Strong Unique Neighbor Property of a Set). *The set S has the strong unique neighbor property if, for each $v \in S$, we have $U(S, v) \neq \emptyset$.*

Definition 6 (f -Strong Unique Neighbor Property of a Set). *The set S has the f -strong unique neighbor property if, for each $v \in S$, we have $|U(S, v)| > f$.*

Remark 1. The latter is equivalent to the requirement that, for an arbitrary selection of f vertices from R , the set S has the strong unique neighbor property in the bipartite subgraph G' obtained from G by removing this selection of f vertices from R and by removing their incident edges from E .

Remark 2. Unlike the unique neighbor property, the (f -)strong unique neighbor property is *anti-monotonous* in the following sense. If S has the (f -)strong unique neighbor property and if $S' \subset S$ (and if $f' \leq f$), then S' has the (f' -)strong unique neighbor property. This follows trivially by exploiting that fact that, by definition, “intersection with S can be prescribed.”

Definition 7 ((f_1, f_2) -Strong Unique Neighbor Property of a Graph G). *The bipartite graph $G = (L, R, E)$ has the (f_1, f_2) -strong unique neighbor property if each set $S \subset L$ with $|S| = f_1$ has the f_2 -strong unique neighbor property.*

By an earlier remark, it follows that this property is anti-monotonous in the sense that the (f_1, f_2) -strong unique neighbor property implies the (f'_1, f'_2) -strong unique neighbor property if $f'_1 \leq f_1$ and $f'_2 \leq f_2$.

The unique neighbor property has been widely considered before and it has many known applications. There are also several applications of an *approximate*

version of the strong unique neighbor property, namely where the property is only guaranteed to hold for a given fraction of each set S .

The following lemma collects some immediate, useful consequences of the definitions.

Lemma 1. *Let $G = (L, R, E)$ be a d' -right-bounded bipartite graph. Suppose there are nonnegative integers f_1, f_2 and a cover of L consisting of sets $S \subset L$ such that $|S| = f_1$ such that S has the f_2 -strong unique neighbor property. Then each of the following holds.*

- (1) $|R| \geq N(S) \geq f_1(f_2 + 1)$, for each S in the cover.
- (2) For each $v \in L$, it holds that $|N(\{v\})| \geq f_2 + 1$.
- (3) $d' \geq (f_2 + 1) \frac{|L|}{|R|}$ if $R \neq \emptyset$.

PROOF. Fix an arbitrary $v \in L$. Let $S \subset L$ be such that $v \in S$, $|S| = f_1$ and S has the f_2 -strong unique neighbor property. Such S exists by the cover condition. Since we have $U(S, v) \subset N(\{v\})$ in general and since we have $|U(S, v)| \geq f_2 + 1$ by the choice of S , the second claim follows. As to the third claim, we have

$$d'|R| \geq |E| \geq (f_2 + 1)|L|,$$

where the inequality on the left follows by the definition of d' -right-boundedness and where the inequality on the right follows from the second claim. As to the first claim, since the sets $U(S, v) \subset R$ with $v \in S$ are pairwise disjoint in general and since each of them satisfies $|U(S, v)| \geq f_2 + 1$ by the choice of S , we have that

$$|R| \geq |N(S)| \geq f_1(f_2 + 1).$$

△

Of course, the lemma holds if the graph has the (f_1, f_2) -unique neighbor property. But its actual formulation under the weaker cover condition is convenient for a purpose later on.

4.2 Details of the Proof

We show the following theorem, which immediately implies Theorem 3 by the correspondence between bi-partite graphs and the balls-and-buckets game explained in the introduction.

Theorem 5. *There is an effective construction that, for each $k > 1$, gives a bipartite graph $G = (L, R, E)$ such that*

- (1) $|L| = |R| = ck^2$ where $4 < c < 16$,
- (2) G is d' -right-bounded with $d' = k$
- (3) G has the (f_1, f_2) -strong unique neighbor property with $f_1 = f_2 = k$.

Moreover, under our conditions that $f_1, f_2 \in \Omega(k)$ and that $|R| = O(|L|)$, each of the achieved parameters for $|L|$ and d' is asymptotically optimal.

To prove this theorem, we now show the claimed construction and provide its analysis. The optimality claim is an immediate consequence of Lemma 1; by substitution of the conditions (dictated by our application to Sigma-protocols), we get $|L| \in \Omega(k^2)$ and we get $d' \in \Omega(k)$.

Now let \mathcal{H} be a ρ -universal family of hash functions $h : X \rightarrow Y$. Thus, for each $x, x' \in X$ with $x \neq x'$, the collision probability that $h(x) = h(x')$ is at most ρ if $h \in \mathcal{H}$ is selected uniformly random.⁸

We define a bipartite graph $G = (X, \mathcal{H} \times Y, E)$ as follows. For a pair

$$(x, (h, y)) \in X \times (\mathcal{H} \times Y),$$

we declare

$$(x, (h, y)) \in E \quad \text{if and only if} \quad h(x) = y.$$

We also define

$$d' = \max_{(h,y) \in \mathcal{H} \times Y} |\{h^{-1}(y)\}|,$$

the maximum preimage size. Thus, the graph G is d' -right-bounded. Note that each of the degrees on the left equals $|\mathcal{H}|$. Thus, the graph G is $|\mathcal{H}|$ -left-regular.

Before proceeding, we first argue why we may exclude the case $\rho = 0$. This case arises if and only if each of the functions is injective. Now, even if some $h \in \mathcal{H}$ is injective, this implies that $|Y| \geq |X|$. So, under our condition that $|R| = O(|L|)$, it should be the case that $|\mathcal{H}|$ is constant. But this leads to a contradiction. Namely, since G is $|\mathcal{H}|$ -left-regular, it follows that G is left-bounded by a constant. But, by Lemma 1, each of the left-degrees is greater than f_2 and $f_2 \in \Omega(k)$ by our condition. So we assume $\rho \neq 0$.

Lemma 2. *Let $S \subset X$ be nonempty. Then, for each $x \in S$, it holds that*

$$(1 - \rho(|S| - 1)) |\mathcal{H}| \leq |U(S, x)| \leq |\mathcal{H}|$$

PROOF. The inequality on the RHS follows from the facts that $U(S, x) \subset N(\{x\})$ in general and that, by $|\mathcal{H}|$ -left-regularity of G , we have $|N(\{x\})| = |\mathcal{H}|$. As to the inequality on the LHS, fix S . In the case that $|S| = 1$, we have $U(S, x) = N(\{x\})$ and, once again by $|\mathcal{H}|$ -left-regularity, we have $|N(\{x\})| = |\mathcal{H}|$. So the inequality follows. Now assume $|S| > 1$ and fix $x \in S$. Consider the neighborhood of x , i.e., the set

$$N(\{x\}) = \{(h, h(x)) : h \in \mathcal{H}\} \subset \mathcal{H} \times Y.$$

It is clear at once that

$$|U(S, x)| = |\{h \in \mathcal{H} : \text{for each } x' \in S \setminus \{x\}, \text{ it holds that } h(x) \neq h(x')\}|$$

Fixing $x' \in S \setminus \{x\}$ for now, there are at most $\rho|\mathcal{H}|$ hash functions h such that $h(x) = h(x')$, by definition of collision probability. Hence, the number of hash functions h such that $h(x) = h(x')$ for *some* $x' \in S \setminus \{x\}$ is at most $\rho|\mathcal{H}|(|S| - 1)$.

⁸ Note that $\rho = 0$ only if each $h \in \mathcal{H}$ is injective.

In conclusion, the number of hash functions h such that $h(x) \neq h(x')$ for each $x' \in S \setminus \{x\}$ is at least $(1 - \rho(|S| - 1)) |\mathcal{H}|$ and the claim follows. \triangle

Note that the lemma only gives a nontrivial result if $|S| < 1 + 1/\rho$.

Let p be a prime number with $p \geq 2k + 1$. By Bertrand's Postulate, there exists such prime p with $p < 4k$. Now consider the family with

$$\mathcal{H} = \mathbb{F}_p, X = \mathbb{F}_p^2, Y = \mathbb{F}_p$$

such that, for $h \in \mathbb{F}_p$, the corresponding hash function is defined as

$$\begin{aligned} h : \mathbb{F}_p^2 &\rightarrow \mathbb{F}_p \\ (x_0, x_1) &\mapsto x_0 h + x_1. \end{aligned}$$

One verifies directly that for this family we can take

$$\rho = 1/p \text{ and } d' = p.$$

Setting $|S| = k$, it follows by Lemma 2 that, for each $x \in S$, we have

$$|U(S, x)| \geq (1 - (k - 1)/p)p = p - k + 1.$$

Therefore, $|U(S, x)| > k$ if the prime p satisfies $p \geq 2k + 1$. This concludes the proof of Theorem 5.

4.3 Alternative Approaches and Generalization

An alternative constructive approach can be based on graphs G with ‘‘excellent expansion,’’ a basic concept from the theory of expander graphs. We say that a d -left-bounded graph G *expands excellently* on a set $S \subset L$ if the neighborhood $N(S) \subset R$ of S satisfies

$$|N(S)| \geq (1 - \epsilon)d|S|$$

where ϵ is a nonnegative real number with

$$\epsilon < 1/2.$$

Excellent expansion is well-known to imply the unique neighbor property. We adapt the arguments so as to imply the (f_1, f_2) -strong unique neighbor property instead, in certain parameter regimes. Then we discuss elementary construction of suitable expander graphs. We elaborate below.

The following lemma is well-known.

Lemma 3. *Suppose G is d -left-bounded. If $|N(S)| \geq (1 - \epsilon)d|S|$, then*

$$|U(S)| \geq (1 - 2\epsilon)d|S|.$$

PROOF. Since G is d -left-bounded, there are at most $d|S|$ edges “emanating” from S and “arriving” at $N(S)$. Write m_1 for the number of vertices $w \in N(S)$ with $|S \cap N(\{w\})| = 1$. Then we have the obvious bound

$$m_1 + 2(|N(S)| - m_1) \leq d|S|.$$

Therefore,

$$m_1 \geq 2|N(S)| - d|S|.$$

Since $|N(S)| \geq (1 - \epsilon)d|S|$, it follows that

$$m_1 \geq (1 - 2\epsilon)d|S|,$$

as desired. △

Using a “greedy argument” the f -strong unique neighbor property for a set is implied by a large unique neighbor set, as follows. Let δ be a real number with $0 < \delta \leq 1$.

Lemma 4. *Suppose that G is d -left-bounded ($d > 0$) and that $S \subset L$ is nonempty. Write $|U(S)| \geq (1 - \delta)d|S|$, where δ is a real number with $0 \leq \delta \leq 1$. If*

$$\delta|S| < 1 - \frac{f}{d},$$

the set S has the f -strong unique neighbor property.

PROOF. If $|S| = 1$, say $S = \{v\}$, then it follows at once that $|U(S, v)| = |N(\{v\})| > f$ and the claim follows. So now assume $|S| > 1$. Using a pigeon-hole argument, we see that, if

$$\frac{(1 - \delta)d|S| - f}{|S| - 1} > d, \quad (*)$$

then the set S has the f -strong unique neighbor property. Indeed, consider the subgraph G' obtained by removing some f vertices from R and by removing their incident edges from E . Towards a contradiction, suppose S does not have the strong unique neighbor property in G' . Say it fails on some $v \in S$. Then the inequality implies that there is some $v' \in S \setminus \{v\}$ with degree greater than d , which contradicts the fact that, just as the graph G , its subgraph G' is d -left-bounded. The proof is finalized by observing that the inequality (*) is equivalent to the inequality $\delta|S| < 1 - f/d$. △

By combining Lemmas 3 and 4 we get the following sufficient condition for the f -strong unique neighbor property of a set $S \subset L$.

Corollary 1. *Suppose G is d -left-bounded ($d > 0$) and suppose $S \subset L$ is nonempty. If, for some nonnegative real number ϵ and for some nonnegative integer f , it holds that*

(1) $|N(S)| \geq (1 - \epsilon)d|S|$ and

$$(2) \ 2\epsilon|S| < 1 - \frac{f}{d},$$

then S has the f -strong unique neighbor property.

Remark 3. In order to satisfy the conditions, it is necessary that $\epsilon < 1/2$, i.e., expansion is excellent.

We now discuss constructions based on this excellent expansion approach. Recall that, under the constraints that $f_1, f_2 \in \Omega(k)$ and that $|R| = O(|L|)$, we wish to minimize $|L|$ (the size of the set of left-vertices) and d' (the right-degree). From the conditions in Corollary 1, we then have that $1/\epsilon \in \Omega(k)$ and that $d \in \Omega(k)$.

Observe that the construction in Theorem 5 gives excellent expansion for all sets of size k . Namely, by Lemma 1, the size of the neighborhood of a set of size k equals $(p - k + 1)k$, where $p = c'k$ for some constant $c' > 2$. Therefore, in this case, $\epsilon = (1 - 1/k) \cdot 1/c' < 1/2$ but $1/\epsilon = c'k/(k - 1) \in O(1)$. In conclusion, the result of Theorem 5 cannot also be obtained by application of Corollary 1, except for less favorable parameter settings. Namely, it would require setting p super-linear in k , thus rendering $|L|$ super-quadratic. Furthermore, since $d \in \Omega(k)$, excellent constant left-degree expander graphs [CRVW02] do not apply here. A (well-known) variation on the greedy counting arguments above shows that a combination of excellent expansion and constant left-degree *does* imply an *approximate* version of the f -strong unique neighbor property, i.e., it holds for a certain fraction of each S . But this notion is not sufficient for our present purposes.

To illustrate this approach based on excellent expansion, we show a construction from *random* permutations instead. This is in contrast with the deterministic approach in Theorem 5 where permutations had to be excluded. We use a classical result by Bassalygo [Bas81] who showed a *Monte Carlo* construction of bipartite graphs with excellent expansion. Basically, a (d, d) -bi-bounded bipartite graph with $|L| = |R|$ is constructed by “taking the union” of d random perfect bipartite matchings (or, equivalently, permutations). In general, the probability of success of this procedure is high but not exponentially close to 1. Therefore, it is not sufficient for our purposes. However, choosing convenient parameters in the procedure, one can show that each individual set S of size k has the required expansion with probability of success exponentially (in k) close to 1. It is not hard to see that this weaker “probabilistic, set-wise” property is sufficient for our purposes as well. The downside, in addition to being Monte Carlo, is that $|L|$ here is *cubic* instead of quadratic. All in all, this leads to the following theorem.

Theorem 6. *There is an efficient construction that, for each $k \geq 1$, gives a bipartite graph $G = (L, R, E)$ such that*

- (1) $|L| \in O(k^3)$ and $|R| = |L|$,
- (2) G is $O(k)$ -right-bounded,
- (3) for each fixed set $S \subset L$ with $|S| = k$, it holds that S has the k -strong unique neighbor property, except with exponentially small (in k) probability,

Remark 4. Lemma 1 implies that such a probabilistic approach obeys the same lower bounds that $|L| \in \Omega(k^2)$ and $d' \in \Omega(k)$ as in the deterministic case, conditioned on $f_1, f_2 \in \Omega(k)$ and $|R| = O(|L|)$. In a nutshell, there is a small cover of L by sets S of size f_1 such that, by a union-bound argument, each set S in this cover has the f_2 -strong unique neighbor property, with probability still extremely close to 1.

We will prove Theorem 6 by combining Corollary 1 with Proposition 1 below. Suppose $|L| = |R| = n$. Write $L = \{v_1, \dots, v_n\}$ and $R = \{w_1, \dots, w_n\}$. For a permutation π on $\{1, \dots, n\}$, define $E(\pi) \subset L \times R$ as the set of edges

$$\{(v_1, w_{\pi(1)}), \dots, (v_n, w_{\pi(n)})\}.$$

Suppose $1 \leq d \leq n$. For a d -vector $\Pi = (\pi_1, \dots, \pi_d)$ of (not-necessarily distinct) permutations on $\{1, \dots, n\}$, define the set

$$E(\Pi) = \bigcup_{j=1}^d E(\pi_j) \subset L \times R$$

and define the bipartite graph

$$G(\Pi) = (L, R, E(\Pi)).$$

Note that G is a (d, d) -bi-bounded (undirected) bipartite graph (without multi-edges). We have the following proposition.

Proposition 1. *Let $G = (L, R, E)$ be a random (d, d) -bi-bounded bipartite graph with $|L| = |R| = n$ as described above. Let α be a real number with $0 < \alpha < 1$. Then, for any fixed set $S \subset L$ with $|S| = \alpha n$, it holds that*

$$N(S) \geq (d - 2)|S|,$$

except with probability

$$p'_S \leq \left(\frac{d^2 \alpha e}{2(1 - \alpha)} \right)^{2\alpha n},$$

where e denotes Euler's constant.

PROOF. Choose the d permutations π_1, \dots, π_d sequentially. For convenience, write $S = \{1, \dots, s\}$. For $i = 1, \dots, s$ and $j = 1, \dots, d$, consider the random variables

$$X_i^j,$$

the image of $i \in S$ under the permutation π_j . We now think of these as “ordered” $X_1^1, \dots, X_s^1, X_1^2, \dots, X_s^2, \dots$, “increasing” from left to right.

For given X_i^j , condition on all “prior” random variables in the ordering. The probability that X_i^j is a repeat, i.e., it lands in what is $N(S)$ -so-far is at most

$$\frac{d|S|}{n - i + 1} \leq \frac{d|S|}{n - |S|}.$$

Here the denominator on the LHS is due to the fact that when choosing the image of i , the $i - 1$ distinct images of $1, \dots, i - 1$ are already taken. Hence, the probability p'_S that the event $|N(S)| \leq (d - 2)|S|$ occurs is at most the probability of the event that there are $2|S|$ repeats. By the union bound, the latter probability is clearly at most

$$\binom{d|S|}{2|S|} \left(\frac{d|S|}{n - |S|} \right)^{2|S|}$$

Therefore,⁹

$$p'_S \leq \binom{d|S|}{2|S|} \left(\frac{d|S|}{n - |S|} \right)^{2|S|} \leq \left(\frac{de}{2} \right)^{2|S|} \left(\frac{d|S|}{n - |S|} \right)^{2|S|} = \left(\frac{d^2 \alpha e}{2(1 - \alpha)} \right)^{2\alpha n}.$$

△

The proposition and its proof are adapted from the classical expander graph construction due to Bassalygo [Bas81]. Our exposition follows (part of) the proof of Theorem 4.4 in [Vad12]. The reason we do not apply the Bassalygo result directly is that the success probability of the construction of an excellent expander is high (i.e., constant) but still much too small for our purposes. Fortunately, we can do with the slightly weaker requirement on G that, for any *fixed* set S of precisely the dictated size, the probability that the set S does not expand excellently is negligibly small. As this saves two applications of the union bound, one to quantify over all sets S of the dictated size and one to quantify over the subsets of size smaller than the dictated size, we get exponentially small failure probability instead of constant.

Now let c_1, c_2 be arbitrary positive integers. Set

- (1) $f_1 = c_1 k, f_2 = c_2 k.$
- (2) $d = c_3 k$ with $c_3 = c_1 + c_2 + 1.$
- (3) $\alpha = \frac{1}{d^2 e + 1}.$
- (4) $n = m = \frac{c_1}{\alpha} k = (d^2 e + 1)c_1 k = (c_3^2 e k^2 + 1)c_1 k = c_1 c_3^2 e k^3 + c_1 k.$

Then, for each fixed set $S \subset L$ with $|S| = f_1$, it holds that S has the f_2 -strong unique neighbor property, except with exponentially small (in k) probability

$$p' \leq \left(\frac{1}{2} \right)^{2c_1 k}$$

Namely, for each set S of size $K = \alpha n = c_1 k = f_1$, it holds that $N(S) \geq (d - 2)|S|$. Note that $\epsilon = 2/d$ here. This means that the second condition for the f_2 -strong unique neighbor property of sets of this size is $f_1 + f_2 < d$. This is satisfied by definition. Efficiency of the construction is obvious. This concludes the proof of Theorem 6.

⁹ Note that $\left(\frac{r}{s} \right)^s \leq \binom{r}{s} \leq \left(\frac{re}{s} \right)^s.$

4.4 A Generalized Construction and Proof of Theorem 4

We now generalize the construction from Theorem 6 to get one where the number of nodes can be much smaller compared to the size of the special set S at the price that the unique neighbour property holds only in a weaker sense.

Recall that $U(S, v)$ is the set of all $w \in R$ such that

$$N(\{w\}) \cap S = \{v\}.$$

The set $U(S)$ is the union of $U(S, v)$ for all $v \in S$.

Definition 8 (*(s, r) -approximate Unique Neighbour Property of a Set*). *Given a bipartite graph $G = (L, R, E)$, the set $S \subseteq L$ has (s, r) -approximate unique neighbour property if there exists a subset $S_1 \subseteq S$ of size s such that for any set $T \subseteq R$ of size r , we have*

$$|U(S, v) - T| > 0 \quad \forall v \in S_1.$$

We may ask whether such set exists. Our following lemma answers this question.

Lemma 5. *Suppose $G = (L, R, E)$ is d -left-bounded. If $U(S) \geq (1 - \epsilon)d|S|$, then the set S has (s, r) -approximate unique neighbour property for $s = (1 - \epsilon)|S| - \frac{r}{d}$.*

Proof. Let $T \subseteq R$ of size r . Let $T_1 = U(S) - T$ and $S_1 \subseteq S$ be the set such that for any $v \in S_1$, $N(v) \cap T_1 \neq \emptyset$. Since $\deg(v) \leq d$, there are at least $\frac{|T_1|}{d}$ vertices contained in S_1 . We are done.

Combining Lemmas 5 and 3, we get a sufficient condition for the unique neighbour property of a set $S \subseteq L$.

Corollary 2. *Suppose G is d -left-bounded. If $N(S) \geq (1 - \epsilon)d|S|$, then the set $S \subseteq L$ has (s, r) -approximate unique neighbour property for $s = (1 - 2\epsilon)|S| - \frac{r}{d}$.*

Now, if in Proposition 1, we set $|S| = k$, $n = O(k^{1+2c})$, $\epsilon = \frac{2}{d}$ and $d = k^c$ for a constant $0 < c < 1$, we can proceed in a way similar to the proof of Theorem 6, and get the following theorem, which immediately implies Theorem 4.

Theorem 7. *There is an efficient construction that, for each $k \geq 1$ and for a constant $0 < c < 1$, gives a bipartite graph $G = (L, R, E)$ such that*

- (1) $|L| = |R| = O(k^{1+2c})$,
- (2) G is $O(k^c)$ -right-bounded,
- (3) for each fixed set $S \subseteq L$ with $|S| = k$, it holds that S has the $(k - 5k^{1-c}, k)$ -approximate unique neighbour property, except with exponentially small (in k) probability.

5 Acknowledgements

We are grateful for feedback and suggestions we received after we circulated a preprint of this work on July 6, 2016. Omer Reingold [Rei16] suggested the approach from Theorem 5 as an improvement to our “strong unique neighbor” Theorem 6, which was also stated in this preprint. This suggestion not only gave a deterministic construction instead of our Monte Carlo one but it also improved the left-vertex-set parameter from cubic to quadratic. We thank him for allowing us to incorporate his suggestion. Independently from Omer Reingold, Gilles Zémor [Z16] suggested an alternative improvement to our “strong unique neighbor” Theorem 6 which removed the probabilism as well but left the parameters essentially unchanged. His suggestion was based on combining our excellent expansion approach with an argument involving the girth of certain graphs and an application of Turán’s Theorem. Furthermore, we thank Gilles Zémor for several helpful discussions and pointers to the literature (also at an earlier stage). Finally, thanks to Amin Shokrollahi and Salil Vadhan for answering questions about the literature.

References

- Bas81. Leonid Bassalygo. Asymptotically optimal switching circuits. *Problems in Information Transmission*, 17(3):81–88, 1981.
- BDLN16. Carsten Baum, Ivan Damgård, Kasper Larsen, and Michael Nielsen. How to prove knowledge of small secrets. Cryptology ePrint Archive, Report 2016/538, 2016. To appear in Crypto 2016.
- BDOZ11. Rikke Bendlin, Ivan Damgård, Claudio Orlandi, and Sarah Zakarias. Semi-homomorphic encryption and multiparty computation. In *Advances in Cryptology – EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 169–188. Springer Berlin Heidelberg, 2011.
- BG93. Mihir Bellare and Oded Goldreich. On defining proofs of knowledge. In *Advances in Cryptology—CRYPTO’92*, pages 390–420. Springer, 1993.
- BGV12. Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCS ’12*, pages 309–325, New York, NY, USA, 2012. ACM.
- BV14. Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) lwe. *SIAM Journal on Computing*, 43(2):831–871, 2014.
- CD09. Ronald Cramer and Ivan Damgård. On the amortized complexity of zero-knowledge protocols. In *Advances in Cryptology-CRYPTO 2009*, pages 177–191. Springer, 2009.
- CDK14. Ronald Cramer, Ivan Damgård, and Marcel Keller. On the amortized complexity of zero-knowledge protocols. *Journal of cryptology*, 27(2):284–316, 2014.
- CRVW02. Michael Capalbo, Omer Reingold, Salil Vadhan, and Avi Wigderson. Randomness conductors and constant-degree lossless expanders. In *STOC*, pages 659–668, 2002.

- DF02. Ivan Damgård and Eiichiro Fujisaki. A statistically-hiding integer commitment scheme based on groups with hidden order. In *Advances in Cryptology—ASIACRYPT 2002*, pages 125–142. Springer, 2002.
- DKL⁺13. Ivan Damgård, Marcel Keller, Enrique Larraia, Valerio Pastro, Peter Scholl, and Nigel Smart. Practical covertly secure mpc for dishonest majority - or: Breaking the spdz limits. In *ESORICS*, pages 1–18, 2013.
- DPSZ12. Ivan Damgård, Valerio Pastro, Nigel Smart, and Sarah Zakarias. Multi-party computation from somewhat homomorphic encryption. In *Advances in Cryptology - CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 643–662. Springer, Berlin, Germany, 2012.
- GGH96. Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Collision-free hashing from lattice problems. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 3, pages 236–241, 1996.
- GSW13. Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *Advances in Cryptology—CRYPTO 2013*, pages 75–92. Springer, 2013.
- KL14. Jonathan Katz and Yehuda Lindell. *Introduction to modern cryptography*. CRC Press, 2014.
- LMPR08. Vadim Lyubashevsky, Daniele Micciancio, Chris Peikert, and Alon Rosen. Swift: A modest proposal for fft hashing. In *Fast Software Encryption*, pages 54–72. Springer, 2008.
- Lyu08. Vadim Lyubashevsky. Lattice-based identification schemes secure under active attacks. In *Public Key Cryptography—PKC 2008*, pages 162–179. Springer, 2008.
- Lyu09. Vadim Lyubashevsky. Fiat-shamir with aborts: Applications to lattice and factoring-based signatures. In *Advances in Cryptology—ASIACRYPT 2009*, pages 598–616. Springer, 2009.
- Rei16. Omer Reingold. private communication to the authors, July 2016.
- Vad12. Salil Vadhan. *Pseudorandomness*. Now publishers, 2012.
- Z16. Gilles Zémor. private communication to the authors, July 2016.